Model predictive control for MR-guided ultrasound hyperthermia in cancer therapy

Model Predictive Control for MR-guided Ultrasound Hyperthermia in Cancer Therapy

Daniel Deenen

About the cover

The crab is internationally used as a symbol for cancer. This relation originated in ancient Greece, and is generally credited to the famous physician Hippocrates of Kos (460 – 370 BC), who is often referred to as the "Father of Medicine." He used the word *καρκίνος* (karkinos), which in (ancient) Greek is synonymous to *crab*, to describe malignant tumors after encountering breast tumors consisting of a hard swelling with finger-like spreading projections that infiltrated the surrounding tissue, which reminded him of the body and feet of a crab. This analogy was later notably perpetuated by the Roman encyclopedist Aulus Cornelius Celsus (25 BC – AD 50), who in his celebrated work *De Medicina* translated *καρκίνος* into *cancer*, the Latin word for *crab*, of which the use extends into modern English.

Moreover, the cover depicts a blue crab (Callinectes sapidus), which by nature actually has a mottled brownish body with blue legs and claws, but has turned red after being cooked, which for this dissertation is a reference to hyperthermia therapy involving heat to treat cancer.

Photo: freepik.com/wirestock

disc

The author has successfully completed the educational program of the Graduate School of the Dutch Institute of Systems and Control (DISC).

The work described in this thesis was carried out at the Eindhoven University of Technology (TU/e), and is part of the research project "Control-2-Act" with project number 15195. This project is made possible by the Dutch Cancer Society (KWF) and the Dutch Research Council (NWO) as part of their joint Partnership Program: "Technology for Oncology," and is partially financed by the PPP Allowance made available by Top Sector Life Sciences & Health.

A catalogue record is available from the Eindhoven University of Technology Library.
ISBN: 978-90-386-5165-1

Reproduction: Ipskamp Printing, Enschede, the Netherlands

FSC keurmerk hier

Model Predictive Control for MR-guided Ultrasound Hyperthermia in Cancer Therapy

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus prof.dr.ir. F.P.T. Baaijens, voor een
commissie aangewezen door het College voor
Promoties, in het openbaar te verdedigen
op donderdag 3 december 2020 om 11.00 uur

door

Daniel Andreas Deenen

geboren te Leonberg, Duitsland

Dit proefschrift is goedgekeurd door de promotoren en de samenstelling van de promotiecommissie is als volgt:

voorzitter:	prof.dr. H. Nijmeijer
1^{e} promotor:	prof.dr.ir. W.P.M.H. Heemels
2^{e} promotor:	dr.ir. A.G. de Jager
leden:	dr.ir. M.M. Paulides (Erasmus MC Kanker Instituut)
	prof.dr. H. Grüll (Uniklinik Köln)
	prof.dr.ir. D. Limón Marruedo (Universidad de Sevilla)
	prof.dr.ir. B. De Schutter (Technische Universiteit Delft)
	dr. M. Lazar

Het onderzoek dat in dit proefschrift wordt beschreven is uitgevoerd in overeenstemming met de TU/e Gedragscode Wetenschapsbeoefening.

Prediction is difficult, especially about the future.

Nostradamus

Summary

Model Predictive Control for MR-guided Ultrasound Hyperthermia in Cancer Therapy

With more than 18 million new cases and 9 million deaths each year, cancer is the second leading cause of death globally. Due to the ageing of the world population, combined with the dramatically higher risk for cancer at older age, the incidence and morbidity rates of cancer are expected to keep increasing in the near future. Currently, cancer treatments typically consist of a combination of surgery, radiotherapy, and chemotherapy, which due to their invasive nature or toxicity bring on severe side effects for the patients.

In mild local hyperthermia therapy, tumors are heated to approximately 42 °C for a prolonged period of time, while preserving non-elevated temperatures in the healthy tissue. Clinical evidence shows that this significantly enhances the therapeutic efficacy of radio- and chemotherapies in the tumor without inducing additional toxicity in the healthy tissue. Consequently, hyperthermia enables substantially improved treatment success, as well as the use of lower systemic doses of radiation and drugs to reduce the severity of the negative side effects typically associated with cancer treatment. Finally, hyperthermia treatments can be performed noninvasively, which eliminates the need for additional surgery. Therefore, hyperthermia therapy is an urgently needed addition to conventional cancer treatment modalities.

However, to unlock the full therapeutic potential of hyperthermia, the tumor temperature must be accurately induced and maintained throughout the entire treatment. Magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU) is a rapidly developing and highly promising technology for noninvasive feedback-controlled hyperthermia, enabling powerful heating with millimeter-scale accuracy guided by near-real-time thermometry via an MRI scanner. Using model predictive control (MPC), the tissue's future thermal behavior and restrictive control constraints can be taken into account explicitly while optimizing the heating inputs. It is the aim of this thesis to enable safe and high-quality MR-

HIFU hyperthermia treatments using MPC-based temperature feedback control strategies.

The first contribution of this thesis is the development of offset-free MPC for ultrasound-based hyperthermia, enabling robust temperature control performance in the presence of plant-model mismatch. This is particularly important for hyperthermia treatments, as modeling errors are typically inevitable due to the natural variation of patients' tissue properties, some of which even exhibit time- and temperature-dependent behavior as a result of the body's thermoregulatory response.

Second, we address the issue that HIFU applicators typically have a small ($\leq$2 cm) local heating range. For (MR-)HIFU hyperthermia treatments of tumors larger than this range, the applicator itself must be relocated mechanically, which results in non-negligible actuator downtime with respect to the system dynamics. In this case, optimally controlling the temperature requires the online solving of both a discrete actuator allocation problem, and a continuous optimization problem for the local heating profile. Therefore, in this thesis a novel modeling framework tailored to systems with such actuator switches is developed, specifically designed to allow for user-friendly and systematic modeling, and to yield a compact model that results in efficient mixed-integer (MI-)MPC schemes for large-volume MR-HIFU hyperthermia.

As the third contribution, a target-conformal optimal actuator placement method is proposed. Based on the specific tumor geometry and tissue properties, this procedure distributes the admissible heating locations and mechanical applicator positions in such a manner that the achievable treatment quality is optimized. In addition, this method allows for limiting the controller complexity, thereby facilitating the real-time application of control strategies that offer superior control potential, but are also more computationally demanding, such as (MI-)MPC.

As a fourth contribution, this thesis presents a hierarchical (H)MI-MPC architecture for large-volume MR-HIFU hyperthermia. This HMI-MPC significantly reduces the computational burden compared to MI-MPC, without sacrificing control performance and treatment quality. This is done by first solving the challenging actuator switching problem using a low-complexity model, followed by the computationally cheap optimization of the heating profiles using a high-fidelity model.

The developed controllers are demonstrated using numerical case studies, phantom experiments, and in-vivo experiments, paving the way for enhanced MR-HIFU hyperthermia treatments in clinics.

Contents

What medicines cannot cure, iron cures;
what iron cannot cure, fire cures;
but what fire cannot cure, must be considered incurable.

Hippocrates of Kos

Chapter 1

Introduction

In local mild hyperthermia therapy for cancer treatment, tumors are heated to temperatures ranging between 40 °C and 43 °C for an extended period of time. This stimulates a range of biological mechanisms, boosting the body's own immunological response against cancer cells, and synergistically enhancing the efficacy of radio- and chemotherapy. Mild hyperthermia is non-toxic for the body, and thus does not introduce any additional undesired side effects such as typically associated with cancer treatment, rendering it a valuable adjuvant cancer treatment modality. Especially considering the growing proportion of people suffering from cancer, partly as a result of our ageing society, improved cancer treatment options form an urgently needed development in today's world. However, to unlock the full therapeutic potential of hyperthermia therapy, accurately controlled heating is required. In this respect, magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU) offers large potential for noninvasive feedback-controlled hyperthermia, as it a rapidly developing technology, enabling powerful ultrasound heating with millimeter-scale accuracy guided by near-real-time MR-based thermometry. Currently, however, most MR-HIFU treatments are performed either by following pre-planned protocols, or using simple feedback schemes, and as a result the realized treatment quality is far from optimal. It is the objective of this thesis to provide a significant step toward achieving high-quality treatments by enabling more accurate and uniform heating using model predictive control (MPC) strategies, which in turn facilitates more widespread clinical application of and research into hyperthermia. To further introduce and motivate the research in this thesis, this chapter elaborates on MR-HIFU hyperthermia in oncology, distills important research objectives from relevant challenges we face today, and outlines the individual contributions in this thesis.

1.1 Hyperthermia for cancer therapy

I

This section provides an introduction to hyperthermia therapy as a cancer treatment modality, providing the necessary background for motivating the research in this thesis.

1.1.1 Cancer treatment

With more than 18 million new cases and over 9 million deaths in 2018, cancer is the second most important cause of death globally [1]. In the Netherlands, it is the leading cause, with approximately one out of every three deaths being related to cancer [2]. Due to the continuing ageing of the world population, combined with the dramatically higher risk for cancer at older age, the incidence and morbidity rates of cancer are expected to keep increasing in the near future. Currently, conventional cancer treatments consist of some combination of surgically resecting the tumor, radiation therapy, and chemotherapy. Due to tremendous advances in oncology over the past decades, it is now well established that these methods can successfully be applied to achieve local control and improve long-term survival rates. However, surgical intervention may exert a significant burden on the patient due to its invasive nature and the resulting need for post-interventional wound care. Moreover, chemo- and radiotherapies may bring on severe toxicity-based side effects, considerably impacting the patient's quality of life during treatment. As a consequence, their dose is limited, which in turn imposes an upper bound on their achievable therapeutic effects. Finally, not all tumors can be successfully treated using the aforementioned methods. For example, metastases or tumors in high-risk locations may not be suitable for surgical resection, and badly perfused tumors may exhibit decreased sensitivity to radiation or chemotherapy due to the resulting hypoxia (i.e., the deprivation of oxygen) or poor drug penetration, respectively [3, 4]. As a result of these drawbacks, there exists a strong need for improved cancer treatment modalities.

1.1.2 Hyperthermia therapy

The medical benefits of heating body tissues to above their normal temperature have been long known by mankind. In fact, applying heat to the body for medical purposes most likely dates back to at least 3000 BC, when the ancient Egyptians used hot irons (which they called "fire drills") for the treatment of breast tumors, as documented by the Edwin Smith Papyrus [5], and Indian physicians used local and whole body heating for the treatment of a variety of ailments [6]. Thereafter, methods similar to the Egyptians' have also been used by Hippocrates of Kos in ancient Greece (460 – 370 BC) and Galen of Pergamon (AD 129 – 210) in ancient Rome, for example, and described by Aulus Cornelius Celsus (25 BC – AD 50) in his famous encyclopedia *De Medicina* [7].

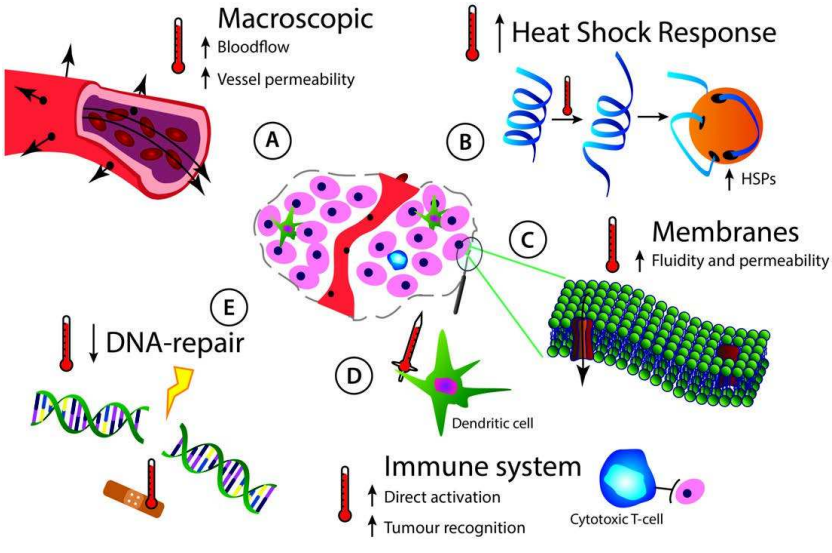

Figure 1.1: Overview of hyperthermia-induced biological mechanisms in the tumor (center), being (A) increased blood perfusion and vessel permeability, (B) protein unfolding and chaperoning by heat-shock proteins (HPSs), (C) increased cell membrane fluidity and permeability, (D) stimulated immune response, and (E) inhibited DNA repair [4].

Over the past decades in particular, as a consequence of the previously mentioned limitations recognized in the conventional cancer treatment options, researchers have shown renewed interest in thermal therapies for the treatment of cancer. In general, two distinct treatment types can be distinguished, namely hyperthermia and thermal ablation [3, 8]. Thermal ablation occurs for temperatures above 45 °C, sometimes even reaching 70 °C or higher, inducing coagulative necrosis of the treated tissue very rapidly [9]. Hyperthermia, on the other hand, is achieved when tissues are heated to temperatures between roughly 40 °C and 45 °C, and is ideally maintained for prolonged periods of time such as 60 to 90 minutes or more [10–12]. For temperatures above 43 °C, this may still cause direct cell killing as a result of heat-induced cytotoxicity. Mild hyperthermic temperatures up to 43 °C, however, stimulate a variety of biological mechanisms on molecular, cellular, and tissue level, which in turn boost the body's own immunological response to cancer cells, inhibit DNA repair, increase blood flow and oxygenation, and enhance drug uptake as a result of increased permeability of the blood vessels and cell membranes, see [3, 4, 12–17] for details. A schematic overview is given in Figure 1.1.

As a consequence, hyperthermia significantly increases the efficacy of radio- and chemotherapies, thereby enabling similar or improved treatment outcomes

using lower doses of radiation and drugs, and, in turn, reducing the severity of the negative side effects typically associated with cancer treatment. Indeed, clinical trials have provided ample evidence that hyperthermia is a potent adjuvant therapy for cancer treatment, see, for example, [3, 16, 18–24] and the references therein. When applying the heat locally, and preserving non-elevated temperatures in healthy tissue, hyperthermia allows for selectively sensitizing only the tumor region, while leaving the healthy tissue unaffected. Moreover, using temperature-sensitive liposomes, it enables targeted drug delivery, allowing for a further reduction of the required systemic drug concentrations and its side effects [25–27]. Most importantly, mild hyperthermia itself is non-toxic, thereby introducing no additional toxicity-related side effects, and can be applied noninvasively, rendering it highly appealing for clinical application to improve treatment outcome as well as the the patient's quality of life.

1.1.3 MR-guided high-intensity focused ultrasound

The two most commonly used methods for noninvasively inducing hyperthermia in clinics are via electromagnetic (EM) waves, either in the radiofrequency (3 Hz to 300 MHz) or microwave (300 MHz to 300 GHz) range, or by means of ultrasound, which are high-frequency (300 kHz to 10 MHz) acoustic waves [8, 14, 16]. To moderately heat larger volumes such as limbs or organs, also referred to as (loco)regional hyperthermia, EM-based applicators are currently considered more suitable, as these generate much wider power deposition profiles [28, 29]. Consequently, however, such applicators typically lack the spatial accuracy to compensate for centimeter- or millimeter-scale temperature variations resulting from, for example, blood vessels or other tissue inhomogeneities, and as a result may be unable to achieve and maintain a uniform temperature throughout the tumor. By contrast, high-intensity focused ultrasound (HIFU) offers unparalleled millimeter-accurate heating, but comes with a significantly smaller heating range of a few centimeters at most [30, 31]. For local hyperthermia, however, where it is the objective to heat only the tumor, while avoiding the unintentional heating of the surrounding healthy tissue, HIFU potentially enables tighter temperature control and, thereby, higher treatment quality.

HIFU has already been proposed as a noninvasive means for delivering heat inside the body over a century ago, see [32] and [33] as some of the earliest works. However, the lack of proper visualization technologies was one of the main reasons hampering further development for widespread clinical use. This changed, however, when more advanced imaging methods emerged, of which the three major modalities useful for noninvasive thermometry are X-ray computed tomography (CT) [34], diagnostic ultrasound [35, 36], or magnetic resonance imaging (MRI) [37]. From these, MRI provides the highest soft-tissue contrast, and is hence considered to be the best available option for high-performance temperature feedback [30]. Several techniques exist for obtaining volumetric temperature

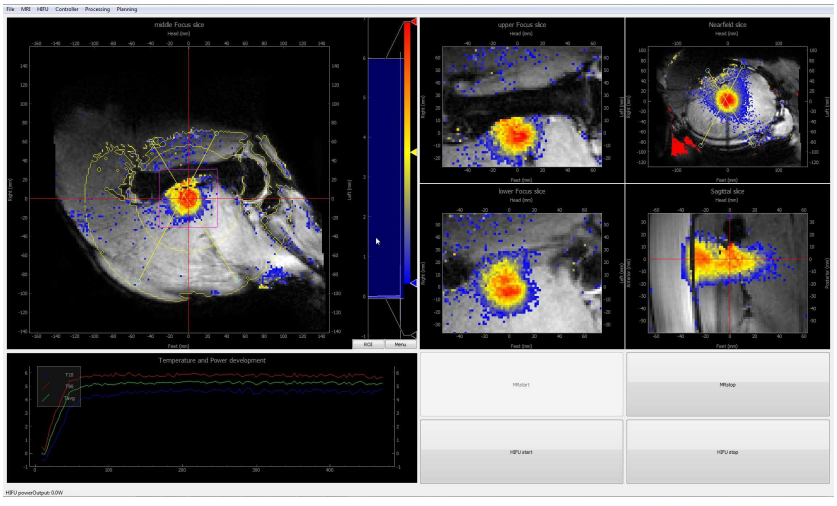

Figure 1.2: Example of MR-HIFU hyperthermia interface, monitoring near-real-time temperature cross sections perpendicular to the beam axis at target depth (left: top), directly above and below target depth (middle: top and bottom), well below target depth in the near field (right: top), and along the beam axis (right: bottom), and the overall target temperature evolution (left: bottom) [46].

maps by MRI, which is referred to as magnetic resonance (MR) thermometry, one of which is based on exploiting the temperature-dependent proton resonance frequency shift (PRFS) [38–40]. Today, PRFS-based MR thermometry is among the best available measurement technologies, providing near-real-time temperature data with excellent spatial and thermal resolution [41–45], an example of which is shown in Figure 1.2. Integrating this technology into HIFU treatments has led to MR-guided (MR-)HIFU as a valuable tool for image-guided or feedback-controlled thermal therapies [30, 31, 43, 47, 48].

1.1.4 Clinical MR-HIFU systems

The two most widely used clinical noninvasive MR-HIFU systems are the Insightec Exablate® (Insightec Ltd., Haifa, Israel) and the Profound Sonalleve® (Profound Medical Corp., Mississauga, Canada) [8, 43, 49], of which the latter is shown in Figure 1.3. These MR-HIFU setups contain an extracorporeal phased-array transducer integrated into a dedicated MR-compatible treatment table, as schematically depicted in Figure 1.4. The transducers consist of many small acoustic elements (256 in case of the Sonalleve). For each element, the amplitude and phase of its driving signal can be individually modulated, see Figure 1.4, such that by constructive interference of the ultrasound waves a focal

I

Figure 1.3: Profound Sonalleve® MR-HIFU therapy platform and Philips MRI scanner.

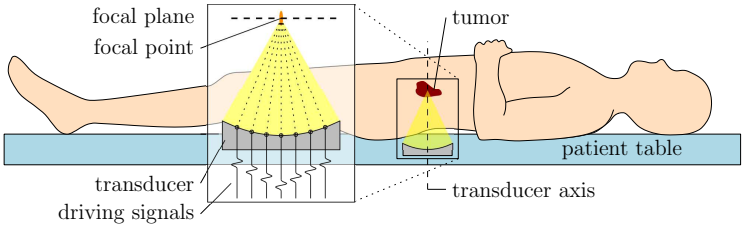

Figure 1.4: Schematic of the transducer embedded in the treatment table, and a HIFU beam and corresponding focal spot inside the tumor via electronic beam steering.

spot of a few millimeters in size can be created and steered with high spatial and temporal resolution, which is commonly termed electronic beam steering [30]. The maximum lateral beam deflection that can be achieved via electronic beam steering, however, is small. For the Sonalleve, for example, the focal spot can be set anywhere within a 16 mm diameter circle centered around the transducer axis [50,51], beyond which deterioration of the focus would occur. To treat larger tumors, the transducer itself can be mechanically repositioned with five degrees of freedom (all except the rotation about the beam axis) by its robotic carrier arm, shown in top-view perspective in Figure 1.5, which is also embedded in the therapy platform. This allows for sequentially heating different parts of the tumor, as shown in [52] for thermal ablation using the Exablate and in [53] for hyperthermia using the Sonalleve, for instance.

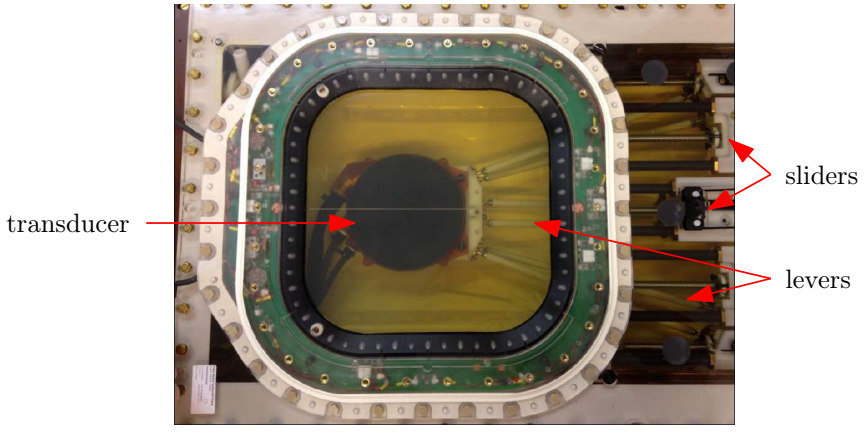

Figure 1.5: Top view of a circular phased-array HIFU transducer, and its mechanical carrier system consisting of multiple levers and sliding actuators.

1.1.5 MR-HIFU feedback control

Accurately inducing and maintaining the desired temperature distribution by MR-HIFU over the course of an entire treatment in a clinical setting is, however, no simple task. Currently, most implementations of MR-HIFU in thermal therapies use largely predetermined sonication plans, possibly extended with basic feedback controllers for some online adjustments. Examples include binary strategies to scale the sonication power and/or the heating intervals in the sonication protocol [50, 51, 53], proportional-integral-derivative (PID) controllers [54–58], or some hybrid form of PID and bang-bang control [59]. A major drawback of such designs is their inability to take into account the tissues's future thermal behavior and the restrictions of the actuator when computing the control inputs, which negatively affects treatment quality and duration. Since the achieved therapeutic benefits of hyperthermia are strongly related to the temperature realized during treatment, see [17, 60–62], there is a dire need for more sophisticated feedback controllers for MR-HIFU, as also argued in [31]. Besides directly improving the temperature control performance, this also enables clinicians to achieve more reliable and reproducible hyperthermia treatments, and may thereby serve as an enabler for more widespread application of and quantitative research into hyperthermia as cancer therapy.

In this respect, model predictive control (MPC) [63–68] offers potential for superior control performance in hyperthermia treatments. That is, in MPC a dynamic model is used to generate predictions of the controlled system's future behavior in response to the available inputs. Based on these predictions, a cost function, which is typically chosen to reflect the treatment objective, is optimized online over a finite future time interval, while explicitly taking into

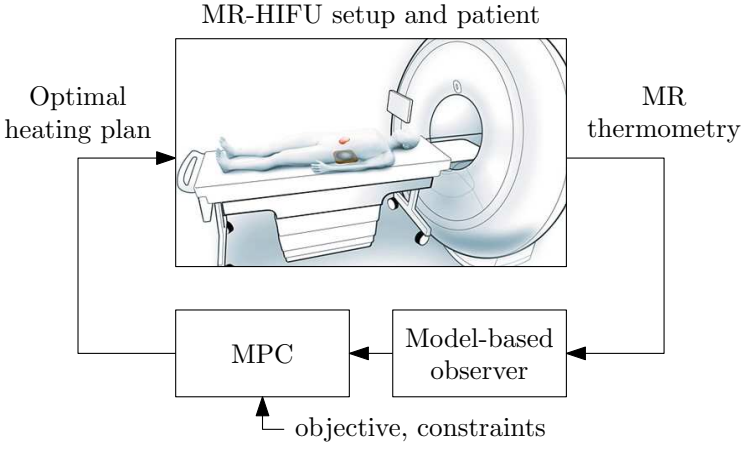

Figure 1.6: Block scheme of the envisioned MPC-based feedback loop for high-quality MR-HIFU hyperthermia treatments.

account the constraints present in, or deliberately imposed on, the system and its inputs. As a result, MPC is a promising control framework for realizing high-quality MR-HIFU hyperthermia treatments. In fact, for thermal ablation this has already been recognized by other researchers. For example, MPC is used for thermal ablation therapy in [69] to minimize treatment time while administering a desired scalar-valued thermal dose in a single point or over a one-dimensional model, and in [70] for minimum-time treatments with a prescribed thermal dose for two-dimensional systems described by heavily simplified models and a fixed spatial heating trajectory. In [71], an MPC scheme is developed that tries to deliver a lethal thermal dose to the tumor, while avoiding damage to healthy tissue. More recently, for MR-HIFU hyperthermia temperature control, in [72] we provided a preliminary MPC scheme using two-dimensional models, which is able to fully exploit the flexibility of the HIFU transducer's electronic beam steering by freely adapting online both the heating location and the acoustic power, resulting in promising performance improvements compared to the binary controller from [53].

1.2 Research objectives

The main objective of this thesis is to provide a significant step forward in achieving high-quality local MR-HIFU hyperthermia treatments by enabling more accurate and uniform heating using MPC-based temperature control strategies. To this end, we envision a feedback control loop as schematically depicted in Figure 1.6, of which from a control perspective the main components are the plant, consisting of the MR-HIFU system and the patient, a model-based observer

that transforms the noise-corrupted MR-based temperature measurements into reliable temperature estimates (and possibly other system parameters), and a high-performance MPC strategy that computes the optimal heating plan. To accomplish the main objective, different subproblems have been identified, and will be described in this section.

1.2.1 Robust performance

The works mentioned in Section 1.1.5 clearly illustrate the potential of MPC for enhancing treatment quality in hyperthermia therapies compared to the existing binary or PID-based control schemes. Still, achieving optimal control performance in a clinical setting is not straightforward. This is due to the large variation of the tissue properties observed in reality, which are typically patient- and tumor-specific, and can additionally vary dramatically as functions of time, space, and temperature [73, 74]. Perhaps the most notable example of these is the drastic increase in blood perfusion as a result of the body's thermoregulatory response [75, 76], of which the thermal dynamics are notoriously difficult to model [74, 77]. Capturing all tissue inhomogeneities and time/temperature-dependencies via extensive pre-treatment model identification is undesirable and impracticable from a clinical perspective, where shorter treatments and avoiding unnecessary heating is preferred. Consequently, model discrepancies are inevitable in practice, and may grow over time, potentially leading to significantly deteriorated control performance and treatment quality. Therefore, the first objective of this thesis can be formulated as:

Objective 1 *Enable the accurate and robust regulation of the temperature in the tumor over a prolonged period of time, despite the presence of considerable and varying model uncertainty.*

1.2.2 Treatment of larger tumors

As previously mentioned in Section 1.1.3, one of the key advantages of using HIFU-based applicators for hyperthermia is their excellent spatial accuracy and fast switching of the focus location by electronic beam steering, which allows for compensating temperature deviations in the millimeter range to achieve uniform tumor temperatures. However, as via electronic beam steering alone the maximum heating range is limited to a few centimeters, for the treatment of larger tumors, it is necessary to also mechanically displace the transducer itself using its robotic carrier system [52, 53], an example of which is shown in Figure 1.5, such that for different transducer positions different parts of the tumor can be heated sequentially, as schematically depicted in Figure 1.7. The mechanical motion of the transducer is slow in comparison to the electronic steering rate, and impairs the MR thermometry's reliability due to intra-scan distortions of the magnetic field. As a consequence, for safety no sonication is allowed during

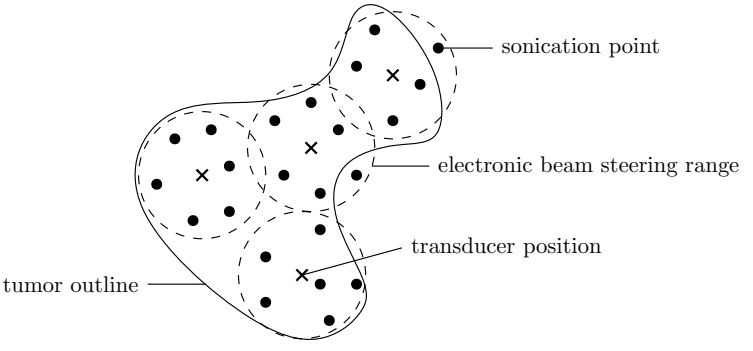

Figure 1.7: Large irregular tumor outline (solid), four transducer positions (×) and their corresponding local heating ranges via electronic beam steering (dashed), and multiple sonication points (•).

transducer repositioning, resulting in considerable heating downtime, which may significantly affect the temperature distribution. Therefore, it is important that, in addition to controlling the sonication, also the transducer displacements are optimally selected online. Correspondingly, the second objective of this thesis reads:

Objective 2 *Enable high treatment quality for large-volume MR-HIFU hyperthermia by simultaneously optimizing online the sonication plan for electronic beam steering and the mechanical transducer trajectory.*

1.2.3 Target-conformal controller design

To enable MR-HIFU feedback controller synthesis, the admissible transducer positions (or position when treating sufficiently small tumors) and the sonication points, i.e., the locations to which the focal spot can potentially be steered, see Figure 1.7, must be chosen before treatment. The sonication points and transducer positions have a key influence on the achievable temperature distribution. However, choosing them in such a manner that the achievable temperature distribution is as close to the desired reference as possible is not straightforward, as this strongly depends on the shape and tissue properties of the specific tumor. At the same time, choosing them optimally is of vital importance, as the realized tumor temperatures have a major effect on the therapeutic effects of hyperthermia [60–62], as also stated in Section 1.1.5. As a solution, one could attempt to simply include an abundance of sonication points and transducer positions in the hope that this provides sufficient controller flexibility, and thereby enables adequate control performance. However, as the sonication powers and the mechanical transducer trajectory must be determined online, this also significantly increases the complexity of the feedback controller. Consequently, such an ap-

proach may jeopardize the real-time feasibility of the feedback setup, especially for more advanced control strategies such as MPC, which offer potential for superior control performance, but in doing so often entail a higher computational burden. Therefore, to optimize the achievable temperature distribution, and at the same time unlock the advantages of more sophisticated high-performance control schemes for MR-HIFU, it is of interest to strategically select the sonication points and transducer positions based on the specific tumor properties, resulting in the following research objective:

Objective 3 *Enable the automated and target-conformal distribution of transducer positions and sonication points for optimal treatment quality.*

1.2.4 Experimental validation

All the research in this thesis is ultimately intended to improve the effectiveness of cancer treatments, such that the patients have a better chance at long-term and disease-free survival. We desire to accomplish this by developing feedback control strategies that achieve more accurate and uniform temperature regulation in local MR-HIFU hyperthermia treatments, which translates to higher treatment quality and, thereby, enhanced therapeutic efficacy and treatment outcome. However, in order to truly enable improved cancer treatments, we must also facilitate the implementation of the novel methods in clinics. Hence, it is essential that the feasibility and performance of the controllers developed in this thesis are demonstrated in practice. Indeed, this provides the first step toward their clinical validation, which in turn may pave the way for their widespread application, thus contributing to improving the treatment of patients suffering from cancer. Therefore, the fourth objective of this thesis is given by:

Objective 4 *Provide an experimental validation of the novel feedback control concepts developed in this thesis on a clinical MR-HIFU system.*

1.3 Contributions

This thesis contains four contributions, addressing the objectives formulated in the previous section. In this section, these contributions are described.

1.3.1 Offset-free model predictive control

The first contribution of this thesis deals with *Objective 1* discussed in Section 1.2.1. In particular, we develop an offset-free MPC setup [78, 79] for MR-HIFU hyperthermia, which incorporates a disturbance estimator that essentially adds integrator functionality to the MPC scheme, such that constant and slowly varying disturbances and plant-model mismatches can be identified during treatment. By including the estimated model discrepancies in the prediction model,

the MPC is made aware of these inaccuracies, and as such is able to compensate for their effects. As result, the steady-state error otherwise resulting from modeling errors can be removed, thereby recovering the optimal performance as achievable in case of zero parameter mismatch. Since the majority of a mild hyperthermia treatment consists of a temperature maintenance phase, in which the tumor temperature is ideally equal to its constant reference value, and model uncertainties are inevitable in practice (as discussed in Section 1.2.1), but typically do not change rapidly in time, offset-free MPC enables significantly improved temperature control in a clinical setting. To demonstrate the controller's robust performance in practice, thereby also addressing *Objective 4*, we implement the offset-free MPC scheme on a clinical setup, and perform validating in-vivo porcine experiments. Correspondingly, the first contribution of this thesis can be summarized as:

Contribution 1 *Development and in-vivo experimental validation of an offset-free MPC scheme for achieving optimal treatment quality in the presence of realistic and varying model uncertainty.*

1.3.2 Switched-actuator systems with setup times

The second contribution of this thesis aims to achieve *Objective 2*. To this end, we first recognize that to simultaneously optimize the mechanical transducer trajectory and the sonication plan, a mixed-integer program (MIP) must be solved [68, 80]. In particular, we interpret the optimization of the mechanical transducer motions as a discrete actuator allocation problem, whereas finding the optimal sonication powers involves continuous variables. However, deriving a model that describes the thermal dynamics in combination with the transducer switching for the purpose of MPC is not a trivial task. That is, the transducer repositioning shifts the applicator's local heating range via electronic beam steering, and as a result effectively determines at which sonication points heating can be applied. We realize that the resulting system can be viewed essentially as a type of switched system [81] of which the input model differs as a function of the discrete transducer position. However, contrary to typical switched systems, which consider instantaneous switches, in large-volume MR-HIFU the switching takes significant time due to the mechanical transducer motions being slow, as discussed in Section 1.2.2. Moreover, recall that during transducer relocation no heating can be applied. To deal with these special features in a manner that is user-friendly and that ensures that the resulting mixed-integer (MI-)MPC [68] algorithms are computationally efficient, a tailored approach is required. To this end, we first introduce the general class of switched-actuator systems with setup times (SAcSSs), which provides the user with a compact and intuitive modeling format. Based on this SAcSS framework, we develop a simple and systematic controller synthesis procedure, which allows for deriving efficient MI-MPC schemes. We demonstrate that this general framework can be used for

high-performance temperature control in large-volume MR-HIFU hyperthermia. In summary, this contribution can be formulated as:

Contribution 2 *Development of a modeling framework tailored to the general class of switched-actuator systems with setup times, which enables the user-friendly modeling and synthesis of computationally efficient mixed-integer MPC algorithms, and is applicable to large-volume MR-HIFU hyperthermia.*

1.3.3 Target-conformal optimal actuator placement

The third contribution addresses *Objective 3*, but also contributes to achieving *Objective 2*. That is, a method is developed that takes into account the tumor geometry and its thermal dynamics to automatically distribute the sonication points and transducer positions for optimal treatment quality. Here, the underlying challenge is to maximize the controller's expected performance, based on the thermal model, and its robustness against unmodeled dynamics and disturbances, while minimizing its online computational complexity by using as few as possible transducer positions and sonication points. However, solving this problem in a single optimization is a complicated task, as the relation between the design parameters and the mentioned objectives is highly nonlinear, and may be difficult to express explicitly. Moreover, the resulting optimization problem contains many decision variables, which are nonlinearly constrained via the finite electronic beam steering range. Therefore, to improve computational tractability and to provide the user with more insightful tuning parameters, a three-step approach is developed that maximizes treatment quality for a limited number of sonication points and transducer positions. In particular, the proposed method allows for optimizing the transducer positions first, based on an intuitive balancing between the achievable steady-state temperature distribution and the heating coverage of the potentially irregularly shaped tumor. Second, using the thermal model, the sonication points that are anticipated to be most crucial for steady-state performance can be selected, which in the final step are complemented by an additional set of points that are maximally distributed throughout the target to enable faster heat-up and improve robust performance in case of spatially localized model uncertainties. In summary, this contribution is given by:

Contribution 3 *Development of a method for the target-conformal optimal placement of transducer positions and sonication points to maximize treatment quality while limiting controller complexity.*

Clearly, such a procedure enables the optimal target-conformal selection of transducer positions and sonication points, thereby accomplishing *Objective 3*. In this thesis, however, this method is also used for the control design of a MI-MPC scheme in a large-volume MR-HIFU hyperthermia case study, and as such also

directly contributes to realizing *Objective 2*. In addition, in this thesis the sonication point selection steps of the actuator placement method are employed in the hierarchical MPC architecture for large-volume MR-HIFU hyperthermia comprising *Contribution 4*, which is discussed next in Section 1.3.4, and thus *Contribution 3* again assists in achieving *Objective 2*.

1.3.4 Hierarchical mixed-integer MPC

Contribution 2 provides an important step toward accomplishing *Objective 2*. However, despite formulating the MI-MPC in a computationally advantageous manner via the SAcSS framework, solving the corresponding optimization problem online within a time interval sufficiently short for real-time feedback control may remain challenging. This is due to two main reasons. First, the models required for describing the tissue's thermal dynamics are typically large. Second, MIPs are inherently complex from a computational perspective. Indeed, such problems are generally classified as NP-hard [80], which means that the worst-case solution time may grow exponentially with the problem size. As a result, for large-volume MR-HIFU hyperthermia only MI-MPCs with limited complexity, either by a sufficiently small input dimension or a short prediction horizon, can typically be implemented in practice. Such limitations, however, impose an upper bound on the achievable control performance, even when using an optimal input selection method such as the target-conformal procedure comprising *Contribution 3*. Therefore, the fourth contribution of this thesis also addresses *Objective 2* by aiming to realize the control performance of highly complex MI-MPC schemes for large-volume MR-HIFU hyperthermia, including adequately many inputs and sufficiently long prediction horizons for optimal treatment quality, while also respecting the solver time constraints for real-time application. This is accomplished using a two-layer hierarchical MPC approach [82]. In particular, we exploit the key observation that the optimal transducer trajectories can be found swiftly using a reduced-complexity MIP in the higher control layer, which is obtained via model order reduction and/or input reduction (where for the latter we employ the target-conformal actuator placement procedure discussed in Section 1.3.3), together with selecting the most appropriate (sparse or condensed) MPC formulation. This subsequently allows for determining the optimal high-dimensional sonication inputs via a computationally efficient quadratic program (QP) in the low control layer. Thereby, the overall computation time of the hierarchical (H)MI-MPC scheme is substantially reduced with respect to the unreduced MI-MPC on which it is based, without sacrificing control performance. We verify the controller's real-time feasibility and performance in phantom experiments on a clinical setup, thereby also addressing *Objective 4*. Hence, the fourth contribution of this thesis is given by:

Contribution 4 *Development and experimental validation of a real-time feasible hierarchical mixed-integer MPC architecture that enables high-quality large-volume MR-HIFU hyperthermia in practice.*

1.4 Organization of the thesis

The remainder of this thesis is organized as follows. Each chapter is based on a journal publication, which relates to a single contribution, and consequently can be read independently from the other chapters. Chapter 2 presents the offset-free MPC design and its in-vivo experimental validation related to *Contribution 1*. In Chapter 3, the SAcSS modeling framework of *Contribution 2* is proposed and applied for setting up a MI-MPC scheme for large-volume MR-HIFU hyperthermia. Then, in Chapter 4 the target-conformal optimal actuator placement method corresponding to *Contribution 3* is given and demonstrated in a large-volume MR-HIFU case study utilizing the SAcSS framework of Chapter 3. This is followed by Chapter 5, which provides *Contribution 4* by discussing the high-performance, yet real-time feasible, HMI-MPC scheme for large-volume MR-HIFU hyperthermia treatments, and the results obtained via phantom experiments on a clinical setup. Although not necessary for the HMI-MPC architecture in general, in Chapter 5 of this thesis we build the HMI-MPC using the SAcSS framework from Chapter 3 and the target-conformal input selection of Chapter 4. Finally, Chapter 6 summarizes the key achievements of this thesis, and discusses a few interesting directions for future research. A schematic overview of the objectives, contributions, and chapters in this thesis is given in Figure 1.8.

I

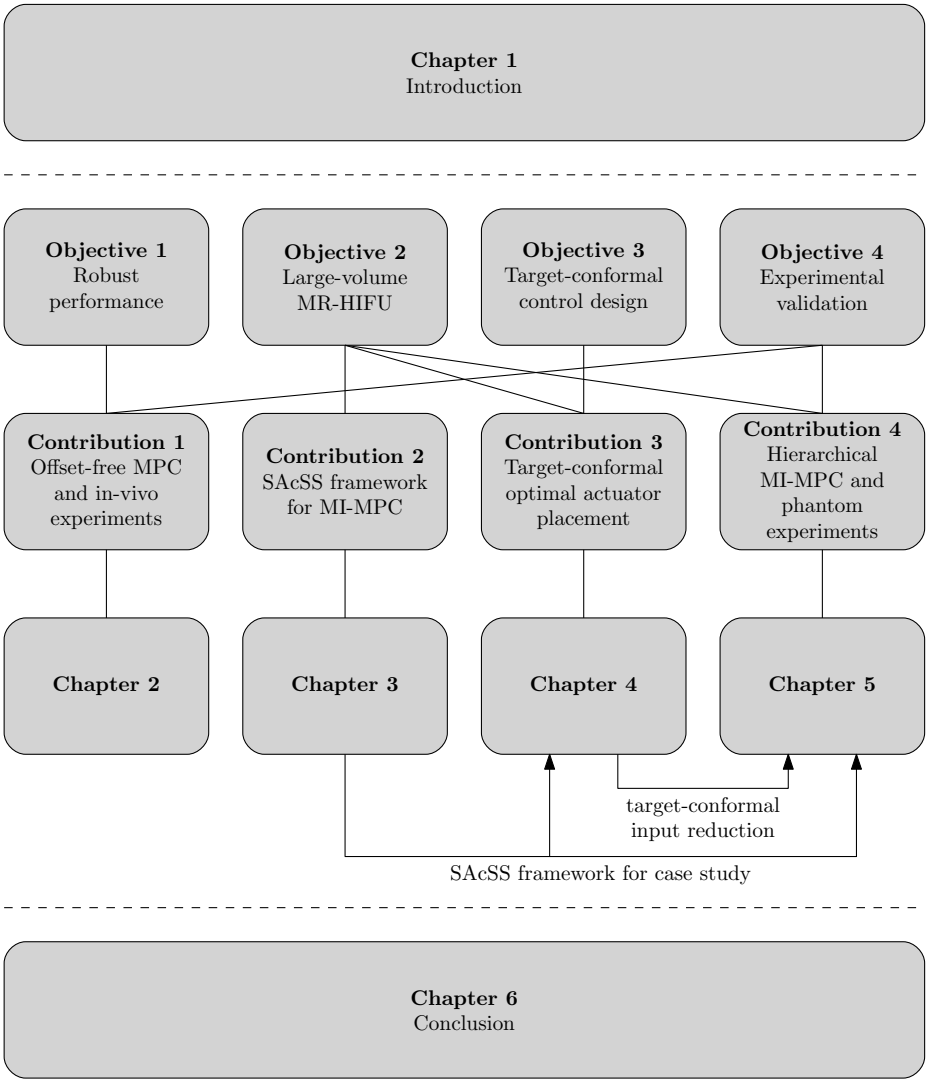

Figure 1.8: Organization of the thesis.

Study the past if you would divine the future.

Confucius

Chapter 2

Offset-Free Model Predictive Temperature Control for Ultrasound-Based Hyperthermia Cancer Treatments

Heating cancer cells over an extended period of time, referred to as hyperthermia, has proven to enhance the effects of chemo- and radiotherapy without inducing additional toxicity or undesirable side effects, and is therefore considered a highly valuable adjuvant therapy in cancer treatment. In this chapter, a model predictive control (MPC) setup is developed for improving performance and robustness in regulating the temperature for magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU) hyperthermia treatments. The proposed control design incorporates a disturbance estimator as encountered in offset-free MPC that is able to remove the steady-state temperature error caused by plant-model mismatch. For the considered healthcare application, such modeling errors are inevitable in practice due to the high variability of tissue properties in patients, some of which even exhibit time- and temperature-dependent behavior due to the body's thermoregulatory response, combined with the fact that extensive model identification is undesirable in the clinic. The controller's performance is demonstrated by means of in-vivo experiments on a porcine thigh muscle using a clinical MR-HIFU treatment setup.

This chapter is based on D.A. Deenen, E. Maljaars, L.C. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Offset-free model predictive temperature control for ultrasound-based hyperthermia cancer treatments," *IEEE Transactions on Control Systems Technology*, in press.

2.1 Introduction

Mild local hyperthermia involves the heating of a specific target volume inside the body, typically containing the tumor, to temperatures of 39-45 °C for up to about 90 minutes, while preserving non-elevated temperatures in healthy tissue. Clinical trials have provided ample evidence that hyperthermia is a potent adjuvant therapy for cancer treatment, see, for example, [3, 16, 22], and the references therein. It locally sensitizes the treated tissue to the effects of chemo- and radiotherapy, while leaving the untreated (healthy) tissue unaffected. Consequently, similar or improved treatment effectiveness can be achieved using lower doses of radiation and drugs, thereby reducing the severity of the negative side effects typically associated with cancer treatment [15, 19]. Moreover, using temperature-sensitive liposomes, it enables localized heat-mediated delivery of anticancer drugs, allowing for a further reduction of the required systemic drug concentrations [25, 26]. Most importantly, hyperthermia itself is non-toxic, and therefore introduces no additional toxicity-related side effects, making it highly appealing for clinical application to increase treatment success rates and improve quality of life for the patients.

A particularly well-suited technology for hyperthermia is magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU). This entails the combined use of powerful and millimeter-accurate heating by means of ultrasound waves, and real-time volumetric thermometry using an MRI scanner, see [30, 83]. Using an external HIFU applicator, this technology allows for a completely noninvasive treatment, which is highly desirable for patient comfort, and eliminates the need for post-interventional wound care. 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, resulting in a need for the development of adequate *feedback* controllers. This is particularly true for mild hyperthermia, since its beneficial effects have been found to be strongly correlated with the tissue temperatures truly achieved during treatment, see [24, 60, 61].

Currently, most implementations of MR-HIFU for temperature control in thermal therapies use predetermined sonication plans, possibly extended with simple feedback controllers that make (minor) online adjustments. Examples include binary strategies to scale the sonication power and/or length of the heating intervals in the sonication protocol [51, 53], ad-hoc proportional-integral-derivative (PID)-based methods [57, 58], or some hybrid form of PID and bang-bang control [59]. In addition to not exploiting the full potential control freedom offered by the setup, the major drawback of such designs is their inability to take the body's future thermal behavior and the restrictive actuator constraints into account when computing the control inputs, which negatively affects treatment quality and duration.

In this respect, we believe that model predictive control (MPC) offers superior potential for temperature control in hyperthermia treatments. This has already

been recognized by other researchers, see, for example, [69] for the control of a scalar thermal dose parameter based on a single-point or one-dimensional model, and [70] (using heavily simplified models and a fixed sonication trajectory) and [71] for thermal dose control in two-dimensional systems. More recently, in [72] we have developed a preliminary MPC design for temperature control in two-dimensional systems with online adaptation of the heating location and power.

These works illustrate clearly the potential of MPC for hyperthermia. However, for accurate temperature control in mild hyperthermia treatments in a clinical setting, a more sophisticated MPC design is required. One important reason for this is the large variation observed in the tissue properties in reality, which are typically patient- and tumor-specific, and can additionally vary dramatically as functions of time, space, and temperature [74, 75, 77]. Extensive pre-treatment model identification is undesirable from a clinical perspective, since, besides improving treatment quality, the goal is to also reduce treatment time and avoid unnecessary heating. Consequently, model discrepancies are typically inevitable in practice, and may result in significantly deteriorated control performance. It is the objective of this work to provide a significant step forward in the development of an MPC setup that is able to robustly regulate the temperature in the tumor over a prolonged period of time, despite the presence of considerable and possibly varying model uncertainty.

To this end, as the first main contribution of this chapter, we develop a temperature controller for MR-HIFU hyperthermia based on an offset-free MPC approach. Using a disturbance model, the effects of constant and slowly varying plant-model mismatch can be captured [78, 79]. An observer is employed to identify the corresponding disturbance, whilst simultaneously providing temperature estimates that improve upon the noise-corrupted MR thermometry measurements. This enables the MPC to compensate for the model-error-induced temperature offset in steady state, recovering the optimal performance as achievable in case of zero plant-model mismatch, thereby allowing for significantly enhanced temperature control and treatment quality.

In [84], we have presented a preliminary version of this control setup. Using simulations it was shown that the proposed solution is able to identify and compensate for model discrepancies. In this work, we present results obtained in porcine in-vivo experiments, demonstrating the developed feedback setup's ability of achieving desired steady-state heating despite the presence of realistic plant-model mismatch in a clinical setting. Furthermore, to improve modeling accuracy without increasing computation time, in this chapter we extend the modeling procedure for the controller and observer separately, additionally accounting for computation/communication delays to reduce modeling error in the temporal discretization of the state-space models. Finally, more details are provided regarding the design considerations of the cost function, which has been changed with respect to [84].

The remainder of this chapter is organized as follows. First, Section 2.2

Figure 2.1: Philips MRI scanner and Profound Sonalleve® MR-HIFU therapy platform.

discusses the MR-HIFU hyperthermia setup and treatment, motivating the development of offset-free MPC for this application. In Section 2.3, we describe the thermal model and its spatial discretization, which is used in the observer and controller design presented in Section 2.4. In Section 2.5, the results of the in-vivo experiments are discussed to illustrate the algorithm's potential for clinical application. Finally, Section 2.6 summarizes the key achievements and corresponding observations.

2.2 System and treatment description

In this section, we introduce the hyperthermia treatment and setup, and motivate the use of MPC in this application.

2.2.1 MR-HIFU hyperthermia setup

Although the MPC setup we propose is generic in nature, in this work it is designed for a clinical MR-HIFU system consisting of a Profound Sonalleve®, which is shown in Figure 2.1, and a Philips 3T Achieva®. The former is a dedicated trolley-tabletop in which an MR-compatible HIFU transducer is integrated, and the latter is an MRI scanner which we use for noninvasive near-real-time thermometry. This system is already being used in clinics to non-invasively treat uterine fibroids and for incision- and radiation-free palliative treatment of pain associated with bone metastases. A custom software layer based on [72] has been developed to connect the MPC algorithms designed in this work to the MR-HIFU setup.

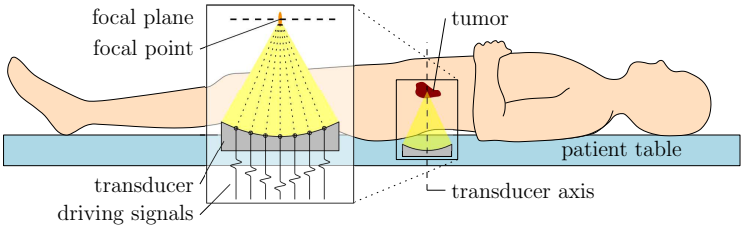

Figure 2.2: Schematic of a HIFU beam into the focal plane in the tumor, with the focal point by electronic beam steering.

2.2.2 HIFU applicator

This system uses a phased-array HIFU transducer to generate the ultrasound waves. It consists of 256 elements of which the phases and amplitudes can individually be chosen such that by interference a focal spot is created, thereby enabling powerful and millimeter-accurate heat delivery to internal tissues as depicted in Figure 2.2. By modulating these settings, which is referred to as electronic beam steering, the focal spot can be repositioned to up to 20 distinct locations per second within a 16 mm diameter circle around the transducer axis. By moving the focal spot through the tumor area, all of the cancerous tissue can be heated. Treating larger regions additionally requires mechanically displacing the transducer, see, e.g., [53]. In this work, however, we limit ourselves to the development of a controller for the treatment of small tumors, and will therefore consider electronic beam steering only.

The focal spot has an axial length of approximately 7 mm, which is significantly larger than its narrow radial width of about 2 mm, as depicted in Figure 2.3. Here, it is also shown how for different focus locations in the focal plane, the acoustic beams (each individually indicated in light gray) significantly overlap each other before and behind the focal plane in the direction of the beam axis (resulting in darker gray in the figure), see also [57, 85]. Consequently, the resulting temperature distribution is fairly homogeneous in axial direction near the focal plane, whereas the in-plane temperature gradients may be large [53]. Since in this work we consider only small treatment volumes, the temperature map in the focal plane characterizes the tissue temperature in the entire treatment volume, allowing us to design our MPC setup for temperature control only in the focal plane using two-dimensional models. Note that, additionally, all (constant and slowly varying) neglected out-of-plane effects, e.g., heat conduction in case of a nonzero axial temperature gradient, will be identified and compensated for by the offset-free algorithm.

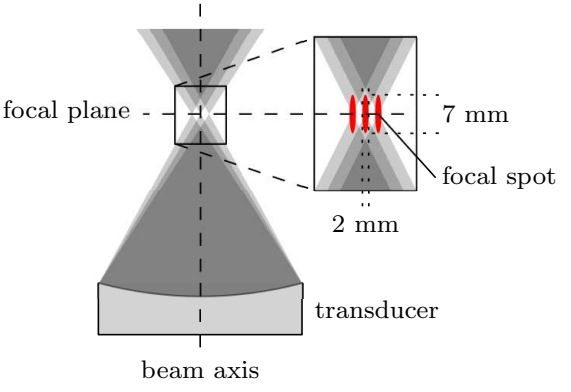

Figure 2.3: Schematic of three HIFU beams (light gray) with different focus location in the focal plane to illustrate the beam overlap (darker gray) outside the focal plane in axial direction. The focal spot (red) and its dimensions are depicted in detail.

2.2.3 Motivation for offset-free MPC in hyperthermia

The key to successful application of mild local hyperthermia is to maintain a steady temperature elevation above 41 °C in the region of interest (ROI), typically the tumor and some adjacent tissue, during the entire treatment. The temperature sweet spot for treatment quality is at 42 °C. This is required to fully benefit from the desired heat-induced effects such as increased blood perfusion, aiding drug delivery to otherwise poorly perfused parts of the tumor when combined with chemotherapy [16, 73], or the inhibition of DNA repair mechanisms during the crucial time window after radiation therapy [15, 24]. On the other hand, overheating can also be detrimental for successful treatment, since some of these mechanisms exhibit reversal effects at higher temperatures (above 43 °C) [73, 76]. Furthermore, temperature elevations above 41 °C outside the ROI must be avoided to prevent sensitization of (or damage to) healthy tissue.

Compared to existing binary or PID-based strategies applied in MR-HIFU hyperthermia, we believe that MPC-based approaches can deliver superior closed-loop properties, as they are able to achieve faster and more uniform heating by exploiting beneficial (future) behavior, e.g., heat transfer by conduction, and can explicitly take into account actuator constraints such as the inability to actively remove heat from inside the body using HIFU. Unfortunately, model-based strategies are also inherently accompanied by the possibility for modeling errors.

As discussed in [74], accurate thermal and thermoregulatory modeling for hyperthermia treatments is especially difficult due to the high variability of the spatially distributed and time/temperature-varying tissue properties [75, 77]. In

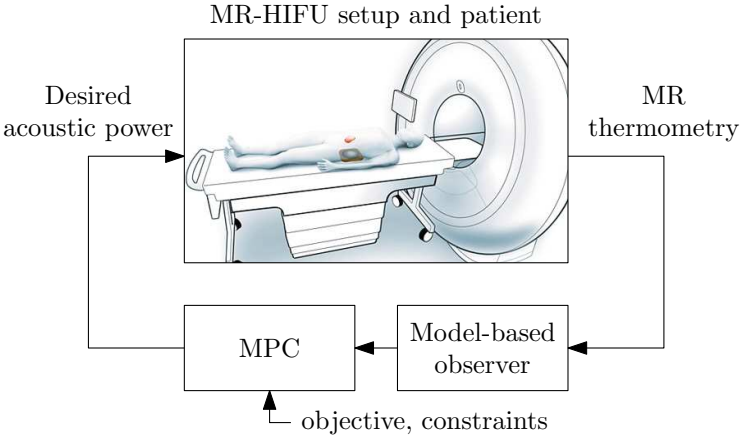

Figure 2.4: Block scheme of the envisioned fully automated MR-HIFU feedback loop for patient treatment.

addition, attempting to capture all these effects using extensive personalized pre-treatment model identification is highly undesirable in the clinic, where, besides improving treatment quality, the aim is also to reduce treatment times and to improve safety by avoiding unnecessary (over)heating. As a result, plant-model mismatch is typically inevitable in practice. If inadequately accounted for, however, such modeling errors may result in insufficient heating of the tumor or in the overheating of healthy tissue, which would significantly deteriorate treatment quality and safety.

Given this situation, we propose an MPC setup novel for hyperthermia inspired by offset-free MPC [78, 79]. Figure 2.4 depicts the resulting feedback scheme, where the observer provides the temperature and disturbance estimates from the noninvasively sampled MR thermometry data, such that the MPC scheme can compute the optimal power distribution to be generated by the extracorporeal HIFU actuator.

2.3 Thermal model

In this section, a model of the body's thermal response will be discussed and spatially discretized, yielding the state-space model from which the observer and controller models can be derived.

2.3.1 Bioheat model

The tissue's thermal behavior is modeled using Pennes' bioheat equation, see [86], given by

$$\rho(r)c(r)\frac{\partial T(r,t)}{\partial t} = \nabla(\kappa(r)\nabla T(r,t)) + Q_a(r,t) - w_b(r,t,T)c_b(T(r,t)-T_b), \quad (2.1)$$

where $T : \Omega \times \mathbb{R}_{\geq 0} \to \mathbb{R}$ is the temperature profile. In particular, $T(r,t)$ 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. Furthermore, $\rho : \Omega \to \mathbb{R}_{>0}$ denotes the tissue's volumetric mass density, $c : \Omega \to \mathbb{R}_{>0}$ the specific heat capacity, $\kappa : \Omega \to \mathbb{R}_{>0}$ the thermal conductivity, and $w_b : \Omega \times \mathbb{R}_{>0} \times \mathbb{R} \to \mathbb{R}_{\geq 0}$ the blood perfusion coefficient. The blood's specific heat capacity and temperature are given by $c_b \in \mathbb{R}_{>0}$ and $T_b \in \mathbb{R}$, respectively. Note that these tissue properties are typically spatially varying in reality due to tissue heterogeneity or blood vessels, for example. In addition, the blood flow coefficient w_b can also depend nonlinearly on time and temperature due to the body's thermoregulatory response, which will be patient- and tumor-specific [74,75,77,87]. Although low-power pre-treatment test sonications may serve as a practical solution to obtain initial estimates of the tissue properties, see [72], fully accurate models are nearly impossible to obtain due to the unavailability of extensive model identification for the reasons previously mentioned. We therefore propose to assume spatially homogeneous tissue properties that remain constant in time, reducing (2.1) to

$$\rho c\frac{\partial T(r,t)}{\partial t} = \kappa\nabla^2 T(r,t) + Q_a(r,t) - w_b c_b(T(r,t)-T_b), \quad (2.2)$$

and design a feedback controller that is able to identify and compensate for the resulting model mismatch. Note, however, that our MPC setup can be directly be applied to a fully inhomogeneous tissue model as well, in case such a model would be available.

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

$$Q_a(r,t) = F(r,t)P(t), \quad r \in \Omega,\ t \in \mathbb{R}_{\geq 0}. \quad (2.3)$$

In reality, for a given focus location $r_f(t)$ at some time t, the intensity $F(r,t)$ also depends on spatially varying factors such as the acoustic properties of the (intermediate) tissue. In our approach, 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

$$F(r,t) = \frac{\alpha}{2\pi\sigma_f^2}\exp\left(-\frac{\|r - r_f(t)\|^2}{2\sigma_f^2}\right), \quad (2.4)$$

with $r \in \Omega$, $t \in \mathbb{R}_{\geq 0}$, and acoustic energy absorption coefficient $\alpha \in \mathbb{R}_{>0}$.

2.3.2 Spatially discretized state-space model

Recall that we only consider the two-dimensional focal plane and assume zero out-of-plane interaction, justified by the temperature distribution being fairly homogeneous in axial direction as discussed in Section 2.2.2. Spatial discretization of (2.2)-(2.4) is done using the central difference scheme, which in [72] has been concluded to best combine model simplicity with predictive power for the considered system, on a 44×44 grid with a voxel size of 1.85×1.85 mm^2. As boundary condition, we prescribe the outward heat flux corresponding to a fixed temperature T_b just outside our grid's edges, and we shift the origin of (2.2) to the blood temperature T_b. The resulting continuous-time state-space dynamics are given by

$$\dot{x}(t) = A_c x(t) + B_c u(t), \tag{2.5a}$$

where the states $x(t) \in \mathbb{R}^{n_x}$ with $n_x = 44^2 = 1936$ represent the voxel's temperature elevations with respect to T_b at time $t \in \mathbb{R}_{\geq 0}$. Due to the central difference scheme, A_c is a sparse matrix containing at most five nonzero elements per row/column. For the input, we choose to allow sonication at the voxel centers within the ROI, which will be referred to as sonication points. The input $u(t) \in \mathbb{R}^{n_u}$, $n_u = 60$, represents the applied acoustic power at the sonication points, which may contain at most one nonzero element at any time t, corresponding to the sonication point that coincides with the focus location $r_f(t)$ using electronic beam steering. Each individual column of B_c therefore captures the rate of increase of the voxel temperatures when applying unit sonication power at the corresponding individual sonication point.

MR thermometry does not yield continuous measurement, but instead induces a sample time of $T_s = 3.7$ s when employing the acquisition protocol used in this work. The spatial discretization in (2.5a) is specifically chosen such that the voxel centers coincide with the points measured by MR thermometry. Hence, the measurements can be modeled by

$$y_k = x(t_k) + v_k, \tag{2.5b}$$

with $k \in \mathbb{N}$ connecting to real time $t_k = kT_s$, and $y_k = y(t_k) \in \mathbb{R}^{n_x}$ representing the measured voxel temperatures corrupted by $v_k \in \mathbb{R}^{n_x}$, where in the monitored area of interest the latter consists mainly of sensor noise, which has been experimentally determined to typically be zero-mean and Gaussian distributed.

2.4 Controller design

In this section, we first introduce the disturbance model and discuss the sampled-data setup, based on which we then derive the discrete-time state-space models used by the controller and observer. The cost function is designed using the control objectives previously described in Section 2.2.3, which in turn is used to formulate the constrained optimization problem for MPC.

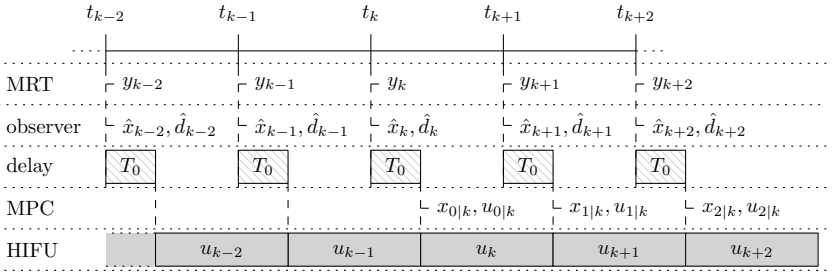

Figure 2.5: Sampled-data timing with nonzero computation/communication delay, and shifted MPC prediction instants synchronized with the control input application times.

2.4.1 Disturbance model

To incorporate integral action in the MPC setup, a disturbance model together with a disturbance estimator can be used [78, 79]. The disturbance estimate is provided to the controller, enabling it to compensate for the steady-state offset induced by model mismatch or slowly varying disturbances. To facilitate the derivation of the required discrete-time models, let us augment the continuous-time dynamics (2.5a) to

$$\dot{x}(t) = A_c x(t) + B_c u(t) + d(t), \tag{2.6a}$$

$$\dot{d}(t) = 0, \tag{2.6b}$$

such that it includes the additive disturbance $d(t) \in \mathbb{R}^{n_x}$, which is assumed to be constant in time. This particular choice of disturbance model allows for the construction of a stable discrete-time estimator, as will be shown in Section 2.4.4, which in turn enables the compensation of the steady-state offset resulting from model error.

2.4.2 Sampled-data setup

We design a discrete-time controller for a continuous-time system. The sampled-data workflow is schematically visualized in Figure 2.5. Here, the sampled output y_k is given by (2.5b). The discrete-time observer variables $\hat{x}_k, \hat{d}_k \in \mathbb{R}^{n_x}$ represent the estimates of the state $x(t_k)$ and disturbance $d(t_k)$, respectively, which become available at (approximately) the same time as the measurement y_k due to negligible observer computation time. As indicated in the figure, for the considered system there is a delay of $T_0 \approx 1.1$ s between obtaining a temperature map and applying the corresponding new heating plan, which is caused mainly by the numerical optimization in MPC and a significant communication and data processing overhead between the MPC and the dedicated MR-HIFU software.

As discrete-time input, we define $u_k \in \mathbb{R}^{n_u}$, which represents the applied sonication power averaged over one sampling interval at each of the sonication points. This is achieved by using electronic beam steering to rapidly switch the focal spot $r_f(t)$ between all sonication points at which heating is desired in such a manner that, on average over one sampling interval, the power as requested by u_k is injected. Recall that the input $u(t)$ in (2.6) can contain at most only one nonzero element. However, due to the high steering frequency with respect to the slow sampling and system dynamics, the input in (2.6) can be well approximated using a zero-order hold (ZOH) given by $u(t) = u_k$ for $t_k + T_0 \leq t < t_{k+1} + T_0$ (note the shift of the interval by T_0 due to the computation/communication delay), where u_k may contain more than one nonzero element.

Finally, in temporally discretizing the continuous-time dynamics (2.6) to obtain the controller state-space models, we account for the delay T_0 without increasing the computational complexity of the numerical optimization in MPC. That is, we first compute the prediction sequence's initial conditions $x_{0|k}, d_{0|k} \in \mathbb{R}^{n_x}$, representing the predicted state and disturbance at time $t_k + T_0$, and then synchronize the prediction instants $t_{i|k}$, $i, k \in \mathbb{N}$, in the MPC horizon with the (future) time instants $t_{k+i} = t_k + T_0 + iT_s$ at which the control inputs will be updated, as indicated in Figure 2.5. As a result, a state update using the prediction model derived in the next subsection requires only a single input value, as opposed to two in case the predictions and controls were not synchronized.

2.4.3 Prediction model

As discussed in the previous subsection, the discrete-time state-space prediction model is then given by

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k} + B_d d_{i|k}, \tag{2.7a}$$
$$d_{i+1|k} = d_{i|k}, \tag{2.7b}$$
$$y_{i|k} = x_{i|k}, \tag{2.7c}$$

where $x_{i|k}, d_{i|k}, y_{i|k} \in \mathbb{R}^{n_x}$ and $u_{i|k} \in \mathbb{R}^{n_u}$ denote the predicted states, disturbance, outputs, and control inputs, respectively, at $i \in \mathbb{N}$ time steps ahead of the prediction sequence's starting time $k \in \mathbb{N}$. Note that we assume the disturbance to be constant over the prediction horizon, and that we do not incorporate measurement noise in the prediction model since (approximately) $\mathbb{E}(v_k) = 0$. Here we use the forward Euler method to temporally discretize (2.6), i.e., $A = I_{n_x} + T_s A_c$, $B = T_s B_c$, and $B_d = T_s I_{n_x}$, as this fully preserves the sparseness present in the continuous-time state-space matrices, which is crucial to reduce computation time of the numerical optimization step in MPC, while also yielding sufficient predictive accuracy. In particular, for the considered system, this yields a (numerically) stable approximation for $0 < T_s < -2/\underline{\lambda}(A_c) = 8.35$ s (where we use

the fact that all eigenvalues of A_c are real and negative, and where $\underline{\lambda}(A_c) < 0$ denotes the most negative eigenvalue of A_c), and when simulating entire treatments given typical input trajectories for $T_s = 3.7$ s it induces an approximation error in the order of 0.01 °C in steady state and at most around 0.1 °C during the transient. Moreover, note that since we have synchronized the time instants of input application and state prediction, (2.7a) indeed contains only a single control input term, which is beneficial for computation time.

In addition to (2.7), we derive the model that computes the prediction's initial conditions as shown in Figure 2.5, which is given by

$$x_{0|k} = A_0\hat{x}_k + B_0u_{k-1} + B_{d,0}\hat{d}_k, \tag{2.8a}$$

$$d_{0|k} = \hat{d}_k. \tag{2.8b}$$

Since these need only be computed once before numerical optimization, matrix sparseness is of negligible importance. To avoid numerical approximation errors, we obtain the matrices by exact discretization using $A_0 = e^{T_0A_c}$, $B_0 = \int_0^{T_0} e^{A_c\tau}d\tau B_c$, and $B_{d,0} = \int_0^{T_0} e^{A_c\tau}d\tau$, of which the latter two expressions exploit ZOH on the corresponding (disturbance) inputs.

2.4.4 State and disturbance estimation

We use an observer to obtain temperature estimates in which the effects of measurement noise are reduced with respect to the MR thermometry readings, and to identify the disturbance that captures the offset induced by plant-model mismatch. The state and disturbance estimator is given by

$$\hat{x}_k = \hat{x}_k^- + L_x(y_k - \hat{y}_k^-), \tag{2.9a}$$

$$\hat{d}_k = \hat{d}_{k-1} + L_d(y_k - \hat{y}_k^-), \tag{2.9b}$$

where

$$\hat{y}_k^- = \hat{x}_k^- = \hat{A}\hat{x}_{k-1} + \hat{B}_1u_{k-1} + \hat{B}_2u_{k-2} + \hat{B}_d\hat{d}_{k-1} \tag{2.9c}$$

denote the model-based output and state estimates at time k before applying the measurement-based correction to obtain $\hat{x}_k$. For the same reasons as for (2.8), the observer model's discrete-time matrices are derived by exact discretization of (2.6) as given by $\hat{A} = e^{T_sA_c}$, $\hat{B}_1 = \int_0^{T_s-T_0} e^{A_c\tau}d\tau B_c$, $\hat{B}_2 = \int_{T_s-T_0}^{T_s} e^{A_c\tau}d\tau B_c$, and $\hat{B}_d = \int_0^{T_s} e^{A_c\tau}d\tau$, where the expressions for $\hat{B}_1$ and $\hat{B}_2$ again follow from the fact that we use ZOH on the inputs, see also [88].

We use Proposition 1 of [78] to verify the observability of the augmented discrete-time state-space model on which the observer is based (i.e., (2.9) with $L_x = L_d = 0$), which is derived using the disturbance model (2.6). That is, we first verify that the nominal system (excluding the disturbance) is observable

due to the full (noisy) state measurement (2.5b). Second, we evaluate the rank condition in Proposition 1 of [78], for which in our case it holds that

$$\operatorname{rank} \begin{bmatrix} I - \hat{A} & -\hat{B}_d \\ I & 0 \end{bmatrix} = 2n, \tag{2.10}$$

from which we conclude the augmented model to be observable. Consequently, an asymptotically stable estimator is known to exist for 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, see Theorem 1 of [89].

By observability of the augmented model, L_x and L_d can in principle be used for arbitrary pole placement of the estimation error system resulting from (2.9). However, determining the observer gains using exact pole placement methods or computing them as the Kalman gains typically results in a prohibitively large computational burden due to the large augmented state dimension $2n_x = 3872$. Therefore, we set L_x to be a diagonal matrix instead, which, due to y_k in (2.5b) being the noisy full state measurement, is a practical solution with good performance for reducing the propagation of noise into the state estimate $\hat{x}_k$. Similarly, we set L_d diagonal to achieve low-pass filtering of the measured plant-model mismatch, which is possible since $\hat{d}_k$ is chosen to contain an estimate of the constant and low-frequency mismatch for each state element individually. Using a polyacrylamide tissue-mimicking phantom as in [90], but with the ink replaced by water, the matrices have been tuned experimentally, resulting in

$$L_x = 0.25 I_{n_x}, \quad L_d = 0.01 I_{n_x}, \tag{2.11}$$

which yield desirable estimator behavior with stable closed-loop poles.

2.4.5 Cost function

The temperature objectives have been discussed in Section 2.2.3 and are schematically depicted in Figure 2.6 in cross-section perspective. The concentric circular regions $\mathcal{R}$ and $\mathcal{S}$ inside the patient domain Ω, for which it holds that $\mathcal{R} \subset \mathcal{S} \subset \Omega$, denote the ROI and the region outside of which healthy tissue must be safeguarded from overheating, respectively. $\underline{T} : \mathcal{R} \to \mathbb{R}$ and $\overline{T} : \Omega \to \mathbb{R}$ represent the location-dependent (desired) lower and upper temperature bounds, respectively, defining the desired temperature range (green). Under- and overheating occurs for temperatures within the (light) blue and (light) red areas in the figure, respectively. $\underline{T}$ is nonzero and defined 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 lower value on $\Omega \setminus \mathcal{S}$ to protect healthy tissue. Within the aforementioned temperature range,

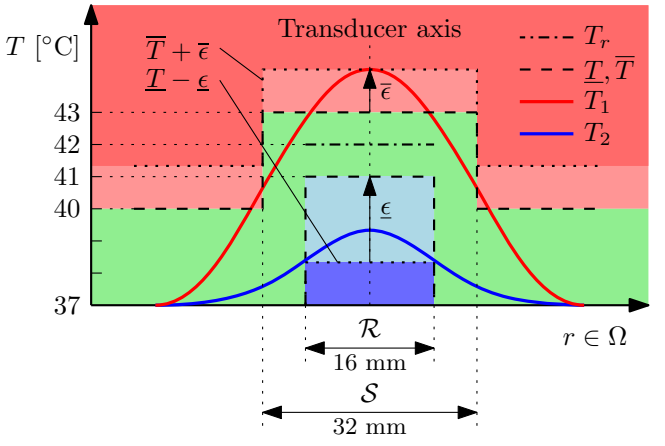

Figure 2.6: Schematic cross section of the temperature objectives corresponding to the circular regions $\mathcal{R}$ and $\mathcal{S}$ centered at the transducer axis. The maximum violations $\underline{\epsilon}$ and $\overline{\epsilon}$ are shown for some over- and underheated temperature distributions such that $T_1 \leq \overline{T} + \overline{\epsilon}$ and $T_2 \geq \underline{T} - \underline{\epsilon}$.

optimal treatment efficacy is achieved when a flat temperature distribution is realized in the ROI, which is described by the reference temperature $T_r : \mathcal{R} \to \mathbb{R}$.

To formulate these objectives mathematically, let us denote the $n_z < n_x$ elements of x_k corresponding to a point within $\mathcal{R}$ by $z_k = Hx_k \in \mathbb{R}^{n_z}$, where $H \in \{0,1\}^{n_z \times n_x}$ is a matrix with one 1 in each row (and at most one 1 per column). Furthermore, let $z_r \in \mathbb{R}^{n_z}$, $\underline{z} \in \mathbb{R}^{n_z}$, and $\overline{x} \in \mathbb{R}^{n_x}$ denote the voxel-wise temperature elevation reference, lower bound, and upper bound corresponding to the values of T_r, $\underline{T}$, and $\overline{T}$ on the discrete voxel locations, respectively. This allows for the ROI voxel temperature deviations with respect to the reference to be given by $z_k - z_r$. The maximum violations of the temperature bounds are defined by $\overline{\epsilon}_k = \overline{\epsilon}(x_k) = \|\max\{x_k - \overline{x}, 0_{n_x}\}\|_\infty \in \mathbb{R}_{\geq 0}$ and $\underline{\epsilon}_k = \underline{\epsilon}(z_k) = \|\max\{\underline{z} - z_k, 0_{n_z}\}\|_\infty \in \mathbb{R}_{\geq 0}$, where 0_n denotes a zero-vector of length n and the maximum operator is used element-wise, and are collected in $\epsilon_k = [\underline{\epsilon}_k \ \overline{\epsilon}_k]^\top \in \mathbb{R}^{n_\epsilon}_{\geq 0}$, $n_\epsilon = 2$.

The cost function is given by

$$V_N(\boldsymbol{z}_k, \boldsymbol{\epsilon}_k) = \sum_{i=0}^{N} \ell(z_{i|k}, \epsilon_{i|k}) \tag{2.12}$$

where $\boldsymbol{z}_k = (z_{0|k}, \ldots, z_{N|k})$ is the sequence of predicted performance variables $z_{i|k} = Hx_{i|k}$ at time instant $k \in \mathbb{N}$, and similarly $\boldsymbol{\epsilon}_k = (\epsilon_{0|k}, \ldots, \epsilon_{N|k})$ the prediction sequence of the slack variable vector $\epsilon_{i|k} = [\underline{\epsilon}_{i|k} \ \overline{\epsilon}_{i|k}]^\top$. The stage cost is given by

$$\ell(z_{i|k}, \epsilon_{i|k}) = (z_{i|k} - z_r)^\top Q (z_{i|k} - z_r) + f_\epsilon^\top \epsilon_{i|k}, \tag{2.13}$$

where Q is a positive definite weighting matrix and $f_\epsilon \in \mathbb{R}^2_{>0}$. The critical importance of achieving temperatures within the desired range $\underline{T} \leq T \leq \overline{T}$ is expressed by using a linear penalty on ϵ, such that its contribution to the cost does not vanish quadratically when approaching zero. Additionally, the weighting in f_ϵ is typically chosen relatively high compared to the weighting in Q, such that any (significant) violation of the temperature range will be the dominant component of the cost. When the soft constraints are (nearly) satisfied, i.e., when $\epsilon_k \approx 0$, the contribution of the quadratic term will be dominant, enforcing temperature tracking of the voxels within the ROI to the optimal treatment temperature T_r.

Although it is a natural choice to use slack variables and soft constraints to incorporate the desire for temperatures within the range $\underline{T} \leq T \leq \overline{T}$ in the cost function, the manner in which these are implemented deserves considerable attention. We choose to penalize the ROI voxels' magnitude-wise largest violations using the ∞-norm, as opposed to introducing a slack variable for each voxel individually and using a penalty on the 1-norm of all slack variables. Compared to the latter, the former has two main advantages. First, it requires only two slack variables per predicted time step (one for the upper bound and one for the lower bound) compared to $2n_x$ slack variables per time step (n_x per bound), thereby introducing less complexity in the MPC optimization problem. Second, it yields superior temperature homogeneity, which is desirable from a clinical point of view. Although this effect is important during the entire treatment, it is especially visible during the initial heat-up phase, and hence we will use this interval to illustrate the behavioral advantage of the ∞-norm approach. To this end, compare the temperature cross sections in Figure 2.7a obtained through simulation with $v_k = 0$ using the ∞-norm to those in Figure 2.7b corresponding to using the 1-norm. In the former, relatively more heating is applied towards the edge of the ROI, where the temperature is naturally lowest due to outward heat diffusion, leading to a more uniform temperature distribution within the ROI (gray area), whereas in Figure 2.7b a much larger difference exists between the maximum ROI voxel temperatures (■) and the minima at the ROI edge (●). Additionally, in Figure 2.7c we show the temperature profiles obtained when omitting the slack variable penalty altogether, i.e., when using (2.13) with $f_\epsilon = 0$. Besides lacking the ability to explicitly counteract temperatures outside the desired range, using $f_\epsilon = 0$ results in similar heat-up behavior as with the 1-norm approach. Finally, in Figure 2.7d, it is clear to see that the resulting temperature range inside the ROI is significantly narrower during heat-up when including the cost on the uniform lower bound following from the ∞-norm penalty (and even also slightly better directly thereafter), from which we conclude that using the ∞-norm penalty is indeed the most favorable of the considered options.

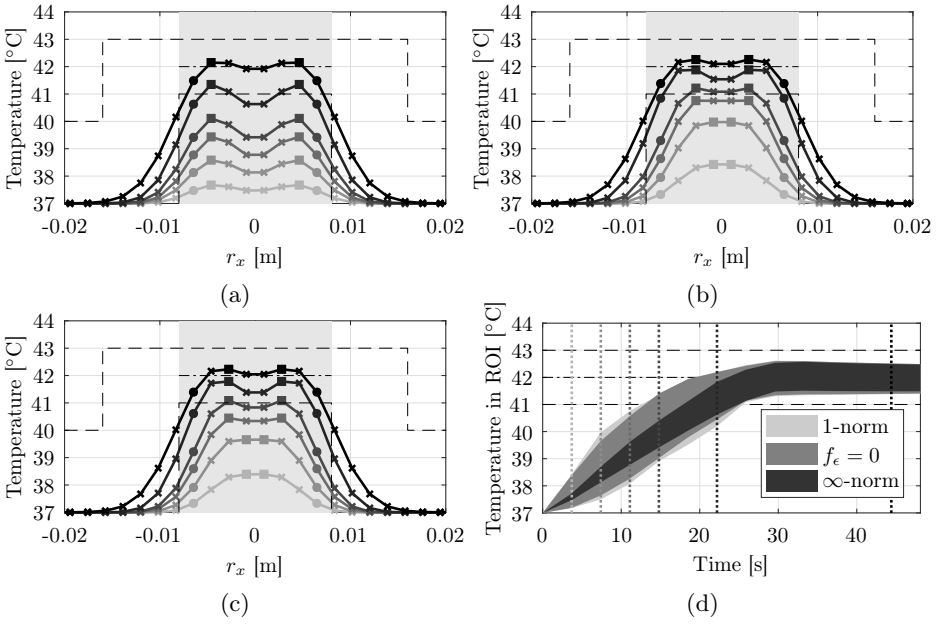

Figure 2.7: The evolution of the voxel temperatures (×, •, and ■) of the cross section at $r_y = 0$ m during heat-up when using (a) an ∞-norm penalty, (b) a 1-norm penalty, and (c) no penalty on the slack variables, at the time instants indicated in (d) (dotted), alongside the ROI (gray area), and the temperature range's upper and lower bounds (dashed) and setpoint temperature (dash-dotted). In (d), the difference between the maximum (■) and minimum (•) temperature in the ROI is shown for all cases.

2.4.6 Optimization problem

The control objective can now be formulated as the constrained optimization problem

$$\min_{\boldsymbol{u}_k} V_N(\boldsymbol{z}_k, \boldsymbol{\epsilon}_k) \tag{2.14a}$$

where $\boldsymbol{u}_k = (u_{0|k}, \ldots, u_{N-1|k})$, subject to

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k} + B_d d_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{2.14b}$$
$$d_{i+1|k} = d_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{2.14c}$$
$$x_{0|k} = A_0 \hat{x}_k + B_0 u_{k-1} + B_{d,0} \hat{d}_k, \tag{2.14d}$$
$$d_{0|k} = \hat{d}_k, \tag{2.14e}$$
$$z_{i|k} \geq \underline{z} - 1_{n_z} \underline{\epsilon}_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N]}, \tag{2.14f}$$
$$x_{i|k} \leq \overline{x} + 1_{n_x} \overline{\epsilon}_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N]}, \tag{2.14g}$$

$$0_{n_\epsilon} \leq \epsilon_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N]}, \tag{2.14h}$$

$$0_{n_u} \leq u_{i|k} \leq 1_{n_u}\overline{u}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{2.14i}$$

$$1_{n_u}^\top u_{i|k} \leq \overline{u_\Sigma}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{2.14j}$$

where $\mathbb{N}_{[a,b]} = \{a, a+1, \ldots, b\}$, $a, b \in \mathbb{N}$ and $a \leq b$, denotes the set of natural numbers from a through b, and 1_n represents an all-ones vector of length n. Equality constraints (2.14b)-(2.14c) ensure satisfaction of the dynamics according to (2.7), with the prediction sequence's initial conditions given by (2.14d)-(2.14e), which correspond to (2.8) and can be computed after first performing the observer step (2.9) to obtain $\hat{x}_k$. Inequalities (2.14f)-(2.14g) capture the temperature bounds as soft constraints using the slack variables bounded from below in (2.14h). Finally, (2.14i)-(2.14j) describe the actuator constraints, where $\overline{u} = 15$ W and $\overline{u_\Sigma} = 60$ W are the maximum allowable power applied to a single sonication point and to the entire treatment region, respectively.

Furthermore, a horizon of $N = 5$ is used, and for the weights we choose

$$Q = \frac{1}{n_z} I, \quad f_\epsilon = \begin{bmatrix} 10 \\ 100 \end{bmatrix}, \tag{2.15}$$

which are normalized with respect to the number of corresponding variables (note that $\underline{\epsilon}$ and $\overline{\epsilon}$ are scalar) for more intuitive balancing of the objectives' relative contribution to the cost function.

We choose to incorporate the desired temperature range objective as soft constraints to prevent the optimization problem from becoming infeasible, as proper controller operation is crucial for patient safety and treatment quality. Note that this is of particular importance for the considered application, as the minimum temperature bound will be directly violated during heat-up. In addition, violations may occur as a result of unexpected heat sinks or sources, or due to measurement artifacts that may occur in MR-based thermometry due to patient motion or magnetic field drift.

Remark 2.1. *The stage cost* (2.13) *is improved with respect to [84] by excluding the cost on the predicted input power $u_{i|k}$. That is, removing this penalty slightly increases control effort, but consequently also removes the associated downward shift of the tumor temperature away from the reference z_r. This is desirable for hyperthermia treatments, where achieving optimal tumor temperatures, and thereby optimal treatment quality, is of the utmost importance, in comparison to which the desire to reduce the control energy is negligible. In addition, note that the control effort is in fact upper bounded by the hard input constraints* (2.14i)-(2.14j).

Similarly, we chose not to include an input rate penalty in (2.13)*, since reducing changes in the input signal is also of little to no importance compared to the temperature objectives. Moreover, the input oscillations resulting from the propagation of noise into the state and disturbance estimates are reduced*

to a negligible level by adequate filtering in the observer, as will be shown in Figure 2.12 in the upcoming discussion of the in-vivo experiment.

Remark 2.2. *The cost function* (2.12) *describes an economic cost, as the set-point z_r does not correspond to a reachable steady-state solution of the considered system* (2.7)*. This is due to the combination of the ROI radius being equal to the lateral range using electronic beam steering, meaning all sonication points are inside the ROI, the finitely narrow power deposition in* (2.4)*, and the fact that the nonnegativity constraints on the inputs* (2.14i) *are active at steady state. As a result, although having satisfied* (2.10)*, we cannot guarantee that $z_k \to z_r$ for $k \to \infty$. Instead, using the fact that the estimator is asymptotically stable and the disturbances are (nearly) constant, we only claim to be able to remove the part of the steady-state offset resulting from plant-model mismatch, recovering the performance as optimally achievable without modeling error.*

An alternative approach would be to formulate the objective as a tracking problem with respect to a feasible steady-state setpoint. To this end, a target selector as discussed in [67, 91, 92] could be used to determine the economically optimal steady state and corresponding input (x_s^*, u_s^*) as a function of the current disturbance estimate by solving at each time k

$$(x_s^*, u_s^*) = \arg\min_{x_s, u_s} \ell(z_s, \epsilon_s), \tag{2.16}$$

with $z_s = Hx_s$ and $\epsilon_s = [\underline{\epsilon}(z_s), \overline{\epsilon}(x_s)]^\top$, subject to

$$\begin{bmatrix} I - A & -B \end{bmatrix} \begin{bmatrix} x_s \\ u_s \end{bmatrix} = B_d \hat{d}_k, \tag{2.17}$$

and inequality constraints similar to (2.14f)-(2.14j)*. However, such a setup has been observed to deteriorate performance compared to the current design during the transient of $\hat{d}_k$ and in the presence of measurement noise, while yielding at best similar performance after convergence of the disturbance estimate, see also [93]. In addition, this increases complexity (and thereby computation time), which is particularly unwanted in the clinic, where any additional complexity is regarded as a potential source of errors. We have therefore chosen to omit such a target selector, and instead optimize the economic cost in* (2.12) *directly.*

2.5 In-vivo experiments

In this section, the experimental results obtained on an in-vivo porcine thigh muscle are presented to demonstrate the effectiveness of the designed feedback control setup to remove the part of the treatment temperature steady-state offset caused by plant-model mismatch.

2.5.1 Implementation

The MR-HIFU system, consisting of the combination of an MRI scanner and a dedicated HIFU therapy platform, has been discussed in Section 2.2.1 and 2.2.2. The MPC and estimator algorithms were implemented in Matlab 2017b and executed in real time using Python 2.7 via the Matlab engine API, running on a HP Z800 Workstation with Intel Xeon X5650 CPU @ 2.67 GHz and 12 GB RAM. The communication with the MR-HIFU system was achieved using the matMRI and matHIFU toolboxes [94]. For fast numerical optimization, we used Gurobi 7.0.2 to solve a sparse formulation of the MPC problem (2.14), which after comparative testing with a dense formulation was concluded to result in significantly smaller computation times.

In particular, the solver times recorded during the experiments were fairly consistent, averaging at approximately 0.26 s when the disturbance estimator was disabled, i.e., when running MPC without offset-free capabilities, and around 0.33 s for the offset-free MPC (2.14). Since the computation time of the observer (2.9) is negligible (≤ 5 ms in both cases), this difference illustrates the additional computation time associated with using an offset-free approach. Recalling that the sample time is $T_s = 3.7$ s, we find that the resulting computation times for both MPC strategies are still well within the bounds for real-time implementation. Given the significant performance improvement demonstrated by the experimental results shown below, this strongly motivates the offset-free MPC approach.

2.5.2 Tissue parameters

Using our proposed approach, we assume the treatment area to have homogeneous tissue properties, and rely on the disturbance estimator to identify the effects of (inhomogeneous) parameter error. Baseline values for the relevant muscle tissue parameters have been obtained from the IT'IS Foundation database [95], which have then been slightly adapted according to the identification procedure in [72] based on low-power test sonications. For the thermal model equations (2.2)-(2.4), from which the continuous-time state-space model (2.6) and in turn the discrete-time controller and observer models (2.7), (2.8) and (2.9) are derived, however, we have intentionally introduced additional parameter mismatch that results in a negative temperature offset in closed loop. In particular, corresponding to the observations in [72], this is achieved by underestimating the thermal conductivity κ in (2.2) and overestimating the acoustic energy absorption coefficient α in (2.4), resulting in the parameters listed in Table 2.1, and a plant-model mismatch as visualized in Section 2.7 by an open-loop model validation.

This is done for two reasons. First, for the purpose of this work it allows for better illustration of the controller's ability to achieve the desired treatment temperature despite significant model error. Second, although any plant-model

Table 2.1: Tissue parameters.

Parameter	Value	Unit
ρ	1090	$\mathrm{kg/m^3}$
c	3421	$\mathrm{J/(kg{\cdot}K)}$
κ	0.3773	$\mathrm{W/(m{\cdot}K)}$
w_b	3.5289	$\mathrm{kg/(s{\cdot}m^3)}$
c_b	3617	$\mathrm{J/(kg{\cdot}K)}$
α	0.06141	-

mismatch is obviously undesirable in general, in case it is inevitable due to large model uncertainty, for patient safety it is preferred to have model bias towards underheating instead of towards overheating. That is, in case of the former, no harm is unintentionally inflicted to the tissue, and, after proper identification of the mismatch, increasing the sonication power will directly contribute to achieving the desired temperature. Contrarily, overheating the tissue may induce hyperthermic reversal effects detrimental to treatment success, or directly cause unwanted tissue damage, and, after identification, cannot be actively counteracted using HIFU (as expressed by the nonnegativity constraints on the inputs $u_k \geq 0$).

2.5.3 Initial disturbance estimator inactivity

In line with the desire to demonstrate the benefits of using the proposed offset-free implementation, the disturbance estimator is disabled during the initial heat-up phase by setting $\hat{d}_k = 0$ and $L_d = 0$ for all $k \in \mathbb{N}$ corresponding to $t_k < 180$ s. This is, however, also desirable from a clinical point of view, as hereby we avoid a potential and unnecessarily large build-up of $\hat{d}_k$ that would otherwise lead to increased temperature over- or undershoot, see *Remark 2.3*.

Remark 2.3. *Initially disabling the disturbance estimator is especially useful in case of a mismatch in the input matrix B_c in* (2.6) *(from which the observer and controller input matrices are derived), e.g., resulting from an incorrect value for the acoustic absorption coefficient α, combined with rapid and substantial changes in the input u_k such as typically occurring in the transient behavior during heat-up (especially when transitioning from heat-up towards steady state). This effect is due to two reasons. First, a model error in B_c is* multiplicative *in nature, whereas we use an* additive *disturbance in the augmented model* (2.6) *and corresponding observer* (2.9)*. Second, our disturbance estimator* (2.9) *is designed to estimate* steady-state *(or slowly varying) errors, while the input u_k (and thus the mismatch in the terms $\hat{B}_i u_k$, $i = 1, 2$) can vary rapidly. Although this could possibly be tackled by using different disturbance models or a more sophisticated adaptive controller including online parameter estimation, in the*

context of MR-HIFU hyperthermia the resulting performance improvement is believed to be minor for the following reasons. First of all, due to its general form, the design proposed in this work is able to fully capture the effects of any constant or slowly varying model mismatch and disturbance in steady state, rather than only specific types of model errors. Furthermore, our control setup already achieves heat-up times of less than one minute, which is short with respect to the envisioned 90-minute treatment time, after which the disturbance estimator can be enabled. In fact, the maximum power limit is already typically reached during heat-up, indicating that there is relatively not much room for shortening the heat-up phase.

2.5.4 Results

To test the performance of the developed offset-free controller and observer setup under realistic clinical conditions, in-vivo experiments have been conducted on an in-vivo porcine thigh muscle. The focal plane MR image is shown in Figure 2.8, indicating the treatment area (green) and ROI (red). We set $r_x = r_y = 0$ m at the center of the ROI, with positive r_x- and r_y-axis directions towards the ventral (belly) and cranial (head) side, respectively. In Figure 2.9, the results of a 30-minute MR-HIFU hyperthermia treatment are presented. Note that the baseline body/blood temperature of the treated pig was $T_b \approx 38$ °C. The mean (solid) and minimum/maximum (dashed) voxel temperatures are shown corresponding to, with a temperature shift of T_b, both the state estimate in the ROI $\hat{z}_k = H\hat{x}_k$ (thick black) and the measurement in the ROI $z_{y,k} = Hy_k$ (thin red), together with the ROI's reference temperature T_r and lower and upper bounds of the desired temperature range $\underline{T}$ and $\overline{T}$, respectively. The disturbance estimator is disabled for $t_k < 180$ s indicated by the gray area. In Figure 2.9a, the temperature indicators are given for the full length of the treatment, showing that after heat-up and disturbance estimator convergence the entire ROI is heated to within the desired temperature range, with the mean temperature settling at T_r. Moreover, it illustrates the benefit of using a model-based observer to obtain temperature estimates that contain significantly reduced noise effects compared to the corrupted measurements.

Closer inspection of the interval where $0 \leq t_k \leq 400$ s in Figure 2.9b reveals the effect of the existing plant-model mismatch on the treatment temperature in closed loop. Due to overestimating the achievable temperature increase, it appears from the observer as if the minimum ROI temperature has exceeded the desired 41 °C at $t_k \approx 26$ s, whereas the measured (noisy) minimum ROI temperature is around 39.7 °C at this time. It can also be seen that with the disturbance estimator disabled, the modeling error results in a discrepancy between the measured and estimated temperature that settles around 0.3 °C on average at $t = 180$ s. Consequently, a negative mean temperature offset of almost 0.5 °C can be observed, assuming that the effect of v_k on the mean of $z_{y,k}$

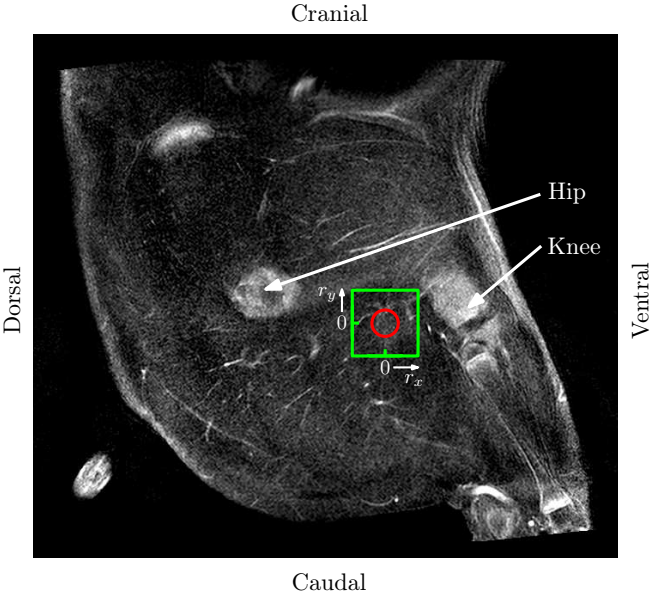

Figure 2.8: Sagittal focal plane MR image of the porcine right thigh area, ranging from halfway the abdomen (cranial) to the ankle (caudal), indicating the location of the ROI (red circle).

is sufficiently suppressed by averaging and therefore $z_{y,k}$ is representative of the actual mean ROI temperature. After activating the disturbance estimator, i.e., setting $L_d = 0.01 I_{n_x}$ as defined in (2.11) for $t_k \geq 180$ s, the state and disturbance estimates converge within approximately one minute, allowing for the feedback setup to achieve the desired temperature distribution, as would typically be the result in case of no model error.

Temperature cross sections at $r_x = 0$ m, i.e., through the center of the ROI in r_x-direction, are shown in Figure 2.10 corresponding to the averages of the state estimate $\bar{\hat{x}}$ (black $\circ$) and measurement $\bar{y}$ (red $\times$) over the 7-sample intervals $154 \leq t_k \leq 180$ s (Figure 2.10a) and $374 \leq t_k \leq 400$ s (Figure 2.10b), together with the corresponding $2\sigma_x$ and $2\sigma_y$ confidence intervals (gray/red areas), where $\sigma_{\hat{x}}$ and σ_y denote the standard deviations over the corresponding 7-sample intervals. The values shown in Figure 2.10a correspond to the last seven samples before activating the disturbance estimator for $t_k \geq 180$ s. First of all, the confidence intervals illustrate the need for an observer, as the measurements y_k clearly suffer from noise effects, whereas these are filtered down to a negligible level in the estimates $\hat{x}_k$. Second, a clear discrepancy can be observed between $\bar{\hat{x}}$ and $\bar{y}$, particularly in the ROI where heat loss effects are

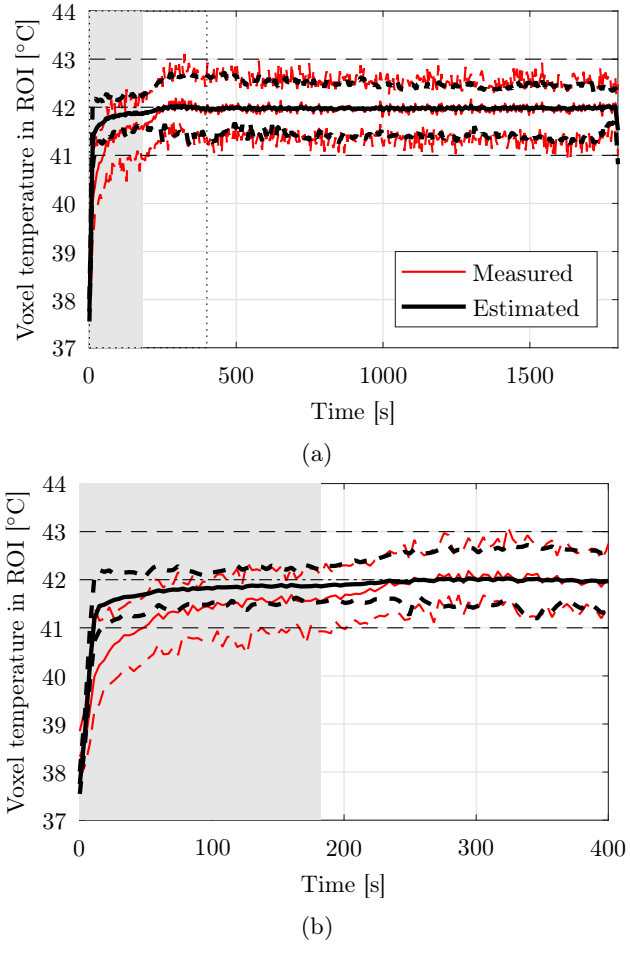

Figure 2.9: The mean (solid) and minimum/maximum (dashed) voxel temperatures of the ROI corresponding to the measurements $z_{y,k} = Hy_k$ (thin red) and state estimates $\hat{z}_k = H\hat{x}_k$ (thick black) with the disturbance estimator disabled for $t < 180$ s (gray area), shown for (a) the entire experiment, and (b) zoomed in on the first 400 s of the experiment (dotted rectangle in (a)).

most strongly present, indicating an offset between $\hat{x}_k$ and y_k, which is the result of plant-model mismatch. The results over the last seven samples before $t_k \approx 400$ s, depicted in Figure 2.10b, exemplify the offset-free implementation's ability to enable convergence of the estimator and enforce the desired heating behavior.

To benchmark the performance of the offset-free MPC scheme developed

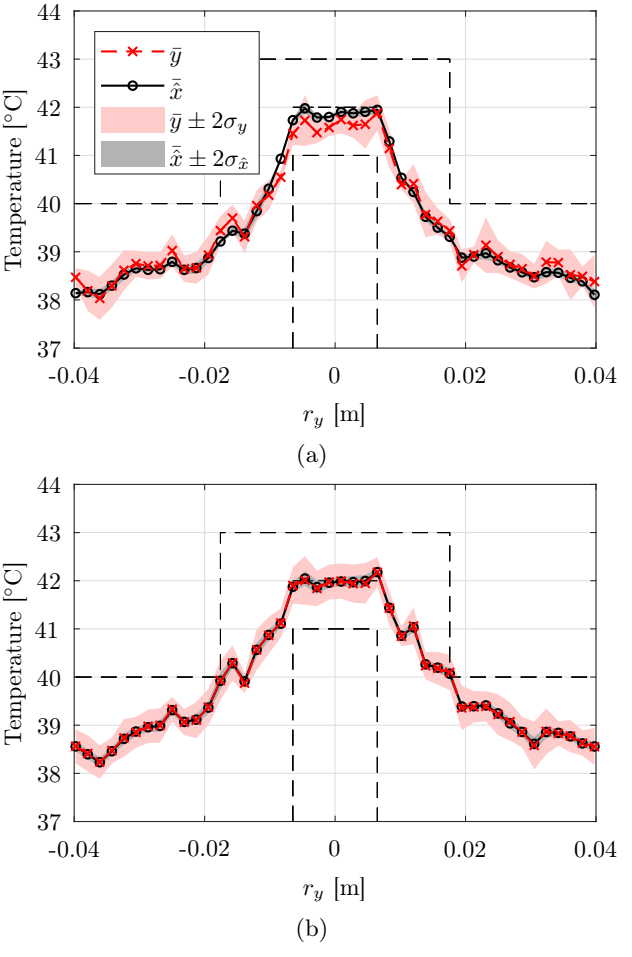

Figure 2.10: The 7-sample averaged temperature cross sections at $r_x = 0$ m corresponding to $\bar{\hat{x}}$ (black $\circ$) and the measurement $\bar{y}$ (red $\times$) with corresponding 2σ confidence intervals (red/gray areas) when the disturbance estimator is (a) inactive over the interval $154 \leq t_k \leq 180$ s, and (b) active during $374 \leq t_k \leq 400$ s, alongside the $\underline{T}$ and $\overline{T}$ (dashed black) and T_r (dash-dotted black).

in this work, consider Figure 2.11. Herein, we show temporal average of the ROI temperature measurements' deviation with respect to the setpoint T_r (i.e., $z_{y,k} = Hy_k$ with respect to z_r) obtained using the proposed MPC design with the disturbance estimator and, thereby, the offset-free capabilities disabled (MPC 1) and enabled (Offset-free MPC 1), which are computed over the intervals $154 \leq t_k \leq 180$ and $300 \leq t_k \leq 1800$ during the experiment, respectively. In

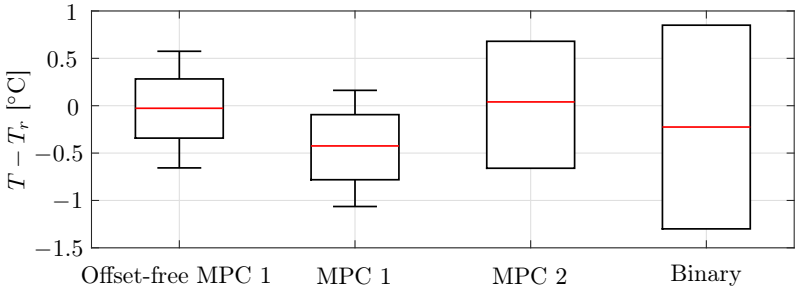

Figure 2.11: Measured ROI temperature mean (red), 10th percentile $T90$ and 90th percentile $T10$ (box edges), and extrema (whiskers) with respect to the ROI setpoint T_r in steady state for the MPC design proposed in this work with the offset-free scheme enabled (Offset-free MPC 1) and disabled (MPC 1), compared to the in-vivo porcine experiment results from [46] using a preliminary MPC scheme (MPC 2) and [53] using a state-of-the-art binary controller (Binary).

addition, the figure depicts the temperature distributions corresponding to the results[1] reported in [46, Table 2] obtained using our preliminary MPC design (MPC 2), and in [53, Table II] using a clinical state-of-the-art binary controller (Binary), both of which also involve in-vivo porcine thigh muscle experiments, and target a similar ROI (18 mm diameter circle) as considered in this work. This figure clearly illustrates, first of all, the benefit of using an MPC-based strategy instead of the binary controller, since the former achieves a more accurate and homogeneous temperature distribution, as shown by the ROI temperatures' 10th and 90th percentiles, which are denoted by $T90$ and $T10$, respectively. To be precise, the $T90$-$T10$ temperature ranges of Offset-free MPC 1 (0.62 °C) and MPC 1 (0.69 °C) corresponding to the MPC presented here, and of MPC 2 in [46] (1.34 °C), are only 29%, 32%, and 62% of the $T90$-$T10$ range obtained using the binary controller in [53] (2.15 °C). Second, it demonstrates the improved control performance of the proposed feedback setup (Offset-free MPC 1 and MPC 1) with respect to that used in [46] (MPC 2), which differs regarding the MPC design and, in addition, does not employ a state or disturbance observer. In particular, the $T90$-$T10$ temperature ranges obtained using Offset-free MPC 1 and MPC 1 (in the presence of considerable plant-model mismatch) are approximately half of the $T90$-$T10$ temperature range achieved using MPC 2 in [46] (without such significant modeling error), and even the minimum-maximum temperature range of Offset-free MPC 1 and MPC 1 (both 1.23 °C) are only 92% of the $T90$-$T10$ range of MPC 2. Finally, the figure once more visualizes the steady-state offset resulting from large plant-model mismatch when using a

[1] The minimum and maximum ROI temperatures were not reported in [46, 53] and are therefore not shown here.

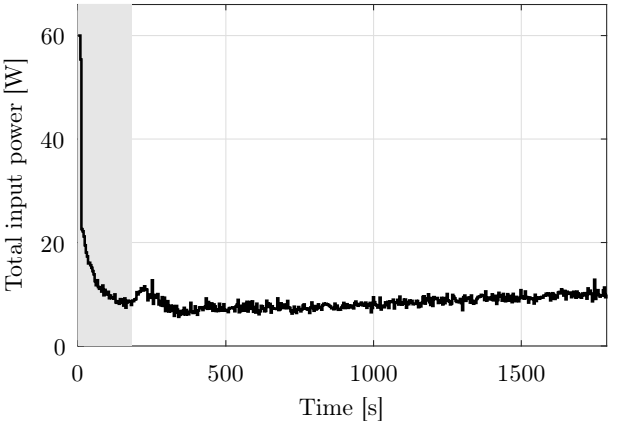

Figure 2.12: The total input power per time step with the disturbance estimator disabled for $t < 180$ s (gray area).

standard tracking MPC strategy without offset-free capabilities (MPC 1), which can be eliminated by using the offset-free MPC scheme developed in this work (Offset-free MPC 1), strongly motivating the use of an offset-free MPC strategy in MR-HIFU hyperthermia treatments.

Next, we show the total input power per time step applied during the experiment in Figure 2.12. The benefit of using an observer, reducing the propagation of noise effects into the control input, is expressed in the small magnitude of the power fluctuations occurring during the treatment. The figure also clearly visualizes the large acoustic power required for initial heat-up limited only by the input constraint (2.14j), and the increased power demand after activating the disturbance estimation that eventually brings the tissue in the ROI to its desired temperature. Interestingly, the requested sonication power in this experiment also exhibits a steady increase for $t_k \geq 800$ s during which the tissue temperature does not significantly change. This is most likely due to the tissue's thermoregulatory response, increasing perfusion when having been exposed to elevated temperatures for an extended period of time, as will be discussed in more detail in Section 2.5.5. This clearly demonstrates the controller's ability to maintain optimal performance also in the presence of slowly varying model mismatch.

Possibly even more remarkable, however, is that the power increase can be observed to be seemingly linear. This is of special importance for medical science, since measuring and modeling the transient behavior of the body's thermoregulatory response to thermal stress has proven to be difficult [74]. By providing the means for accurately controlling the hyperthermic conditions and identifying tissue changes, our controller may serve as an enabler for research into better

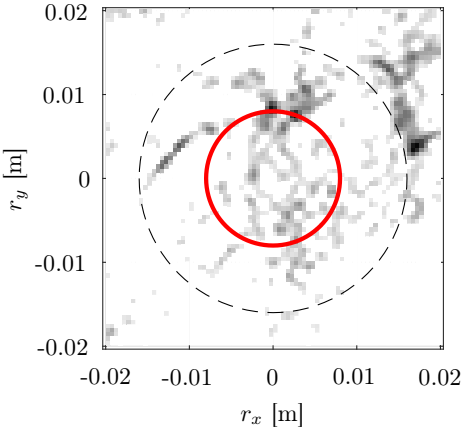

Figure 2.13: MR image of the treatment area (green box in Figure 2.8), highlighting the intensity of the vascularization (grayscale) in and around the ROI $\mathcal{R}$ (red circle) and $\mathcal{S}$ (dashed black circle).

physiological tissue models.

2.5.5 Spatially and temporally varying disturbance estimation

For the considered experiment, the ROI contained significantly vascularized tissue, as can be seen from the focal plane MR image shown in Figure 2.13, which corresponds to the area marked in Figure 2.8 (green box). The outer edges of the ROI $\mathcal{R}$ and of the area including the surrounding tissue $\mathcal{S}$ are marked by the solid red circle and the dashed black circle, respectively. Tissue areas with stronger perfusion, typically containing networks of small blood vessels and the tissue in their vicinity, are visualized in darker gray. In Figure 2.14, the estimated voxel-wise temperature mismatch $\hat{B}_d\hat{d}_k$ is shown for two different time instants. Additionally, the subfigures herein also include the MR image in the background (gray, with increased contrast for visibility). Both Figure 2.14a and Figure 2.14b illustrate that the proposed disturbance estimator is able to capture spatially varying model discrepancies. Moreover, when comparing the temperature mismatch estimate before the input power increases at $t_k \approx 800$ s in Figure 2.14a to the estimated mismatch at $t_k \approx 1800$ s in Figure 2.14b, it is clear to see that the estimated local heat loss has increased in some areas, demonstrating the estimator's ability to identify model mismatches that (slowly) change over time.

Especially inside the ROI, where heating is applied and consequently the temperature is elevated, negative values for the estimated temperature mismatch (cyan/blue) can be seen to largely coincide with the more strongly perfused tissue (gray). Also at the edge of and outside the ROI, a striking spatial correlation can

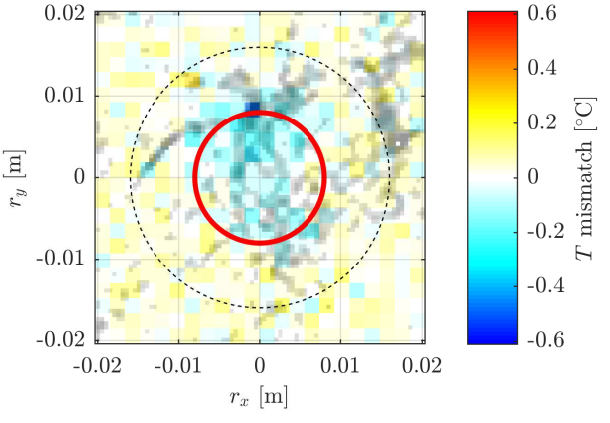

(a)

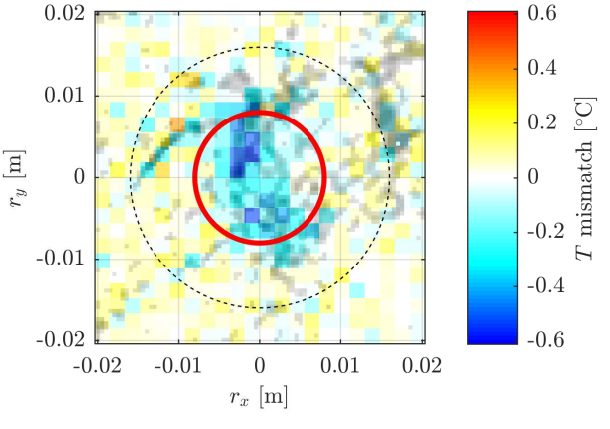

(b)

Figure 2.14: The estimated voxel-wise temperature mismatch $\hat{B}_d\hat{d}_k$ at time (a) $t_k \approx 800$ s and (b) $t_k \approx 1800$ s for the tissue area containing the ROI $\mathcal{R}$ (solid red circle) and $\mathcal{S}$ (dashed black circle).

be observed, most notably so in the areas around the coordinates $(-0.01, 0.003)$, $(0.004, -0.008)$ and $(0.007, 0.006)$. In the regions outside the ROI where neither sonication occurs, and thereby modeling errors in the overestimated B are not expressed, nor significant perfusion is present, the disturbance estimate mostly takes on small positive values (yellow). This is the result of the estimator identifying the effects of underestimating the thermal conductivity κ, by which more heat is transported outwards from the ROI than expected by the nominal model.

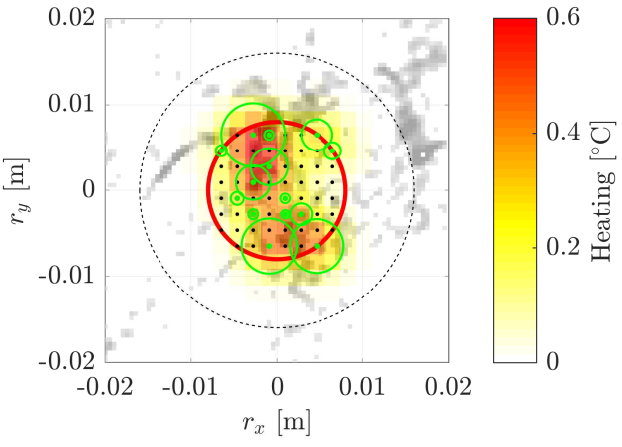

Figure 2.15: The heating Bu_k at $t_k \approx 1800$ s for the tissue area containing the ROI $\mathcal{R}$ (solid red circle) and $\mathcal{S}$ (dashed black circle). Active sonication points (green dots) are distinguished from inactive sonication points (black dots). For the former, the relative magnitude of the applied acoustic power is indicated by the radius of the corresponding green circle.

In Figure 2.15, a spatial resemblance can also be observed between the disturbance estimate and the heating that is applied using this value for $\hat{d}_k$ at time $t_k \approx 1800$ s. Here, the active and inactive sonication points (green and black dots, respectively) are shown inside the ROI (solid red circle). The radius of the green circles drawn around the active sonication points represent the relative magnitude of the applied acoustic power. The resulting spatial distribution of Bu_k and the active sonication points can be seen to correspond well to the highlighted tissue in Figure 2.13 and the temperature mismatch in Figure 2.14b.

From this, we draw several conclusions. First, neglecting perfusion in (2.2) may indeed be the cause of significant unmodeled spatially varying heat loss when the ROI contains (considerably) vascularized tissue. Second, the effects of perfusion increasing over time due to the body's thermoregulatory response can be of sufficient magnitude to require an increasing steady-state power demand as seen in Figure 2.12. Next, note that although the identified heat loss shows obvious similarity to the MR perfusion image, it is not an exact match. Additionally, relating the MR image to the magnitude of the perfusion-related heat loss is not straightforward. In fact, even the relative magnitude of these effects does not exactly correspond to the MRI data, as exemplified in Figure 2.14b by the disturbance estimates at $(-0.001, -0.005)$ and $(-0.002, 0.002)$ being larger in magnitude than the estimate at $(-0.001, 0.008)$, while from Figure 2.13 the latter area is seen to be most strongly perfused. Therefore, we conclude that accurately modeling the thermal effects of perfusion a priori based on an MR im-

age as in Figure 2.13 may be considered as extremely difficult, even more so due to the unavailability of extensive model identification (as previously mentioned). These conclusions motivate our offset-free MPC approach, involving the use of a simple homogeneous thermal model (2.2), and incorporating a disturbance estimator to allow for the identification and compensation of the potentially time-, space-, and temperature-dependent unmodeled thermal effects.

2.6 Conclusion

In this chapter, an offset-free MPC scheme is presented to enhance performance and robustness in temperature control using an MR-HIFU hyperthermia system for cancer treatment. Since plant-model mismatch inevitably occurs when treating different patients and tumors, and accurate thermal modeling is impracticable for MR-HIFU hyperthermia, we propose to assume a simple model with homogeneous tissue properties and (near) database parameter values, and then include a disturbance estimator to cope with the resulting model mismatch. We showed that this estimator is able to identify the effects of constant and slowly varying modeling errors (and disturbances), enabling the controller to eliminate the steady-state offset otherwise resulting from such errors. By means of experiments on the thigh muscle of a living anaesthetized pig using a pre-clinical MR-HIFU hyperthermia system, we have verified the feedback setup's effectiveness in terms of recovering optimal treatment temperatures despite the presence of realistic plant-model mismatch.

2.7 Appendix: Open-loop model validation

The plant-model mismatch in the in-vivo experiment is visualized by feeding the inputs applied during the experiment, of which the total power is shown in Figure 2.12, to the nominal model (2.5) in open loop. The nominal model's resulting minimum/maximum (dashed) and mean (solid) tumor temperatures are shown in Figure 2.16 (red), and compared to the thermal response of the porcine thigh muscle measured during the experiment (black). Clearly, there exists a severe mismatch, which, in addition, can be seen to grow over time due to the increased perfusion caused by the body's thermoregulatory response, see also Figure 2.12 and Figure 2.15. This illustrates the degree of model error that our offset-free MPC was able to compensate for during the in-vivo experiment, see Figure 2.9.

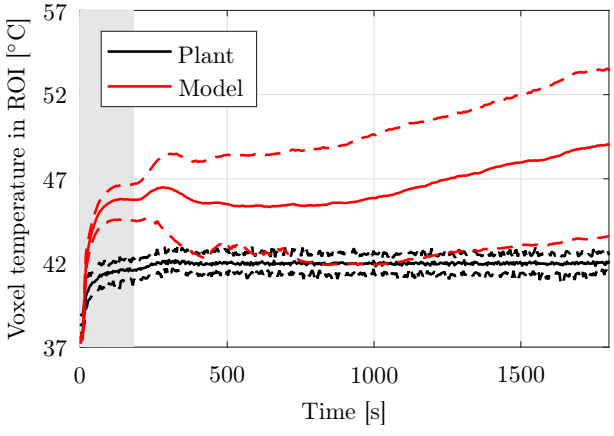

Figure 2.16: The mean (solid) and minimum/maximum (dashed) voxel temperatures of the ROI corresponding to the experiment measurements (black) and generated by the nominal model (red) when supplied with the inputs applied during the in-vivo experiments. The gray background indicates when the disturbance estimator was disabled during the in-vivo experiments.


The greatest challenge to any thinker is stating the problem in a way that will allow a solution.
Bertrand Russell

Chapter 3

Switched-Actuator Systems with Setup Times: Efficient Modeling, MPC, and Application to Hyperthermia Therapy

In many control applications, the actuator configuration has to be switched during operation. This often requires non-negligible actuator downtime, resulting in switched-actuator systems with setup times (SAcSSs). Consequently, optimally controlling SAcSSs requires the online solving of both a discrete actuator allocation problem, in which the switch-induced actuator downtime is taken into account, as well as an optimization problem for the (typically continuous) control inputs. Mixed-integer model predictive control (MI-MPC) offers a powerful framework for tackling such problems. However, the efficient modeling of SAcSSs for MI-MPC is not straightforward, and real-time feasibility is often a major hurdle in practice. Therefore, in this chapter we first formalize the general description of the class of SAcSSs. Then, we propose an intuitive and systematic modeling procedure tailored to SAcSSs, which is specifically designed to allow for user-friendly controller synthesis, and to yield efficient MI-MPCs. We apply these new results in a case study of large-volume magnetic-resonance-guided high-intensity focused ultrasound hyperthermia, which involves the heating of tumors (using real-valued local heating controls, as well as discrete range-extending actuator relocation during which no heating is allowed) to enhance the efficacy of radio- and chemotherapy.

This chapter is based on D.A. Deenen, E. Maljaars, L.C. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Switched-Actuator Systems with Setup Times: Efficient Modeling, MPC, and Application to Hyperthermia Therapy," *Submitted for journal publication.*

3.1 Introduction

A major motivation for the system-theoretic contributions of this chapter is large-volume magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU) hyperthermia [83]. This is a novel technology for noninvasive cancer treatment in which ultrasound waves are used to locally heat up cancer tissues to 42 °C for 60 to 90 minutes, based on real-time temperature measurements obtained with an MRI scanner. The heating effectively sensitizes the tissue to the effects of chemo- and radiotherapy without adding any undesirable (toxic) side effects [3, 15]. Thereby, hyperthermia enables significantly higher success rates and allows for a considerable reduction of the unwanted side effects of chemo- and radiotherapeutic cancer treatments [16, 22].

MR-HIFU allows for powerful heating with millimeter-accurate steering. Using MPC, HIFU-mediated thermal therapies of optimal quality can be realized, while respecting actuator and safety constraints, see [69–72]. Unfortunately, this high accuracy comes with limited heating range, implying that only small tumors ($\leq$16 mm diameter) can be treated using a stationary actuator. For larger tumors, the actuator itself must be relocated, during which no heating can occur [53]. As the set of admissible positions must be discrete and finite, the resulting system can be described as a switched system, where the input model differs depending on the actuator location. However, a first distinctive property of this hyperthermia system, compared to typical switched systems, is that the time required for actuator relocation, and thus a mode switch, is non-negligible. Second, any mode switch induces nonzero actuator downtime, during which the system itself keeps evolving in time according to its unforced dynamics. Clearly, for the optimal control of the heating process in large-volume MR-HIFU hyperthermia, these two features have to be incorporated in the controller design, determining online both the continuous local heating and the discrete actuator positioning.

The essential features recognized in large-volume MR-HIFU hyperthermia, i.e., dealing with dynamical systems in which the (finite number of) actuator configurations have to be switched and where the switching takes significant time during which no active control is possible, are of a general nature as they can be identified in many other applications. Indeed, one could think about manufacturing systems in which machine reconfiguration takes time, see [96–98] for comprehensive surveys, but also agents in agriculture (e.g., drones, fertilizers, irrigation systems) which must serve multiple (sub)fields resulting in significant field-to-field travel times [99, 100], or the coordinated deployment of fire-fighting units for wildfire management [101–103]. Motivated by this range of applications, we formally introduce the class of *switched-actuator systems with setup times* (SAcSSs), exhibiting the mentioned features, as the first contribution of this work.

As a second contribution, we address the design of easy-to-derive and efficient

MPC schemes for SAcSSs. In particular, we desire a natural and systematic modeling procedure for SAcSSs, yielding a compact model, i.e., with a small number of integer variables, that can be directly integrated into a mixed-integer MPC (MI-MPC) [68] setup. The small number of integers is motivated by the desire to keep the computational burden limited, thereby facilitating real-time implementation of the resulting MPC schemes. The mixed-integer programming (MIP) compatibility could be achieved by describing a SAcSS as a mixed logical dynamical (MLD) system [104]. However, efficiently capturing the actuator switching behavior including setup times is not straightforward. Clearly, one could attempt to describe a SAcSS as a discrete hybrid automaton (DHA) [105], which using the modeling language HYSDEL could then be automatically converted into MLD form amenable for online optimization. Unfortunately, due to the generic nature of the DHA framework, it may be unclear how to best incorporate SAcSSs' actuator switching, thereby potentially leading to models containing too many Boolean variables, which would have negative consequences for the computational complexity of the online MPC problem. For example, a SAcSS could first be cast as a constrained switched (linear) system [81, 106] before deriving its equivalent DHA and (using HYSDEL) MLD forms, yielding a similar result as the lifting approach in [107]. However, incorporating an actuator switch's setup time would in these cases require the inclusion of as many additional "transitioning modes" (i.e., duplicates of the zero-input mode corresponding to the unforced dynamics) as the switch's setup time. In turn, this results in a model with many integer (Boolean) variables, leading indeed to high-complexity representations. Therefore, as a second contribution, we propose an intuitive and convenient modeling procedure tailored to SAcSSs, specifically designed to yield compact models that are directly suitable for efficient MI-MPC setups.

The final contribution of this work consists of a large-volume MR-HIFU hyperthermia cancer therapy case study. Herein, we apply the novel methods to obtain a SAcSS model and MI-MPC, of which the performance and computational efficiency are validated.

A preliminary version of this work was published in [108], which discusses only a simulation study of MI-MPC for large-volume MR-HIFU hyperthermia (different from the one included here) as an illustrative proof-of-concept, but does not provide the general modeling procedure for SAcSSs presented in this chapter. In fact, the class of SAcSSs is formally defined here for the first time. Moreover, in this work we formalize the key concepts of the proposed modeling framework, provide rigorous proofs, discuss the general MI-MPC setup for SAcSSs, and investigate its improved computational efficiency, all of which was not included in [108].

The remainder of the chapter is organized as follows. In Section 3.2 we formally define SAcSSs, for which in Section 3.3 we present the intuitive and compact modeling procedure. The resulting MI-MPC will be given in Section 3.4.

In Section 3.5, the large-volume MR-HIFU hyperthermia case study is presented. Finally, Section 3.6 summarizes the key contributions of this work.

Notation. The real, integer, and natural numbers (including zero) are denoted by $\mathbb{R}$, $\mathbb{Z}$, and $\mathbb{N}$, respectively. Given a set $\mathcal{S} \subseteq \mathbb{R}$ and values $a, b \in \mathcal{S}$, we use $\mathcal{S}_{>a}$, $\mathcal{S}_{\geq a}$, $\mathcal{S}_{<a}$, and $\mathcal{S}_{\leq a}$ to denote the subset of $\mathcal{S}$ of which the elements satisfy the condition in the subscript, and we define $\mathcal{S}_{[a,b]} = \{s \in \mathcal{S} \mid a \leq s \leq b\}$ (hence $\mathcal{S}_{[a,b]} = \emptyset$ if $a > b$). $\mathcal{P}(\mathcal{S})$ denotes the power set of $\mathcal{S}$, being the collection of all subsets of $\mathcal{S}$, and $|\mathcal{S}|$ denotes the cardinality of $\mathcal{S}$. For $n \in \mathbb{N}_{>0}$, we denote the n-dimensional identity matrix by I_n, all-ones vector by 1_n, and all-zeros vector by 0_n (subscript may be omitted if the dimension is clear from context).

3.2 Switched-actuator system with setup times

In this section, we formally introduce the class of SAcSSs using a state-space system, a weighted graph, and an actuator selector function.

Definition 3.1. *[109] A simple (i.e., without self-loops or multi-arcs) arc-weighted directed graph (digraph) is defined by the triple $\Gamma = (\mathcal{Q}, \mathcal{E}, s)$, where $\mathcal{Q} = \mathbb{N}_{[1,N_q]}$ is the set of $N_q \in \mathbb{N}_{>0}$ nodes, $\mathcal{E} \subseteq \mathcal{Q}^2 \setminus \{(q,q) \mid q \in \mathcal{Q}\}$ is the set of arcs (directed edges), and $s : \mathcal{E} \to \mathbb{N}$ is an arc-weighting function. If $\mathcal{E} = \mathcal{Q}^2 \setminus \{(q,q) \mid q \in \mathcal{Q}\}$, we say that the simple digraph is complete.*

Example 3.2. *A complete simple arc-weighted digraph Γ with $N_q = 4$ nodes and weighting function*

$$s(q,\tilde{q}) = \begin{cases} 1, & \text{if } (q,\tilde{q}) \in \{(1,3),(3,1)\}, \\ 3, & \text{if } (q,\tilde{q}) \in \{(2,4),(4,2)\}, \\ 2, & \text{otherwise}, \end{cases} \tag{3.1}$$

is depicted in Figure 3.1. Note that although here $s(q,\tilde{q}) = s(\tilde{q},q)$ for all $(q,\tilde{q}) \in \mathcal{E}$, this is not necessary.

Definition 3.3. *A switched-actuator system with setup times (SAcSS) is defined by the quintuple $\Sigma = (\mathcal{X}, \mathcal{U}, f, \Gamma, \Phi)$, where $\mathcal{X} \subseteq \mathbb{R}^{n_x}$ with $n_x \in \mathbb{N}_{>0}$ is the state space, $\mathcal{U} \subseteq \mathbb{R}^{n_u}$ with $n_u \in \mathbb{N}_{>0}$ is the input space, $f : \mathcal{X} \times \mathcal{U} \to \mathcal{X}$ is a state transition function, Γ is a simple arc-weighted digraph, and $\Phi : \mathcal{Q} \cup \mathcal{E} \to \mathcal{P}(\mathbb{N}_{[1,n_u]})$ is an actuator selector function for which it holds that*

$$\Phi(\sigma) \subseteq \begin{cases} \mathbb{N}_{[1,n_u]}, & \text{if } \sigma \in \mathcal{Q}, \\ \emptyset, & \text{if } \sigma \in \mathcal{E}, \end{cases} \tag{3.2}$$

where $\sigma \in \mathcal{Q} \cup \mathcal{E}$ indicates the actuator state.

Some explanation regarding *Definition 3.3* is in order. In a SAcSS, the plant dynamics are given by

$$x_{k+1} = f(x_k, u_k) \tag{3.3}$$

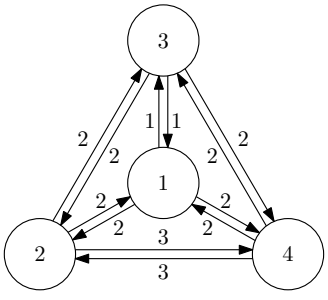

Figure 3.1: Example of a complete simple arc-weighted digraph. The numbered circles represent the nodes, and the numbered arrows represent the weighted arcs.

with state $x_k \in \mathcal{X}$ and input $u_k \in \mathcal{U}$ at discrete time $k \in \mathbb{N}$. Note that a solution of (3.3) is completely defined by an initial condition $x_0 \in \mathcal{X}$ and an input sequence $\boldsymbol{u}_{[0,K]} = (u_0, \ldots, u_K) \in \mathcal{U}^{K+1}$ with $K \in \mathbb{N}$. The input sequence, however, cannot be freely chosen, but must satisfy the restrictions resulting from the actuator switching, as will be defined in the remainder of this section. To this end, the arc-weighted digraph Γ describes the actuator switching behavior. That is, the nodes in Γ represent the N_q *operational* actuator modes $q \in \mathcal{Q}$, and an arc $(q, \tilde{q}) \in \mathcal{E}$ represents the possible actuator mode switch from $q \in \mathcal{Q}$ to $\tilde{q} \in \mathcal{Q}$. Correspondingly, the digraph contains no self-loops (since remaining in an actuator mode requires no switch) or multi-arcs (since we define each possible switch only once), and hence the digraph is simple. The arc weight $s(q, \tilde{q}) \in \mathbb{N}$ represents the setup time associated with the mode switch $(q, \tilde{q}) \in \mathcal{E}$, where the setup time is the number of discrete time instants during which active control is disabled for (3.3), i.e., during which all inputs must be zero. Hence, the weighting function s describes the setup times. We denote the actuator state at time k by $\sigma_k \in \mathcal{Q} \cup \mathcal{E}$, for which it holds that $\sigma_k = q$ when the actuator mode $q \in \mathcal{Q}$ is active, and $\sigma_k = (q, \tilde{q}) \in \mathcal{E}$ when the SAcSS is in the process of switching from mode $q \in \mathcal{Q}$ to $\tilde{q} \in \mathcal{Q}$. We denote the *destination* mode corresponding to σ_k by $\mathrm{post}(\sigma_k)$, where $\mathrm{post} : \mathcal{Q} \cup \mathcal{E} \to \mathcal{Q}$ is defined as $\mathrm{post}(\sigma) = q$ if $\sigma = q \in \mathcal{Q}$ or $\sigma = (\tilde{q}, q) \in \mathcal{E}$ (note that in case of no switch, the destination mode is simply defined as the current mode). If the actuator state describes a mode switch, i.e., $\sigma_k \in \mathcal{E}$, we denote the set of possible actuator states upon finishing the switch by $\mathrm{Post}(\sigma_k)$, where $\mathrm{Post} : \mathcal{E} \to \mathcal{Q} \cup \mathcal{E}$ is defined as $\mathrm{Post}(q, \tilde{q}) = \{\tilde{q}\} \cup \{(\tilde{q}, \bar{q}) \in \mathcal{E} \mid s(\tilde{q}, \bar{q}) > 0\} \cup \{\bar{q} \in \mathcal{Q} \mid (\tilde{q}, \bar{q}) \in \mathcal{E} \text{ and } s(\tilde{q}, \bar{q}) = 0\}$. In other words, given $\sigma_k = (q, \tilde{q}) \in \mathcal{E}$, the SAcSS can either stay in the destination mode $\tilde{q}$ upon arrival at time $l \in \mathbb{N}_{>k}$, indicated by $\sigma_l = \tilde{q}$, or immediately start another switch $(\tilde{q}, \bar{q}) \in \mathcal{E}$, resulting in $\sigma_l = (\tilde{q}, \bar{q})$ or $\sigma_l = \bar{q}$ in case of nonzero or zero setup time $s(\tilde{q}, \bar{q})$, respectively.

Example 3.4. *Consider a SAcSS with* Γ *as in Example 3.2. Then for* $\sigma = (1, 2)$*, the destination mode is given by* $\mathrm{post}(\sigma) = 2$*, and the destination actuator state set is* $\mathrm{Post}(\sigma) = \{2, (2, 1), (2, 3), (2, 4)\}$.

To ensure proper switching behavior, we allow only Σ-*admissible* (also referred to as *admissible* in case Σ is clear from context) actuator sequences, as defined next.

Definition 3.5. *An actuator sequence* $\boldsymbol{\sigma}_{[0,K]} = (\sigma_0, \ldots, \sigma_K) \in (\mathcal{Q} \cup \mathcal{E})^{K+1}$*,* $K \in \mathbb{N}$*, is called* Σ-admissible*, if*

(A) $\sigma_0 \in \mathcal{Q}$*, and,*

(B) if $\sigma_k \neq \sigma_{k-1}$ *with* $\sigma_{k-1} \in \mathcal{Q}$ *and* $\sigma_k \in \mathcal{Q} \cup \mathcal{E}$ *for some* $k \in \mathbb{N}_{[1,K]}$*, then* $\sigma_l = \tilde{\sigma} = (\sigma_{k-1}, \mathrm{post}(\sigma_k))$ *for all* $l \in \mathbb{N}_{[k,\min\{k+s(\tilde{\sigma})-1,K\}]}$*, and, if* $k + s(\tilde{\sigma}) \leq K$*, also* $\sigma_{k+s(\tilde{\sigma})} \in \mathrm{Post}(\tilde{\sigma})$.

In other words, (A) specifies that the SAcSS must initialize in an operational actuator mode $\sigma_0 \in \mathcal{Q}$. By (B), any actuator mode switch that is started must also be completed while adhering to its setup time. In particular, if at some time $k \in \mathbb{N}_{[1,K]}$ the SAcSS starts the switch from $q \in \mathcal{Q}$ to $\tilde{q} \in \mathcal{Q}$, we have $\sigma_{k-1} = q$, and either $\sigma_k = \tilde{q}$ or $\sigma_k = (q, \tilde{q})$ in case of zero or nonzero setup time $s(q, \tilde{q})$, respectively, and hence $\sigma_k \neq \sigma_{k-1}$. Moreover, since $\sigma_k \in \{\tilde{q}, (q, \tilde{q})\}$, we have $\mathrm{post}(\sigma_k) = \tilde{q}$, and hence the considered mode switch can be denoted by $\tilde{\sigma} = (\sigma_{k-1}, \mathrm{post}(\sigma_k))$ with setup time $s(\tilde{\sigma})$. Thus, (B) indeed states that, once started, this switch $\tilde{\sigma}$ must be completed, by defining that the SAcSS must be in the process of switching $\sigma_l = \tilde{\sigma}$ for all times $l \in \mathbb{N}_{[k,k+s(\tilde{\sigma})-1]}$ if $k + s(\tilde{\sigma}) \leq K$ (note that $\mathbb{N}_{[k,k+s(\tilde{\sigma})-1]} = \emptyset$ if $s(\tilde{\sigma}) = 0$, corresponding to an instantaneous mode switch), or for all $l \in \mathbb{N}_{[k,K]}$ if $k + s(\tilde{\sigma}) > K$. In case of the former, upon reaching the destination mode $\tilde{q}$ at time $k + s(\tilde{\sigma}) \leq K$, all possible options for $\sigma_{k+s(\tilde{\sigma})}$ are described by $\mathrm{Post}(\tilde{\sigma})$.

Finally, the actuator selector function Φ in (3.2) selects the subset of input channels that are allowed to be nonzero on the basis of the actuator state. In particular, when at time k the actuator is operational in mode $q \in \mathcal{Q}$, i.e., $\sigma_k = q$, the input channels $u_{k,i}$ with $i \in \Phi(\sigma_k) = \Phi(q) \subseteq \mathbb{N}_{[1,n_u]}$ can be nonzero, whereas $u_{k,j} = 0$ for all $j \in \mathbb{N}_{[1,n_u]} \setminus \Phi(q)$. When, on the other hand, the SAcSS is in the process of switching from a mode $q \in \mathcal{Q}$ to another mode $\tilde{q} \in \mathcal{Q}$, and thus $\sigma_k = (q, \tilde{q}) \in \mathcal{E}$, all inputs are disabled, i.e., $u_{k,i} = 0$ for all $i \in \mathbb{N}_{[1,n_u]} \setminus \Phi(\sigma_k) = \mathbb{N}_{[1,n_u]}$ in this case as $\Phi(\sigma_k) = \Phi(q, \tilde{q}) = \emptyset$. We define the Σ-admissibility of an actuator sequence in combination with a control input sequence as follows.

Definition 3.6. *Let a SAcSS* Σ *be given. Then, a pair* $(\boldsymbol{\sigma}_{[0,K]}, \boldsymbol{u}_{[0,K]})$*,* $K \in \mathbb{N}$*, consisting of an actuator sequence* $\boldsymbol{\sigma}_{[0,K]} \in (\mathcal{Q} \cup \mathcal{E})^{K+1}$ *and a control input sequence* $\boldsymbol{u}_{[0,K]} = (u_0, \ldots, u_K) \in \mathcal{U}^{K+1}$*, is called* Σ-admissible *if* $\boldsymbol{\sigma}_{[0,K]}$ *is a* Σ*-admissible actuator sequence, and for all* $k \in \mathbb{N}_{[0,K]}$ *it holds that* $u_{k,i} = 0$ *for all* $i \in \mathbb{N}_{[1,n_u]} \setminus \Phi(\sigma_k)$.

To facilitate the automated model generation proposed in Section 3.3, we use the following standing assumptions.

Assumption 3.7. *The simple arc-weighted digraph* Γ *is complete, i.e.,* $\mathcal{E} = \mathcal{Q}^2 \setminus \{(q,q) \mid q \in \mathcal{Q}\}$.

Assumption 3.8. *The digraph* Γ *satisfies the triangle inequality [109], which is equivalent to stating that for each sequence* $(q_1, \dots, q_N) \in \mathcal{Q}^N$ *with* $(q_i, q_{i+1}) \in \mathcal{E}$ *for all* $i \in \mathbb{N}_{[1,N-1]}$*, it holds that* $s(q_1, q_N) \leq \sum_{i=1}^{N-1} s(q_i, q_{i+1})$.

Assumption 3.9. *For the actuator selector function* Φ*, it holds that* $\Phi(q) \cap \Phi(\tilde{q}) = \emptyset$ *for all* $q, \tilde{q} \in \mathcal{Q}$*,* $q \neq \tilde{q}$.

Assumption 3.7 describes that a switch from any actuator mode $q \in \mathcal{Q}$ to any other mode $\tilde{q} \in \mathcal{Q}$ is possible, and can in fact be weakened to the assumption of the SAcSS' true mode-switching behavior being described by a strongly connected digraph (i.e., where for every pair $q, \tilde{q} \in \mathcal{Q}$, $q \neq \tilde{q}$, there exists a directed path connecting q to $\tilde{q}$ and vice versa [109]), as explained in *Remark 3.10*.

Remark 3.10. *In case a direct switch from mode* $q \in \mathcal{Q}$ *to* $\tilde{q} \in \mathcal{Q}$ *does not exist, i.e.,* $(q, \tilde{q}) \notin \mathcal{E}$*, the strong connectivity of the digraph guarantees the existence of at least one sequence* $\boldsymbol{q}_{[1,N]} = (q_1, \dots, q_N)$*, where* $q_1 = q$*,* $q_N = \tilde{q}$*, and* $(q_i, q_{i+1}) \in \mathcal{E}$ *for all* $i \in \mathbb{N}_{[1,N-1]}$*. Thus, we can simply define the additional arc* $(q, \tilde{q})$ *with weight* $s(q, \tilde{q}) \coloneqq \min_{N, \boldsymbol{q}_{[1,N]} \in \mathcal{Q}^N} \sum_{i=1}^{N-1} s(q_i, q_{i+1})$*, corresponding to the minimum-time mode sequence by which* q *and* $\tilde{q}$ *can be connected. In this way, a strongly connected digraph can be transformed into a complete digraph without loss of generality.*

Assumption 3.8 states that any switch $(q_1, q_N) \in \mathcal{E}$ cannot induce a setup time larger than that along any other sequence that connects q_1 to q_N, representing the underlying assumption that each switch corresponds to changing the actuator configuration in the fastest way possible.

Finally, *Assumption 3.9* states that the SAcSS has disjoint input channels per actuator mode. This is without loss of generality, as this can always be achieved by simple duplication of the shared input channels, see *Remark 3.11*.

Remark 3.11. *Consider a SAcSS with some input* $w \in \mathcal{W} \subseteq \mathbb{R}^{n_w}$*,* $n_w \in \mathbb{N}_{>0}$ *and corresponding actuator selector function* $\Phi_w : \mathcal{Q} \cup \mathcal{E} \to \mathcal{P}(\mathbb{N}_{[1,n_w]})$*. In case of shared input channels, i.e.,* $\Phi_w(q) \cap \Phi_w(\tilde{q}) \neq \emptyset$ *for some* $q, \tilde{q} \in \mathcal{Q}$*,* $q \neq \tilde{q}$*, we can define the augmented input* $u \in \mathcal{U} \subseteq \mathbb{R}^{n_u}$*, where* $n_u = \sum_{q \in \mathcal{Q}} |\Phi_w(q)| > n_w$*, by duplicating each input channel in* w *as many times as the number of distinct modes in which it is used according to* Φ_w*, and subsequently defining* Φ *such that* u *contains disjoint input channels. For example, we can define* u *as*

$$u_{i+\sum_{j=1}^{q-1} |\Phi_w(j)|} = w_{\Phi_{w,i}(q)} \tag{3.4}$$

for all $i \in \mathbb{N}_{[1,|\Phi_w(q)|]}$ *and* $q \in \mathcal{Q}$*, where* $\Phi_{w,i}(q)$ *denotes the* i*-th element of* $\Phi_w(q)$*, and correspondingly* Φ *as*

$$\Phi(\sigma) = \begin{cases} \mathbb{N}_{[1+\sum_{i=1}^{q-1}|\Phi_w(i)|,\sum_{i=1}^{q}|\Phi_w(i)|]}, & \text{if } \sigma = q \in \mathcal{Q}, \\ \emptyset, & \text{if } \sigma \in \mathcal{E}. \end{cases} \tag{3.5}$$

If w *does not contain shared input channels, then* (3.4)-(3.5) *can be used to reorder the input channels per mode.*

3.3 SAcSS modeling framework

In this section, we present a convenient and systematic procedure for modeling SAcSSs, which directly yields a compact MIP-compatible model. This procedure is specifically designed to (a) avoid the need for introducing many auxiliary discrete (Boolean) variables, which would otherwise increase model complexity and thereby (likely) the corresponding MI-MPC's computational burden, and (b) allow for user-friendly model specification of a SAcSS Σ as in *Definition 3.3*, on the basis of which the MIP-compatible model can be automatically generated. In this new format, the mode-switching behavior is incorporated using mixed-integer linear inequality constraints, requiring only N_q Boolean variables.

3.3.1 Input bounds

By *Assumption 3.9*, the SAcSS has disjoint inputs. Without loss of generality, and in accordance with *Remark 3.11*, let us also assume the input channels in u are ordered per mode according to

$$u = \begin{bmatrix} u^1 \\ \vdots \\ u^{N_q} \end{bmatrix} \in \mathcal{U} \subseteq \mathbb{R}^{n_u}, \tag{3.6}$$

where u^q represents the input of dimension $n_u^q = |\Phi(q)|$ corresponding to mode $q \in \mathcal{Q}$. Note that $\sum_{q\in\mathcal{Q}} n_u^q = n_u$. We define $\underline{u}^q \in \mathbb{R}^{n_u^q}$ and $\overline{u}^q \in \mathbb{R}^{n_u^q}$ as the corresponding lower and upper bounds, resulting in the input sets

$$\mathcal{U}_q = \{u^q \in \mathbb{R}^{n_u^q} \mid \underline{u}^q \leq u^q \leq \overline{u}^q\},\ q \in \mathcal{Q}. \tag{3.7}$$

We collect these bounds in the stacked vectors $\underline{u} = [\underline{u}^{1\top} \ \cdots \ \underline{u}^{N_q\top}]^\top$ and $\overline{u} = [\overline{u}^{1\top} \ \cdots \ \overline{u}^{N_q\top}]^\top$, yielding

$$\mathcal{U} = \{u \in \mathbb{R}^{n_u} \mid \underline{u} \leq u \leq \overline{u}\}. \tag{3.8}$$

3.3.2 Setup time matrix

We construct the *setup time matrix*, being the weighted adjacency matrix of Γ, as

$$S \coloneqq \begin{bmatrix} s_{11} & \cdots & s_{1N_q} \\ \vdots & \ddots & \vdots \\ s_{N_q 1} & \cdots & s_{N_q N_q} \end{bmatrix} \in \mathbb{N}^{N_q \times N_q}, \tag{3.9}$$

with $s_{q\tilde{q}} \coloneqq s(q, \tilde{q})$ for all $(q, \tilde{q}) \in \mathcal{E}$, and where we additionally define $s_{qq} = 0$ for all $q \in \mathcal{Q}$, reflecting the fact that staying in the same mode induces no setup time.

3.3.3 Integer variables and admissibility conditions

3.3.3.1 Activator

To encode the actuator state $\sigma_k \in \mathcal{Q} \cup \mathcal{E}$, we use a one-hot Boolean vector (i.e., with one element equal to 1 and all others 0 [110]), which we refer to as the *activator*. To this end, let $\mathcal{B} = \{0,1\}^{N_q}$ denote the set of N_q Boolean variables, and consider the function $\Delta : \mathcal{Q} \cup \mathcal{E} \to \mathcal{B}$ defined as $\Delta(\sigma) = [\Delta_1(\sigma) \ \cdots \ \Delta_{N_q}(\sigma)]^\top$, where

$$\Delta_q(\sigma) = \begin{cases} 1, & \text{if } \operatorname{post}(\sigma) = q, \\ 0, & \text{otherwise.} \end{cases} \tag{3.10a}$$

The activator δ_k at time k is related to σ_k according to

$$\delta_k = \begin{bmatrix} \delta_k^1 \\ \vdots \\ \delta_k^{N_q} \end{bmatrix} = \begin{bmatrix} \Delta_1(\sigma_k) \\ \vdots \\ \Delta_{N_q}(\sigma_k) \end{bmatrix} = \Delta(\sigma_k) \in \mathcal{B}. \tag{3.10b}$$

The activator indicates which operational mode is (being) activated, i.e., $\delta_k^q = 1$ if $\operatorname{post}(\sigma_k) = q$, meaning at time k the SAcSS is either in the corresponding actuator mode ($\sigma_k = q$) or switching towards this mode ($\sigma_k = (\tilde{q}, q)$), and $\delta_k^q = 0$ otherwise. Clearly, it holds that the activator is a one-hot vector, and hence

$$\sum_{q \in \mathcal{Q}} \delta_k^q = 1. \tag{3.10c}$$

Definition 3.12. *An activator sequence $\boldsymbol{\delta}_{[0,K]} = (\delta_0, \ldots, \delta_K) \in \mathcal{B}^{K+1}$, $K \in \mathbb{N}$, is called* Σ-admissible *if there is a Σ-admissible actuator sequence $\boldsymbol{\sigma}_{[0,K]}$ with $\delta_k = \Delta(\sigma_k)$ for all $k \in \mathbb{N}_{[0,K]}$.*

Definition 3.13. *A pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$, $K \in \mathbb{N}$, is called* Σ-admissible *if there exists a Σ-admissible actuator sequence $\boldsymbol{\sigma}_{[0,K]}$ such that $\delta_k = \Delta(\sigma_k)$ for all $k \in \mathbb{N}_{[0,K]}$ (i.e., $\boldsymbol{\delta}_{[0,K]}$ is a Σ-admissible activator sequence) and the pair $(\boldsymbol{\sigma}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ is Σ-admissible.*

Note that while the actuator state σ_k differentiates between the SAcSS having an operational actuator mode activated ($\sigma_k \in \mathcal{Q}$) or being in transition between two modes ($\sigma_k \in \mathcal{E}$), the activator δ_k effectively lumps these two actuator states together due to the function Δ considering only the destination mode $\text{post}(\sigma_k)$. This property is instrumental in reducing the number of required Boolean variables in the model. That is, explicitly modeling also the states $\sigma_k \in \mathcal{E}$, as in a constrained switched (linear) system [81,106] in mixed logical dynamical (MLD) form [104] for instance, would require additional Boolean variables. Moreover, since the switching is to be determined by MPC, where a performance-related stage cost must be evaluated at each prediction instant, switches $(q, \tilde{q}) \in \mathcal{E}$ with $s(q, \tilde{q}) > 1$ even require multiple additional Boolean variables to model the different levels of switching progress, which is referred to as lifting [107]. As a result, such approaches require $\sum_{(q,\tilde{q})\in\mathcal{E}} s_{q\tilde{q}}$ Boolean variables (which grows for increasing setup times) in addition to the N_q Boolean variables that model the operational actuator states $\sigma_k \in \mathcal{Q}$. By contrast, our approach leads to a more compact model using only the N_q Boolean variables in the activator, independent of the setup times. Since, however, the activator δ_K at some time $K \in \mathbb{N}$ does not indicate the actuator state σ_K directly, see (3.10), we instead infer σ_K, and thereby the admissible input u_K according to *Definition 3.6*, from a *sequence* of current and past activators in $\boldsymbol{\delta}_{[0,K]}$. Note that this essentially exchanges the need for many Boolean variables per time instant against the need for a few Boolean variables over multiple instants, and that an activator δ_k, $k \in \mathbb{N}$, is required to determine σ_l at multiple (equal and later) times $l \in \mathbb{N}_{\geq k}$. In the context of MPC, however, the decision variables such as the activators must be considered at all time instants in the prediction horizon anyway. Hence, including only the N_q-dimensional Boolean activator per time step, and effectively "reusing" each (predicted) activator to determine multiple (predicted) actuator states, leads to a MIP problem with significantly less Boolean variables, thereby improving efficiency.

3.3.3.2 Admissibility conditions

To explain the logic underlying the relation between $\boldsymbol{\sigma}_{[0,K]}$ and $\boldsymbol{\delta}_{[0,K]}$, let us first introduce

$$\mathcal{Q}^q_{<\tau} = \{\tilde{q} \in \mathcal{Q} \mid s_{\tilde{q}q} < \tau\}, \quad \tau \in \mathbb{N}_{>0},\ q \in \mathcal{Q}, \tag{3.11}$$

denoting the set of modes $\tilde{q} \in \mathcal{Q}$ from which the switch to mode $q \in \mathcal{Q}$ induces a setup time smaller than τ instants. Note that $q \in \mathcal{Q}^q_{<1}$ as $s_{qq} = 0$ by definition, $\mathcal{Q}^q_{<\tau_1} \subseteq \mathcal{Q}^q_{<\tau_2}$ for positive times $\tau_1 < \tau_2$, and $\mathcal{Q}^q_{<\tau} = \mathcal{Q}$ for all $\tau > \max_{\tilde{q}\in\mathcal{Q}} s_{\tilde{q}q}$ (due to *Assumption 3.7* and *Assumption 3.8*). Using (3.11), we can express the equality $\sigma_K = q$ in terms of the activator sequence $\boldsymbol{\delta}_{[0,K]}$ as follows.

Lemma 3.14. *For a Σ-admissible actuator sequence $\boldsymbol{\sigma}_{[0,K]}$ and corresponding activator sequence $\boldsymbol{\delta}_{[0,K]}$, it holds that $\sigma_K = q$ if and only if:*

(i) $\delta^q_K = 1$, *and,*

(ii) for all $\tau \in \mathbb{N}_{[1,\min\{K,\overline{s}_q\}]}$, *where* $\overline{s}_q = \max_{\tilde{q}\in\mathcal{Q}} s_{\tilde{q}q}$, *there exists* $\tilde{q} \in \mathcal{Q}^q_{<\tau}$ *such that* $\delta^{\tilde{q}}_{K-\tau} = 1$.

Proof 3.15. *Note that by Assumption 3.7, combined with Assumption 3.8, we have that given some* $\tilde{q} \in \mathcal{Q} \setminus \mathcal{Q}^q_{<\tau}$ *there exists no admissible actuator sequence connecting* $\tilde{q}$ *to* q *in less than* τ *steps. Consequently, in the remainder of this proof, inspecting which admissible past actuator sequences can lead to* $\sigma_K = q$ *can be done entirely on the basis of* $\mathcal{Q}^q_{<\tau}$, *and thereby on only the actuator sequences describing "direct" switches to* q, *while disregarding all actuator sequences by which the SAcSS reaches* $\sigma_K = q$ *"indirectly" (i.e., by first switching to some other mode before finally switching to* q*).*

(Only if.) *We will first show that for admissible* $\boldsymbol{\sigma}_{[0,K]}$ *and* $\boldsymbol{\delta}_{[0,K]}$, *the equality* $\sigma_K = q$ *implies conditions* (i)-(ii) *in the lemma. By admissibility of* $\boldsymbol{\delta}_{[0,K]}$, *we can use* (3.10) *to find*

$$\sigma_K = q \;\Rightarrow\; \delta^q_K = \Delta_q(\sigma_K) = 1, \tag{3.12}$$

and thus condition (i) *is satisfied. Next, by admissibility of* $\boldsymbol{\sigma}_{[0,K]}$, $\sigma_K = q$ *implies that*

$$\sigma_{K-\tau} \in \mathcal{Q}^q_{<\tau} \cup \{(\tilde{q},\bar{q}) \in \mathcal{E} \mid \bar{q} \in \mathcal{Q}^q_{<\tau}\} \tag{3.13}$$

for all $\tau \in \mathbb{N}_{[1,K]}$, *i.e, the SAcSS must at all times* $k = K - \tau$, $\tau \in \mathbb{N}_{[1,K]}$, *have been in or transitioning towards a mode from which the switch to mode* q *would induce a setup time of less than* τ *time instants (as otherwise the SAcSS would not be able to reach mode* q *at time* K*). By admissibility of* $\boldsymbol{\delta}_{[0,K]}$, *we can apply the function* Δ *from* (3.10) *to find that* (3.13) *implies that for all* $\tau \in \mathbb{N}_{[1,K]}$ *there exists* $\tilde{q} \in \mathcal{Q}^q_{<\tau}$ *such that* $\delta^{\tilde{q}}_{K-\tau} = 1$, *and thus condition* (ii) *holds.*

(If.) *By admissibility of* $\boldsymbol{\delta}_{[0,K]}$, *we use* (3.10) *to derive from condition* (i)

$$\delta^q_K = 1 \;\Leftrightarrow\; \mathrm{post}(\sigma_K) = q \;\Leftrightarrow\; \sigma_K \in \{q\} \cup \{(\hat{q},q) \in \mathcal{E}\},$$

i.e., the SAcSS was either in or switching towards mode q *at time* K*. We will complete this proof by showing that condition* (ii) *excludes the latter possibility. To this end, note that for any admissible actuator sequence with* $\sigma_K = (\hat{q},q) \in \mathcal{E}$, *it must hold that* $\mathrm{post}(\sigma_{K-\tau}) = \hat{q} \in \mathcal{Q} \setminus \mathcal{Q}^q_{<\tau}$ *for some* $\tau \in \mathbb{N}_{[1,\min\{K,s_{\hat{q}q}\}]}$, *followed by* $\sigma_k = (\hat{q},q)$ *for all* $k \in \mathbb{N}_{[K-\tau+1,K]}$ *(describing that the SAcSS was in or arriving at mode* $\hat{q}$ *at time* $K-\tau$, *and that the switch* $(\hat{q},q)$ *started at time* $K-\tau+1$ *and is not yet completed at time* K*). Using* (3.10b) *by admissibility of* $\boldsymbol{\delta}_{[0,K]}$, *this is equivalent to* $\delta^{\hat{q}}_{K-\tau} = 1$ *for some* $\tau \in \mathbb{N}_{[1,\min\{K,s_{\hat{q}q}\}]}$ *and* $\hat{q} \notin \mathcal{Q}^q_{<\tau}$ *(and* $\delta^q_k = 1$ *for all* $k \in \mathbb{N}_{[K-\tau+1,K]}$*). Using* (3.10c), *however, we find that this is contradicted by condition* (ii), *which completes the proof.* □

Remark 3.16. *In the first (only if) part of the above proof, we show that condition* (ii) *holds for all* $\tau \in \mathbb{N}_{[1,K]}$*, whereas in Lemma 3.14 condition* (ii) *only includes all* $\tau \in \mathbb{N}_{[1,\overline{s}_q]}$*,* $q \in \mathcal{Q}$*. Note that these conditions are in fact equivalent, since for all* $\tau > \overline{s}_q$ *it holds that* $\mathcal{Q}^q_{<\tau} = \mathcal{Q}$*, and thus the existence of a mode* $\tilde{q} \in \mathcal{Q}$ *such that* $\delta^{\tilde{q}}_{K-\tau} = 1$ *for all such* τ *is satisfied by definition, see* (3.10c).

To further clarify *Lemma 3.14*, and in particular the first (only if) part of its proof, consider *Example 3.17*.

Example 3.17. *Consider a SAcSS with* Γ *as in Example 3.2. For an admissible actuator sequence with* $\sigma_K = 1$*, the possible actuator states at the preceding time instants are depicted in Figure 3.2, where the possible actuator states at each time are highlighted (black). If* $\sigma_K = 1$ *(Figure 3.2a), then at time* $K-1$ *by admissibility the SAcSS must have either already been in mode 1 or in the last time instant of switching towards this mode, i.e.,* $\sigma_{K-1} \in \{1\} \cup \{(q,1) \in \mathcal{E}\}$ *(Figure 3.2b). Consequently, taking the setup times into account, at time* $K-2$ *the SAcSS may have been in modes 1 or 3 or in the last step of switching towards these modes, i.e.,* $\sigma_{K-1} \in \{1,3\} \cup \{(q,\tilde{q}) \in \mathcal{E} \mid \tilde{q} \in \{1,3\}\}$ *(Figure 3.2c), by which in turn at time* $K-3$ *the SAcSS may have been in any actuator state, i.e.,* $\sigma_{K-3} \in \mathcal{Q} \cup \mathcal{E}$ *(Figure 3.2d). Let us now realize that this in fact demonstrates* (3.13)*, which is the key result in the "only if" part of the proof of Lemma 3.14, since for mode 1 we find* $\mathcal{Q}^1_{<1} = \{1\}$*,* $\mathcal{Q}^1_{<2} = \{1,3\}$*, and* $\mathcal{Q}^1_{<3} = \mathcal{Q}$*.*

For the second (if) part of the proof, the conditions (i)-(ii) may be interpreted as follows. Satisfying (ii) ensures that any switch $(\tilde{q},q) \in \mathcal{E}$ with setup time $s_{\tilde{q}q}$ can only have started at time $K - s_{\tilde{q}q} - 1$ at the latest. By admissibility, any such switch is then completed at time K. Condition (i) is equivalent to $\mathrm{post}(\sigma_K) = q$, thus excluding the possibility that the SAcSS initiates another switch away from mode q at time K, resulting in conditions (i)-(ii) together to imply $\sigma_K = q$.

We have shown that the conditions in *Lemma 3.14* are equivalent to the statement $\sigma_K = q$, which by *Definition 3.6* is directly related to the admissible values of u_K. Therefore, using *Lemma 3.14* we can formulate explicit admissibility conditions on $\boldsymbol{u}_{[0,K]}$ in relation to $\boldsymbol{\delta}_{[0,K]}$, as is done in the following theorem.

Theorem 3.18. *Given an admissible activator sequence* $\boldsymbol{\delta}_{[0,K]}$*, the pair* $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ *is admissible if and only if for all* $k \in \mathbb{N}_{[0,K]}$ *it holds that* $u^{\tilde{q}}_k = 0_{n^{\tilde{q}}_u}$ *for all* $\tilde{q} \in \mathcal{Q} \setminus \{q\}$*, where* $q \in \mathcal{Q}$ *satisfies*

(I) $\delta^q_k = 1$*, and,*

(II) for all $\tau \in \mathbb{N}_{[1,\min\{k,\overline{s}_q\}]}$*, there exists* $\tilde{q} \in \mathcal{Q}^q_{<\tau}$ *such that* $\delta^{\tilde{q}}_{k-\tau} = 1$*.*

Proof 3.19. *By admissibility of* $\boldsymbol{\delta}_{[0,K]}$ *there exists a corresponding admissible actuator sequence* $\boldsymbol{\sigma}_{[0,K]}$*. Using Lemma 3.14, the conditions* (I)-(II) *are equivalent to* $\sigma_k = q$*, and consequently Theorem 3.18 restricts* u_k *exactly to its admissible set as in Definition 3.6. Hence, satisfying the conditions in Theorem 3.18*

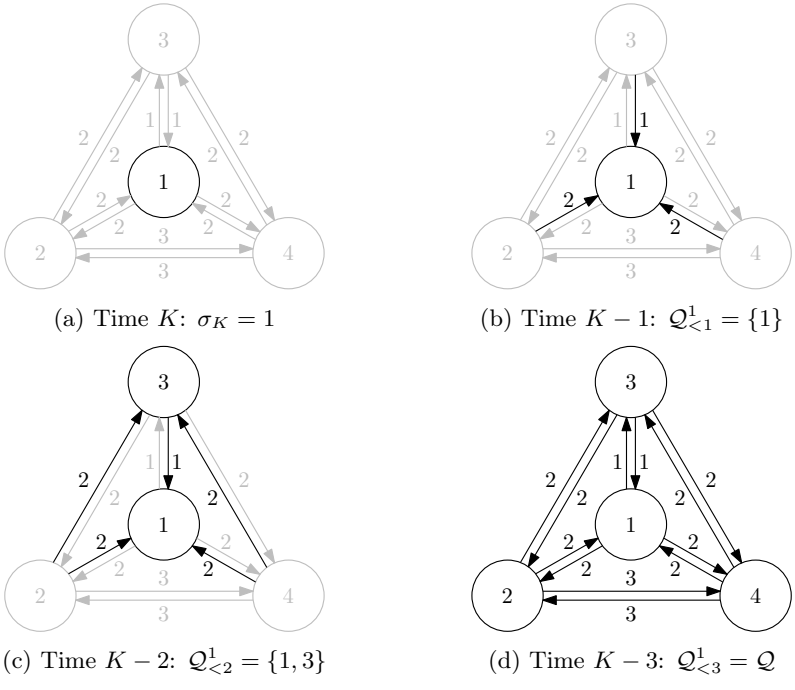

Figure 3.2: Possible actuator states (black) at different times, given an admissible actuator sequence with $\sigma_K = 1$.

on $\boldsymbol{u}_{[0,K]}$ *in relation to an admissible* $\boldsymbol{\delta}_{[0,K]}$ *is equivalent to the admissibility of the pair* $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$. □

3.3.4 Mixed-integer linear inequality constraints

In this subsection, we use *Theorem 3.18* to formulate the mixed-integer linear inequality constraints (in terms of the inputs and activators) that incorporate the actuator switching behavior in a manner that ensures admissibility of the pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ for a *given* $\boldsymbol{\delta}_{[0,K]}$. Then, we discuss a significant simplification of these constraints for systems with one-sided controls.

3.3.4.1 General form

Recall that $\underline{u}^q$ and $\overline{u}^q$ denote the lower and upper input bounds of the inputs in mode $q \in \mathcal{Q}$ as in (3.7). For some $k \in \mathbb{N}$, the linear inequality constraints that ensure that $u_k^{\tilde{q}} = 0_{n_u^{\tilde{q}}}$ for all $\tilde{q} \in \mathcal{Q} \setminus \{q\}$, where $q \in \mathcal{Q}$ satisfies condition (I) of

Theorem 3.18, are then given by

$$\delta_k^q \underline{u}^q \leq u_k^q \leq \delta_k^q \overline{u}^q \quad \text{for all } q \in \mathcal{Q}, \tag{3.14}$$

indeed allowing for nonzero $u_k^q \in \mathcal{U}_q$ only when $\delta_k^q = 1$. Collecting all modal input bounds (3.14) results in the constraints on u_k corresponding to condition (I) to read

$$\underline{U}_0 \delta_k \leq u_k \leq \overline{U}_0 \delta_k, \tag{3.15a}$$

where the $n_u \times N_q$-dimensional block diagonal input constraint matrices are given by

$$\underline{U}_0 = \begin{bmatrix} \underline{u}^1 & & \\ & \ddots & \\ & & \underline{u}^{N_q} \end{bmatrix}, \quad \overline{U}_0 = \begin{bmatrix} \overline{u}^1 & & \\ & \ddots & \\ & & \overline{u}^{N_q} \end{bmatrix}. \tag{3.15b}$$

Note that the products $\underline{U}_0 \delta_k$ and $\overline{U}_0 \delta_k$ contain mostly zeros, except for the entries corresponding to $\delta_k^q = 1$.

By the same reasoning, the input restrictions resulting from condition (II) of *Theorem 3.18* for some $k \in \mathbb{N}$ can equivalently be expressed as requiring for all $q \in \mathcal{Q}$ that

$$\sum_{\tilde{q} \in \mathcal{Q}^q_{<\tau}} \delta_{k-\tau}^{\tilde{q}} \underline{u}^q \leq u_k^q \leq \sum_{\tilde{q} \in \mathcal{Q}^q_{<\tau}} \delta_{k-\tau}^{\tilde{q}} \overline{u}^q \tag{3.16}$$

for all $\tau \in \mathbb{N}_{[1,\min\{k,\overline{s}_q\}]}$. To automate the generation of the inequality constraints (3.16) for the entire input vector, we first use the setup time matrix S from (3.9) to construct the setup time constraint matrices

$$S_\tau \coloneqq \begin{bmatrix} s_{\tau,11} & \cdots & s_{\tau,1N_q} \\ \vdots & \ddots & \vdots \\ s_{\tau,N_q1} & \cdots & s_{\tau,N_qN_q} \end{bmatrix} \in \{0,1\}^{N_q \times N_q}, \quad \tau = 1, \ldots, \overline{s}, \tag{3.17a}$$

with largest setup time $\overline{s} \coloneqq \max_{(q,\tilde{q}) \in \mathcal{E}} s_{q\tilde{q}}$, and

$$s_{\tau,q\tilde{q}} \coloneqq \begin{cases} 1, & \text{if } s_{q\tilde{q}} < \tau, \\ 0, & \text{otherwise.} \end{cases} \tag{3.17b}$$

The input constraints from (II) are then given by

$$\underline{U}_\tau \delta_{k-\tau} \leq u_k \leq \overline{U}_\tau \delta_{k-\tau} \text{ for all } \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}\}]}, \tag{3.18a}$$

with input constraint matrices

$$\underline{U}_\tau = \begin{bmatrix} s_{\tau,11}\underline{u}^1 & \cdots & s_{\tau,N_q1}\underline{u}^1 \\ \vdots & \ddots & \vdots \\ s_{\tau,1N_q}\underline{u}^{N_q} & \cdots & s_{\tau,N_qN_q}\underline{u}^{N_q} \end{bmatrix}, \tag{3.18b}$$

$$\overline{U}_\tau = \begin{bmatrix} s_{\tau,11}\overline{u}^1 & \cdots & s_{\tau,N_q1}\overline{u}^1 \\ \vdots & \ddots & \vdots \\ s_{\tau,1N_q}\overline{u}^{N_q} & \cdots & s_{\tau,N_qN_q}\overline{u}^{N_q} \end{bmatrix}. \tag{3.18c}$$

In summary, (3.15) and (3.18) translate the conditions in *Theorem 3.18* into mixed-integer linear inequality constraints that are directly suitable for MI-MPC. Precisely, if we have an admissible pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$, then this pair satisfies the constraints in (3.15) and (3.18). Moreover, these constraints are constructed automatically on the basis of $\underline{u}$, $\overline{u}$, Φ, and S, and result in $n_u(\overline{s}+1)$ (two-sided) linear inequalities.

Remark 3.20. *In case of non-uniform maximum setup times per mode, i.e., if $\overline{s}_q \neq \overline{s}_{\tilde{q}}$ for some $q, \tilde{q} \in \mathcal{Q}$ and thus $\overline{s}_q < \overline{s}$ for some $q \in \mathcal{Q}$, the number of inequality constraints can be reduced by omitting the constraints that follow from all-ones columns in the matrices S_τ, $\tau = 1, \ldots, \overline{s}$. That is, these constraints are a byproduct of automatically constructing* (3.16) *in the form of* (3.18)*, and in fact correspond to* (3.16) *for $\tau > \overline{s}_q$, which are not included in condition* (II) *(and which are satisfied by definition, see Remark 3.20). In many numerical solvers, however, these constraints are automatically removed using a presolve algorithm, which eliminates the need to explicitly remove them during our modeling procedure.*

Remark 3.21. *If the system dynamics* (3.3) *are linear, then the combination of* (3.3) *with the linear inequality constraints* (3.15) *and* (3.18) *in terms of continuous and Boolean variables in fact describes a SAcSS as a MLD system [104]. Contrary to the general-purpose discrete hybrid automaton (DHA) framework (and accompanying modeling language HYSDEL which automates the conversion of a DHA into MLD form) [105], our method is tailored to SAcSSs. As a result, it allows for easier high-level specification of SAcSSs, and enables the automated generation of integer-wise compact (MLD) models.*

3.3.4.2 Simplifications for one-sided control

The number of input constraints in (3.18) may be reduced in case of one-sided input channels, i.e., when $\underline{u}_i = 0$ or $\overline{u}_i = 0$ for some $i \in \mathbb{N}_{[1,n_u]}$. This is the case in MR-HIFU hyperthermia, for example, where due to physical limitations heat cannot actively be removed from the system, and thus $u_{k,i} \geq \underline{u}_i = 0$ for all $k \in \mathbb{N}$. For ease of exposition, we discuss here the simplification for the case that all inputs are nonnegative, i.e., $\underline{u} = 0_{n_u}$. All other cases can be treated mutatis mutandis.

The key property that enables the reduction of the number of inequality constraints is that by the one-sidedness of the controls, the logic of condition (II) can be imposed on the *sum* of the one-sided inputs per mode, as opposed to on the individual input elements. That is, if $\underline{u}^q = 0$ for all $q \in \mathcal{Q}$, and additionally

(3.14) is imposed to restrict the individual input elements, then (II) (and thereby (3.16)) is equivalent to requiring for all $q \in \mathcal{Q}$ that

$$1_{n_u^q}^\top u_k^q \leq \sum_{\tilde{q} \in \mathcal{Q}_{<\tau}^q} \delta_{k-\tau}^{\tilde{q}} 1_{n_u^q}^\top \overline{u}^q \text{ for all } \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}_q\}]}.$$

Consequently, we may replace (3.18a) by the inequalities

$$J_u u_k \leq J_u \overline{U}_\tau \delta_{k-\tau} \text{ for all } \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}\}]}, \tag{3.19a}$$

where J_u represents the block diagonal matrix

$$J_u = \begin{bmatrix} 1_{n_u^1}^\top & & \\ & \ddots & \\ & & 1_{n_u^{N_q}}^\top \end{bmatrix} \in \{0,1\}^{N_q \times n_u}. \tag{3.19b}$$

Together, (3.15) and (3.19) comprise the reduced set of inequality constraints that incorporate the SAcSS' switching including setup times.

Remark 3.22. *Note that if, in addition, all modes have a common bound on the sum of their inputs, i.e.,* $1_{n_u^q}^\top u^q \leq \overline{u_\Sigma}$ *with* $\overline{u_\Sigma} \in \mathbb{R}_{>0}$ *must hold for all* $q \in \mathcal{Q}$*, then* (3.19) *can be written as*

$$J_u u_k \leq \overline{u_\Sigma} S_\tau^\top \delta_{k-\tau} \quad \textit{for all } \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}\}]}. \tag{3.20}$$

Summarizing, the $n_u \overline{s}$ two-sided inequality constraints in (3.18a) can be replaced by only $N_q \overline{s}$ single inequalities in (3.19) (or (3.20)), achieving a reduction of $(2n_u - N_q)\overline{s}$ single inequalities. Especially for SAcSSs with high input dimension and only few modes, such as in MR-HIFU hyperthermia, and in the case of large maximum setup time $\overline{s}$, this leads to a significant reduction.

Remark 3.23. *In many numerical solvers, one needs to specify the upper and lower bounds of the search space for the optimization variables, by which the constraints* $\underline{u} \leq u_k \leq \overline{u}$ *are already included. Then, the above simplification procedure may also be applied to* (3.15)*, yielding*

$$J_u u_k \leq \overline{u_\Sigma} \delta_k. \tag{3.21}$$

In this case, the $n_u(\overline{s}+1)$ *two-sided inequalities in* (3.15) *and* (3.18) *are reduced to the* $N_q(\overline{s}+1)$ *one-sided inequalities in* (3.19) *and* (3.21)*, resulting in a reduction of* $(2n_u - N_q)(\overline{s}+1)$ *constraints.*

3.3.5 Admissibility assurance

The inequalities (3.15) and (3.18) (or, if applicable, (3.21) and (3.19) or (3.20)) provide a set of admissibility conditions on $\boldsymbol{u}_{[0,K]}$ in relation to $\boldsymbol{\delta}_{[0,K]}$. Being

based on *Theorem 3.18*, these constraints rely on the assumption that the activator sequence $\boldsymbol{\delta}_{[0,K]}$ is *admissible*. To ensure the admissibility of $\boldsymbol{\delta}_{[0,K]}$ *a priori* in the context of MPC, where the pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ is to be determined by optimization, one could try to impose restrictions on $\boldsymbol{\delta}_{[0,K]}$. However, for compatibility with MI-MPC, this must then be done using additional mixed-integer linear inequality constraints, of which the derivation is not straightforward and due to linearization may require the introduction of many auxiliary integer variables. To avoid the need for this, we instead opt to ensure the admissibility of $\boldsymbol{\delta}_{[0,K]}$ *a posteriori*. In particular, we will use only (3.15) and (3.18) directly in the MPC to obtain a pair $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ that satisfies the conditions in *Theorem 3.18*, where $\tilde{\boldsymbol{\delta}}_{[0,K]}$ need not be admissible. Subsequently, on the basis of $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ we construct (using a simple procedure) a potentially different activator sequence $\boldsymbol{\delta}_{[0,K]}$ such that $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ is admissible for the same $\boldsymbol{u}_{[0,K]}$. We call any such pair $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ Σ-*feasible* (or *feasible* if Σ is clear from context), as defined next.

Definition 3.24. *A pair* $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$, $K \in \mathbb{N}$, *is called* Σ-feasible *if it satisfies the conditions on* $\boldsymbol{u}_{[0,K]}$ *in relation to* $\boldsymbol{\delta}_{[0,K]}$ *given in Theorem 3.18.*

The construction of an alternative sequence $\boldsymbol{\delta}_{[0,K]}$ such that the pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ is Σ-admissible, given a Σ-feasible pair $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$, is based on the following theorem.

Theorem 3.25. *Let a* Σ*-feasible pair* $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ *be given. Define* $\mathcal{K} = \{k_1, \dots, k_L\} = \{k \in \mathbb{N}_{[0,K]} \mid u_k \neq 0\}$, $L \in \mathbb{N}$, *where* $k_l < k_{l+1}$ *for all* $l \in \mathbb{N}_{[1,L-1]}$, *and construct the activator sequence* $\boldsymbol{\delta}_{[0,K]}$ *as*

$$\delta_k = \begin{cases} \tilde{\delta}_{k_1}, & \text{if } k \in \mathbb{N}_{[0,k_1]}, \\ \tilde{\delta}_{k_l}, & \text{if } k \in \mathbb{N}_{[k_{l-1}+1,k_l]},\ l \in \mathbb{N}_{[2,L]}, \\ \tilde{\delta}_{k_L}, & \text{if } k \in \mathbb{N}_{[k_L+1,K]}. \end{cases} \tag{3.22}$$

Then, the pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ *is* Σ*-admissible.*

Proof 3.26. *Note that* $\mathcal{K}$ *denotes the chronologically ordered set of* L *time instants at which the control inputs are nonzero. Let* q_l *with* $l \in \mathbb{N}_{[1,L]}$ *denote the mode with nonzero input at time* $k_l \in \mathcal{K}$, *i.e., such that* $u_{k_l}^{q_l} \neq 0$. *By feasibility of* $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$, *condition* (II) *of Theorem 3.18 holds for all* $k \in \mathbb{N}_{[0,K]}$. *Combined with* (3.10c), *this implies that* $q_l \in \mathcal{Q}^{q_{l+1}}_{<k_{l+1}-k_l}$ *for all* $l \in \mathbb{N}_{[1,L-1]}$, *and therefore*

$$s_{q_l q_{l+1}} < k_{l+1} - k_l \quad \text{for all } l \in \mathbb{N}_{[1,L-1]}, \tag{3.23}$$

by which we can construct the actuator sequence $\boldsymbol{\sigma}_{[0,K]}$ *as*

$$\sigma_k = \begin{cases} q_1, & \text{if } k \in \mathbb{N}_{[0,k_1]}, \\ (q_l, q_{l+1}), & \text{if } k \in \mathbb{N}_{[k_l+1,k_l+s_{q_l q_{l+1}}]},\ l \in \mathbb{N}_{[1,L-1]}, \\ q_l, & \text{if } k \in \mathbb{N}_{[k_{l-1}+s_{q_{l-1}q_l}+1,k_l]},\ l \in \mathbb{N}_{[2,L]}, \\ q_L, & \text{if } k \in \mathbb{N}_{[k_L,K]}. \end{cases}$$

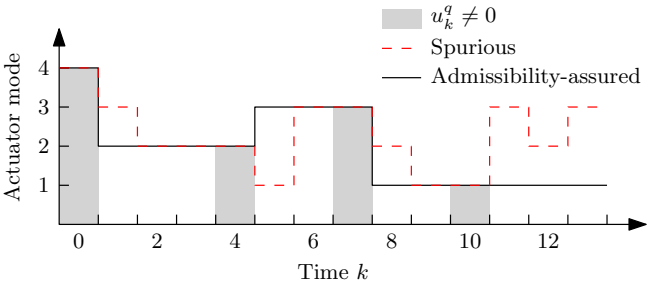

Figure 3.3: Example of $\text{post}(\sigma_k)$ related to a spurious pair $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ (dashed red) and the corresponding admissibility-assured pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ (solid black), alongside the mode with nonzero inputs in $\boldsymbol{u}_{[0,K]}$ (gray).

First, note that this sequence $\boldsymbol{\sigma}_{[0,K]}$ is admissible according to Definition 3.5, describing the SAcSS (a) starting and staying in mode $q_1 \in \mathcal{Q}$ during the times $k \in \mathbb{N}_{[0,k_1]}$, (b) performing each switch (q_l, q_{l+1}), $l \in \mathbb{N}_{[1,L-1]}$, during the interval $k \in \mathbb{N}_{[k_l+1, k_l+s(q_l,q_{l+1})]}$ in case $q_l \neq q_{l+1}$, or remaining in the same mode during the interval $k \in \mathbb{N}_{[k_l,k_{l+1}]}$ if $q_l = q_{l+1}$ (since then $s_{q_l q_{l+1}} = 0$ in the above construction of $\boldsymbol{\sigma}_{[0,K]}$), and (c) staying in the last active mode q_L during $k \in \mathbb{N}_{[k_L,K]}$. Since, additionally, $\sigma_{k_l} = q_l$ for all $l \in \mathbb{N}_{[1,L]}$, the pair $(\boldsymbol{\sigma}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ is admissible according to Definition 3.6. Finally, note that for $\boldsymbol{\delta}_{[0,K]}$ from (3.22) it holds that $\delta_k = \Delta(\sigma_k)$ for all $k \in \mathbb{N}_{[0,K]}$, which proves the admissibility of the pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ according to Definition 3.13. □

We call the procedure in *Theorem 3.25* the *admissibility assurance*. Note that this procedure is simple, and thus computationally lightweight, and can be performed automatically. The activator sequence constructed using the admissibility assurance corresponds to the actuator sequence that enables $\boldsymbol{u}_{[0,K]}$ in an admissible manner, using the minimum number of mode switches, and where each switch occurs as early as possible.

Example 3.27. *For a SAcSS with Γ as in Example 3.2 and initial state $\sigma_0 = 4$, Figure 3.3 shows the destination mode $\text{post}(\sigma_k)$ corresponding to a pair $(\tilde{\boldsymbol{\delta}}_{[0,K]}, \boldsymbol{u}_{[0,K]})$, $K = 13$, that is feasible but not admissible (dashed red), referred to as* spurious*, and the destination mode $\text{post}(\sigma_k)$ corresponding to the resulting admissibility-assured pair $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ obtained via (3.22) (solid black). By inspecting the figure in relation to the setup times defined in the previous examples, we find that the admissibility-assured sequence allows the SAcSS to apply the input sequence $\boldsymbol{u}_{[0,K]}$ (gray) while switching as little as possible (see all switches and the last three time instants in Figure 3.3) and as soon as possible (see the last required switch at $k \in \mathbb{N}_{[7,10]}$).*

One may note that this is only one particular choice out of a potentially larger

set of activator sequences that render $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ admissible, as discussed in *Remark 3.28.*

Remark 3.28. *Consider the integer programming problem*

$$\min_{\boldsymbol{\delta}_{[0,K]}} \sum_{k=0}^{K-1} d_k(\delta_{k+1} - \delta_k)^\top (\delta_{k+1} - \delta_k), \tag{3.24}$$

subject to the constraints (3.15) *and* (3.18) *(or* (3.19) *or* (3.20)*) for all* $k \in \mathbb{N}_{[0,K]}$ *with* $\boldsymbol{u}_{[0,K]}$ *given, where* $d_k \in \mathbb{R}_{>0}$, $k \in \mathbb{N}_{[0,K]}$, *represent scalar weights. The activator sequence constructed using the admissibility assurance* (3.22) *is in fact the optimal solution to* (3.24) *in case* $d_k < d_{k+1}$ *for all* $k \in \mathbb{N}_{[0,K-1]}$. *Indeed, the constraints* (3.15) *and* (3.18) *then ensure feasibility, the minimization of the number of switches by* (3.24) *ensures admissibility, and the strict increase of the weights* d_k *(with respect to* k*) favors switching as early as possible. Since, however,* $d_k \neq d_l$ *for all* $k, l \in \mathbb{N}_{[1,K]}$, $k \neq l$, *is a sufficient condition for the minimizer of* (3.24) *to be unique,* (3.24) *is a generalization of the admissibility assurance, and can be used to uniquely construct other minimum-switching activator sequences by which* $(\boldsymbol{\delta}_{[0,K]}, \boldsymbol{u}_{[0,K]})$ *is admissible, reflecting different switching time preferences (e.g.,* $d_k > d_{k+1}$ *for all* $k \in \mathbb{N}_{[0,K-1]}$ *for switching as late as possible). Note that for such different switching preferences, the procedure* (3.22) *can also be easily augmented to quickly construct the corresponding minimum-switching activator sequence.*

3.4 Mixed-integer MPC for SAcSSs

In this section, we first show how a SAcSS model derived using the procedure discussed in the previous section can be directly integrated into MI-MPC. Then, we present the control algorithm, taking into account that re-optimizing during a switch is redundant, as all switches must be completed by definition.

3.4.1 Optimal control problem

The MI-MPC based on the previously derived SAcSS model is given by

$$\min_{\boldsymbol{u}_k, \boldsymbol{\delta}_k} V_N(x_k, \boldsymbol{u}_k) \tag{3.25a}$$

subject to $\boldsymbol{u}_k = (u_{0|k}, \dots, u_{N-1|k}) \in \boldsymbol{\mathcal{U}}_N(x_k)$, $\boldsymbol{\delta}_k = (\delta_{0|k}, \dots, \delta_{N-1|k}) \in \mathcal{B}^N$, and

$$\left[0_{N \times N n_u} \;\middle|\; O_1(k) \quad I_N \otimes 1_{N_q}^\top\right] \left[\frac{\mathbf{u}_k}{\mathbf{d}_k}\right] = 1_N, \tag{3.25b}$$

$$\left[\begin{array}{c|ccc} -I_{N n_u} & O_2(k,\tau) & I_{N-\tau} \otimes \underline{U}_\tau & O_3(\tau) \\ I_{N n_u} & O_2(k,\tau) & -I_{N-\tau} \otimes \overline{U}_\tau & O_3(\tau) \end{array}\right] \left[\frac{\mathbf{u}_k}{\mathbf{d}_k}\right] \leq 0_{2 N n_u}, \tag{3.25c}$$

for all $\tau \in \mathbb{N}_{[0,\overline{s}]}$. Here, $V_N : \mathcal{X} \times \mathcal{U}^N \to \mathbb{R}_{\geq 0}$ is the cost function over the prediction horizon $N \in \mathbb{N}_{>0}$ in terms of the predicted states (based on the initial state x_k and the model dynamics (3.3)) and inputs, which reflects the performance-related control objective. Moreover, $u_{i|k} \in \mathcal{U}$ and $\delta_{i|k} \in \mathcal{B}$ denote the predicted inputs and activators at $i \in \mathbb{N}$ time instants ahead of the prediction's starting time $k \in \mathbb{N}$. The set $\boldsymbol{\mathcal{U}}_N(x_k) \subseteq \mathcal{U}^N$ represents all input sequences by which the predicted states (which depend on x_k and $\boldsymbol{u}_k$ according to (3.3)) and inputs satisfy their corresponding constraints. Hence, the set membership $\boldsymbol{u}_k \in \boldsymbol{\mathcal{U}}_N(x_k)$ imposes all general state and input constraints. For the switching constraints specific to our SAcSSs model, (3.25b) incorporates (3.10c), and (3.25c) corresponds to (3.15) and (3.18). Here, $\mathbf{u}_k = [u_{0|k}^\top \ \dots \ u_{N-1|k}^\top]^\top$ denotes the stacked vector containing all elements of $\boldsymbol{u}_k$, and $\mathbf{d}_k = [\delta_{-\min\{k,\overline{s}\}|k}^\top \ \dots \ \delta_{N-1|k}^\top]^\top$, where $\delta_{k-\tau|k} = \delta_{k-\tau}$ for $\tau \in \mathbb{N}_{>0}$, contains the $\min\{k, \overline{s}\}$ activators prior to time k and all activators in $\boldsymbol{\delta}_k$. Finally, $\otimes$ denotes the Kronecker product, and for ease of presentation we use the zero matrix shorthand notations $O_1(k) = 0_{N \times N_q \min\{k,\overline{s}\}}$, $O_2(k,\tau) = 0_{Nn_u \times N_q \min\{k,\overline{s}-\tau\}}$ and $O_3(\tau) = 0_{Nn_u \times N_q \tau}$.

Thus, due to the structure in (3.25b)-(3.25c), the constraints (3.10c), (3.15), and (3.18) that describe the mode switching and setup times can directly and automatically be used to formulate MI-MPCs for SAcSSs. Moreover, (3.25b)-(3.25c) are linear in terms of the activators and inputs. Consequently, in case V_N is quadratic and $\boldsymbol{u}_k \in \boldsymbol{\mathcal{U}}_N(x_k)$ translates into linear constraints, (3.25) is a MIQP, for which efficient solvers are available.

Next, note that by *Definition 3.13* and *Definition 3.24* the admissibility of a pair $(\boldsymbol{\delta}_k, \boldsymbol{u}_k)$ implies its feasibility. The converse does not hold, but given a feasible pair $(\tilde{\boldsymbol{\delta}}_k, \boldsymbol{u}_k)$ the existence of an admissible pair with the same input sequence is guaranteed (and it can be constructed using *Theorem 3.25*), from which we conclude the set of admissible pairs to be contained in the set of feasible pairs. Consequently, since in solving (3.25) to find an optimal pair $(\tilde{\boldsymbol{\delta}}_k^*, \boldsymbol{u}_k^*)$ we minimize over the set of feasible solutions, the optimal objective function value $V_N^* = V_N(x_k, \boldsymbol{u}_k^*)$ is a lower bound for the minimization performed over the admissible solutions only. However, now note that in $V_N(x_k, \boldsymbol{u}_k)$ (3.25a) we assume no explicit cost on the predicted activator sequence $\boldsymbol{\delta}_k$, which is valid for many applications where the control-related performance measure is mainly based on the system states and the required control effort. As a result, applying the admissibility assurance to a feasible pair $(\tilde{\boldsymbol{\delta}}_k^*, \boldsymbol{u}_k^*)$ that optimizes (3.25) yields an admissible pair $(\boldsymbol{\delta}_k^*, \boldsymbol{u}_k^*)$ (corresponding to the preferred minimum-switching sequence, see *Remark 3.28*) that also optimizes (3.25) with the same objective function value, and thus truly minimizes the optimal control problem evaluated over the set of admissible solutions.

Remark 3.29. *Instead of optimizing* (3.25) *and subsequently applying the admissibility assurance* (3.22)*, one could consider adding a switching penalty in the form of* (3.24) *to* V_N*. The optimal activator sequence* $\boldsymbol{\delta}_k^*$ *would then always*

correspond to a minimum-switching actuator sequence, and as a result the optimal pair $(\boldsymbol{\delta}_k^*, \boldsymbol{u}_k^*)$ *would be admissible. However, for control applications where there is no inherent desire to reduce the number of mode switches, or where less switching is generally preferred, but not at the expense of increasing the state- and input-related costs, introducing such a bias in the cost function could result in suboptimal behavior with respect to the true control objective. For such applications, our approach is considered favorable, as it allows for optimizing the true (unbiased) control objective in* V_N*, while additionally using as few mode switches as possible to achieve it. Additionally, when suboptimally solving the MIP problem* (3.25)*, as is typically done in practice to reduce computation time (e.g., by terminating the optimization in branch-and-bound [68, 111] or branch-and-cut [112] algorithms when a solution is found of which the cost is closer to the best theoretical cost than some threshold [80]), the obtained near-optimal activator sequence is not guaranteed to be admissible even when including a switching penalty of the form* (3.24) *in the cost function. Then,* (3.22) *again offers a way to quickly and easily construct an activator sequence by which* $(\boldsymbol{\delta}_k^*, \boldsymbol{u}_k^*)$ *is guaranteed to be admissible.*

3.4.2 MI-MPC algorithm

Recall that if a mode switch is initiated, for admissibility it must also be completed (see (B) of *Definition 3.5*), and that during a switch all control inputs are zero (see *Definition 3.6*). Therefore, in the context of MPC there is no need for re-optimizing the control inputs and activators while the SAcSS is in the process of switching. When incorporating this practical consideration, the proposed MI-MPC setup is given by *Algorithm 1.* The variable s_t stores the setup time of the last/current switch, and the counter k_t represents the number of elapsed time steps since starting this switch. Initially, we set $k = k_t = s_t = 0$, perform an admissibility-assured MI-MPC optimization by consecutively solving (3.25) and (3.22), and apply the optimal pair $(\delta_{0|k}^*, u_{0|k}^*)$ to the system. If at some time $k-1$ the SAcSS was in mode $q \in \mathcal{Q}$, and at time k it is optimal to start transitioning towards another mode $\tilde{q} \in \mathcal{Q}$, $q \neq \tilde{q}$, indicated by $\delta_k \neq \delta_{k-1}$, s_t stores the setup time $s_{q\tilde{q}}$, and k_t is reset to zero. At the following time steps, we iteratively increase k_t by one, using which we can evaluate whether the setup time required for the current transition has elapsed. If so, i.e., $k_t \geq s_t$, we perform another admissibility-assured MI-MPC optimization by consecutively solving (3.25) and (3.22). If not, i.e., $k_t < s_t$, we simply continue the mode switch ($\delta_k = \delta_{k-k_t}$) while disabling all inputs ($u_k = 0_{n_u}$) until this transition has completed.

3.5 Case study

This section first discusses the MR-HIFU hyperthermia setup and treatment. Then, we apply the procedure proposed in Section 3.3 to derive a compact

Algorithm 1 MI-MPC for for SAcSSs

Require: SAcSS-based MI-MPC (3.25), and initial conditions including $k = k_t = s_t = 0$

```
while SAcSS feedback control do
    if k_t ≥ s_t then
        Solve (3.25) and (3.22) to find (δ*_k, u*_k)
        Apply (δ_k, u_k) ← (δ*_{0|k}, u*_{0|k})
        if δ_k ≠ δ_{k−1} then
            k_t ← 0
            s_t ← s(q, q̃), where q, q̃ ∈ Q correspond to
            δ^q_{k−1} = 1 and δ^q̃_k = 1
        end if
    else
        Apply (δ_k, u_k) ← (δ_{k−k_t}, 0_{n_u})
    end if
    k ← k + 1
    k_t ← k_t + 1
end while
```

SAcSS model, set up the MI-MPC, and perform validating simulations.

3.5.1 MR-HIFU hyperthermia therapy platform

The MR-HIFU system considered in this case study consists of a Profound Sonalleve® HIFU platform, depicted in Figure 3.4, and a Philips 3T Achieva® MRI scanner. The former is a dedicated trolley-tabletop in which an MR-compatible HIFU transducer and its carrier system are integrated. The latter noninvasively provides near-real-time temperature measurements. This setup is already being used in clinics for the treatment of uterine fibroids and for palliative therapy in patients with painful bone metastases.

3.5.1.1 MR thermometry

The temperature maps are obtained noninvasively using the proton resonance frequency shift (PRFS) method [38, 44], providing the temperature difference with respect to some baseline by comparing the current MR image to a reference image. The reference image is typically acquired before treatment, such that the baseline corresponds to zero treatment-induced temperature elevation. Due to the relative nature of this MR-based thermometry, combined with the fact that the HIFU transducer itself distorts the magnetic field, for accurate measurements given a certain transducer position, a baseline image is required that was obtained with the transducer in the same position. Therefore, we con-

Figure 3.4: Philips MRI scanner and Profound Sonalleve® MR-HIFU therapy platform.

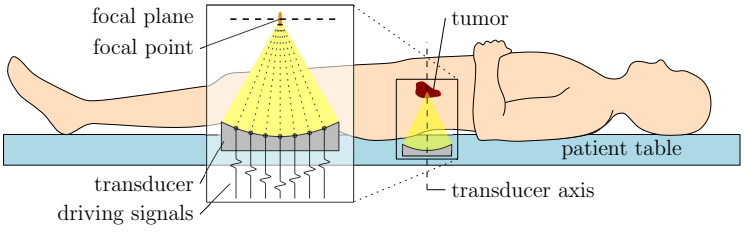

Figure 3.5: Schematic of a focal point by electronic beam steering, with the focal plane in the tumor volume.

strain the possible mechanical transducer positions to a 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 during treatment.

3.5.1.2 HIFU transducer

The Sonalleve contains a phased-array transducer consisting of 256 acoustic elements, each of which is able to generate high-intensity ultrasound waves. By coordinated modulation of the individual elements' phases and amplitudes, referred to as electronic beam steering, a focal spot can be created, see Figure 3.5, and steered through a circular area with a diameter of 16 mm, which we call a treatment cell. To treat larger regions, the transducer itself must be mechanically relocated, see [53], enabling the heating of multiple treatment cells throughout the treatment area. Due to the previously mentioned limitations of the MR thermometry, the treatment cell locations are limited to a predefined discrete set. In this case study, we allocate $N_q = 4$ treatment cells as depicted in Figure 3.6, which we will explain further below.

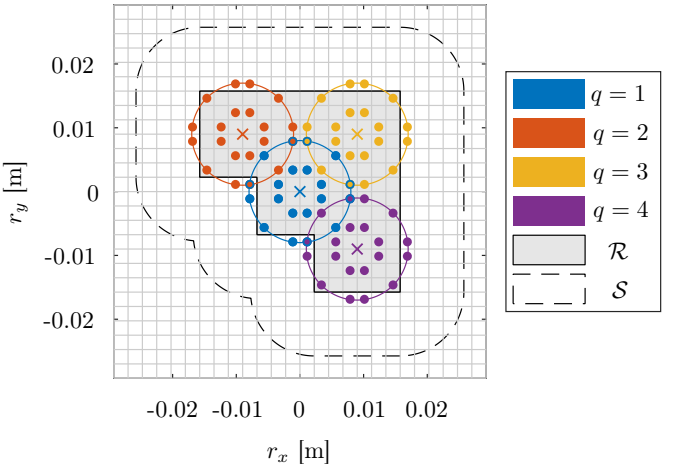

Figure 3.6: The centers (×), electronic beam steering ranges (circles), and sonication points (•) of the four treatment cells, over a large ROI $\mathcal{R}$. The tissue outside $\mathcal{S}$ must be safeguarded against overheating.

3.5.2 Hyperthermia treatment

The main goal in a hyperthermia treatment is to create and maintain a controlled and homogeneous temperature elevation in the region of interest (ROI) $\mathcal{R}$ shown in Figure 3.6 containing the tumor, as this sensitizes the treated tissue to the effects of chemo- and radiotherapy. The tissue sensitization gradually increases for rising temperatures starting around 40 °C [12], which is why the healthy tissue outside $\mathcal{S}$ in Figure 3.6 will be safeguarded against temperatures above 40 °C. In the ROI, adequate sensitization is considered to occur around 41 °C, and optimal treatment quality is achieved at 42 °C, while overheating above 43 °C reduces the beneficial heat-induced effects and should therefore be avoided.

3.5.3 SAcSS model

3.5.3.1 Arc-weighted digraph

The complete simple arc-weighted digraph $\Gamma = (\mathcal{Q}, \mathcal{E}, s)$ describing the system considered in this case study corresponds to *Example 3.2*. Here, the setup time defined by s represents the number of samples required to mechanically relocate the transducer from one cell to another.

3.5.3.2 Plant dynamics, state space, input space, and actuator selector function

To model the tissue's thermal dynamics, we follow a procedure similar to Chapter 2. In summary, this entails first describing the tissue's thermal dynamics

using the Pennes bioheat equation [86] using tissue parameters from [95] and then spatially discretizing this partial differential equation on a two-dimensional square grid with 2.25×2.25 mm^2 voxels (depicted in Figure 3.6 by the grid lines) using the central difference scheme, which in [72] was verified to be the method that best balances model simplicity with descriptive accuracy. Finally, the model is temporally discretized with the MR thermometry sample time $T_s = 3.2$ s using the forward Euler method, which preserves model sparsity, while providing sufficient accuracy for the considered system, as discussed in Chapter 2.

The resulting discrete-time state-space model in the form of (3.3) is given by

$$x_{k+1} = f(x_k, u_k) = Ax_k + Bu_k, \tag{3.26a}$$

$$y_k = \begin{cases} x_k + v_k, & \text{if } \sigma_{k-1} \in \mathcal{Q}, \\ \emptyset, & \text{if } \sigma_{k-1} \in \mathcal{E}, \end{cases} \tag{3.26b}$$

where the states $x_k \in \mathcal{X} = \mathbb{R}^{n_x}$, $n_x = 36^2 = 1296$, represent the temperature elevations with respect to the baseline of the voxels in the focal plane at real time $t_k = kT_s$. The matrix A captures 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 $y_k \in \mathbb{R}^{n_x}$ to consist of full state measurements corrupted by MR measurement noise $v_k \in \mathbb{R}^{n_x}$, which can be well approximated by spatially uncorrelated zero-mean Gaussian noise with a standard deviation of 0.4 °C, i.e., $v_k \sim \mathcal{N}(0, 0.4^2 I_{n_x})$, when the transducer was not moving during the interval from t_{k-1} until t_k indicated by $\sigma_{k-1} \in \mathcal{Q}$. In case of transducer motion $\sigma_{k-1} \in \mathcal{E}$, no measurement is available at time t_k. Per treatment cell, and thereby per actuator mode $q \in \mathcal{Q}$, we use $n_u^q = n_{u,\text{cell}} = 20$ 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 Figure 3.6 by the markers '•'. Using our method, we only need to model the N_q *operational* input modes. To this end, let the input u_k^q describe the average acoustic power applied at each of the sonication points in cell $q \in \mathcal{Q}$ over the course of one sampling interval. The inputs are collected in the input vector $u_k \in \mathcal{U}$ as in (3.6), and hence the input matrix is given by

$$B = \begin{bmatrix} B^1 & \cdots & B^{N_q} \end{bmatrix} \in \mathbb{R}^{n_x \times n_u}, \tag{3.27}$$

where each submatrix $B^q \in \mathbb{R}^{n_x \times n_{u,\text{cell}}}$ describes the system's temperature change in response to the heating power u_k^q applied at the sonication points in cell $q \in \mathcal{Q}$. Since by HIFU we can only deposit heat in the system, but not extract it, the inputs are nonnegative. Additionally, for patient safety, the maximum input power per sonication point is limited to $u_{\max} = 15$ W. Consequently, the input space is given by $\mathcal{U} = \mathbb{R}^{n_u}_{[0,u_{\max}]}$, $n_u = \sum_{q \in \mathcal{Q}} n_u^q = N_q n_{u,\text{cell}} = 80$. Next, note that even though six sonication points on the outer ring of cell 1 are shared by multiple cells, to satisfy *Assumption 3.9* we have defined the $n_{u,\text{cell}}$

sonication points for each cell individually, resulting in disjoint inputs per actuator mode. Correspondingly, in accordance with *Remark 3.11*, the inputs are ordered as in (3.6), and the actuator selector function is given by

$$\Phi(\sigma) = \begin{cases} \mathbb{N}_{[1+n_{u,\mathrm{cell}}(q-1),n_{u,\mathrm{cell}}q]}, & \text{if } \sigma = q \in \mathcal{Q}, \\ \emptyset, & \text{if } \sigma \in \mathcal{E}, \end{cases} \tag{3.28}$$

describing the fact that heating can only occur in the cell where the transducer is located, and that no heating may occur during transducer motion.

3.5.4 Input constraints

In this case study, we use Gurobi 8.1.1, which requires explicit specification of the input's upper and lower limits. From $\mathcal{U}$ as in Section 3.5.3.2, according to (3.8) we find

$$0_{n_u} = \underline{u} \leq u_k \leq \overline{u} = 1_{n_u} u_{\max}. \tag{3.29}$$

In addition, as a safety measure in the considered MR-HIFU hyperthermia setup we constrain the total applied power in a single treatment cell to 100 W, i.e.,

$$1_{n_u^q}^\top u^q \leq \overline{u_\Sigma} = 100 \quad \text{for all } q \in \mathcal{Q}. \tag{3.30}$$

To formulate the inequality constraints that incorporate the switching and setup times in the upcoming MI-MPC, of which the general form is given by (3.15) and (3.18), based on Γ we first construct the setup time matrix (3.9) given by

$$S = \begin{bmatrix} 0 & 2 & 1 & 2 \\ 2 & 0 & 2 & 3 \\ 1 & 2 & 0 & 2 \\ 2 & 3 & 2 & 0 \end{bmatrix}, \tag{3.31}$$

and the corresponding setup time constraint matrices (3.17), which read

$$S_1 = I_4, \quad S_2 = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad S_3 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{bmatrix}. \tag{3.32}$$

However, since by $\mathcal{U} \subset \mathbb{R}_{\geq 0}$ the controls are nonnegative, see also (3.29), and a common upper bound in the sum of the inputs is imposed by (3.30), we do not need to formulate the constraints (3.18), but we can instead use the simplified form (3.20) discussed in *Remark 3.22*. Moreover, by imposing (3.29) we can also use the simplified form (3.21) instead of (3.15), see *Remark 3.23*. Thus, introducing the Boolean one-hot activators $\delta_k \in \mathcal{B}$ as in (3.10) for $N_q = 4$, we find the linear inequality constraints on the inputs u_k related to the setup times to be given by

$$J_u u_k \leq \overline{u_\Sigma} \delta_k, \tag{3.33a}$$

$$J_u u_k \leq \overline{u_\Sigma} S_\tau^\top \delta_{k-\tau} \text{ for all } \tau \in \mathbb{N}_{[1,\min\{k,3\}]}. \tag{3.33b}$$

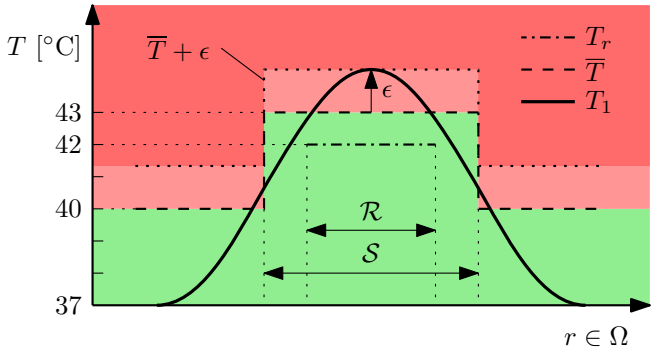

Figure 3.7: Schematic cross section of the temperature objectives corresponding to $\mathcal{R}$ and $\mathcal{S}$. The maximum violation ϵ is shown for some overheated temperature distribution T_1 such that $T_1 \leq \overline{T} + \epsilon$.

3.5.5 MI-MPC for large-volume MR-HIFU hyperthermia

3.5.5.1 Prediction model

Based on (3.26), we define the prediction model as

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k}, \tag{3.34}$$

where $x_{i|k} \in \mathcal{X}$ and $u_{i|k} \in \mathcal{U}$ denote the predicted states and inputs, respectively, at $i \in \mathbb{N}$ time steps ahead of the prediction sequence's starting time $k \in \mathbb{N}$.

3.5.5.2 Observer model

To reduce the propagation of the measurement noise into the desired input u_k computed by MPC, we use a Luenberger observer given by

$$\hat{x}_k = \begin{cases} A\hat{x}_{k-1} + Bu_{k-1} + L(y_k - \hat{y}_k^-), & \text{if } \sigma_{k-1} \in \mathcal{Q}, \\ A\hat{x}_{k-1} + Bu_{k-1}, & \text{if } \sigma_{k-1} \in \mathcal{E}, \end{cases} \tag{3.35}$$

where $\hat{y}_k^- = A\hat{x}_{k-1} + Bu_{k-1}$ denotes the estimated output at time k before the measurement-based correction step in (3.35) if possible by $\sigma_{k-1} \in \mathcal{Q}$. The observer gain matrix is given by $L = 0.25I_{n_x}$, which has been tuned such that the estimator is exponentially stable, see Section 3.7, and exhibits desirable convergence behavior.

3.5.5.3 MI-MPC optimization problem

The temperature objectives of our MI-MPC setup are depicted schematically in Figure 3.7 in cross-section perspective. On $\Omega \subset \mathbb{R}^2$, denoting the patient domain

in the focal plane, $T_r : \mathcal{R} \to \mathbb{R}$ (dash-dotted) is the reference temperature of 42 °C in the ROI $\mathcal{R}$, and $\overline{T} : \Omega \to \mathbb{R}$ is the location-dependent upper temperature bound (dashed) used to prevent overheating (red). To translate these objectives to the state-space representation (3.26), we introduce the performance variables $z_k = Hx_k \in \mathbb{R}^{n_z}$, with $H \in \{0,1\}^{n_z \times n_x}$ being a matrix with one 1 per row (and at most one 1 per column), which are the temperatures of the $n_z < n_x$ voxels inside $\mathcal{R}$. Furthermore, we use $z_r \in \mathbb{R}^{n_z}$ and $\overline{x} \in \mathbb{R}^{n_x}$ to denote the voxel-wise temperature reference and upper bounds corresponding to the values of T_r and $\overline{T}$, respectively. The maximum violation of the upper temperature bound is measured by the slack variable $\epsilon_k = \epsilon(x_k) = \|\max\{x_k - \overline{x}, 0_{n_x}\}\|_\infty \in \mathbb{R}_{\geq 0}$, where the maximum operator is used element-wise. The predicted performance and slack variables are denoted by $z_{i|k}$ and $\epsilon_{i|k}$, respectively.

The resulting MI-MPC, which can be easily written in the form of (3.25), is then given by

$$\min_{\boldsymbol{\delta}_k, \boldsymbol{u}_k} \sum_{i=0}^{N} (z_{i|k} - z_r)^\top Q (z_{i|k} - z_r) + f_\epsilon \epsilon_{i|k}, \tag{3.36a}$$

subject to

$$\begin{aligned}
x_{i+1|k} &= Ax_{i|k} + Bu_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && \text{(3.36b)}\\
x_{0|k} &= \hat{x}_k, && && \text{(3.36c)}\\
1_{N_q}^\top \delta_{i|k} &= 1, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && \text{(3.36d)}\\
x_{i|k} &\leq \overline{x} + 1_{n_x} \epsilon_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N]}, && \text{(3.36e)}\\
0 &\leq \epsilon_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N]}, && \text{(3.36f)}\\
0_{n_u} &\leq u_{i|k} \leq \overline{u}, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && \text{(3.36g)}\\
J_u u_{i|k} &\leq \overline{u_\Sigma} \delta_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && \text{(3.36h)}\\
J_u u_{i|k} &\leq \overline{u_\Sigma} S_\tau^\top \delta_{i-\tau|k}, && \begin{cases} \forall\, \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}\}]}, \\ \forall\, i \in \mathbb{N}_{[0,N-1]}. \end{cases} && \text{(3.36i)}
\end{aligned}$$

In (3.36a), we choose $Q = \frac{1}{n_z} I_{n_z}$ and $f_\epsilon = 10$, which are normalized with respect to the number of weighted variables (note that ϵ is scalar) for more intuitive tuning, and use horizon $N = 8$. The weighting f_ϵ on the linear term in (3.36a), incorporating the upper temperature bound as a soft constraint, is large relative to Q, reflecting the fact that the prevention of overheating is prioritized. The quadratic term in (3.36a) weighted by Q enforces tracking of the ROI temperature towards the optimal treatment temperature. The equality constraints (3.36b) capture the system dynamics as modeled by (3.34), with the initial condition (3.36c) determined by the observer (3.35). The equality (3.36d) captures (3.10c) throughout the horizon. The inequalities (3.36e) encode the upper temperature bound as a soft constraint using the slack variable, which is restricted to be nonnegative by (3.36f). Finally, (3.36g) corresponds to the

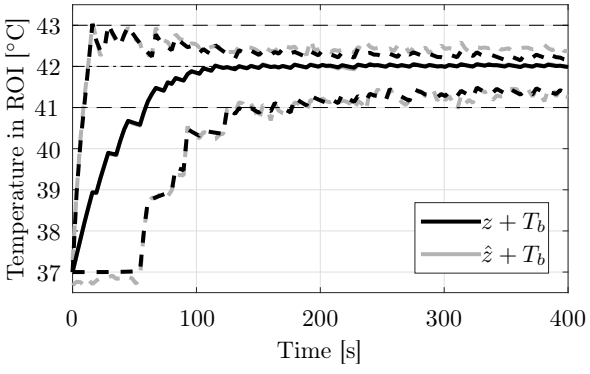

Figure 3.8: The mean (solid) and maximum/minimum (dashed) temperature inside the ROI $\mathcal{R}$ of the plant (black) and as estimated by the observer (gray).

actuator constraints (3.29), and the inequalities (3.36h)-(3.36i) incorporate the mode and setup time constraints (3.33). Note that in (3.36i), $\delta_{i-\tau|k} = \delta_{k+i-\tau}$ for $i < \tau$, following from the system's activator sequence prior to time $k \in \mathbb{N}$.

3.5.6 Simulation results

The temperature control performance of the MI-MPC setup is verified by means of simulation. We initialize the transducer in cell 1, set the plant, observer, and controller states to zero, corresponding to the monitored tissue to be at the blood temperature $T_b = 37$ °C before treatment, and initialize the noise sequence. The temperature evolution of the ROI voxels of the plant (gray) and the observer/controller (black) are visualized in Figure 3.8 using their mean (solid) and maximum/minimum values (dashed). It can be seen that after approximately 150 s the entire ROI is heated to 41 °C, which is the temperature above which the beneficial hyperthermia-related effects are adequately triggered. Furthermore, during heat-up the upper temperature limit is observed to be reached, but no overheating occurs, and the average ROI temperature converges to the optimum 42 °C.

In Figure 3.9, the total acoustic input power per cell is shown, together with the destination mode $\text{post}(\sigma_k)$. First, note that $\text{post}(\sigma_k)$ (which is directly related to δ_k by (3.10)) and u_k^q indeed correspond to an admissible sequence of actuator-input pairs, describing the SAcSS to heat only in the active cell, perform the minimum number of mode switches required to allow for the desired nonzero inputs, and respect the setup time induced by each switch. From this, we verify the admissibility-assured MI-MPC setup to be functioning as intended. Next, regarding the input power, the figure shows that the controller initially requests maximum power for fast heat-up, reaching the upper limit of the total cell power constraint (3.36h). Moreover, in this period it often heats a certain

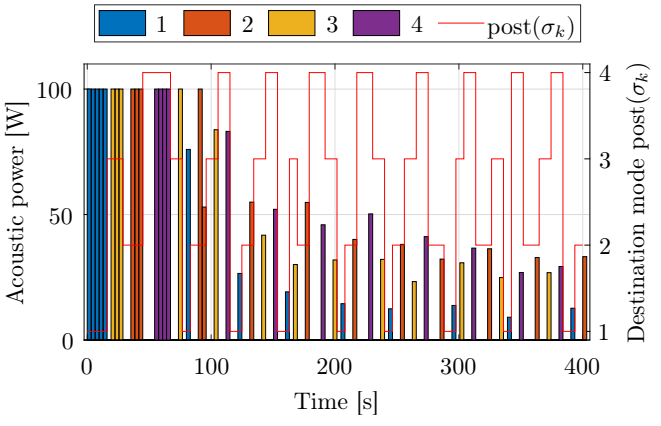

Figure 3.9: The total acoustic power per cell, and the destination mode post(σ_k) (red line).

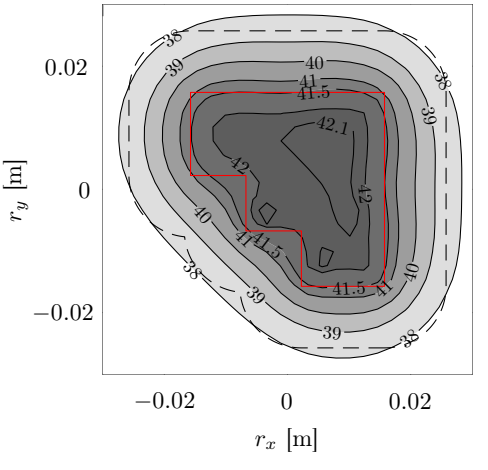

Figure 3.10: Contour plot of the focal plane temperature distribution in the ROI $\mathcal{R}$ (red) and $\mathcal{S}$ (dashed) at $t_k = 400$ s.

cell for several consecutive samples, thereby reducing actuator downtime, thus contributing to achieving a short heat-up phase. For $t_k > 100$ s, however, each cell is only heated for one sample before continuing to the next cell, as this allows for maintaining a ROI temperature distribution that is as homogeneous as possible.

The plant temperature at $t_k = 400$ s is shown in Figure 3.10, where additionally the perimeters of $\mathcal{R}$ (red) and $\mathcal{S}$ (dashed black) are plotted. This figure clearly illustrates the temperature homogeneity in $\mathcal{R}$. The maximum tempera-

ture (42.1 °C) is observed in all cells except cell 2 (top-left). In Figure 3.9, it can be seen that at time $t_k = 400$ s the transducer has just arrived at cell 2, and that the MPC has computed significant sonication power (approximately 33 W) to be applied, which will result in a temperature increase in this cell. This exemplifies the advantage of a *predictive* controller for large-volume MR-HIFU hyperthermia. That is, the MPC is able to anticipate the diffusive heat loss during the future time interval in which a cell cannot be heated due to the transducer having moved to another cell. To preemptively counteract this heat loss, which is most severe at the ROI's corners, additional heat is injected especially at the corners just before relocating the transducer, resulting in temperature peaks (slightly) above the reference temperature of 42 °C. Then, after the subsequent heat diffusion during the time period that a cell cannot be heated, minor temperature peaks remain near the ROI's corners, while the corner temperatures are still within the desired range (≥ 41 °C), as can be seen in the corners of cells 1, 3, and 4 in Figure 3.10.

3.5.7 Computational efficiency

To evaluate the improved computational efficiency of the online MI-MPC problem when using the modeling approach proposed in this chapter, we compare the resulting computation time to that of the MI-MPC setup obtained by modeling the SAcSS using a lifting approach [107], effectively yielding a constrained switched linear system (cSLS) [81, 106] in MLD form [104]. Similar to our method, for the purpose of MI-MPC this also requires duplicating shared input channels such that the resulting input vector is the collection of the disjoint inputs, and hence the discrete-time dynamics are described by (3.26). Contrary to our method, however, is the fact that we must now define $s_{q\tilde{q}}$ actuator *transition* modes for each switch $(q, \tilde{q}) \in \mathcal{E}$, in addition to the $N_q = 4$ operational modes, which for this case study results in $s_\Sigma \coloneqq \sum_{(q,\tilde{q})\in\mathcal{E}} s_{q\tilde{q}} = 24$ additional modes. The constrained switching is graphically represented in Figure 3.11, where the operational and transition modes are indicated by the numbered and (smaller) unnumbered nodes, respectively. The arcs represent possible switches. Additionally, switches are possible from each color-filled node to each node with an outline of the same color, modeling the possibility for back-to-back actuator switches.

Correspondingly, we define the N_q Boolean actuator-operation states

$$\beta_k^q \in \{0, 1\}, \quad q \in \mathcal{Q}, \tag{3.37}$$

for which $\beta_k^q = 1$ if at time k the SAcSS is in operational mode $q \in \mathcal{Q}$, and $\beta_k^q = 0$ otherwise. In addition, we must for each actuator switch $(q, \tilde{q}) \in \mathcal{E}$ with

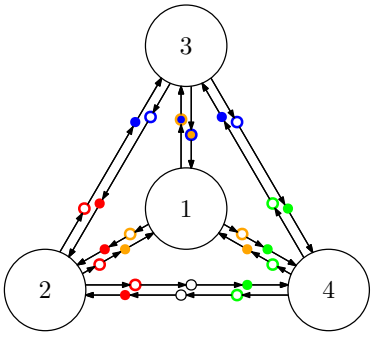

Figure 3.11: Graph representation of the switching when modeled as a cSLS. Switches are possible as indicated by the arrows, and from each color-filled node to each node with an outline of the same color.

nonzero setup time $s_{q\tilde{q}}$ define the Boolean actuator-transition states

$$\beta_k^{q\tilde{q}} = \begin{bmatrix} \beta_k^{q\tilde{q},1} \\ \vdots \\ \beta_k^{q\tilde{q},s_{q\tilde{q}}} \end{bmatrix} \in \{0,1\}^{s_{q\tilde{q}}}, \tag{3.38}$$

where $\delta_k^{q\tilde{q},i} = 1$ if at time k the SAcSS is at setup time instant $i \in \mathbb{N}_{[1,s_{q\tilde{q}}]}$ of the switch $(q,\tilde{q}) \in \mathcal{E}$. Since a SAcSS' actuator can only be in one mode or at one stage of a switch at any given time, it must hold that

$$1_{N_q+s_\Sigma}^\top \beta_k = 1, \tag{3.39}$$

where β_k is the collection of all actuator-operation and actuator-transition states in (3.37) and (3.38) at time k. The mode-dependent input constraints can be expressed as

$$\underline{u}^q \beta_k^q \leq u_k^q \leq \overline{u}^q \beta_k^q \quad \text{for all } q \in \mathcal{Q}, \tag{3.40}$$

or, in case the input bounds are explicitly imposed by (3.29) and using $\overline{u_\Sigma}$ from (3.30), as the simplified form

$$1_{n_u^q}^\top u_k^q \leq \overline{u_\Sigma} \beta_k^q \quad \text{for all } q \in \mathcal{Q}. \tag{3.41}$$

Next, the switching progression along the successive actuator-transition states is prescribed by

$$\begin{bmatrix} \beta_{k+1}^{q\tilde{q},2} \\ \vdots \\ \beta_{k+1}^{q\tilde{q},s_{q\tilde{q}}} \end{bmatrix} = \begin{bmatrix} 1 & 0 & & \\ & \ddots & \ddots & \\ & & 1 & 0 \end{bmatrix} \beta_k^{q\tilde{q}}, \tag{3.42}$$

for all $(q, \tilde{q}) \in \mathcal{E}$ with $s_{q\tilde{q}} > 1$. Furthermore, recall that a SAcSS can only be in operational actuator mode $q \in \mathcal{Q}$ when at the previous time instant it was either already in this mode, in its last setup time instant towards this mode, or in an operational mode $\tilde{q} \in \mathcal{Q}$ from which the switch towards q induces zero setup time. This can be described by the inequality

$$\beta_k^q \leq \beta_{k-1}^q + \sum_{\tilde{q} \in \mathcal{Q}_0^q} \beta_{k-1}^{\tilde{q}} + \sum_{\tilde{q} \in \mathcal{Q}_{>0}^q} \beta_{k-1}^{\tilde{q}q, s_{\tilde{q}q}}, \tag{3.43}$$

where similarly to (3.11) we define for $q \in \mathcal{Q}$ the sets

$$\mathcal{Q}_0^q = \{\tilde{q} \in \mathcal{Q} \mid s_{\tilde{q}q} = 0\},$$
$$\mathcal{Q}_{>0}^q = \{\tilde{q} \in \mathcal{Q} \mid s_{\tilde{q}q} > 0\},$$

as the set of actuator modes $\tilde{q} \in \mathcal{Q}$ from which the switch towards mode $q \in \mathcal{Q}$ induces zero and nonzero setup time, respectively. Similarly, an actuator switch $(q, \tilde{q}) \in \mathcal{E}$ can only start at time k when at $k-1$ the SAcSS was in or arriving at mode $q \in \mathcal{Q}$, and hence

$$\beta_k^{q\tilde{q},1} \leq \beta_{k-1}^q + \sum_{\hat{q} \in \mathcal{Q}_{>0}^q} \beta_{k-1}^{\hat{q}q, s_{\hat{q}q}}. \tag{3.44}$$

Note that, contrary to (3.43), in (3.44) we do not need to include the possibility that at time $k-1$ the SAcSS was in some mode $\hat{q} \in \mathcal{Q}_0^q$, since by *Assumption 3.7* and *Assumption 3.8* switching towards mode $\tilde{q}$ at time k would then be described by the direct switch $(\hat{q}, \tilde{q}) \in \mathcal{E}$ (i.e., $\beta_k^{\hat{q}\tilde{q}} = 1$, not $\beta_k^{q\tilde{q}} = 1$).

Finally, to obtain the MLD system form of the cSLS model of the considered SAcSS, all of the above constraints must be combined and rewritten in matrix-vector form in terms of the inputs and the actuator-operation and actuator-transition states, and combined with the discrete-time plant dynamics (3.26). However, as these are not easy to write down compactly, the MLD system description is not explicitly given here.

The resulting MI-MPC consists of the optimization problem (3.36a) (except for optimizing over $\boldsymbol{\beta}_k = (\beta_{0|k}, \ldots, \beta_{N-1|k})$ instead of $\boldsymbol{\delta}_k$) subject to the thermal dynamics and initial condition constraints (3.36b) and (3.36c), the upper temperature bound soft constraints (3.36e) and (3.36f), the input bounds (3.36g), and additionally (after substituting u_k and β_k by their prediction counterparts $u_{i|k}$ and $\beta_{i|k}$) for all $i \in \mathbb{N}_{[0,N-1]}$ the one-hot encoding constraint (3.39), the mode-dependent input constraints (3.41), the switch progression constraints (3.42), and the constraints (3.43) and (3.44) describing the actuator possibilities at the end and beginning of an actuator switch. Note that (for all $i \in \mathbb{N}_{[0,N-1]}$) we must impose (3.41) and (3.43) for all $q \in \mathcal{Q}$, and (3.42) and (3.44) for all $(q, \tilde{q}) \in \mathcal{E}$ with $s_{q\tilde{q}} > 1$ and $s_{q\tilde{q}} > 0$, respectively.

Compared to (3.36), the continuous part of the optimization problem is unchanged, while the number of Boolean decision variables increases from

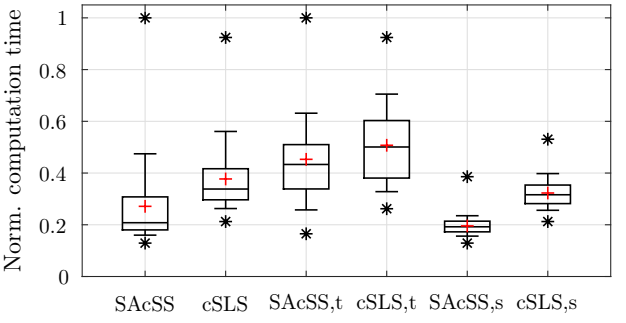

Figure 3.12: Distribution of the normalized computation times for the SAcSS MI-MPC (3.36) and cSLS MI-MPC over the entire simulation, and over the transient (t) and steady-state (s) intervals only. We indicate the mean (red +), median (central mark), 25th and 75th percentiles (box edges), 10th and 90th percentiles (whiskers), and extrema ($*$).

$N_q N = 32$ ($\delta_{i|k}$ in (3.10b) over the horizon N) to $(N_q + s_\Sigma)N = 224$ ($\beta_{i|k}$ in (3.37)-(3.38) over N). The number of integer equality constraints increases from $N = 8$ in (3.36d) to $(1 + \sum_{(q,\tilde{q})\in\mathcal{E}} \max\{s_{q\tilde{q}} - 1, 0\})N = 104$ in (3.39) and (3.42) over N. Finally, the number of mixed-integer inequality constraints increases from $N_q(1 + \overline{s})N = 128$ in (3.36h)-(3.36i) (actually $N_q(1 + \overline{s} - 2/4)N = 112$, since 2 out of 4 constraints generated using S_3 from (3.32) for each i in (3.36i) are redundant, see *Remark 3.20*) to $(2N_q + \sum_{(q,\tilde{q})\in\mathcal{E}} \min\{s_{q\tilde{q}}, 1\})N = 160$ in (3.41) and (3.43)-(3.44) over N.

In simulation, the control inputs and resulting temperatures using the cSLS-based MI-MPC are found to indeed be exactly the equal to those obtained using the MI-MPC (3.36) up to numerical tolerances. The computation times for both MI-MPC setups have been recorded using five different noise realizations and five runs per realization, i.e., twenty-five simulations for each MI-MPC. The distributions of the computation times are depicted in Figure 3.12 (normalized with respect to the largest observed computation time) over the entire simulation, and over the transient ($t_k < 100$) and steady-state ($t_k \geq 100$) intervals separately. The MI-MPC (3.36) designed specifically for SAcSS clearly outperforms the MI-MPC derived using the cSLS model, with the latter showing 39% and 63% larger overall mean and median computation times. In fact, in the constant-temperature phase, which comprises the majority of a hyperthermia treatment, the cSLS MI-MPC computation times' 10th percentile is larger than the 90th percentile of (3.36), and the mean and median computation times are 65% larger than those obtained using our novel method.

3.6 Conclusion

In this chapter, the class of switched-actuator systems with setup times (SAcSSs) has been formally introduced. As key contribution, a modeling framework for SAcSSs has been presented, which is specifically tailored to (a) allow for user-friendly system specification, modeling, and controller synthesis, and (b) yield compact models leading to efficient MI-MPCs. In particular, the resulting model is in MLD form, consisting of a state-space representation of the SAcSS' dynamics in its N_q operational modes, combined with systematically derived mixed-integer linear inequality constraints on the inputs (thus achieving compatibility with MI-MPC) to incorporate the mode switching and setup times. A distinctive property of our method is that instead of explicitly modeling the zero-input "switching modes" as in a constrained switched linear system (cSLS) or lifting approach, which would require many auxiliary Boolean variables per time step, it uses only the N_q Boolean variables corresponding to the operational modes, and infers the actuator activity from a sequence of these variables in the prediction horizon. This reduces the MI-MPC's computational complexity. The proposed modeling procedure and corresponding MI-MPC setup have been validated in a large-volume MR-HIFU hyperthermia case study. It is demonstrated that the desired temperature distribution can be achieved and maintained in a large tumor by coordinated heating and mechanical transducer displacement, and that the MI-MPC's computational efficiency is improved with respect to a cSLS/lifting approach.

3.7 Appendix: Proof of observer stability

Consider the observer (3.35), and note that A is Schur, i.e., with all eigenvalues strictly inside the unit disc of the complex plane, due to the tissue's stable (first-order) thermal dynamics (3.26). The corresponding estimation error dynamics are given by

$$e_k = x_k - \hat{x}_k = \begin{cases} (A - LA)e_{k-1}, & \text{if } \sigma_{k-1} \in \mathcal{Q}, \\ Ae_{k-1}, & \text{if } \sigma_{k-1} \in \mathcal{E}, \end{cases} \tag{3.45}$$

Since A is Schur, there exists a positive definite matrix $P > 0$ such that

$$A^\top P A - P < 0, \tag{3.46}$$

which ensures global exponential stability of the $\sigma_{k-1} \in \mathcal{E}$ subsystem of (3.45) when using $V(e_k) = e_k^\top P e_k$ as Lyapunov function. Next, for $L = \alpha_l I$ with $0 \leq \alpha_l \leq 1$, it holds that

$$(A - LA)^\top P(A - LA) - P = (1 - \alpha_l)^2 A^\top P A - P \leq A^\top P A - P < 0, \tag{3.47}$$

which guarantees global exponential stability of the $\sigma_{k-1} \in \mathcal{Q}$ subsystem of (3.45) using the same Lyapunov function $V(e_k) = e_k^\top P e_k$. As $V(e_k)$ is a *common*

Lyapunov function for the two subsystems, (3.45) is globally exponentially stable for arbitrary switching sequences of σ_k.

Acknowledgements

The authors thank Karthik Mohan and Jesper van Wordragen for their contributions to the case study, and Jim Rawlings for his interesting insights and suggestions.


Good design is a lot like clear thinking made visual.
Edward Tufte

Chapter 4

Target-Conformal Optimal Actuator Placement for Ultrasound-Mediated Hyperthermia in Cancer Treatments

In hyperthermia for cancer treatments, tumors are heated to improve the outcome of radio- and chemotherapies. Using extracorporeal phased-array high-intensity focused ultrasound (HIFU) transducers, the heating can be applied noninvasively, accurately, and with high acoustic power. The heating location can be changed by electronic beam steering, within the transducer's local heating range, and by mechanically repositioning the transducer, which effectively shifts the local range to enable the treatment of larger tumors. The sonication points, i.e., the discrete locations reachable by electronic beam steering, and the admissible transducer positions must be chosen before treatment. However, their locations have a major impact on the achievable tumor temperatures, and their number is directly related to the controller's complexity and, thereby, real-time feasibility. To address this challenging issue, this chapter presents a target-conformal optimal actuator placement procedure, designed to enable maximum treatment quality using only a limited number of sonication points and transducer positions. The method is computationally tractable, and takes into account the specific tumor geometry and tissue properties. In a numerical case study, the actuator placement procedure is applied and combined with the design of various mixed-integer model predictive controllers for large-volume magnetic-resonance-guided HIFU hyperthermia, demonstrating the achievable control performance and computation time using significantly fewer, but strategically selected, control inputs.

This chapter is based on D.A. Deenen, L.C. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Target-conformal optimal actuator placement for ultrasound-mediated hyperthermia in cancer treatments," *In preparation for journal submission.*

4.1 Introduction

It has been shown in many clinical trials that heating tissues to approximately 42 °C over an extended period of time, referred to as mild hyperthermia, significantly enhances the desired effects of chemo- and radiotherapy [3,4,16,19–24]. As a major advantage, hyperthermia is non-toxic and can be induced locally, allowing for selectively sensitizing only the cancer tissue, while leaving healthy tissue unaffected. Consequently, hyperthermia treatments as an adjuvant therapy enable improved clinical outcome and lowered chemo- and radiotherapeutic doses, the latter of which alleviates the unwanted and often severe side effects. In fact, using temperature-sensitive drug carriers, local hyperthermia enables targeted drug release, allowing for further reducing the systemic toxicity in chemotherapy [25–27].

Using high-intensity focused ultrasound (HIFU) as heating modality, hyperthermic temperatures can be induced in tissue. In recent years, this has been successfully combined with near-real-time temperature measurements obtained using an MRI scanner, resulting in magnetic-resonance-guided (MR-)HIFU as a valuable technology for feedback-controlled thermal therapies [30, 43, 47, 48, 83]. When using extracorporeal applicators, MR-HIFU hyperthermia treatments become fully noninvasive, which significantly contributes to the patients' quality of life by avoiding the need for additional surgery. Modern MR-HIFU setups typically use phased-array transducers, which enable powerful heating with millimeter-scale accuracy, allowing for high-quality treatments. In particular, a focal spot can be generated by constructive interference of the ultrasound waves, and steered up to a maximum deflection via coordinated phase shifts of the acoustic driving signals, which is referred to as electronic beam steering. To treat tumors larger than the electronic beam steering range, the transducer itself can also be repositioned mechanically [52, 53]. The transducer's mechanical point-to-point motion, however, takes a non-negligible amount of time, during which no heating is allowed.

The discrete locations to which the focal spot can be steered during treatment are called sonication points. Together with the admissible transducer positions, these must be chosen before treatment, such that subsequently a feedback controller can be used to determine the desired transducer position and sonication point powers during treatment. However, choosing the sonication points and admissible transducer positions in such a manner that the achievable temperature distribution is as close to the desired reference as possible is not straightforward, as this strongly depends on the shape and tissue properties of the specific tumor. At the same time, choosing them optimally is of vital importance, as the therapeutic effects of hyperthermia are directly linked to the realized tumor temperatures [60–62]. Moreover, including more points and positions generally allows for better control performance, but also increases controller complexity, which especially for more advanced, and thereby often computationally demand-

ing, control strategies may hamper real-time feasibility.

A particularly interesting example of such a strategy is model predictive control (MPC) [66]. Indeed, researchers have recognized MPC to be highly suitable for MR-HIFU treatments, as it allows for high-dimensional multi-input multi-output control to optimize a cost function representing the treatment objective, while explicitly taking into account the tissue's thermal dynamics and the restrictive input and temperature constraints. Examples include [69, 70] for minimizing the treatment time of thermal ablation therapies, and [71] for maximizing the thermal dose applied to the cancer cells without harming the surrounding healthy tissue. More recently, we have developed a preliminary MPC scheme for MR-HIFU hyperthermia temperature control using a stationary transducer in [72], which is shown to significantly outperform the current clinical state-of-the-art (binary) controller from [53] in phantom experiments. To also enable such improved control performance for the treatment of larger tumors requiring multiple transducer positions, in Chapter 3 a mixed-integer (MI-)MPC design for switched-actuator systems with setup times (SAcSSs) is proposed, which allows for simultaneously optimizing the discrete transducer positioning and the continuous acoustic powers of the sonication points in range.

Due to its model- and optimization-based nature, MPC is among the most promising control solutions for high-precision thermal therapies in the future. Unfortunately, due to these properties, MPC also easily leads to increased computational complexity. Especially for MI-MPCs with long prediction horizons and many transducer positions and sonication points, combined with the large models typically used in MR-HIFU treatments, the resulting computation times may be too large for practical application. In this case, optimally selecting only a small number of transducer positions and sonication points may enable superior treatment quality, while also satisfying computation time constraints for real-time feasibility.

In line with the above, as the first and main contribution of this chapter, we propose a target-conformal optimal actuator placement method for MR-HIFU cancer treatments, which aims to maximize treatment quality using only a limited number of sonication points and transducer positions, while taking into account the specific tumor geometry and tissue properties. In particular, the method first optimizes the transducer positions based on a weighted balance between the achievable steady-state temperature distribution and target coverage. Then, it identifies the sonication points critical in steady state, and subsequently allows for including additional sonication points that maximize target coverage for more uniform heating and enhanced robustness in case of disturbances or unmodeled tissue inhomogeneities.

As a second contribution, the resulting control performance and computation times are investigated for various actuator settings in a large-volume (requiring multiple admissible transducer positions) MR-HIFU hyperthermia numerical case study. Although the actuator placement procedure presented in this chap-

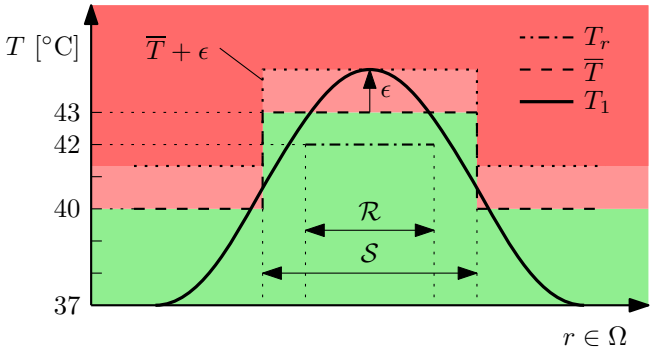

Figure 4.1: Schematic cross section of the temperature objectives corresponding to $\mathcal{R}$ and $\mathcal{S}$. The maximum violation ϵ is shown for some overheated temperature distribution T_1 such that $T_1 \leq \overline{T} + \epsilon$.

ter is considered to be applicable to MR-HIFU control design in general, in this case study it is used for MI-MPC, as this control strategy is of particular interest for MR-HIFU due to the reasons mentioned previously.

The remainder of the chapter is organized as follows. First, we discuss the MR-HIFU hyperthermia setup and treatment in Section 4.2. Second, in Section 4.3 we derive the SAcSS model describing the system dynamics, and provide the problem formulation. Next, the target-conformal transducer location and sonication point optimization is presented in Section 4.4. In Section 4.5, the MI-MPC setup and corresponding simulation results are given, and in Section 4.6 the key contributions of this chapter are summarized.

4.2 Hyperthermia treatment and setup

In this section, we discuss the MR-HIFU hyperthermia treatment and therapy platform.

4.2.1 Hyperthermia treatment

Consider a target region of interest (ROI) $\mathcal{R}$ outlining the tumor, and its corresponding safety boundary region $\mathcal{S}$, see Figure 4.4, in the patient domain Ω, where $\mathcal{R} \subset \mathcal{S} \subset \Omega$. The temperature objectives for hyperthermia therapy are schematically depicted in cross-section perspective in Figure 4.1. The main goal is to achieve constant and uniform heating of $\mathcal{R}$, with optimal treatment quality being achieved at 42 °C [4, 12], indicated by the reference temperature $T_r : \mathcal{R} \to \mathbb{R}$ in the figure. However, temperatures above 43 °C inside the ROI must be prevented, as this significantly deteriorates treatment quality and may cause direct tissue damage. In addition, the healthy tissue, defined as the area

Figure 4.2: MR-HIFU system consisting of a Philips MRI scanner and Profound Sonalleve® therapy platform.

outside $\mathcal{S}$, must remain cooler than 40 °C, since above this temperature the heat-induced tissue sensitization gradually increases. Correspondingly, the temperature upper bound $\overline{T} : \Omega \to \mathbb{R}$ is defined. The maximum violation of this bound is measured by $\epsilon \in \mathbb{R}_{\geq 0}$.

4.2.2 MR-HIFU hyperthermia system

In this work, we consider a MR-HIFU hyperthermia treatment setup similar to Figure 4.2. Note, however, that the principles on which this work's optimal actuator placement method is based can also be used for the control design of other MR-HIFU setups. The main components are a Philips 3T Achieva® MRI scanner, using which volumetric temperature maps can be obtained noninvasively and in near-real time, and a Profound Sonalleve® HIFU applicator platform, consisting of a patient bed in which an MR-compatible HIFU transducer and its mechanical robotic carrier are integrated. Currently, this setup is already used in clinics for the treatment of uterine fibroids and for palliative treatment of painful bone metastases.

4.2.2.1 MR thermometry

By exploiting the temperature-dependent proton resonance frequency shift (PRFS) in water molecules, the temperature of tissues inside the body can be measured noninvasively using an MRI scanner [38, 44]. To be precise, the temperature change with respect to a reference can be measured. By acquiring the reference image before treatment, the actual temperature can be reconstructed during treatment, with superior accuracy and resolution compared to other imaging methods [30, 39, 40]. However, due to its relative nature, it is sensitive to changes in the magnetic field, such as those resulting from repositioning the transducer, resulting in non-temperature-related measurement artifacts. To

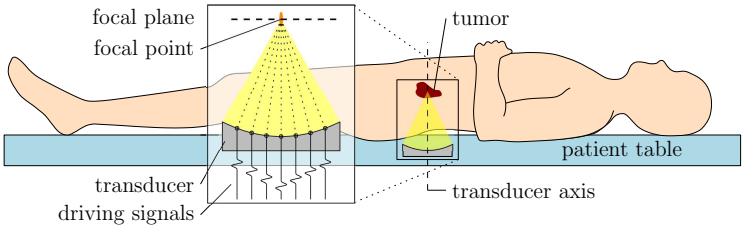

Figure 4.3: Schematic of a HIFU beam into the focal plane in the tumor, with the focal point by electronic beam steering.

avoid this, a multi-baseline approach is employed, where before treatment reference images are obtained for each of the transducer positions, such that during treatment the baseline corresponding to the current transducer position can be used for accurate thermometry. In case of intra-scan transducer motion, however, the obtained temperature map is considered unreliable.

4.2.2.2 HIFU transducer

By constructively interfering ultrasound waves generated by the 256 acoustic elements of the phased-array HIFU transducer inside the Sonalleve system, a focal spot can be created, see Figure 4.3. Using coordinated phase shifts in the driving signals, the focus location can be moved, which is called electronic beam steering, up to a lateral deflection of $R_{\mathrm{EBS}} = 8$ mm. The focus is relatively long along the beam axis (7 mm), but narrow in lateral direction (2 mm), and additionally significant acoustic beam overlap occurs directly in front of and behind the focal plane, see Chapter 2 and [53, 57, 85]. Consequently, the temperature distribution inside the focal plane typically characterizes the temperature in its surrounding volume. Correspondingly, we choose the focal plane such that it bisects the target volume, and consider only the patient domain in this plane, i.e., $\Omega \subset \mathbb{R}^2$, in designing the temperature controller.

We call the 16 mm diameter ($= 2R_{\mathrm{EBS}}$) circular area around the transducer axis, corresponding to the maximum electronic beam steering range, a treatment cell. For larger tumors, multiple (possibly overlapping) treatment cells must be defined. An example is shown in Figure 4.4, depicting a target ROI $\mathcal{R}$, safety region $\mathcal{S}$, and four transducer positions (red $\times$) and corresponding treatment cells (red circles). Inside a treatment cell, the focus can quickly be switched between all corresponding sonication points ($\bullet$), which in the figure are distributed using a fixed pattern. The different cells can be heated sequentially by mechanically switching between the transducer positions [53].

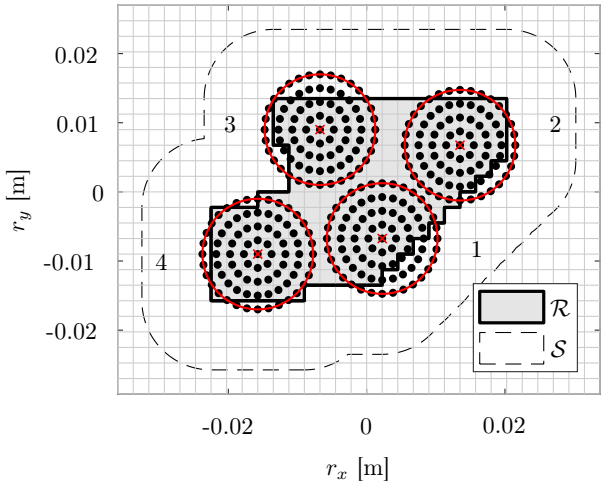

Figure 4.4: Four transducer positions (red ×) and corresponding treatment cells (red circles) and sonication points (•) spread over a large irregularly shaped ROI $\mathcal{R}$ with corresponding safety boundary region $\mathcal{S}$.

4.3 Modeling for control

In this section, we derive a model of the tissue's thermal dynamics and the transducer switching, and formulate the objective of the optimal actuator placement problem.

4

4.3.1 Thermal model

4.3.1.1 Transducer positioning

Following the switched-actuator system with setup times (SAcSS) framework proposed in Chapter 3, we first model the mechanical transducer switching using an arc-weighted digraph $\Gamma = (\mathcal{Q}, \mathcal{E}, s)$ as depicted in Figure 4.5 for $N_q = 4$ nodes. Here, $\mathcal{Q} = \mathbb{N}_{[1,N_q]}$ (nodes) represents the set of N_q operational actuator modes, resulting from the transducer being in each of its admissible positions. The weighted arcs $\mathcal{E} = \mathcal{Q}^2 \setminus \{(q,q) \in \mathcal{Q}^2\}$ (numbered arrows) correspond to all possible transducer relocations. The arc weights $s_{q\tilde{q}} \in \mathbb{N}$, $(q, \tilde{q}) \in \mathcal{E}$, represent the number of MR thermometry sampling periods required to complete the mechanical transducer displacement from position q to $\tilde{q}$. In this work, the sample time is $T_s = 3.2$ s. Moreover, let us remark that although in Figure 4.5 the setup times are independent of the switching direction between two modes, this is not necessary. Next, the actuator state is denoted by $\sigma_k \in \mathcal{Q} \cup \mathcal{E}$, where $\sigma_k = q \in \mathcal{Q}$ means that at discrete time $k \in \mathbb{N}$, connected to real time via $t_k = kT_s$, the

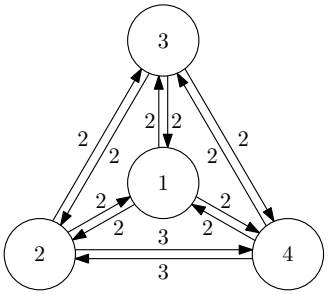

Figure 4.5: Four-nodal arc-weighted digraph example.

transducer is in its q-th admissible position, and $\sigma_k = (q, \tilde{q}) \in \mathcal{E}$ indicates that the transducer is being relocated from position q to $\tilde{q}$.

4.3.1.2 Thermal dynamics

Next, continuing the SAcSS modeling framework of Chapter 3, we model the tissue's thermal response using the discrete-time state-space system

$$x_{k+1} = Ax_k + B(\mathcal{P})u_k, \tag{4.1a}$$

$$y_k = \begin{cases} x_k + v_k, & \text{if } \sigma_{k-1} \in \mathcal{Q}, \\ \emptyset, & \text{if } \sigma_{k-1} \in \mathcal{E}, \end{cases} \tag{4.1b}$$

which is derived by spatially and temporally discretizing the Pennes bioheat equation [86] with the homogeneous tissue parameters from [95] using the MR thermometry grid, consisting of 2.25×2.25 mm^2 voxels (grid lines in Figure 4.4), and the MR sample time $T_s = 3.2$ s mentioned previously. Consequently, $x_k \in \mathbb{R}^{n_x}$ represents the voxels' treatment-induced temperature elevation, where typically $n_x = 36^2 = 1296$ for tumors up to 4×4 cm^2, at discrete time $k \in \mathbb{N}$. Correspondingly, the matrix A has the form of a graph Laplacian matrix, capturing the effects of heat conduction inside the focal plane, and the heat loss resulting from blood perfusion, using the database muscle tissue parameters from [95]. Since we have chosen the voxel grid to correspond to the MR thermometry grid, the output $y_k \in \mathbb{R}^{n_x}$ consists of the state, corrupted by the measurement noise $v_k \sim \mathcal{N}(0, 0.4^2 I_{n_x})$, when measurements are available due to the transducer being at standstill during the scanning interval ($\sigma_{k-1} \in \mathcal{Q}$). In case the transducer was moving ($\sigma_{k-1} \in \mathcal{E}$), no reliable temperature map can be acquired, as expressed by $y_k = \emptyset$ in (4.1b). The input vector

$$u_k = \begin{bmatrix} u_k^1 \\ \vdots \\ u_k^{N_q} \end{bmatrix} \in \mathbb{R}^{n_u}_{[0, u_{\max}]} \tag{4.2}$$

represents the acoustic powers applied at the set of sonication points $\mathcal{P} = \{r_p^1, \ldots, r_p^{n_u}\}$, in which the points $r_p^i \in \Omega$ are assumed to be ordered per treatment cell, such that in (4.2) $u_k^q \in \mathbb{R}^{n_u^q}$ denotes the sonication point powers corresponding to cell $q \in \mathcal{Q}$. To be precise, u_k denotes the average power applied to the sonication points $\mathcal{P}$ during the interval from t_k to t_{k+1}, resulting from rapidly switching the focal spot over all points at which nonzero power is requested, which is justified by the electronic beam steering rate being considerably faster than the timescale of the thermal dynamics. In (4.2), the lower input bound reflects that we cannot actively cool tissue, and the upper bound $u_{\max} = 15$ W is enforced for safety. In addition, the total input power is upper bounded by $\|u_k\|_1 \leq \overline{u_\Sigma} = 100$ W. Using 0_n to denote the zero-vector of length n, the transducer switching is incorporated by imposing that

$$u_k^{\tilde{q}} = 0_{n_u^{\tilde{q}}} \text{ for all } \tilde{q} \in \mathcal{Q} \setminus \{\sigma_k\}, \tag{4.3}$$

i.e., heating can only occur at the sonication points of cell q if the transducer is at the corresponding position by $\sigma_k = q$. If, on the other hand, $\sigma_k \in \mathcal{E}$, i.e., the transducer is moving, we have in (4.3) that $\mathcal{Q} \setminus \{\sigma_k\} = \mathcal{Q}$, by which no sonication is allowed according to $u_k = 0_{n_u}$. Finally, the input matrix corresponding to the set $\mathcal{P} = \{r_p^1, \ldots, r_p^{n_u}\} \subset \Omega$ is given by

$$B(\mathcal{P}) = \begin{bmatrix} b(r_p^1) & \cdots & b(r_p^{n_u}) \end{bmatrix} \in \mathbb{R}^{n_x \times n_u}, \tag{4.4}$$

where $b(r_p^i) \in \mathbb{R}^{n_x}$ describes the system's temperature change in response to unit heating at sonication point r_p^i, which is obtained by spatial and temporal discretization of the corresponding power deposition intensity. The latter is modeled as a radially symmetric Gaussian distribution centered around the considered sonication point $r_p \in \Omega$ by

$$F(r) = \frac{\alpha_f}{2\pi\sigma_f^2} \exp\left(-\frac{\|r - r_p\|^2}{2\sigma_f^2}\right), \tag{4.5}$$

with spatial coordinate $r \in \Omega$, and where in this work we use $\sigma_f = 2.4$ mm to determine the profile width and $\alpha_f = 0.0267$ as the acoustic energy absorption coefficient.

4.3.2 Problem formulation

The set of admissible transducer positions, denoted by $\mathcal{T} \subset \Omega$, and the sonication points $\mathcal{P}$ must be chosen before treatment to allow for controller synthesis and acquiring the MR thermometry reference library. Choosing them optimally is crucial for achieving high-quality hyperthermia treatments, as via the input matrix $B(\mathcal{P})$ in (4.1), the achievable temperature distribution of (4.1) strongly depends on the points $\mathcal{P}$, which in turn are constrained by $\mathcal{T}$ due to the finite

electronic beam steering range. Corresponding to the temperature objectives discussed in Section 4.2.1 and the spatial discretization of (4.1), optimizing the treatment quality can be expressed as minimizing the cost

$$\ell(z_k, \epsilon(x_k)) = (z_k - z_r)^\top Q(z_k - z_r) + f_\epsilon \epsilon(x_k), \tag{4.6a}$$

where $z_k = Hx_k \in \mathbb{R}^{n_z}$, $n_z < n_x$, represents the target voxel temperatures, using the matrix $H \in \{0,1\}^{n_z \times n_x}$ containing a single 1 per row and at most one 1 per column, and $z_r \in \mathbb{R}^{n_z}$ is the reference temperature corresponding to T_r on the voxel grid. Similarly, $\epsilon(x_k) = \|\max\{x_k - \overline{x}, 0_{n_x}\}\|_\infty \in \mathbb{R}_{\geq 0}$, where the maximum operator is used element-wise, corresponds to the maximum violation of the temperature upper bound $\overline{x} \in \mathbb{R}^{n_x}$ obtained by spatially discretizing $\overline{T}$. Using

$$Q = \frac{1}{n_z} I_{n_z}, \quad f_\epsilon = 10, \tag{4.6b}$$

which are normalized with respect to the corresponding number of weighted variables (note that ϵ is scalar), (4.6a) measures the temperature error with respect to the reference, and incorporates the temperature upper bound as a soft constraint.

Unfortunately, finding the sets $\mathcal{P}$ and $\mathcal{T}$ that allow for the system (4.1) to achieve high treatment quality according to (4.6) is not straightforward, especially for irregularly shaped tumors and in case of inhomogeneous tissue properties, e.g., due to chaotic vascularization. As a solution, one could attempt to simply define many sonication points and transducer positions, as this generally allows for adequate control performance due to the controller having sufficient degrees of freedom. However, as the sonication powers and mechanical transducer motions have to be determined by the feedback controller, this also increases the online complexity, which is potentially problematic for real-time implementation, especially for optimization-based strategies such as MPC. As a solution, we propose a target-conformal optimal actuator placement method, of which the objective can be formulated as: *Find a set of sonication points and transducer positions that is small, thereby avoiding excessive controller complexity, but which is distributed in such a manner that the hyperthermia treatment quality* (4.6) *is optimized.*

4.4 Target-conformal actuator placement

This section proposes a three-step procedure for optimizing the set of transducer positions and sonication points before treatment.

4.4.1 Outline

Ideally, the optimization problem to be solved here would balance computational effort, control performance, and robustness, using a variable number of

treatment cells and sonication points. Unfortunately, this leads to many decision variables, which are also nonlinearly constrained due to the finite electronic beam steering range. Moreover, the relation between these design parameters and the mentioned objectives is highly nonlinear, and may be non-trivial to obtain. For example, deriving an explicit expression of the computational burden for the purpose of penalizing it relative to some performance measure is already not straightforward. If, additionally, the sonication points are restricted to a discrete set of locations, as will be done in this work, the optimization is of the mixed-integer type and the (nonlinear) cost function becomes discontinuous, which further complicates the problem.

To make the procedure computationally tractable and, in addition, conceptually more insightful, it is split into three steps. In particular, we first optimize the transducer positions, based on which subsequently a large set of candidate sonication points within electronic beam steering range $\mathcal{P}_c$ is determined. Then, using the thermal model, we select from the candidate sonication points $\mathcal{P}_c$ a substantially smaller subset $\mathcal{P}_s$ that is expected to be critical for steady-state performance. Finally, to improve the transient performance and the robustness against unmodeled heat losses, a small set of additional sonication points $\mathcal{P}_u$ are selected from $\mathcal{P}_c \setminus \mathcal{P}_u$ to improve target coverage.

4.4.2 Procedure

4.4.2.1 Transducer positions

The transducer positions are determined by solving an optimization problem that seeks to maximize both steady-state performance and target coverage. However, since the heating occurs by sonication, but the sonication points are not yet known and, due to the limited electronic beam steering, are constrained by the to-be-determined transducer positions, we first set up the intermediate system

$$x_{k+1} = Ax_k + B(\mathcal{P}_t(\mathcal{T}))u_t, \tag{4.7}$$

where the input matrix is explicitly parameterized in terms of the set of transducer positions $\mathcal{T} = \{r_t^1, \ldots, r_t^{N_q}\}$, with $r_t^q \in \Omega$ for $q \in \mathcal{Q}$, and where N_q is user-defined. For a given $\mathcal{T}$, the set $\mathcal{P}_t(\mathcal{T})$ represents the intermediate sonication points, which follow directly from the transducer positions by using a fixed sonication point pattern per cell, such that changing $\mathcal{T}$ immediately leads to a different input matrix $B(\mathcal{P}_t(\mathcal{T}))$. This avoids the need for including the sonication points as decision variables, and instead requires only the optimization of the $2N_q$ planar coordinates of the transducer positions $\mathcal{T}$, thus considerably improving the computational tractability of the upcoming nonlinear transducer optimization problem. To optimize $\mathcal{T}$, the fixed sonication point pattern must be able to capture the full heating potential of a cell, such that in the subsequent actuator placement steps a small, but strategically selected, set of sonication points

can be determined that achieves similar heating performance. To this end, we overpopulate each cell using the pattern in Figure 4.4, consisting of $N_{pp} = 81$ sonication points divided over one center point and four concentric rings of diameters 4, 8, 12, and 16 mm containing 8, 16, 24, and 32 evenly distributed points, respectively. The corresponding intermediate input $u_t \in \mathbb{R}^{N_{pt}}_{[0,u_{\max}]}$, with $N_{pt} = N_q N_{pp}$, denotes the sonication powers at $\mathcal{P}_t(\mathcal{T})$.

Using (4.7), we formulate the nonlinear program

$$\min_{\mathcal{T}\subset\Omega} \left(\frac{\alpha}{\alpha_p} g_p(\mathcal{T}) + \frac{1-\alpha}{\alpha_c} g_c(\mathcal{T}) \right), \tag{4.8a}$$

with performance cost

$$g_p(\mathcal{T}) = \min_{x,u_t} \ \ell(z, \epsilon(x)) \tag{4.8b}$$

$$\text{s.t.} \quad x = Ax + B(\mathcal{P}_t(\mathcal{T}))u_t, \tag{4.8c}$$

$$0_{N_{pt}} \leq u_t \leq 1_{N_{pt}} u_{\max}, \tag{4.8d}$$

$$1_{N_{pt}}^\top u_t \leq \overline{u_\Sigma}, \tag{4.8e}$$

and coverage cost

$$g_c(\mathcal{T}) = \sum_{r\in\mathcal{R}\cap\mathcal{V}} d(r, \mathcal{C}(\mathcal{T})). \tag{4.8f}$$

The performance cost $g_p(\mathcal{T})$ (4.8b), where $z = Hx$, is based on the treatment quality cost from (4.6), which by minimization subject to (4.8c)-(4.8e) is a measure for the optimally achievable steady-state temperature for a given $\mathcal{T}$. To be precise, (4.8c) ensures that the intermediate system (4.7) is in steady state, and by (4.8d)-(4.8e), where 1_n denotes the all-ones vector of length n, the input u_t satisfies the acoustic power limits discussed in Section 4.3.1.2.

Remark 4.1. *The constraints* (4.8c)-(4.8e) *do not incorporate the input restrictions resulting from the transducer switching as described in* (4.3)*. Instead, they represent the "averaged" system, allowing for sonication at all points $\mathcal{P}_t(\mathcal{T})$ simultaneously, but with the total heating power limited to $\overline{u_\Sigma}$ instead of $N_q\overline{u_\Sigma}$. This is done for computational tractability, as incorporating the switching behavior would typically result in a periodic steady state (sequentially heating all treatment cells at which steady-state heat losses must be counteracted) of which the period is unknown a priori and depends on $\mathcal{T}$, and hence must be determined during optimization.*

The cost $g_c(\mathcal{T})$ (4.8f) is used as a measure of target coverage, which can be interpreted as a type of continuous maximal covering location problem [113]. Here, $\mathcal{V} = \{r_v^i, \ldots, r_v^{n_x}\}$ denotes the set of voxel coordinates $r_v^i \in \Omega$, $i \in \mathbb{N}_{[1,n_x]}$, and consequently $\mathcal{R} \cap \mathcal{V}$ is the collection of all target voxel coordinates. Next,

the set $\mathcal{C}(\mathcal{T}) \subset \mathbb{R}^2$ denotes the union of the treatment cell areas corresponding to the transducer positions $\mathcal{T}$, i.e.,

$$\mathcal{C}(\mathcal{T}) = \bigcup_{q \in \mathcal{Q}} \{r_t^q\} \oplus \mathcal{D}(R_{\mathrm{EBS}}), \tag{4.9}$$

where $\mathcal{D}(R_{\mathrm{EBS}}) = \{r \in \mathbb{R}^2 \mid \|r\|_2 \leq R_{\mathrm{EBS}}\}$ is the disk centered at zero with radius equal to the electronic beam steering range R_{EBS}, and $\oplus$ denotes the Minkowski sum given by $\mathcal{A} \oplus \mathcal{B} = \{a + b \mid a \in \mathcal{A}, b \in \mathcal{B}\}$. Also,

$$d(r, \mathcal{C}(\mathcal{T})) = \min\{\|r - r_c\|_2 \mid r_c \in \mathcal{C}(\mathcal{T})\} \tag{4.10}$$

measures the shortest Euclidian distance from some point $r \in \mathbb{R}^2$ to $\mathcal{C}(\mathcal{T})$. As a result, in finding the set of positions $\mathcal{T}$ that minimizes $g_p(\mathcal{T})$, we effectively seek to cover as many ROI voxels $r \in \mathcal{R} \cap \mathcal{V}$ by the treatment cell areas $\mathcal{C}(\mathcal{T})$ (as $d(r, \mathcal{C}(\mathcal{T})) = 0$ for $r \in \mathcal{C}(\mathcal{T})$), while also minimizing the distance from all non-covered target voxels to the nearest point in $\mathcal{C}(\mathcal{T})$. The latter improves the ability to heat the non-covered points by diffusion, and helps to distribute the cells more evenly throughout the target.

The performance and coverage objectives are combined using a scaling variable $\alpha \in \mathbb{R}_{[0,1]}$, determining the balance between these objectives. Due to the nonlinearity of (4.8), α does not provide a linear balancing. We do, however, make the scaling more intuitive using the normalization constants

$$\alpha_p = \frac{g_p(\mathcal{T}_p) + g_p(\mathcal{T}_c)}{2}, \quad \alpha_c = \frac{g_c(\mathcal{T}_p) + g_c(\mathcal{T}_c)}{2}, \tag{4.11}$$

representing the "mean" performance and coverage costs, respectively, between the performance-optimal transducer positions $\mathcal{T}_p = \arg\min_{\mathcal{T} \subset \Omega} g_p(\mathcal{T})$ and the coverage-optimal positions $\mathcal{T}_c = \arg\min_{\mathcal{T} \subset \Omega} g_c(\mathcal{T})$.

4.4.2.2 Sonication points for steady state

For a given set of admissible transducer positions $\mathcal{T}$, we first determine the N_{pc} candidate sonication points $\mathcal{P}_c$. These can in principle be chosen as any arbitrary set of points within the cell areas $\mathcal{C}(\mathcal{T})$, for example using the fixed pattern in the previous step, in which case $\mathcal{P}_c = \mathcal{P}_t(\mathcal{T})$. In this work, however, we select $\mathcal{P}_c$ as the voxel centers inside the union of the optimized cell areas, i.e., $\mathcal{P}_c = \mathcal{C}(\mathcal{T}) \cap \mathcal{V}$, as this leads to better model and measurement accuracy. Note that, nevertheless, the remainder of this procedure can be applied to any finite and discrete set $\mathcal{P}_c \subset \mathcal{C}(\mathcal{T})$.

To identify the sonication points critical for steady-state performance, denoted by $\mathcal{P}_s \subseteq \mathcal{P}_c$, we first compute the optimal steady-state temperature distribution for the candidate system

$$x_{k+1} = Ax_k + B_c u_c. \tag{4.12}$$

where the input $u_c \in \mathbb{R}^{N_{pc}}$ and matrix $B_c \coloneqq B(\mathcal{P}_c)$ correspond to using all candidate points $\mathcal{P}_c$. To this end, we solve

$$\min_{x,u_c} \ell(z, \epsilon(x)), \tag{4.13a}$$

subject to

$$x = Ax + B_c u_c, \tag{4.13b}$$
$$0_{N_{pc}} \leq u_c \leq 1_{N_{pc}} u_{\max}, \tag{4.13c}$$
$$1_{N_{pc}}^\top u_c \leq \overline{u_\Sigma}. \tag{4.13d}$$

which can be written as a computationally efficient quadratic program, and yields the optimal steady-state input u_c^*. Subsequently, we select only the sonication points with the $N_{ps} \leq N_{pc}$ largest acoustic powers, where N_{ps} is user-defined or follows from selecting only the points with sonication powers above some user-defined threshold, such as a fraction of the largest optimal sonication power $\alpha_s \|u_c^*\|_\infty$ with $\alpha_s \in \mathbb{R}_{[0,1]}$. As (4.13) is based on the thermal model via (4.13b), the points $\mathcal{P}_s$ are typically around the ROI edge and in areas with increased perfusion according to the model (4.1), such as in the vicinity of blood vessels, to counteract the anticipated diffusive and perfusive steady-state heat losses, respectively.

Remark 4.2. *By using $\mathcal{P}_c = \mathcal{C}(\mathcal{T}) \cap \mathcal{V} \neq \mathcal{P}_t(\mathcal{T})$, we effectively introduce quantization in selecting the candidate sonication points. As* (4.8) *is a* continuous *optimization, this may cause some mismatch between the optimal heating profiles $B(\mathcal{P}_t(\mathcal{T}))u_t^*$ in* (4.8b) *and $B_c u_c^*$ in* (4.13)*, corresponding to $\mathcal{P}_t(\mathcal{T})$ and $\mathcal{P}_c$, respectively. To reduce this quantization-induced mismatch, a minor shift can be applied to $\mathcal{T}$ such that the optimal heating profiles $B(\mathcal{P}_t(\mathcal{T}))u_t^*$ and $B_c u_c^*$ are more similar.*

4.4.2.3 Sonication points for uniform coverage

The final step is to include $N_{pu} \leq N_{pc} - N_{ps}$ additional sonication points throughout the ROI, with N_{pu} being user-defined, to improve transient performance during the initial heat-up phase and to enhance the controller's robustness against unexpected spatially localized disturbances or unmodeled tissue inhomogeneities. Especially for the latter, we cannot use the thermal model to determine which sonication points must be selected. Therefore, we instead desire the additional points $\mathcal{P}_u \subseteq \mathcal{P}_c \setminus \mathcal{P}_s$ to be distributed in such a manner that, combined with the steady-state-critical sonication points in $\mathcal{P}_s$, they are spread over the target as uniformly as possible. This can be recognized to be a conditional k-median (with $k = N_{pu}$) facility location problem [113], in which the objective is to select from $\mathcal{P}_c \setminus \mathcal{P}_s$ the points $\mathcal{P}_u$ that minimize the mean distance from each

target voxel $r \in \mathbb{R} \cap \mathcal{V}$ to its nearest sonication point in $\mathcal{P}_s \cup \mathcal{P}_u$. This can be solved efficiently in the form of the mixed-integer linear program

$$\min_{\gamma,\beta} \sum_{p\in\mathcal{P}_c} \sum_{r\in\mathcal{R}\cap\mathcal{V}} d_{pr}\beta_{pr}, \tag{4.14a}$$

subject to

$$\gamma_p = 1, \qquad \text{for all } p \in \mathcal{P}_s, \tag{4.14b}$$

$$\sum_{p\in\mathcal{P}_u} \gamma_p = N_{pu}, \tag{4.14c}$$

$$\sum_{p\in\mathcal{P}_c} \beta_{pr} = 1, \qquad \text{for all } r \in \mathcal{R} \cap \mathcal{V}, \tag{4.14d}$$

$$\beta_{pr} \leq \gamma_p, \qquad \text{for all } p \in \mathcal{P}_c \text{ and } r \in \mathcal{R} \cap \mathcal{V}, \tag{4.14e}$$

where $\gamma = \{\gamma_p\}_{p\in\mathcal{P}_c}$ is the collection of the N_{pc} Boolean selection variables $\gamma_p \in \{0, 1\}$, with $\gamma_p = 1$ indicating that candidate point p is selected, i.e., $p \in \mathcal{P}_s \cup \mathcal{P}_u$, and $\gamma_p = 0$ otherwise. Using (4.10), the weights are defined as $d_{pr} = d(p, \{r\}) \in \mathbb{R}_{\geq 0}$, each representing the Euclidian distance between candidate sonication point $p \in \mathcal{P}_c$ and target voxel center $r \in \mathcal{R} \cap \mathcal{V}$, and which are computed before solving (4.14). Next, $\beta = \{\beta_{pr}\}_{p\in\mathcal{P}_c, r\in\mathcal{R}\cap\mathcal{V}}$ denotes the collection of the $N_{pc}n_z$ assignment variables, where $\beta_{pr} \in \mathbb{R}_{[0,1]}$ represents the fraction of d_{pr} that is penalized, for which by minimization it holds that if $\beta_{pr} = 1$ then p is (one of) the selected sonication point(s) closest to target voxel $r \in \mathcal{R} \cap \mathcal{V}$. The equalities (4.14b) fix the set of steady-state-critical points $\mathcal{P}_s$ to be selected, and (4.14c) requires that N_{pu} additional points must be assigned. By (4.14d), the distance from each ROI voxel to its closest selected sonication point (or combinations of distances to multiple equally close selected points) must be fully accounted for, where (4.14e) ensures that only the distances to selected sonication points can be penalized. Thus, (4.14) represents minimizing the sum of distances, and thereby indeed the mean distance, from each target voxel center to its nearest selected sonication point. The number N_{pu} can be updated iteratively until the ROI is considered to be adequately covered, which can be measured in terms of the mean distance between each target voxel and its closest sonication point, i.e., the mean of $d(r, \mathcal{P}_s \cup \mathcal{P}_u)$ over all $r \in \mathcal{R} \cap \mathcal{V}$.

4.4.3 Results

As an example, we demonstrate the above procedure to derive an optimal actuator configuration for the ROI in Figure 4.5.

4.4.3.1 Transducer positions

In Figure 4.6, we show the treatment cells obtained with $\alpha \in \{0, 0.4, 1\}$ in (4.8) to balance the steady-state performance and target coverage in optimizing the

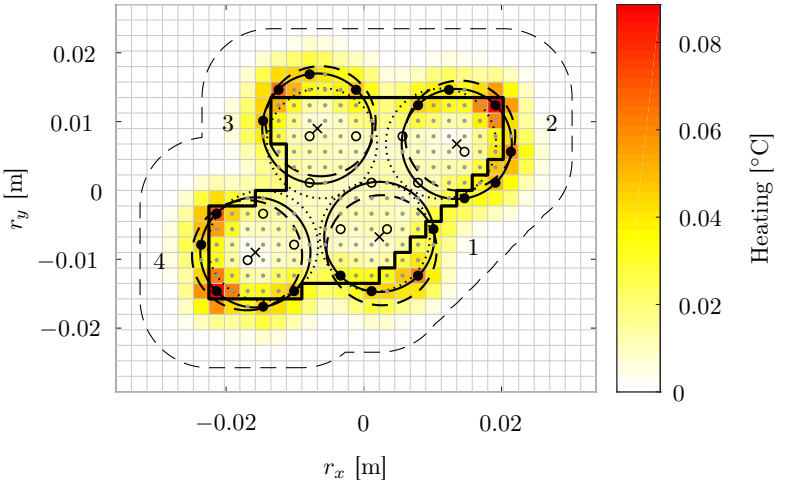

Figure 4.6: The numbered treatment cells for $\alpha = 0$ (dotted circles), $\alpha = 1$ (dashed circles), and $\alpha = 0.4$ (solid circles), for a large irregularly shaped ROI $\mathcal{R}$ (solid outline) and corresponding safety boundary $\mathcal{S}$ (dashed outline). For $\alpha = 0.4$, we also show the transducer positions $\mathcal{T}$ ($\times$), the optimal heating Bu_c^* (color map), and the sonication points $\mathcal{P}_s$ for $N_{ps} = 19$ ($\bullet$), $\mathcal{P}_u$ for $N_{pu} = 12$ ($\circ$), and $\mathcal{P}_c \setminus (\mathcal{P}_s \cup \mathcal{P}_u)$ (gray dots).

4

admissible transducer positions. It is clear to see that using the treatment cells for $\alpha = 0$ (dotted circles) maximum target coverage is achieved, but that the target corners and edge regions are outside the heating range as a result of not taking into account the achievable temperature distribution. For the cells corresponding to $\alpha = 1$ (dashed circles), on the other hand, an outward shift can be observed, which allows for a better steady-state temperature of the averaged system (in the absence of spatially localized modeling errors). For a desirable balance, we set $\alpha = 0.4$, resulting in the admissible transducer positions $\mathcal{T}$ ($\times$) with treatment cells (solid circles).

4.4.3.2 Sonication points for steady state

Using $\mathcal{T}$ from the previous step, we first define the candidate sonication points $\mathcal{P}_c = \mathcal{C}(\mathcal{T}) \cap \mathcal{V}$ ($\bullet$, $\circ$, and gray dots), from which we find $N_{pc} = 176$. Next, we solve (4.13), yielding the optimal steady-state input u_c^* corresponding to using all candidate points $\mathcal{P}_c$, of which the fifty largest normalized sonication powers are shown in Figure 4.7 in descending order. Based on this figure, we choose to select the $N_{ps} = 19$ most strongly heated sonication points, which corresponds to setting an acoustic power threshold at $0.27\|u_c^*\|_\infty$. The resulting set of steady-state-critical sonication points $\mathcal{P}_s$ ($\bullet$ in Figure 4.6) are responsible

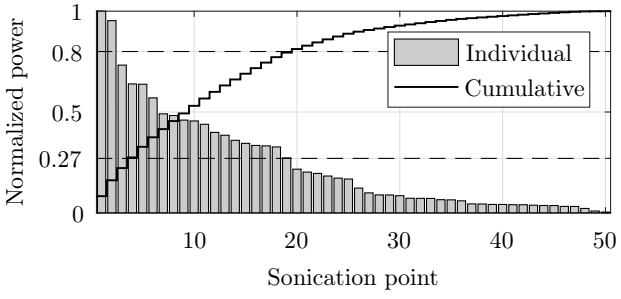

Figure 4.7: The 50 largest sonication powers in u_c^* in descending order (normalized with respect to $\|u_c^*\|_\infty$) and their cumulative sum (normalized with respect to $\|u_c^*\|_1$).

for almost 80% of the total steady-state input power. By visual comparison with the optimal power deposition $B_c u_c^*$ (color map in Figure 4.6), we verify $\mathcal{P}_s$ to correspond to the key sonication points in steady state. Note that due to using a homogeneous model, these generally lie as far as possible towards the edge of the ROI $\mathcal{R}$, or even outside $\mathcal{R}$. This allows for more effectively counteracting the diffusive steady-state heat loss, as was also observed in optimizing the transducer positions for $\alpha = 1$.

4.4.3.3 Sonication points for uniformity

Finally, given $\mathcal{P}_s$, we can solve (4.14) for a user-defined N_{pu}. As an example, we set $N_{pu} = 12$, resulting in the sonication points $\mathcal{P}_u$ ($\circ$ in Figure 4.6) corresponding to a mean distance of 2.55 mm between each target voxel and its nearest sonication point in $\mathcal{P}_s \cup \mathcal{P}_u$.

4.5 Mixed-integer model predictive control

In this section, we discuss a MI-MPC setup for large-volume MR-HIFU hyperthermia, with which simulations are performed to investigate the control performance and computation time for different actuator placement settings.

4.5.1 Observer

Due to the presence of measurement noise in (4.1b), we employ a Luenberger observer given by

$$\hat{x}_k = \begin{cases} A\hat{x}_{k-1} + Bu_{k-1} + L(y_k - \hat{y}_k^-), & \text{if } \sigma_{k-1} \in \mathcal{Q}, \\ A\hat{x}_{k-1} + Bu_{k-1}, & \text{if } \sigma_{k-1} \in \mathcal{E}, \end{cases} \tag{4.15a}$$

where

$$\hat{y}_k^- = A\hat{x}_{k-1} + Bu_{k-1}, \tag{4.15b}$$

denotes the model-based output estimate at time k before performing the measurement-based correction step in (4.15a) (if available by $\sigma_{k-1} \in \mathcal{Q}$). Using the observer gain matrix $L = 0.25I_{n_x}$, combined with A being a Schur matrix, the estimation error dynamics are exponentially stable, see Section 3.7, with desirable convergence properties.

4.5.2 Controller

As discussed in Chapter 3, we can easily set up a MI-MPC based on the SAcSS model described in Section 4.3.1, which is here given by

$$\min_{\boldsymbol{\delta}_k, \boldsymbol{u}_k} \sum_{i=0}^{N} \ell(z_{i|k}, \epsilon_{i|k}), \tag{4.16a}$$

subject to

$$\begin{aligned}
x_{i+1|k} &= Ax_{i|k} + Bu_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && (4.16b)\\
x_{0|k} &= \hat{x}_k, && && (4.16c)\\
u_{0|k} &= u^*_{1|k-1}, && && (4.16d)\\
\delta_{0|k} &= \delta^*_{1|k-1}, && && (4.16e)\\
1_{N_q}^\top \delta_{i|k} &= 1, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && (4.16f)\\
x_{i|k} &\leq \overline{x} + 1_{n_x}\epsilon_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N]}, && (4.16g)\\
0 &\leq \epsilon_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N]}, && (4.16h)\\
0_{n_u} &\leq u_{i|k} \leq \overline{u}, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && (4.16i)\\
J_u u_{i|k} &\leq \overline{u_\Sigma}\delta_{i|k}, && \forall\, i \in \mathbb{N}_{[0,N-1]}, && (4.16j)\\
J_u u_{i|k} &\leq \overline{u_\Sigma} S_\tau^\top \delta_{i-\tau|k}, && \begin{cases} \forall\, \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}\}]}, \\ \forall\, i \in \mathbb{N}_{[0,N-1]}. \end{cases} && (4.16k)
\end{aligned}$$

together with the so-called admissibility assurance (of which the computation time is negligible) to guarantee that the optimal solution $(\boldsymbol{\delta}_k^*, \boldsymbol{u}_k^*)$ is physically realizable by the HIFU setup, see Chapter 3. In (4.16a), we use the treatment quality measure from (4.6) as the stage cost, where $z_{i|k} = Hx_{i|k}$ and $\epsilon_{i|k} \coloneqq \epsilon(x_{i|k})$, and for this case study we set the prediction horizon to $N = 7$. Moreover, the sequence $\boldsymbol{u}_k = (u_{0|k}, \ldots, u_{N-1|k})$ contains the predicted sonication inputs $u_{i|k}$ at $i \in \mathbb{N}$ steps ahead of time $k \in \mathbb{N}$. Similarly, $\boldsymbol{\delta}_k = (\delta_{0|k}, \ldots, \delta_{N-1|k})$ denotes the predicted sequence of the so-called activators $\delta_{i|k}$, which are Boolean variables that encode the transducer steering in a manner compatible with mixed-integer

optimization, see Chapter 3. In summary, at discrete time $k \in \mathbb{N}$, the activator $\delta_k = [\delta_k^1 \ \dots \ \delta_k^{N_q}]^\top \in \{0,1\}^{N_q}$ follows directly from the actuator state σ_k according to

$$\delta_k^q = \begin{cases} 1, & \text{if } \sigma_k = q \in \mathcal{Q} \text{ or } \sigma_k = (\tilde{q}, q) \in \mathcal{E}, \\ 0, & \text{otherwise,} \end{cases} \tag{4.17}$$

i.e., $\delta_k^q = 1$ if at time k the transducer is located at or moving towards cell q, and $\delta_k^q = 0$ otherwise. Note that, as a consequence, δ_k must always consist of one element equal to 1, and all others 0, as is reflected by (4.16f). Regarding the other constraints, (4.16b) incorporates the thermal dynamics of (4.1), of which the initial condition is provided by the observer according to (4.16c). To allow for computation times up to one sampling period, we adopt a one-step-ahead optimization using (4.16d) and (4.16e), by which we effectively apply $u_k = u_{0|k}^* = u_{1|k-1}^*$ and $\delta_k = \delta_{0|k}^* = \delta_{1|k-1}^*$ (and for initialization we use $\delta_{0|0} = \delta_0$, corresponding to the transducer starting position, and $u_{0|0} = 0_{n_u}$). Inequalities (4.16g)-(4.16h) include the temperature upper bound as a soft constraint in such a manner that (4.16) is a mixed-integer quadratic program. Finally, as discussed in detail in Chapter 3, (4.16i)-(4.16k) impose the sonication power limits per sonication point and per treatment cell, and encode all input restrictions resulting from the transducer switching as described in Section 4.3.1. To this end, we use the matrix

$$J_u = \begin{bmatrix} 1_{n_u^1}^\top & & \\ & \ddots & \\ & & 1_{n_u^{N_q}}^\top \end{bmatrix} \in \{0,1\}^{N_q \times n_u},$$

such that $J_u u_{i|k}$ is a vector containing the total sonication powers per treatment cell, and we use the matrices

$$S_\tau = \begin{bmatrix} s_{\tau,11} & \cdots & s_{\tau,1N_q} \\ \vdots & \ddots & \vdots \\ s_{\tau,N_q1} & \cdots & s_{\tau,N_qN_q} \end{bmatrix} \in \{0,1\}^{N_q \times N_q}, \quad \tau = 1, \dots, \overline{s},$$

of which the diagonal elements are given by $s_{\tau,qq} = 0$ for $q \in \mathcal{Q}$, the off-diagonal elements follow from the arc weights $s_{q\tilde{q}}$ of the digraph in Figure 4.5 according to

$$s_{\tau,q\tilde{q}} = \begin{cases} 1, & \text{if } s_{q\tilde{q}} < \tau, \\ 0, & \text{otherwise,} \end{cases}$$

and $\overline{s} = \max_{(q,\tilde{q}) \in \mathcal{E}} s_{q\tilde{q}}$ denotes the transducer's largest cell-to-cell travel time.

4.5.3 Simulation results

To demonstrate the performance and computation time of the MI-MPC (4.16) derived using the actuator placement method proposed in this chapter, simula-

tions are performed for the ROI in Figure 4.6. To visualize the control performance, we compute the time evolution of the M-sample centered moving average of the stage cost (4.6) $\mathrm{MA}_M(\ell(z_k, \epsilon_k))$, which for odd integers $M = 1 + 2m$, $m \in \mathbb{N}$, is defined as

$$\mathrm{MA}_M(\ell(z_k, \epsilon_k)) = \frac{1}{M} \sum_{l=\max\{k-\frac{M}{2},0\}}^{k+\frac{M}{2}} \ell(z_l, \epsilon_l), \tag{4.18}$$

which smooths the switching effects over an interval of $(M-1)T_s$ for better readability. In analyzing the results of this simulation study, adequate smoothing was observed using $M = 11$. Correspondingly, for $\mathcal{T}$ optimized with $\alpha = 0.4$ ($\times$ in Figure 4.6, $\mathcal{P}_c$ determined using $N_{ps} = 19$ ($\bullet$ in Figure 4.6), and using $M = 11$, Figure 4.8a visualizes the control performance obtained using $\mathcal{P}_u$ optimized for various values of N_{pu} (solid, see legend), where $N_{pu} = 157$ corresponds to selecting all candidate sonication points by $\mathcal{P}_s \cup \mathcal{P}_u = \mathcal{P}_c$. The figure shows that the control performance is unsatisfactory for $N_{pu} \leq 2$, as this does not allow for sufficiently heating the interior of $\mathcal{R}$ for fast heat-up and for counteracting the perfusive heat loss inside $\mathcal{R}$. Moreover, we observe that the performance clearly improves by including more sonication points in $\mathcal{P}_u$. However, the added benefit for both the transient and steady-state behavior becomes progressively smaller for increasing N_{pu}, while the corresponding solver times continue to grow, see Figure 4.8b, reaching a maximum computation time of almost 20 s for $N_{pu} = 157$. An example where the maximum computation time is smaller than $T_s = 3.2$ s, but which exhibits adequate control performance, is recognized in $N_{pu} = 12$ (red in Figure 4.8a, and $\circ$ in Figure 4.6 for the corresponding $\mathcal{P}_u$). This illustrates that the target-conformal actuator design method presented in this chapter can enable real-time feasibility of more advanced (thereby often more computationally challenging) control methods such as (MI-)MPC without sacrificing significant performance, by strategically selecting only the key sonication points.

The tumor temperatures obtained with $N_{pu} \in \{12, 157\}$, from hereon also referred to as input reduction cases A and B (see Table 4.1), are depicted in Figure 4.9a, again visualizing that larger N_{pu} enables better heat-up. However, the observed temperature difference is small due to case A containing sufficiently

Table 4.1: Input reduction cases.

Case	α	N_{ps}	N_{pu}
A	0.4	19	12
B	0.4	19	157
C	0.4	0	31
D	0	19	12

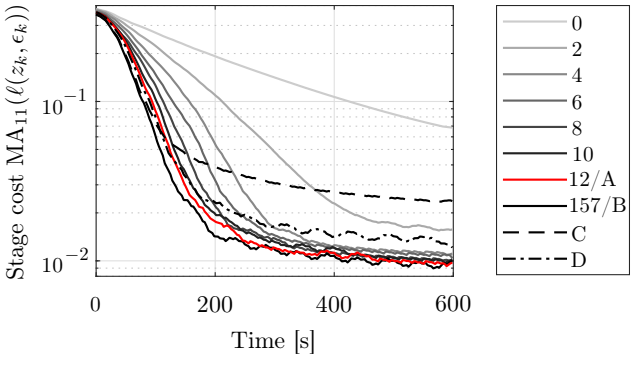

(a)

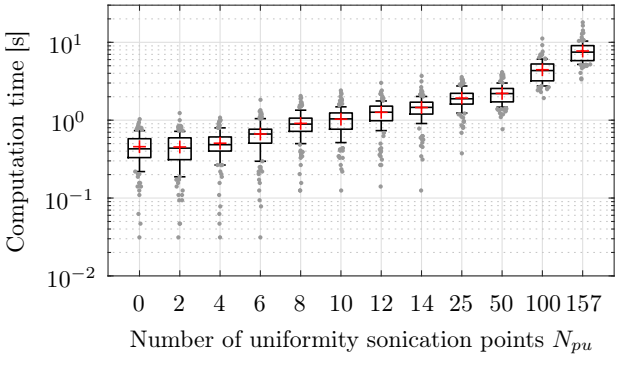

(b)

Figure 4.8: In (a), the 11-sample moving average of the stage cost $\mathrm{MA}_{11}(\ell(z_k, \epsilon_k))$ versus time t_k for $(\alpha, N_{ps}) = (0.4, 19)$ and various N_{pu} (solid, see legend), and for cases C (dashed) and D (dash-dotted) as in Table 4.1. In (b), for $(\alpha, N_{ps}) = (0.4, 19)$, the computation time versus N_{pu} indicating the mean (red +), median (central mark), 25th and 75th percentiles (box edges), 10th and 90th percentiles (whiskers), and outliers (gray dots).

many sonication points, and due to the upper power limit being active, see Figure 4.9b and Figure 4.9c, where post : $\mathcal{Q} \cup \mathcal{E} \to \mathcal{Q}$ denotes the *destination* mode according to $\mathrm{post}(\sigma_k) = q$ if $\sigma_k = q \in \mathcal{Q}$ (transducer is at position q) or $\sigma_k = (\tilde{q}, q) \in \mathcal{E}$ (transducer is moving toward position q). After approximately 200 s the ROI is found to be adequately heated for both cases. In fact, the resulting steady-state temperatures ($350 \leq t_k \leq 600$) are very similar, see Table 4.2 (A and B). This verifies that using the proposed actuator placement method, we can select from a large (case B: $N_{ps} + N_{pu} = N_{pc} = 176$) set of candidate sonication points $\mathcal{P}_c$, a significantly smaller (case A: $N_{ps} + N_{pu} = 19 + 12 = 31$)

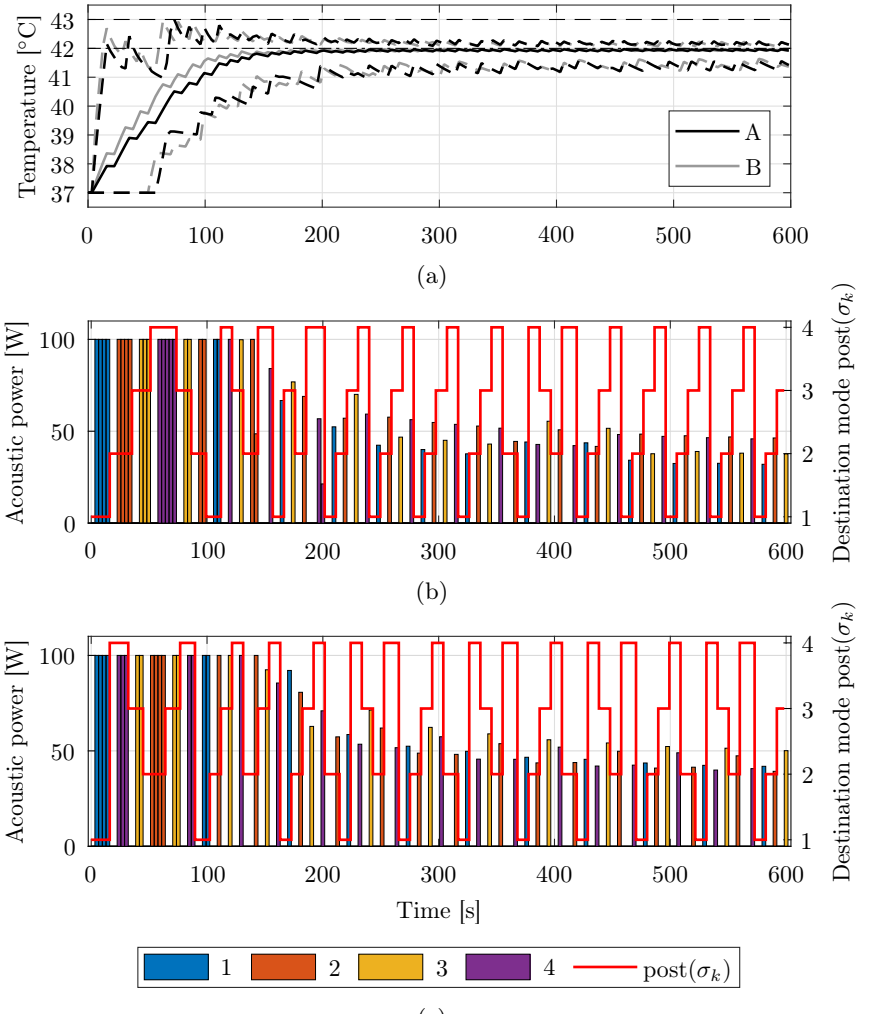

Figure 4.9: Simulation results for cases A and B, see Table 4.1, using $(\alpha, N_{ps}) = (0.4, 19)$. In (a), the mean (solid) and extremum (dashed) temperatures in the ROI $\mathcal{R}$ for $N_{pu} = 12$ (A, black) and $N_{pu} = 157$ (B, gray), and the temperature reference (thin dash-dotted) and upper bound (thin dashed). In (b) and (c), the total acoustic power per cell (legend) and the destination mode post(σ_k) (red line) for A and B, respectively.

subset $\mathcal{P}_s \cup \mathcal{P}_u$ that achieves similar performance, by taking into account the specific tumor properties during the actuator optimization.

Next, we exemplify the influence of improperly choosing $\mathcal{P}_s$ by comparing in Figure 4.8a for $\alpha = 0.4$ the smoothed stage cost for $(N_{ps}, N_{pu}) = (0, 31)$ (dashed), referred to as case C (see Table 4.1), to the cost for case A with $(N_{ps}, N_{pu}) = (19, 12)$ (red). Note that both cases contain $N_{ps} + N_{pu} = 31$ sonication points, but for case C these are all in $\mathcal{P}_u$, and hence are distributed as evenly throughout $\mathcal{R}$ using (4.14). As a consequence, in the figure case C exhibits slightly better initial heat-up due to the mean distance between the ROI voxels and their closest sonication point being smaller (1.96 mm for C versus 2.55 mm for A). However, case C fails to achieve a satisfactory steady state due to $\mathcal{P}_s = \emptyset$ lacking the sonication points at the critical locations around the target edge, eventually resulting in a significantly larger steady-state ROI temperature range, see Table 4.2 (A compared to C).

In a similar fashion, we show an example of the performance loss resulting from suboptimally choosing the transducer positions $\mathcal{T}$. To this end, compare in Figure 4.8a for $(N_{ps}, N_{pu}) = (19, 12)$ the results obtained with $\alpha = 0$ (dash-dotted), which we call case D (see Table 4.1), to the results of case A with $\alpha = 0.4$ (red). Here, for both sets of transducer positions, the time required for the mechanical cell-to-cell transducer motion is modeled by Figure 4.5. Moreover, the number of steady-state and uniformity-related sonication points are equal for both cases. However, by selecting $\mathcal{T}$ based only on the coverage-related cost in (4.8) in case D, the corresponding treatment cells (dotted circles in Figure 4.6) do not cover as much of the regions at which steady-state heating is required. As a result, case A outperforms case D for $t_k \geq 130$ in Figure 4.8a and regarding steady-state temperature in Table 4.2 (A outperforms D).

4.6 Conclusion

We proposed a target-conformal optimal actuator placement procedure for MR-HIFU hyperthermia therapies in cancer treatment. For computational tractability and conceptual simplicity, the proposed method consists of three steps. First, it uses the specific tumor properties to automatically distribute the transducer positions based on a weighted balance between the steady-state performance and target coverage by the resulting treatment cells. Second, the sonication points critical for steady-state heating are selected. Finally, additional evenly

Table 4.2: ROI temperature during $350 \leq t_k \leq 600$ for cases A, B, C, and D as in Table 4.1.

Temperature	A	B	C	D
Maximum	42.28	42.25	42.82	42.62
Mean	41.93	41.94	41.89	41.92
Minimum	41.14	41.14	40.24	40.81

distributed sonication points are included to enable more heating throughout the tumor for faster heat-up and improved robustness. In a case study involving an irregularly shaped tumor, the method has been demonstrated, and subsequently used for setting up a MI-MPC scheme. In a numerical case study, the resulting control performance and computation times have been investigated, illustrating that using our method high-quality temperature control can be achieved using a limited number of strategically selected sonication points and transducer positions, which in addition translates to limited controller complexity and thereby sufficiently small solver times for real-time feasibility.


Nothing is particularly hard
if you divide it into small jobs.
Henry Ford

Chapter 5

Hierarchical Mixed-Integer MPC for Real-Time Large-Volume Ultrasound Hyperthermia in Cancer Therapy

In oncological applications of mild hyperthermia, tumors are heated to enhance the efficacy of chemo- and radiotherapies. For magnetic-resonance-guided high-intensity focused ultrasound hyperthermia, treating tumors larger than the transducer's local heating range is a challenge, as then the transducer itself must also be mechanically relocated. Using mixed-integer model predictive control (MI-MPC) to simultaneously steer the local heating and the discrete mechanical transducer positioning, optimal treatment quality can be achieved in theory. In practice, however, MI-MPC may require prohibitively large computation times, which hampers clinical implementation. This chapter describes a two-layer hierarchical MI-MPC scheme designed to substantially decrease the computational burden, thereby enabling real-time feasibility, without sacrificing performance. This is accomplished by swiftly solving a lower-complexity mixed-integer program in the higher control layer, based on a model obtained via model order reduction and/or (target-conformal) input reduction, to determine the optimal mechanical transducer path. Subsequently, given the discrete transducer path, the lower control layer optimizes the local heating profiles using a high-fidelity model, which despite large state and input dimensions can be computed efficiently in the form of a quadratic program. For the high-level controller, both sparse and condensed MPC formulations are used to investigate which is preferable in terms of computation time for different state- and input-reduction settings. In experiments on a clinical MR-HIFU system using a tissue-mimicking phantom, the strengths of the proposed controller are demonstrated.

This chapter is based on D.A. Deenen, J. van Wordragen, L.C. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Hierarchical mixed-integer MPC for real-time large-volume ultrasound hyperthermia in cancer therapy," *In preparation for journal submission.*

5.1 Introduction

The developments presented in this chapter are aimed at enabling clinical application of large-volume magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU) hyperthermia for cancer therapy. Mild local hyperthermia in cancer treatments is the heating of tumors inside the body to around 42 °C for a duration of up to 90 minutes, while avoiding temperature elevations in healthy tissue. The prolonged exposure to mild heating has in numerous clinical studies been shown to sensitize the tissue to the desired effects of chemo- and radiotherapy, see [3, 4, 16, 19–22, 24] and the references therein, without introducing additional toxicity and hence without aggravating the severe side effects typically associated with cancer therapy. In addition, using thermo-sensitive liposomes, local hyperthermia enables targeted drug delivery, resulting in even more effective chemotherapy with reduced systemic drug concentrations [25–27]. As a consequence, hyperthermia is considered a valuable adjuvant therapy in cancer treatment, enabling significantly improved treatment outcome and survival rates. However, since the achieved therapeutic benefits are strongly related to the temperature realized during treatment [60–62], there is a dire need for hyperthermia temperature feedback controllers.

Using an extracorporeal applicator, MR-HIFU is a promising feedback technology for *noninvasive* hyperthermia therapy [30, 43, 47, 48], which is highly desirable for patients' quality of life due to the absence of additional surgery and post-interventional wound care. An MRI scanner is used for near-real-time volumetric thermometry, and by means of the coordinated modulation of constructively interfering ultrasound waves, referred to as electronic beam steering, powerful heating with spatial accuracy in the millimeter-range is applied. Using model predictive control (MPC) [66, 67], ultrasound-based thermal therapies of optimal quality can be realized, while explicitly satisfying actuator and safety constraints. For example, MPC is used for thermal ablation therapy in [69] to minimize treatment time while administering a desired scalar-valued thermal dose in a single point or over a one-dimensional model, and in [70] for minimum-time treatments with a prescribed thermal dose for two-dimensional systems described by heavily simplified models and a fixed spatial heating trajectory. In [71], an MPC scheme is developed that tries to deliver a lethal thermal dose to the tumor, while avoiding damage to healthy tissue. More recently, in [72] we have designed a preliminary MPC algorithm for temperature control in two-dimensional systems, which exploit the full flexibility of the HIFU transducer's beam steering by freely adapting online both the heating location and the acoustic power.

Unfortunately, the methods developed in [72] are only applicable to the treatment of small-size tumors due to the limited electronic beam steering range. To treat larger tumors, the transducer itself must also be moved mechanically, allowing for all parts of the tumor to be sonicated sequentially, see [52] for a

pre-planning method for HIFU treatments involving both electronic beam steering and mechanical transducer displacements, and see [53] for a binary controller with fixed beam steering trajectories and simple logic-based transducer positioning. Due to MR thermometry limitations [44], the set of admissible mechanical transducer positions is discrete, finite, and must be chosen before treatment. Consequently, the resulting system can be described as a linear switched-actuator system with setup times (SAcSS), see Chapter 3, which is a special type of switched linear system [81]. In particular, the input model depends on the transducer location, but, contrary to typical switched systems, it cannot switch instantaneously, as the time required for the transducer to complete its mechanical point-to-point motion is significant compared to the timescale of the heating dynamics. Moreover, during transducer relocation no heating is allowed, meaning every actuator mode switch induces non-negligible actuator downtime during which the system itself keeps evolving in time according to its unforced dynamics.

Mixed-integer (MI-)MPC [68] allows for optimally controlling both the range-extending discrete mechanical transducer displacements and the continuous sonication inputs. In Chapter 3, a modeling framework tailored to SAcSSs is presented, specifically designed to (a) allow for user-friendly system specification and effortless MI-MPC synthesis, and (b) yield a model that includes only a small number of Boolean variables, such that the resulting MI-MPC schemes are computationally efficient. However, for large-volume MR-HIFU hyperthermia temperature control, real-time feasibility of the obtained MI-MPC remains a major issue in practice, see [108] for a numerical case study, due to the large state and input dimensions of the required thermal models obtained by spatially discretizing Pennes' bioheat equation [86].

To alleviate the computational burden of the online MI-MPC, one could resort to model order reduction techniques [114]. Although such methods are effective in reducing model complexity, for MPC their benefits are always accompanied by decreased predictive accuracy, potentially leading to deteriorated control performance. Alternatively, the controller could be simplified by decreasing the input dimension. Unfortunately, this typically results in reduced controller flexibility, and thereby suboptimal performance and robustness, even when using an optimal input selection procedure such as in Chapter 4. However, by recognizing that for large-volume MR-HIFU hyperthermia the full-scale MI-MPC can be decomposed into a two-layer controller with different levels of accuracy, a real-time feasible *hierarchical* (H)MI-MPC scheme [82] with optimal performance can be designed. In particular, for optimizing the electronic beam steering and acoustic power, temperature predictions with high spatial accuracy are needed, whereas to solve the discrete transducer allocation problem, the predictions only need to be sufficiently detailed to determine which tumor region should be heated.

Therefore, as the main contribution of this work, we propose a HMI-MPC ar-

chitecture of which the high-level controller solves a mixed-integer quadratic program (MIQP) using a reduced-complexity thermal model, obtained via model order and/or input reduction, to find the optimal transducer path in much shorter time than the full-scale MI-MPC. The discrete transducer path is then passed on to a low-level controller that, for the given path, determines the optimal sonication plan using a high-fidelity thermal model. By fixing the integer part of the problem, the low-level optimization boils down to solving a QP, which despite large state and input dimensions is computationally lightweight. As a result, the total computational efforts are substantially reduced compared to the MI-MPC setup, without sacrificing significant performance.

The second contribution of this chapter consists of a numerical simulation study to investigate the control performance and computation time of the proposed HMI-MPC for different high-level model reductions in comparison to the unreduced MI-MPC. Here, two methods for reducing the model complexity are considered. The first involves an input reduction, for which in this work we employ the target-conformal procedure proposed in Chapter 4 to approximate the HIFU setup's full heating capabilities using significantly fewer inputs by strategically selecting only the most important inputs based on the specific tumor properties. This effectively reduces the number of required high-level MPC input decision variables to decrease its computational complexity. As a second method, we reduce the number of state decision variables by using a reduced-order model derived via proper orthogonal decomposition (POD) [114]. In addition, the high-level controller is formulated both using a sparse and a condensed approach [67, 115], to examine which is favorable in terms of computation time for the different model reduction cases.

Finally, as the third contribution, we verify the control performance and real-time feasibility of the developed HMI-MPC design via experiments on a tissue-mimicking phantom using a clinical MR-HIFU system. We compare the results obtained using our novel HMI-MPC setup to those reported in [53], which represent a clinical state-of-the-art binary controller.

The remainder of the chapter is organized as follows. In Section 5.2, we discuss the MR-HIFU hyperthermia setup and treatment. Section 5.3 provides the problem statement based on the system model and MI-MPC setup. In Section 5.4, we describe the proposed HMI-MPC architecture, followed by a comparison of the performance and computational burden using different reduced-order high-level models in Section 5.5. The results of the phantom experiments are presented Section 5.6. Finally, Section 5.7 summarizes the key contributions of this work.

5.2 Large-volume MR-HIFU hyperthermia

This section first discusses the key aspects of a hyperthermia treatment, followed a description of the MR-HIFU hyperthermia setup considered in this work.

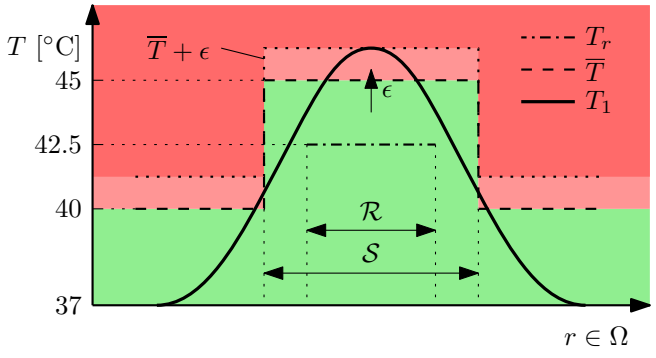

Figure 5.1: Schematic cross section of the temperature objectives corresponding to $\mathcal{R}$ and $\mathcal{S}$. The maximum violation ϵ is shown for some overheated temperature distribution T_1 such that $T_1 \leq \overline{T} + \epsilon$.

5.2.1 Hyperthermia treatment

We first define a target area $\mathcal{R}$ containing the tumor, also referred to as a region of interest (ROI), and a safety boundary region $\mathcal{S}$ in the patient domain Ω, where $\mathcal{R} \subset \mathcal{S} \subset \Omega$, as schematically shown in cross-section perspective in Figure 5.1. In a hyperthermia treatment, we strive for a constant and homogeneous temperature elevation in $\mathcal{R}$ to optimally sensitize the tumor tissue to the chemo- and radiotherapeutic cancer-killing effects [4, 12, 13, 116]. To enable a performance comparison with the binary controller results from [53] in the upcoming experiment, we set the ROI reference temperature $T_r : \mathcal{R} \to \mathbb{R}$ to 42.5 °C inside $\mathcal{R}$. Overheating the ROI ($\geq$45 °C) causes significant reversal effects and hence should be avoided. Also, to ensure that no healthy tissue is sensitized, we aim to prevent temperatures above 40 °C outside $\mathcal{S}$. Combined, these constraints lead to the temperature upper bound $\overline{T} : \Omega \to \mathbb{R}$. The maximum violation of this bound is measured by $\epsilon \in \mathbb{R}_{\geq 0}$.

5.2.2 MR-HIFU hyperthermia treatment setup

First, let us note that the HMI-MPC design proposed in this chapter is generally applicable to any SAcSS, as defined in Chapter 3, in which an accuracy-based separation can be identified between the discrete and continuous parts of the mixed-integer program describing the true optimal control problem. For the scope of this work in particular, however, the developed HMI-MPC setup is designed for an MR-HIFU hyperthermia system consisting of a Philips 3T Achieva® MRI scanner, used for near-real-time noninvasive thermometry, and a Profound Sonalleve® HIFU platform shown in Figure 5.2, which is a dedicated trolley-tabletop with an embedded MR-compatible HIFU transducer and

Figure 5.2: Philips MRI scanner and Profound Sonalleve® MR-HIFU therapy platform.

its mechanical carrier system. This system is already being used clinically for the treatment of uterine fibroids and for palliative treatment of painful bone metastases.

5.2.2.1 MR thermometry

The temperature maps are obtained noninvasively by MRI using the proton resonance frequency shift (PRFS) method [38, 44], which provides the temperature difference with respect to a baseline by comparing the current MR image to a reference image. The reference is typically acquired before treatment, such that the baseline corresponds to zero treatment-induced temperature elevation. The PRFS method is among the best available measurement technologies for image-guided hyperthermia feedback control [30, 39, 40]. However, due to it being a relative method, its accuracy may significantly deteriorate in case of distortions of the magnetic field, such as those caused by changing the HIFU transducer's position. Consequently, for accurate measurements given a certain transducer position, a baseline image is required which was obtained with the transducer in the same position. For this reason, the discrete set of admissible mechanical transducer positions must be chosen before treatment. Subsequently, a reference image can be acquired for each position, and during treatment a lookup table is used to select the correct reference image for accurate MR thermometry.

5.2.2.2 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 coordinating the individual elements' phases and amplitudes, referred to as electronic beam steering, a focal spot with variable intensity and location can be created, as schematically depicted in Figure 5.3. The resulting power depo-

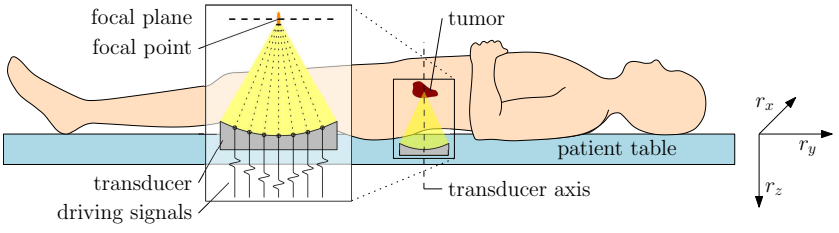

Figure 5.3: Schematic of a patient table with integrated HIFU transducer, generating an acoustic beam with the focal point inside in the tumor by electronic beam steering, and the coordinate axes as used in this work when the patient enters the MRI with the feet first.

sition profile is relatively wide in the beam's axial direction, but very narrow radially. In addition, for different in-plane focus locations, the acoustic beams are substantially overlapping everywhere along the beam axis except in the focal plane, see Chapter 2 and [53, 57, 85]. As a consequence, the axial temperature profile is fairly uniform near the focal plane, and hence accurate volumetric heating of the target tissue can be achieved by setting the focal plane at a suitable depth inside the target volume and controlling only the in-plane temperature distribution. Accordingly, the MR-HIFU hyperthermia temperature controller developed in this chapter is based on a thermal model of the patient domain in the two-dimensional focal plane $\Omega \subset \mathbb{R}^2$.

Using a stationary transducer, electronic beam steering allows for heating a 16 mm diameter circular area centered around the transducer axis, which is referred to as a treatment cell. By moving the focal spot rapidly, compared to the timescale of the tissue's thermal dynamics, the heating applied in a cell can be considered to be (quasi-)continuous in time. For the treatment of tumors larger than the electronic beam steering range, multiple treatment cells must be defined throughout the target area, which can be heated sequentially by exploiting the transducer's robotic carrier system to mechanically relocate the transducer itself [52, 53]. Note that since the treatment cells are directly coupled to the admissible transducer positions, they must also be chosen before treatment due to the MR thermometry limitations mentioned previously. As an example, Figure 5.4 depicts a large circular ROI $\mathcal{R}$ with 58 mm diameter (red circle), safety boundary $\mathcal{S}$ with 78 mm diameter (dashed circle), and seven transducer positions ($\times$) and their corresponding treatment cells (black circles), which corresponds to the largest setup programmed into the clinical Sonalleve software and presented in [53]. The remainder of the figure will be explained later in this chapter.

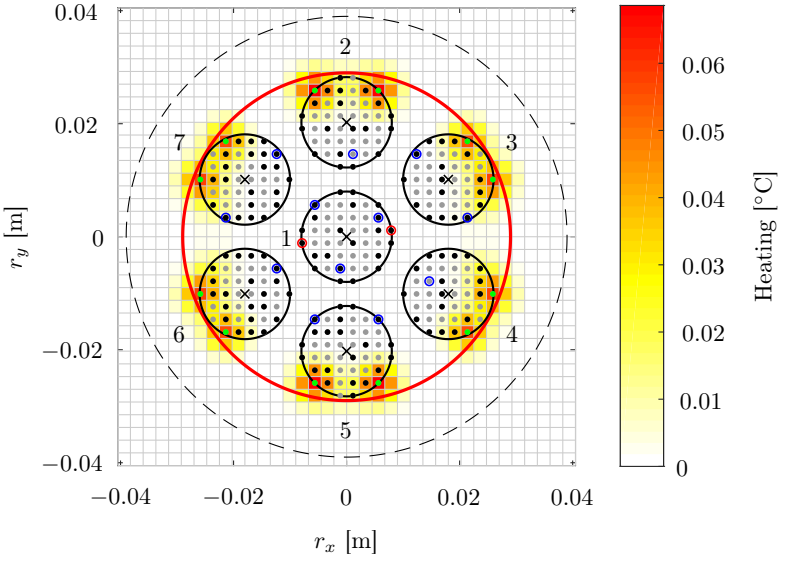

Figure 5.4: A large circular ROI $\mathcal{R}$ (red circle) and corresponding safety boundary $\mathcal{S}$ (dashed circle), the transducer positions ($\times$) and corresponding numbered treatment cells (black circles), and all candidate sonication points $\mathcal{P}_c$ (dots) with corresponding optimal "averaged" steady-state heating Bu^*_{ss} from (5.13) (color map). For $N_{ps} = 12$ steady-state-critical sonication points $\mathcal{P}_s$ (green dots), we indicate the coverage-related points $\mathcal{P}_u$ for N_{pu} equal to 2 (red $\circ$), 12 (blue $\circ$), 140 (black dots), and 280 (black and gray dots) corresponding to input reduction settings A, B, C, and D, respectively, see Table 5.1.

5.3 Problem description

In this section, we derive a compact mixed-integer model of the system using the SAcSS framework introduced in Chapter 3, with which we can directly formulate the full-scale MI-MPC that forms the key element of the problem statement.

5.3.1 Thermal Model in SAcSS Framework

5.3.1.1 Actuator modes and setup times

The transducer switching behavior is modeled using an arc-weighted digraph $\Gamma = (\mathcal{Q}, \mathcal{E}, s)$, which for the setup in Figure 5.4 is given by Figure 5.5. The $N_q = 7$ nodes model the operational actuator modes $\mathcal{Q} = \mathbb{N}_{[1,N_q]}$ corresponding to the transducer being in each of its N_q admissible positions. The set of arcs $\mathcal{E} = \mathcal{Q}^2 \setminus \{(q,q) \in \mathcal{Q}^2\}$ represents all possible transducer displacements. The arc weights $s_{q\tilde{q}} \in \mathbb{N}$, $(q,\tilde{q}) \in \mathcal{E}$, correspond to the number of sampling periods

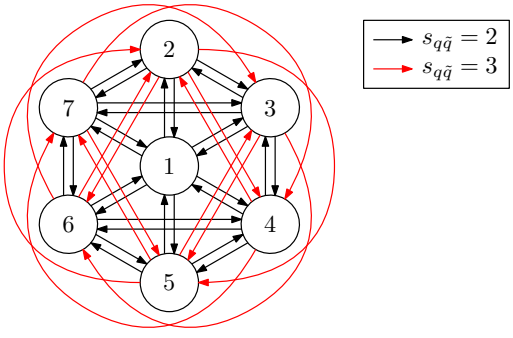

Figure 5.5: Seven-nodal arc-weighted digraph.

required for mechanically moving the transducer from position q to $\tilde{q}$, and are collected in the setup time matrix

$$S \coloneqq \begin{bmatrix} s_{11} & \cdots & s_{1N_q} \\ \vdots & \ddots & \vdots \\ s_{N_q 1} & \cdots & s_{N_q N_q} \end{bmatrix} \in \mathbb{N}^{N_q \times N_q}, \tag{5.1}$$

for which we additionally define $s_{qq} = 0$ for all $q \in \mathcal{Q}$, since staying in the same mode does not require transducer relocation. Note that although $S = S^\top$ for Figure 5.5, this is not necessary for the upcoming (H)MI-MPC setup in general. Next, we introduce the actuator state $\sigma_k \in \mathcal{Q} \cup \mathcal{E}$, where $\sigma_k = q \in \mathcal{Q}$ indicates that at discrete time $k \in \mathbb{N}$ the transducer is in its q-th admissible position, and $\sigma_k = (q, \tilde{q}) \in \mathcal{E}$ means that the transducer is moving from position q to $\tilde{q}$.

5.3.1.2 Plant dynamics

The thermal discrete-time state-space model is derived using a similar approach as discussed in Chapter 2, which essentially entails discretizing the Pennes bioheat equation [86] spatially on the two-dimensional MR thermometry grid with 2.25×2.25 mm^2 voxels (grid in Figure 5.4), and temporally using the corresponding MR thermometry sampling time $T_s = 3.2$ s. The resulting model is given by

$$x_{k+1} = Ax_k + Bu_k, \tag{5.2a}$$

$$y_k = \begin{cases} x_k + v_k, & \text{if } \sigma_{k-1} \in \mathcal{Q}, \\ \emptyset, & \text{if } \sigma_{k-1} \in \mathcal{E}, \end{cases} \tag{5.2b}$$

where the states $x_k \in \mathbb{R}^{n_x}$, with typically $n_x = 44^2 = 1936$ for tumors up to 6×6 cm^2, represent the temperature elevations (with respect to the baseline) of the voxels in the focal plane at time instant $k \in \mathbb{N}$ connecting to real time

via $t_k = kT_s$. The matrix A is a two-dimensional Laplacian matrix, describing the diffusion of heat inside the focal place and the heat loss resulting from blood perfusion. By choosing the voxel grid equal to the MR thermometry grid, the output $y_k \in \mathbb{R}^{n_x}$ is the sum of the voxel temperatures x_k and the measurement noise $v_k \in \mathbb{R}^{n_x}$, when the measurements are available due to the transducer being at standstill during the MR acquisition interval ($\sigma_{k-1} \in \mathcal{Q}$). The input $u_k \in \mathbb{R}^{n_u}$ is the acoustic power applied at the sonication points, i.e., the locations to which the focal spot may be steered during treatment. For maximum treatment quality, these points are chosen as all voxel centers inside the union of the cell areas (green, black, and gray dots in Figure 5.4). To be precise, u_k represents the sonication point powers averaged over the sampling interval from t_k to t_{k+1}, achieved by rapidly switching the focus over all sonication points at which nonzero power is requested, which is justified by the electronic beam steering being fast with respect to the thermal dynamics and sample time. Using $u_k^q \in \mathbb{R}^{n_u^q}$ to denote the acoustic powers at the n_u^q sonication points inside treatment cell $q \in \mathcal{Q}$, the input is given by

$$u_k = \begin{bmatrix} u_k^1 \\ \vdots \\ u_k^{N_q} \end{bmatrix} \in \mathbb{R}^{n_u}_{[0,u_{\max}]}, \tag{5.3}$$

where $n_u = \sum_{q\in\mathcal{Q}} n_u^q = 292$, the input space's lower bound is due to the inability to actively cool using HIFU, and the upper bound $u_{\max} = 15$ W is enforced for safety. The corresponding input matrix is given by

$$B = \begin{bmatrix} B^1 & \cdots & B^{N_q} \end{bmatrix} \in \mathbb{R}^{n_x \times n_u}, \tag{5.4}$$

where each submatrix $B^q \in \mathbb{R}^{n_x \times n_{u,\text{cell}}}$ describes the system's temperature change in response to u_k^q.

5

5.3.1.3 Input constraints

From (5.3), we find

$$0_{n_u} \leq u_k \leq \overline{u} = 1_{n_u} u_{\max}. \tag{5.5}$$

Next, we incorporate the input restrictions that describe the fact that heating can only occur in the cell where the transducer is located, i.e., $u_k^q = 0_{n_u^q}$ for all $q \neq \sigma_k \in \mathcal{Q}$, and the actuator downtime during transducer repositioning, i.e., $u_k = 0_{n_u}$ if $\sigma_k \in \mathcal{E}$. To this end, we introduce the Boolean activator $\delta_k \in \{0,1\}^{N_q}$ of which the elements δ_k^q are related to the actuator state σ_k by

$$\delta_k^q = \begin{cases} 1, & \text{if } \sigma_k = q \in \mathcal{Q} \text{ or } \sigma_k = (\tilde{q}, q) \in \mathcal{E}, \\ 0, & \text{otherwise,} \end{cases} \tag{5.6}$$

meaning that $\delta_k^q = 1$ if at time k the transducer is located at or moving towards cell q, and zero otherwise. Note that, as a consequence, δ_k is a one-hot vector [110], i.e.,

$$\sum_{q \in \mathcal{Q}} \delta_k^q = 1. \tag{5.7}$$

Due to imposing (5.5), combined with additionally desiring a maximum total power limit of $\overline{u_\Sigma} = 180$ W per cell, the mode-switching input constraints can be formulated as

$$J_u u_k \leq \overline{u_\Sigma} \delta_k, \tag{5.8a}$$

$$J_u u_k \leq \overline{u_\Sigma} S_\tau^\top \delta_{k-\tau}, \text{ for all } \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}\}]}, \tag{5.8b}$$

as discussed in Chapter 3 in more detail. Here, J_u represents the block diagonal matrix

$$J_u = \begin{bmatrix} 1_{n_u^1}^\top & & \\ & \ddots & \\ & & 1_{n_u^{N_q}}^\top \end{bmatrix} \in \{0,1\}^{N_q \times n_u}, \tag{5.8c}$$

such that $J_u u_k$ is a vector containing the total input powers of each cell, $\overline{s} = \max_{(q,\tilde{q}) \in \mathcal{E}} s_{q\tilde{q}}$ denotes the largest travel time, and the setup time constraint matrices S_τ follow directly from S in (5.1) according to

$$S_\tau = \begin{bmatrix} s_{\tau,11} & \cdots & s_{\tau,1N_q} \\ \vdots & \ddots & \vdots \\ s_{\tau,N_q1} & \cdots & s_{\tau,N_qN_q} \end{bmatrix} \in \{0,1\}^{N_q \times N_q}, \quad \tau = 1, \ldots, \overline{s}, \tag{5.9a}$$

with

$$s_{\tau,q\tilde{q}} = \begin{cases} 1, & \text{if } s_{q\tilde{q}} < \tau, \\ 0, & \text{otherwise.} \end{cases} \tag{5.9b}$$

The combination of (5.2), and (5.5)-(5.8) can be recognized to be a mixed logical dynamical (MLD) system [104], and hence is directly compatible with MI-MPC.

5.3.2 Mixed-integer model predictive control

5.3.2.1 Prediction model

The MI-MPC prediction model follows from (5.2a), and is thus given by

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k}, \tag{5.10}$$

where $x_{i|k}$ and $u_{i|k}$ denote the predicted states and inputs, respectively, at $i \in \mathbb{N}$ time steps ahead of the prediction starting time $k \in \mathbb{N}$.

5.3.2.2 State estimator

For model-based attenuation of the measurement noise, we use a Luenberger observer given by

$$\hat{x}_k = \begin{cases} A\hat{x}_{k-1} + Bu_{k-1} + L(y_k - \hat{y}_k^-), & \text{if } \sigma_{k-1} \in \mathcal{Q}, \\ A\hat{x}_{k-1} + Bu_{k-1}, & \text{if } \sigma_{k-1} \in \mathcal{E}, \end{cases} \tag{5.11a}$$

where

$$\hat{y}_k^- = A\hat{x}_{k-1} + Bu_{k-1}, \tag{5.11b}$$

denotes the model-based output estimate at time k before performing the measurement-based correction step in (5.11a) (if available by $\sigma_{k-1} \in \mathcal{Q}$). The observer gain matrix $L_x = 0.25 I_{n_x}$ has been tuned such that the estimation error dynamics are exponentially stable, see Section 3.7, with desirable convergence properties.

5.3.2.3 Optimal control problem

To transform the temperature objectives discussed in Section 5.2.1 to the state space of (5.2), we use $z_k = Hx_k \in \mathbb{R}^{n_z}$, $n_z < n_x$, to denote the temperature elevations of the ROI voxels, where $H \in \{0,1\}^{n_z \times n_x}$ is a matrix with a single 1 in each row, and at most one 1 per column. Next, let $z_r \in \mathbb{R}^{n_z}$ and $\overline{x} \in \mathbb{R}^{n_x}$ denote the voxel-wise temperature reference and upper bound corresponding to the spatial discretizations of T_r and $\overline{T}$, respectively. The maximum violation of the upper bound is given by $\epsilon_k = \epsilon(x_k) = \|\max\{x_k - \overline{x}, 0_{n_x}\}\|_\infty \in \mathbb{R}_{\geq 0}$, where 0_{n_x} denotes a zero-vector of length n_x and the maximum operator is used element-wise.

Following the MI-MPC synthesis for SAcSSs described in Chapter 3, we obtain the optimization problem

$$\min_{\boldsymbol{\delta}_k, \boldsymbol{u}_k} \sum_{i=0}^{N} \ell(z_{i|k}, \epsilon_{i|k}), \tag{5.12a}$$

with prediction horizon $N \in \mathbb{N}$ and stage cost

$$\ell(z_{i|k}, \epsilon_{i|k}) = (z_{i|k} - z_r)^\top Q (z_{i|k} - z_r) + f_\epsilon \epsilon_{i|k}, \tag{5.12b}$$

subject to

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{5.12c}$$
$$x_{0|k} = \hat{x}_k, \tag{5.12d}$$
$$u_{0|k} = u^*_{1|k-1}, \tag{5.12e}$$
$$\delta_{0|k} = \delta^*_{1|k-1}, \tag{5.12f}$$
$$1_{N_q}^\top \delta_{i|k} = 1, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{5.12g}$$

$$x_{i|k} \leq \overline{x} + 1_{n_x}\epsilon_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N]}, \tag{5.12h}$$
$$0 \leq \epsilon_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N]}, \tag{5.12i}$$
$$0_{n_u} \leq u_{i|k} \leq \overline{u}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{5.12j}$$
$$J_u u_{i|k} \leq \overline{u_\Sigma}\delta_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{5.12k}$$
$$J_u u_{i|k} \leq \overline{u_\Sigma} S_\tau^\top \delta_{i-\tau|k}, \qquad \begin{cases} \forall\, \tau \in \mathbb{N}_{[1,\min\{k,\overline{s}\}]}, \\ \forall\, i \in \mathbb{N}_{[0,N-1]}. \end{cases} \tag{5.12l}$$

In (5.12a), the sequences $\boldsymbol{u}_k = (u_{0|k}, \ldots, u_{N-1|k})$ and $\boldsymbol{\delta}_k = (\delta_{0|k}, \ldots, \delta_{N-1|k})$ denote the predicted inputs $u_{i|k}$ and activators $\delta_{i|k}$, respectively, at $i \in \mathbb{N}$ time steps ahead of the prediction sequence's starting time $k \in \mathbb{N}$. The cost function contains a tracking term to steer the ROI states $\tilde{z}_k$ towards z_r, and a linear cost on $\epsilon_{i|k}$ to incorporate the upper temperature bound as a soft constraint. The weights are chosen as

$$Q = \frac{1}{n_z} I_{n_z}, \quad f_\epsilon = 10, \tag{5.12m}$$

which are normalized with respect to the corresponding number of decision variables (note that ϵ is scalar) for more intuitive tuning. Consequently, the weighting by f_ϵ exceeds that of Q, such that the prevention of detrimental overheating is prioritized above the reference tracking. The equality constraints (5.12c) account for the system dynamics (5.10), which start at the initial condition (5.12d) following from the observer (5.11). By (5.12e) and (5.12f), the first predicted input $u_{0|k}$ and activator $\delta_{0|k}$ are fixed to their optimal values computed at the previous time step (and for initialization $\delta_{0|0} = \delta_0$, corresponding to the transducer starting position, and $u_{0|0} = 0_{n_u}$). Consequently, we in fact apply $u_k = u^*_{0|k} = u^*_{1|k-1}$ and $\delta_k = \delta^*_{0|k} = \delta^*_{1|k-1}$, meaning that solving (5.12) at time k is effectively the one-step-ahead optimization of the inputs for time $k+1$. This allows for computational delays up to the sample time T_s. Hence, T_s will be regarded as the solver times' upper limit, which we will aim to satisfy for real-time feasibility. Next, (5.12g) includes the one-hot encoding constraint (5.7). Regarding the inequalities, (5.12h) and (5.12i) incorporate the temperature upper bound as a soft constraint with nonnegative slack variable, and (5.12j), (5.12k), and (5.12l) follow directly from the input constraints (5.5), (5.8a), and (5.8b), respectively. Note that in (5.12l), $\delta_{i-\tau|k} = \delta_{k+i-\tau}$ for $i < \tau$, which follows from the past activator values. As discussed in Chapter 3, after solving (5.12) we apply a so-called admissibility assurance procedure, of which the execution time is negligible, and which ensures that the optimal pair $(\boldsymbol{\delta}^*_k, \boldsymbol{u}^*_k)$ is physically realizable by the HIFU system.

5.3.3 Problem statement

Due to being based on a mixed-integer model with large state and input dimension, the MI-MPC (5.12) is computationally demanding. As a result, the

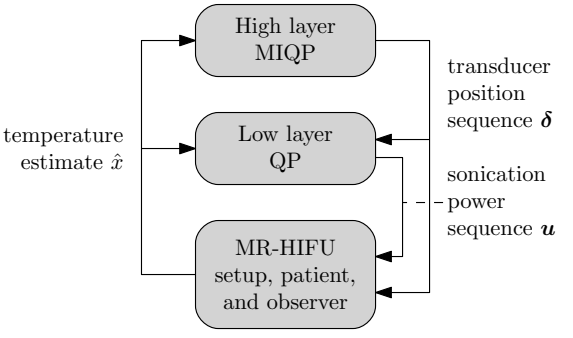

Figure 5.6: Hierarchical control architecture.

required solver times are often larger than the sample time T_s, which hampers practical implementation. To circumvent this issue, one could allow for larger computation times, for instance by increasing T_s or by re-optimizing the MPC only at discrete times $k = al$ with $l \in \mathbb{N}$ and $a \in \mathbb{N}_{\geq 1}$. Due to the increased delay, however, such approaches are prone to reduced control performance. As an alternative that does not suffer from these drawbacks, we propose a hierarchical control approach in this chapter. Hence, the key objective of this work can be formulated as: *Enable the real-time feasibility of the control principles of the MI-MPC* (5.12) *using a hierarchical MI-MPC scheme with substantially reduced computational complexity, and thereby solver times, but without significantly sacrificing performance.*

5.4 Hierarchical mixed-integer MPC

This section presents the hierarchically decomposed MI-MPC, consisting of a low-complexity high-level MI-MPC coupled to a high-accuracy low-level MPC.

5.4.1 Hierarchical control architecture

The hierarchical control structure is schematically depicted in Figure 5.6. The full-scale MI-MPC (5.12) is decomposed into two control layers. In the higher layer, a MIQP is solved to find the optimal mechanical transducer path. Since this requires only sufficient predictive accuracy to determine which region of the tumor should be heated at which time, it is expected that a reduced-complexity thermal model can be used, and as a consequence the MIQP can be solved in considerably less time than (5.12). Subsequently, we feed the discrete transducer path to the lower layer, thus reducing the low-level optimization problem to a QP, which is computationally lightweight also when using a model with large state and input dimensions. Hence, in the low layer we use the unreduced prediction

model (5.10) to optimize the sonication powers, without inducing large computation time. As a result, the hierarchical controller yields high-performance sonication and transducer inputs while requiring significantly reduced overall computational efforts compared to (5.12).

5.4.2 High control layer

5.4.2.1 Input reduction

As a first method for alleviating the computational burden of the high-level controller, we decrease the number of MPC decision variables by reducing the input dimension. In principle, the HMI-MPC setup proposed in this work is compatible with any input reduction method. In this work, however, we use the sonication point selection steps of the target-conformal optimal actuator placement procedure proposed in Chapter 4. This allows for identifying the most important sonication points, based on the specific tumor geometry and tissue properties. As a result, we can approximate the HIFU setup's heating capability with sufficient accuracy to determine a (near-)optimal transducer path in the higher control layer, while using substantially fewer sonication points than in the unreduced MI-MPC (5.12) for computational efficiency.

To this end, we first select a subset of sonication points critical for steady-state performance $\mathcal{P}_s \subseteq \mathcal{P}_c$, from the set of n_u "candidate" sonication points $\mathcal{P}_c$ used in the unreduced MI-MPC (green, black, and gray dots in Figure 5.4). In the high-level model, the points $\mathcal{P}_s$ characterize the steady-state heating of the full-scale MI-MPC (5.12), which is key for achieving high treatment quality. As described in Chapter 4, this can be done by computing the optimal steady state of (5.2a) when disregarding the mode-switching constraints, i.e., when allowing all elements of u_k to be nonzero, using

$$(x_{ss}^*, u_{ss}^*) = \arg\min_{x_{ss}, u_{ss}} \ell(Hx_{ss}, \epsilon(x_{ss})), \tag{5.13a}$$

with the stage cost from (5.12b), and subject to

$$x_{ss} = Ax_{ss} + Bu_{ss}, \tag{5.13b}$$
$$0_{n_u} \leq u_{ss} \leq \overline{u}, \tag{5.13c}$$
$$1_{n_u}^\top u_{ss} \leq \overline{u_\Sigma}, \tag{5.13d}$$

which can be written as a simple QP. Then, based on u_{ss}^*, we define $\mathcal{P}_s$ as the sonication points with the user-defined $N_{ps} \in \mathbb{N}_{[0,n_u]}$ largest steady-state powers. Typically, these are located near the edge of $\mathcal{R}$ and in the vicinity of blood vessels, counteracting the diffusive and perfusive heat losses, respectively. Second, we select a set of $N_{pu} \in \mathbb{N}_{[0,n_u-N_{ps}]}$ additional sonication points $\mathcal{P}_u \subseteq \mathcal{P}_c \setminus \mathcal{P}_s$ in such a manner that the combined set $\mathcal{P}_s \cup \mathcal{P}_u$ is distributed throughout the ROI as uniformly as possible. In the higher control layer, the points $\mathcal{P}_u$ are needed to

better approximate the unreduced MI-MPC's heating behavior during heat-up and when counteracting heat losses in areas not in the vicinity of $\mathcal{P}_s$. The set $\mathcal{P}_u$ can be found by solving, for a user-defined number N_{pu}, an efficient mixed-integer linear program that minimizes the mean distance from each ROI voxel center to its closest sonication point in $\mathcal{P}_s \cup \mathcal{P}_u$, see Chapter 4 for details. With respect to the MI-MPC (5.12), the resulting high-level input reads

$$\tilde{u}_k = G u_k \in \mathbb{R}^{n_{\tilde{u}}} \tag{5.14}$$

where the matrix $G \in \{0,1\}^{n_{\tilde{u}} \times n_u}$ selects the $n_{\tilde{u}} = N_{ps} + N_{pu} \leq n_u$ elements from the unreduced input u_k in (5.2), corresponding to selecting the points $\mathcal{P}_s \cup \mathcal{P}_u$ from $\mathcal{P}_c$.

5.4.2.2 State reduction

The second manner in which we decrease the solver time of the high-level controller is by reducing the number of state decision variables in the optimization problem, which is achieved by using a reduced-order model. Although such an approach typically induces prediction errors due to using lower-dimensional state approximations, this effect may be small as a result of the hierarchical control structure (with high-accuracy low-level controller) if the reduction is applied properly. As also recognized in [117], model order reduction by proper orthogonal decomposition (POD) [114] is particularly well suited for MR-HIFU treatments, as it is a data-based method achieving accurate lower-order approximations when the realized output trajectories resemble the (measured or model-generated) data used for generating the POD. This is the case for high-quality hyperthermia treatments, in which typically an initial heat-up phase with gradually rising tumor temperatures is followed by a steady-state phase where the temperature is close to optimal. As an additional advantage, due to being data-based, POD allows for tumor-specific reduced-order modeling when a patient has to be treated multiple times, as is often true in reality, by using the data recorded during previous treatments. In fact, as argued in [117], the model could even be updated during treatment after sufficient measurements have been collected, which may be particularly useful in case the tissue properties have changed over time, for instance due to the body's thermoregulatory response.

Correspondingly, we write the high-level model as

$$\tilde{x}_{k+1} = \tilde{A}\tilde{x}_k + \tilde{B}\tilde{u}_k, \tag{5.15}$$

with high-level state $\tilde{x}_k = \tilde{U}^\top x_k \in \mathbb{R}^{n_{\tilde{x}}}$ and matrices $\tilde{A} = \tilde{U}^\top A \tilde{U}$ and $\tilde{B} = \tilde{U}^\top B G^\top$ (with G from (5.14)). In case of POD-based model order reduction, the high-level state transformation matrix is given by $\tilde{U} = U_{n_{\tilde{x}}}$, where $U_{n_{\tilde{x}}}$ denotes the first $n_{\tilde{x}} < n_x$ columns of U, which follows from the singular value

decomposition $X = U\Sigma V^\top$ of the N_X temperature snapshots collected in

$$X = \begin{bmatrix} x_0 & \dots & x_{N_X} \end{bmatrix} \in \mathbb{R}^{n_x \times N_X}. \tag{5.16}$$

The snapshots can be obtained from previous treatments or generated by simulation of the full-scale model (5.2). If, on the other hand, no model order reduction is applied, the high-level transformation matrix is given simply by $\tilde{U} = I_{n_x}$, such that $\tilde{x}_k = x_k$ and $n_{\tilde{x}} = n_x$.

5.4.2.3 Optimal control problem

Since the control objectives are defined in the state space of (5.2), see Section 5.2.1 and Section 5.3.2.3, we write the high-level optimization problem in terms of the potentially lower-dimensional approximation of the unreduced state $\tilde{U}\tilde{x}_k$, with correspondingly $\tilde{z}_k = H\tilde{U}\tilde{x}_k$ and $\tilde{\epsilon}_k = \epsilon(\tilde{U}\tilde{x}_k)$, yielding

$$\min_{\boldsymbol{\delta}_k, \tilde{\boldsymbol{u}}_k} \sum_{i=0}^{N} (\tilde{z}_{i|k} - z_r)^\top Q (\tilde{z}_{i|k} - z_r) + f_\epsilon \tilde{\epsilon}_{i|k}, \tag{5.17a}$$

subject to

$$\tilde{x}_{i+1|k} = \tilde{A}\tilde{x}_{i|k} + \tilde{B}\tilde{u}_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{5.17b}$$

$$\tilde{x}_{0|k} = \tilde{U}\hat{x}_k, \tag{5.17c}$$

$$0_{n_u} \leq \tilde{u}_{i|k} \leq \alpha_{\tilde{u}} \overline{u}, \qquad \forall\, i \in \mathbb{N}_{[0,N-1]}, \tag{5.17d}$$

$$\tilde{U}^\top \tilde{x}_{i|k} \leq \overline{x} + 1_{n_x} \tilde{\epsilon}_{i|k}, \qquad \forall\, i \in \mathbb{N}_{[0,N]}, \tag{5.17e}$$

and the constraints (5.12i) on $\tilde{\boldsymbol{\epsilon}}_k$, and (5.12e)-(5.12g) and (5.12k)-(5.12l) on $\boldsymbol{\delta}_k$ and $\tilde{\boldsymbol{u}}_k$. Here, (5.17b)-(5.17e) represent the high-level system's dynamics (5.15), the initial condition, the amplified individual sonication point power range with $\alpha_{\tilde{u}} \geq 1$ to compensate for the potentially smaller number of sonication points, and the temperature upper bound (where the reduced-order states are projected back onto the full-order state space), respectively. Similar to the MI-MPC (5.12), we first solve (5.17), and subsequently apply the admissibility assurance described in Chapter 3 to ensure that the optimal pair $(\boldsymbol{\delta}_k^*, \tilde{\boldsymbol{u}}_k^*)$ is physically realizable.

The high-level MPC (5.17) can generally be formulated using a sparse or a condensed approach [67,115]. In the sparse formulation, the predicted states are included as optimization variables, and the system dynamics (5.17b) are incorporated as equality constraints, whereas using the condensed approach, the state decision variables are eliminated by expressing the state predictions in terms of the initial prediction state $\tilde{x}_{0|k}$ and the input sequence $\tilde{\boldsymbol{u}}_k$. Although the sparse method leads to more decision variables, it is generally preferable regarding computational efficiency when using large prediction horizons N, due to the

structure and sparsity of the matrices defining the optimization problem. In this case, using a POD-reduced model may speed up computations by significantly reducing the state dimension. Contrarily, for smaller N and in case of small input dimension relative to the state dimension $n_{\tilde{u}} \ll n_{\tilde{x}}$, the latter of which can be achieved by input reduction, eliminating the state decision variables may outweigh the loss of sparsity in the cost and constraint matrices, resulting in the condensed formulation to yield smaller solver times. Although condensing the MPC does not preclude using a reduced-order model, combining these is not sensible, since in a condensed formulation the computational complexity does not depend on the (reduced) state dimension due to all state decision variables being eliminated. In fact, if for sparse unreduced models the state-space matrices of their reduced-order approximations are full, as is often the case, using the reduced models in a condensed formulation typically results in a denser problem, and thereby larger computation times, than when using the unreduced models. Since, in addition, reduced-order models usually lead to approximation errors, a condensed POD-based MPC is typically inferior in terms of accuracy as well as computational efficiency. Therefore, in this work we consider only the condensed form of (5.17) when using $\tilde{U} = I_{n_x}$, i.e., without model order reduction.

5.4.3 Low control layer

After solving (5.17) to find $\boldsymbol{\delta}_k^*$, representing the optimal transducer path, the optimal sonication powers are determined in the lower control layer. To enable maximum control performance and robustness, this is done using the high-accuracy prediction model (5.10), which also includes all sonication points $\mathcal{P}_c$ within electronic beam steering range, thus offering the same sonication flexibility as the unreduced MI-MPC (5.12). The corresponding low-level controller is given by the QP that results from fixing $\boldsymbol{\delta}_k = \boldsymbol{\delta}_k^*$ in (5.12), which by solving yields the optimal input sequence $\boldsymbol{u}_k^*$, such that the pair $(\boldsymbol{\delta}_k^*, \boldsymbol{u}_k^*)$ constitutes the optimized control actions for the MR-HIFU system.

5.5 Numerical simulation study

This section presents simulation results obtained with various reduced-complexity high-level models to illustrate the achievable control performance and computation time of the HMI-MPC in comparison to the unreduced MI-MPC.

5.5.1 High-level model complexity reduction

We first follow the input reduction method proposed in Chapter 4, as discussed in Section 5.4.2.1. To determine the steady-state-critical sonication points $\mathcal{P}_s$, we solve (5.13) (with (A, B) corresponding to the phantom for the upcoming

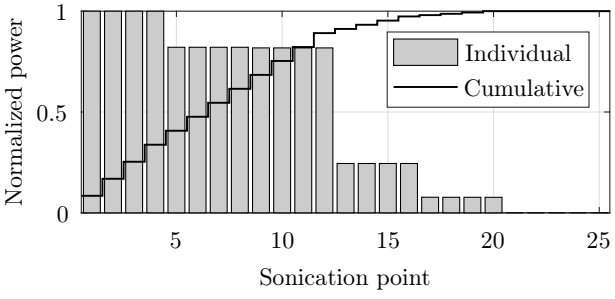

Figure 5.7: The 25 largest sonication powers in u^*_{ss} in descending order (normalized with respect to $\|u^*_{ss}\|_\infty$) and their cumulative sum (normalized with respect to $\|u^*_{ss}\|_1$).

experiments in Section 5.6) to find u^*_{ss}, of which the 25 largest sonication powers are shown in Figure 5.7 in descending order and normalized with respect to the largest power $\|u^*_{ss}\|_\infty$. Based on this, we choose $\mathcal{P}_s$ to consist of the $N_{ps} = 12$ sonication points that are heated most by u^*_{ss}, and together apply 89% of the total steady-state input power $\|u^*_{ss}\|_1$. These points are also indicated in Figure 5.4 (green dots), where they can also be visually verified to correspond to the key sonication points for the optimal steady-state heating Bu^*_{ss} (color map). Next, we must determine the set of N_{pu} additional uniformly distributed sonication points $\mathcal{P}_u$. To evaluate the effect on performance and solver time when including a different number of sonication points, simulations are performed for various $\mathcal{P}_u$. These are indicated in Figure 5.4, and correspond to N_{pu} equal to 2 (red ∘), 12 (blue ∘), 140 (black dots), and 280 (black and gray dots). Note that since $n_u = 292$ and $N_{ps} = 12$, the case with $N_{pu} = 280$ in fact yields $\mathcal{P}_s \cup \mathcal{P}_u = \mathcal{P}_c$, which means no input reduction is applied.

Second, to apply the POD-based high-level model order reduction, we generate the snapshot matrix X (5.16) by running seven simulations of the unreduced model (5.2) in closed-loop feedback with the MI-MPC (5.12), with each simulation starting from a different treatment cell for a richer data set, especially regarding the heat-up phase. The resulting singular values are shown in Figure 5.8, suggesting that the first 10 modes contain the most dominant behavior. Correspondingly, to exemplify the achievable results using POD-based reduced-order models, and to show the effects of using too few modes, simulations are performed using $n_{\tilde{x}} \in \{5, 10, 15\}$.

As mentioned previously, when using a reduced-order high-level model, the MPC is formulated using a sparse approach. When using the full-order model by $\tilde{U} = I_{n_x}$, on the other hand, we investigate the computation times obtained using both a sparse and condensed MPC approach. For convenience, all high-level controller settings considered in this case study are summarized in Table 5.1.

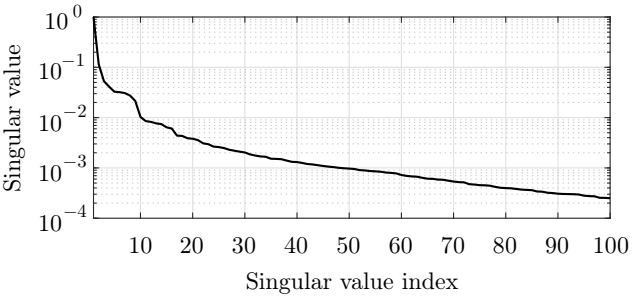

Figure 5.8: The 100 largest singular values of X (normalized).

Note that by $N_{pu} = 280$ in input setting D and $\tilde{U} = I_{n_x}$ in state settings 1 and 2, the high-level controllers of setups D{1,2} in fact represent the full-scale MI-MPC (5.12) in sparse and condensed form, respectively.

5.5.2 Control performance and computation time

Simulations have been performed using Matlab R2017b and Gurobi 8.1.1 on a laptop with Intel Core i7 @ 2.60 GHz CPU and 8 GB RAM. The resulting computation times, for a prediction horizon $N = 8$, are presented in Figure 5.9. Clearly, the reduced-complexity high-level controllers achieve much smaller computation times than the unreduced setups D{1,2}. In particular, we observe that for small input dimensions the condensed formulation enables the fastest computations on average, indicating that in this case eliminating the state decision variables outweighs the loss of sparsity in the optimization problem's small input component. This is visualized in Figure 5.10a-5.10c, depicting the nonzero elements of the Hessian matrices of the different HMI-MPCs' high-level stage costs (note the different axis scaling). For increasing $n_{\tilde{u}}$, however, the input component of the optimal control problem grows, resulting in the significant filling of the Hessian matrices, see Figure 5.10d-5.10f, which causes the complexity of the dense controllers to grow more rapidly than that of the sparse controllers.

Table 5.1: HMI-MPC high-level reduction parameters.

Input setting	A	B	C	D
N_{ps}	12	12	12	12
N_{pu}	2	12	140	280

State setting	1	2	3	4	5
MPC formulation	Sparse	Dense	Sparse	Sparse	Sparse
Transformation matrix $\tilde{U}$	I_{n_x}	I_{n_x}	U_5	U_{10}	U_{15}

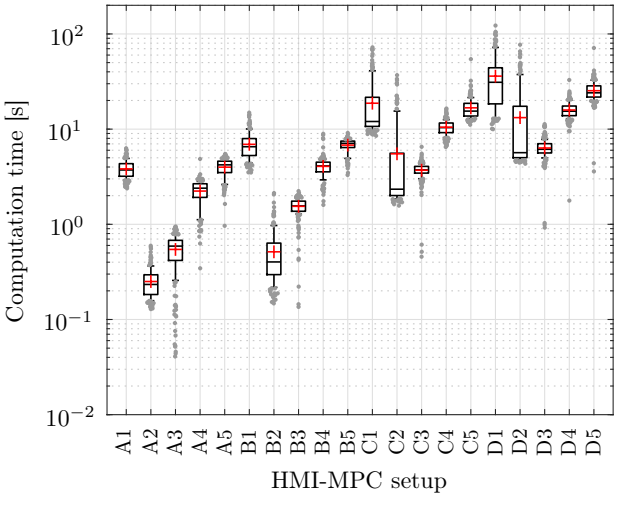

Figure 5.9: High-level computation time for the HMI-MPCs (see Table 5.1 for parameters), indicating the mean (red +), median (central mark), 25th and 75th percentiles (box edges), 10th and 90th percentiles (whiskers), and outliers (gray dots).

Consequently, the densely formulated HMI-MPCs (state setting 2) are mostly outperformed in terms of maximum computation time by the sparse-form POD-reduced HMI-MPCs (state settings 3 to 5) for $N_{pu} \geq 140$ (input settings C and D).

For the lower control layer, which is equal for all considered HMI-MPC setups, the sparse formulation was found to be less computationally demanding than the condensed form. When including the resulting solver times, with mean, 90th percentile, and maximum at 0.38 s, 0.58 s, and 1.07 s, respectively, real-time feasibility (here defined as the maximum computation time of the HMI-MPC being smaller than the sample time $T_s = 3.2$) is achieved using controllers {A,B}{2,3}.

To evaluate the control performance, we compute the M-sample centered moving average of the stage cost (5.12b) $\mathrm{MA}_M(\ell(z_k, \epsilon_k))$, which for odd integers $M = 1 + 2m$, $m \in \mathbb{N}$, is defined as

$$\mathrm{MA}_M(\ell(z_k, \epsilon_k)) = \frac{1}{M} \sum_{l=\max\{k-\frac{M}{2},0\}}^{k+\frac{M}{2}} \ell(z_l, \epsilon_l), \tag{5.18}$$

effectively smoothing the stage cost over an interval of $(M-1)T_s$ to reduce the oscillations caused by relocating the transducer, thereby improving readability. In Figure 5.11a, using $M = 11$, we indicate the range in which the smoothed

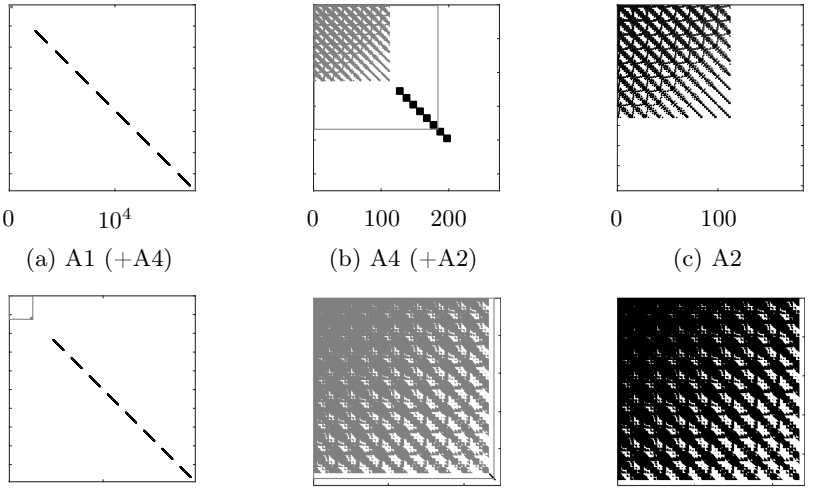

(a) A1 (+A4) (b) A4 (+A2) (c) A2

(d) D1 (+D4) (e) D4 (+D2) (f) D2

Figure 5.10: High-level cost function Hessian matrix sparsity pattern for setting (a) A1 (black) and A4 with matrix outline (gray, top left), (b) A4 (black) and A2 with matrix outline (gray, top left), (c) A2, (d) D1 (black) and D4 with matrix outline (gray, top left), (e) D4 (black) and D2 with matrix outline (gray, top left), (f) D2.

stage costs lie for all the high-level input settings when using state setting 3 (red) and when using settings {1,2,4,5} (gray). In addition, the results obtained with A{1,2} are highlighted individually (black). Clearly, using state setting 3, the HMI-MPC does not achieve satisfactory performance, from which we conclude that using only the first five modes of the considered POD by $\tilde{U} = U_5$ does not provide the higher control layer with sufficient model information to properly optimize the mechanical transducer path. As a consequence, the transducer is repositioned too infrequently, and the ROI cannot be adequately heated. For all other high-level controller settings, on the other hand, this figure verifies that in this case study the corresponding reduced-complexity HMI-MPC controllers achieve similar performance as the full-scale MI-MPC (5.12). For example, the HMI-MPC with high-level reduction setting A2 achieves similar tumor temperatures as the unreduced MI-MPC, see Figure 5.11b, while exhibiting substantially smaller computation times. In particular, compared to the fastest full-scale MI-MPC implementation (the dense-form high-level controller in D2, see Figure 5.9), the HMI-MPC A2 reduces the mean solver time from 36.05 s to 0.76 s, and the maximum from 122.54 s to 1.27 s, which corresponds to a reduction of 98% and 99%, respectively.

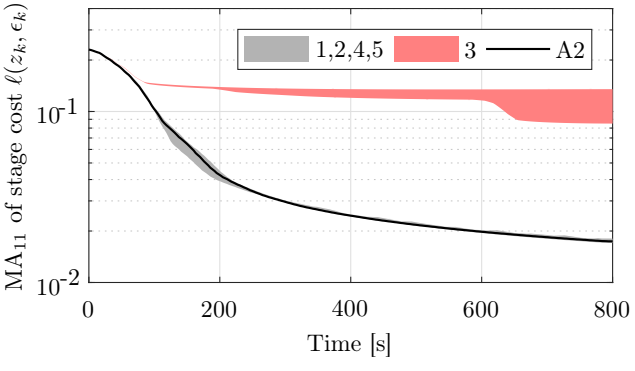

(a)

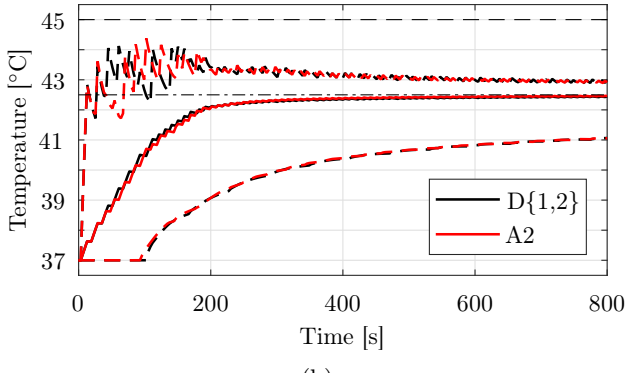

(b)

Figure 5.11: In (a), the range of the 11-sample moving average stage costs $\mathrm{MA}_{11}(\ell(z_k, \epsilon_k))$ versus time t_k for settings {A,B,C,D}3 (red), {A,B,C,D}{1,2,4,5} (gray), and A2 individually (black). In (b), the mean (solid) and extremum (dashed) tumor temperatures for the HMI-MPC A2 and the unreduced MI-MPC D{1,2}, and the temperature reference (thin dash-dotted) and upper bound (thin dashed).

5

5.5.3 Discussion

These results strongly motivate the proposed HMI-MPC (5.17) for substantially decreasing the computation time with respect to the unreduced MI-MPC (5.12), without sacrificing significant performance. In the considered case study, using the input selection method of Chapter 4, the HIFU system's full heating capability using $n_u = 292$ sonication points can be approximated using only $N_{ps} + N_{pu} = 12 + 12 = 24$ (A2) points with sufficient accuracy to determine a (near-)optimal mechanical transducer path. As a result, using a dense MPC formulation for the high-level controller achieves the fastest computations, and,

in addition, is considered preferable over using a POD-reduced model due to the absence of state approximation errors (although due to the hierarchical architecture this effect is shown to be small for properly selected POD-reduced models). If, however, such a considerable input reduction would not be possible, for instance due to the thermal system having a very low diffusion coefficient, model order reduction by POD in the high control layer provides an alternative for reducing the solver time with respect to the unreduced MI-MPC, while achieving similar control performance. In this case, however, one must be sure to include enough POD modes.

5.6 Phantom experiments

This section presents the results obtained from phantom experiments using a clinical MR-HIFU setup to validate its control performance and real-time feasibility.

5.6.1 Setup

In these experiments, we aim to heat a polyacrylamide tissue-mimicking phantom as in [90], but with the ink replaced by water, using our HMI-MPC (5.17)To allow for a comparison with the results reported in [53] obtained using a state-of-the-art binary controller, we use the same ROI and the treatment cells, which are depicted in Figure 5.4. However, as we are heating a phantom, of which the baseline corresponds to the ambient temperature, we use as temperature objectives the values described in Section 5.2.1 relative to 37 °C, i.e., we use as ROI temperature reference an elevation of 5.5 °C, and the upper bound equals 8 °C elevation inside $\mathcal{S}$ and 3 °C elevation outside $\mathcal{S}$. For the HMI-MPC, we use setting A2 from Table 5.1 for the higher control layer, and a prediction horizon $N = 10$. In the binary controller of [53], each treatment cell contains 80 sonication points (distributed over four concentric circles with diameters 4, 8, 12, and 16 mm, where each circle contains 8, 16, 24, and 32 sonication points, respectively), which is approximately twice as many as in the lower layer of our HMI-MPC, where each cell contains only 40 or 44 points, see Figure 5.4. Thus, regarding the number of control inputs, the binary controller offers more controller flexibility, meaning that if the HMI-MPC outperforms the binary controller, it is truly due to possessing a superior control *strategy*.

5.6.2 Results

5.6.2.1 Performance and computation time

The distribution of the HMI-MPC's high-level, low-level, and overall computation times are shown in Figure 5.12, verifying that the proposed setup achieves real-time feasibility on the clinical setup. In Figure 5.13, we show the mean

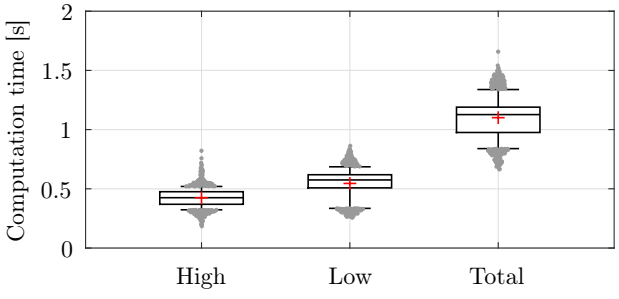

Figure 5.12: High-level, low-level, and overall computation time of the HMI-MPC during the phantom experiments, indicating the mean (red +), median (central mark), 25th and 75th percentiles (box edges), 10th and 90th percentiles (whiskers), and outliers (gray dots).

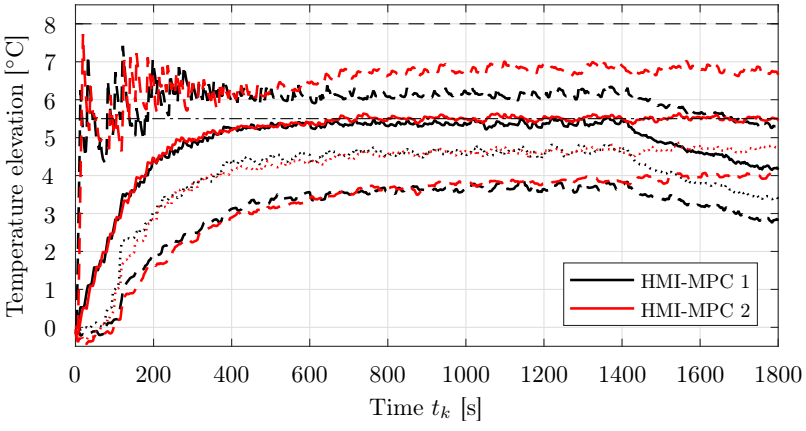

Figure 5.13: The mean (solid), 10th percentile $T90$ (dotted), and extremum (dashed) tumor temperatures for two HMI-MPC experiments (red and black)and the temperature reference (thin dash-dotted) and upper bound (thin dashed).

(solid), 10th percentile (dash-dotted), which is commonly denoted as $T90$, and extremum (dashed) ROI temperatures for two experiments using the HMI-MPC with heating until $t_k = 1400$ (black) and $t_k = 1800$ (red). These results indicate that using the HMI-MPC, the large ROI can in general be adequately heated, reaching satisfactory a steady-state temperature. For closer inspection regarding the steady state, in Figure 5.14 we show the temporal averages of the mean value (red), 10th percentile $T90$ and 90th percentile $T10$ (box edges), and extrema (whiskers) of the measured ROI temperature distribution's deviation from its setpoint (i.e., Hy_k with respect to z_r) over the interval from $t_k \geq 600$ until the sonication is stopped. Additionally, this figure shows the mean, $T90$, and

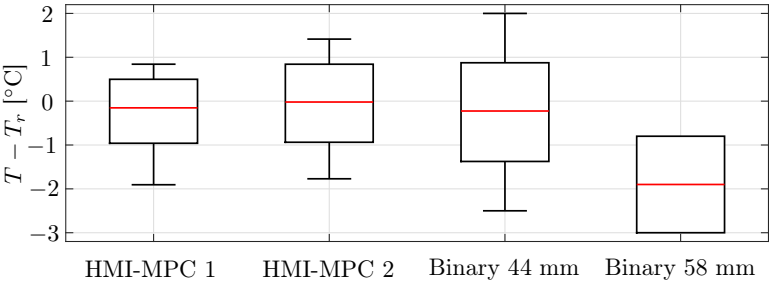

Figure 5.14: ROI temperature mean (red), 10th percentile $T90$ and 90th percentile $T10$ (box edges), and extrema (whiskers) with respect to the ROI setpoint T_r in steady state for HMI-MPC experiments 1 and 2, compared to results reported in [53] using a state-of-the-art binary controller and ROI of 44 mm and 58 mm diameter.

$T10$ as reported in [53] using a binary state-of-the-art clinical controller for the same 58 mm diameter ROI and for a smaller 44 mm diameter ROI[1], although it must be remarked that [53] concerns an in-vivo experiment. Clearly, the temperature distributions achieved in this work using HMI-MPC 1 (1.46 °C) and HMI-MPC 2 (1.78 °C) are significantly more uniform than those obtained in [53] using the binary controller for the 44 mm and 58 mm diameter ROIs (2.25 °C and 2.20 °C, respectively). Moreover, the mean temperature deviations from the desired setpoint for HMI-MPC 1 (−0.15 °C) and HMI-MPC 2 (−0.02 °C) are also smaller than for the Binary 44 mm (−0.23 °C) and 58 mm (−1.9 °C). These results already demonstrate the enhanced control performance of the proposed MPC design. However, as discussed in more detail in Section 5.6.2.2 below, the HMI-MPC's temperature control performance is actually considered to be even better, as the results presented in Figure 5.13 and Figure 5.14 are corrupted by thermometry artifacts, causing the measured ROI temperature distribution to seem less homogeneous than the actual ROI temperature in the phantom.

Next, we analyze the spatial variation of the temperature distribution. We find that the coldest 10 percent of the ROI lies at its periphery, as shown in the temperature contour plots in Figure 5.15, corresponding to the areas outside the treatment cells (dashed green), meaning they cannot be heated directly. This could be improved by redefining the transducer positions and treatment cells, or including more, such that they better cover the areas critical for steady-state

[1]The only 58 mm diameter ROI experiment presented in [53] (see [53, Data set 8]) was mentioned not to be representative, since excessive unforeseen near-field heating at a bone-tissue interface prevented further heating. Therefore, we also show results for a 44 mm diameter ROI computed from the average of [53, Data sets 1 to 4]) for comparison with our HMI-MPC. The temperature mean, and 10th and 90th percentiles are obtained from [53, Table II]. For the 44 mm ROI, the extrema are estimated from [53, Figure 3a]. For the 58 mm ROI, extrema indications were not presented.

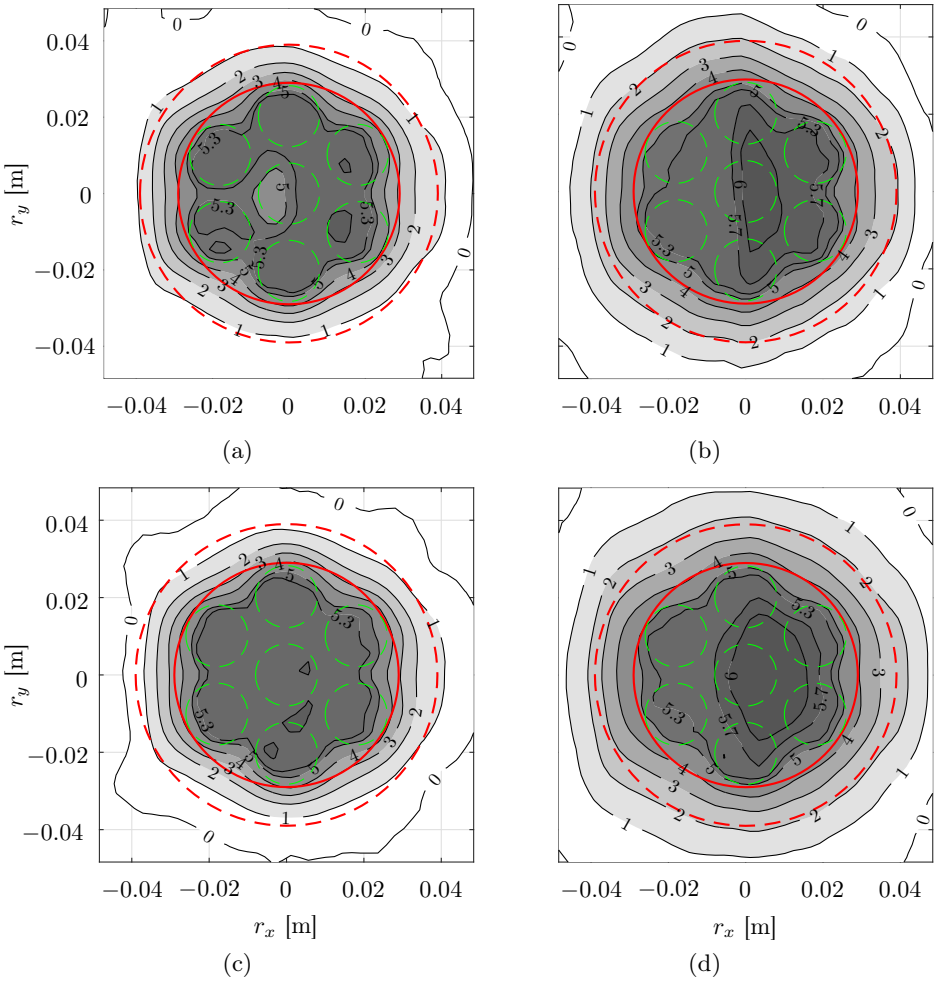

Figure 5.15: Contour plots of the focal plane temperature elevation for HMI-MPC experiment 1 averaged over (a) $320 \leq t_k \leq 400$ and (b) $1300 \leq t_k \leq 1400$, and for experiment 2 over (c) $320 \leq t_k \leq 400$ and (d) $1700 \leq t_k \leq 1800$, and the outline of $\mathcal{R}$ (solid red), $\mathcal{S}$ (dashed red), and the treatment cells (dashed green).

heating, which could be done via the optimal transducer placement method from Chapter 4. Figure 5.15c shows that for $320 \leq t_k \leq 400$, most of the ROI is between 5.3 °C and 5.7 °C elevation, i.e., within 0.2 °C from the reference. To relate it to (the incremental versions of) the treatment objectives visualized in Figure 5.1, consider the temperature cross sections at $r_y = 0$ in Figure 5.16a and

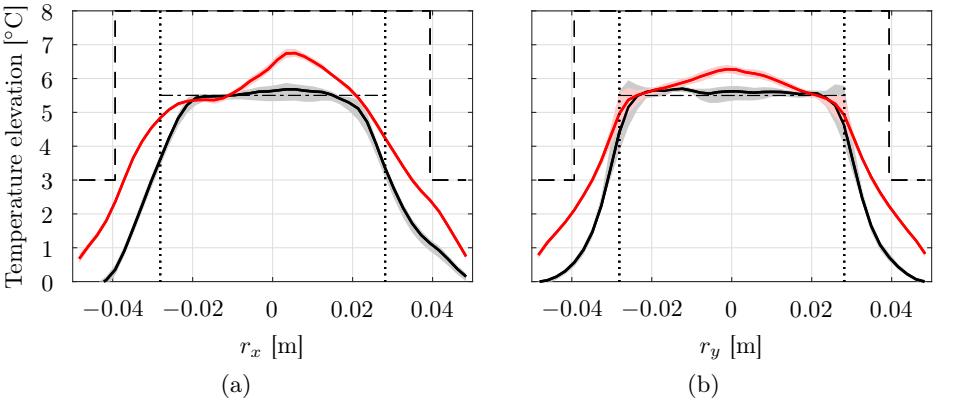

Figure 5.16: The mean (line) and extremum (area) tumor temperature estimates of HMI-MPC experiment 2 during $320 \leq t_k \leq 400$ (black) and $1700 \leq t_k \leq 1800$ (red) of the cross section at (a) $r_y = 0$ and (b) $r_x = 0$, and the temperature reference (dash-dotted), upper bound (dashed), and ROI edge (dotted).

$r_x = 0$ in Figure 5.16b for experiment 2 during $320 \leq t_k \leq 400$ (black), clearly depicting the controller's ability to achieve a uniform temperature inside $\mathcal{R}$ near the reference of 5.5 °C (dash-dotted). In Figure 5.16b, the temperature fluctuations (gray area) can be seen to be largest at the ROI edge, which is the result of constant diffusive heat loss being periodically counteracted by direct sonication, which is possible due to these areas being within electronic beam steering range, see the treatment cells in Figure 5.15 for $r_x = 0$, and which enables tighter temperature control. By contrast, this is not the case in Figure 5.16a due to the periphery of $\mathcal{R}$ at $r_y = 0$ being outside the treatment cells, as mentioned earlier, see Figure 5.15. Upon closer inspection of the applied heating during $t_k \geq 400$, which for experiment 2 is visualized in Figure 5.17 (and which is also representative for experiment 1), we find that indeed the only significant heating occurs at the ROI edge to counteract diffusive heat loss, as one would expect for a non-perfused phantom. Moreover, note the similarity of the averaged steady-state heating in Figure 5.17 compared to its model-based counterpart in Figure 5.4 used in finding the steady-state-critical sonication points $\mathcal{P}_s$ in the high control layer. This verifies that $\mathcal{P}_c$ is adequately selected, indeed enabling the high-level controller (5.17) to approximate the required steady-state heating with sufficient accuracy for determining the optimal mechanical transducer path.

5.6.2.2 Thermometry artifacts

Unfortunately, the conducted experiments suffered from measurement artifacts, resulting in the false registration of an additional ROI temperature increase for

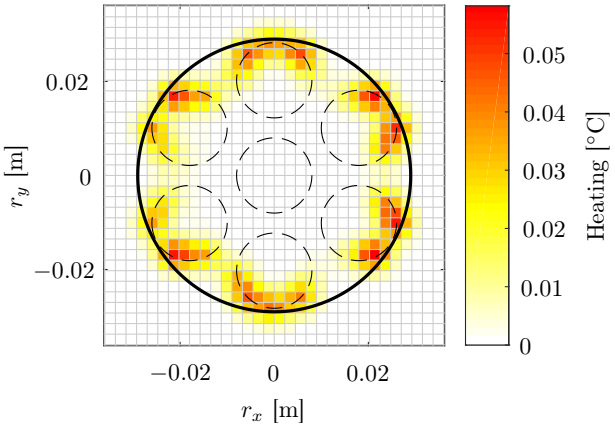

Figure 5.17: The heating Bu_k averaged over $400 \leq t_k \leq 1800$ (color map), the ROI $\mathcal{R}$ (solid black), and the treatment cells (dashed black).

approximately $t_k \geq 400$ s, expressed by a larger maximum ROI temperature in Figure 5.13 than truly in the phantom. The artifacts are centered at approximately $(r_x, r_y) = (0.005, 0)$, and are shorter in the r_x-direction than along the r_y-axis, see Figure 5.15b for HMI-MPC experiment 1, and Figure 5.15d and Figure 5.16 (red) for experiment 2. Consider again the applied heating for $t_k \geq 400$, an example of which is shown in Figure 5.17. Clearly, the temperature evolutions from Figure 5.15a to Figure 5.15b, and from Figure 5.15c to Figure 5.15d (or equivalently from the black to the red profiles in Figure 5.16), do not correspond to the applied heating, and are therefore concluded to be caused by thermometry artifacts. This is further supported by the strongly heterogeneous (especially along the r_x-axis) temperature maps that remain after the phantom has fully cooled down, see Figure 5.18 for an example, which in the absence of artifacts should be uniformly equal to zero. Consequently, we expect the truly achievable temperature homogeneity of the ROI to be even better than Figure 5.13 suggests.

The artifacts were most likely caused by the occurrence of B_0 field drift, i.e., changes over time of the polarizing magnetic field in r_y-direction along the MRI scanner bore (see Figure 5.3), which could not be adequately compensated for by the second-order drift correction algorithm due to significant heterogeneous phase drift on a relatively small spatial scale. That is, the phantom is relatively small (17 cm diameter), and its magnetic susceptibility is considerably different from air. This results in strong localized field inhomogeneities in and around the phantom, which in case of B_0 drift may lead to inhomogeneous phase drift in the phantom, especially in the r_x-direction perpendicular to the field, thereby leading to thermometry artifacts [40,44]. In addition, due to the heated volume being large (8 to 10 cm diameter, see Figure 5.15 and Figure 5.16) compared

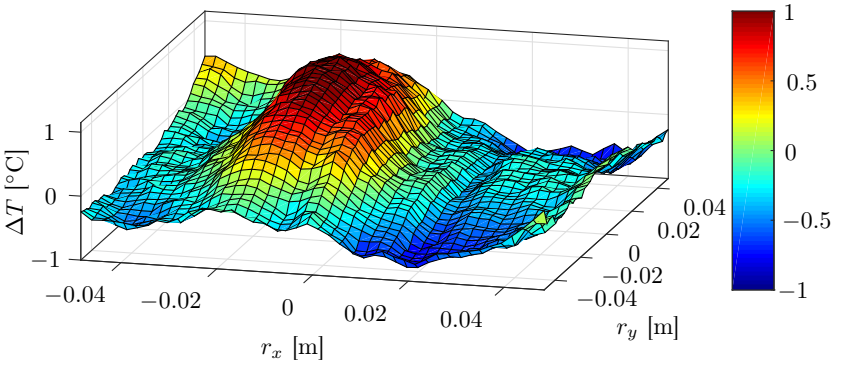

Figure 5.18: Artifact-corrupted temperature elevation ΔT in the focal plane after cooling down.

to the phantom, only a relatively small unheated phantom volume remains as background reference for the drift estimation [118]. Finally, using the phantom test rig, the phantom is not centered in the MRI scanner bore in r_x-direction, which could explain the r_x-asymmetry of the measurement artifacts.

To reduce the occurrence of the observed thermometry artifacts, further development of the drift correction algorithms could be pursued [118–120]. Another option, relating more directly to actual patient treatment, would be to perform the experiment using a larger (possibly anthropomorphic) phantom and a redesigned test rig that allows for centering the phantom in the MRI scanner bore. The increased object volume, compared to the phantom considered in this work, typically leads to a more uniform magnetic susceptibility distribution, and results in more unheated tissue being available for drift estimation, both of which are beneficial for the accuracy of the MR thermometry and the drift correction. Alternatively, an offset-free MPC scheme [78, 79] could be employed, which essentially estimates online the constant or slowly varying mismatch between the measurements and model. For MR-HIFU hyperthermia using a mechanically stationary transducer, this has already been shown to be able to identify significant plant-model mismatch, see Chapter 2. However, two important remarks are in order. First, particular attention must be paid in selecting the disturbance model, as this determines the manner in which the offset-free algorithm captures the input or output disturbances. Second, when switching the transducer position, the steady-state temperature and heating are (close to) periodic, instead of constant, which requires appropriate modifications to the conventional offset-free MPC architectures.

5.7 Conclusion

This chapter presented a hierarchical MI-MPC approach for improved temperature control in large-volume MR-HIFU hyperthermia. First, the integer part of the problem, representing the range-extending mechanical transducer switching, is determined by swiftly solving a reduced-complexity MIQP in the higher control layer. Here, the complexity reduction is achieved by (target-conformal) input reduction and/or (POD-based) model order reduction, and by wisely choosing a sparse or condensed MPC approach depending on the specific reduction settings. Subsequently, in the lower control layer, the sonication plan is determined using an unreduced high-fidelity model for optimal control performance. In numerical simulation, the proposed HMI-MPC architecture is shown to enable tremendous computational benefit without sacrificing significant performance. Using phantom experiments on a clinical MR-HIFU setup, the HMI-MPC's performance and real-time feasibility are verified.


The more I learn, the more I realize how much I don't know.
Albert Einstein

Chapter 6

Conclusion

Local mild hyperthermia is a potent adjuvant cancer treatment modality, able to significantly enhance the desired effects of radio- and chemotherapies, without introducing additional toxicity and its unwanted side effects. To unlock the full therapeutic potential of hyperthermia therapy, the temperature elevation must be induced accurately and uniformly over a prolonged period of time. MR-HIFU is a promising technology for noninvasively applying powerful heating with millimeter-range accuracy based on near-real-time temperature measurements. However, currently most MR-HIFU treatments are performed using predetermined treatment plans or simple feedback controllers, resulting in suboptimal treatment quality. In this thesis, several relevant control challenges have been identified, in response to which MPC-based control solutions have been developed, aimed at enabling local MR-HIFU hyperthermia treatments of optimal quality. In this chapter, the key results of the individual contributions in this thesis will be summarized, and recommendations for interesting future research topics will be given.

6.1 Conclusions

This thesis provides a significant contribution toward realizing higher-quality MR-HIFU hyperthermia treatments in clinics by enabling more accurate temperature control using advanced MPC-based control solutions. In Chapter 1, three relevant research objectives have been formulated, which for convenience are repeated here.

Objective 1 *Enable the accurate and robust regulation of the temperature in the tumor over a prolonged period of time, despite the presence of considerable and varying model uncertainty.*

Objective 2 *Enable high treatment quality for large-volume MR-HIFU hyperthermia by simultaneously optimizing online the sonication plan for electronic beam steering and the mechanical transducer trajectory.*

Objective 3 *Enable the automated and target-conformal distribution of transducer positions and sonication points for optimal treatment quality.*

Objective 4 *Provide an experimental validation of the novel feedback control concepts developed in this thesis on a clinical MR-HIFU system.*

This thesis accomplishes these objectives via four individual contributions, each of which relates to a single chapter, see also Chapter 1, of which in this section the key results will be discussed and related to the objectives above.

6.1.1 Offset-free MPC

To achieve *Objective 1*, in Chapter 2 an offset-free MPC scheme is presented, enabling MR-HIFU hyperthermia treatments of optimal quality despite the presence of (severe) plant-model mismatch, which is inevitably present when treating patients due to the large uncertainty in thermal and thermoregulatory modeling, and due to the impracticability of extensive model identification. In particular, this approach allows for assuming a time-invariant plant model with homogeneous tissue properties, thereby avoiding the necessity for accurate thermal modeling. Subsequently, a disturbance observer is used to identify constant and slowly varying model errors, allowing for the MPC to take these discrepancies into account when computing the optimal control inputs to recover the steady-state control performance as achievable in case of no model mismatch. To also address *Objective 4*, this feedback setup has been implemented on a clinical MR-HIFU system, with which experiments have been performed on an in-vivo porcine thigh muscle with considerable vascularization in the target region. These experiments verified the offset-free MPC scheme's effectiveness for successfully estimating, and subsequently mitigating, the effects of time-invariant parameter mismatch and of the spatially and temporally varying heat loss by

blood perfusion. To be precise, it was shown that when using a standard MPC design without offset-free capabilities, a negative steady-state offset of 0.5 °C in the mean target temperature occurred, which using the presented offset-free MPC scheme is eliminated within approximately one minute. As we have recently shown that a standard MPC already outperforms a clinical state-of-the-art binary controller, the proposed offset-free MPC design is concluded to be suitable for enabling significantly more robust and accurate temperature control for MR-HIFU hyperthermia treatments in clinics.

6.1.2 Switched-actuator systems with setup times

In Chapter 3, the class of switched-actuator systems with setup times (SAcSSs) is introduced, which can be used to conveniently describe the tissue's thermal dynamics together with the transducer repositioning. A modeling framework tailored to SAcSSs is provided, which offers two key advantages. First, it allows for the user-friendly and systematic derivation of a mixed-integer state-space model that is directly suitable for formulating MI-MPCs. Second, using this framework the resulting model contains significantly less Boolean variables than via other typical modeling approaches, which consequently leads to more computationally efficient MI-MPC schemes. As these MI-MPC setups are suitable for simultaneously optimizing online the discrete mechanical transducer motions as well as the continuous sonication powers in large-volume MR-HIFU hyperthermia, as verified via numerical case studies, Chapter 3 provides an important step toward achieving *Objective 2*.

6.1.3 Target-conformal optimal actuator placement

In Chapter 4, a method is developed for the automatic distribution of transducer positions and sonication points, aimed at optimizing the achievable treatment quality for the specific tumor shape and tissue properties, thereby realizing *Objective 3*. The procedure is computationally tractable and provides the user with insightful tuning parameters to balance the importance of optimizing the steady-state temperature as expected from the thermal model, against the desire for distributing the transducer positions and sonication points as uniformly as possible throughout the target region to enable faster heat-up and generally improve robust performance in the presence of spatially varying disturbances or modeling errors. The procedure's effectiveness for optimally selecting the transducer positions and sonication points was demonstrated via a numerical case study involving an irregularly shaped target region, illustrating its usefulness for MR-HIFU control design in general. Regarding the sonication points, the case study revealed that near-optimal performance can be achieved using a strategically selected small set of points, and thereby limited controller complexity, thus paving the way for the practical implementation of more advanced control solu-

tions such as (MI-)MPC, which offer potential for superior control performance compared to simple feedback schemes, but in doing so are also often more prone to increased computational complexity. Besides accomplishing *Objective 3*, the developments in Chapter 4 also addressed *Objective 2* in two ways. First, in the case study in Chapter 4, the method was used to set up a MI-MPC scheme for large-volume MR-HIFU hyperthermia (derived using the SAcSS framework from Chapter 3), thereby directly contributing to *Objective 2*. Second, in Chapter 5 the sonication point selection steps of the actuator placement method were employed for the high-level input reduction in the hierarchical MI-MPC architecture for large-volume MR-HIFU hyperthermia, as will be discussed next in Section 6.1.4, thus again assisting in achieving *Objective 2*.

6.1.4 Hierarchical mixed-integer MPC

Chapter 5 is dedicated to fully realizing *Objective 2*. That is, although the MI-MPC schemes derived using the SAcSS framework introduced in Chapter 3 are less computationally demanding than via typical other modeling approaches, real-time implementation may still be challenging when using long prediction horizons and many transducer positions and sonication points, which may be required for optimal treatment quality. Moreover, while the target-conformal actuator placement method proposed in Chapter 4 allows for optimizing the control performance for a limited number of control inputs, this will always lead to some loss of performance and robustness. Therefore, in Chapter 5 a two-layer hierarchical (H)MI-MPC architecture is presented, in which in the higher layer the discrete transducer positioning is swiftly determined by a reduced-complexity MI-MPC, and subsequently in the low level the sonication powers are optimized via a QP based on a high-fidelity model for optimal performance and robustness. Consequently, the HMI-MPC achieves the same treatment quality as highly complex MI-MPCs, but with substantially smaller solver times, thereby satisfying the computational requirements for real-time feasibility. This is verified by phantom experiments on a clinical MR-HIFU setup, thereby also addressing *Objective 4*. Moreover, by comparing the experimental results to previously published results of a state-of-the-art binary controller, the HMI-MPC is shown to enable significantly more accurate and homogeneous heating over a large tumor region than the control design currently used in clinics.

6.2 Recommendations

There remain many open challenges in improving hyperthermia therapy for cancer treatment. Based on the results in this thesis and on observations made in parallel with this work, this section discusses several relevant directions for future research.

6.2.1 Hyperthermia therapy

Over the past decades in particular, tremendous research efforts have been made regarding hyperthermia in oncology, resulting in valuable knowledge about its underlying heat-induced mechanisms that are beneficial for cancer treatment [3, 4, 12–17], and substantial pre-clinical and clinical evidence of its potency as a chemo- and radiosensitizer [3, 16, 18–24]. These results provided the main motivation for the research in this thesis. However, further research into hyperthermia for cancer therapy remains urgently required today, as a deeper quantitative understanding of its effects, and their relation to treatment parameters such as temperature and duration, may lead to a significant improvement of the clinical success of hyperthermia-based treatments. To this end, it is imperative that more (quantitative) clinical trials regarding hyperthermia therapy will continue to be performed. Up to now, one of the key obstacles for exact quantitative research was the inability to accurately induce the temperature elevation throughout the tumor, which contributed to, for example, lack of clarity regarding which thermal parameter is best suited as a prognostic for treatment outcome [17, 62]. Regarding this problem, the control solutions in this thesis could serve as an enabling factor, as they can achieve more accurate, uniform, and reproducible treatment temperatures. In addition, more research and clinical trials involving hyperthermia therapy is desired, as this continues to build evidence for its safety and effectiveness, which contributes to more widespread acceptance and clinical application, and as new discoveries may translate into promising novel applications, such as hyperthermia in conjunction with proton irradiation [121–123], or indicate the applicability of hyperthermia therapy to different types of tumors.

6.2.2 HIFU applicator

The development of ultrasound applicators has led to highly advanced and versatile phased-array HIFU transducers with an electronically steerable focus [30]. This is a core technology for achieving high-performance MR-HIFU temperature control. Further developing these instruments and integrating them in clinical MR-HIFU setups enables improved actuator flexibility, which may lead to better control performance [124]. For example, compared to the spherically curved, random sparse array designs [85] such as in the Profound Sonalleve or Insightec Exablate, of which the former contains 256 acoustic elements randomly distributed (without overlap) over the 13 cm diameter transducer area [30, 125], planar transducers populated with significantly more (5000 to 8000) elements on a similar surface area are expected to exhibit a considerably larger electronic beam steering range [126, 127]. This allows for defining larger treatment cells inside which the focal spot can be rapidly steered, and consequently reduces the need for the mechanical repositioning of the transducer, during which no heating is allowed, in large-volume MR-HIFU treatments. As a result, such highly

populated planar transducers would significantly improve the ratio of heating versus non-heating time in each treatment cell (which for cycling steady-state behavior such as observed in Chapters 3 to 5 can be roughly approximated by $1{:}(s_{\text{cycle}} + N_{\text{cycle}})$, with s_{cycle} and N_{cycle} denoting the total transducer travel time and the number of treatment cells, respectively, in the cycle), which would positively affect the achievable homogeneity of the temperature distribution. Alternatively, the latter issue can be addressed by redesigning the transducer's robotic carrier system in such a manner that it allows for faster mechanical displacements.

6.2.3 MR measurement and estimation

As a result of outstanding developments in the area of MR thermometry, the current spatial, temporal, and thermal resolution is sufficient for feedback control via MR-HIFU [40, 42, 44, 45]. However, when using more high-performance controllers, such as the MPC algorithms presented in this thesis, the potential controller accuracy has surpassed the measurement accuracy. Hence, further improvements in the thermometry's reliability, quality, sampling frequency, and data processing time are expected to contribute to enabling tighter temperature control. Interesting ongoing developments include novel MR acquisition protocols for an improved balance between resolution, speed, and measurement accuracy [128], and more sophisticated data processing methods such as multi-baseline [129], referenceless [130], hybrid [119, 131], or various other [118, 120] techniques to reduce thermometry artifacts resulting from temperature-independent magnetic field changes such as due to motion (e.g., patient, organ, or respiratory motion) or strong spatial variations of the magnetic susceptibility in the measurement volume [44]. Indeed, such improved methods may mitigate artifacts such as encountered in the phantom experiments of Chapter 5.

Due to the spatial density of MR imaging, i.e., providing data on a planar or volumetric grid, for thermometry noise effects can to a great extent be attenuated using simple state observers, as done throughout this thesis. However, to increase functionality, extending the feedback designs in this thesis with more sophisticated estimation or image recognition strategies may prove to be useful. Examples of such extensions include (variations of) Kalman filters [132, 133] for adaptive balancing between model and measurement data when combined with online estimation schemes of the measurement noise and process noise covariances, or motion tracking schemes for the treatment of moving organs [131, 134–136] when combined with adaptive MPC setups that allow for model updating [137–139].

6.2.4 Experimental and clinical evaluation

The novel MPC schemes for MR-HIFU hyperthermia developed in this thesis have been verified regarding temperature control performance in simulations, phantom experiments, and in-vivo experiments. The results clearly indicate more accurate, robust, and homogeneous heating when compared to results reported previously for current state-of-the-art controllers, see Chapters 2 and 5. However, more experimental evaluation is warranted to provide further evidence demonstrating and quantifying the performance improvement of the proposed MPC-based strategies. When combined with the relation between the (thermal) treatment parameters and clinical outcome (which as discussed above in Section 6.2.1 is highly desirable to continue studying via quantitative clinical trials for hyperthermia therapy), such rigorous experimental proof may in turn be used as a solid indication of the potential clinical impact of the presented MPC designs, providing a stronger motivation for their clinical application. Moreover, further experimental and clinical testing of the developed control strategies under various different circumstances may allow for a better assessment the controller sensitivity with respect to, for example, modeling errors, measurement uncertainty, or organ motion. Such an analysis is important to identify which aspects of the feedback scheme are most practically relevant for further development, which allows for prioritizing possible future research directions, such as the measurement and estimation technologies mentioned above in Section 6.2.3 or the robust and adaptive MPC methods discussed below in Section 6.2.5.1 for dealing with parameter mismatch and organ motion.

6.2.5 MPC design

This thesis provides a clear motivation for using MPC to regulate the temperature in hyperthermia therapy. The MPC designs in this work can be extended or augmented to provide more rigorous theoretical support, enhance their practical performance, and diversify their applicability to different situations. Several interesting research directions are listed here.

6.2.5.1 Model uncertainty

In Chapter 2, it is argued that modeling errors are inevitable when treating patients, and correspondingly an offset-free MPC design is proposed to improve the control performance in the presence of constant or slowly varying plant-model mismatch. Such a feature is lacking from the (H)MI-MPC designs in Chapters 3 to 5 for large-volume MR-HIFU, and consequently their performance may be suboptimal in case of modeling uncertainty. The offset-free architecture from Chapter 2 is a powerful general-purpose solution, compounding the steady-state effects of all unmodeled behavior into an additive disturbance vector. However, such an approach is not suitable for controlling switched systems, such as the

large-volume MR-HIFU setup, due to the states and inputs not converging to a fixed point in steady state, but instead exhibiting cyclic or periodic behavior due to the transducer switching. To enhance the robust performance of the (H)MI-MPC schemes for large-volume MR-HIFU in a clinical setting, the development of an offset-free MPC design for switched systems could be pursued. Here, a particularly interesting challenge is the selection of the disturbance model, i.e., what type of uncertainty each disturbance variable should capture and how these should be incorporated in the state-space model in such a manner that convergence of the estimates is guaranteed. For example, by properly including output disturbances, measurement artifacts as encountered in Chapter 5 could be identified. Alternatively, especially regarding parametric uncertainty or to account for patient and organ motion, researchers may turn to the recently developed adaptive MPC methods incorporating online parameter estimation and robust recursive model updates [137–139], possibly employing machine learning techniques [140, 141].

6.2.5.2 Stability and convergence guarantees

The controllers developed in this thesis exhibit stable closed-loop behavior and desirable control performance, partly due to the inherent stability of the tissue's first-order thermal dynamics, but are not guaranteed to converge to an optimal equilibrium. Note that due to the stability of the thermal system and the absence of hard state constraints, we could easily guarantee asymptotic stability of the optimal nominal steady-state pair (x_s, u_s) via a standard tracking MPC formulation employing a terminal penalty, see [67, Section 2.5.3]. However, we opted not to use such a scheme for two reasons. First, a tracking MPC with terminal penalty was observed to exhibit significantly deteriorated transient performance, as it aims to minimize the distance to (x_s, u_s), resulting in excessive heating of the tumor region to achieve faster convergence of the temperature in the healthy tissue via diffusion, which is exactly opposite to the true treatment objective of heating the tumor to the optimum temperature while avoiding temperature elevations in healthy tissue. Second, in case of plant-model mismatch, which is inevitable in practice as discussed in Chapter 2, the optimal steady-state pair (x_s, u_s) of the nominal model may not be an equilibrium for the actual plant due to the nonnegativity input constraints $u_k \geq 0_{n_u}$, by which the stability proof of the tracking MPC does not hold anymore. The second issue can be resolved using an offset-free MPC scheme [78,79] together with a target selector [91,92] to determine at each time $k \in \mathbb{N}$ the augmented system's reachable optimal steady-state pair $(x_s(\hat{d}_k), u_s(\hat{d}_k))$ corresponding to the current disturbance estimate $\hat{d}_k$. However, when using a tracking cost with $(x_s(\hat{d}_k), u_s(\hat{d}_k))$ as target, the undesirable transient behavior persists. In addition, as discussed in Chapter 2, due to noise combined with input constraints being active in steady state, the optimal target $(x_s(\hat{d}_k), u_s(\hat{d}_k))$ may be far away from the reference tumor temperature

42 °C in case of an unfortunate disturbance estimate $\hat{d}_k$, which causes the target-tracking MPC to steer the system towards the undesirable target, see also [93], resulting in deteriorated treatment quality.

Therefore, in all MPC designs in this thesis we chose to directly optimize the true treatment objective, which involves minimizing the tumor temperature deviation from the typically unreachable 42 °C reference while protecting healthy tissue and avoiding overheating the tumor (and in Chapter 2 also penalizing the underheating of the tumor). Such an approach in fact corresponds to the rationale of *economic* (E)MPC [142–145]. Therefore, a meaningful future research direction would be to cast the MPC schemes in this thesis into the formal EMPC framework to assess whether asymptotic stability can be guaranteed, or whether adaptations are required and what their impact on the control performance is. For example, it may be interesting to investigate if EMPC architectures with terminal (set) constraints and penalties [93, 146–148] may guarantee stability without being overly restrictive or detrimental to performance, or if EMPC schemes with only simple linear end penalties [149–151] or even without any terminal conditions (but, instead, approximate optimality for sufficiently long horizons) are better suited for the problem at hand [152–154]. Notable extensions that may prove useful regarding the MPC algorithms in this thesis include EMPC with offset-free functionalities [79, 155–157], periodic cost functions [158–161], and hierarchical structures [162].

On a final note regarding stability guarantees, for the HMI-MPC scheme in Chapter 5 a closer look may be required into accounting for the prediction error in the high control layer resulting from the complexity reduction. A potential solution could be found in tube-based robust MPC methods [163–165] as applied in the hierarchical design of [166]. Note that also for tube-based robust MPC, integration exists toward EMPC [167] with additionally periodic [159] or offset-free [157] capabilities.

6.2.5.3 Input parameterization

In this thesis, a zero-order hold is assumed for the control inputs, which during optimization are allowed to vary for each discrete time step in the MPC prediction horizon for maximum control flexibility, which generally implies maximum performance. However, this also leads to (relatively) rapidly growing controller complexities for longer horizons, thereby essentially imposing an upper limit on the maximum horizon length, which in turn may negatively affect control performance. In this thesis, these negative effects have been observed to be inconsequential, partially as a result of the tissue's stable first-order thermal dynamics together with the hyperthermia temperature control problem essentially being a regulator problem due to the constant optimal tumor temperature. Nevertheless, in case longer horizons are desired, one may employ input parameterizations such as move blocking [67]. This allows for including longer predictions in the MPC

optimization, without increasing the number of independent decision variables and, thereby, the controller complexity. This may be useful when, for example, using an EMPC framework of which the stability proof requires a sufficiently long horizon [152, 153].

6.2.6 Applications

This thesis focused on MR-HIFU for hyperthermia therapy, clearly demonstrating its strengths for highly accurate feedback-controlled heating. Therefore, MPC is also considered a promising control solution for enhancing the quality of other thermal therapies.

6.2.6.1 Thermal ablation

As mentioned in Chapter 1, MR-HIFU is also highly suitable for thermal ablation. In fact, currently most clinical applications of (MR-)HIFU concern ablation, rather than (mild) hyperthermia, for the (experimental) treatment of tumors of the liver, breast, uterus, pancreas, bone, connective tissue, (para)thyroid, kidney, and brain, and for functional neurosurgery aimed at treating Parkinson's disease, essential tremor, epilepsy, and chronic neuropathic pain [8,30,47,83,168]. As this thesis clearly showcases the suitability of MPC for accurately controlled MR-HIFU treatments, and other researchers have also already recognized the potential of MPC for thermal ablation [69–71], continuing the development of such control strategies may further enhance the quality, safety, effectiveness, and time efficiency of MR-HIFU thermal ablation treatments. When translating the MPC designs in this thesis to thermal ablation, specific attention must be paid to the formulation of the cost function and constraints. For example, when the cost function involves the thermal dose, which is commonly defined as depending exponentially on time and temperature [17, 69, 116], a nonlinearity is introduced in the controller. However, based on [69–71], this is not expected to be a major hurdle, as in terms of control design it is one of the key strengths of MPC (and especially EMPC) to allow for intuitively specifying the cost function and constraints in such a manner that they reflect the underlying treatment objective.

6.2.6.2 MR-guided radiofrequency and microwave hyperthermia

Besides HIFU, electromagnetic (EM) waves are highly suitable for noninvasively heating internal body tissues. As discussed in Chapter 1, this is typically done using radiofrequency (RF) or microwave (MW) applicators. Actually, today hyperthermia treatments are more often performed using EM-based methods than via ultrasound [16, 29]. Although HIFU offers certain advantages, especially regarding its extraordinary spatial accuracy, RF/MW heating should not be considered an inferior technology for hyperthermia. To the contrary, RF/MW is complementary to HIFU, as the former is generally better suited for the heating

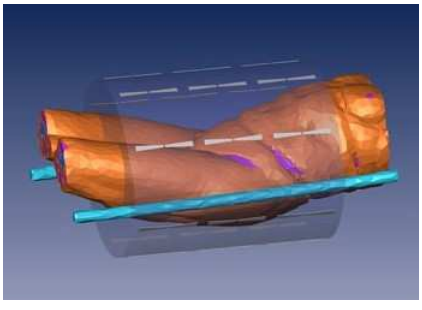

Figure 6.1: Visualization of the antennas (gray) around the patient (orange) for the BSD Sigma-Eye EM applicator [169].

of larger regions as in deep (loco)regional hyperthermia, or to treat areas that contain significant air pockets or which are obstructed by bone. With the ongoing development of more advanced MR-compatible multi-antenna phased-array EM applicators that enable three-dimensional focus steering, see Figure 6.1 for an example, there is major potential for high-quality MR-guided (MR-)RF/MW hyperthermia treatments. The current clinical state of the practice is to determine the optimal relative antenna settings offline using hyperthermia treatment planning [170–172], based on extensive EM and possibly thermal simulations using patient-specific models, and then scale the total heating power online via manual intervention. This practically comes down to using only one degree of freedom for controlling the temperature, which in addition can be assumed to not be optimally exploited. That is, the power steering is performed manually, but, as will be discussed in more detail below, the mapping from antenna settings to temperature is high-dimensional, multi-input multi-output, and nonlinear, and therefore too complex for manual optimization. Promising results have already been obtained in research evaluating online simulation-based re-optimization of the antenna settings, especially regarding the reduction of unwanted hot spots in healthy tissue, which are currently often a limiting factor in achieving adequate heating of the tumor [173–178]. Therefore, realizing a fully automated MPC feedback loop is expected to enable significantly enhanced treatment quality, as this allows for optimally controlling all antenna phases and amplitudes individually, which for the 12-antenna applicator in Figure 6.1 leads to 23 degrees of controller freedom[1], for instance. In addition, eliminating the clinician from the feedback loop provides a more standardized and reproducible treatment approach, which in principle (clinical acceptance and regulatory approval or certification aside) facilitates more reliable and widespread application of MR-RF/MW

[1]The constructive interference of EM waves depends on the antenna phases relative to each other. Therefore, one antenna phase is fixed to zero, leaving 11 phases and 12 amplitudes as independent control variables when using a 12-antenna applicator.

hyperthermia. Therefore, the development of high-performance MPC schemes for MR-RF/MW deserves considerable future research efforts. Although many principles of the MR-HIFU MPC designs in this thesis are easily translated to MR-RF/MW, new problems will also arise in closing the loop for MR-RF/MW, due to essential practical differences with MR-HIFU. Two notable challenges that are anticipated are discussed next.

A first key topic that may require significant attention is the large system scale inherent to MR-RF. That is, the modern advanced phased-array EM-applicators with steerable focus consist of multiple rings of antennas placed around the patient, as shown in Figure 6.1. Consequently, there is a large patient volume inside the applicator, which potentially receives heat and therefore must be monitored. Combined with the sufficiently high MR thermometry spatial resolution required to detect hot spots in the (sub)centimeter-range, especially at bone-tissue interfaces [171, 177], this may lead to the state and output dimensions of the thermal state-space model being over 10^5. To enable computational tractability and real-time feasibility of a corresponding MPC scheme, these incredibly large system dimensions must be dealt with properly. Possible solutions may be found in MPC and observer schemes based on reduced-order models [114, 133], condensed MPC formulations [67, 115], and graphical processing unit (GPU-)based parallel computing. In addition, due to the large number of measurements required for the entire monitored volume, the full MR thermometry acquisition sequence may take a significant amount time. Sampling intervals of around 90 seconds, for example, are not uncommon, resulting in considerable latency from a feedback control perspective. To reduce the acquisition time without sacrificing accuracy, an option would be to abandon the fixed full-volume measurement sequence, and instead process each slice[1] individually. This could be especially beneficial when combined with an algorithm that adaptively selects the slice at which a new measurement is needed most, which may be determined based on the estimated observer error or some weighting that indicates the performance-critical region. Moreover, an adaptive online switching between slow accurate measurements and fast coarse measurements may be employed [180]. Note, however, that such enhancements must be accompanied by developments of the MRI console to facilitate the online switching between different scanning protocols. In parallel, an observer design exploiting sensor fusion may improve the temperature estimation accuracy by combining the temporally sparse and spatially dense MR thermometry data with the temporally dense and spatially sparse readings of the (minimally) invasive or intraluminal thermocouple probes often used in MR-RF hyperthermia.

A second important challenge from a control perspective is the nonlinearity of the input model. That is, in EM-based heating, the specific absorption rate (SAR) indicates the power deposition density. Contrary to HIFU, where

[1] In MR imaging, measurement data is collected per slice, which is a thin volumetric layer centered at a planar cross section [179].

the control inputs enter the control-oriented state-space model linearly (describing the acoustic power at the sonication points, which in turn is realized by determining the corresponding acoustic driving signals' phases and amplitudes separately), using EM the SAR cannot be treated as a simple linear control input. In particular, although the SAR enters the state update equation linearly, the SAR itself depends nonlinearly on the antenna amplitudes and phases, and not all SAR profiles are physically realizable, as evidenced by the many works on SAR-based optimization of the antenna settings [170, 181–183]. Consequently, the MPC should either optimize the antenna settings directly, which leads to a nonlinear input model, or a suitable input linearization or parameterization must be performed, which may require additional input constraints to ensure its validity.

Bibliography

[1] F. Bray, J. Ferlay, I. Soerjomataram *et al.*, "Global cancer statistics 2018: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries," *CA: A Cancer Journal for Clinicians*, vol. 68, no. 6, pp. 394–424, 2018.

[2] CBS, "54 procent sterfgevallen in 2019 door kanker of hart- en vaatziekten," https://www.cbs.nl/nl-nl/nieuws/2020/27/54-procent-sterfgevallen-in-2019-door-kanker-of-hart-en-vaatziekten, 2020.

[3] M. Mallory, E. Gogineni, G. C. Jones *et al.*, "Therapeutic hyperthermia: The old, the new, and the upcoming," *Critical Reviews in Oncology/Hematology*, vol. 97, pp. 56–64, 2016.

[4] N. van den Tempel, M. R. Horsman, and R. Kanaar, "Improving efficacy of hyperthermia in oncology by exploiting biological mechanisms," *International Journal of Hyperthermia*, vol. 32, no. 4, pp. 446–454, 2016.

[5] J. H. Breasted, "The Edwin Smith Surgical Papyrus," Chicago, Tech. Rep., 1930.

[6] B. B. Singh, "Hyperthermia: An ancient science in India," *International Journal of Hyperthermia*, vol. 7, no. 1, pp. 1–6, 1991.

[7] M. H. Seegenschmiedt and C. C. Vernon, "A Historical Perspective on Hyperthermia in Oncology," in *Thermoradiotherapy and Thermochemotherapy*, M. H. Seegenschmiedt, P. Fessenden, and C. C. Vernon, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, ch. 1, pp. 3–44.

[8] H. P. Kok, E. N. Cressman, W. Ceelen *et al.*, "Heating technology for malignant tumors: a review," *International Journal of Hyperthermia*, vol. 37, no. 1, pp. 711–741, 2020.

[9] K. F. Chu and D. E. Dupuy, "Thermal ablation of tumours: Biological mechanisms and advances in therapy," *Nature Reviews Cancer*, vol. 14, no. 3, pp. 199–208, 2014.

[10] P. Wust, B. Hildebrandt, G. Sreenivasa *et al.*, "Hyperthermia in combined treatment of cancer," *The Lancet Oncology*, vol. 3, no. 8, pp. 487–497, 2002.

[11] J. Van Der Zee, "Heating the patient: A promising approach?" *Annals of Oncology*, vol. 13, no. 8, pp. 1173–1184, 2002.

[12] R. D. Issels, "Hyperthermia adds to chemotherapy," *European Journal of Cancer*, vol. 44, no. 17, pp. 2546–2554, 2008.

[13] M. W. Dewhirst, Z. Vujaskovic, E. Jones, and D. Thrall, "Re-setting the biologic rationale for thermal therapy," *International Journal of Hyperthermia*, vol. 21, no. 8, pp. 779–790, 2005.

[14] S. Toraya-Brown and S. Fiering, "Local tumour hyperthermia as immunotherapy for metastatic cancer," *International Journal of Hyperthermia*, vol. 30, no. 8, pp. 531–539, 2014.

[15] A. L. Oei, L. E. Vriend, J. Crezee *et al.*, "Effects of hyperthermia on DNA repair pathways: One treatment to inhibit them all," *Radiation Oncology*, vol. 10, no. 1, pp. 1–13, 2015.

[16] N. R. Datta, S. G. Ordóñez, U. S. Gaipl *et al.*, "Local hyperthermia combined with radiotherapy and-/or chemotherapy: recent advances and promises for the future." *Cancer treatment reviews*, vol. 41, no. 9, pp. 742–53, 2015.

[17] G. C. van Rhoon, "Is CEM43 still a relevant thermal dose parameter for hyperthermia treatment monitoring?" *International Journal of Hyperthermia*, vol. 32, no. 1, pp. 50–62, 2016.

[18] R. D. Issels, L. H. Lindner, J. Verweij *et al.*, "Neo-adjuvant chemotherapy alone or with regional hyperthermia for localised high-risk soft-tissue sarcoma: A randomised phase 3 multicentre study," *The Lancet Oncology*, vol. 11, no. 6, pp. 561–570, 2010.

[19] N. R. Datta, E. Puric, D. Klingbiel *et al.*, "Hyperthermia and Radiation Therapy in Locoregional Recurrent Breast Cancers: A Systematic Review and Meta-analysis," *International Journal of Radiation Oncology*Biology*Physics*, vol. 94, no. 5, pp. 1073–1087, 2016.

[20] N. R. Datta, S. Rogers, S. G. Ordóñez *et al.*, "Hyperthermia and radiotherapy in the management of head and neck cancers: A systematic

review and meta-analysis," *International Journal of Hyperthermia*, vol. 32, no. 1, pp. 31–40, 2016.

[21] N. R. Datta, S. Rogers, D. Klingbiel *et al.*, "Hyperthermia and radiotherapy with or without chemotherapy in locally advanced cervical cancer: a systematic review with conventional and network meta-analyses," *International Journal of Hyperthermia*, vol. 32, no. 7, pp. 809–821, 2016.

[22] R. D. Issels, L. H. Lindner, J. Verweij *et al.*, "Effect of Neoadjuvant Chemotherapy Plus Regional Hyperthermia on Long-term Outcomes Among Patients With Localized High-Risk Soft Tissue Sarcoma," *JAMA Oncology*, vol. 4, no. 4, pp. 483–492, 2018.

[23] M. Kroesen, H. T. Mulder, J. M. L. van Holthe *et al.*, "The Effect of the Time Interval Between Radiation and Hyperthermia on Clinical Outcome in 400 Locally Advanced Cervical Carcinoma Patients," *Frontiers in Oncology*, vol. 9, p. 134, 2019.

[24] M. Kroesen, H. T. Mulder, J. M. van Holthe *et al.*, "Confirmation of thermal dose as a predictor of local control in cervical carcinoma patients treated with state-of-the-art radiation therapy and hyperthermia," *Radiotherapy and Oncology*, vol. 140, no. July 2005, pp. 150–158, 2019.

[25] N. Hijnen, S. Langereis, and H. Grüll, "Magnetic resonance guided high-intensity focused ultrasound for image-guided temperature-induced drug delivery," *Advanced Drug Delivery Reviews*, vol. 72, pp. 65–81, 2014.

[26] R. M. Staruch, K. Hynynen, and R. Chopra, "Hyperthermia-mediated doxorubicin release from thermosensitive liposomes using MR-HIFU: Therapeutic effect in rabbit Vx2 tumours," *International Journal of Hyperthermia*, vol. 31, no. 2, pp. 118–133, 2015.

[27] N. Hijnen, E. Kneepkens, M. de Smet *et al.*, "Thermal combination therapies for local drug delivery by magnetic resonance-guided high-intensity focused ultrasound," *Proceedings of the National Academy of Sciences*, vol. 114, no. 24, pp. E4802–E4811, 2017.

[28] F. Adibzadeh, K. Sumser, S. Curto *et al.*, "Systematic review of preclinical and clinical devices for magnetic resonance-guided radiofrequency hyperthermia," *International Journal of Hyperthermia*, vol. 37, no. 1, pp. 15–27, 2020.

[29] M. M. Paulides, H. Dobsicek Trefna, S. Curto, and D. B. Rodrigues, "Recent technological advancements in radiofrequency- andmicrowave-mediated hyperthermia for enhancing drug delivery," *Advanced Drug Delivery Reviews*, 2020.

[30] K. Hynynen and R. M. Jones, "Image-guided ultrasound phased arrays are a disruptive technology for non-invasive therapy," *Physics in Medicine and Biology*, vol. 61, no. 17, pp. R206–R248, 2016.

[31] L. Zhu, M. B. Altman, A. Laszlo *et al.*, "Ultrasound Hyperthermia Technology for Radiosensitization," *Ultrasound in Medicine and Biology*, vol. 45, no. 5, pp. 1025–1043, 2019.

[32] R. Wood and A. L. Loomis, "The physical and biological effects of high-frequency sound-waves of great intensity," *Philosophical Magazine and Journal of Science*, vol. 4, no. 22, pp. 417–436, 1927.

[33] J. G. Lynn, R. L. Zwemer, A. J. Chick, and A. E. Miller, "A new method for the generation and use of focused ultrasound in experimental biology," *The Journal of general physiology*, vol. 26, no. 2, pp. 179–93, 1942.

[34] F. Fani, E. Schena, P. Saccomandi, and S. Silvestri, "CT-based thermometry: An overview," *International Journal of Hyperthermia*, vol. 30, no. 4, pp. 219–227, 2014.

[35] M. A. Lewis, R. M. Staruch, and R. Chopra, "Thermometry and ablation monitoring with ultrasound," *International Journal of Hyperthermia*, vol. 31, no. 2, pp. 163–181, 2015.

[36] E. S. Ebbini and G. Ter Haar, "Ultrasound-guided therapeutic focused ultrasound: Current status and future directions," *International Journal of Hyperthermia*, vol. 31, no. 2, pp. 77–89, 2015.

[37] K. Hynynen, "MRI-guided focused ultrasound treatments," *Ultrasonics*, vol. 50, no. 2, pp. 221–229, 2010.

[38] Y. Ishihara, A. Calderon, H. Watanabe *et al.*, "A precise and fast temperature mapping using water proton chemical shift," *Magnetic Resonance in Medicine*, vol. 34, no. 6, pp. 814–823, 1995.

[39] B. Quesson, J. A. De Zwart, and C. T. W. Moonen, "Magnetic resonance temperature imaging for guidance of thermotherapy," *Journal of Magnetic Resonance Imaging*, vol. 12, no. 4, pp. 525–533, 2000.

[40] V. Rieke and K. Butts Pauly, "MR thermometry," *Journal of Magnetic Resonance Imaging*, vol. 27, no. 2, pp. 376–390, 2008.

[41] O. I. Craciunescu, P. R. Stauffer, B. J. Soher *et al.*, "Accuracy of real time noninvasive temperature measurements using magnetic resonance thermal imaging in patients treated for high grade extremity soft tissue sarcomas." *Medical physics*, vol. 36, pp. 4848–4858, 2009.

[42] L. Lüdemann, W. Wlodarczyk, J. Nadobny *et al.*, "Non-invasive magnetic resonance thermography during regional hyperthermia," *International Journal of Hyperthermia*, vol. 26, no. 3, pp. 273–282, 2010.

[43] Y. S. Kim, "Advances in MR image-guided high-intensity focused ultrasound therapy," *International Journal of Hyperthermia*, vol. 31, no. 3, pp. 225–232, 2015.

[44] L. Winter, E. Oberacker, K. Paul *et al.*, "Magnetic resonance thermometry: Methodology, pitfalls and practical solutions," *International Journal of Hyperthermia*, vol. 32, no. 1, pp. 63–75, 2016.

[45] Curto, Aklan, Mulder *et al.*, "Quantitative, Multi-institutional Evaluation of MR Thermometry Accuracy for Deep-Pelvic MR-Hyperthermia Systems Operating in Multi-vendor MR-systems Using a New Anthropomorphic Phantom," *Cancers*, vol. 11, no. 11, p. 1709, 2019.

[46] L. C. Sebeke, P. Rademann, A. C. Maul *et al.*, "Visualization of Thermal Washout due to Spatiotemporally Heterogeneous Perfusion in the Application of a Model-Based Control Algorithm for MR-HIFU Mediated Hyperthermia," *Submitted for journal publication.*

[47] D. Schlesinger, S. Benedict, C. Diederich *et al.*, "MR-guided focused ultrasound surgery, present and future," *Medical Physics*, vol. 40, no. 8, p. 080901, 2013.

[48] W. Chu, R. M. Staruch, S. Pichardo *et al.*, "Magnetic resonance-guided high-intensity focused ultrasound hyperthermia for recurrent rectal cancer: MR thermometry evaluation and preclinical validation," *International Journal of Radiation Oncology Biology Physics*, vol. 95, no. 4, pp. 1259–1267, 2016.

[49] F. Siedek, S. Y. Yeo, E. Heijman *et al.*, "Magnetic Resonance-Guided High-Intensity Focused Ultrasound (MR-HIFU): Technical Background and Overview of Current Clinical Applications (Part 1)," *RöFo - Fortschritte auf dem Gebiet der Röntgenstrahlen und der bildgebenden Verfahren*, vol. 191, no. 06, pp. 522–530, 2019.

[50] J. Enholm, M. Köhler, B. Quesson *et al.*, "Improved Volumetric MR-HIFU Ablation by Robust Binary Feedback Control," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 1, pp. 103–113, 2010.

[51] A. Partanen, P. S. Yarmolenko, A. Viitala *et al.*, "Mild hyperthermia with magnetic resonance-guided high-intensity focused ultrasound for applications in drug delivery," *International Journal of Hyperthermia*, vol. 28, no. 4, pp. 320–336, 2012.

[52] B. E. O'Neill, C. Karmonik, and K. C. P. Li, "An optimum method for pulsed high intensity focused ultrasound treatment of large volumes using the InSightec ExAblate® 2000 system," *Physics in Medicine and Biology*, vol. 55, no. 21, pp. 6395–6410, 2010.

[53] M. Tillander, S. Hokland, J. Koskela *et al.*, "High intensity focused ultrasound induced in vivo large volume hyperthermia under 3D MRI temperature control," *Medical Physics*, vol. 43, no. 3, pp. 1539–1549, 2016.

[54] R. Salomir, F. C. Vimeux, J. A. de Zwart *et al.*, "Hyperthermia by MR-guided focused ultrasound: Accurate temperature control based on fast MRI and a physical model of local energy deposition and heat conduction," *Magnetic Resonance in Medicine*, vol. 43, no. 3, pp. 342–347, 2000.

[55] C. Mougenot, R. Salomir, J. Palussière *et al.*, "Automatic spatial and temporal temperature control for MR-guided focused ultrasound using fast 3D MR thermometry and multispiral trajectory of the focal point," *Magnetic Resonance in Medicine*, vol. 52, no. 5, pp. 1005–1015, 2004.

[56] B. D. De Senneville, C. Mougenot, and C. T. W. Moonen, "Real-time adaptive methods for treatment of mobile organs by MRI-controlled high-intensity focused ultrasound," *Magnetic Resonance in Medicine*, vol. 57, no. 2, pp. 319–330, 2007.

[57] C. Mougenot, B. Quesson, B. D. de Senneville *et al.*, "Three-dimensional spatial and temporal temperature control with MR thermometry-guided focused ultrasound (MRgHIFU)," *Magnetic Resonance in Medicine*, vol. 61, no. 3, pp. 603–614, 2009.

[58] R. Staruch, R. Chopra, and K. Hynynen, "Localised drug release using MRI-controlled focused ultrasound hyperthermia," *International Journal of Hyperthermia*, vol. 27, no. 2, pp. 156–171, 2011.

[59] C. Bing, J. Nofiele, R. Staruch *et al.*, "Localised hyperthermia in rodent models using an MRI-compatible high-intensity focused ultrasound system," *International Journal of Hyperthermia*, vol. 31, no. 8, pp. 813–822, 2015.

[60] M. Franckena, D. Fatehi, M. de Bruijne *et al.*, "Hyperthermia dose-effect relationship in 420 patients with cervical cancer treated with combined radiotherapy and hyperthermia," *European Journal of Cancer*, vol. 45, no. 11, pp. 1969–1978, 2009.

[61] T. Ohguri, Y. Harima, H. Imada *et al.*, "Relationships between thermal dose parameters and the efficacy of definitive chemoradiotherapy plus

regional hyperthermia in the treatment of locally advanced cervical cancer: data from a multicentre randomised clinical trial," *International Journal of Hyperthermia*, vol. 34, no. 4, pp. 461–468, 2018.

[62] A. Bakker, J. van der Zee, G. van Tienhoven *et al.*, "Temperature and thermal dose during radiotherapy and hyperthermia for recurrent breast cancer are related to clinical outcome and thermal toxicity: a systematic review," *International Journal of Hyperthermia*, vol. 36, no. 1, pp. 1024–1039, 2019.

[63] J. Maciejowski, *Predictive Control with Constraints*. Prentice Hall, 2000.

[64] E. F. Camacho and C. Bordons, *Model Predictive control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer London, 2007.

[65] J. Rossiter, *Model-Based Predictive Control*, J. Rossiter, Ed. CRC Press, jul 2017.

[66] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.

[67] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, LLC, 2017.

[68] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017.

[69] D. Arora, M. Skliar, D. Cooley, and R. B. Roemer, "Constrained Predictive Control of Thermal Therapies for Minimum-Time Delivery of Thermal Dose," *IEEE Transactions on Control Systems Technology*, vol. 15, no. 6, pp. 1030–1037, 2007.

[70] J. de Bever, N. Todd, A. Payne *et al.*, "Adaptive model-predictive controller for magnetic resonance guided focused ultrasound therapy," *International Journal of Hyperthermia*, vol. 30, no. 7, pp. 456–470, 2014.

[71] D. Hensley, R. Orendorff, E. Yu *et al.*, "Model predictive control for treating cancer with ultrasonic heating," in *2015 American Control Conference (ACC)*. IEEE, jul 2015, pp. 220–225.

[72] L. Sebeke, D. A. Deenen, E. Maljaars *et al.*, "Model predictive control for MR-HIFU-mediated, uniform hyperthermia," *International Journal of Hyperthermia*, vol. 36, no. 1, pp. 1039–1049, 2019.

[73] P. W. Vaupel and D. K. Kelleher, "Pathophysiological and vascular characteristics of tumours and their importance for hyperthermia: Heterogeneity is the key issue," *International Journal of Hyperthermia*, vol. 26, no. 3, pp. 211–223, 2010.

[74] M. M. Paulides, K. Sumser, I. Vilas-Boas Ribeiro *et al.*, "Challenges and opportunities in thermal tissue modelling for electromagnetic applications," in *13th European Conference on Antennas and Propagation (EuCAP)*. IEEE, 2019, pp. 1–4.

[75] C. W. Song, A. Lokshina, J. G. Rhee *et al.*, "Implication of Blood Flow in Hyperthermic Treatment of Tumors," *IEEE Transactions on Biomedical Engineering*, vol. BME-31, no. 1, pp. 9–16, 1984.

[76] C. W. Song, "Effect of local hyperthermia on blood flow and microenvironment: A review," *Cancer Research*, vol. 44, no. 10 Supplement, pp. 4721s–4730s, 1984.

[77] H. Arkin, L. Xu, and K. Holmes, "Recent developments in modeling heat transfer in blood perfused tissues," *IEEE Transactions on Biomedical Engineering*, vol. 41, no. 2, pp. 97–107, 1994.

[78] U. Maeder, F. Borrelli, and M. Morari, "Linear offset-free Model Predictive Control," *Automatica*, vol. 45, no. 10, pp. 2214–2222, 2009.

[79] G. Pannocchia, M. Gabiccini, and A. Artoni, "Offset-free MPC explained: Novelties, subtleties, and applications," *IFAC-PapersOnLine*, vol. 48, no. 23, pp. 342–351, 2015.

[80] G. Appa, L. Pitsoulis, and H. P. Williams, Eds., *Handbook on Modelling for Discrete Optimization*, ser. International Series in Operations Research & Management Science. Boston, MA: Springer US, 2006, vol. 88.

[81] D. Liberzon, *Switching in Systems and Control*, ser. Systems & Control: Foundations & Applications. Boston, MA: Birkhäuser Boston, 2003.

[82] R. Scattolini, "Architectures for distributed and hierarchical Model Predictive Control – A review," *Journal of Process Control*, vol. 19, no. 5, pp. 723–731, 2009.

[83] E. Maloney and J. H. Hwang, "Emerging HIFU applications in cancer therapy," *International Journal of Hyperthermia*, vol. 31, no. 3, pp. 302–309, 2015.

[84] D. A. Deenen, E. Maljaars, L. Sebeke *et al.*, "Offset-free model predictive control for enhancing MR-HIFU hyperthermia in cancer treatment," in *IFAC-PapersOnLine*, vol. 51, no. 20, 2018, pp. 191–196.

[85] J. W. Hand, A. Shaw, N. Sadhoo *et al.*, "A random phased array device for delivery of high intensity focused ultrasound," *Physics in Medicine and Biology*, vol. 54, no. 19, pp. 5675–5693, 2009.

[86] H. H. Pennes, "Analysis of Tissue and Arterial Blood Temperatures in the Resting Human Forearm," *Journal of Applied Physiology*, vol. 1, no. 2, pp. 93–122, 1948.

[87] D. Akyürekli, L. H. Gerig, and G. P. Raaphorst, "Changes in muscle blood flow distribution during hyperthermia," *International Journal of Hyperthermia*, vol. 13, no. 5, pp. 481–496, 1997.

[88] M. Cloosterman, L. Hetel, N. van de Wouw *et al.*, "Controller synthesis for networked control systems," *Automatica*, vol. 46, no. 10, pp. 1584–1594, 2010.

[89] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[90] A. H. Negussie, A. Partanen, A. S. Mikhail *et al.*, "Thermochromic tissue-mimicking phantom for optimisation of thermal tumour ablation," *International Journal of Hyperthermia*, vol. 32, no. 3, pp. 239–243, 2016.

[91] K. R. Muske and J. B. Rawlings, "Model predictive control with linear models," *AIChE Journal*, vol. 39, no. 2, pp. 262–287, 1993.

[92] K. R. Muske and T. A. Badgwell, "Disturbance modeling for offset-free linear model predictive control," *Journal of Process Control*, vol. 12, no. 5, pp. 617–632, 2002.

[93] J. B. Rawlings, D. Bonne, J. B. Jorgensen *et al.*, "Unreachable Setpoints in Model Predictive Control," *IEEE Transactions on Automatic Control*, vol. 53, no. 9, pp. 2209–2215, 2008.

[94] B. Zaporzan, A. C. Waspe, T. Looi *et al.*, "MatMRI and MatHIFU: software toolboxes for real-time monitoring and control of MR-guided HIFU." *Journal of therapeutic ultrasound*, vol. 1, no. 1, p. 7, 2013.

[95] IT'IS Foundation, "Tissue Properties Database V4.0," 2018.

[96] A. Allahverdi, J. N. Gupta, and T. Aldowaisan, "A review of scheduling research involving setup considerations," *Omega*, vol. 27, no. 2, pp. 219–239, 1999.

[97] A. Allahverdi, C. T. Ng, T. C. Cheng, and M. Y. Kovalyov, "A survey of scheduling problems with setup times or costs," *European Journal of Operational Research*, vol. 187, no. 3, pp. 985–1032, 2008.

[98] A. Allahverdi, "The third comprehensive survey on scheduling problems with setup times/costs," *European Journal of Operational Research*, vol. 246, no. 2, pp. 345–378, 2015.

[99] A. T. Cobbenhagen, D. J. Antunes, M. J. Van De Molengraft, and W. P. Heemels, "Coverage control for outbreak dynamics," *2017 IEEE 56th Annual Conference on Decision and Control, CDC 2017*, vol. 2018-Janua, no. Cdc, pp. 984–989, 2018.

[100] L. P. Schoonen, A. T. Cobbenhagen, and W. P. Heemels, "Optimal irrigation management for large-scale arable farming using model predictive control," *IFAC-PapersOnLine*, vol. 52, no. 30, pp. 56–61, 2019.

[101] G. H. Donovan and D. B. Rideout, "An integer programming model to optimize resource allocation for wildfire containment," *Forest Science*, vol. 49, no. 2, pp. 331–335, 2003.

[102] R. G. Haight and J. S. Fried, "Deploying wildland fire suppression resources with a scenario-based standard response model," *Infor*, vol. 45, no. 1, pp. 31–39, 2007.

[103] N. Petrovic, D. L. Alderson, and J. M. Carlson, "Dynamic resource allocation in disaster response: Tradeoffs in wildfire suppression," *PLoS ONE*, vol. 7, no. 4, 2012.

[104] A. Bemporad and M. Morari, "Control of systems integrating logic, dynamics, and constraints," *Automatica*, vol. 35, no. 3, pp. 407–427, 1999.

[105] F. Torrisi and A. Bemporad, "HYSDEL - A Tool for Generating Computational Hybrid Models for Analysis and Synthesis Problems," *IEEE Transactions on Control Systems Technology*, vol. 12, no. 2, pp. 235–249, 2004.

[106] M. Philippe, R. Essick, G. E. Dullerud, and R. M. Jungers, "Stability of discrete-time switching systems with constrained switching sequences," *Automatica*, vol. 72, pp. 242–250, 2016.

[107] K. Subramanian, C. T. Maravelias, and J. B. Rawlings, "A state-space model for chemical production scheduling," *Computers and Chemical Engineering*, vol. 47, pp. 97–110, 2012.

[108] D. A. Deenen, E. Maljaars, L. Sebeke *et al.*, "Mixed-integer model predictive control for large-area MR-HIFU hyperthermia in cancer therapy," in *IFAC World Congress 2020*.

[109] J. Gross, J. Yellen, and P. Zhang, Eds., *Handbook of Graph Theory, Second Edition*, 2nd ed., ser. Discrete Mathematics and Its Applications. New York: Chapman and Hall/CRC, dec 2013.

[110] A. Bemporad and S. Di Cairano, "Model-Predictive Control of Discrete Hybrid Stochastic Automata," *IEEE Transactions on Automatic Control*, vol. 56, no. 6, pp. 1307–1321, 2011.

[111] E. L. Lawler and D. E. Wood, "Branch-and-Bound Methods: A Survey," *Operations Research*, vol. 14, no. 4, pp. 699–719, 1966.

[112] M. Padberg and G. Rinaldi, "A Branch-and-Cut Algorithm for the Resolution of Large-Scale Symmetric Traveling Salesman Problems," *SIAM Review*, vol. 33, no. 1, pp. 60–100, 1991.

[113] G. Laporte, S. Nickel, and F. Saldanha da Gama, Eds., *Location Science*. Cham: Springer International Publishing, 2015.

[114] A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems*. Society for Industrial and Applied Mathematics, jan 2005.

[115] J. L. Jerez, E. C. Kerrigan, and G. A. Constantinides, "A sparse and condensed QP formulation for predictive control of LTI systems," *Automatica*, vol. 48, no. 5, pp. 999–1002, 2012.

[116] M. W. Dewhirst, B. L. Viglianti, M. Lora-Michiels *et al.*, "Basic principles of thermal dosimetry and thermal thresholds for tissue damage from hyperthermia," *International Journal of Hyperthermia*, vol. 19, no. 3, pp. 267–294, 2003.

[117] R. Niu and M. Skliar, "Identification of reduced-order thermal therapy models using thermal MR images: Theory and validation," *IEEE Transactions on Medical Imaging*, vol. 31, no. 7, pp. 1493–1504, 2012.

[118] C. Bing, R. M. Staruch, M. Tillander *et al.*, "Drift correction for accurate PRF-shift MR thermometry during mild hyperthermia treatments with MR-HIFU," *International Journal of Hyperthermia*, vol. 32, no. 6, pp. 673–687, 2016.

[119] W. A. Grissom, V. Rieke, A. B. Holbrook *et al.*, "Hybrid referenceless and multibaseline subtraction MR thermometry for monitoring thermal therapies in moving organs," *Medical Physics*, vol. 37, no. 9, p. 5014, 2010.

[120] M. Wu, H. T. Mulder, P. Baron *et al.*, "Correction of motion-induced susceptibility artifacts and B0 drift during proton resonance frequency shift-based MR thermometry in the pelvis with background field removal methods," *Magnetic Resonance in Medicine*, vol. 84, no. 5, pp. 2495–2511, 2020.

[121] N. Grosse, A. O. Fontana, E. B. Hug *et al.*, "Deficiency in Homologous Recombination Renders Mammalian Cells More Sensitive to Proton Versus Photon Irradiation," *International Journal of Radiation Oncology*Biology*Physics*, vol. 88, no. 1, pp. 175–181, 2014.

[122] N. R. Datta, E. Puric, R. Schneider *et al.*, "Could hyperthermia with proton therapy mimic carbon ion therapy? Exploring a thermo-radiobiological rationale," *International Journal of Hyperthermia*, vol. 30, no. 7, pp. 524–530, 2014.

[123] S. Tran, E. Puric, M. Walser *et al.*, "Early results and volumetric analysis after spot-scanning proton therapy with concomitant hyperthermia in large inoperable sacral chordomas," *The British Journal of Radiology*, vol. 93, no. 1107, p. 20180883, 2020.

[124] P. Ramaekers, M. de Greef, R. Berriet *et al.*, "Evaluation of a novel therapeutic focused ultrasound transducer based on Fermat's spiral," *Physics in Medicine and Biology*, vol. 62, no. 12, pp. 5021–5045, 2017.

[125] M. O. Köhler, C. Mougenot, B. Quesson *et al.*, "Volumetric HIFU ablation under 3D guidance of rapid MRI thermometry," *Medical Physics*, vol. 36, no. 8, pp. 3521–3535, 2009.

[126] N. Ellens, A. Pulkkinen, J. Song, and K. Hynynen, "The utility of sparse 2D fully electronically steerable focused ultrasound phased arrays for thermal surgery: a simulation study," *Physics in Medicine and Biology*, vol. 56, no. 15, pp. 4913–4932, 2011.

[127] N. P. Ellens, B. B. Lucht, S. T. Gunaseelan *et al.*, "A novel, flat, electronically-steered phased array transducer for tissue ablation: Preliminary results," *Physics in Medicine and Biology*, vol. 60, no. 6, pp. 2195–2215, 2015.

[128] M. Marx and K. Butts Pauly, "Improved MRI thermometry with multiple-echo spirals," *Magnetic Resonance in Medicine*, vol. 76, no. 3, pp. 747–756, 2016.

[129] S. Pichardo, M. Köhler, J. Lee, and K. Hynnyen, "In vivo optimisation study for multi-baseline MR-based thermometry in the context of hyperthermia using MR-guided high intensity focused ultrasound for head and neck applications," *International Journal of Hyperthermia*, vol. 30, no. 8, pp. 579–592, 2014.

[130] V. Rieke, K. K. Vigen, G. Sommer *et al.*, "Referenceless PRF shift thermometry," *Magnetic Resonance in Medicine*, vol. 51, no. 6, pp. 1223–1231, 2004.

[131] B. D. De Senneville, S. Roujol, C. Moonen, and M. Ries, "Motion correction in MR thermometry of abdominal organs: A comparison of the referenceless vs. the multibaseline approach," *Magnetic Resonance in Medicine*, vol. 64, no. 5, pp. 1373–1381, 2010.

[132] S. Roujol, B. D. De Senneville, S. Hey *et al.*, "Robust adaptive extended kalman filtering for real time MR-thermometry guided HIFU interventions," *IEEE Transactions on Medical Imaging*, vol. 31, no. 3, pp. 533–542, 2012.

[133] R. Hendrikx, S. Curto, B. De Jager *et al.*, "POD-Based Recursive Temperature Estimation for MR-Guided RF Hyperthermia Cancer Treatment: A Pilot Study," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, dec 2018, pp. 5201–5208.

[134] M. Ries, B. D. De Senneville, S. Roujol *et al.*, "Real-time 3D target tracking in MRI guided focused ultrasound ablations in moving tissues," *Magnetic Resonance in Medicine*, vol. 64, no. 6, pp. 1704–1712, 2010.

[135] P. Bour, V. Ozenne, F. Marquet *et al.*, "Real-time 3D ultrasound based motion tracking for the treatment of mobile organs with MR-guided high-intensity focused ultrasound," *International Journal of Hyperthermia*, vol. 34, no. 8, pp. 1225–1235, 2018.

[136] K. Kuroda, "MR techniques for guiding high-intensity focused ultrasound (HIFU) treatments," *Journal of Magnetic Resonance Imaging*, vol. 47, no. 2, pp. 316–331, 2018.

[137] M. Tanaskovic, L. Fagiano, and V. Gligorovski, "Adaptive model predictive control for linear time varying MIMO systems," *Automatica*, vol. 105, pp. 237–245, 2019.

[138] M. Lorenzen, M. Cannon, and F. Allgöwer, "Robust MPC with recursive model update," *Automatica*, vol. 103, pp. 461–471, 2019.

[139] K. Zhang and Y. Shi, "Adaptive model predictive control for a class of constrained linear systems with parametric uncertainties," *Automatica*, vol. 117, p. 108974, 2020.

[140] C. J. Ostafew, A. P. Schoellig, T. D. Barfoot, and J. Collier, "Learning-based Nonlinear Model Predictive Control to Improve Vision-based Mobile Robot Path Tracking," *Journal of Field Robotics*, vol. 33, no. 1, pp. 133–152, 2016.

[141] E. D. Klenske, M. N. Zeilinger, B. Scholkopf, and P. Hennig, "Gaussian Process-Based Predictive Control for Periodic Error Correction," *IEEE Transactions on Control Systems Technology*, vol. 24, no. 1, pp. 110–121, 2016.

[142] J. B. Rawlings and R. Amrit, "Optimizing Process Economic Performance Using Model Predictive Control," in *Nonlinear Model Predictive Control: Towards New Challenging Applications*, L. Magni, D. M. Raimondo, and

F. Allgöwer, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 119–138.

[143] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*. IEEE, dec 2012, pp. 3851–3861.

[144] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

[145] T. Faulwasser, L. Grüne, and M. A. Müller, "Economic Nonlinear Model Predictive Control," *Foundations and Trends® in Systems and Control*, vol. 5, no. 1, pp. 224–409, 2018.

[146] M. Diehl, R. Amrit, and J. B. Rawlings, "A Lyapunov Function for Economic Optimizing Model Predictive Control," *IEEE Transactions on Automatic Control*, vol. 56, no. 3, pp. 703–707, 2011.

[147] R. Amrit, J. B. Rawlings, and D. Angeli, "Economic optimization using model predictive control with a terminal cost," *Annual Reviews in Control*, vol. 35, no. 2, pp. 178–186, 2011.

[148] D. Angeli, R. Amrit, and J. B. Rawlings, "On average performance and stability of economic model predictive control," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1615–1626, 2012.

[149] T. Faulwasser and D. Bonvin, "Exact turnpike properties and economic NMPC," *European Journal of Control*, vol. 35, pp. 34–41, 2017.

[150] T. Faulwasser and M. Zanon, "Asymptotic Stability of Economic NMPC: The Importance of Adjoints," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 157–168, 2018.

[151] M. Zanon and T. Faulwasser, "Economic MPC without terminal constraints: Gradient-correcting end penalties enforce asymptotic stability," *Journal of Process Control*, vol. 63, pp. 1–14, 2018.

[152] L. Grüne, "Economic receding horizon control without terminal constraints," *Automatica*, vol. 49, no. 3, pp. 725–734, 2013.

[153] L. Grüne and M. Stieler, "Asymptotic stability and transient optimality of economic MPC without terminal conditions," *Journal of Process Control*, vol. 24, no. 8, pp. 1187–1196, 2014.

[154] T. Faulwasser and D. Bonvin, "On the design of economic NMPC based on approximate turnpike properties," in *2015 54th IEEE Conference on Decision and Control (CDC)*. IEEE, dec 2015, pp. 4964–4970.

[155] M. Vaccari and G. Pannocchia, "Implementation of an economic MPC with robustly optimal steady-state behavior," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 92–97, 2018.

[156] G. Pannocchia, "An economic MPC formulation with offset-free asymptotic performance," *IFAC-PapersOnLine*, vol. 51, no. 18, pp. 393–398, 2018.

[157] Z. Dong and D. Angeli, "Homothetic tube-based robust offset-free economic Model Predictive Control," *Automatica*, vol. 119, p. 109105, 2020.

[158] M. Zanon, S. Gros, and M. Diehl, "A Lyapunov function for periodic economic optimizing Model Predictive Control," in *52nd IEEE Conference on Decision and Control*. IEEE, dec 2013, pp. 5107–5112.

[159] T. J. Broomhead, C. Manzie, R. C. Shekhar, and P. Hield, "Robust periodic economic MPC for linear systems," *Automatica*, vol. 60, pp. 30–37, 2015.

[160] M. A. Müller and L. Grüne, "Economic model predictive control without terminal constraints for optimal periodic behavior," *Automatica*, vol. 70, pp. 128–139, 2016.

[161] M. Zanon, L. Grune, and M. Diehl, "Periodic Optimal Control, Dissipativity and MPC," *IEEE Transactions on Automatic Control*, vol. 62, no. 6, pp. 2943–2949, 2017.

[162] W. C. Clarke, C. Manzie, and M. J. Brear, "Hierarchical economic MPC for systems with storage states," *Automatica*, vol. 94, pp. 138–150, 2018.

[163] D. Mayne, M. Seron, and S. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.

[164] D. Mayne, "Robust and stochastic model predictive control: Are we going in the right direction?" *Annual Reviews in Control*, vol. 41, pp. 184–192, 2016.

[165] M. B. SaltÄśk, L. Özkan, J. H. Ludlage *et al.*, "An outlook on robust model predictive control algorithms: Reflections on performance and computational aspects," *Journal of Process Control*, vol. 61, pp. 77–102, 2018.

[166] M. Farina, X. Zhang, and R. Scattolini, "A hierarchical MPC scheme for interconnected systems," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 12 021–12 026, 2017.

[167] F. A. Bayer, M. A. Müller, and F. Allgöwer, "Tube-based robust economic model predictive control," *Journal of Process Control*, vol. 24, no. 8, pp. 1237–1246, 2014.

[168] S. Ellis, V. Rieke, M. Kohi, and A. C. Westphalen, "Clinical applications for magnetic resonance guided high intensity focused ultrasound (MRgHIFU): Present and future," *Journal of Medical Imaging and Radiation Oncology*, vol. 57, no. 4, pp. 391–399, 2013.

[169] R. A. M. Canters, M. M. Paulides, M. Franckena *et al.*, "Benefit of replacing the Sigma-60 by the Sigma-Eye applicator: A Monte Carlo-based uncertainty analysis," *Strahlentherapie und Onkologie*, vol. 189, no. 1, pp. 74–80, 2013.

[170] M. M. Paulides, P. R. Stauffer, E. Neufeld *et al.*, "Simulation techniques in hyperthermia treatment planning," *International Journal of Hyperthermia*, vol. 29, no. 4, pp. 346–357, 2013.

[171] H. Kok, P. Wust, P. Stauffer *et al.*, "Current state of the art of regional hyperthermia treatment planning: a review," *Radiation Oncology*, vol. 10, no. 1, p. 196, 2015.

[172] H. P. Kok, A. N. Kotte, and J. Crezee, "Planning, optimisation and evaluation of hyperthermia treatments," *International Journal of Hyperthermia*, vol. 33, no. 6, pp. 593–607, 2017.

[173] R. a. M. Canters, M. M. Paulides, M. F. Franckena *et al.*, "Implementation of treatment planning in the routine clinical procedure of regional hyperthermia treatment of cervical cancer: An overview and the Rotterdam experience," *International Journal of Hyperthermia*, vol. 28, no. 6, pp. 570–581, 2012.

[174] Z. Rijnen, J. F. Bakker, R. A. Canters *et al.*, "Clinical integration of software tool VEDO for adaptive and quantitative application of phased array hyperthermia in the head and neck," *International Journal of Hyperthermia*, vol. 29, no. 3, pp. 181–193, 2013.

[175] H. P. Kok, S. Ciampa, R. De Kroon-Oldenhof *et al.*, "Toward online adaptive hyperthermia treatment planning: Correlation between measured and simulated specific absorption rate changes caused by phase steering in patients," *International Journal of Radiation Oncology Biology Physics*, vol. 90, no. 2, pp. 438–445, 2014.

[176] H. P. Kok, L. Korshuize-van Straten, A. Bakker *et al.*, "Feasibility of on-line temperature-based hyperthermia treatment planning to improve tumour temperatures during locoregional hyperthermia," *International Journal of Hyperthermia*, vol. 34, no. 7, pp. 1–10, 2017.

[177] H. P. Kok, L. Korshuize- van Straten, A. Bakker *et al.*, "Online Adaptive Hyperthermia Treatment Planning During Locoregional Heating to Suppress Treatment-Limiting Hot Spots," *International Journal of Radiation Oncology Biology Physics*, vol. 99, no. 4, pp. 1039–1047, 2017.

[178] H. P. Kok, G. Schooneveldt, A. Bakker *et al.*, "Predictive value of simulated SAR and temperature for changes in measured temperature after phase-amplitude steering during locoregional hyperthermia treatments," *International Journal of Hyperthermia*, vol. 35, no. 1, pp. 330–339, 2018.

[179] R. W. Brown, Y.-c. N. Cheng, E. M. Haacke *et al.*, *Magnetic Resonance Imaging*, R. W. Brown, Y.-C. N. Cheng, E. M. Haacke *et al.*, Eds. Chichester, UK: John Wiley & Sons Ltd, apr 2014.

[180] E. P. van Horssen, D. Antunes, and M. Heemels, "Switched LQG control for linear systems with multiple sensing methods," *Automatica*, vol. 103, pp. 217–229, 2019.

[181] T. Köhler, P. Maass, P. Wust, and M. Seebass, "A fast algorithm to find optimal controls of multiantenna applicators in regional hyperthermia," *Physics in Medicine and Biology*, vol. 46, no. 9, pp. 2503–2514, 2001.

[182] K.-S. Cheng, V. Stakhursky, P. Stauffer *et al.*, "Online feedback focusing algorithm for hyperthermia cancer treatment," *International Journal of Hyperthermia*, vol. 23, no. 7, pp. 539–554, 2007.

[183] R. M. C. Mestrom, J. P. Van Engelen, M. C. Van Beurden *et al.*, "A refined eigenvalue-based optimization technique for hyperthermia treatment planning," *8th European Conference on Antennas and Propagation, EuCAP 2014*, no. EuCAP, pp. 2010–2013, 2014.

List of publications

Peer-reviewed journal articles

- D.A. Deenen, E. Maljaars, L.C. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Offset-free model predictive temperature control for ultrasound-based hyperthermia cancer treatments," *IEEE Transactions on Control Systems Technology*, in press.

- D.A. Deenen, E. Maljaars, B. de Jager, L.C. Sebeke, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Switched-Actuator Systems with Setup Times: Efficient Modeling, MPC, and Application to Hyperthermia Therapy," *Submitted for journal publication.*

- D.A. Deenen, L.C. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Target-conformal optimal actuator placement for ultrasound-mediated hyperthermia in cancer treatments," *In preparation for journal submission.*

- D.A. Deenen, J. van Wordragen, L.C. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Hierarchical mixed-integer MPC for real-time large-volume ultrasound hyperthermia in cancer therapy," *In preparation for journal submission.*

- L. Sebeke, D.A. Deenen, E. Maljaars, E. Heijman, B. de Jager, W.P.M.H. Heemels, and H. Grüll, "Model predictive control for MR-HIFU-mediated, uniform hyperthermia," *International Journal of Hyperthermia*, vol. 36, no. 1, pp. 1040-1050, 2019.

- L.C. Sebeke, P. Rademann, A.C. Maul, S.Y. Yeo, J.D. Castillo-Gómez, D.A. Deenen, P. Schmidt, B. de Jager, W.P.M.H. Heemels, H. Grüll, and E. Heijman, "Visualization of Thermal Washout due to Spatiotemporally Heterogeneous Perfusion in the Application of a Model-based Control Algorithm for MR-HIFU Mediated Hyperthermia," *Submitted for journal publication.*

- M.F. Heertjes, N. Irigoyen Perdiguero, and D.A. Deenen, "Robust control and data-driven tuning of a hybrid integrator-gain system with applications to wafer scanners," *International Journal of Adaptive Control and Signal Processing*, vol. 33, no. 2, pp. 371-387, 2019.

- D.A. Deenen, B. Sharif, S.J.A.M. van den Eijnden, H. Nijmeijer, W.P.M.H. Heemels, and M.F. Heertjes, "Projection-Based Integrators for Improved Motion Control: Formalization, Well-posedness and Stability of Hybrid Integrator-Gain Systems," *Provisionally accepted for journal publication.*

Peer-reviewed conference articles

- D.A. Deenen, E. Maljaars, L. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Mixed-integer model predictive control for large-area MR-HIFU hyperthermia in cancer therapy," *1st Virtual IFAC World Congress*, 2020.

- D.A. Deenen, E. Maljaars, L. Sebeke, B. de Jager, E. Heijman, H. Grüll, and W.P.M.H. Heemels, "Offset-free model predictive control for enhancing MR-HIFU hyperthermia in cancer treatment," *IFAC-PapersOnLine (6th IFAC Conference on Nonlinear Model Predictive Control (NMPC) 2018, Madison, Wisconsin, USA)*, vol. 51, no. 20, pp. 191-196, 2018.

- D.A. Deenen, M.F. Heertjes, W.P.M.H. Heemels, and H. Nijmeijer, "Hybrid integrator design for enhanced tracking in motion control," *2017 IEEE American Control Conference (ACC), Seattle, Washington, USA*, pp. 2863-2868, 2017.

- M.F. Heertjes, N. Irigoyen Perdiguero, and D.A. Deenen, "Data-driven tuning of a hybrid integrator-gain system," *IFAC-PapersOnLine (20th IFAC World Congress, Toulouse, France)*, vol. 50, no. 1, pp. 10851-10856, 2017.

Dankwoord

Deze pagina's vormen het laatste deel van mijn proefschrift, maar symboliseren ook de laatste fase van mijn tijd als promovendus. Terugkijkend naar de afgelopen vier jaar kan ik niets anders zeggen dan dat ik er enorm van heb genoten. Dit allemaal was natuurlijk niet mogelijk geweest, of met veel minder plezier, zonder de fantastische mensen om mij heen. Ik wil bij deze iedereen bedanken die op enige manier heeft bijgedragen aan mijn promotie of de weg hiernaartoe. Een aantal van jullie benoem ik graag in het bijzonder.

Als eerste wil ik mijn promotoren bedanken. Maurice, bedankt voor het vertrouwen dat je in mij legde toen je me de kans gaf om een promotie te starten. Jouw enthousiasme en optimisme werken aanstekelijk, waardoor hard werken en hard willen werken vaak gekoppeld waren in een positieve feedback loop, en je kritische blik en diepgaande kennis in combinatie met je fanatieke gebruik van een rode pen hebben zowel dit proefschrift als ook mijn onderzoekend vermogen naar een hoger niveau getild. Je was een inspirerende mentor op wetenschappelijk en menselijk vlak, maar wist dit altijd in een uiterst aangenaam amicaal jasje te verpakken. Bram, bedankt voor de geweldige begeleiding; jij wist je als de perfecte aanvulling op Maurice te positioneren. Met je bijna encyclopedisch brede kennis, je oog voor detail en je directe manier van vragen kreeg je vaak binnen enkele zinnen mijn blinde vlek gevonden, wat me regelmatig tot nieuwe inzichten heeft gebracht. Ook je speciale kijk op de wereld en unieke gevoel voor humor heb ik erg gewaardeerd, en waren een verrijking voor onze meetings.

Ik wil ook graag NWO en KWF bedanken voor het verenigen van hun krachten voor het Technology for Oncology programma, waar dit onderzoek via het Control-2-Act project deel van uitmaakte. Het raakvlak tussen technologische innovatie en medische applicatie is ongelofelijk belangrijk, en ik hoop dat de ontwikkelingen in dit proefschrift een steentje kunnen bijdragen in de strijd tegen kanker. I also wish to thank our Control-2-Act partner Erasmus MC Cancer Institute and all the members of the user committee for their support. In particular, I thank the Uniklinik Köln and Philips Research for granting me access to their MR-HIFU hyperthermia expertise and setup.

Next, I would like to thank everyone whom I have had the honor of collab-

orating with. Lukas, your work has laid the foundation on which I could build my research, and your help was indispensable for experimentally validating my control designs. Edwin and Holger, thank you for sharing your vast MR-HIFU experience and clinical perspective. Maarten and Iva, unfortunately our collaborative efforts have not led to a chapter in this thesis. Nevertheless, working with you has taught me much about MR-RF hyperthermia, for which hopefully this thesis can serve as a source of inspiration. Bert, bedankt voor de ontzettend fijne samenwerking. Juist toen ik nog wat zoekende was naar richting, was het jouw praktische insteek die een waardevol katalyserend effect had op de vaart in mijn onderzoek. Stefan en Tomas, als enige hyperthermie-control-promovendus was de korte samenwerking met jullie een welkome afwisseling. Sven, bedankt dat je deze onderzoekslijn doorzet, waardoor ik de afgelopen paar maanden niet meer de enige was. I also thank the (former) students Jeroen, Soufyan, Jiaheng, Bob, Jip, Ronnie, and in particular Karthik and Jesper for their efforts.

I would also like to express my gratitude to the other members of my PhD committee Daniel Limón, Bart De Schutter, and Mircea Lazar for reviewing this dissertation and taking part in my PhD defense ceremony.

To all my (former) CST and D&C colleagues, thank you for making the office more than just a place to work. It was your company that made me experience not only the group outings, but also the coffee breaks, conferences, and Benelux meetings as the enjoyable perks of pursuing a PhD. Ik benoem in het bijzonder mijn (oud-)kantoorgenoten van -1.115 Leroy, Ruud, Roeland en Annemiek. Bedankt voor een ongeëvenaarde werksfeer, de dartlesjes, en het helpen maximaliseren van gratis koffie. Ook bedank ik mijn oud-kantoorgenoten van 0.07a toen ik als bureauloze promovendus startte, waaronder mijn startdatum/startkantoordrielingbroers en Las Vegas-kamergenoten Roy en Enzo.

A special thanks also goes out to the HIGS team for allowing me to avoid any HIGS-related separation anxienty. Marcel, bedankt voor jouw ongeremde enthousiasme, positieve instelling en scherpe geest, welke me al sinds mijn afstudeerproject inspireren, en destijds ook over de streep hebben getrokken om te willen promoveren. Ik ben blij dat mijn keuze voor een ander promotieonderwerp niet heeft geleid tot een einde aan onze samenwerking. Sebastiaan and Bardia, even though our rooms were at opposite ends of a long hallway, I have truly enjoyed crashing your office every now and then for discussions about the HIGS (and the Hybrid Systems course). Maurice, ook aan dit onderwerp hebben we al sinds mijn afstudeerproject mogen samenwerken. Bedankt dat je me de vrijheid gaf om hieraan te blijven hobbyen naast mijn promotieonderzoek.

Nancy en Roos, bedankt dat ik altijd op jullie kon bouwen als er iets geregeld moest worden, voor het opzetten van de leuke groepsuitjes, en voor de gezelligheid wanneer ik jullie kantoor weer eens als wachtkamer gebruikte. Geertje, bedankt dat ik ook als "overloper" nog altijd bij jou terecht kon voor organisatorische hulp en een gezellig praatje bij de koffie.

Als volgende wil ik al mijn vrienden bedanken. Ze zeggen wel eens dat je

vrienden de familie zijn die je kiest, en dat is ook zeker hoe ik het ervaar. In het bijzonder bedank ik mijn Kroepoeksclan-genoten. Mede-pinda's Leroy en Robin, en adoptie-pinda Alex, ik voel me bevoorrecht dat ik de afgelopen jaren ook op het werk drie van mijn beste vrienden naast me had. Al zijn onze promoties officieel individuele trajecten, door jullie voelde ik me nooit alleen. Bedankt voor de late avonden vol zinnige en onzinnige discussies, het tolereren van mijn "logica", en het het mij laten winnen in de gym, maar van mij winnen met bordspellen. Ik ook ben mijn oud-huisgenoten, waaronder Niels, Sanne, Tobias en Vera, dankbaar voor het bieden van een vertrouwde thuishaven waar altijd wel iets te beleven viel, en vriendschappen die ook na het studentenhuis nog doorgroeien.

Ik wil ook mijn familie bedanken. Aan papa en mama, bedankt voor de onvoorwaardelijke liefde en aanmoediging sinds zo lang ik me kan herinneren. Jullie hebben me altijd de vrijheid gegeven voor het maken van mijn eigen keuzes, om deze vervolgens volledig te steunen, wat ongetwijfeld een essentiële rol heeft gespeeld in de persoon die ik ben en alles wat ik bereik. Martin en Patrick, we hebben (gelukkig) onze verschillen, maar lijken desondanks ontzettend op elkaar. Ik ben dankbaar dat ik hierdoor niet alleen met twee broers, maar ook met twee extra vrienden door het leven ga. Arie, Cisca en Jeroen, bedankt voor jullie warmte en betrokkenheid. Ik kan me geen betere bonusfamilie wensen.

Als laatste bedank ik jou, Patricia. Al meer dan elf jaar breng jij geluk in mijn leven. Je bent mijn partner, maar ook mijn beste vriendin. Jij zorgt ervoor dat ik niet vergeet te genieten wanneer ik weer eens te lang en te diep met mijn neus in de boeken zit, en je staat altijd voor me klaar nog voor ik me realiseer dat ik dat nodig heb, wat een onschatbaar positieve invloed heeft gehad op dit proefschrift en alles daarbuiten. Na al onze tijd samen ben ik nog elke dag blij dat we elkaar hebben gevonden, en kijk ik enthousiast uit naar wat de toekomst ons zal brengen.

Daniel Deenen
Eindhoven, 27 oktober 2020

About the author

Daniel Deenen was born on December 19, 1990, in Leonberg, Germany. After finishing his bilingual (Dutch and English) secondary education in 2008 (cum laude) at the Jan Tinbergen College in Roosendaal, the Netherlands, he studied Mechanical Engineering at the Eindhoven University of Technology, the Netherlands, from which in 2012 he received his Bachelor of Science degree. He continued studying Mechanical Engineering in the Dynamics and Control group at the Eindhoven University of Technology, receiving his Master of Science degree in 2016 (cum laude). As part of his master's program, he did an internship at the University of California San Diego, USA. During his master's thesis research, which was conducted at ASML, the Netherlands, under the supervision of Henk Nijmeijer and Marcel Heertjes, he developed a hybrid integrator-gain system for improved wafer stage motion control.

In 2016, Daniel started his PhD research in the Control Systems Technology group at the Department of Mechanical Engineering of the Eindhoven University of Technology under the supervision of Maurice Heemels and Bram de Jager, and in close collaboration with the University Hospital of Cologne, Germany, and the Erasmus MC Cancer Institute, the Netherlands. This work has been made possible by the Dutch Cancer Society (KWF) and the Dutch Research Council (NWO) as part of their joint Partnership Program "Technology for Oncology," and is partially financed by the PPP Allowance made available by Top Sector Life Sciences & Health. This research focused on model predictive control for magnetic-resonance-guided ultrasound-based hyperthermia in cancer treatment, and its main results are presented in this dissertation.