

On-line MRI guidance for radiotherapy

Sjoerd Crijs

Colophon

This text was set using the freely available L^AT_EX typesetting and text formatting system.

Cover: In order to create the cover picture I wrote wordQloud, based on the basic word layout principles introduced with Wordle [Viégas et al., 2009], the flexibility of ManiWordle [Koh et al., 2010] and the picture-based layout possibilities from Tagxedo [Leung, 2012]. After futzing around until things looked ‘good’ and worked ‘well’ the word layout on the cover was created in wordQloud, using words from the thesis.

Cover design	Sjoerd Crijns
Print	Ipskamp Drukkers
ISBN	978-90-393-5920-4

Copyright © Sjoerd Crijns 2013

The copyright of chapters 2, 3, 5 and 6 has been transferred to IOP Publishing Ltd.

This research is supported by the Dutch Technology Foundation STW, which is the applied science division of NWO, and the Technology Programme of the Ministry of Economic Affairs, Agriculture and Innovation (project OTP-07793). Financial support for publication of this thesis was kindly provided by Philips Healthcare, Elekta and Röntgenstichting Utrecht.

On-line MRI guidance for radiotherapy

On-line MRI geleiding voor radiotherapie

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de
Universiteit Utrecht
op gezag van de rector magnificus, prof. dr. G.J. van der Zwaan,
ingevolge het besluit van het college voor promoties
in het openbaar te verdedigen
op donderdag 21 maart 2013 des middags te 12.45 uur

door
Sjoerd Peter Maria Crijs
geboren op 8 december 1984 te Sittard

Promotor: Prof. dr. ir. J.J.W. Lagendijk
Co-promotor: Dr. B.W. Raaymakers

Contents

1	Introduction	1
2	A post-processing algorithm for on-line, real-time correction of field inhomogeneity induced MR image distortions	15
3	Distortion-free MR images: a fundamental acquisition solution	25
4	An MRI/linac interfacing framework	35
5	MRI-guided gated radiation delivery	49
6	MRI-guided tracked radiation delivery	59
7	Summary and discussion	69
8	Samenvatting	79
	References	83
	Publications	93
	Dankwoord	97
	Curriculum Vitae	101

Introduction

In radiotherapy, the aim is to eradicate malignant cells in a tumour by exposing them to a high dose of ionising radiation. Exposed cells suffer damage to their genetic material. In cases where cells have unreparable damage, the body will eliminate them. As a side effect, the radiation dose required to damage tumour cells is high enough to damage normal cells as well. Although most normal cells have a larger repair potential than tumour cells, the radiation dose to the surrounding organs (so called organs at risk (OARs)) must be limited as much as possible to reduce adverse effects.

Typically, the preparation for treatment of a patient is started by acquiring a pre-treatment computed tomography (CT) scan. This scan is used to visualise the anatomy, delineate relevant anatomical structures (i.e., tumour and OARs) and to define the material properties of the tissue needed for dose calculations [ICRU 50, 1993; ICRU 62, 1999]. Based on this delineation and a prescription, a treatment plan is derived that aims to deliver a designated dose to the tumour while keeping the dose to the OARs at a minimum.

A radiation treatment is typically not delivered at once. Fractionated radiotherapy can be used to allow the healthy tissue to recover between fractions of treatment. For example, in our department, prostate cancer is treated with 35 fractions of 2.2 Gy, spread out over seven weeks. Before delivery of each fraction the patient is positioned on the treatment couch, using in-room lasers and tattooed marks on the patient body, to exactly reproduce the state of the anatomy represented by the pre-treatment CT. Moreover, patient weight loss as well as tumour growth or shrinkage produce inevitable variations and constitute (together with patient positioning errors) motion on an inter-fraction scale. Organ motion (e.g, breathing motion, cardiac motion, bowel peristalsis, bladder filling, (in)voluntary muscle contractions) cause more uncertainties regarding shape and position of both target and OARs on an intra-fraction time-scale. Traditionally, all of these problems are overcome using the margin concept: the volume to be irradiated is derived by expanding the tumour volume with a margin of a few millimetres up to centimetres

in all directions, to ensure irradiation of the tumour volume with the prescribed dose. Margin sizes are derived from population statistics (e.g., van Herk et al. [2000]). Although the margin approach takes care of the under dosage problem, it results in irradiation of regions containing mainly healthy tissue at least part of the time. This increases the risk of complications and thus limits the dose that can safely be delivered to the tumour.

1.1 The use of imaging in radiotherapy

Image guidance plays a vital role in reducing geometrical uncertainties associated with tumour position and shape [Verellen et al., 2007]. Five different purposes of imaging related to radiation treatment can be distinguished: treatment volume definition and characterisation, set-up correction, treatment plan adaptation, response monitoring and intra-fraction motion compensation.

Treatment volume definition and characterisation

In modern radiotherapy, combined information from a myriad of imaging modalities can be used in the pre-treatment phase to improve the definition of the tumour volume, such as CT, positron emission tomography (PET), single photon emission tomography (SPECT), ultrasound and magnetic resonance imaging (MRI). These modalities may at the same time be employed to picture tumour tissue characteristics, like cell density, hypoxia and perfusion. Improved treatments of different kinds of malignancies may be implemented based on such information [Bentzen, 2005; Ling et al., 2000].

Periodic target and OAR motion may be quantified and characterised in the pre-treatment phase as well. For example, having such knowledge for lung tumours helps to derive a mid-ventilation position for planning and appropriate non-uniform margins that optimally account for organ motion [Hughes et al., 2008; Wolthaus et al., 2008]. A different strategy is to take motion into account via 4D planning strategies, as was proposed for example by Unkelbach and Oelfke [2004] and Keall et al. [2005]. Modalities for characterisations like these should provide information at a sufficiently high spatiotemporal resolution. Current studies are based on information from fluoroscopy [Tsunashima et al., 2004], 4D CT [Colgan et al., 2008] and cine-MRI [Feng et al., 2009; Gwynne et al., 2009; Koch et al., 2004].

The ideal imaging modality for treatment volume definition and characterisation can depict the tumour in high contrast to its surroundings. It is highly beneficial to have one modality that can picture a series of tumour characteristics. This decreases the patient burden of being subjected to many different examinations, shortens the characterisation process and decreases the possibility of registration errors. Moreover, this modality should be capable of producing 3D data sets at high spatial resolution to aid the delineation process and high temporal resolution to support the motion characterisation process.

Set-up correction

Imaging just before delivery of each fraction is used to verify patient set-up and, wherever necessary, correct for daily variation, to ensure a more accurate dose delivery. An increased delivery precision enables reduction of treatment margins and allows an increase of dose to the tumour volume, which respectively reduces treatment toxicity and increases tumour control probability. Protocols for set-up correction have been devised (e.g., Bel et al. [1993]; de Boer and Heijmen [2001] and variations on that theme), that need input from regular position measurements. Several in-room imaging modalities with which such measurements can be acquired are currently in clinical use. The most commonly used ones are discussed in this section.

The first available in-room imaging modality for set-up verification makes use of the treatment beam itself. Mounting a flat panel imager on the accelerator gantry, opposite to the MV source, enables acquisition of MV transmission images, which is called portal imaging (PI). A portal image has fairly low contrast due to the high energy of MV radiation. It is possible however to distinguish bones and implanted markers from soft tissues (see Figure 1.1a). Matching the bony anatomy [Nederveen et al., 2003; Van de Bunt et al., 2008] or markers [Van der Heide et al., 2007] on portal images to the same landmarks on pre-treatment images enables accurate set-up correction whenever their positions are closely correlated to the tumour position. In principle, PI can be used to perform volumetric imaging (MV cone-beam computed tomography (CBCT)) as well [Pouliot et al., 2005].

Imaging with kV radiation beams gives much more contrast compared to MV beams due to the lower energy. Two options exist to acquire volumetric CT data sets just prior to treatment. Uematsu et al. [1996] suggested an in-room CT scanner positioned at a 180° angle with respect to the treatment position, with a rotating treatment couch that brings the patient from CT to linac and vice versa. Alternatively, by suspending the scanner from a rail in the ceiling it can be brought to the patient [Kuriyama et al., 2003]. Instead of using an ordinary CT scanner, Jaffray et al. [2002] suggested mounting a kV source on the linac gantry, orthogonal to the MV source, in combination with a flat panel detector on the other side of the patient. Hence an on-board CBCT apparatus was realised. Although it produces data of lower quality than a free-standing CT scanner (cf. Figure 1.1b and c), integrated CBCT systems have a larger potential than in-room CT systems, as the precision gained from the latter suffers from motion induced by transferring the patient from imaging to treatment modality. In principle, in-room (CB)CT, like PI, enables set-up correction based on bony anatomy and markers, but the higher contrast facilitates use of soft tissue landmarks as well.

Alternative but less common modalities for set-up correction include ultrasound (which is mainly used for prostate [McNair et al., 2006; Mohan et al., 2000], although feasibility was also demonstrated for upper abdominal organs [Fuss et al., 2004]), plain radiographs and electromagnetic (EM) transponders.

In conclusion, the ideal guidance modality for set-up corrections acquires high resolution 3D images with sufficient soft-tissue contrast to distinguish the target. The modality must not only be in-room but integrated with the treatment modality




Figure 1.1: Four potential in-room imaging modalities for image guided prostate radiotherapy; (a) portal image showing pelvic bones and the implanted markers in the prostate, (b) pre-treatment CT, (c) CBCT acquired just prior to fraction delivery, (d) pre-treatment MRI as used in our department as a delineation guide. (a) courtesy of Alexis Kotte; (b)–(d) courtesy of Ellen Kerkhof.

as well, to minimise any possible changes of patient position.

Treatment plan adaptation

Instead of adapting the patient to the treatment plan to mimic the planned situation in the best possible way, the plan may be adapted to the patient instead. It is important to distinguish two paradigms for plan adaptations.

The information from the pre-treatment imaging sessions can be used to identify a number of different possible geometrical variations occurring for one patient. For example, the patient bladder may or may not be filled at the time of delivery of one fraction. A set of treatment plans is prepared for each of the possible anatomies. On-line imaging is then used to acquire images of the anatomy at the moment of treatment, in order to select the plan that treats the patient anatomy for that day in the best possible way. In this thesis, this paradigm is referred to as the ‘library of plans’ technique [Murthy et al., 2011]. The main advantage of the library of plans paradigm is that all treatment plans are prepared before the treatment. Hence, the established quality assurance (QA) process can be applied to all plans made ahead of treatment. One of the drawbacks, however, is that anatomical variations that are not captured in the library of plans cannot be accommodated.

This can be solved by interpolation between closest matching anatomical states [Meijer et al., 2012] or by creating (as part of the set of plans) a treatment plan that is robust to many anatomical variations.

Alternatively, instead of preparing a set of plans in advance, a plan may be generated just before the fraction delivery takes place [Ahunbay et al., 2009; Mohan et al., 2005]. In this thesis, this paradigm is referred to as the plan of the day approach. Although this technique brings high flexibility and adaptivity, it puts high demands on on-line imaging modalities and QA. Moreover, it requires robust (deformable) registration, in conjunction with dose accumulation tools, to make sure that the sum of all treatment plans administered to the patient matches the prescription. Once the plan of the day paradigm has been introduced, the timescale at which new plans are generated may be decreased to the intra-fraction level. This requires ultra-fast plan optimisation.

The optimal image guidance modality for treatment plan adaptations acquires images of the internal patient anatomy, with sufficiently high soft-tissue contrast to distinguish the tumour from its surroundings and to facilitate deformable registration with pre-treatment images. Ideally, the modality is able to re-acquire one of the pre-treatment data sets used in the planning phase.

Intra-fraction motion management

Radiation treatment of cancer of the abdominal organs such as kidney, pancreas, stomach and liver, is hampered by respiratory motion since it causes the target to make large semi-periodic excursions [Bortfeld et al., 2004; Kitamura et al., 2003; Van Sörnsen de Koste et al., 2006]. To deliver an adequate dose to such targets while at the same time sparing surrounding organs, the intra-fraction motion should be taken into account. Margins, that encompass the region where the tumour could potentially be, become very large. Instead of using concepts like the mid-ventilation approach or 4D planning as discussed above, intra-fraction motion compensation strategies should be adopted, like gated or tracked radiation delivery. Adequate on-line image guidance lies at the basis of the success of such strategies.

In motion compensated radiation delivery it is not uncommon to keep track of markers *outside* a patient's body and correlating their time-resolved position to internal organ motion to obtain the actual compensation signal. The process of establishing the relation between external and internal motion may be aided by motion/position models [Seppenwoolde et al., 2007]. A better correlation between motion of markers and anatomy to be tracked results in more reliable treatment. However, for some sites the correlation is poor or very difficult to establish and may vary during the treatment [Feng et al., 2009; Koch et al., 2004]. Hence, the best treatment is realised when the internal organ motion itself would be depicted directly [Minn et al., 2009].

In gated radiation delivery, the target position signal is employed to switch the treatment beam on when the target is in a favourable location, and off when the target moves away from this position. This strategy has proved to be useful for margin reduction since it diminishes target motion induced errors [Hugo et al., 2003]. Korreman et al. [2008] have shown that it must be applied under adequate

image guidance to obtain good results. Gated radiation delivery allows for dose escalation [Wagman et al., 2003] as well. Since the beam is switched off for a large portion of the treatment time, the technique is associated with patient throughput issues [Fox et al., 2007]. Recently it has been shown that the mean target position concept may also be incorporated in this strategy in order to increase the treatment efficiency [Guckenberger et al., 2011]. Voluntary or involuntary breath hold, as well as audiovisual feedback [George et al., 2006] may be applied to realise further optimisation. Since gated radiation delivery mainly relies on signals derived from markers outside the body, optical systems are mostly used for image guidance, for example as presented by Ford et al. [2002]; Wagman et al. [2003].

Currently, the ultimate way of intra-fraction motion compensation is to keep the radiation beam aligned with the moving target using dynamic motion compensation. This can be achieved, for example, with robotic linear accelerator control [Schweikard et al., 2000]. Motion of the treatment couch in opposite direction of the target is a possibility as well [D’Souza et al., 2005], although this is rather uncomfortable for the patient. Recently, tracked radiation delivery was realised with a gimbals based system [Depuydt et al., 2011; Kamino et al., 2006]. On multi-leaf collimator (MLC) based systems, the collimator aperture has to be adapted to the target position. *MV* imaging with the treatment beam can provide the required feedback signal [Poulsen et al., 2012]. Combining orthogonal *kV* and *MV* imaging facilitates compensations in 3D as well [Cho et al., 2009b, 2011]. *kV* imaging may be supplemented by information from optical systems [Cho et al., 2009a]. X-ray based imaging, just like in set-up correction must be used in conjunction with implanted radiopaque markers. Instead of using more radiation for imaging, wired [Ravkilde et al., 2011] or wireless [Sawant et al., 2009] EM transponders can be implanted in or near targets to get real-time 3D position feedback. Hsu et al. [2005] have studied the feasibility of using ultrasound target tracking as a non-invasive alternative.

The disadvantage of current intra-fraction motion compensation guidance modalities is that they are marker based; by tracking one or a couple of markers they still look at surrogates for the actual target motion. OARS may display uncorrelated motion and are not visualised. Moreover, marker based techniques hardly clear the way for intra-fraction adaptations for target deformations and are invasive and not suitable for all sites.

The ultimate image guidance modality for intra-fraction motion compensation acquires images of the internal patient anatomy, with sufficiently high soft-tissue contrast to distinguish the tumour from its surroundings and to detect, besides translations and rotations, deformations as well. To attain adequate delivery and compensation precision, these images must have a decent spatial resolution and give a complete picture to compensate in all three dimensions. Temporal resolution must be sufficiently high, to follow the process accounted for. Moreover, the imaging system must be tightly integrated with the treatment modality, to minimise latencies introduced by data transport and subsequent processing.

Treatment response monitoring

Imaging modalities are used in radiotherapy to assess whether the body responds to the treatment, and if so, in what way. The requirements are almost the same as in the tumour characterisation phase. Images have to be acquired from which tissue characteristics, like cell density, hypoxia and perfusion can be derived.

Additionally, treatment response monitoring requires accurate dose accumulation. Both inter- and intra-fraction variations should be taken into account, and not only rigid transformations, but especially organ deformations. QA in this case is a complicated task, but a key requirement as well.

A good treatment response imaging modality provides a way to record a series of tumour characteristics. Ideally, a data set is acquired after each fraction, and during a follow-up period as well, without an additional radiation burden for the patient. The recorded data must represent absolute tissue characteristics since it is impossible to find day-to-day variations otherwise. Moreover the data has to be registered to accumulated dose maps. These maps should be derived from the treatment plans, in combination with adequate on-line imaging to account for intra-fraction motion.

1.2 Use of MRI in radiotherapy

Compared to commonly used imaging modalities in radiotherapy, MRI provides superior soft-tissue contrast and thus a clear view of both tumorous tissue and OARs (see Figure 1.1). Magnetic resonance (MR) imaging sequences have numerous parameters, that can be varied to obtain a multitude of contrasts reflecting tissue characteristics, like tissue oxygenation, cell density and tissue perfusion, which in turn reflect characteristic parameters of a tumour.

Both volumetric/3D and 2D data sets can be obtained in arbitrary orientations. Moreover, cine-MRI sequences exist that are able to depict the patient anatomy at a rate of multiple frames per second, satisfying the criteria for a good intra-fraction motion compensation guidance modality. MR imaging uses no ionising radiation, such that imaging can be repeated without increasing the patient radiation burden. MRI is therefore very suitable not only to characterise the tumour in advance of the radiation treatment but during and after treatment as well in order to obtain treatment response characteristics. These properties make MRI a prime image guidance candidate for all purposes stated in the previous section.

On-line MRI-guidance is not yet clinically available, however. Karlsson et al. [2009] have proposed a next door MRI scanner for the radiotherapy clinic; a 1.5 T MRI system was installed close to a linear accelerator and patients were transported on a custom made trolley from the treatment couch to the imaging device and back. A more sophisticated rail-mounted MRI system that moves in and out of the operating room has been introduced to the neurosurgery field by IMRIS Inc. (Winnipeg, Canada). Although having an in-room MRI scanner is an attractive concept and a step in the right direction, the arguments against in-room CT solutions apply here as well. Moreover, these systems do not provide feedback during radiation delivery, necessitating the use of yet other modalities such as discussed in the

previous section. It was therefore proposed to integrate a linear accelerator for radiotherapy with an MRI scanner [Legendijk et al., 2008].

MRI-linac design

Three different designs for integrated MRI/radiotherapy devices have been proposed that recently reached at least a prototype status.

The Cross Cancer Institute (Edmonton, Alberta, Canada) team have proposed a design composed of a linac and a low field MRI magnet. Their proposed system consists of a 6 MV linear accelerator for radiation production, which is situated on the open end of a bi-planar 0.5 T MRI magnet. Flat gradient coils are placed on each magnet pole, facing the subject. In this design, both the MR system and the linac are carried on a gantry that is able to rotate the entire system around the cranio-caudal patient axis. In principle this facilitates both transverse and longitudinal configurations (Figure 1.2). A fully functional transverse configura-




Figure 1.2: The Cross Cancer Institute MRI-linac design; (a) transverse configuration (radiation travels perpendicular to the magnetic field lines). (b) longitudinal configuration (radiation travels along the magnetic field lines). Courtesy of the Cross Cancer Institute.

tion prototype, with an 0.2 T permanent magnet and a bore size large enough to encompass a human head, was built and tested in December 2008 [Fallone et al., 2009].

ViewRay (Oakwood Village, Ohio) have proposed a design based on 3 ^{60}Co sources for radiation production, which are mounted on a circular gantry (Figure 1.3, Dempsey et al. [2005]). The ring-based gantry, that facilitates rotation over a limited range of angles (140°), rotates around a 0.32 T full-body cylindrical MRI magnet. The full magnet system is split in two halves to enable radiation beam passage; both the main magnet coils and the gradient coils are composed of two separate pieces [Low et al., 2011]. The ViewRay magnet makes use of passive shielding to reduce the magnitude of the fringe field away from the magnet. Since the design does not incorporate a linear accelerator, it seems that there is hardly any need to worry about interference of the radiation producing chain with the magnetic field and no adaptations have to be made to the magnet specifically for

this purpose. However, the ^{60}Co sources may become magnetic in the magnet fringe field, thereby compromising the homogeneity of the main magnetic field and the geometric fidelity of the acquired MR images. ViewRay built a working




Figure 1.3: Rendered image of the ViewRay design. Note the ring in the centre that carries the ^{60}Co source vaults (two are visible). Image courtesy of ViewRay Inc..

prototype over 2010 and 2011, receiving FDA 510(k) clearance for their apparatus in the spring of 2012.

The research described in this thesis centres on the Utrecht MRI-linac design. This concept features a $\sim 30\text{ cm}$, 6 MV standing wave linear accelerator (Elekta AB, Stockholm, Sweden) for radiation production and a 1.5 T cylindrical MRI scanner (Achieva, Philips, The Netherlands). The magnet is not split and hence the radiation must go through its cryostat. The amount of material in this part of the cryostat is therefore minimised; the central 15 cm are free of coils. This gap allows a maximum irradiation field of 24 cm in the head-feet direction at the magnet isocentre. The gradient coil is split in two halves to allow beam passage. Active shielding is used, not only to suppress the magnet fringe field, but to create a toroidal low magnetic field zone around the magnet as well. The linear accelerator is positioned such that its electron gun end is always in this field void, to minimise Lorentz forces on the accelerated electrons. Radio frequency (RF) shielding of the MRI scanner from the noise produced by the accelerator is accomplished by attaching two copper plated Faraday cages to it. A first prototype in a static configuration and using cerrosafe blocks for beam collimation was constructed in 2009, by which the proof of concept for this design was delivered (see Figure 1.4 and Raaymakers et al. [2009]).

Over 2011 a drum and base type gantry was designed to create a proof of concept system suitable for intensity modulated radiotherapy (IMRT) type treatment. This gantry carries the whole chain of radiation producing components, including a 160 leaf MLC, as well as the controller electronics and cooling equipment (see Figure 1.5). Slip ring technology is employed to transfer the power needed for radiation production, as well as interlock signals, from the static world onto the




Figure 1.4: The first Utrecht MRI-linac prototype. (a) rendered CAD drawing showing the MR scanner and the linac positioned in the zero magnetic field region; (b) photograph of the prototype showing the accelerator components on a wooden table and copper-plated Faraday cages around the magnet.

rotating system. Communication with on-board electronics is carried out by way of wireless ethernet. The current gantry prototype facilitates continuous rotation at a maximum speed of 10 *rpm*.




Figure 1.5: The current MRI-linac prototype built in Utrecht; (a) rendered CAD drawing with a human figure to be able to appreciate the size of the machine and (b) photograph of the prototype, showing the gantry around the magnet cryostat.

1.3 Challenges

A few interesting challenges must be solved for the MRI-linac concept to be adopted in routine clinical practise. Obviously there are many clinical challenges, since the machine is thought to bring many new possibilities compared to conventional linear

accelerators. On the technical side, treatment plan adaptation and planning in magnetic fields, as well as QA, represent complicated issues that have to be dealt with.

Solutions to two technical challenges are investigated in this thesis, which are introduced in more detail below.

Geometrical accuracy

One of the aims of the MRI-linac is to introduce the stereotactic precision used in radiotherapy of brain tumours to body applications. This would make margin reduction a possibility, which in turn can facilitate hypo fractionation as well as dose escalation. Moreover, having stereotactic precision for body applications may lead to new targets for radiotherapy, that require high precision therapy delivery and small margins. Considering the intended use of an MRI-linac and the currently established levels of accuracy of radiotherapy hardware, a targeting accuracy on the order of 0.5 mm was specified for the MRI-linac.

MRI has since long been associated with inadequate geometric fidelity, especially in applications that require an accuracy of 1 mm or smaller. An MRI examination begins by placing the subject to be imaged in a homogeneous static magnetic field \mathbf{B}_0 . As a result the nuclear spins in the object start to precess around this field with a specific frequency, called the Larmor frequency. This frequency depends on the magnetic field strength B_0 according to $\omega = \gamma B_0$, where γ denotes the gyromagnetic ratio (which for protons is given as $\gamma = 2.675 \cdot 10^8\text{ rad T}^{-1}\text{s}^{-1}$). By precessing, the nuclear spins generate an electromagnetic signal at the Larmor frequency, that can be recorded using coils. Moreover, gradient magnetic fields along the three axis of the laboratory coordinate system (\mathbf{G}_x , \mathbf{G}_y and \mathbf{G}_z) are used, that produce linear variations of the magnetic field magnitude with the spatial coordinate. These gradient magnetic fields lie at the basis of spatial localisation of the signals originating from the subject; by applying gradient magnetic fields, spins at varying locations precess at varying frequencies¹.

The spatial localisation scheme employed becomes erroneous as soon as the assumption of spatial linearity of the gradient magnetic fields is incorrect. Moreover, gradient magnetic fields originating from other sources than the gradient coils themselves, such as inhomogeneity of the main magnetic field \mathbf{B}_0 , disturb the spatial localisation as well. Ironically, inhomogeneity of the main magnetic field can be induced by the subject; the non-constant magnetic susceptibility in a subject perturbs the field, which makes this a subject-specific problem.

These discrepancies from the ideal system behaviour present themselves as image distortions. As a special case, a global offset of the magnetic field results in a shift of the image, thus affecting the positioning of the image in the laboratory coordinate frame. The more localised distortions cause local image warps. The severity of the distortions depends not only on the magnitude of the magnetic field disturbances, but on the imaging sequence and its parameters as well. Fortunately the errors due to gradient field non-linearity can be accounted for by careful calibra-

¹For a more detailed explanation regarding the theory of MRI physics, the reader is referred to Haacke et al. [1999] and references therein.

tion [Doran et al., 2005; Moerland, 1996] for most cases (this subject is revisited in Chapter 7).

This leaves us with the image distortions due to field inhomogeneity. When one needs to rely on the images for the correctness of interventional procedures, all systematic errors due to image distortion must be ruled out. More particular, in light of the design criteria and anticipated operating condition for an MRI-linac, this challenge has to be overcome.

Interfacing

Radiation treatment machines tend to become increasingly complex, with add-ons for each new treatment option. In our institute, *MV* imaging, *kV* imaging, the oncology information system, position verification tools and machine control interfaces each run at their own computer system, with minimal automated interaction between the systems.

From an operator perspective, access through one PC station reduces the number of steps required to perform a certain task and thereby promotes smooth and efficient work flows. Most importantly, having integrated systems reduces possibilities of errors due to manual transfer of data from one system to the next. As an important example, due to the constraint of the narrow magnet bore, set-up correction already requires a daily re-planning step based on daily MR scans. A smooth work flow is key for this case, which is aided by presenting the user with one common interface to the subsystems.

An MRI-linac is believed to add new and improved treatment strategies to the therapy toolbox as well. Gated and tracked radiation delivery techniques require the imaging and treatment modalities to be merged. This minimises latency due to transfer of large amounts of data (multiple images per second) of one subsystem to the other, thereby improving the machine performance and delivery precision. A tight integration is therefore needed between imaging and treatment systems in an MRI-linac.

On the other hand, the MRI-linac is a machine that will be under constant development for the next few years. Although from an operator and performance perspective the subsystems must be presented as one system with one common look and feel, the design must be flexible as well.

The design and implementation of a software framework capable of interfacing between MRI scanner and linear accelerator, while providing the flexibility required for development, experimentation and growth, is a major challenge that must be solved to make the MRI-linac to a success.

1.4 Thesis outline

The purpose of this thesis is to present strategies that deal with the geometrical accuracy challenge. Then, taking geometrically accurate images as a starting point, the interfacing challenge is met and the technical feasibility of using MRI feedback for intra-fraction motion compensation is demonstrated.

Chapter 2 describes a post-processing algorithm that corrects image distortions induced by magnetic field inhomogeneity based on field measurements. The implementation of this algorithm is validated based on phantom measurements. This technique is subsequently used to obtain a distortion corrected body contour of the pelvic region for prostate radiotherapy treatment planning.

Chapter 3 presents pure phase encoding as a fundamental solution to the image distortion problem. Moreover, the inherently long acquisition times in this technique are mitigated by under-sampling and regularised reconstruction. To illustrate the performance of this acquisition in the presence of field inhomogeneity phenomena, images obtained with canonical spin-echo and pure phase encoding techniques of phantoms and healthy volunteers are compared.

In Chapter 4 an MRI/linac interfacing framework is introduced that was implemented in order to overcome the interfacing challenge. It includes code for low latency data communication with the MR scanner, as well as controllers and controller clients to manipulate the gantry, the radiation beam and the MLC in a dynamic way. Next, we take a closer look at intra-fraction motion compensation in radiotherapy using on-line MRI-guidance.

In Chapter 5, the feasibility of gated radiation delivery based on MRI-guidance is demonstrated. The linear accelerator is modified to allow beam on/off during delivery. Real-time communication with the MR scanner and control of the radiation beam are established in an MRI/linac interfacing framework. The position of a phantom subjected to one-dimensional periodic translation is tracked with the MR scanner. Based on the time-resolved position of the phantom, gated radiation delivery to the phantom is realised. Dose distributions for dynamic delivery conditions with varying gating windows are recorded on GafChromic film. The gating implementation is validated by comparing measured dose profiles to theoretical profiles calculated using knowledge of the imposed motion pattern. At the same time, this illustrates that real-time on-line reconstruction of the actual accumulated dose can be performed using time-resolved target position information.

In Chapter 6, the feasibility of tracked radiation delivery based on MRI-guidance is demonstrated in a set-up analogue to that used in Chapter 5. Without motion a sharply defined dose distribution is obtained, whereas considerable blur occurs for delivery to a moving phantom. When synchronising the MLC aperture to the phantom motion derived from fast 1D MR images, the sharpness of the dose distribution is nearly restored. Some residual blur remains that can be explained in terms of mechanical latency of the MLC. Knowledge of the relative motion between target and aperture again enables accurate reconstruction of the delivered dose distribution.

A post-processing algorithm for on-line, real-time correction of field inhomogeneity induced MR image distortions

This chapter is based on:

S.P.M. Crijs, B.W. Raaymakers and J.J.W. Lagendijk. Real-time correction of magnetic field inhomogeneity induced image distortions for MRI-guided conventional and proton radiotherapy *Physics in Medicine and Biology* 2011 **56**(1):289–297.

2.1 Introduction

In radiotherapy, image guidance decreases the geometrical uncertainties associated with tumour position and shape [Verellen et al., 2007]. This enables simultaneous reduction of treatment margins and increase of dose (e.g., with a local boost), thereby enhancing tumour control probability and lowering normal tissue complication probability. MRI is a prime candidate for image guidance since it provides superior soft-tissue contrast and thus a clear view of both tumorous tissue and OARs.

Lagendijk et al. [2008] have presented the benefit of using MRI, not only as a position verification system, but as an on-line treatment guidance modality. On-line MR images enable development of new treatment strategies, even for patient groups that were not eligible for radiotherapy so far (e.g., Kerkhof et al. [2011]). Moreover, cine-MRI enables real-time tracking or characterisation of the motion of a 3D volume of interest. Recently, we have reported on a prototype integrated MRI-linac that aims to implement such MR guidance. Simultaneous irradiation and MR imaging of a phantom was achieved without loss of image quality or accelerator malfunction [Raaymakers et al., 2009]. Other designs have been proposed that also reached the prototyping phase [Dempsey et al., 2005; Fallone et al., 2009]. Integrated MRI scanners can be used to guide patient irradiation, provided that the images produced are a faithful representation of the patient anatomy [Lagendijk et al., 2008].

Image distortions induced by magnetic field inhomogeneities may pose a problem for direct application of MR images in radiotherapy [Moerland et al., 1995]. A field deviation of only parts per million of the nominal field value potentially causes image distortions [Haacke et al., 1999]. These image distortions subsequently cause targeting inaccuracies when images are used as a basis for treatment. Besides a static (system-specific) field inhomogeneity, the main source of magnetic field deviations is the susceptibility of the patient itself, thereby making this a patient specific problem. Using strong imaging gradients, which is a feasible solution at lower field strengths, results in poor signal-to-noise ratios (SNRs) and imposes restrictions on imaging sequences at 1.5 T [Wachowicz et al., 2010].

It has been shown that distortions caused by system specific (i.e., hardware related) magnetic field inhomogeneities and gradient non-linearities can be greatly reduced by calibrating for their influence [Doran et al., 2005; Moerland, 1996]. Several post-processing approaches have been proposed to deal with patient specific field-inhomogeneity induced image distortions. Most of these methods are based on magnetic field maps (e.g., Baldwin et al. [2009]; Jezard and Balaban [1995]; Moerland et al. [1995]; Munger et al. [2000]). Although good results have been obtained with these methods, their speed is inadequate for on-line, let alone real-time, use in MRI-guided radiotherapy. Specifically, Baldwin et al. [2009] have reported a processing time of 30 minutes to several hours.

The current work concerns a validated implementation of a fast image-based distortion correction method that represents the steps taken to come to nearly real-time MR-based feedback during irradiation. In the following the procedure is validated with a phantom, where the correction performance is quantified based on landmark positions extracted from the images. Then the technique is applied

to body contour extraction from MR images of the pelvic anatomy to demonstrate its use in vivo. Finally, we discuss implications for targeting and dose accuracy of MR-guided conventional and proton radiotherapy.

2.2 Materials and Methods

Mapping image distortions

For all gradient recalled echo (GRE) and spin echo (SE) sequences where one line of k-space is filled after each excitation, an object at point (x, y) that experiences a field inhomogeneity $\Delta B_0(x, y)$ is depicted in the image at

$$x' = x - \frac{\Delta B_0(x, y)}{G_x}. \quad (2.1)$$

Here x denotes the readout direction and G_x is the readout gradient strength. (2.1) provides the relation between a measured field inhomogeneity and local (in-slice) image distortion. Hence, when $\Delta B_0(x', y)$ is known, the amount of distortion in each image pixel is known and the image can in principle be corrected.

The correction method as proposed by Jezzard and Balaban [1995] is adopted to first map the image distortions and subsequently correct them. A magnetic field map is obtained from two GRE images, where the signal phase at the centre of an echo is given by

$$\phi = \phi_0 + \gamma \Delta B_0 T_E.$$

Here ϕ_0 denotes the undisturbed phase (i.e., the phase in an ideal situation without any field inhomogeneities), γ is the gyromagnetic ratio, ΔB_0 is the field inhomogeneity and T_E is the echo time. The echo time is varied from one image to the other, such that the field inhomogeneity can be obtained from their phases according to

$$\Delta B_0 = \frac{\phi_2 - \phi_1}{\gamma(T_{E,2} - T_{E,1})}. \quad (2.2)$$

The phase as obtained from the MR image reconstruction process is wrapped in the interval $[-\pi, \pi]$ and must be unwrapped before further use.

Image processing steps

The following steps are taken to correct the distortions that occurred in an MR image.

1. A field map of the object is acquired using the echo time difference method (2.2).
2. The field map is unwrapped using an in-house C++ implementation of Jenkinson's widely used algorithm [Jenkinson, 2003]. The algorithm divides the map into regions which do not contain any wraps. An error function that penalises phase differences between regions is subsequently minimised by adding integer multiples of 2π to each region, thereby creating an unwrapped

phase image. Regions of low signal intensity (i.e., air or bone) are filled in by linear interpolation to guarantee continuity.

3. Unwarping is carried out on a per-row (or per-column) basis using 1D linear interpolation according to the inverse of Equation (2.1). Since a coordinate transformation is applied (pixels may be stretched or compressed), the Jacobian of this transformation is used for local intensity corrections [Doran et al., 2005].

All image processing is done using in-house developed C++ programs on a PC running Ubuntu Linux, equipped with an Intel®Core™2 Duo 3.0 GHz processor.

Experiments

First, the correction method is validated and quantitative measures of its performance are obtained. MR images are acquired of a phantom using a GRE sequence with a particularly weak read-out gradient ($T_E/T_R = 15/50$ ms, flip angle 25° , field of view (FOV) 195×195 mm², image matrix 128×128 pixels, readout gradient 0.93 mT/m, ΔT_E (for field map) 4.6 ms). The phantom is a plastic box measuring $150 \times 150 \times 30$ mm³ filled with water, containing a regular rectangular plastic two-dimensional grid (Figure 2.1). It is positioned in the coronal plane at a 45° angle with respect to the feet-head direction. Locations of the central 64 grid points are determined manually before and after correction and compared to the expected grid point locations to assess the deviations and obtain performance measures.

Next we demonstrate the use of this method *in vivo*. In our department, MR images are already used during treatment planning of prostate cancer. Accurate delineation of the tumour as well as surrounding structures is achieved based on various MR contrasts [Groenendaal et al., 2010] to allow for dose painting (i.e., tumour boosting). For planning purposes, we still rely on CT for geometrically accurate body contours; body contours are situated at the edges of the field of view where field inhomogeneities effects are largest. As an *in vivo* application, the correction scheme is used to extract geometrically accurate body contours for treatment planning of the prostate region from a distorted MR image. Two spin-echo images are acquired of a healthy volunteer ($T_E/T_R = 25/400$ ms, flip angle 90° , FOV 400×400 mm² in the transverse plane, image matrix 192×192 pixels, acquisition time 76.8s, phase encoding direction AP). The readout gradient used is 1.0 mT/m. As a ground-truth for comparison a second SE image is obtained with the same imaging parameters as before, but with a read-out gradient of 15.0 mT/m. A field map is acquired using a GRE sequence ($T_E/T_R = 15/30$ ms, flip angle 20° , FOV 400×400 mm², image matrix 192×192 pixels, ΔT_E 4.6 ms). Body contours are obtained via an adaptive thresholding procedure and subsequent edge detection with a Prewitt filter. Edges inside and outside the body contour are discarded.

2.3 Results and discussion

The phantom is examined with CT and MR. Figure 2.1 shows the images before and after correction.

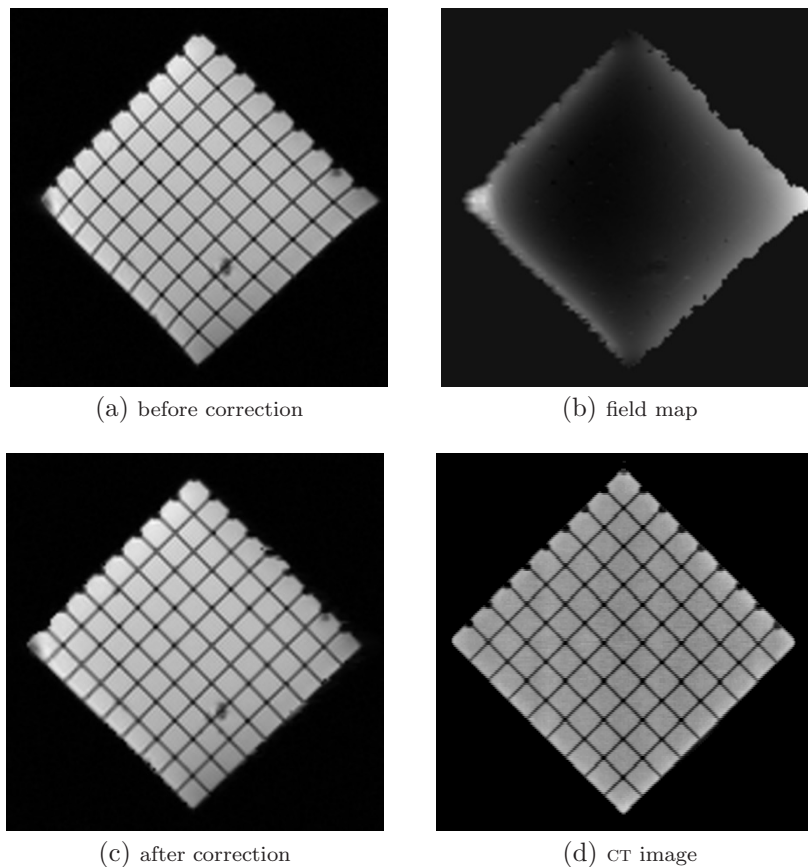


Figure 2.1: (a)–(c) MR images of the phantom with the measured (unwrapped) field map used for correction. (d) Reference CT image of the phantom.

Results for the deviations of obtained grid point locations with respect to their real coordinates are shown in Figure 2.2. The average difference found before correction is 0.65 mm ; after correction this is reduced to 0.30 mm . The distributions of differences reveal a maximum of 2.9 mm without correction, while a maximum displacement of 0.63 mm is obtained after correction (figures 2.2b and (d)). Near the left and right corners of the object, the deformations are largest, as can be observed from Figure 2.2a, but also there the image is corrected properly. The residual error in the corrected images partly arises from uncertainties associated with point clicking, since the acquired voxel size is relatively large ($1.52 \times 1.52\text{ mm}^2$ compared to grid plastic thickness of 1 mm).

The results for the body contour extraction experiment are displayed in Figure




Figure 2.2: Deviations of grid positions (in mm) in the phantom before (top row) and after (bottom row) correction.

2.3. From the displacement map (Figure 2.3b, derived from Equation (2.1)) it is deduced that distortions in the image taken at high gradient strength are within 0.5 mm . This image is therefore taken as a ground truth to compare the image taken at low gradient strength before and after correction. The same field map shows that before correction of the low gradient image, distortions are as large as 6 mm , which is well in agreement with the contours observed in Figure 2.3a (denoted by the arrows). Discrepancies in AP direction (the phase encoding direction, where no field-inhomogeneity related errors are present) do occur but are very small and can be attributed to slight movements of the volunteer. After correction, the differences as compared to the ground truth contour are reduced to a voxel at most.

All results discussed here are provided within a lag time of 8 ms after acquisition of the field map. Nonetheless, the performance of our implementation in terms of correction accuracy is comparable to previously established measures (cf. Baldwin et al. [2009]). It is therefore considered a viable tool for use in interven-




Figure 2.3: Pelvic contours for the prostate region: contour subtraction images. (a) truth - distorted; (c) truth - corrected. (b) shows the displacement map (in mm) used for distortion correction; the left scale is for acquisition with a read out gradient of $1.0\ mT/m$ whereas the right scale is for acquisition with $15.0\ mT/m$.

tion settings, such as that of MRI-guided radiotherapy, where ideally plans should be changed directly after obtaining new images. Sequences used in this study have acquisition times too long for use in these settings and are therefore subject of further investigations. Note that the time for phase unwrapping mainly depends on the complexity of the phase map, while the time for the other steps scales linearly with image size. For larger images, the dewarping step is easily parallelised in case an application would demand even higher speed, since the distortion is in 1D. On the other hand, solving the problem on the graphics processing unit (GPU) could be a tractable approach. The linear interpolation step to dewarp the image can be implemented extremely efficiently using GPU texture-lookup operations. GPU implementations of phase unwrapping algorithms have been presented by several authors [Gao et al., 2009; Karasev et al., 2007; Mistry et al., 2009].

Accuracy of treatment starts with hitting a target as planned. The targeting accuracy in MR-guided photon therapy is compromised if the depiction of the treatment situation is different from the actual treatment situation due to field inhomogeneity effects. This is especially important for small targets in heterogeneous parts of the anatomy. For instance, in head and neck radiotherapy targets of few cubic centimetres may be treated with small beams. As can be observed from Figure 2.4b the tissue/air/bone transitions here cause much field inhomogeneity and thereby image distortions that have to be dealt with before use in treatment guidance.

Figure 2.4b also implies that when using a low gradient strength of $1\ mT/m$, distortions are as large as $9\ mm$; to yield an accuracy of $1\ mm$ a gradient strength of $15\ mT/m$ must be used. For small beams as used in this area, a targeting inaccuracy readily causes under dosage of the target and over dosage of surrounding OARS.

It should be pointed out that in practise, correcting an image for distortions may not always be a necessity. For example, the displacement map in Figure 2.3b indicates that, at the higher gradient strength (which for this sequence was feasible), image distortions are only as large as $0.4\ mm$, which we find very acceptable for our application. For the field differences depicted in Figure 2.4b this accuracy




Figure 2.4: (a) CT of the head and neck area and (b) the magnetic field distribution in *ppm* of the nominal value as calculated analog to Marques and Bowtell [2005].

would only be obtained when working with gradients as large as 38 mT/m (at 1.5 T) which is only just feasible with the usual imaging hardware. In cases where geometrical accuracy is adequate, acquisition of the field map (which took approximately 15 s in the experiments presented here) may then be omitted. Careful investigation for the need for image correction on an application and sequence specific basis is thus recommended in cases where geometrical accuracy is of importance. The tool we have presented here may be used in such an assessment.

Surface coils (as used in most MRI examinations) exert some pressure on the patient thereby imposing small deformations on the body contours relative to the treatment situation. Even if correction algorithms are employed, pre-treatment MR images should thus be used cautiously for planning (especially in proton treatment). For the integrated MRI-linac this is of course different. Since the aim here is to acquire images on-line and deliver radiation through the receiver coils, corrected MR images can be used for on-line decision making regarding treatment plans.

The method as presented here is very suitable for position verification of a static part of the anatomy in radiotherapy with on-line MR guidance. The field map is reliable provided that the anatomy under observation is not moving. However, slight movements are very likely to occur in treatment situations, as are movements due to breathing. In general, these lead to inaccurate magnetic field maps, since the pixel-to-pixel correspondence of the two images used in the field mapping process is no longer valid. The consistency of the field map may be improved by the use of image registration techniques in such cases. On the other hand, acquiring MR images of moving parts of the anatomy requires use of short repetition times and thereby strong gradients, which already minimises images distortions. Correction of MR image distortions in images of more mobile and deformable parts of the anatomy are the subject of further investigations.

2.4 Conclusion

A method is demonstrated that can provide corrections to MR images distorted by field inhomogeneity effects. It yields results comparable to previously established methods, but is far less time consuming since it needs only a few milliseconds of processing time. Body contour distortions can be accurately corrected as well. Our results open new possibilities to investigate on-line MRI-based plan adaptations and ultimately MRI-based treatment planning based on images from an MRI-linac.

Distortion-free MR images: a fundamental acquisition solution

This chapter is based on:

S.P.M. Crijs, C.J.G. Bakker, P.R. Seevinck, H. de Leeuw, J.J.W. Lagendijk, B.W. Raaymakers. Towards inherently distortion-free MR images for image guided radiotherapy on an MRI-accelerator *Physics in Medicine and Biology* 2012 **57**(2):289–297.

3.1 Introduction

In MR-guided interventions, it is necessary to establish a solid relationship between the imaging coordinate system and world coordinates. This is generally crucial in MR applications that put high demands on geometric accuracy such as biopsies, stereotactic neurosurgery and multi-modal imaging. As radiation cannot be visualised on MR images it is particularly important in on-line MR-guided radiotherapy. Legendijk et al. [2008] discussed the benefit of using MRI in radiotherapy for both position verification based on soft tissue contrast and as an on-line treatment guidance modality. On-line MR images provide superior soft-tissue contrast compared to conventional X-ray based imaging modalities and enable development of new treatment strategies, even for patient groups that were not eligible for radiotherapy so far (e.g., renal cancer patients [Kerkhof et al., 2011]). Therefore integration of an MR scanner and a linear accelerator is pursued. Recently, we have reported on a prototype integrated MRI-linac and demonstrated good image quality irrespective of the presence of the photon beam [Raaymakers et al., 2009]. To guarantee correct treatment based on on-line MR images, a rigorous relation must exist between the treatment and imaging coordinate systems [Moerland et al., 1995]. Discrepancies between the two are generally caused by non-linearity of the gradient magnetic fields, that are used for position encoding and local off-resonances [Haacke et al., 1999]. The former effect can be effectively dealt with by careful calibration [Doran et al., 2005; Moerland, 1996]. When frequency encoding is used (as is done conventionally), off-resonance effects like field inhomogeneities, chemical shift differences, influence each point in k-space along the frequency encoding direction in a different way, which manifests itself as image distortion [Reichenbach et al., 1997].

Many strategies have been developed to mitigate image distortions. Since the distortions are inversely proportional to the strength of the encoding gradient used, a simple solution is to make these gradients as strong as is needed to keep distortions within the dimensions of a voxel [Haacke et al., 1999]. This approach can only partially resolve the problems and leads to lower SNRs. Retrospective compensation strategies that measure the distortion and then correct for it have been developed as well [Baldwin et al., 2009; Chang and Fitzpatrick, 1992; Crijns et al., 2011a; Jezzard and Balaban, 1995]. Moreover, some methods exist to fully counteract the distortion for some specific application or known field inhomogeneity, by employing compensating field gradients [Cho et al., 1988; Glover, 1999]. Although they can be effective, these retrospective compensation methods require prior knowledge of the distortion or measurement of the underlying effects and are only effective for distortions of limited magnitude.

In this study, a method is presented that is used in solid-state imaging and MR microscopy that inherently produces correct images, where the signal is encoded exclusively with phase encoding. This method is called single point imaging (SPI) and it has been used in medicine related applications, mainly to study (quasi-)solid tissue samples [Gruwel et al., 2007; Seo et al., 1998] and for imaging prosthetic materials [Ramos-Cabrer et al., 2004]. The SPI method has not been used for *in vivo* imaging experiments so far, mainly since the associated acquisition times are relatively long. Recently, the compressed sensing (CS) paradigm has been used

to perform MR image reconstruction from a randomly under-sampled data set [Lustig et al., 2007]. The principles of CS are well-suited for application in SPI since the under-sampling can be performed here on a point-by-point basis. We therefore propose to apply this paradigm in reconstruction from under-sampled SPI data to reduce acquisition times.

In this paper, the aim is to show that with purely phase encoded MRI we can overcome all off-resonance related geometrical image artifacts and thus establish an accurate relation between physical coordinates and imaging coordinates. This is highly advantageous both in the treatment planning phase and for on-line image guidance. Furthermore, the benefit that CS brings to this case is demonstrated and we argue that with techniques like these, SPI can also yield a time-efficient and practical solution.

3.2 Methods

Theory

For a 1D spin echo (SE) experiment with frequency encoding along the x-axis (Figure 3.1a) the MRI signal equation (ignoring relaxation effects) for a signal originating in volume V is given by

$$S(t) \propto \int_V \rho(x) e^{-i\gamma G_x x t} dx, \quad (3.1)$$

where $\rho(x)$ is the spin density, γ is the proton gyromagnetic ratio and G_x the frequency encoding gradient which is assumed to have a perfectly rectangular profile in time. The influence of field inhomogeneities, susceptibility transitions, chemical shifts and resonance frequency-shifts (f_0 -shifts) can be expressed in terms of local deviations $\Delta B_0(x) = B(x) - B_0$ from the nominal field strength B_0 , such that the signal becomes

$$S(t) \propto \int_V \rho(x) e^{-i\gamma(G_x x + \Delta B_0(x))t} dx. \quad (3.2)$$

Whenever all N k -space points are acquired after one excitation, the term $\Delta B_0(x)t$ in Equation (3.2) introduces warping of the k -space along the read-out direction. Therefore (by application of the Fourier shift relation) the reconstructed image is distorted according to

$$\hat{\rho}(x) = \rho(x - \Delta B_0(x)/G_x). \quad (3.3)$$

f_0 -shift can be regarded as a special case where ΔB_0 is constant over the imaging volume. The distortion then presents itself as a translation of the imaged object in the frequency encoding direction. In practise, B_0 is obtained from f_0 via the relation $B_0 = 2\pi f_0/\gamma$, whereas f_0 is measured over a region in the sample.

When using phase encoding instead of frequency encoding, as is shown in Figure 3.1b, the recorded signal in the presence of field inhomogeneities becomes

$$S(m) \propto \int_V \rho(x) e^{-i\gamma(m\Delta G_x x \tau + \Delta B_0(x)\delta t)} dx, \quad (3.4)$$

with τ the (constant) duration of the phase encoding step $m\Delta G_x$ and δt the time duration between signal acquisition and the sequence echo time T_E . Note that after excitation and phase encoding, only one k -space point is acquired. The field inhomogeneity term $\Delta B_0(x)\delta t$ now has a constant influence from one k -space point to another, such that the geometrical fidelity of the reconstructed image is preserved. Moreover, in the canonical SPI sequence the echo is acquired at the echo time, such that the field inhomogeneity term does not influence spatial encoding at all. The sequence of Figure 3.1b is repeated N times to measure the full signal over k -space.




Figure 3.1: Sequence diagrams for 1D SE experiments with (a) frequency and (b) phase encoding. In the phase encoding sequence, only one k -space point is acquired after each excitation, ideally at $t = T_E$.

Phantoms

Images are acquired of two phantoms to illustrate SPI performance in various classes of off-resonance related phenomena. Phantom I is an 8 *cm* diameter glass cylinder filled with copper sulfate-doped water. Phantom II is a co-axial glass cylinder. The outer cylinder has a diameter of 8 *cm* and is filled with copper sulfate-doped water. The inner cylinder has a diameter of 2 *cm* and is filled with sunflower oil (phantom IIa) to create a chemical shift difference of approximately 3.5 *ppm*. In phantom IIb, the inner cylinder is filled with air to create a susceptibility transition of 9 *ppm*.

Both phantoms are positioned upright in the geometrical centre of the scanner using the integrated laser-guided positioning system. Hence the long axis of the phantom is oriented perpendicular to the magnetic field.

Imaging

Experiments are performed on a 1.5T whole body Achieva scanner (Philips, Best, The Netherlands). In all imaging experiments, gradient non-linearity effects are corrected for by calibration. All images are acquired using an eight-channel head coil. Coronal cross-sectional images of the phantoms are acquired. The SE sequences are created such that possible artifacts are clearly visible.

SE images are acquired with an FOV of $128 \times 128 \text{ mm}^2$, acquisition matrix 128×128 , flip angle 90° , T_E/T_R 30/100 *ms*, acquisition time 12.8 *s*, bandwidth per pixel 120 *Hz*. Frequency encoding is performed in the craniocaudal direction.

Purely phase encoded images are acquired by means of a spectroscopic imaging sequence. Images are taken with an FOV of $128 \times 128 \text{ mm}^2$, acquisition matrix 92×92 , flip angle 90° , T_E/T_R $30/100 \text{ ms}$, acquisition time 11.8 min , spectral bandwidth 8000 Hz . 64 points are acquired in the spectral dimension during each repetition, of which only one point is used to produce an image.

The method is applied *in vivo* to image the head and neck region of a healthy volunteer. A SE image is acquired with FOV $220 \times 220 \text{ mm}^2$, acquisition matrix 128×128 , flip angle 90° , T_E/T_R $22/100 \text{ ms}$, acquisition time 12.8 s , bandwidth per pixel 120 Hz . Purely phase encoded images are taken with an FOV $220 \times 220 \text{ mm}^2$, while the other sequence parameters are the same as in the phantom studies.

Image reconstruction from incomplete data

To increase imaging speed, we investigate the impact of under-sampling. Data obtained from phantom experiments is retrospectively under-sampled with varying under-sampling factors. Subsequent reconstruction is carried out by zero-filling and Fourier transformation. Furthermore, regularised iterative reconstruction is implemented in MATLAB (The Mathworks, Natick, USA) by means of non-linear conjugate gradient optimisation as outlined in Lustig et al. [2007], which enables inclusion of arbitrary regularisation terms. In such iterative reconstruction the optimisation problem

$$\operatorname{argmin}_x \{|Mx - y|_2^2 + \lambda R(x)\} \quad (3.5)$$

is solved, where M is the measurement operator (i.e., the Fourier operation and subsequent under-sampling), y the measured data, x the image we are looking for, $|\cdot|_2$ the L_2 -norm of a function, R an appropriate regularisation function and λ a regularisation weight. By including a regularisation function in the reconstruction, the incoherent image artifacts that arise from under-sampling are suppressed. For further details the reader is referred to [Lustig et al., 2007].

3.3 Results

SE and SPI reference images are acquired of phantom I with the same shim settings, that were determined by the scanner with the phantom in place. Results are displayed in Figure 3.2. Both SE and SPI yield geometrically reliable images in this case, where field inhomogeneities are absent.

Subsequently, the resonance frequency is manually set to create a 500 Hz global off-resonance. This results in a translation of the object along the frequency encoding direction in the SE image as is confirmed by Figure 3.3a. Phase encoding is insensitive to this global off-resonance and thus localisation accuracy persists in the SPI image (Figure 3.3b). Next we examine phantom IIa, where a water/fat shift artifact is clearly visible in the SE image (Figure 3.4a). In purely phase encoded imaging, the difference between water and fat resonance frequencies does not lead to signal shift and thus the inner cylinder is properly depicted in the outer cylinder (Figure 3.4b). The susceptibility difference in phantom IIb causes a susceptibility artifact that is clearly visible in the conventional SE image (Figure 3.5a). However, in the image acquired using only phase encoding, the artifact is




Figure 3.2: Reference SE (a) and SPI (b) images of phantom I.




Figure 3.3: SE (right) and SPI (left) images of phantom I acquired with a manually set resonance frequency-shift. The red cross denotes the scanner centre in each image.




Figure 3.4: The chemical shift artifact between water and fat is visible in a SE image (a), while proper localisation persists in an image taken with the SPI method (b).

completely suppressed and the co-axial cylindrical geometry is displayed properly. (Figure 3.5b). At last, we apply our technique *in vivo* to image the head and neck region in a healthy volunteer. In the SE image (Figure 3.6a), water/fat




Figure 3.5: Field inhomogeneity induced by a susceptibility transition causes image distortions in a SE image (a), while the SPI image (b) remains unaffected.

shift and susceptibility artifacts are visible. Due to the low resolution (voxel size $2.4 \times 2.4 \text{ mm}^2$) the SPI image appears rather unsharp. The artifacts are clearly mediated effectively however.




Figure 3.6: Image of the head and neck region of a healthy volunteer. (a) SE image. The arrows denote water/fat-shift (w/f) occurring near the skin and susceptibility artifacts (χ) arising from air in the inner ear and the nasal cavity. (b) corresponding SP image, where the artifacts are greatly suppressed.

The influence of under-sampling is investigated using data acquired with phantom IIb, which is retrospectively under-sampled with high acceleration factors according to a piecewise polynomial probability density function as proposed in Lustig et al. [2007]. Images are reconstructed using zero-filling and regularised iterative reconstruction. Results are shown in Figure 3.7. Images of phantom IIb reconstructed using the iterative approach appear crisp even under high acceleration factors, while the simple zero-filling produces noisy images that become increasingly unrecognisable. Reconstruction took approximately 1 *min.* per image. The *in vivo* SPI data is under-sampled retrospectively as well, but now with much more modest acceleration factors of up to four. Results of iterative reconstructions are shown in Figure 3.8. At higher acceleration factors details are lost and stair-casing effects produced by the total variation regularisation function are visible. Nevertheless the image geometry remains a truthful depiction of reality.




Figure 3.7: Reconstruction from under-sampled data with under-sampling factors of 10, 20, 30 and 40 (left to right) for an object that can be represented sparsely in the finite difference domain. Upper row: zero-filling reconstruction; lower row: iterative reconstruction.




Figure 3.8: Reconstruction from under-sampled data for data acquired from a volunteer. (a) reference image, (b)–(d) reconstruction from data sampled at acceleration factors (A) from 2 to 4.

3.4 Discussion

We have demonstrated that, using pure phase encoding, geometrically correct MR images can be obtained, under conditions where a conventional SE image (obtained using frequency encoding) shows f_0 -shift, chemical shift and susceptibility artifacts. As such, SPI presents a fundamental solution to yield undistorted images with one-to-one correspondence to world coordinates. After establishment of this correspondence this method is a valuable tool in QA procedures concerning geometrical relations.

Strictly speaking any errors that arise in the slice selection process have been disregarded here. Since the geometry of objects used in our phantom experiments does not vary along the slice selection direction this is not a problem. In the *in vivo* images slice selection errors cannot be avoided unless 3D phase encoding would be applied. The acquisition time associated with such a scan is prohibitively long however, even when additional parallel imaging strategies are adopted. At this

moment, the use of strong slice selection gradients is the most viable solution to this problem.

The current implementation of SPI through a spectroscopic imaging sequence is rather inefficient. Moreover, operating in spectroscopic imaging mode imposes restrictions on acquisition matrix size and hence on resolution, leading to blurred results for *in vivo* application such as the image presented for the head-and-neck case (Figure 3.6). Additionally, for moving parts of the human anatomy, the SPI technique presents difficulties and causes motion artifacts. Introduction of the purely phase encoded imaging paradigm in the clinic requires implementation of faster sampling schemes, for example as proposed for solid-state MRI [Balcom et al., 1996; Beyea et al., 2000]. Still there will be a large benefit from under-sampling and regularised iterative reconstruction such as proposed here.

The use of a spectroscopic imaging sequence is at the same time an advantage. From such a sequence, a purely phase encoded image can be reconstructed for every time-point in the read-out, yielding many images without increasing acquisition time. All images are geometrically accurate but affected by local field disturbances which lead to signal loss from one image to another. Hence from these images properly localised field maps and T_2^* maps can be derived.

Regularised iterative reconstruction is a good candidate to reconstruct images obtained with the proposed technique and thus speed up the acquisition process. There is a larger benefit as more samples are acquired. However, for phantoms such as used here, faithful reconstruction can be performed from very few samples. This is due to the fact that the object has low total variation and is sparse in the wavelet transform domain. The regularisation functions effectively suppress artifacts arising from the under-sampling. Reconstruction of images of geometry QA phantoms could further benefit from the assumption that the image should be binary to yield good image quality under extremely high acceleration factors [Mao, 2011]. Moreover, as proposed by Parasoglou et al. [2008], an SNR optimal k-space sampling pattern may be derived from knowledge of the geometry of the object. Application to *in vivo* images, which have much higher information content, can benefit from combination with parallel imaging strategies like SENSE.

3.5 Conclusion

Pure phase encoding offers a fundamental solution to the problem of acquiring undistorted images with one-to-one correspondence between image coordinates and world coordinates, which is highly valuable in radiotherapy planning as well as in intervention settings. Regularised iterative reconstruction can be used to speed up the acquisition while retaining geometrically accurate images. High acceleration factors can be reached, especially for images which are extremely sparse in some appropriate transform domain, such as phantoms for calibration and QA of geometrical relations. Further developments may make SPI suitable for clinical introduction in the near future.

An MRI/linac interfacing framework

In this chapter, an MRI/linac interfacing framework to facilitate real-time MRI-guided radiotherapy is described, that is used in the experiments presented in the following two chapters. It aims to tackle the interfacing challenge presented in Chapter 1. The goal is not only to implement communication between the various subsystems in the MRI-linac, but to present an end-user with the look-and-feel of a completely integrated application as well. Moreover, the framework should facilitate investigation of advanced treatment strategies like gated and tracked radiation delivery, based on MRI feedback.

To this end, a client/server architecture is chosen. Controllers are implemented as servers, which each have the responsibility over one of the subsystems (i.e., gantry, beam or MLC) of the linear accelerator. Clients are created, which may request the corresponding server to have the device perform a certain action. Moreover, a server is implemented that is capable of communication with the MR scanner. The latter is responsible for the distribution of received image data over subscribed clients, as well as for relaying requested scan geometry updates. Clients are subsequently wrapped in user interfaces, which may be combined into one interfacing treatment application.

The whole framework is implemented in C++ [Stroustrup, 1991], making use of object oriented programming (OOP) techniques. In OOP, logically coupled data is combined with the functions operating on them to form a class [Booch, 1994]. Through inheritance, a class can be derived from another class. The derived class adds information and functionality to its base class. The classes often represent real-life entities such that the relationships between them easily follow from a good understanding of the problem that must be solved.

In the following, a description is given of the technology on which the framework is built. The interfacing code that facilitates data communication between MR scanner and a third party is subsequently introduced. Then the software components

for linear accelerator control are discussed and an outline is presented of how the implemented client side and server side building blocks may be linked together to form interfacing applications.

4.1 Used technology

In the realisation of a solution to the interfacing challenge, we rely on some existing enabling technology. A short discussion of the two most important existing techniques used is given in this section.

Common Object Request Broker Architecture

For inter-process (client/server) communication, the common object request broker architecture (CORBA) is adopted [OMG, 2012]. CORBA is a standardised middleware that enables programs to communicate among each other. These programs may run on the same computer, but can also run on computers far apart from each other, as long as there exists an open route between them. From an end-user perspective, it does not matter where these applications run. CORBA is supported on a host of hardware/operating system combinations. Moreover, standardised mappings exist for a fair number of programming languages, including C++ [Henning and Vinoski, 1999].

Development of a CORBA-based application framework starts by specifying the interfaces that classes expose. A standardised interface definition language (IDL)¹ exists for this purpose. Subsequently, a code generation tool is used to create client and server side code to respectively access and implement CORBA objects that support the declared interfaces. The generated server side code represents an abstract base class from which a class is derived that contains the actual implementation of the interface methods (see Figure 4.1).

The implemented classes are subsequently used in a server application. Part of this server program is the object request broker (ORB), that actually performs the communication. A client application is implemented, that also contains an ORB. The client program obtains a reference to the server side object, through which it can make calls (or requests) on that object, mediated by the ORBs on both sides. In Figure 4.1 a schematic overview is given of the auto-generated and user implemented parts of the application framework code, as well as of the calls over a network mediated by ORBs. For a more rigorous explanation of CORBA programming with C++ the reader is referred to Henning and Vinoski [1999].

It is important to note that CORBA does not give any guarantees for real-time performance. However, it does facilitate the use of constraints on the communication; for example, policies can be set regarding the round-trip time allowed for a request and its reply. In case this time is exceeded, the ORB generates a time-out exception and hence notifies the client of a ‘problem’. As such, CORBA provides some means to take measures in case performance is inadequate. Moreover, in a research environment, where network traffic and processes running on the con-

¹Not to be confused with any popular mathematics computer program.

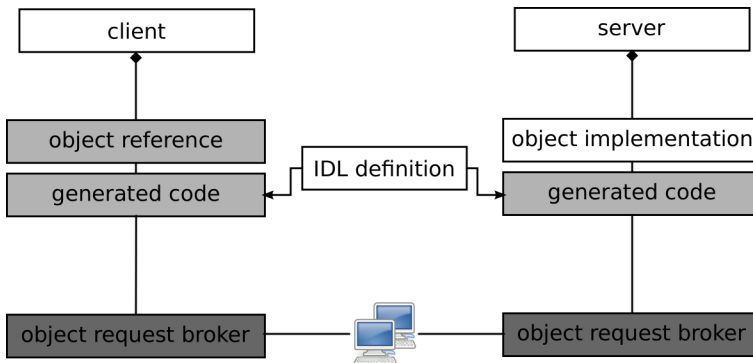


Figure 4.1: Auto-generated and user-implemented code in a CORBA-based client and server architecture. White blocks represent code implemented by the developer, light grey blocks are generated by tools provided by the CORBA-vendor and dark grey blocks are implemented by the CORBA-vendors themselves.

cerned computers can be strictly controlled and performance failures do not have large implications, the efficiency of CORBA is sufficient.

Controller area network bus

For communication with hardware, controller area network (CAN) technology is employed. The CAN protocol is a message based protocol for communication over a serial line. Originally it was designed for the vehicle industry but it now finds applications in industrial automation and medical technology as well.

The CAN protocol defines packets as logical messages. In principle all nodes in a CAN network have the possibility to tranceive messages; simultaneous broadcasting results in an exception, however. Arbitration of messages is done based on identifiers; the lower the message ID, the higher its priority. A CAN message further contains a data field of a length between zero and eight bytes, in which any kind of information may be stored. In general, data transmission is performed on an autonomous basis with the data source node sending out messages. The CAN bus protocol supports bit rates up to 1 *Mbit/s* at network lengths below 40 *m*.

4.2 Communication with the MR scanner

The code used within the interfacing framework to communicate with the MR scanner was pioneered by the IMF lab at the of the Université de Bordeaux, in collaboration with Philips. Their development of a real-time interface that mediates transfer of data regarding images, raw samples, patient physiology and pencil beam navigators has been the basis of a number of proof of concept experiments in motion adapted high intensity focused ultrasound therapy [Köhler et al., 2011; Ries et al., 2010].

The scanner to third party interface consists of a server that receives the previously mentioned data from the MR scanner as soon as it is recorded and reconstructed. Clients are implemented that can register itself with the server as being interested in a particular type of data. The server is responsible for distributing the data among the clients upon arrival. Communication between clients and server is mediated by CORBA.

4.3 Linac control utility and support classes

A library is created with common utility and support classes, which take care of file input/output, parameter files, bit-wise operations, computations of data statistics, logging and exception handling.

Additionally, some support classes are created that implement safety mechanisms, take care of CAN devices and communication, as well as classes that aim to hide the details of CORBA mediated client/server communication from an application programmer. These are discussed below in more detail.

Watchdogs

Since the proposed framework is modular, but also distributed, we must account for possible risks due to lost connections between the building blocks.

In case a connection is dropped, server side controller applications can no longer receive input from clients. For example, a server side application managing the radiation beam could keep the beam on, while the client has information that indicates it must be switched off. More general, all server side applications in the framework, that manage devices, should bring the device into a safe state in case of headless (i.e., without an attached client) operation.

On the client side, a malfunctioning connection has the consequence that a user tries to move forward with a procedure, while the device he or she is trying to control has become unresponsive. The user should thus at least be notified (audio)visually of a lost connection so that additional measures can be taken. Additionally, a client application may interface with multiple systems; for example, a client that realises advanced radiation delivery techniques may provide control commands to beam, gantry and MLC at the same time. In case such a client loses the connection to one subsystem, other subsystems should be brought to a halt safely.

A watchdog (see Figure 4.2) provides a mechanism by which a client can regularly test the connection to the server. At the same time, this mechanism periodically notifies the server that a client is still attached. The server side portion of the watchdog mechanism is implemented in a `Watchdog` class. The client side portion is delegated to the `CORBAProxyBase` class (see below).

Server orb handler

All server applications in the suggested framework must instantiate and manage at least one ORB to manage client connections. The best performance is obtained

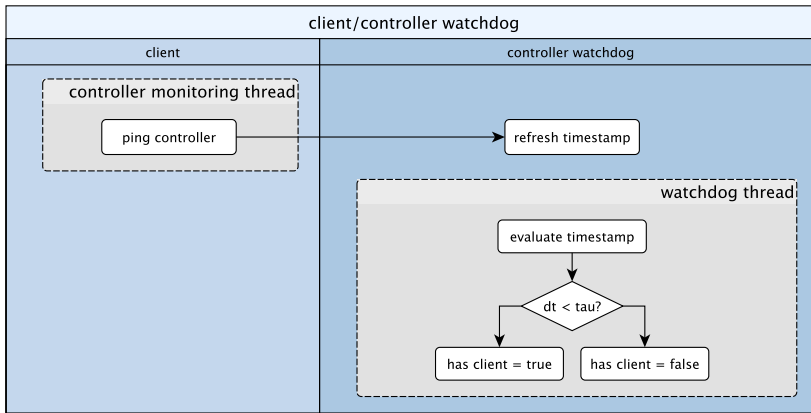


Figure 4.2: Watchdog mechanism for safe(r) operation.

when each CORBA-enabled object has its own ORB [Henning and Vinoski, 1999], which for a small number of server side objects is feasible.

A base class for server side exposure of CORBA enabled classes is implemented. This `ORBHandlerBase` class sets a few common default ORB arguments and provides a method to set other ORB arguments, as well as to run the ORB and shut it down. The class is abstract, in the sense that it requires derived classes to implement methods that expose an object reference, set the connection policies and perform a clean-up on shutdown. Derived classes are implemented, that expose references to CORBA enabled classes for the devices to be controlled (Figure 4.3). These classes are introduced in Section 4.4.




Figure 4.3: Generalised class diagram for an `ORBHandler` in the framework. Closed diamonds denote composition, whereas the open arrow denotes inheritance.

Proxy classes

In the suggested interfacing framework, CORBA takes care of the technical details of the inter process communication. Unfortunately that does not mean that the

developer of a client application should not be concerned about the communication. Code has to be written every time a method on a CORBA object is called, to map C++ data types to CORBA data types and to manage possibly thrown exceptions. Moreover, the client side code of the watchdog mechanism should be implemented for every connection the client application creates. To hide CORBA-related code from the application programmer and to provide a common class for the client side watchdog mechanism, proxy classes have been implemented. The common functionality of testing the server connection periodically and attempting to restore the connection in case it is dropped, is wrapped in a `CORBAProxyBase` class (Figure 4.4). It is an abstract class that enforces the derived classes to implement a method to connect to obtain an object reference, as well as a shutdown method that performs a cleanup. More specialised proxy classes are implemented to give a client access to the devices to be controlled. Each of these proxy classes




Figure 4.4: Generalised class diagram for a Proxy class in the interfacing framework.

is derived from the `CORBAProxyBase` class and exposes a set of methods to control the device concerned, which manages the data type conversions required to call CORBA code. All these methods return an indicator that tells whether a call has made its way to the server side object (i.e., an indicator of the connection status, which is derived from any exception caught on making the CORBA call).

CAN communication

CAN communication is the basis for control of both gantry and MLC. Classes are implemented to wrap and re-use the common functionality needed to communicate with these devices (Figure 4.5).




Figure 4.5: Overview of CAN-related classes.

The message sent or received by CAN hardware is represented by the `CANMessageBase` class. It manages the array of actual message bytes and provides routines to access these in a thread-safe manner.

The `CANControlledDevice` class takes care of initialisation of the CAN hardware with the right parameters and shuts it down at the end of the object life cycle. Moreover, it implements message send/receive routines, that respectively take `CANTransmitMessage` and `CANReceiveMessage` classes as their arguments.

The `CANReceiveMessage` is a specialisation of the `CANMessageBase` class, where the byte array always has length eight. It requires derived classes to implement an `unpack` routine that either buffers the message bytes for later use or assigns values to member variables based on the message byte content. Moreover, it provides variables and routines to perform time stamping of the received data.

The `CANTransmitMessage` inherits the `CANMessageBase` class as well. It is an abstract class that requires derived classes to implement a `pack` method that copies the values of member variables into the message byte array. Alternatively, the member variables may be stored directly in the byte array.

Last but not least, CAN related exception classes are implemented as defined by the API for the device used. All CAN related classes are packaged into a library for re-use.

4.4 Controlling the linear accelerator

An important part of the framework deals with the sub-devices that a radiotherapy linear accelerator is composed of. In this section, the implementation of the device classes introduced in Figure 4.3 is dealt with. Device classes to control gantry, MLC and radiation beam are presented, that must be wrapped in a corresponding `ServerORBHandler` class, following the pattern in Section 4.3. Each of the `ServerORBHandlers` is then used in a server application for the respective device.

Controlling the gantry

The gantry in the MRI-linac prototype is a drum-and-base type gantry powered by an electric motor. A physical layout diagram for gantry control is shown in Figure 4.6. A gantry drive inside the treatment shelter is in control of the gantry; it can




Figure 4.6: Physical layout for gantry control.

be requested to move the gantry to a certain position, or to perform rotation at a certain speed in a defined direction. To make these requests, a set of CAN messages implementing commands was defined. A server must be implemented that is able to relay client requests to the gantry drive. The actual control intelligence remains in the gantry drive, however. Furthermore, the gantry drive notifies the server of gantry variables (position, velocity, etc.) every millisecond.

The `Gantry` class is the CORBA-enabled class for communications regarding the gantry device (Figure 4.7). A `CANControlledGantry` class inherits from the

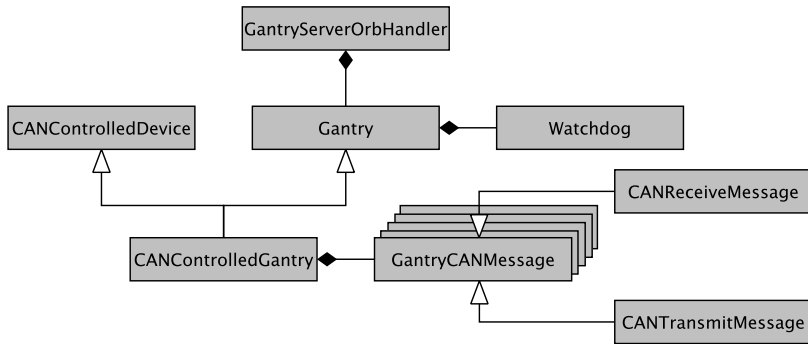


Figure 4.7: Class diagram for gantry related classes. Gantry related CAN message classes inherit either from `CANReceiveMessage` or from `CANTransmitMessage`.

`Gantry` class and from the `CANControlledDevice` class; it implements CAN communication to the gantry drive. Specialised gantry related CAN message classes are implemented, that are derived from the previously introduced CAN message utility classes.

The `CANControlledGantry` class implements a monitoring loop (see Figure 4.8). On each loop iteration, a keep alive signal is send to the gantry drive, status




Figure 4.8: Schematic overview of the gantry monitoring as implemented in the `CANControlledGantry` class.

information received from the gantry drive is analysed and the gantry is stopped in the event that rotation takes place unsupervised (i.e., without an attached client). The loop runs every two milliseconds. Client requests are directly relayed to the gantry drive over the CAN connection, since this does not happen at high temporal resolution and requires only minimal time.

Controlling the MLC

The collimator in our MRI-linac prototype is a 160 leaf MLC, derived from the Agility MLC (Elekta AB, Stockholm, Sweden). The leaf shape was modified to account for the different source-to-surface distance and the leaf and yoke materials were changed to be able to keep forces due to the magnetic field on the leaves to a minimum. Communication with the MLC electronics that control the various MLC subsystems is through the CAN bus. In contrast to the gantry, the MLC electronics do not contain intelligence for leaf positioning; this must therefore be implemented into the controller. For feedback of leaf positions, the MLC has a video camera based system running at 25 Hz. The video camera looks at reflectors on the leaf tips, which are illuminated by LEDs. A schematic overview of the systems in the MLC control chain is given in Figure 4.9. To minimise latency in the control




Figure 4.9: Physical layout for MLC control.

algorithms, the MLC server/controller application is required to run on a PC on board the gantry. Additionally, the controller must run on a computer connected to the MLC electronics via CAN and with access to a frame grabber card to acquire pictures taken by the MLC video camera.

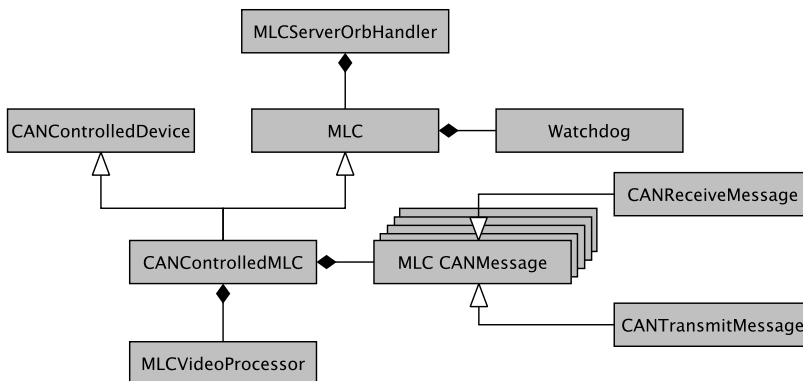


Figure 4.10: Class diagram for MLC related classes.

The classes implemented for control of the MLC are shown in Figure 4.10. The MLC class is the CORBA-enabled class for communication regarding the MLC device. The `CANControlledMLC` class inherits from MLC as well as from `CANControlledDevice`. It takes care of the actual communication with the MLC electronics. MLC related CAN message classes are implemented (e.g., to enable and disable leaves, to set their speed and to vary the intensity of the illumination LEDs), that inherit from the previously defined `CANReceiveMessage` or `CANTransmitMessage` utility classes. The multitude of parameters and commands that can be given to the MLC is reflected by the number of MLC CAN message classes (22 in total). The `MLCVideoProcessor` class provides access to a frame grabber card. It grabs a frame as soon as it is acquired by the camera, and subsequently analyses it to find the leaf tip positions.

Besides handling communication, the `CANControlledMLC` class also implements a leaf positioning control loop (see Figure 4.11). Within the control loop, the




Figure 4.11: Simplified schematic overview of the leaf positioning loop as implemented in the `CANControlledDevice` class. The video analysis thread is spawned by the `MLCVideoProcessor` object.

electronics status is monitored and controlled, a keep alive signal is sent to the electronics, leaves are frozen in case the watchdog mechanism fails and a servo algorithm is carried out. The servo compares the leaf tip positions obtained from the `MLCVideoProcessor` to leaf tip positions as requested by a client and adjusts leaf speeds to compensate for position discrepancies. After running the servo,

leaf velocity updates are sent to the MLC electronics and some additional status information is acquired. As a result of the acquisition frequency of the video camera, the control loop runs at 25 Hz , or every 40 ms .

Controlling the radiation beam

The last sub-device in the linear accelerator that we establish control classes for is the radiation beam. Figure 4.12 gives a schematic overview of the physical layout regarding beam control. A real-time computer (RTC) controls the beam generation

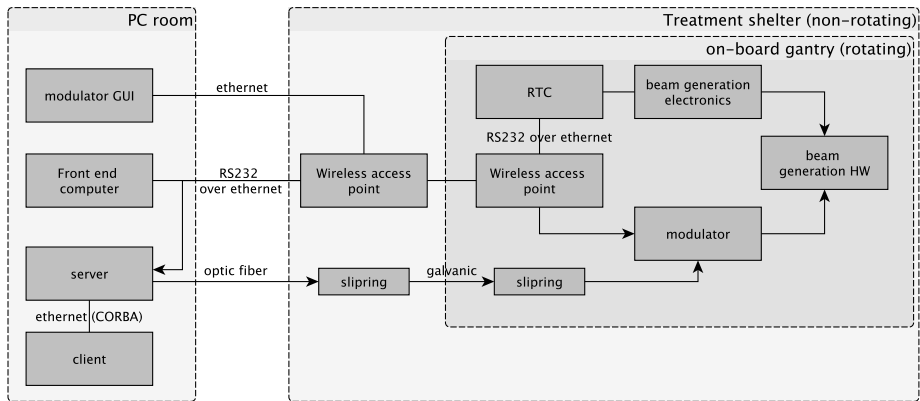


Figure 4.12: Physical layout for beam control and monitoring

electronics. In the current prototype this does not include the modulator. Using the front-end computer (FEC), a prescription (i.e., a number of monitor units at a certain dose rate) can be uploaded to the RTC. To start the beam, the modulator must be activated from the modulator graphical user interface (GUI). A server application is implemented that runs on a PC in the PC room, which is connected to the modulator via an optic fibre and a hard-wired connection through a slip ring. This connection can be used to keep the modulator from producing high voltage pulses, hence turning the radiation beam off and back on after it has been turned on by the RTC. Furthermore, the same PC listens in on the RS232 communication between FEC and RTC, to monitor the dose rate and accumulated delivered dose as measured by the RTC.

To implement the beam controlling/monitoring server application, the classes in Figure 4.13a are introduced. The `Beam` class is the CORBA enabled class for communication with clients. It has a `Watchdog` to prevent beam on without an attached client. No derived class is created to implement the actual beam control, since the required additional code is minimal. Instead, this is implemented directly into the `Beam` class. A `CallbackAsyncSerial` helper class is implemented that handles communication over the serial port (i.e., RS232), to be able to listen to RTC/FEC communication.

The control loop implemented in the `Beam` class is rather basic; it monitors the `Watchdog` and prevents beam on in case no client is present (Figure 4.13a). All




Figure 4.13: Code written to control the beam and monitor its status.

requests from clients regarding beam status are directly carried out, since turning the beam on or off requires a minimal amount of time and does not (or should not) happen too often.

4.5 User interfaces

For a front end computer experience, clients are wrapped in user interfaces. These interfaces give an end user control over the device and present visual information regarding the connection and the device status. The most important examples are presented in this section.

In Figure 4.14 a user interface for the control and monitoring of the gantry is depicted. It contains a video panel, where a live video feed from inside the treatment




Figure 4.14: A user interface for control of the gantry.

shelter is displayed. Furthermore, it provides information regarding the position, velocity and controller status on the left. With the controls on the right side, a user can set position and velocity of rotation or initiate a homing sequence, for example.

Figure 4.15 shows a part of a composite user interface to control the MLC. On the right, a panel is shown that displays the video as recorded by the MLC camera. The




Figure 4.15: A user interface for control of the MLC.

white rectangles are the leaf tip reflectors. Leaf position as requested by the user are depicted with a dark green overlay rectangle. The left panel contains means to switch the servo on and off, and to set the position of individual leaves. The user interface contains a series of other widgets that present status data, provide interfaces to upload a field, and parametrise the video processor, amongst others. Last but not least some user interfaces are implemented to display data from the MR scanner and help to set scan geometry. Figure 4.16 shows a composite user




Figure 4.16: MR scanner user interfaces.

interface that was created for this purpose. On the left, it has two panels that display magnitude and phase images acquired by the scanner. On the right, a scan plane manipulation tool is displayed. By dragging the mouse over the green plane, the scan plane can be rotated interactively.

4.6 Concluding remarks

Control over the linear accelerator sub devices, as well as communication with the MR scanner have been established. The implemented dynamic controllers and clients are suitable for research on IMRT delivery (Figure 4.17), provided that the MLC is calibrated, as well as intra-fraction motion management strategies.



Figure 4.17: IMRT fields delivered on the MRI-linac prototype, that together form the acronym UMCU. The radiation fields were recorded on GafChromic EBT2 radiation sensitive film.

Using the building blocks introduced in this chapter, it is relatively easy to build client applications that provide an interface between the MRI scanner and the linac. MRI guided control of the linear accelerator subsystems is thereby realised. Moreover, with the created user interfaces an integrated look-and-feel can be presented to a user, while keeping the actual application modular.

MRI-guided gated radiation delivery

This chapter is based on:

S.P.M. Crijns, J.G.M. Kok, J.J.W. Lagendijk and B.W. Raaymakers. Towards MRI-guided linear accelerator control: gating on an MRI accelerator *Physics in Medicine and Biology* 2011 **56**(15):4815–4825.

5.1 Introduction

In radiotherapy one aims to deliver a radiation dose to a tumour as precisely as possible. Image guidance decreases the geometrical uncertainties associated with target position and shape [Verellen et al., 2007]. MRI is a prime candidate for image guidance since it provides superior soft-tissue contrast and thus a clear view of both tumorous tissue and OARs. Legendijk et al. [2008] have discussed the advantage of using MRI, not only as a position verification system, but as an on-line treatment guidance modality as well. On-line MR images enable development of new treatment strategies, even for patient groups that were not eligible for radiotherapy so far (e.g., Kerkhof et al. [2011]). Moreover, cine-MRI enables real-time tracking or characterisation of the motion of a 3D volume of interest. Recently, we have reported on a prototype integrated MRI-linac that aims to implement such MR-guidance [Raaymakers et al., 2009]. Other designs have been proposed that also reached the prototyping phase [Dempsey et al., 2005; Fallone et al., 2009]. For an integrated MRI-linac, the ultimate goal is to provide on-line MR image guidance for therapy plan adaptations based on the current state of the anatomy and the accumulated delivered dose. Treatment of cancer of the abdominal organs such as kidney, pancreas and liver, is hampered by respiratory motion since it causes the target to make large semi-periodic excursions [Bortfeld et al., 2004; Kitamura et al., 2003; Van Sörnsen de Koste et al., 2006]. To deliver an adequate dose to such targets while at the same time sparing surrounding organs, this intra-fraction motion should be taken into account.

Lujan et al. [1999] have shown that doses delivered under dynamic conditions may be calculated from planned dose distributions via position probability density functions (PPDFs). Unkelbach and Oelfke [2004] subsequently demonstrated that this information can be taken into account via probabilistic inverse planning. These approaches assume reproducible periodic motion that is known already in the planning phase. Moreover, dose blurring still occurs and potentially causes an increased dose to the OARs. These issues may be resolved by adopting on-line compensation strategies such as gating or tracking [Ozhasoglu and Murphy, 2002].

Gating, like tracking strategies, can be used for margin reduction provided that it is applied under image guidance [Korreman et al., 2008] since it diminishes target motion induced errors [Hugo et al., 2003], which may allow for dose escalation [Wagman et al., 2003]. Recently it has been shown that the mean target position concept [Wolthaus et al., 2008] may also be incorporated in this strategy in order to increase the treatment efficiency [Guckenberger et al., 2011]. The challenges encountered in the implementation of gating (i.e., target position monitoring, communication infrastructure) are generally applicable to tracking. Moreover, full tracking radiation delivery may not be viable in all cases and combination with gating strategies could then be a tractable alternative.

In this paper we present gated radiation delivery where, for the first time, the gating signal is derived from on-line MR imaging. Hereby we show that the issues and technical challenges around the hypothesis that MR guidance and control is possible can be overcome. We establish real-time communication with the MR scanner and modify the accelerator to facilitate gating. Subsequently we evaluate this implementation using methods and concepts that are well-established in the

gating community. A phantom is subjected to 1D periodic motion. We monitor its position in time using fast 1D MR imaging and radiation delivered to the phantom is gated based on this signal. The time-resolved position information is processed to yield a PPDF for comparison with the theoretical PPDF derived from the imposed motion pattern. Since communication latency and other unpredictable effects may introduce errors, we validate our gating implementation by comparing the measured dose profiles to profiles obtained from the combination of a theoretical PPDF and a measured statically delivered dose. With this comparison we illustrate the feasibility of calculating dose distributions delivered to targets exhibiting motion using time-resolved position information which should ultimately facilitate the accumulation of dose maps for plan adaptation.

5.2 Materials and Methods

Setup

Experiments are carried out using our MRI-linac prototype, which comprises a table-mounted 6 MV Elekta AB (Stockholm, Sweden) accelerator and a modified 1.5 T Philips (Best, The Netherlands) Achieva MRI system [Raaymakers et al., 2009].

Our experimental setup consists of a cart positioned in the scanner isocentre. It carries dosimetric film between two layers of water equivalent (with respect to attenuation of radiation) build-up material as well as a box-like MR phantom filled with CuSO_4 -doped water (Figure 5.1). The cart is driven by an electric motor outside the scanner bore (at approximately 2.5 m from the isocentre) via a rigid crankshaft mechanism. The motor is fed by a programmable power supply attached to a digital-to-analog converter to allow prescription of arbitrary periodic one-dimensional motion of the cart. Motion of the cart is along the long axis of




Figure 5.1: Cart (green) with dosimetric setup (red) and MR phantom (blue). The imaging volume covered by the pencil-beam navigators is indicated in yellow. The cart is propelled by an electric motor (purple) via a crankshaft mechanism.

the scanner bore and perpendicular to the radiation beam.

Tracking protocol

Since the prescribed motion in this setup is strictly 1D translation, we can fully track it using a pencil-beam navigator. Navigators are 1D acquisitions that are conventionally used to gate the acquisition of 2 or 3D images or to track organ position as an input for retrospective phase sorting and motion compensation. They

take approximately 15 *ms* to acquire and are relatively easy to analyse compared to higher dimensional images. To realise real-time feedback in the gating experiments, pencil-beam navigators with a resolution of 1 *mm* are recorded at a rate of approximately 25 per second. The position of the phantom is inferred from each navigator by edge detection to yield experimental PPDFs for dose accumulation.

MRI/accelerator interface

MRI-based gating requires extraction of acquired data from the MR scanner and control of the accelerator.

Real-time communication with the MR scanner is realised through the External Control interface [Ries et al., 2010; Smink et al., 2011]. This interface relies on the ACE ORB (TAO) implementation of CORBA [Schmidt et al., 1998] for the underlying computer-to-computer communications. It facilitates communication of the acquired data to an external PC with minimal lag-time (~ 10 *ms*).

The accelerator is modified to handle gating. We use pulse rate frequency (PRF) muting, similar to the method previously described by Evans et al. [2010]. However, no off-the-shelf gating interface could be used with our dedicated accelerator. Instead an optical interface from the serial port of a PC to the circuit board that controls the magnetron PRF is realised. The beam can thus be turned off by muting the PRF generator, with the consequence that the magnetron is inhibited from producing RF pulses. Automatic frequency compensation (AFC), which during normal operation counteracts temperature effects in the accelerator tube, has to be disabled to prevent erratic compensation during beam-off periods.

Acquired navigators, as well as the derived beam status and the current gating window are visualised in real-time in a GUI (Figure 5.2). Both width and centre of the gating window can be manipulated interactively, allowing manual compensation of drifting effects if desired. To permit future extensions, separate processes




Figure 5.2: GUI for gating. Navigators are visualised and overlaid with the derived position (pink dots). The gating window is displayed as a blue bar and the background colour (either red or green) indicates the current beam on/off status.

are created for scanner communication and data analysis and to control the accelerator. Inter process communication is again handled by CORBA.

Experiments

The described setup is subjected to a $5 \times 5 \text{ cm}^2$ radiation beam under several conditions.

- Delivery to a non-moving phantom to acquire a ground-truth for comparison.
- Delivery to a phantom that moves according to a perfect sine function with a period of 5 s. Besides ungated irradiation, gating with windows between 5 and 25 mm is evaluated. Gating is performed at one of the extremes of the motion.
- Delivery while the phantom is prescribed a motion that resembles movement in the abdomen caused by breathing patterns (as described by Equation (2) in Lujan et al. [1999]). This motion has a period of 7 s and it is characterised by a relatively long plateau-phase. Here gating windows of 5 and 10 mm are investigated, as well as ungated irradiation. Gating is performed in the plateau phase.

The motion amplitude for the dynamic cases is fixed at 25 mm. The phantom positions according to the tracking protocol are recorded, as well as the beam status. For all situations a prescribed dose of 2 Gy is delivered to the phantom, which for the average treatment plan coincides with the dose given in one fraction. Resulting dose distributions are measured with GafChromic EBT2 film (ISP, Wayne, USA).

Data processing

The GafChromic films are scanned and processed with in-house developed software. The resulting dose distributions are averaged over a 1 cm wide band to obtain dose profiles. These are subsequently normalised by height and centred for comparison.

The recorded phantom positions are combined with the beam status information to yield beam-on position probability density functions (PPDFs) by binning the positions where the beam is on and normalising the resulting distribution. We compare the measured distributions (see Section 5.2) to the theoretical PPDF for sinusoidal translation along the x -axis, which is given by

$$P(x) = \begin{cases} \frac{C}{\sqrt{A^2 - x^2}}, & |x - x_{gate}| \leq w, \\ 0, & \text{elsewhere,} \end{cases} \quad (5.1)$$

with C a scaling constant such that $\int_{-\infty}^{\infty} P(x) dx = 1$, A the motion amplitude, x_{gate} the gating position and w the size of the window [Bortfeld et al., 2002].

Position information can be employed to calculate the delivered dose from the planned dose under the assumption of purely rigid motion as

$$D_{delivered}(\vec{r}) = P(\vec{r}) \otimes D_{planned}(\vec{r}), \quad (5.2)$$

where $D(\vec{r})$ is the dose at position \vec{r} and \otimes denotes a convolution [Lujan et al., 1999]. Equations (5.1) and (5.2) are then used to calculate the expected delivered dose which is compared to the measured dose profile in order to validate our gating implementation.

All off-line data processing is carried out in MATLAB (The Mathworks, Natick, USA).

5.3 Results and discussion

Figure 5.3 shows an impression of navigator signals of the phantom exhibiting both motion patterns. The position inferred from the navigators is overlaid on the signals themselves. No apparent tracking errors are observed.




Figure 5.3: Example of acquired navigator signals for (a) a perfect sine (b) respiratory-like motion. The green line in both images denotes the position as obtained from each navigator.

Dose profiles are reconstructed from the GafChromic film measurements for each of the motion patterns and gating windows (Figure 5.4). As expected, the simi-




Figure 5.4: Measured normalised dose profiles obtained with several gating window widths w (in mm) for (a) sinusoidal motion and (b) respiratory-like motion as displayed in Figure 5.3b. $w = \infty$ denotes delivery without gating.

ilarity between dynamically and statically obtained profiles gradually increases as the gating window is decreased. The time for delivery of the full planned dose

increases from 2.5 *min* for ungated irradiation to approximately 20 *min* for the 5 *mm* window. This would be an unacceptably long time in clinical practise since the amount of radiation delivered is approximately equivalent to one treatment fraction. The size of the gating window is apparently not very critical as can be seen in Figure 5.4a. There is hardly any difference between the curves obtained with windows of 5 and 10 *mm* and hence time can be gained by enlarging the window, thereby increasing the duty cycle. The importance of the window size may vary for more complex motion patterns however and thus trade-offs have to be made [Dietrich et al., 2005]. Influencing the PPDF by coached breathing (e.g., as proposed by Masselli et al. [2009]) or breath-hold strategies in combination with on-line imaging is an effective way of boosting the duty-cycle during treatment [Kerckhof et al., 2011].

The PPDF for delivery without gating during sinusoidal motion is derived from the logged position data (Figure 5.5b). From this distribution, it is immediately obvious that gating in the extremes of the motion curve is most sensible. As such, PPDFs derived from (potentially patient specific) motion models may aid in the determination of the best gating window position and width for more complex cases.

The theoretical PPDF in Figure 5.5b is modified according to the applied gating window (i.e., varying w in (5.1)) to obtain a beam-on-PPDF. The statically acquired profile is convolved with this latter distribution to obtain predictions for the dose delivered with gating windows of 25 and 5 *mm* (figures 5.5c and 5.5d). Especially for the former case, excellent correspondence with the measured profile is observed. This clearly illustrates the potential of real-time on-line imaging during radiotherapy: since the location of the target is known during beam on, we are able to reconstruct the delivered dose profile from the planned dose profile. Moreover, having time-resolved position information enables real-time on-line dose accumulation. This not only permits adaptation on a fraction-to-fraction scale, but also during one fraction, which is especially valuable in highly accelerated treatment strategies.

The profile obtained with a window of 5 *mm* is slightly different from the calculated dose in Figure 5.5d. More specifically there is a drop-off in the plateau of the profile, as could also be noted from the curve for a 10 *mm* window in Figure 5.4a, albeit less pronounced. We hypothesise that this is caused by repeated start-up effects of the accelerator; the beam is switched on as soon as the target arrives in the gating window, but the dose rate increases for a short period of time before it reaches its maximum. Since the AFC of the accelerator is turned off, the dose rate is much more instable than during normal operation. Consequently the dose rate varies over the period spent in the gating window, which is more evident in smaller gating windows. Since there is a long plateau phase in the respiration-like motion, the dose rate is more stable during the beam on time and the plateau stays more flat (see Figure 5.4b). Although the improvised gating interface works well enough to carry out the current experiments, a more solid implementation has to be realised in the future. Start-up effects would be reduced if the diode electron gun in our accelerator would be substituted with a triode gun. Additionally, automatic frequency compensation can be performed, hence boosting the performance even more.




Figure 5.5: Measured and predicted dose profiles for gating windows of (c) 25 and (d) 5 mm obtained by convolving the profile recorded during delivery without motion (a) with the corresponding theoretical PPDF for sinusoidal motion (b). In (b), the shaded area denotes the beam on probability density function for a 25 mm gating window. $P(x)$ is the probability to find the target at position x . Measured profiles are shown by dashed lines whereas the calculated profiles are shown in solid lines. All profiles are normalised for comparison.

Controlling gantry rotation and MLC settings can be realised within the feedback framework created here. It is possible to use methods such as presented to provide feedback for the MLC hence realising MRI-guided tracking applications. Correction for baseline-drift during gating is a first step in implementing more adaptive treatment strategies, that requires control of both beam and MLC [Pepin et al., 2011]. Magnetic field inhomogeneities lead to distorted MR images which, if used as an input for gating procedures, could introduce errors in the dose delivery process. These distortions can be compensated for with methods such as presented in Crijns et al. [2011a], which is also applicable to the pencil-beam navigator acquisitions presented here.

Motion of abdominal organs is generally much more complicated than the 1D rigid translation we have used here. 3D motion patterns, deformations and hysteresis are to be expected *in vivo*. The convolution approach employed here to calculate dose from time-resolved target positions is not generally applicable in cases with deformation. Non-rigid registration facilitates dose accumulation under these

conditions, as was recently demonstrated by Velec et al. [2011] for treatment of liver tumours. Köhler et al. [2011] have successfully demonstrated the use of 2D imaging at 10 Hz in combination with pencil-beam navigators to compensate for 3D motion in MR-guided high-intensity focused ultrasound therapy of abdominal organs. With this promising technique it should be possible to realise radiation beam tracking for complex motion patterns as well.

5.4 Conclusion

We have demonstrated gated radiation delivery based on target position feedback from on-line MR imaging. Monitoring of the phantom position was achieved with a 1D pencil-beam navigator sequence which resulted in exact reproduction of the applied motion pattern. Resulting dose distributions confirm excellent control of the beam based on the MR signal acquired. Moreover, we have shown reconstruction of dose delivered to a moving target using the latter measurement. Adapting the therapy plan for the next fraction based on the accumulated dose delivered so far is thereby within reach. With the currently established framework and the planned improvements we will realise more intricate MRI-guided linear accelerator control in the near future.

MRI-guided tracked radiation delivery

This chapter is based on:

S.P.M. Crijs, B.W. Raaymakers, and J.J.W. Lagendijk. Proof of concept of MRI-guided tracked radiation delivery: tracking one-dimensional motion. *Physics in Medicine and Biology*, 2012 **56**(15):7863–7872.

6.1 Introduction

In radiotherapy one aims to deliver a radiation dose to a tumour with high geometrical accuracy while sparing OARs. Image guidance decreases the geometrical uncertainties associated with target position and shape [Verellen et al., 2007]. Moreover, treatment of cancer of the abdominal organs, such as kidney, pancreas and liver, is hampered by respiratory motion since it causes the target to make large semi-periodic excursions [Bortfeld et al., 2004; Kitamura et al., 2003; Van Sörnsen de Koste et al., 2006]. Delivery of an adequate dose to such targets while at the same time sparing surrounding organs requires compensation for such motion.

4D planning strategies cope with intra-fraction motion by explicitly taking it into account (e.g., Keall et al. [2005]; Unkelbach and Oelfke [2004]). However, these approaches assume complete pre-treatment knowledge of the motion patterns and dose blurring still occurs and potentially causes an increased dose to the OARs. Alternatively, gated radiation delivery, when applied under image guidance, can be used to reduce margins, while preserving the quality of the delivered dose distribution [Crijs et al., 2011b; Korreman et al., 2008]. As such this strategy may also allow for dose escalation [Wagman et al., 2003]. Although gated delivery is a promising and relatively simple technique for intra-fraction motion compensation, its efficiency can be rather low depending on target motion statistics.

The ultimate way of intra-fraction motion compensation is to keep the radiation beam aligned with the moving target using dynamic motion compensation. This can be achieved with robotic linear accelerator control [Schweikard et al., 2000], couch motion [D'Souza et al., 2005], a gimbals based system [Depuydt et al., 2011; Kamino et al., 2006] or by adapting the MLC aperture, (e.g., as demonstrated by Keall et al. [2001]; Papiez and Rangaraj [2005]; Papiez et al. [2005]; Suh et al. [2004]). Guidance modalities for dynamic MLC (dMLC) based tracked radiation delivery explored so far include kV [Poulsen et al., 2010] and MV imaging and combinations thereof [Cho et al., 2009b, 2011], combined kV and optical systems [Cho et al., 2009a] and wired [Ravkilde et al., 2011] or wireless [Sawant et al., 2009] EM transponders, which have been developed mainly for prostate cancer treatment. Good performance is generally obtained depending mainly on latency of the feedback signal employed. However, all these feed-back modalities track one or a couple of points (markers) and thereby look at surrogates for the actual target motion. This is a clear disadvantage when looking at aperture adaptations to facilitate target deformations. Moreover, only the target is tracked, whereas OARs may exhibit (uncorrelated) motion as well. Finally, marker implantation is an invasive procedure which is not suitable for any tumour site.

The use of MRI may allow for alleviation these shortcomings by providing images of both targets and OARs directly and with a superior soft-tissue contrast [Legendijk et al., 2008], while remaining non-invasive. Moreover, cine MR imaging (where MR images are acquired multiple times per second) enables tracking of specific anatomical features at high temporal resolution, as was recently demonstrated for lung tumours [Yun et al., 2012]. We have constructed a prototype integrated MRI-linac that aims to implement such MR-guidance [Raaymakers et al., 2009]. This prototype was recently upgraded from a table-mounted accelerator with cast

block collimators to a continuously rotating gantry-based system with a 160 leaves MLC derived from Agility (Elekta AB, Stockholm, Sweden) to shape the radiation beam.

In this paper we show the technical feasibility of tracked radiation delivery where, for the first time, the tracking signal is derived from on-line MR imaging. To this end, the previously established MRI/linac interfacing framework [Crijns et al., 2011b] is extended with a dynamic MLC controller. The updated framework is then used to perform tracked radiation delivery to a phantom that exhibits 1D motion, with feedback from fast 1D MR imaging.

6.2 Materials and Methods

Experimental set-up

All experiments were carried out using our MRI-linac prototype. It consists of a modified 1.5 T Achieva (Philips Best, The Netherlands) MRI system [Raaymakers et al., 2009], integrated with a gantry-based 6 MV linear accelerator (Elekta AB, Stockholm, Sweden). The source-to-surface distance in our prototype is approximately 1.47 m.

Our experimental setup consisted of a cart positioned in the scanner isocentre. It carried radiation sensitive GafChromic EBT2 film under a layer of water equivalent build-up material as well as a box-like MR phantom filled with CuSO₄-doped water (Figure 6.1). The cart was driven by an electric motor outside the scanner bore (at approximately 2.5 m from the isocentre) via a rigid crankshaft mechanism. The motor was fed by a programmable power supply attached to a digital-analog converter to allow prescription of arbitrary periodic 1D motion of the cart. Motion

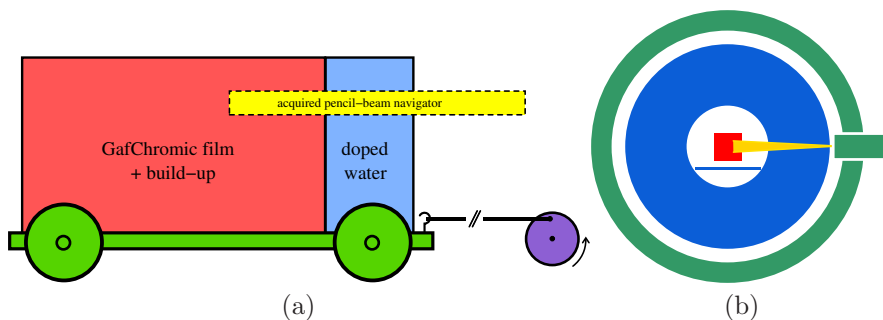


Figure 6.1: Schematic depiction of the experimental set-up used. (a) Cart (green) with dosimetric setup (red) and MR phantom (blue), imaging volume covered by the pencil-beam navigators (yellow) and electric motor/crankshaft mechanism (purple). (b) Set-up inside the scanner (view from front into the cylindrical scanner bore). The scanner and table are depicted in blue, the phantom is indicated in red, the gantry-mounted linac is shown in green and the beam is denoted by the yellow triangle. The motion of the cart is perpendicular to the plane of the drawing.




Figure 6.2: Example of pencil-beam navigator signal acquired with a moving phantom; each column corresponds to one navigator. The green line denotes the phantom edge detected by the forward difference method.

of the cart was along the long axis of the scanner bore and the MLC leaf travel direction, while being perpendicular to the radiation beam (see Figure 6.1b). Motion tracking via the MR scanner was realised through edge detection on acquired pencil-beam navigators (Figure 6.2). To this end, a forward difference computation is carried out on each navigator; the position of the maximum forward difference found is regarded the edge position of the phantom.

MRI/linac interface

Real-time communication with the MR scanner was realised through the External Control interface [Ries et al., 2010; Smink et al., 2011]. This interface relies on the TAO implementation of CORBA [Schmidt et al., 1998] for the underlying computer-to-computer communications. It facilitates communication of the acquired data to an external PC with minimal lag-time (~ 10 ms).

We implemented our own MLC controller that runs on an MLC computer (which has an Intel®Core™2 Duo 3.00GHz CPU, 2GB of RAM and runs Ubuntu Linux) on the gantry (see Figure 6.3). This controller communicates with the MLC electronics and runs a dynamic positioning feed-back loop for all MLC leaves, relying on a video stream (25 fps) to obtain the instantaneous positions of MLC leaves.

An interfacing application (running on an interfacing computer that has an Intel®Core™ i5-2400 3.10GHz CPU, 8GB of RAM and runs Ubuntu Linux) was implemented to process data acquired by the MR scanner to derive the position of the phantom in the scanner bore and subsequently send a rigid aperture translation to the MLC controller. The interfacing computer and the MLC computer were synchronised with respect to their system times to enable investigation of latencies based on timestamped data. These computers communicate through a wireless ethernet connection (see Figure 6.3). CORBA was used to facilitate network communication.

Experiments

The phantom was subjected to a 4×2 cm² radiation field

- without imposed motion, as a reference;




Figure 6.3: Physical layout of the MR scanner to MLC interfacing architecture used in our experiments.

- while the cart is moving (sinusoidal motion, amplitude 2.5 *cm* and period 5 *s*), but without using any motion compensation;
- while the cart is moving and with motion compensation by aperture translation based on the signal derived from pencil-beam navigators.

For all situations 2000 monitor units of radiation were delivered to the phantom. Our set-up is such that the motion can be compensated by aperture translation along the leaf-travel direction. By motion compensation, we mean that the MLC controller is notified of an update in desired aperture position as soon as this information is available; it will then try to realise the desired aperture as fast as possible.

6.3 Results and discussion

Figure 6.4(a) shows the dose distribution obtained on a static phantom. A rectangular dose distribution is obtained, that has a penumbra in leaf-travel direction of approximately 4 *mm* (defined between the 20 and 80% points).

Unsurprisingly the dose distribution is smeared when delivery takes place on a moving target without motion compensation (Figure 6.4(b)). Dose is spread over almost double the area of the statically delivered distribution, increasing the dose locally outside the area of the prescription by up to 150% (Figure 6.5a, top). The flat dose plateau delivered on the static phantom is lost and the maximum delivered dose at any point within the original field is reduced to approximately two thirds of the original value (see also Figure 6.7a).

When tracked radiation delivery is applied to account for the phantom motion, a much sharper dose distribution is obtained (Figure 6.4(c)). The plateau is restored and is at a level very close to that delivered to the static phantom. A penumbra of approximately 9 *mm* (a little more than twice that of the statically acquired distribution) remains but it is obviously much smaller than the smearing induced by completely uncompensated motion (Figure 6.5a).

We have implemented tracking as an aperture shift carried out as soon as new information on the target position is available. Evidently, this will cause latency between phantom position and MLC aperture; indeed, a close look at the




Figure 6.4: Results for MRI-guided tracked radiation delivery. Dose distributions delivered on a (a) static phantom, (b) moving phantom without motion compensation and (c) moving phantom while tracking the motion with the MLC aperture. Motion was in the left/right direction. Colour coding is the same for all three figures and given by the colour bar on the right.

MLC controller log files reveals a latency of approximately 150 *ms* in the tracking of the sinusoidal motion applied here (Figure 6.5b), on top of the acquisition/communication lag of circa 50 *ms*. This latency still introduces an over dosage to the enlarged penumbra and an under dosage in the plateau area of the distribution. Note, however, that we have not enlarged any margins in this case, which is what one would usually do to cope with motion.

The delivered dose distribution in the case of motion can be calculated by combining the statically delivered dose distribution with a PPDF. For the uncompensated case, the target motion is a sinusoid and the PPDF is detailed by

$$P(x) = \begin{cases} \frac{C}{\sqrt{A^2-x^2}}, & |x| \leq A, \\ 0, & \text{elsewhere,} \end{cases} \quad (6.1)$$




Figure 6.5: (a) Relative dose difference maps of uncompensated (top) and tracked (bottom) delivery compared to static delivery (scale is for both images). (b) Required and recorded MLC leaf position from the MLC log file during a tracking experiment.

with C a scaling constant such that $\int_{-\infty}^{\infty} P(x)dx = 1$ and A the motion amplitude [Bortfeld et al., 2002]. For the tracked delivery case, the relative motion between target and aperture (i.e., residual uncompensated motion) is what causes blurring. Assuming that the lag causes only a phase shift, the relative motion function is again a sinusoid, which has the same frequency but a lower amplitude than the target motion. In this case the residual motion amplitude was 5 mm . The resulting PPDFs are shown in Figure 6.6a. The dose distribution delivered on the static target is convolved with these probability density functions to obtain a prediction for the dose delivered in the uncompensated and tracked delivery cases. Results are shown in figures 6.6b and 6.6c respectively. To be able to compare

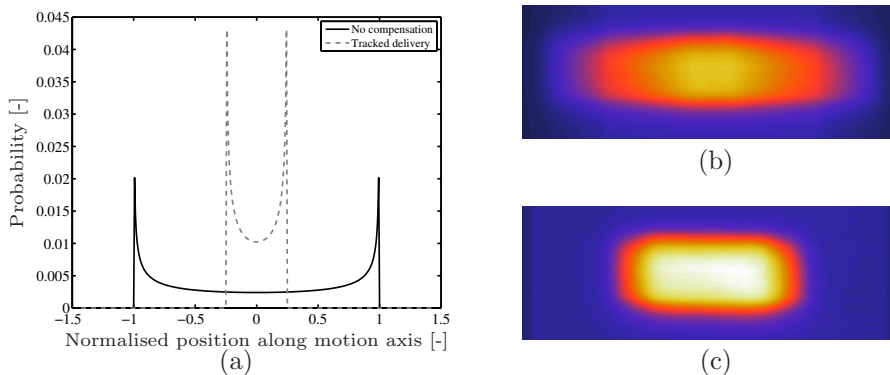


Figure 6.6: (a) motion induced PPDFs for full range motion as in the uncompensated case (solid line) and residual motion in the tracked delivery case (dashed line). (b) calculated delivered dose for the case without compensation. (c) calculated delivered dose for tracked delivery.

these calculated distributions to the measured ones, profiles are obtained along the centre of the distributions in the direction of phantom motion. Results are given in Figure 6.7. Close agreement with measurements is obtained in both cases which




Figure 6.7: Profiles along the centre of the radiation field for the (a) un-compensated and (b) tracked cases. The solid lines represent measured data whereas the dashed lines represent calculated dose distributions.

illustrates the value of recording the target motion during delivery and validates the latency found in the tracked delivery case.

As motion of abdominal organs is generally much more complicated than the 1D rigid translation used here, more intricate forms of MRI-based feedback are needed, that are able to provide information on deformations and 3D motion patterns. The high intensity focussed ultrasound (HIFU) community pursues optimised imaging strategies in order to tackle this problem (cf. Yuan et al. [2012] and references therein). Alternatively, since acquiring more data generally causes longer latencies, a 3D view of the instantaneous anatomy may be reconstructed based on a model and some lower dimensional data acquired at high temporal resolution. For example, the 3D target location may be derived from registration of 2D images against pre-treatment 3D data sets [Bjerre et al., 2011], or 3D models can be updated using information derived from 1D measurements [Nguyen et al., 2009]. This is especially useful for semi-periodic motion such as that caused by respiration. Such a model-based approach to 3D tracking may benefit from other inputs like the marker-based methods mentioned in the introduction or even ultrasound [Schwartz and McDannold, 2012].

To improve the tracking performance, latency compensation algorithms can be employed that predict the target position based on its trajectory so far. Having such predictor becomes especially important when the speeds in the system become larger. There are several options for such prediction algorithms, including state-based [Kalet et al., 2010; Yang et al., 2009] and kernel density estimation-based [Ruan, 2010] prediction, on-line support vector regression [Ernst and Schweikard, 2009], neural networks and Kalman filtering [Sharp et al., 2004]. The focus so far has been on predictions based on data acquired at high sampling rates compared to the time lag to compensate. In the case of MRI-guided radiotherapy however,

we may have to deal with predictions for latencies on the order of the image acquisition time, especially when using 2D and 3D imaging.

Accommodating tracking for 2D motion is much more complicated than the 1D tracking studied here, since leaves have to be opened and closed completely as the aperture is shifted [McQuaid and Webb, 2006; Rangaraj et al., 2008]. The head rotation in our prototype MRI-linac is fixed with the leaf travel direction along the cranio-caudal axis of the patient. For most targets, this coincides fairly well with the main axis of translation induced by respiratory motion. In general there will be modes of the motion that are oriented more perpendicular to the leaf travel direction. Clever strategies are needed to make the impact of fixed head rotation as small as possible, for example as proposed by McMahon et al. [2008].

In this paper we have looked at aperture tracking for only one field. In clinical practise IMRT plans are used for most sites in curative treatments, which consist of a myriad of fields distributed over several angles. When the anatomy exhibits rigid body translation as a whole, synchronised rigid shifts of each IMRT field will cause the resulting dose distribution to be equal to the planned distribution (equivalent to patient set-up correction). This is not at all self-evident in case of deformations and not even for rotations. In such cases we must keep track of the dose delivered up to a certain time-point and account for it during the rest of the delivery process. This greatly advocates the use of on-line MR imaging for dose accumulation as shown for rigid motion in Figure 6.6. As on-line compensation for deformations is a very challenging undertaking, a first step would be to accumulate the dose delivered in one fraction and compensate as good as possible for the deformation induced delivery errors in the next fraction.

6.4 Conclusion

We have demonstrated the feasibility of tracked radiation delivery to a moving phantom, based on target position feedback from on-line MR imaging. More generic tracking scenarios require advanced MR imaging, latency compensation, and on-line dose accumulation and (re-)planning tools. These are subjects of future investigations.

Summary and discussion

Image-guided radiotherapy has the potential to increase success of treatment by decreasing uncertainties concerning tumour position and shape. MRI is the modality of choice when it comes to imaging for tumour delineation and characterisation, set-up correction, treatment plan adaptation, response monitoring and intra-fraction motion compensation.

To boost the possibilities of image guidance by providing images with superior soft-tissue contrast *during* treatment, integrated diagnostic quality MRI functionality with a radiotherapy linear accelerator is pursued. To this end, an integrated MRI-linac was built in cooperation with Elekta AB (Stockholm, Sweden) and Philips (Best, The Netherlands). It features a 6 MV standing wave linear accelerator supported by a circular, continuously rotating, gantry around a custom 1.5 T cylindrical MRI scanner with a split gradient coil.

Solutions to two technical challenges related to MRI/linac integration are investigated in this thesis. First, the use of MR images in radiotherapy can be hindered by geometrical distortions due to magnetic field inhomogeneity problems, which is the subject of chapters 2 and 3. Secondly, the MRI and linac must be interfaced seamlessly to allow advanced types of radiation delivery which compensate for target motion based on feedback from MR images. This subject is discussed in chapters 4 to 6.

In Chapter 2 a method for fast retrospective corrections of magnetic field inhomogeneity induced image distortions, based on acquired magnetic field maps, is presented and implemented.

Using a square phantom containing a regular grid, a measure of residual deformation after correction was established. Next, the method was applied *in vivo* to the case of pelvic body contour extraction for prostate radiotherapy treatment planning. Gradient echo images were acquired with varying acquisition bandwidths and field maps were acquired using a multi-echo gradient echo approach.

With the phantom experiments, at very low gradient strength (which leads to large deformations) a maximum displacement of 2.9 *mm* was shown to be reduced to 0.63 *mm* after correction was applied. In the body contour experiments, again with low gradient strengths, distortions of up to 6 *mm* could be reduced to 2 *mm*. All results were provided within a lag time of 8 *ms*, facilitating use of the method in real-time on-line applications.

An alternative to the post-processing strategy is given in Chapter 3. In conventional MR acquisitions that use frequency encoding, off-resonance effects cause discrepancies between imaging and physical coordinate systems. A proposal to mitigate this effect by using only phase encoding and to reduce the longer acquisitions by under-sampling and regularised reconstruction is demonstrated.

To illustrate the performance of this acquisition in the presence of off-resonance phenomena (global magnetic field offset, chemical shift and magnetic field inhomogeneity induced by susceptibility variations) phantom and *in vivo* images were acquired using SE and purely phase encoded sequences. The latter was implemented by a SE type spectroscopic imaging sequence. Data was retrospectively under-sampled and reconstructed iteratively.

We observe accurate geometries in purely phase encoded images for all three investigated off-resonance phenomena, whereas SE images of the same phantoms display image distortions. Regularised reconstruction yields accurate phantom images under high acceleration factors, which is possible due to the simple nature of the phantom used. *In vivo* images were reconstructed faithfully while using acceleration factors up to four.

With the proposed technique inherently undistorted images with one-to-one correspondence to world coordinates can be obtained. It is a valuable tool in geometry QA, treatment planning and on-line image guidance. Under-sampled acquisition combined with regularised reconstruction can be used to accelerate the acquisition while retaining geometrical accuracy.

In Chapter 4, an architecture is proposed to realise a solution to the interfacing problem. Building on CORBA for inter-process communication, a class structure is implemented that facilitates data transfer between the MR scanner and a third-party, as well as scan geometry updates. Controllers are implemented for the linac sub devices (gantry, radiation beam, MLC), together with clients that can request the controllers over ethernet to perform certain actions. The clients are wrapped in user interfaces to present an end-user with one integrated look-and-feel.

Large respiration-induced semi-periodic target excursions hamper treatment of cancer of the abdominal organs. Methods to compensate for such motion in real-time are gated and tracked radiation delivery.

In Chapter 5 gated radiation delivery is realised in a proof-of-concept experiment, based on the previously established interfacing framework.

The position of a phantom subjected to 1D periodic translation (sinusoidal motion with a 25 *mm* amplitude) was tracked with the MR scanner by means of pencil-beam navigators. Based on the time-resolved position of the phantom, gated radiation delivery to the phantom was realised. Dose distributions for dy-

dynamic delivery conditions with gating windows varying between 5 and 25 *mm* were recorded on GafChromic film.

The similarity between dynamically and statically obtained dose profiles gradually increased as the gating window was decreased. With gating windows of 5 *mm* we obtained sharp dose profiles. The gating implementation was validated by comparing measured dose profiles to theoretical profiles calculated using knowledge of the imposed motion pattern. Excellent correspondence was observed. At the same time, we showed that real-time on-line reconstruction of the accumulated dose can be performed using time-resolved target position information. This facilitates plan adaptation not only on a fraction-to-fraction scale but also during one fraction, which is especially valuable in highly accelerated treatment strategies. Moreover, reconstruction of the delivered dose is required to establish accurate dose-response relations.

Finally, the feasibility of tracked radiation delivery is demonstrated in a similar proof-of-concept experiment in Chapter 6.

Again, the position of a phantom subjected to 1D periodic translation was tracked using a pencil-beam navigator MR sequence. Synchronisation of the motion of a small MLC aperture to the motion of the phantom was realised by means of the previously established interfacing framework. Dose distributions for various delivery conditions were recorded on GafChromic film.

Without motion a sharply defined dose distribution is obtained, whereas considerable blur occurs for delivery to a moving phantom. With compensation for motion (i.e., tracked delivery), the sharpness of the dose distribution greatly improved. The total latency in our motion management architecture is approximately 200 *ms*. The combination of the recorded phantom and aperture positions with the planned dose distribution enabled reconstruction of the delivered dose in all cases, which illustrates the promise of on-line dose accumulation and confirms that latency compensation could further reduce residual blurring.

Overall, a proof of concept was given for some strategies for on-line MRI-guided radiotherapy. The feasibility to overcome several technical issues has been demonstrated. Techniques have been developed to deal with MR image distortions, or to prevent them from occurring at all. A software architecture was implemented to facilitate real-time on-line MRI-guidance. A first implementation of IMRT as well as gated and tracked treatment delivery was given. Further research is needed to advance the technical prototype MRI-linac to the status of a clinical prototype.

7.1 General discussion

Geometric fidelity

The possible treatment accuracy for an MRI-linac depends on a number of factors. Mechanical alignment and calibration of the individual parts is obviously important, as in any linear accelerator. Specific for the MRI-linac is the fact that one must rely on MR images for set-up verification and treatment. The geometric

fidelity of the images is therefore of vital importance.

The post-processing strategy to correct for image distortions induced by field inhomogeneity presented in Chapter 2 has been shown to yield a reasonable performance. Images corrected using such a scheme can comfortably be used for therapy guidance. However, as was pointed out, the method requires the acquisition of a magnetic field map, which causes some overhead, especially for 3D scans.

Whether it is really necessary to correct an image data set for distortions depends on the scenario. In many cases distortions can be mitigated simply by using a sufficiently high acquisition bandwidth (i.e., high read-out gradient strength). For example, in our department uncorrected MR images are used for delineation of prostate tumours and associated OARS. What gradient strength is sufficient in a specific imaging scenario may be explored by measurements (i.e., acquisition of field maps) in volunteers and patients or calculations based on segmented CT scans [Marques and Bowtell, 2005; Wachowicz et al., 2010]. In cases where gradient strength is not a solution (e.g., due to limited SNR or the chosen read-out strategy) the correction strategy may be applied.

On the other hand, having an overhead on the order of one minute to be certain of the accuracy of the images is acceptable, especially in case of 3D scans that may take a few minutes to acquire as well.

An alternative for the post-processing technique described in Chapter 2 is the acquisition solution proposed in Chapter 3. The single point imaging technique as presented there is not yet suitable for clinical use. Especially the long acquisition times prevent a swift clinical introduction. Still, validation for the above mentioned experiments can be done based on the proposed technique. The use of an FID-based sequence (instead of echo-based), as was recently demonstrated by Bakker et al. [2012]; van Gorp et al. [2012a], together with speed-up from parallel imaging and under-sampling in combination with regularised reconstruction brings the technique another step closer to this goal.

So far we have only dealt with in-plane image distortions; discrepancies in slice selection are not considered. These errors have two causes, that have to be dealt with in different ways.

Slice selection is influenced by the non-linearity of the gradient coils. Specifically, off-axis slices are curved towards the edge of the FOV instead of exactly planar (see Figure 7.1 for an example). This is exactly the same effect that also happens in-slice. When the field gradient non-linearity is exactly known, these in-slice distortions can be de-warped [Doran et al., 2005; Moerland, 1996]. MR vendors implement this procedure in their products. For 3D scans, distortions due to this effect in the slice encoding direction can be de-warped as well. However, in the case of 2D imaging this is not possible, since it requires information from outside a slice. When a therapeutic intervention is carried out solely based on image feedback, but the intervention is not visible in the images (as is, so far, the case with MRI-guided radiotherapy), one must account for this effect. Computations that take information from 2D images as a basis should work with 3D coordinates for all the pixel coordinates, where these coordinates do not lie in a single plane.




Figure 7.1: Example of slice select distortions due to non-linearity of the gradient coils. This is an example for a slice at $x = 10\text{ cm}$. The colour coding represents the error from a planar slice at the chosen x -coordinate. All numbers are give in metres.

This is necessary especially for computations on structures near the edge of the FOV, where the effect can amount to centimetres. In the centre of the FOV, gradient non-linearity should not be an issue.

Secondly, similar to in-plane image distortion, the slice selection process is susceptible to resonance frequency variations, caused by chemical shift differences and \mathbf{B}_0 inhomogeneity. Variations of susceptibility in the subject therefore do not only cause distortions in the frequency encoding direction, but warps in the slice direction as well. This effect occurs on top of the curving due to gradient field non-linearity. A solution is to acquire 3D information instead of only a 2D slice and to use a field map to account for slice selection errors.

Radiotherapy planning relying on MRI only

Eventually, MRI might be able to completely replace CT-based planning. But before this becomes a reality, several hurdles must be taken. An important reason to continue taking CT scans of patients, even when geometric fidelity is under control and delineation would take place based on MR images, is that the CT data set represents information regarding the attenuation of radiation, which is necessary for accurate dose computations. Such information is not readily available from MR images.

In an attempt to overcome this limitation and introduce MR-based planning, one could simply ignore the influence of variations in attenuation of tissues. It has recently been demonstrated that this has little influence on planning accuracy for the prostate [Lee et al., 2003] and even for the brain [Prabhakar et al., 2007]. Addi-

tionally it was demonstrated that for prostate cancer treatment planning, physical path length information is more important than knowledge of electron densities [Kerckhof et al., 2010]. In such schemes, abandoning CT for planning in favour of MR becomes feasible, even when electron density information is completely absent.

Alternatively, it may soon become possible to derive attenuation maps from MR images. A similar problem is encountered in hybrid PET/MRI scanners, where an attenuation map for 511 *KeV* photons must be present in order to perform accurate reconstruction of the acquired PET data. It was recently demonstrated that digitally reconstructed radiographs (DRRs) can be generated based on a multi-echo ultra-short T_E (UTE) acquisition combined with Dixon reconstruction [Kotys-Traugber et al., 2012]. 4-Class tissue segmentation was accomplished based on a similar sequence, from which attenuation maps were reconstructed [Berker et al., 2012]. Alternatively a set of standardised T_1 and T_2 weighted scans in combination with UTE acquisitions may be used, in conjunction with a carefully calibrated model, to derive electron density maps [Johansson et al., 2011]. T_2 maps contain information on electron densities as well [Keereman et al., 2010].

Applying the UTE-based strategies to the abdomen is not straightforward. Respiratory and cardiac motion introduce artifacts in relatively long acquisitions. These challenges must be overcome before an MR-based attenuation map can be used for dose computations.

MRI-guidance for hadron therapy

Besides photon beams, proton or heavy ion beams can also be used for radiotherapy. The hadron beam depth dose profile increases to a sharp maximum (the Bragg peak), followed by a rapid drop-off. This is often presented as a major advantage over photon beams, which deposit a large amount of dose relatively close to the entry surface and have a much slower drop-off. However, at the same time range uncertainty is a very important limitation in hadron therapy [Unkelbach et al., 2007]. The sharpness of the Bragg peak has serious consequences when the exact target depth is unknown; a slight over- or underestimation of the depth leads to large under or over dosage. The ability to accurately measure these errors is vital. Some rather preliminary results, although they are in a sequential experiment, indicate that the effect of proton radiation on tissue can be made visible using MRI [Secon et al., 2009]. Without integrated therapy and imaging modalities, the suggested range verification [Gensheimer et al., 2010] can only be done retrospectively however. It is therefore of great value to investigate the potential of on-line MRI-guidance for proton therapy as well as hadron therapy in general. Although Lorentz forces play a significant role in MRI-guided proton therapy [Wolf and Bortfeld, 2012], these effect can be dealt with in the planning phase in the same manner as we do for photon therapy in magnetic fields [Raaymakers et al., 2008]. Having an imaging modality that could visualise the treatment beam in real-time during therapy would permit hadron beam therapy to make an enormous leap forward.

Intra-fraction motion management

One of the important goals in the developments around the MRI-linac is to realise treatment delivery adapted to organ motion. The experiments presented in chapters 5 and 6 are a first exploration of this concept.

Tracked radiation delivery is a complicated task, that gets more difficult as target and OAR motion becomes more complex. We can split intra-fraction organ motion in three categories: rigid and periodic motion (as is for example experienced by the kidney), non-rigid and periodic motion (liver) and non-rigid and non-periodic motion (e.g., as experienced by the pancreas, which moves due to breathing and due to bowel motion as well). For each category, tracked radiation delivery comes with a unique set of computational challenges (i.e., image registration and plan adaptation).

From a computational point of view, tracking rigid motion that repeats itself in a periodic fashion is the easiest. This requires rigid registration of on-line images against pre-treatment images to transfer contours, which is a well-developed and relatively easy to validate technique. The periodicity of the motion can relatively easily be employed by prediction algorithms, that improve the synchronisation of the aperture with the motion. Rigid motion of the target and the OARs together can be compensated by rigidly manipulating the plan in many cases.

Non-rigid motion is more difficult to deal with for a number of reasons. From the image registration point of view, dealing with deformations is non-trivial. Specifically, the validation and QA of deformable registration is difficult. Moreover, the registration algorithm must be robust against various scenarios. Secondly, from a plan adaptation perspective, non-rigid motion is much more difficult to cope with. Most likely, non-rigid motion requires re-planning. When re-planning on-line for tracking, accumulated dose which has been delivered to the patient up to a certain point in time must be accounted for. This requires adequate dose accumulation tools and specialised planning algorithms. Although non-rigid motion is difficult to cope with, it is also the kind of motion for which an MRI-linac is most valuable. Currently implemented tracked delivery strategies rely on markers, detected in kV or MV images, from which deformations are difficult to deduce.

Finally, there is the class of non-rigid and non-periodic motion. Difficulties here arise primarily because the therapy synchronisation (i.e., the MLC control) has to deal with movement that is hard to predict and therefore imposes stringent latency constraints. Pattern recognition strategies may provide a point of departure for further research [Ross et al., 2008].

Although it was argued that tracked delivery is more efficient than gated delivery, the latter has some clear advantages when we think about clinical introduction. Gated radiation delivery is applied by several groups already. QA of this technique is relatively straightforward; in gated delivery the anatomy is assumed frozen in the gating window, such that plan adaptations are not necessary and established QA procedures still apply.

Besides a smoother clinical introduction from a QA perspective, gated delivery is

also easier to realise technically. The imaging required for gating is less complex; for example pencil-beam navigators can be used to monitor the position of a critical organ/target boundary. For the difficult to handle class of non-rigid and non-periodic motion, gated delivery is a very good option to keep margins small while realising steep dose gradients. Trailing approaches may be used to correct for slow baseline variations [Pepin et al., 2011]

High spatiotemporal resolution imaging

In motion compensated therapy delivery, the aim is to completely account for motion induced errors in 3D. This requires high spatiotemporal resolution MR images. Currently a number of methods exist to retrospectively reconstruct such a set, but these are insufficient for therapy guidance. 3D imaging at high spatiotemporal resolution, while maintaining an adequate SNR, is not yet possible. Instead, we may choose to look at a smaller amount of data that can be acquired at high frame rate and use computational methods to obtain the needed 3D information. One way to accomplish this is to acquire high quality 3D data sets in both in- and expiration, on-line but prior to the treatment. During treatment, orthogonal 2D scans are acquired at high temporal resolution. By registering the 2D slices to the pre-acquired 3D data sets, on-line target contour propagation can be realised under a number of assumptions [Bjerre et al., 2011]. Furthermore, Köhler et al. [2011] have successfully demonstrated the use of 2D imaging at 10 Hz in combination with pencil-beam navigators to compensate for through-plane motion in MR-guided HIFU therapy of abdominal organs.

Instead of looking at an information subset in image space, we may also choose to record less information and use a subset of fully sampled Fourier space to reconstruct a full 3D image data set. Parallel imaging permits image reconstruction from a part of the raw data and can hence speed up the acquisition process [Griswold et al., 2002; Pruessmann et al., 1999; Sodickson and Manning, 2005]. This strategy takes advantage of multi-element coil arrays. CS techniques use under-sampling as well, but incorporate sparsity constraints on the image in the reconstruction process to suppress artifacts [Lustig et al., 2007]. Although high acceleration factors were presented for phantom images in Chapter 3, more modest scan time reductions, in the order of those realisable by parallel imaging, can be obtained for *in vivo* imaging (e.g., as illustrated by [Chartrand, 2009; Gamper et al., 2008; Kim et al., 2009; Lustig et al., 2007; Otazo et al., 2010]). It is highly likely that combinations of these approaches (e.g., as presented by Liang et al. [2009]; Liu et al. [2008]; Lustig et al. [2009]) facilitate high spatiotemporal resolution imaging in the future. One should keep in mind, however, that image SNR is proportional to the amount of collected raw data and that image quality may be reduced when using these techniques. A regularisation function may be able to restore the SNR of the image, but this may suppress information as well. Moreover, finding a solution to these large inverse problems requires some time. GPU-based solutions exist that bring acceptable reconstruction times within reach [Uecker et al., 2010]. As long as reconstruction times exceed acquisition

times, CS is not a viable option in intervention scenarios.

When looking at a subset of the complete 3D anatomy, there is a good chance that one overlooks part of the motion or deformation. The information in this subset may be complemented by a selection of pre-acquired imaging data, based on input from other sensors as well. For example, biometric navigators have been derived from ultrasound measurements [Schwartz and McDannold, 2012]. A relation of sensor output a specific anatomy state can be pre-calibrated in a training scan, after which the obtained relation can be used in the actual imaging session. More concretely, image data acquired in real-time may be complemented by a selection of data from a pre-acquired set, which is determined based on real-time ultrasound measurements.

As an alternative sensor system, the RF receive coils used in an MR examination may be used for an additional purpose, other than MR signal reception. Some time ago, it was demonstrated that patient movement within the coil causes changes in the RF impedance match of the coil, which can be measured as variations in the reflected RF power [Buikman et al., 1988]. Contact-free monitoring of respiratory motion is thereby possible, which has already been shown to facilitate motion corrected image reconstruction [Graesslin et al., 2010]. The viability of this solution is boosted by the trend to create receive coils with an increasing number of channels. Each of the coils in the array has a small sensitive volume and thus variations of its impedance can be employed as an indicator of motion in that volume. It is important to note, however, that each of the coil elements has to be radiation transparent to function well in an MRI-linac environment.

Samenvatting

Medische beelden worden in de radiotherapie gebruikt om het behandelingsuccess te vergroten, doordat informatie hieruit onzekerheden met betrekking tot tumorpositie en -vorm verminderen. Als het gaat om tumorkarakterisatie, beeldvorming voor intekenen, response monitoring en intra-fractie bewegingscompensatie is MRI de beste afbeeldingsmodaliteit.

De mogelijkheden van beeldsturing worden nog groter wanneer MRI beelden van diagnostische kwaliteit beschikbaar zouden zijn *tijdens* de behandeling. Om dit mogelijk te maken is een prototype MRI-linac gerealiseerd, in samenwerking met Elekta AB (Stockholm, Zweden) en Philips (Best, Nederland). Dit prototype bestaat uit een 6 MV staande golf lineaire versneller opgehangen in een ring (gantry) die continue rotatie mogelijk maakt. Deze ring is gebouwd rond een gemodificeerd 1.5 T cilindrische MRI scanner met een gradient spoel uit twee delen.

In dit proefschrift worden twee uitdagingen, gerelateerd aan MRI/linac integratie, onder de loep genomen. Als eerste wordt het gebruik van MRI beelden in de radiotherapie bemoeilijkt door geometrische beeldnauwkeurigheden, veroorzaakt door inhomogeniteit van het magneetveld. Hoofdstukken 2 en 3 gaan op de aanpak van dit probleem in. Daarnaast moeten de MR scanner en de linac naadloos worden geïntegreerd om compensatie van beweging van het doelgebied, gebaseerd op MRI beelden, mogelijk te maken. Dit is het onderwerp van hoofdstukken 4 tot en met 6.

In Hoofdstuk 2 wordt een methode geïntroduceerd voor een snelle retrospectieve correctie van beeldverstoringen veroorzaakt door inhomogeniteit van het magneetveld. Deze methode is gebaseerd op het direct meten van die inhomogeniteit. Met behulp van een vierkant fantoom met daarin een regelmatig raster werd de overblijvende beeldverstoring na correctie geschat. Vervolgens werd de methode toegepast om de lichaamscontour van een prostaatankerpatiënt te corrigeren voor beeldverstoringen. In beide experimenten werden gradient echo beelden geacqui-

reerd met verschillende bandbreedtes. Magnetische velden werden afgebeeld met behuld van een multi-echo gradient echo sequentie.

Met de fantoomexperimenten werd geïllustreerd dat, bij zeer lage acquisitiebandbreedte, een maximale beeldverstoring van 2.9 mm gereduceerd kon worden tot 0.63 mm. Fouten in de lichaamscontour tot 6 mm konden worden verminderd tot 2 mm. De duur van de berekeningen was voor alle gevallen maximaal 8 ms, waardoor de methode goed toegepast kan worden in real-time applicaties.

Een alternatief voor deze post-processing methode wordt besproken in Hoofdstuk 3. In conventionele MR acquisities wordt frequentie encoding gebruikt, waardoor veldverstorings kunnen leiden tot verschillen tussen wereld- en beeldcoördinaten. In dit hoofdstuk wordt een demonstratie gegeven van het vervangen van frequentie encoding door fase encoding, waarmee de acquisitie ongevoelig wordt voor magneetveldverstorings. De inherent langere acquisitietijd wordt opgevangen door onderbemonstering van het signaal te combineren met een iteratieve geregulariseerde reconstructie.

Om de mogelijkheden van deze acquisitie bij verschillende veldverstoringsfenomen (f_0 -verschuiving, chemical shift en magneetveldinhomogeniteiten veroorzaakt door lokale susceptibiliteitsverschillen) te illustreren werden zowel van een fantoom als *in vivo* beelden verkregen met behulp van SE en pure fase encoding sequenties. De laatste sequentie werd geïmplementeerd via een SE type spectroscopie sequentie. De data hieruit werd retrospectief undersampled en iteratief gereconstrueerd.

Voor elk van de drie veldverstoringsfenomenen werden met pure fase encoding onverstoorte beelden verkregen, terwijl conventionele SE beelden verstoring lieten zien. Geregulariseerde iteratieve reconstructie leidde tot accurate beelden van het fantoom, zelfs bij hoge versnellingsfactoren. Dit was mogelijk door de eenvoudige geometrie van het gebruikte fantoom. *In vivo* beelden werden goed bevonden tot een versnellingsfactor van vier.

Met behulp van de voorgestelde techniek kunnen beelden geacquireerd worden waarvan het beeldcoördinatensysteem een inherent juiste relatie heeft met de wereldcoördinaten. Door verdere versnellingsstechnieken toe te passen kan de techniek uitgroeien tot een waardevol stuk gereedschap voor kwaliteitszorg, radiotherapieplanning en on-line beeldsturing.

In Hoofdstuk 4 wordt een software architectuur besproken waarmee MRI-geleide radiotherapie mogelijke gemaakt kan worden. Een klassestructuur werd geïmplementeerd waarmee een verbinding gelegd kan worden tussen de MR scanner en een andere PC. Via deze verbinding kan dataoverdracht plaatsvinden en kunnen updates van de scangeometrie worden doorgevoerd. Daarnaast worden in dit hoofdstuk controllers besproken voor de linacsubsystemen (gantry, MLC en bundel), samen met clients die de controllers tot bepaalde acties kunnen laten overgaan. Deze clients worden verpakt in user interfaces met het doel een gebruiker een geïntegreerde look-en-feel te kunnen presenteren.

De behandeling van tumoren wordt bemoeilijkt door semi-periodieke bewegingen veroorzaakt door de ademhaling van de patient, met name in het abdomen. Twee methoden om met dit probleem om te gaan zijn gating, alleen dan bestralen als

het doelgebied zich bevindt op een vooraf gedefinieerde positie, en tracking, het synchroniseren van de MLC beweging met de beweging van het doelgebied. In Hoofdstuk 5 wordt gating gerealiseerd en de technische haalbaarheid van dit concept gedemonstreerd aan de hand van een eenvoudig experiment. De positie van een fantoom dat een periodieke 1D beweging beschrijft (een sinusoïde met een amplitude van 25 mm) werd gevolgd met behulp van een 1D pencil-beam navigator sequentie. Gebaseerd op de terugkoppeling van deze positie werd het fantoom bestraald volgens een gating strategie. Op stralingsgevoelige film werden dosis distributies vastgelegd voor gating windows van 5 tot 25 mm. De overeenkomst tussen dosis distributies afgeleverd zonder en met beweging van het fantoom verbeterde naarmate het gating window kleiner gemaakt werd. Bij een gating window van 5 mm werd een scherp afgebakende dosisverdeling verkregen. De gating implementatie werd gevalideerd door gemeten dosisprofielen te vergelijken met theoretische profielen die werden berekend gebaseerd op kennis van het opgelegde bewegingspatroon. De verschillen hiertussen waren bijzonder klein. Tegelijkertijd werd hiermee geïllustreerd dat real-time reconstructie van de geaccumuleerde dosis mogelijk is op basis van informatie van de positie van het doelgebied. Zo'n reconstructie is bijzonder waardevol in versnelde behandelingschema's en bij het vaststellen van accurate dosis-effectrelaties.

Als laatste demonstreren we de haalbaarheid van tracking in een soortgelijk experiment in Hoofdstuk 6.

Opnieuw werd de positie van een fantoom dat een periodieke 1D beweging beschrijft gevolgd met behulp van een pencil-beam navigator sequentie. Synchronisatie van de beweging van een kleine MLC apertuur met de beweging van het fantoom werd gerealiseerd met behulp van het interfacing framework geïntroduceerd in Hoofdstuk 4. Dosis distributies voor de verschillende bewegingscondities werden vastgelegd op stralingsgevoelige film.

Wanneer het fantoom stilstaat wordt een scherp gedefinieerde dosis distributie verkregen. Bij aflevering op een bewegend fantoom wordt deze dosis aanzienlijk meer verspreid over een grotere regio. Compensatie van de beweging (oftewel, tracking) leidde tot aanmerkelijke verbetering van de scherpte van de dosis distributie. De totale responsietijd van het trackingsysteem is ongeveer 200 ms. De combinatie van de opgenomen fantoom- en apertuurposities maakt in alle bekeken gevallen een accurate reconstructie van de afgegeven dosis mogelijk. Dit illustreert wederom de mogelijkheden voor dosisaccumulatie die real-time beeldacquisitie met zich meebrengt. Bovendien laat dit zien dat ongecompenseerde overblijvende beweging de dosisafwijkingen volledig beschrijven. Het resultaat kan dus verder worden verbeterd door compensatie met behulp van voorspellende modellen en het verkorten van de responsietijd in het systeem.

Samenvattend wordt in dit proefschrift een demonstratie gegeven van enkele strategieën voor on-line MRI geleiding voor radiotherapie. Als eerste zijn er technieken beschreven om om te gaan met patient specifieke veldverstoringen die leiden tot beeldvervalsing, of deze helemaal te voorkomen. Daarnaast is een software architectuur gedefinieerd om MRI-geleide radiotherapie mogelijk te maken. Bewegingsgecompenseerde radiotherapie is gedemonstreerd op basis van MRI beelden in

gating en tracking scenarios. Daarnaast is voor deze gevallen de dosisaccumulatie die real-time beeldacquisitie mogelijk maakt geïllustreerd. Verdere ontwikkeling is nodig om het technische MRI-linac prototype tot een klinisch prototype uit te bouwen.

References

- E. E. Ahunbay, C. Peng, A. Godley, C. Schultz, and X. A. Li (2009). An on-line replanning method for head and neck adaptive radiotherapy. *Med. Phys.*, 36:4776.
- C. J. G. Bakker, H. de Leeuw, G. H. van de Maat, J. S. van Gorp, J. G. Bouwman, and P. R. Seevinck (2012). On the utility of spectroscopic imaging as a tool for generating geometrically accurate mr images and parameter maps in the presence of field inhomogeneities and chemical shift effects. *Magn. Res. Imag.*.
- B. J. Balcom, R. P. MacGregor, S. D. Beyea, D. P. Green, R. L. Armstrong, and T. W. Bremner (1996). Single-point ramped imaging with T1 enhancement (SPRITE). *J. Magn. Res.*, 123(1):131.
- L. N. Baldwin, K. Wachowicz, and B. G. Fallone (2009). A two-step scheme for distortion rectification of magnetic resonance images. *Med. Phys.*, 36:3917–3926.
- A. Bel, M. Van Herk, H. Bartelink, and J. V. Lebesque (1993). A verification procedure to improve patient set-up accuracy using portal images. *Radiother. Oncol.*, 29(2):253–260.
- S. M. Bentzen (2005). Theragnostic imaging for radiation oncology: dose-painting by numbers. *Lancett Oncology*, 6:112–117.
- Y. Berker, J. Franke, A. Salomon, M. Palmowski, H. C. W. Donker, Y. Temur, F. M. Mottaghy, C. Kuhl, D. Izquierdo-Garcia, and Z. A. Fayad (2012). MRI-based attenuation correction for hybrid PET/MRI systems: a 4-class tissue segmentation technique using a combined ultrashort-echo-time/Dixon MRI sequence. *Journal of Nuclear Medicine*, 53(5):796–804.
- S. D. Beyea, B. J. Balcom, I. V. Mastikhin, T. W. Bremner, R. L. Armstrong, and P. E. Grattan-Bellew (2000). Imaging of heterogeneous materials with a turbo spin echo single-point imaging technique. *J. Magn. Res.*, 144(2):255–265.
- T. Bjerre, T. Kim, P. Munk af Rosenschöld, M. Aznar, S. L. R. Larsen, and P. Keall (2011). A Feasible Real-Time Multi-plane MRI-Linac Guidance Strategy. In *Proceedings of the 54rd ASTRO Annual Meeting*. (accepted).
- G. Booch (1994). *Object-Oriented Analysis and Design with Applications*. 2nd edition.
- T. Bortfeld, K. Jokivarsi, M. Goitein, J. Kung, and S. B. Jiang (2002). Effects of intra-fraction motion on IMRT dose delivery: statistical analysis and simulation. *Phys. Med. Biol.*, 47:2203.
- T. Bortfeld, S. B. Jiang, and E. Rietzel (2004). Effects of motion on the total dose distribution. *Seminars in radiation oncology*, 14(1):41–51.
- D. Buikman, T. Helzel, and P. Roschmann (1988). The RF coil as a sensitive motion detector for magnetic resonance imaging. *Magn. Res. Imag.*, 6(3):281–289.
- H. Chang and J. M. Fitzpatrick (1992). A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. *IEEE Trans. Med. Imag.*, 11(3):319–329.
- R. Chartrand (2009). Fast algorithms for nonconvex compressive sensing: MRI reconstruction from very few data. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 262–265.
- B. Cho, P. R. Poulsen, D. Ruan, A. Sawant, S. Povzner, and P. J. Keall (2009a). DMLC tracking with real-time target position estimation combining a single kV imager and an optical respiratory monitoring system. *Int. J. Radiat. Oncol. Biol. Phys.*, 75(3):S24–S25.

- B. Cho, P. R. Poulsen, A. Sloutsky, A. Sawant, and P. J. Keall (2009b). First demonstration of combined kV/MV image-guided real-time dynamic multileaf-collimator target tracking. *Int. J. Radiat. Oncol. Biol. Phys.*, 74(3):859–867.
- B. Cho, P. R. Poulsen, A. Sawant, D. Ruan, and P. J. Keall (2011). Real-time target position estimation using stereoscopic kilovoltage/megavoltage imaging and external respiratory monitoring for dynamic multileaf collimator tracking. *Int. J. Radiat. Oncol. Biol. Phys.*, 79(1):269–278.
- Z. Cho, D. Kim, and Y. Kim (1988). Total inhomogeneity correction including chemical shifts and susceptibility by view angle tilting. *Med. Phys.*, 15:7.
- R. Colgan, J. McClelland, D. McQuaid, P. M. Evans, D. Hawkes, J. Brock, D. Landau, and S. Webb (2008). Planning lung radiotherapy using 4D CT data and a motion model. *Phys. Med. Biol.*, 53:5815.
- S. P. M. Crijs, B. W. Raaymakers, and J. J. W. Lagendijk (2011a). Real-time correction of magnetic field inhomogeneity-induced image distortions for MRI-guided conventional and proton radiotherapy. *Phys. Med. Biol.*, 56:289–297.
- S. P. M. Crijs, J. G. M. Kok, J. J. W. Lagendijk, and B. W. Raaymakers (2011b). Towards MRI-guided linear accelerator control: gating on an MRI accelerator. *Phys. Med. Biol.*, 56(15):4815–4825.
- H. C. J. de Boer and B. J. M. Heijmen (2001). A protocol for the reduction of systematic patient setup errors with minimal portal imaging workload. *Int. J. Radiat. Oncol. Biol. Phys.*, 50(5):1350–1365.
- J. F. Dempsey, D. Benoit, J. R. Fitzsimmons, A. Haghghat, J. G. Li, D. A. Low, S. Mutic, J. R. Palta, H. E. Romeijn, and G. E. Sjoden (2005). A Device for Realtime 3D Image-Guided IMRT. *Int. J. Radiat. Oncol. Biol. Phys.*, 63(Supplement 1):S202. Proceedings of the American Society for Therapeutic Radiology and Oncology 47th Annual Meeting.
- T. Depuydt, D. Verellen, O. Haas, T. Gevaert, N. Linthout, M. Duchateau, K. Tournel, T. Reynders, K. Leysen, M. Hoogeman, G. Storme, and M. De Ridder (2011). Geometric accuracy of a novel gimbals based radiation therapy tumor tracking system. *Radiother. Oncol.*, 98(3):365–372.
- L. Dietrich, T. Tücking, S. Nill, and U. Oelfke (2005). Compensation for respiratory motion by gated radiotherapy: an experimental study. *Phys. Med. Biol.*, 50:2405–2414.
- S. J. Doran, L. Charles-Edwards, S. A. Reinsberg, and M. O. Leach (2005). A complete distortion correction for MR images: I Gradient warp correction. *Phys. Med. Biol.*, 50(7):1343–1361.
- W. D. D’Souza, S. A. Naqvi, and C. X. Yu (2005). Real-time intra-fraction-motion tracking using the treatment couch: a feasibility study. *Physics in medicine and biology*, 50:4021.
- F. Ernst and A. Schweikard (2009). Forecasting respiratory motion with accurate online support vector regression (SVRpred). *Int J CARS*, 4(5):439–447.
- P. M. Evans, J. R. N. Symonds-Taylor, R. Colgan, G. D. Hugo, N. Letts, and C. Sandin (2010). Gating characteristics of an Elekta radiotherapy treatment unit measured with three types of detector. *Phys. Med. Biol.*, 55:N201–N210.
- B. G. Fallone, B. Murray, S. Rathee, T. Stanescu, S. Steciw, S. Vidakovic, E. Blosser, and D. Tyomofichuk (2009). First MR images obtained during megavoltage photon irradiation from a prototype integrated linac-MR system. *Med. Phys.*, 36:2084–2088.
- M. Feng, J. M. Balter, N. D. A. S, C. Y, T. L. Chenevert, and E. Ben-Josef (2009). Characterization of Pancreatic Tumor Motion Using Cine MRI: Surrogates for Tumor Position Should Be Used With Caution. *Int. J. Radiat. Oncol. Biol. Phys.*, 74(3):884–891.
- E. C. Ford, G. S. Mageras, E. Yorke, K. E. Rosenzweig, R. Wagman, and C. C. Ling (2002). Evaluation of respiratory movement during gated radiotherapy using film and electronic portal imaging. *Int. J. Radiat. Oncol. Biol. Phys.*, 52(2):522–531.
- T. Fox, E. L. Simon, E. Elder, R. H. Riffenburgh, and P. A. S. Johnstone (2007). Free breathing gated delivery (FBGD) of lung radiation therapy: Analysis of factors affecting clinical patient throughput. *Lung cancer*, 56(1):69–75.
- M. Fuss, B. J. Salter, S. X. Cavanaugh, C. Fuss, A. Sadeghi, C. D. Fuller, A. Ameduri, J. M. Hevezi, T. S. Herman, and C. R. Thomas (2004). Daily ultrasound-based image-guided targeting for radiotherapy of upper abdominal malignancies. *Int. J. Radiat. Oncol. Biol. Phys.*, 59(4):1245–1256.

- U. Gamper, P. Boesiger, and S. Kozerke (2008). Compressed sensing in dynamic MRI. *Magn. Reson. Med.*, 59(2):365–373.
- W. Gao, N. T. T. Huyen, H. S. Loi, and Q. Kemao (2009). Real-time 2D parallel windowed Fourier transform for fringe pattern analysis using Graphics Processing Unit. *Opt. Express*, 17(25):23147–23152.
- M. F. Gensheimer, T. I. Yock, N. J. Liebsch, G. C. Sharp, H. Paganetti, N. Madan, P. E. Grant, and T. Bortfeld (2010). In Vivo Proton Beam Range Verification Using Spine MRI Changes. *Int. J. Radiat. Oncol. Biol. Phys.*, 78(1):268–275.
- R. George, T. D. Chung, S. S. Vedam, V. Ramakrishnan, R. Mohan, E. Weiss, and P. J. Keall (2006). Audio-visual biofeedback for respiratory-gated radiotherapy: impact of audio instruction and audio-visual biofeedback on respiratory-gated radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 65(3):924–933.
- G. H. Glover (1999). 3D z-shim method for reduction of susceptibility effects in BOLD fMRI. *Magn. Reson. Med.*, 42(2):290–299.
- I. Graesslin, G. Mens, A. Guillaume, H. Stahl, P. Koken, P. Vernickel, P. Harvey, J. Smink, K. Nehrke, and P. Boernert (2010). Advancements in Contact-free Respiration Monitoring using RF Pick-up coils. In *Proceedings of the 18th Annual Meeting of ISMRM, Stockholm, Sweden*.
- M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. Med.*, 47(6):1202–1210.
- G. Groenendaal, M. R. Moman, J. G. Korporaal, P. J. van Diest, M. van Vulpen, M. E. P. Philippens, and U. A. van der Heide (2010). Validation of functional imaging with pathology for tumor delineation in the prostate. *Radiother. Oncol.*, 94(2):145–150.
- M. L. H. Gruwel, P. Latta, M. Tanasiewicz, V. Volotovskyy, M. Šramek, and B. Tomanek (2007). MR imaging of teeth using a silent single point imaging technique. *Applied Physics A: Materials Science & Processing*, 88(4):763–767.
- M. Guckenberger, A. Kavanagh, S. Webb, and M. Brada (2011). A novel respiratory motion compensation strategy combining gated beam delivery and mean target position concept—A compromise between small safety margins and long duty cycles. *Radiother. Oncol.*, 98(3):317–322.
- S. Gwynne, L. Wills, G. Joseph, G. John, J. Staffurth, C. Hurt, and S. Mukherjee (2009). Respiratory Movement of Upper Abdominal Organs and its Effect on Radiotherapy Planning in Pancreatic Cancer. *Clinical Oncology*, 21(9):713–719.
- E. M. Haacke, M. R. T. Thompson, R. W. Brown, and R. Venkatesan (1999). *Magnetic resonance imaging: physical principles and sequence design*. Wiley-Liss.
- M. Henning and S. Vinoski (1999). *Advanced CORBA® Programming with C++*. Addison-Wesley Professional.
- A. Hsu, N. R. Miller, P. M. Evans, J. C. Bamber, and S. Webb (2005). Feasibility of using ultrasound for real-time tracking during radiotherapy. *Med. Phys.*, 32:1500–1512.
- S. Hughes, J. McClelland, A. Chandler, M. Adams, J. Boutland, D. Withers, S. Ahmad, J. Blackall, S. Tarte, D. Hawkes, and D. Landau (2008). A comparison of internal target volume definition by limited four-dimensional computed tomography, the addition of patient-specific margins, or the addition of generic margins when planning radical radiotherapy for lymph node-positive non-small cell lung cancer. *Clinical Oncology*, 20(4):293–300.
- G. D. Hugo, N. Agazaryan, and T. D. Solberg (2003). The effects of tumor motion on planning and delivery of respiratory-gated IMRT. *Med. Phys.*, 30:1052–2525.
- ICRU 50 (1993). *Prescribing, recording and reporting photon beam therapy*. ICRU Bethesda Maryland.
- ICRU 62 (1999). *Prescribing, Recording and Reporting Photon Beam Therapy (Supplement to ICRU Report 50)*. ICRU Bethesda Maryland.
- D. A. Jaffray, J. H. Siewerdsen, J. W. Wong, and A. A. Martinez (2002). Flat-panel cone-beam computed tomography for image-guided radiation therapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 53(5):1337–1349.

- M. Jenkinson (2003). Fast, automated, N-dimensional phase-unwrapping algorithm. *Magn. Reson. Med.*, 49(1):193–197.
- P. Jezzard and R. S. Balaban (1995). Correction for geometric distortion in echo planar images from B_0 field variations. *Magn. Reson. Med.*, 34(1):65–73.
- A. Johansson, M. Karlsson, and T. Nyholm (2011). CT substitute derived from MRI sequences with ultrashort echo time. *Med. Phys.*, 38:2708.
- A. Kalet, G. Sandison, H. Wu, and R. Schmitz (2010). A state-based probabilistic model for tumor respiratory motion prediction. *Phys. Med. Biol.*, 55:7615–7631.
- Y. Kamino, K. Takayama, M. Kokubo, Y. Narita, E. Hirai, N. Kawawada, T. Mizowaki, Y. Nagata, T. Nishidai, and M. Hiraoka (2006). Development of a four-dimensional image-guided radiotherapy system with a gimbaled X-ray head. *Int. J. Radiat. Oncol. Biol. Phys.*, 66(1):271–278.
- P. A. Karasev, D. P. Campbell, and M. A. Richards (2007). Obtaining a 35x speedup in 2D phase unwrapping using commodity graphics processors. In *Radar Conference, 2007 IEEE*, pages 574–578.
- M. Karlsson, M. G. Karlsson, T. Nyholm, C. Amies, and B. Zackrisson (2009). Dedicated magnetic resonance imaging in the radiotherapy clinic. *Int. J. Radiat. Oncol. Biol. Phys.*, 74(2):644–651.
- P. J. Keall, V. R. Kini, S. S. Vedam, and R. Mohan (2001). Motion adaptive X-ray therapy: a feasibility study. *Phys. Med. Biol.*, 46:1–10.
- P. J. Keall, S. Joshi, S. S. Vedam, J. V. Siebers, V. R. Kini, and R. Mohan (2005). Four-dimensional radiotherapy planning for DM-LC-based respiratory motion tracking. *Med. Phys.*, 32:942.
- V. Keereman, Y. Fierens, T. Broux, Y. De Deene, M. Lonnew, and S. Vandenberghe (2010). MRI-based attenuation correction for PET/MRI using ultrashort echo time sequences. *Journal of Nuclear Medicine*, 51(5):812–818.
- E. M. Kerkhof, J. M. Balter, K. Vineberg, and B. W. Raaymakers (2010). Treatment plan adaptation for MRI-guided radiotherapy using solely MRI data: a CT-based simulation study. *Phys. Med. Biol.*, 55(16):N433–N440.
- E. M. Kerkhof, B. W. Raaymakers, M. Van Vulpen, B. A. Zonnenberg, J. L. H. R. Bosch, R. J. A. Van Moerselaar, and J. J. W. Lagendijk (2011). A new concept for non-invasive renal tumour ablation using real-time MRI-guided radiation therapy. *Brit. J. Urol. Int.*, 107(1):63–68.
- Y. C. Kim, S. S. Narayanan, and K. S. Nayak (2009). Accelerated three-dimensional upper airway MRI using compressed sensing. *Magn. Reson. Med.*, 61(6):1434–1440.
- K. Kitamura, H. Shirato, Y. Seppenwoolde, T. Shimizu, Y. Kodama, H. Endo, R. Onimaru, M. Oda, K. Fujita, S. Shimizu, and K. Miyasaka (2003). Tumor location, cirrhosis, and surgical history contribute to tumor movement in the liver, as measured during stereotactic irradiation using a real-time tumor-tracking radiotherapy system. *Int. J. Radiat. Oncol. Biol. Phys.*, 56(1):221–228.
- N. Koch, H. H. Liu, G. Starkschall, M. Jacobson, K. Forster, Z. Liao, R. Komaki, and C. W. Stevens (2004). Evaluation of internal lung motion for respiratory-gated radiotherapy using MRI: Part I: correlating internal lung motion with skin fiducial motion. *Int. J. Radiat. Oncol. Biol. Phys.*, 60(5):1459–1472.
- K. Koh, B. Lee, B. Kim, and J. Seo (2010). Maniwordle: Providing flexible control over wordle. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1190–1197.
- M. O. Köhler, B. D. De Senneville, B. Quesson, C. T. W. Moonen, and M. Ries (2011). Spectrally selective pencil-beam navigator for motion compensation of MR-guided high-intensity focused ultrasound therapy of abdominal organs. *Magn. Reson. Med.*, 66(1):102–111.
- S. S. Korreman, T. Juhler-Nøttrup, and A. L. Boyer (2008). Respiratory gated beam delivery cannot facilitate margin reduction, unless combined with respiratory correlated image guidance. *Radiother. Oncol.*, 86(1):61–68.
- M. S. Kotys-Traughber, B. J. Traughber, M. Meltsner, and R. F. J. Muzic (2012). MR Cortical Bone Imaging Using UTE for Digitally-Reconstructed Radiographs and Attenuation Correction at 3.0T. In *Proceedings of the 20th Annual Meeting of ISMRM, Melbourne, Australia*.
- K. Kuriyama, H. Onishi, N. Sano, T. Komiyama, Y. Aikawa, Y. Tateda, T. Araki, and M. Uematsu (2003). A new irradiation unit constructed of self-moving gantry-CT and linac. *Int. J. Radiat. Oncol. Biol. Phys.*, 55(2):428–435.

- J. J. W. Lagendijk, B. W. Raaymakers, A. J. E. Raaijmakers, J. Overweg, K. J. Brown, E. M. Kerkhof, R. W. Van der Put, B. Hårdemark, M. Van Vulpen, and U. A. Van der Heide (2008). MRI/linac integration. *Radiother. Oncol.*, 86(1):25–29.
- Y. K. Lee, M. Bollet, G. Charles-Edwards, M. A. Flower, M. O. Leach, H. McNair, E. Moore, C. Rowbottom, and S. Webb (2003). Radiotherapy treatment planning of prostate cancer using magnetic resonance imaging alone. *Radiother. Oncol.*, 66(2):203–216.
- H. Leung (2012). Tagxedo wordcloud generator. <http://www.tagxedo.com>.
- D. Liang, B. Liu, J. J. Wang, and L. Ying (2009). Accelerating SENSE using compressed sensing. *Magn. Reson. Med.*, 62(6):1574–1584.
- C. C. Ling, J. Humm, S. Larson, H. Amols, Z. Fuks, S. Leibel, and J. A. Koutcher (2000). Towards multidimensional radiotherapy (MD-CRT): biological imaging and biological conformality. *Int. J. Radiat. Oncol. Biol. Phys.*, 47(3):551–560.
- B. Liu, Y. M. Zou, and L. Ying (2008). Sparsesense: application of compressed sensing in parallel mri. In *International Conference on Information Technology and Applications in Biomedicine*, pages 127–130.
- D. A. Low, R. Stark, and J. F. Dempsey (2011). *Adaptive Motion Compensation in Radiotherapy*, chapter 12, pages 115–128. Imaging in Medical Diagnosis and Therapy. CRC Press.
- A. E. Lujan, E. W. Larsen, J. M. Balter, and R. K. Ten Haken (1999). A method for incorporating organ motion due to breathing into 3D dose calculations. *Med. Phys.*, 26(5):715–721.
- M. Lustig, D. Donoho, and J. M. Pauly (2007). Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.*, 58(6):1182–1195.
- M. Lustig, M. Alley, S. Vasanawala, D. L. Donoho, and J. M. Pauly (2009). L1 spir-it: autocalibrating parallel imaging compressed sensing. In *Proceedings of the 17th Annual Meeting of ISMRM, Honolulu, Hawaii, USA*, page 379.
- Y. Mao (2011). Reconstruction of binary functions and shapes from incomplete frequency information.
- J. P. Marques and R. Bowtell (2005). Application of a Fourier-based method for rapid calculation of field inhomogeneity due to spatial variation of magnetic susceptibility. *Concepts Magnetic Res.B*, 25(1):65–78.
- G. M. P. Masselli, S. Silvestri, S. Ramella, and L. Trodella (2009). Design and evaluation of a methodology to perform personalized visual biofeedback for reducing respiratory amplitude in radiation treatment. *Med. Phys.*, 36:1467–1472.
- R. McMahon, R. Berbeco, S. Nishioka, M. Ishikawa, and L. Papiez (2008). A real-time dynamic-MLC control algorithm for delivering IMRT to targets undergoing 2D rigid motion in the beams eye view. *Med. Phys.*, 35:3875.
- H. A. McNair, S. A. Mangar, J. Coffey, B. Shoulders, V. N. Hansen, A. Norman, J. Staffurth, S. A. Sohaib, A. P. Warrington, and D. P. Dearnaley (2006). A comparison of CT-and ultrasound-based imaging to localize the prostate for external beam radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 65(3):678–687.
- D. McQuaid and S. Webb (2006). IMRT delivery to a moving target by dynamic MLC tracking: delivery for targets moving in two dimensions in the beam’s eye view. *Phys. Med. Biol.*, 51:4819.
- G. J. Meijer, P. P. van der Toorn, M. Bal, D. Schuring, J. Weterings, and M. de Wildt (2012). High precision bladder cancer irradiation by integrating a library planning procedure of 6 prospectively generated sib imrt plans with image guidance using lipiodol markers. *Radiother. Oncol.*
- A. Y. Minn, D. Schellenberg, P. Maxim, Y. Suh, S. McKenna, B. Cox, S. Dieterich, L. Xing, E. Graves, K. A. Goodman, D. Chang, and A. C. Koong (2009). Pancreatic tumor motion on a single planning 4D-CT does not correlate with intrafraction tumor motion during treatment. *American journal of clinical oncology*, 32(4):364.
- P. Mistry, S. Braganza, D. Kaeli, and M. Leeser (2009). Accelerating phase unwrapping and affine transformations for optical quadrature microscopy using CUDA. In *Proceedings of 2nd Workshop on General Purpose Processing on Graphics Processing Units*, pages 28–37.

- M. A. Moerland, R. Beersma, R. Bhagwandien, H. K. Wijrdeman, and C. J. G. Bakker (1995). Analysis and correction of geometric distortions in 1.5 T magnetic resonance images for use in radiotherapy treatment planning. *Phys. Med. Biol.*, 40(10):1651–1664.
- M. A. Moerland (1996). *Magnetic Resonance Imaging in Radiotherapy Treatment Planning (chapter 3)*. PhD thesis, Universiteit Utrecht.
- D. S. Mohan, P. A. Kupelian, and T. R. Willoughby (2000). Short-course intensity-modulated radiotherapy for localized prostate cancer with daily transabdominal ultrasound localization of the prostate gland. *Int. J. Radiat. Oncol. Biol. Phys.*, 46(3):575–580.
- R. Mohan, X. Zhang, H. Wang, Y. Kang, X. Wang, H. Liu, K. K. Ang, D. Kuban, and L. Dong (2005). Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes. *Int. J. Radiat. Oncol. Biol. Phys.*, 61(4):1258–66.
- P. Munger, G. R. Crelier, T. M. Peters, and G. B. Pike (2000). An inverse problem approach to the correction of distortion in EPI images. *IEEE Trans. Med. Imag.*, 19(7):681–689.
- V. Murthy, Z. Master, P. Adurkar, I. Mallick, U. Mahantshetty, G. Bakshi, H. Tongaonkar, and S. Shrivastava (2011). ‘plan of the day’ adaptive radiotherapy for bladder cancer using helical tomotherapy. *Radiother. Oncol.*, 99(1):55–60.
- A. J. Nederveen, H. Dehnad, U. A. Van der Heide, R. J. A. Van Moorselaar, P. Hofman, and J. J. W. Lagendijk (2003). Comparison of megavoltage position verification for prostate irradiation based on bony anatomy and implanted fiducials. *Radiother. Oncol.*, 68(1):81–88.
- T. N. Nguyen, J. L. Moseley, L. A. Dawson, D. A. Jaffray, and K. K. Brock (2009). Adapting liver motion models using a navigator channel technique. *Med. Phys.*, 36(4):1061–1073.
- OMG (2012). CORBA specification. <http://www.omg.org/spec/CORBA/>.
- R. Otazo, D. Kim, L. Axel, and D. K. Sodickson (2010). Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion mri. *Magn. Reson. Med.*, 64(3):767–776.
- C. Ozhasoglu and M. J. Murphy (2002). Issues in respiratory motion compensation during external-beam radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 52(5):1389–1399.
- L. Papiez and D. Rangaraj (2005). DMLC leaf-pair optimal control for mobile, deforming target. *Med. Phys.*, 32:275–284.
- L. Papiez, D. Rangaraj, and P. Keall (2005). Real-time DMLC IMRT delivery for mobile and deforming targets. *Med. Phys.*, 32:3037.
- P. Parasoglou, A. J. Sederman, J. Rasburn, H. Powell, and M. L. Johns (2008). Optimal k-space sampling for single point imaging of transient systems. *J. Magn. Res.*, 194(1):99–107.
- E. W. Pepin, H. Wu, and H. Shirato (2011). Dynamic gating window for compensation of baseline shift in respiratory-gated radiation therapy. *Med. Phys.*, 38:1912–1918.
- J. Pouliot, A. Bani-Hashemi, J. Chen, M. Svatos, F. Ghelmansarai, M. Mitschke, M. Aubin, P. Xia, O. Morin, K. Bucci, M. Roach, P. Hernandez, Z. Zhen, D. Hristov, and L. Verhey (2005). Low-dose megavoltage cone-beam CT for radiation therapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 61(2):552–560.
- P. R. Poulsen, B. Cho, D. Ruan, A. Sawant, and P. J. Keall (2010). Dynamic multileaf collimator tracking of respiratory target motion based on a single kilovoltage imager during arc radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 77(2):600–607.
- P. R. Poulsen, J. Carl, J. Nielsen, M. S. Nielsen, J. B. Thomsen, H. K. Jensen, B. Kjærgaard, P. R. Zepernick, E. Worm, W. Fledelius, B. Cho, A. Sawant, D. Ruan, and P. J. Keall (2012). Megavoltage image-based dynamic multileaf collimator tracking of a NiTi stent in porcine lungs on a linear accelerator. *Int. J. Radiat. Oncol. Biol. Phys.*, 82(2):321–327.
- R. Prabhakar, P. K. Julka, T. Ganesh, A. Munshi, R. C. Joshi, and G. K. Rath (2007). Feasibility of using MRI alone for 3D radiation treatment planning in brain tumors. *Jpn J Clin Oncol*, 37(6):405–411.
- K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger (1999). SENSE: sensitivity encoding for fast MRI. *Magn. Reson. Med.*, 42(5):952–962.

- B. W. Raaymakers, A. J. E. Raaijmakers, and J. J. W. Lagendijk (2008). Feasibility of MRI guided proton therapy: magnetic field dose effects. *Phys. Med. Biol.*, 53(20):5615.
- B. W. Raaymakers, J. J. W. Lagendijk, J. Overweg, J. G. M. Kok, A. J. E. Raaijmakers, E. M. Kerkhof, R. W. van der Put, I. Meijnsing, S. P. M. Crijs, F. Benedosso, M. M Van Vulpen, C. H. W. De Graaff, J. Allen, and K. J. Brown (2009). Integrating a 1.5 T MRI scanner with a 6 MV accelerator: proof of concept. *Phys. Med. Biol.*, 54(229):229–237.
- P. Ramos-Cabrer, J. P. M. van Duynhoven, A. Van der Toorn, and K. Nicolay (2004). MRI of hip prostheses using single-point methods: in vitro studies towards the artifact-free imaging of individuals with metal implants. *Magn. Res. Imag.*, 22(8):1097–1103.
- D. Rangaraj, G. Palaniswaamy, and L. Papiez (2008). DMLC IMRT delivery to targets moving in 2D in beams eye view. *Med. Phys.*, 35:3765.
- T. Ravkilde, P. J. Keall, K. Højbjerg, W. Fledelius, E. Worm, and P. R. Poulsen (2011). Geometric accuracy of dynamic MLC tracking with an implantable wired electromagnetic transponder. *Acta Oncologica*, 50(6):944–951.
- J. R. Reichenbach, R. Venkatesan, D. A. Yablonskiy, M. R. Thompson, S. Lai, and E. M. Haacke (1997). Theory and application of static field inhomogeneity effects in gradient-echo imaging. *J. Magn. Res. Imag.*, 7(2):266–279.
- M. Ries, B. D. De Senneville, S. Roujol, Y. Berber, B. Quesson, and C. Moonen (2010). Real-time 3D target tracking in MRI guided focused ultrasound ablations in moving tissues. *Magn. Reson. Med.*, 64(6):1704–1712.
- J. Ross, R. Tranquebar, and D. Shanbhag (2008). Real-time liver motion compensation for MRgFUS. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2008*, pages 806–813.
- D. Ruan (2010). Kernel density estimation-based real-time prediction for respiratory motion. *Phys. Med. Biol.*, 55:1311–1326.
- A. Sawant, R. L. Smith, R. B. Venkat, L. Santanam, B. Cho, P. Poulsen, H. Cattell, L. J. Newell, P. Parikh, and P. J. Keall (2009). Toward submillimeter accuracy in the management of intrafraction motion: the integration of real-time internal position monitoring and multileaf collimator target tracking. *Int. J. Radiat. Oncol. Biol. Phys.*, 74(2):575–582.
- D. C. Schmidt, D. L. Levine, and S. Mungee (1998). The design of the TAO real-time object request broker. *Computer Communications*, 21(4):294–324.
- B. M. Schwartz and N. J. McDannold (2012). Ultrasound echoes as biometric navigators. *Magn. Reson. Med.*. Epub ahead of print.
- A. Schweikard, G. Glosser, M. Bodduluri, M. J. Murphy, and J. R. Adler (2000). Robotic motion compensation for respiratory movement during radiosurgery. *Computer Aided Surgery*, 5(4):263–277.
- J. Secon, R. T. Seethamraju, G. P. Broussard, H. Kooy, and M. Harisinghani (2009). The application of MRI pulse sequences in the verification of proton beam radiotherapy. In *Proceedings of the 17th Annual Meeting of ISMRM, Honolulu, Hawaii, USA*.
- Y. Seo, H. Takamiya, H. Ishikawa, T. Nakashima, and G. Navon (1998). *Spatially resolved magnetic resonance*, chapter 42. Wiley.
- Y. Seppenwoolde, R. I. Berbeco, S. Nishioka, H. Shirato, and B. Heijmen (2007). Accuracy of tumor motion compensation algorithm from a robotic respiratory tracking system: a simulation study. *Med. Phys.*, 34:2774–2784.
- G. C. Sharp, S. B. Jiang, S. Shimizu, and H. Shirato (2004). Prediction of respiratory tumour motion for real-time image-guided radiotherapy. *Phys. Med. Biol.*, 49:425–440.
- J. Smink, M. Häkkinen, R. Holthuizen, S. Krueger, M. Ries, Y. Berber, C. Moonen, M. Köhler, and E. Vahala (2011). eXTernal Control (XTC): a flexible, real-time, low-latency, bi-directional scanner interface. In *Proceedings of the 19th Annual Meeting of ISMRM, Montréal, Québec, Canada*, page 1755.
- D. K. Sodickson and W. J. Manning (2005). Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn. Reson. Med.*, 38(4):591–603.
- B. Stroustrup (1991). *The C++ Programming Language*. 2nd edition.

- Y. Suh, B. Yi, S. Ahn, J. Kim, S. Lee, S. Shin, and E. Choi (2004). Aperture maneuver with compelled breath (AMC) for moving tumors: a feasibility study with a moving phantom. *Med. Phys.*, 31:760–766.
- Y. Tsunashima, T. Sakae, Y. Shioyama, K. Kagei, T. Terunuma, A. Nohtomi, and Y. Akine (2004). Correlation between the respiratory waveform measured using a respiratory sensor and 3D tumor motion in gated radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 60(3):951–958.
- M. Uecker, S. Zhang, D. Voit, A. Karaus, K. D. Merboldt, and J. Frahm (2010). Real-time mri at a resolution of 20 ms. *NMR in Biomedicine*, 23(8):986–994.
- M. Uematsu, T. Fukui, A. Shioda, H. Tokumitsu, K. Takai, T. Kojima, Y. Asai, and S. Kusano (1996). A dual computed tomography linear accelerator unit for stereotactic radiation therapy: a new approach without cranially fixated stereotactic frames. *Int. J. Radiat. Oncol. Biol. Phys.*, 35(3):587–592.
- J. Unkelbach and U. Oelfke (2004). Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions. *Phys. Med. Biol.*, 49:4005–4029.
- J. Unkelbach, T. C. Y. Chan, and T. Bortfeld (2007). Accounting for range uncertainties in the optimization of intensity modulated proton therapy. *Phys. Med. Biol.*, 52(10):2755.
- L. Van de Bunt, I. M. Jürgenliemk-Schulz, G. A. P. De Kort, J. M. Roesink, R. J. H. A. Tersteeg, and U. A. Van der Heide (2008). Motion and deformation of the target volumes during IMRT for cervical cancer: What margins do we need? *Radiother. Oncol.*, 88(2):233–240.
- U. A. Van der Heide, A. N. T. J. Kotte, H. Dehnad, P. Hofman, J. J. W. Lagendijk, and M. Van Vulpen (2007). Analysis of fiducial marker-based position verification in the external beam radiotherapy of patients with prostate cancer. *Radiother. Oncol.*, 82(1):38–45.
- J. G. van Gorp, J S Bouwman, C. J. G. Bakker, and P. R. Seevinck (2012a). Towards compressed sensing accelerated geometrically undistorted Single Point Images under 10 seconds. In *Proceedings of the 20th Annual Meeting of ISMRM, Melbourne, Australia*.
- M. van Herk, P. Remeijer, C. Rasch, and J. V. Lebesque (2000). The probability of correct target dosage: dose-population histograms for deriving treatment margins in radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 47(4):1121–1135.
- J. R. Van Sörnsen de Koste, S. Senan, C. E. Kleynen, B. J. Slotman, and F. J. Lagerwaard (2006). Renal mobility during uncoached quiet respiration: An analysis of 4DCT scans. *Int. J. Radiat. Oncol. Biol. Phys.*, 64:799–803.
- M. Velec, J. L. Moseley, C. L. Eccles, T. Craig, M. B. Sharpe, L. A. Dawson, and K. K. Brock (2011). Effect of breathing motion on radiotherapy dose accumulation in the abdomen using deformable registration. *Int. J. Radiat. Oncol. Biol. Phys.*, 80(1):265–272.
- D. Verellen, M. D. Ridder, N. Linthout, K. Tournel, G. Soete, and G. Storme (2007). Innovations in image-guided radiotherapy. *Nat Rev Cancer*, 7(12):949–960.
- F. B. Viégas, M. Wattenberg, and J. Feinberg (2009). Participatory visualization with wordle. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1137–1144.
- K. Wachowicz, T. Stanescu, S. D. Thomas, and B. G. Fallone (2010). Implications of tissue magnetic susceptibility-related distortion on the rotating magnet in an MR-linac design. *Med. Phys.*, 37(4):1714–1721.
- R. Wagman, E. Yorke, E. Ford, P. Giraud, G. Mageras, B. Minsky, and K. Rosenzweig (2003). Respiratory gating for liver tumors: use in dose escalation. *Int. J. Radiat. Oncol. Biol. Phys.*, 55(3):659–668.
- R. Wolf and T. Bortfeld (2012). An analytical solution to proton Bragg peak deflection in a magnetic field. *Phys. Med. Biol.*, 57(17):N329.
- J. W. H. Wolthaus, J. J. Sonke, M. van Herk, J. S. A. Belderbos, M. M. G. Rossi, J. V. Lebesque, and E. M. F. Damen (2008). Comparison of different strategies to use four-dimensional computed tomography in treatment planning for lung cancer patients. *Int. J. Radiat. Oncol. Biol. Phys.*, 70(4):1229–1238.
- J. Yang, Z. Zhang, S. Zhou, and H. Yin (2009). Respiratory motion prediction based on maximum posterior probability. In *Proc. ICBBE*, pages 1–4.

J. Yuan, C. S. Mei, L. P. Panych, N. J. McDannold, and B. Madore (2012). Towards fast and accurate temperature mapping with proton resonance frequency-based MR thermometry. *Quant. Imag. Med. Surg.*, 2(1):21–32.

J. Yun, E. Yip, K. Wachowicz, S. Rathee, M. Mackenzie, D. Robinson, and B. G. Fallone (2012). Evaluation of a lung tumor autocontouring algorithm for intrafractional tumor tracking using low-field MRI: A phantom study. *Med. Phys.*, 39(3):1481–1494.

Publications

Papers

- T. Bjerre, S. P. M. Crijs, F. Munck af Rosenschöld, M. Azanar, L. Specht, R. Larsen, and P. Keall (2012). Three-dimensional MRI-Linac Intra-Fraction Guidance Using Multiple Orthogonal Cine-MRI planes. *Medical Physics*. submitted.
- S. P. M. Crijs, B. W. Raaymakers, and J. J. W. Lagendijk (2011a). Real-time correction of magnetic field inhomogeneity-induced image distortions for MRI-guided conventional and proton radiotherapy. *Physics in Medicine and Biology*, 56(1):289–297.
- S. P. M. Crijs, J. G. M. Kok, J. J. W. Lagendijk, and B. W. Raaymakers (2011b). Towards MRI-guided linear accelerator control: gating on an MRI accelerator. *Physics in Medicine and Biology*, 56(15):4815–4825.
- S. P. M. Crijs, B. W. Raaymakers, and J. J. W. Lagendijk (2012a). Proof of concept of MRI-guided tracked radiation delivery: tracking one-dimensional motion. *Physics in Medicine and Biology*, 57(23):7863–7872.
- S. P. M. Crijs, C. J. G. Bakker, P. R. Seevinck, H. de Leeuw, J. J. W. Lagendijk, and B. W. Raaymakers (2012b). Towards inherently distortion-free MR images for image-guided radiotherapy on an MRI accelerator. *Physics in Medicine and Biology*, 57:1349.
- S. J. Hoogcarspel, S. P. M. Crijs, J. J. W. Lagendijk, M. van Vulpen, and B. W. Raaymakers (2013). The feasibility of using a conventional flexible RF-coil for an online MR-guided radiotherapy treatment. *Physics in Medicine and Biology*.
- B. W. Raaymakers, J. J. W. Lagendijk, J. Overweg, J. G. M. Kok, A. J. E. Raaijmakers, E. M. Kerkhof, R. W. Put, I. Meijnsing, S. P. M. Crijs, F. Benedosso, M. van Vulpen, C. H. de Graaff, J. Allen, and K. J. Brown (2009). Integrating a 1.5 T MRI scanner with a 6 MV accelerator: proof of concept. *Physics in Medicine and Biology*, 54(12):N229.
- B. W. Raaymakers, J. C. J. De Boer, C. Knox, S. P. M. Crijs, K. Smit, M. K. Stam, M. R. Van den Bosch, J. G. M. Kok, and J. J. W. Lagendijk (2011). Integrated MegaVoltage portal imaging with a 1.5 T MRI linac. *Physics in Medicine and Biology*, 56(19):N229.
- M. K. Stam, S. P. M. Crijs, B. A. Zonnenberg, M. M. Barendrecht, M. Van Vulpen, J. J. W. Lagendijk, and B. W. Raaymakers (2012). Navigators for motion detection during real-time MRI-guided radiotherapy. *Physics in Medicine and Biology*, 57(21):6797–6805.
- M. K. Stam, M. Van Vulpen, M. M. Barendrecht, B. A. Zonnenberg, M. Intven, S. P. M. Crijs, J. J. W. Lagendijk, and B. W. Raaymakers (2013). Kidney motion during free breathing and breath hold for MR-guided radiotherapy. *Physics in Medicine and Biology*. submitted.
- H. M. M. Ten Eikelder, S. P. M. Crijs, M. N. Steijaert, A. M. L. Liekens, and P. A. J. Hilbers (2009). Computing with Feedforward Networks of Artificial Biochemical Neurons. *Natural Computing*, pages 38–47.

Proceedings and abstracts

- S. P. M. Crijns, B. W. Raaymakers, J. Overweg, and J. J. W. Lagendijk (2010a). Characterizing gradient non-linearities of a split gradient coil in a hybrid MRI-linear accelerator. In *Proceedings of the 18th annual meeting of the ISMRM*.
- S. P. M. Crijns, B. W. Raaymakers, and J. J. W. Lagendijk (2010b). Real-time distortion-free MR images for MR-guided radiotherapy. In *Proceedings of the XVth ICCR*.
- S. P. M. Crijns, B. W. Raaymakers, and J. J. W. Lagendijk (2011a). Distortion free MR images for IGRT on the MRI-accelerator. In *Proceedings of the 11th biannual ESTRO*.
- S. P. M. Crijns, B. W. Raaymakers, and J. J. W. Lagendijk (2011b). Radiation induced RF-coil degradation in hybrid MRI-accelerator systems. In *Proceedings of the 19th annual meeting of the ISMRM*.
- S. P. M. Crijns, J. G. M. Kok, J. J. W. Lagendijk, and B. W. Raaymakers (2011c). Towards MRI-guided linear accelerator control: gating on an MRI accelerator. In *Proceedings of the 28th annual scientific meeting of the ESMRMB*.
- S. P. M. Crijns, H. Bluemink, J. Zwanenburg, F. Visser, H. Dehnad, and M. Philippens (2012a). MR neurography using a 3D T2-FFE sequence with binomial RF-pulse selective excitation. In *Proceedings of the 20th annual meeting of the ISMRM*.
- S. P. M. Crijns, B. W. Raaymakers, and L. J. J. W (2012b). Towards MRI-guided tracked radiation delivery on a 1.5T MRI-linear accelerator. In *Proceedings of the 9th Interventional MRI Symposium*.
- S. P. M. Crijns, B. W. Raaymakers, and J. J. W. Lagendijk (2012c). Towards MRI-guided tracked radiation delivery on an MRI-linac. In *Proceedings of the 29th annual scientific meeting of the ESMRMB*.
- S. P. M. Crijns, A. N. T. J. Kotte, J. J. W. Lagendijk, and B. W. Raaymakers (2012d). Validation of algorithm for real-time target tracking in MRI-guided interventions. In *Proceedings of the 20th annual meeting of the ISMRM*.
- F. M. Lever, I. M. Lips, S. P. M. Crijns, O. Reerink, J. P. Ruurda, A. L. W. H. M. Van Lier, M. A. Moerland, M. Van Vulpen, and M. G (2012). Real-time esophageal tumour motion quantification on cine MRI using an adaptive correlation filter. In *Proceedings of the ESTRO 2012 annual meeting*.
- M. E. P. Philippens, J. J. Bluemink, S. P. M. Crijns, J. J. M. Zwanenburg, F. Visser, F. A. Pameijer, and C. H. J. Terhaard (2012). Tracing the Cranial Nerve Pathways NV and NVII with 3D T2-FFE. In *Proceedings of the 20th annual meeting of the ISMRM*.
- B. W. Raaymakers, J. J. W. Lagendijk, M. K. Stam, K. Smit, and M. Van Vulpen (2011). Radiotherapy systems with integrated MRI functionality for direct, intra-fraction tumour tracking. In *Proceedings of the ESTRO 2011 annual meeting*.
- B. W. Raaymakers, K. Smit, G. H. Bol, M. R. Van den Bosch, S. P. M. Crijns, J. G. M. Kok, and J. J. W. Lagendijk (2012a). Dosimetry in strong magnetic fields; issues and opportunities. In *Proceedings of the ESTRO 2012 annual meeting*.
- B. W. Raaymakers, S. P. M. Crijns, J. G. M. Kok, K. Smit, M. K. Stam, M. Van Vulpen, O. J. C. Knox, B. K. J, and L. J. J. W (2012b). The MRI accelerator status update: real-time MRI guided Radiotherapy. In *Proceedings of the 9th Interventional MRI Symposium*.
- M. K. Stam, S. P. M. Crijns, J. J. W. Lagendijk, and B. W. Raaymakers (2010a). Real-time contour propagation for on-line MRI guided radiotherapy. In *Proceedings of the XVth ICCR*. Submitted.
- M. K. Stam, S. P. M. Crijns, J. J. W. Lagendijk, and B. W. Raaymakers (2010b). Three dimensional real-time organ motion detection for MRI-guided radiotherapy. In *ISMRM workshop on motion correction*.
- M. K. Stam, S. P. M. Crijns, J. J. W. Lagendijk, and B. W. Raaymakers (2011). Navigators for motion detection in the 1.5T MRI-linac. In *Proceedings of the ASTRO 2011 annual meeting*.
- M. K. Stam, M. Van Vulpen, B. A. Zonnenberg, M. Intven, S. P. M. Crijns, J. J. W. Lagendijk, and B. W. Raaymakers (2012). Gated MRI-linac (MRL) treatment for kidneys: inter breath hold variations. In *Proceedings of the ESTRO 2012 annual meeting*.

B. Van Asselen, S. P. M. Crijns, M. B. B. W. Raaymakers, and O. Reerink (2010). Minimizing liver motion using a mini-corset in SBRT liver: quantification of the reduction and repeatability using cine-MRI. In *Proceedings of 29th ESTRO meeting*.

J. S. Van Gorp, S. P. M. Crijns, B. J. G. B. W. Raaymakers, B. C. J. G., and P. R. Seevinck (2012). Compressed Sensing-Accelerated Spectroscopic Imaging to Obtain Geometrically Accurate Water & Fat Images and Field Maps for Use in MR-Guided Intervention. In *Proceedings of the 20th annual meeting of the ISMRM*.

Invited presentations

Towards on-lineMR-guidance for radiotherapy – gating on the MRI-linac *at*: Annual NVKF meeting, Woudschoten, The Netherlands, 2011

Towards MRI-guided intra-fraction motion management *at*: University of Sydney Physics Seminar, Sydney, Australia, 2012

Towards MRI-guided intra-fraction motion management *at*: NKI-AvL Physics Seminar, Amsterdam, The Netherlands, 2012

Quality control in MRI-guided therapy; application to radiotherapy *at*: 29th Annual Scientific Meeting of the ESMRMB, Lisbon, Portugal, 2012

Sequences for MRI-guided external beam radiotherapy *at*: 32th Annual Meeting of the ESTRO, Geneva, Switzerland, 2013

Dankwoord

Een promotieproject kun je niet in je eentje uitvoeren. Ook in mijn geval waren er veel mensen bij betrokken. Waarschijnlijk vergeet ik dan ook een aantal mensen hieronder te noemen, maar ik ga toch een poging wagen. En mocht je je naam hieronder niet kunnen ontwaren maar wel in het rijtje thuishoren: dankjewel!

Jan, als promotor heb je je weinig bemoeid met de dagelijkse gang van zaken tijdens mijn onderzoek, maar je bracht wel een structuur in mijn werk aan en bleef telkens kritisch kijken naar de resultaten. Ik bewonder je visie en overredingskracht; zo'n groot idee als het is, het CBOI is naar mijn gevoel in een mum van een tijd al bijna werkelijkheid geworden. We blijven zoeken naar een manier om de stralingsbundel op MRI beelden te visualiseren.

Bas, hoewel ik solliciteerde op een OiO-plek waarin Monte-Carlo berekeningen centraal stonden, aarzelde jij niet om me aan te nemen als MRI fysicus op het MRI-linac project. Van MRI had ik destijds helemaal geen kaas gegeten. Je optimistische blik en 'omdat het kan' houding heb ik ontzettend kunnen waarderen tijdens het project. Met onze verschillende manieren van werken vullen we elkaar denk ik goed aan; als ik iets van je heb mogen leren, is het wel dat ideeën gewoon uitgeprobeerd moeten worden, in plaats van er over te blijven nadenken. Hoewel ik in het begin mijn ei niet kwijt kon heb je me geholpen mijn plek in het project te vinden. Bedankt!

Jan leerde ik kennen als die man die in de bunker koperen kisten stond te solderen. De laatste kist, ook al was het de allerkleinste die uiteindelijk gebruikt zou worden, maakten we met z'n tweeën. Hierna hebben we nog vaak de kans gekregen om samen te knutselen. Jan, bij het aan de praat krijgen van de versneller, scanner problemen oplossen, talloze kleine fröbelprojectjes en debuggen van het nieuwe prototype vormden we een goed team. Het zal nog wel even duren voor ik een versneller zo begrijp als jij, als dat nog eens mag lukken. Top dat je me ook tijdens de verdediging van mijn proefschrift wil bijstaan.

Het MRI-linac project maakte telkens grote sprongen wanneer Johan Overweg zich ermee bemoeide. Johan, ook jij bent iemand die in het rijtje mensen past die zich bij een out-of-the-box idee gaan afvragen hoe ze het zouden kunnen realiseren, in plaats ervoor terug te schrikken. Ik heb veel opgestoken van je hulp bij mijn strubbelingen om wat meer van magneten te begrijpen; hopelijk blijft het ook hangen. Ik vond het bijzonder leuk om met je samen te werken, dat komt niet in de laatste plaats door je ontzettend droge humor.

Chris, you and your team have made great efforts to turn the idea of a gantry-based

linac around an MRI scanner into reality. You helped me a lot by structuring my ideas on the motion management architectures into readable physical and logical diagrams. Thanks for coping with my whining about delays, interface specs and different pieces of hardware that would not work instantly; it must have driven you mad at times! Fortunately it got us somewhere and I am glad you take up the challenge to turn this crazy idea into clinical reality now. It was and still is a pleasure to work with you.

Kevin, you are among the first believers in the idea of MRI-guided radiotherapy. Thanks for making the project possible and for your ever critical remarks on my results.

Een groot aantal mensen heeft geholpen bij het daadwerkelijk bouwen van de MRI-linac prototypes. Dankzij het installatiewerk van Jan, Arie en Cor werkte het eerste MRI-linac prototype als een zonnetje. Jan, jij bent daarnaast ook het aanspreekpunt geweest om hardware te regelen voor 'het hobbyprojectje'. Gelukkig kunnen we met onze vragen nog steeds bij jullie aankloppen. Bij Jan en Ed van Jurriaans kun je terecht voor schitterend timmerwerk. Ruud regelde regelmatig nieuwe scannerlicenties, aangezien ons systeem in de boeken van Philips onvindbaar was. Niels legde de basis voor de gating modificatie aan de versneller, die ik nodig had voor het werk uit Hoofdstuk 5. De ingenieurs van CCM en NTS ontwierpen de ring en verschillende subsystemen van het prototype. Met Rob heb ik mogen samenwerken om het CAN protocol voor communicatie met de gantry drive vorm te geven. Onder leiding van Eric en Jos werd de gantry veilig geïnstalleerd. Met Johan, Jouke en Ronald hadden we de eerste discussies over communicatie met de MR scanner en zij hadden gelukkig oplossingen klaar liggen.

By sharing the code of their realTI project Yasmina and Mario have helped me to realise the results presented in chapters 5 and 6. Their interface takes care of communication with the MR scanner and I have learned a great deal from their example of the use of the underlying technology to write the control software for the linac presented in Chapter 4. Thanks guys!

Johan Sjöholm built us a great MR-compatible watertank and provided us with ultrasonic motors. It was good fun to have you as a guest Johan.

I would like to thank the members of the STW users committee for their useful input.

I have had the pleasure to discuss MRI-linac concepts with a large amount of people. Paul, Jill, Melanie, Erkki, Robin, Colin, Christopher, Marco, Gert, Romhild (and there must be more of you), thanks for the many discussions and for listening to my ideas.

De leden van de beoordelingscommissie wil ik graag bedanken voor het kritisch lezen van dit proefschrift.

Alexis en Gijs, naast de talloze tools van jullie hand die ik bij het onderzoek heb kunnen gebruiken, had ik aan jullie beide ook een fijne sparring partner als het ging om het opzetten van het interfacing framework. Heel wat features die bij mij als een hersenkronkel begonnen zijn door er met jullie verder over te praten uitgedroefd tot volwaardige ideeën die vervolgens in het framework zijn geïmplementeerd. Jullie zijn samen een onmisbare schat aan informatica-ervaring binnen de afdeling. Ik ben er dan ook trots op nu samen met jullie een team te mogen vormen. Alexis, daarnaast mag ik met jou het sjitterendste kantoor van

de afdeling delen. Ik denk met veel plezier terug aan het frees- en 3D-printuitje. Fijn dat onze samenwerking nu nog net iets verder gaat.

Op de afdeling is veel MRI kennis aanwezig, maar Nico en Mariëlle zijn de echte goeroes als het gaat om deze techniek. Samen hebben we Bernstein's boek doorgewerkt en er werd fel gediscussieerd over een scala aan sequenties. Met het imaging team vertegenwoordigen jullie een belangrijke brok kennis op een MRI-minded afdeling. Ook Wilbert, Peter, Hendrik en Chris wil ik in deze context bedanken.

Natuurlijk heb je ook op de afdeling voor de dagelijkse gang van zaken en speciale wensen mensen nodig. Ric en consorten dragen zorg voor de broodnodige IT infrastructuur. Met mijn administratieve frustraties kan ik terecht bij Leonie, Therèse of Judith op het secretariaat. Bert opent de koffiepauze en heeft de nodige schroeven, moeren en een schat aan gereedschap. Bij de versnellertechnici en fysisch medewerkers liep ik geregeld binnen met technische vragen.

De OiOs van de radiotherapie afdeling en de 7T groep wil ik graag bedanken voor de lunches en koffiepauzes die telkens een fijne afleiding waren. Ook de gezelligheid tijdens conferenties zal mij nog lang bij blijven. Markus also belongs with the group of people with whom I shared an office. Thanks for your good spirits and the nice conversations we had, as well as the great night out for beer. Astrid wil ik graag bedanken voor stelling 9 en voor het delen van haar ervaringen wat betreft het maken van een proefschrift.

UB, wat een gezellig zootje maken we ervan. Of het nu in Eindhoven, de Ardennen of op skivakantie is, gezellig wordt het zeker. Met jullie discussieren over niets is minstens zo belangrijk als de rest van de inhoud van dit boekje. Rik, top dat je me bij wil staan als paranimf. Ook voor jou komt het verdedigingsmoment gelukkig dichterbij. Hanneke, Miriam, Sten, Emmy, Jarno, onze maandelijkse afspraak staat al lang niet meer als een huis, maar altijd kijk ik er naar uit om jullie weer te zien. Waar spelen we het volgende spel? En Arjan, Elise en Maarten, ook al zijn jullie meestal ver weg, als we weer even bij elkaar zijn is het helemaal goed. Jac, Riny, Paulien en Bram, het is super om me ook in Born thuis te voelen. Pap en mam, in Oirsbeek is de basis gelegd. Dankzij jullie was er altijd wat nieuws om mijn nieuwsgierigheid op te richten en altijd wat nieuws om te leren; het is nog lang niet af.

Het ereplekje van dit dankwoord is gereserveerd voor Anne. Ik ben enorm gelukkig met jou op onze nieuwe Utrechtse stek. Dat je daarvoor dagelijks naar Nijmegen reist is een grote investering, dat is echt top! Fijn dat ik altijd iemand heb om op te bouwen, op terug te vallen en die de frustraties die bij het promotie onderzoek hoorden kan relativeren. Gaan we nu snel vakantie houden?

Sjoerd

Curriculum Vitae

Sjoerd Crijns was born on December 8th 1984 in Sittard, The Netherlands. After finishing his secondary school education from the St.-Janscollege in Hoensbroek in 2003, he started with the Bachelor program Biomedical Engineering at the Eindhoven University of Technology. In his subsequent Master in Biomedical Engineering he focussed on modeling and simulations. Sjoerd's master thesis centered on the topic of meso-scale molecular modeling of amphiphilic polymers in aqueous solutions.

After graduating (*cum laude*) in 2008, Sjoerd became more interested in the application of technology in healthcare and particularly in the field of medical imaging. He accepted a PhD position in the MRI-linac group of the department of Radiotherapy at the University Medical Center Utrecht. The research focuss was on geometrical accuracy of MR images, rapid imaging and MRI/linac interfacing. The results of these projects are described in this thesis.

Sjoerd continues his work at the Utrecht radiotherapy department as a computer scientist, combining clinically oriented informatics tasks with research regarding on-line MRI-guidance and intra-fraction motion management for radiotherapy.