ב"ה

Numerical design and optimization of a millimetre-scale absorbed dose graphite calorimeter for use in small-field megavoltage photon beams

Federico M. R. Keszti

Medical Physics Unit
McGill University, Montréal
November 2020

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Master of Science in Medical Radiation Physics

Table of Contents

Abstract

The aim of this work is to numerically characterize and reoptimize the design of a previously developed probe-format absorbed dose graphite calorimeter, termed Aerrow MK7, to be able to perform accurate absolute dosimetry in small high-energy photon beams. Aerrow MK7 has a cylindrical sensitive volume with a radius of 3 mm and a length of 5 mm, making it a practical dosimeter for routine clinical use. Aerrow MK7 can operate in two modes to determine absorbed dose: quasi-adiabatic and isothermal. The latter mode of operation displayed an unexpected overresponse error of 12–31 % in all measurements. A systematic investigation of the steps involved in the absorbed dose determination chain for Aerrow MK7 was conducted, with the aim of (a) reviewing and verifying the methodology used to determine absorbed dose, (b) identifying a problematic step that could potentially account for the overresponse seen in isothermal mode, and (c) proposing feasible design revisions to rectify the response. After establishing that the overresponse was most likely related to a heat transfer effect, a numerical model was constructed with finite element method-based software to simulate the heat transfer processes of the different detector components. The model was subjected to numerous tests to understand the effect of various design factors on the absorbed dose. The results from these tests established that a construction error relating to the assembly of electrical components in the jacket (a graphite body that encloses the sensitive volume) was the most likely source of the malfunction observed in isothermal mode. The effect of the material type and density of the thermal insulation layers on the absorbed dose were also studied. These simulations revealed that, for a fully functional and properly assembled Aerrow MK7, an underresponse is expected instead. A list of possible workarounds, such as modifications to the probe's geometry and a heat transfer correction factor, are proposed to mitigate the underresponse.

Abrégé

L'objectif de cette recherche est de caractériser numériquement et de réoptimiser la sonde calorimètre au graphite précédemment développé (appelé Aerrow MK7) afin de pouvoir effectuer une dosimétrie absolue précise et juste de petits champs en radiothérapie externe. Aerrow MK7 a un volume sensible cylindrique avec un rayon de 3 mm et une longueur de 5 mm, ce qui en fait un dosimètre pratique pour une utilisation clinique. Aerrow MK7 peut fonctionner selon deux modes pour déterminer la dose absorbée: adiabatique et isothermique. Ce dernier mode de fonctionnement a affiché une erreur de surréponse inattendue de 12 à 31% dans toutes les mesures. Une investigation systématique des étapes impliquées dans la chaîne de détermination de la dose absorbée pour Aerrow MK7 a été menée, dans le but de (a) revoir et vérifier la méthodologie utilisée pour déterminer la dose absorbée, (b) identifier une étape problématique qui pourrait potentiellement expliquer la surréponse vue en mode isothermique, et (c) proposer des révisions de conception réalisables pour rectifier la réponse d'Aerrow MK7. Après avoir établi que la surréponse était très probablement liée à un effet de transfert de chaleur, un modèle numérique a été construit avec un logiciel basé sur la méthode des éléments finis pour simuler le transfert d'énergie thermique dans les différents composants du détecteur. Le modèle a été soumis à de nombreux tests pour comprendre l'effet de divers facteurs de conception sur la dose absorbée. Les résultats de ces tests ont établi qu'une erreur de construction liée à l'assemblage des composants électriques dans la jaquette (un corps en graphite qui renferme le volume sensible) était la source la plus probable du dysfonctionnement observé en mode isothermique. L'effet du type de matériau et de la densité des couches d'isolation thermique sur la dose absorbée a également été étudié. Ces simulations ont révélé que, pour un Aerrow MK7 entièrement fonctionnel et correctement assemblé, une sous-réponse est attendue à la place. Une liste de solutions de contournement possibles, telles que des modifications de la géométrie de la sonde et un facteur de correction du transfert de chaleur, est proposée pour atténuer la sous-réponse.

Acknowledgements

This thesis and the work it entailed could not have been completed without the help of my family, friends, colleagues, and supervisors, whose contributions deserve to be acknowledged here.

First, I would like to extend my gratitude to Dr. James Renaud for his infinite patience, his guidance in all matters, and his excellent supervision. My most sincere thanks to Benjamin Côté for introducing me to his project and for benefitting my own with his hard work. I would also like to thank Dr. Jan Seuntjens, a source of seemingly endless knowledge and expertise from whom I continuously learn. It has been a joy to work with such a knowledgeable and dedicated team.

I would like to thank Dr. Arman Sarfehnia for his valuable input throughout this project and for his advice on presentations and writing. My thanks also to Dr. Carl Ross for taking the time to review this thesis and for his sharp and incredibly insightful comments. I would also like to thank my student colleagues, particularly Veng Heng and Julien Bancheri, for their support and for their patient help with some of the computational tasks.

Many thanks to the clinical and academic staff of the Medical Physics Unit. Thanks to Michael Evans for sharing his valuable time and expertise with me on some of the technical aspects of this project. Thanks to Margery Knewstubb and Tatjana Nisić for handling all the administrative items and for their positive, welcoming attitudes.

Finally, I would like to thank my entire family, including my loyal study companion ‘Geezer’, for their support and for helping me to overcome challenges and reach new milestones. Your encouragement has always been strongly felt. A special thank you to my mentor in life and my guiding best friend, Hashem.

Chapter 1. Introduction

1.1. Cancer Incidence and Treatment

Cancer remains the leading cause of death in Canada; 225,800 new cancer cases and 83,300 cancer deaths are expected in Canada for the year 2020. It is estimated that half of all Canadians will develop cancer in their lifetime, and about a quarter of all Canadians will die of cancer. In 2019, 25 Canadians were diagnosed with cancer every hour. Due to the growing and overall aging Canadian population, the number of cancer cases and deaths is expected to rise, and the burden of cancer is unlikely to decline in the coming years. Although the number of total cases and deaths is growing, incidence and mortality rates continue to decline for many cancers. The five-year net survival has increased from 55% in the early 1990's to an estimated 63% in 2020.[1,2]

Cancer is not a singular disease but rather an umbrella of diseases, characterized by uncontrolled and dysfunctional cellular growth. Hence, survival rates vary widely by cancer type. Thyroid and testicular cancer have relatively high survival rates, at 98% and 97%, respectively. Female breast cancer death rates have decreased by approximately 48% since 1986. The diminishing mortality rates for these and other cancers indicate continued success in cancer prevention and treatment. But the relatively low five-year net survival of other types, like esophageal (15%) and pancreatic (8%) cancers, point to the growing need for additional efforts in the advancement of research, screening and treatment modes, and the overall improvement of existing programs.[1,2]

The different types of treatments and treatment schemes existent for cancer constitute many extensive disciplines of research. Some of the most common treatment methods include surgery, immunotherapy, hormone therapy, chemotherapy and radiotherapy. Any one of these treatment methods may be prescribed in conjunction with one or many others. Furthermore, therapy methods can be designated and prescribed as any of the following:

i. Primary, in which the objective is the absolute removal and elimination of the cancer.
ii. Adjuvant, in which the objective is to eliminate any cancer cells that may remain after primary treatment, in order to prevent a recurrence (also called 'neoadjuvant' if done prior to primary treatment).
iii. Palliative, in which the objective is to eliminate symptoms, relieve pain, and/or prolong the well-being of the patient.[3]

The focus of this thesis belongs within the vast field of radiotherapy research.

1.2. Radiotherapy

Radiotherapy (RT), also known as radiation therapy, radiation oncology, and therapeutic radiology, employs ionizing radiation to treat malign cancerous cells. The basic principle of radiotherapy is to destroy or alter the chemical bonds of the deoxyribonucleic acid (DNA) inside the cell's nucleus, using ionizing radiation, in order to kill the cell or halt its reproduction.[4]

To this end, many kinds of radiation can be used, and these are classified into directly and indirectly ionizing radiation. Charged particles like electrons, protons and atomic ions fall in the category of directly ionizing radiation, since they can ionize atoms directly through Coulomb interactions. Indirectly ionizing radiation includes photons and neutrons, which can produce charged particles through various atomic interactions; in turn, these secondary charged particles can ionize atoms if they are sufficiently energetic. Each radiation type offers unique therapeutic characteristics and radically differ from each other in how they interact with biological tissue. To exemplify, an equal dosage (defined as the amount of radiation energy absorbed per unit mass of tissue — see Section 1.3.1.3) of proton therapy and photon therapy will not produce the same biological effects.[5]

Radiotherapy can be further subdivided according to the delivery method. The three main radiation delivery methods for therapeutic purpose are external beam radiation therapy, brachytherapy and radionuclide therapy. In brachytherapy, radioactive sources are placed directly into the patient's affected area in order to kill the cancer cells adjacent to it. This is accomplished via special implants designed to house the radiation source and are usually inserted surgically into the patient. Radionuclide therapy consists of administering a radioactive drug that targets a specific type of cell. This drug is called the radiopharmaceutical, and is primarily metabolized by the cancer cells, hence destroying them preferentially. These two methods involve invasive clinical procedures.[4-7]

1.2.1. External Beam Radiation Therapy

By contrast, external beam radiation therapy (EBRT) uses radiation beams from sources external to the patient and is therefore non-invasive. Partly due to this reason, high-energy EBRT is the most common form of radiation treatment.[5] In the clinic, high-energy denotes the energy range around 1–25 MeV for photon treatments (and 4 to 22 MeV for electrons). Lower photon energy ranges are less commonly used for radiotherapy; energies of 1 MeV and below are generally used to treat superficial lesions or tumours located at relatively low depths.[8]

1.2.1.1. Linacs

Photon beams with energies below 500 keV are generated with diagnostic X-ray tubes and orthovoltage units. High-energy photon and electron beams are generated with linear accelerators (linacs). In simplistic terms, a linear accelerator is a machine that generates high-energy electrons which can be aimed at a target. When treating with electron beams, this target is a dense metal filter that broadens the pencil beam exiting the linac. A second filter equalizes the electron energy distribution over the area of the beam. Photon beams are produced when the electrons are focused on a high-density target, usually a tungsten block. Upon impinging on the target's atoms, the electrons release X-rays primarily through bremsstrahlung interactions, and create a photon beam

with a given energy spectrum. Before exiting the linac, the photon beam can be collimated to a desired shape and size using stationary or movable collimators. Figure 1.1 illustrates the basic components of a modern linear accelerator.[4,5,8]

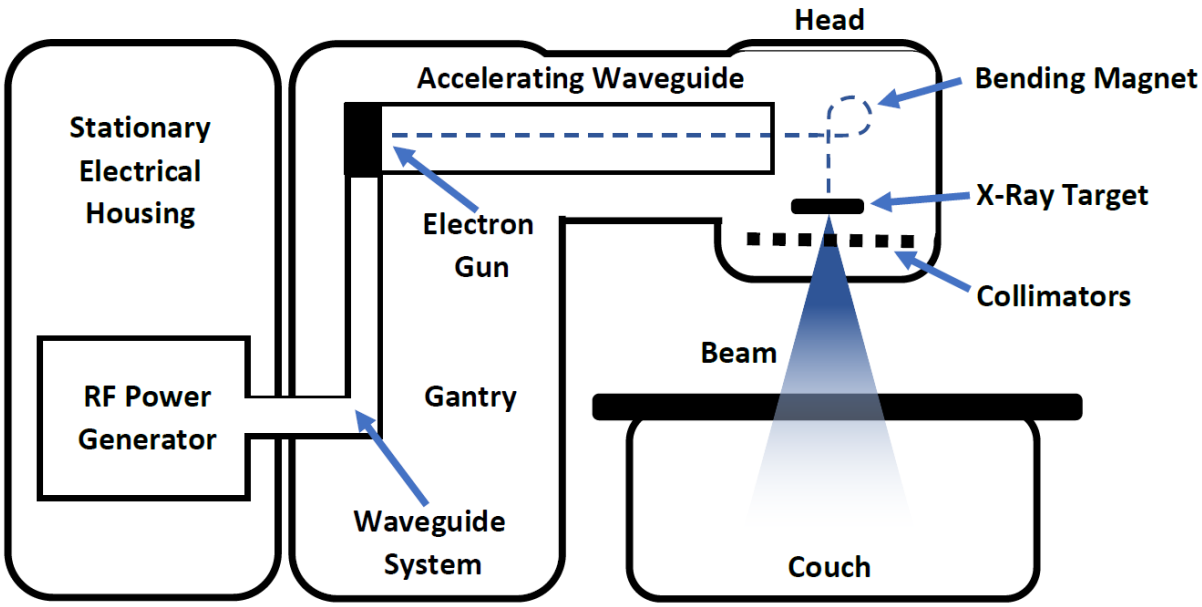

Figure 1.1. Schematic representation of a common linac design. The radiofrequency (RF) power generator converts direct current into megawatt-power microwaves that are transported to the accelerating waveguide, where electrons are accelerated to a specific kinetic energy before striking the X-ray target.[5,8,9]

1.2.1.2. High-Energy Photon Beam Radiotherapy

Photon beams can be used to treat superficial as well as deep-seated tumours. They offer certain advantages over particulate radiation types; photons can provide transmission through thicker objects than electrons, and photon beam linacs are presently more accessible than atomic ion RT machines in terms of cost and necessary infrastructure.[10]

Unlike electron beams, photon beams are virtually never monoenergetic. For this reason, photon beams are specified according to the megavoltage that is used to generate them and labeled as 'MV'. Photon beams can be flexibly modified; they may be filtered (in terms of energy) and

collimated (in terms of dimensions). To treat with a photon beam, the patient is laid down on a couch and the linac's gantry rotates around them, in order to rotate the incidence direction of the beam. To be exact, the linac head rotates around the isocenter, a machine coordinate along the beam axis that is specified by the operator and usually placed at the centre of the tumour. Multiple incidence angles ensure that the tumour gets good coverage and shape conformity from the beam. Irradiation can be done continuously as the source rotates or statically (the 'step-and-shoot' approach) with many different photon fields.[5]

1.3. Dosimetry

Because of the ionizing and ultimately destructive effect of radiation on cellular structures, it is imperative in radiotherapy that the radiation dose delivered be precisely and accurately known. Since radiation is used with the premise of preferentially targeting malignancies while simultaneously sparring the healthy tissues and organs around them, it is also essential to determine the dose distribution spatially. Dosimetry is the sub-field of radiotherapy physics that deals with these concerns; namely, quantifying and reducing uncertainties in the radiation dose delivered by a given RT modality. Dosimetry is key to machine calibration, quality assurance, patient risk assessment and personnel safety.

1.3.1. Dosimetric Quantities and Symbols

To understand dosimetric measurements, it is necessary to introduce a few basic dosimetric quantities and symbols.

1.3.1.1. Fluence

The particle fluence is defined as:

$$\Phi = \frac{dN}{dA} \tag{1.1}$$

Where dN is the number of particles passing through dA, and dA is the cross-sectional area of a sphere at the point of interest.[11] The energy fluence is the product of the particle fluence with the energy of each particle:

$$\Psi = \int \frac{d\Psi}{dE'}\, dE' = \int \frac{d\Phi}{dE'} \cdot E'\, dE' \tag{1.2}$$

Here, $\frac{d\Phi}{dE}$ and $\frac{d\Psi}{dE}$ represent the particle and energy fluence spectrums, respectively, and are both functions of the particle energy, E. Fluence is quantified in [m^{-2}] and energy fluence is quantified in [$J \cdot m^{-2}$].[4,5,11,12]

1.3.1.2. Kerma

Kerma stands for Kinetic Energy Released per unit MAss, and quantifies the average energy transferred ($\bar{E}_{\mathrm{tr}}$) from the incoming photons to the electrons in a medium. This includes energy transfers from all possible interactions (Compton effect, pair production, photoelectric effect, etc.) but it does not account for any energy transferred by the secondary electrons to other particles. For a medium with an infinitesimal mass element dm, then:

$$K = \frac{d\bar{E}_{\mathrm{tr}}}{dm} \tag{1.3}$$

The mass energy transfer coefficient, $\left(\frac{\mu_{\mathrm{tr}}}{\rho}\right)$, is the product of the beam's mass attenuation coefficient and the fraction of the energy transferred to electrons by the incident photons. It is related to the kerma through:

$$K = \int_0^{E_{\max}} \frac{d\Psi}{dE'}\left(\frac{\mu_{tr}}{\rho}(E')\right) dE' = \Psi\left(\frac{\bar{\mu}_{tr}}{\rho}\right) \tag{1.4}$$

Where $E_{\max}$ is the maximum photon energy in the beam. Kerma is usually split into a radiative and a collisional component. The collisional component, K_{col}, accounts for the energy transferred to secondary electrons that will produce ionizations through Coulomb interactions. The radiative component, K_{rad}, comprises the energy transferred to secondary electrons that will produce

photons, which are generated predominantly through bremsstrahlung interactions. The average fraction of the energy imparted onto the electrons that is expended through radiative processes is known as the radiation yield, $\bar{g}$. Hence, K_{rad} and K_{col} are expressed as:

$$K_{\mathrm{rad}} = K\bar{g} \tag{1.5}$$

$$K_{\mathrm{col}} = K(1 - \bar{g}) \tag{1.6}$$

Kerma is measured in [J · kg^{-1}] or [Gy].[4,5,11,12]

1.3.1.3. Absorbed Dose

Absorbed dose is defined as the energy deposited by ionizing radiation per unit mass of an absorbing medium. Mathematically, the absorbed dose is given by:

$$D = \frac{dE_{\mathrm{abs}}}{dm} \tag{1.7}$$

Here, dE_{abs} is the amount of energy absorbed by a medium with a mass unit element dm. After the photons impart energy to the electrons (kerma), the energy from these secondary electrons can be absorbed by the medium in the form of ionizations and excitations, resulting in absorbed dose. Alternatively, the energy of these secondary charged particles can exit the medium in the form of bremsstrahlung or in-flight annihilation photons. Accordingly, the absorbed dose is the kerma minus the energy lost due to radiative processes. Absorbed dose is measured in [Gy] or [cGy].[4,5,11]

1.3.1.4. Percentage Depth Dose

The absorbed dose is not simply an exponential decay, but rather a non-trivial function of depth in the irradiated medium. For megavoltage photon beams, the absorbed dose exhibits a peak at a certain depth, z_{max}, called the depth of maximum dose. The percentage depth dose (PDD) is defined as the percentage of the absorbed dose at a reference depth, z_{ref}, at a depth z along the central beam axis. The PDD is conventionally measured in water phantoms. The reference depth is usually taken as z_{max} to normalize the curve to 100%. Since the PDD is a function of critical

beam parameters like beam energy, field size at the surface of a flat phantom, and the source-to-surface (SSD) distance, medical physicists characterize beams by their PDD curves.[4,5,12]

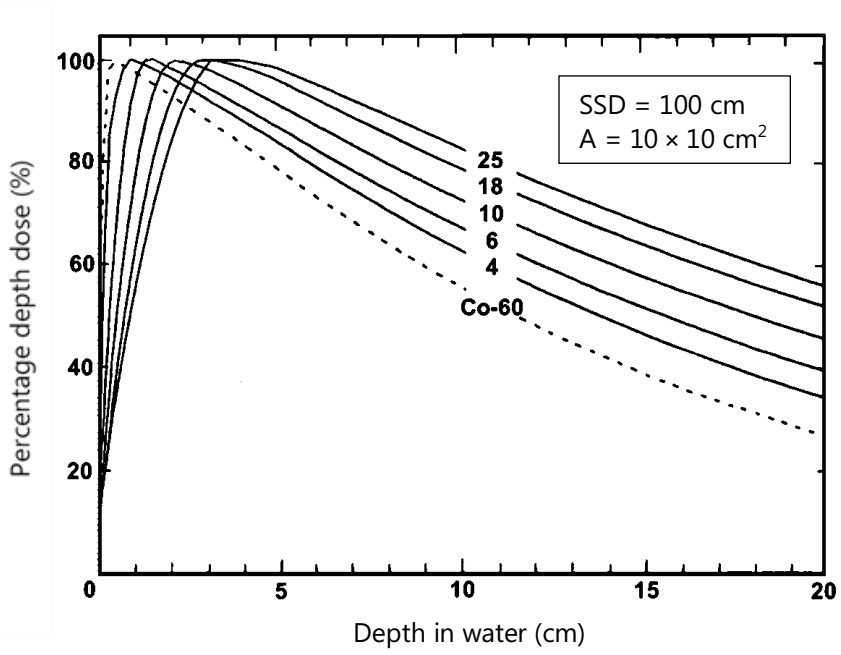

Figure 1.2. PDD curves normalized at $z_{\max}$ for photon beam energies ranging between 4–25 MV. The dashed line is the PDD curve from a ^{60}Co machine, which emits photons at an average energy of approximately 1.25 MeV. Image courtesy of Dr. Jan Seuntjens (McGill University, 2020).[5]

Thus, for a field size with equivalent square dimensions A, an SSD of f, and a beam energy E, the PDD value at an arbitrary depth z is:

$$PDD(z, A, f, E) = 100\left(\frac{D(z)}{D(z_{\max})}\right) = 100\left(\frac{\dot{D}(z)}{\dot{D}(z_{\max})}\right) \tag{1.8}$$

Where $\dot{D}$ denotes the time derivative of the dose, or the dose rate.

1.3.2. Dosimeters

Dosimeters, for the purposes of this thesis, can be conceived as devices through which absorbed radiation dose (or dose rate) is measured. Currently existent dosimeters are incapable of directly measuring absorbed dose. Instead, the operating principle of most dosimeters is to measure a

surrogate quantity and infer the dose from it. For example, calorimeters measure the temperature change in a medium before and after irradiation, and radiochromic films get more opaque as they are irradiated, so the measured quantity in film is the optical density.[13]

Besides the physical quantity which they measure, dosimeters can be characterized through the many variables that affect their response. Some of these include the detector's energy dependence, total dose dependence, dose-rate dependence, medium dependence, orientation dependence, and spatial resolution. Additionally, dosimeters can measure absorbed dose in relative, reference or absolute terms. The difference between these is explained next.

1.3.2.1. Relative, Reference and Absolute Dosimetry

Absolute dosimetry refers to dosimetric measurements which are based on a physical quantity that can be converted to absorbed dose through fundamental principles, i.e. a physical quantity from which the dose can be directly inferred. Typically, absolute dosimetry is done in primary standards dosimetry laboratories (PSDLs) through a primary standard dosimetry method (ionometry, Fricke dosimetry, or absorbed dose calorimetry — these and other dosimetry methodologies will be introduced at the end of this chapter).

In the clinical environment, reference dosimetry usually refers to calibration measurements performed with ionization chambers that have a calibration coefficient ($N_{D,\mathrm{w}}$) traceable to a PSDL. This coefficient is obtained in specific irradiation conditions – with predefined field sizes, source-to-detector distances, and in specific mediums. The calibration is thus only applicable under these well-established reference measurement conditions. Users can recreate these reference conditions at their institution following a reference dosimetry protocol like the AAPM's TG-51.[14,15]

Relative dosimetry refers to any dosimetric measurements performed with a detector used outside of reference conditions, in which the absorbed dose cannot be directly inferred from the physical quantity being measured. In the case of calibrated ionization chambers, since the calibration coefficient is only applicable in reference conditions, conversion and correction factors are needed to arrive at the correct absorbed dose in non-reference conditions. Thus, the measurements are done relative to reference conditions, and the absorbed dose is determined indirectly.

Having presented some of the basic properties used to describe dosimeters and dosimetric practices in radiotherapy, it is now possible to briefly introduce a limited list of some of the most common types of dosimeters encountered in the clinical environment.

1.3.2.2. Ionization Chambers

Ionization chambers can be gas-filled or liquid-filled. A typical ionization chamber will consist of a small sensitive volume with a collecting electrode in the middle, and an oppositely charged electrode surrounding the entire volume. These electrodes create an electric field between them. The polarity of the electrodes can be reversed. When ionizing radiation pierces the medium of the sensitive volume, ions are created and they are attracted towards the central collecting electrode (or the outer electrode, depending on their charge). An electrometer connected to the collecting electrode can measure the charge, effectively measuring the number of ions created, from which the radiation dose can be inferred. The absorbed dose-to-water is given by:

$$D_{\mathrm{w}} = \left(\frac{\overline{W}}{e}\right)_{\mathrm{air}} \left(\frac{Q}{m_{\mathrm{air}}}\right) s_{\mathrm{wall,air}} \prod_i k_i \qquad (1.9)$$

In Eq. (1.9), $\left(\frac{\overline{W}}{e}\right)_{\mathrm{air}}$ is the average energy required to produce a cation-anion pair in air, $\left(\frac{Q}{m_{\mathrm{air}}}\right)$ represents the amount of charge collected per unit mass of air, and $s_{\mathrm{wall,air}}$ is the stopping power ratio between the wall material and air, which is dictated by the thickness of the chamber walls. The k_i factors account for fluence perturbation effects, like those introduced by the presence of

the wall. This equation is used to determine absorbed dose-to-water from the reading of an air-filled ionization chamber. The treatment for other gas-filled ionization chambers is similar.

Due to their higher density, liquid-filled ionization chambers yield stronger signals and their composition is more equivalent to water, requiring smaller corrections. However, they are harder to stabilize and have more complicated dose-rate and temperature dependencies.[16-18]

1.3.2.3. Solid-State Detectors

Solid-state detectors refer to p-n junction diodes like silicon diodes. These semiconductors have two regions with opposite excess electrical charges, referred to as ‘p-type’ and ‘n-type’ materials. This configuration naturally produces an electric field at the interface between the two regions. When exposed to ionizing radiation, electron-hole pairs are created in the diode. These free charges are carried across the junction by the built-in potential between the two surfaces. The absorbed dose can be inferred from the amount of current induced by the radiation, which is measured by an electrometer across the p-n junction. Due to their small sensitive volume, solid state detectors have excellent spatial resolution. However, they suffer from limited lifetimes, as their sensitivity decreases with use.[13,18,19]

1.3.2.4. Diamond Detectors

Diamond detectors are, conceptually, an extension of solid-state detectors. The diamonds used in radiotherapy dosimetry today are artificial, and they are constructed through chemical vapour deposition (CVD). CVD diamonds are built with a known amount of impurities, which introduce an overall charge, making them p- or n-type materials. Because of the high mobility of electrons and holes inside the diamond, these detectors yield a measurable current when irradiated, in the same way semiconductors do. Diamond detectors yield high signals, are waterproof, and are

specially suited for small fields due to their small volume. The principal drawback to diamond detectors is that they have long manufacturing times and high costs.[13]

1.3.2.5. Fricke Dosimeters

Fricke dosimetry consists of measuring a change in optical density, not unlike the principle behind radiochromic film dosimetry. This method is chemical in nature, as the change in optical density is produced by a chemical reaction. Namely, Fricke dosimeters house a ferrous sulfate solution which contains ferrous ions (Fe^{2+}). When exposed to ionizing radiation, these ions can be oxidized into ferric ions (Fe^{3+}), which exhibit an absorption peak in the ultraviolet (UV) spectrum (λ = 304 nm) that the ferrous parent ions do not. The change in molar concentration of Fe^{3+} ions after irradiation will affect the amount of UV light (at 304 nm) that is transmitted through the solution. Consequently, the change in UV transmission can be measured by a spectrophotometer, and this quantity is therefore indicative of the amount of radiation dose absorbed by the solution. Mathematically, the absorbed dose-to-water in Fricke dosimetry is given by:

$$D_w = \frac{\Delta OD}{\rho l}\frac{1}{(\varepsilon G)_{\mathrm{Fe^{3+}}}} f^{D_{\mathrm{sol}}\rightarrow D_{\mathrm{w}}} \prod_i k_i \qquad (1.10)$$

In the above equation, the dosimeter's output is the change in optical density (ΔOD), ρ is the density of the solution, l is the optical path length, $(\varepsilon G)_{\mathrm{Fe^{3+}}}$ is the product of the molar extinction coefficient and the radiation chemical yield (analogous to $\left(\frac{\overline{W}}{e}\right)_{\mathrm{air}}$ for ionization chambers), and $f^{D_{\mathrm{sol}}\rightarrow D_{\mathrm{w}}}$ is the conversion factor from absorbed dose-to-solution to absorbed dose-to-water. The remaining term comprises various correction factors.

When their use is optimized, Fricke dosimeters can achieve reproducibility uncertainties of 0.1 %. However, its low sensitivity requires high doses, and extreme care must be taken to ensure the purity of chemical solutions, not to mention the cost of the readout device (spectrophotometer).[13,18]

1.3.2.6. Ionometry

Previously, ionization chambers were discussed in the context of relative and reference dosimetry, for which most ionization chambers are used. However, ionization chambers can be used to determine absorbed dose-to-water in absolute terms also. This practice is known as ionometry. To do this, the sensitive volume of the ionization chamber and the wall thicknesses must be precisely known. Once these quantities are determined, one can calculate the absorbed dose-to-water directly from first principles, by developing an equation based on cavity theory considerations. This equation is presented below for completeness, but its derivation and meaning are not critically relevant for this thesis. A development of this expression can be found in various works referenced throughout this chapter.[12,13,17,18]

$$D_{\mathrm{w}} = \left(\frac{\overline{W}}{e}\right)_{\mathrm{air}} \left(\frac{Q_{\mathrm{air}}}{V_{\mathrm{air}}\,\rho_{\mathrm{air}}}\right) \left(\frac{\bar{L}_{\Delta}}{\rho}\right)_{\mathrm{air}}^{\mathrm{wall}} \left(\frac{\bar{\mu}_{en}}{\rho}\right)_{\mathrm{wall}}^{\mathrm{w}} \left(\Psi_{\mathrm{wall}}^{\mathrm{w}}\,\beta_{\mathrm{wall}}^{\mathrm{w}}\right) \prod_{i} k_i \qquad (1.11)$$

It is evident from this expression that relative ionization chamber dosimetry is a 'special case' of ionometry, i.e. Eq. (1.11) can be simplified to get Eq. (1.9). Two new quantities appear in Eq. (1.11), the restricted stopping power ratio between the wall and air, $\left(\frac{\bar{L}_{\Delta}}{\rho}\right)_{\mathrm{air}}^{\mathrm{wall}}$, and the ratio of the quotient of dose to collisional kerma for water to wall at the point of measurement, $\beta_{\mathrm{wall}}^{\mathrm{w}}$.

1.3.2.7. Calorimeters

Absorbed dose calorimeters operate on the principle that virtually all the energy from the incoming ionizing radiation is transformed to thermal energy in the medium of interest. Thus, a sufficiently large radiation dose can produce a measurable temperature rise (ΔT) in the absorbing medium. For a graphite calorimeter irradiated by a photon beam, this is about a one mK rise per Gy of radiation. By definition, the specific heat capacity of a given medium ($c_{p,\mathrm{med}}$) is the amount of energy (in joules) required per unit mass (in kilograms) to raise the temperature of the medium by one degree

kelvin. Therefore, the absorbed dose-to-water can be directly obtained from the product of the temperature rise and the specific heat capacity of the medium:

$$D_{\mathrm{w}} = c_{p,\mathrm{med}} \cdot \Delta T \cdot f^{D_{\mathrm{med}} \to D_{\mathrm{w}}} \cdot \prod_i k_i \qquad (1.12)$$

The $f^{D_{\mathrm{med}} \to D_{\mathrm{w}}}$ in Eq. (1.12) is a conversion factor that takes absorbed dose-to-medium to absorbed dose-to-water. Absorbed dose calorimeters are typically made of water or graphite. Of course, the water-to-medium conversion factor is unity for water calorimeters, but it is not for graphite calorimeters. In fact, $f^{D_{\mathrm{med}} \to D_{\mathrm{w}}}$ is one of the principal contributors to absorbed dose uncertainty for graphite calorimeters. The correction factors k_i account for non-ideal experimental conditions, including but not limited to: mass impurities in the medium, heat loss due to energy expended in radiative or chemical processes, and heat loss due to convective motions in water calorimeters.

Water and graphite calorimeters present different benefits. Besides requiring no conversion factors, water calorimeters allow for point-wise absorbed dose measurements due to the low thermal diffusivity of water. However, water calorimeters suffer from large convective motions at room temperature,[20] and care needs to be taken to account for water radiolysis and other radiochemical effects, as well as ensuring that the water remains pure. Instead, graphite has a specific heat capacity that is six times lower than that of water, making its signal six times higher. In addition, it is solid and has a high thermal conductivity, making it useful for operation in isothermal mode. An in-depth discussion of absorbed dose graphite calorimetry, with detailed explanations of isothermal and quasi-adiabatic modes of operation, is provided in the next chapter.[12,13,17,21]

1.4. References

1. Brenner D, Weir H, Demers A, et al. Projected estimates of cancer in Canada in 2020. *Canadian Medical Association Journal*. 2020;192:e199–e205.

2. Canadian Cancer Statistics Advisory Committee. Canadian Cancer Statistics 2019. Toronto, ON: 2019; cancer.ca. Accessed July 2020.

3. America Society of Clinical Oncology. Cancer Terms: Treatment. Alexandria, VA: 2018; cancer.net. Accessed July 2020.

4. Johns HE, Cunningham J. The Physics of Radiology, Fourth Edition. Charles C. Thomas Publisher. Springfield, IL: 1983.

5. Podgorsak E. Radiation Oncology Physics: A Handbook for Teachers and Students. Vienna, AT: International Atomic Energy Agency (IAEA). 2005.

6. Society of Nuclear Medicine & Molecular Imaging. Fact Sheet: Targeted Radionuclide Therapy and Prostate Cancer. Reston, VA: 2019; snmmi.org. Accessed July 2020.

7. International Atomic Energy Agency (IAEA). Yttrium-90 and Rhenium-188 radiopharmaceuticals for Radionuclide Therapy. *Radioisotopes and Radiopharmaceuticals Series*. Vienna, AT: 2015; 5.

8. Pawlicki T, Scanderberg DJ, Starkschall G. Machines for Producing Radiation. In Hendee's *Radiation Therapy Physics*, Fourth Edition. 2016; 35-56.

9. Khan FM, Gibbons JP. Measurement of Ionizing Radiation and Measurement of Absorbed Dose. In Khan's *the Physics of Radiation Therapy*, Fifth Edition. Philadelphia, PA: 2014; 78-101.

10. Jäkel O, Karger CP, Debus J. The future of heavy ion radiotherapy. *Med. Phys.* 2008;35: 5653-5663.

11. ICRU Report 85: Fundamental quantities and units for ionizing radiation. *J. Int. Comm. Radiat. Units Meas*. 2011;11(1):1-31.

12. Seuntjens J, Rogers DW, Cygler JE. Basic Radiation Interactions, Definition of Dosimetric Quantities, and Data Sources. In Siebers' and Hugo's *Clinical Dosimetry Measurements in Radiotherapy*. Madison, WI: 2009; 45-70.

13. Andreo P, Burns DT, Nahum AT, Seuntjens JP, Attix FH. Fundamentals of Ionizing Radiation Dosimetry. Wiley. 2017.

14. Almond PR, Biggs PJ, Coursey BM, et al. AAPM's TG-51 protocol for clinical reference dosimetry of high-energy photon and electron beams. *Med Phys*. 1999;26:1847–1870.

15. McEwen MR, DeWerd LA, Ibbott GS, et al. Addendum to the AAPMs TG-51 protocol for clinical reference dosimetry of high-energy photon beams. *Med Phys*. 2014;41:041501.

16. McEwen M, Rogers DW, Cygler JE. Primary Standards of Air Kerma for ^{60}Co and X-Rays and Absorbed Dose in Photon and Electron Beams. In Siebers' and Hugo's *Clinical Dosimetry Measurements in Radiotherapy*. Madison, WI: 2009.

17. Seuntjens J, Duane S. Photon absorbed dose standards. *Metrologia*. 2009;46:39–58.

18. Podgorsak E. Radiation Physics for Medical Physicists. Springer. Switzerland: 2016.

19. Rogers DW, Cygler JE. General Characteristics of Radiation Dosimeters and a Terminology to Describe Them. In Siebers' and Hugo's *Clinical Dosimetry Measurements in Radiotherapy*. Madison, WI: 2009.

20. Seuntjens J, Kawrakow I, Ross CK. Revisiting convective motion in stagnant water calorimeters operated at room temperatures. National Physical Laboratory Workshop on Water Calorimetry. Teddington, UK: 1999.

21. Renaud J, Palmans H, Sarfehnia A, Seuntjens J. Absorbed dose calorimetry. *Phys Med Biol*. 2020;65(5):05TR02.

Chapter 2. Absorbed Dose Calorimetry

2.1. Calorimetry as an Absorbed Dose Primary Standard

Photon beams in the energy range of 1-30 MV are, in radiotherapy, most usefully characterized through the quantity of absorbed dose-to-water, D_{w}. This is because human tissue is primarily composed of water, and most tissues can be approximated as volumes of water to first order. As discussed in Chapter 1, calorimetry is one of the three main absolute dosimetry methodologies.

2.1.1. Primary Standards in Radiation Dosimetry

Any dosimeter which can determine absorbed dose-to-water in absolute terms can serve as a primary standard, provided it can match the highest accuracy already attainable with other primary standards. A primary standard is any instrument which can obtain the physical quantity of interest (in the case of radiotherapy, D_{w}) without reference to other instruments that measure the same physical quantity. In other words, primary standards in radiation dosimetry can be calibrated in terms of non-radiation standards (e.g. absorbed dose calorimeters rely on temperature standards and/or electrical standards), but they cannot rely on other radiation standards. Primary standards are used to calibrate secondary standards, which are in turn used to calibrate user instruments. To exemplify, absorbed dose calorimeters are generally used to calibrate secondary standard ionization chambers which reside at calibration laboratories. These secondary standards are in turn used to calibrate end-user ionization chambers which are finally used to calibrate the beam output of a linac. This is referred to as the calibration chain, and its purpose is to ensure that all user instruments can be traced back to a national or international primary absorbed dose standard.[1-3]

More generally, the traceability of primary standards and their associated calibration services is organized in the following manner: at a national level, there exist Primary Standards Dosimetry Laboratories (PSDLs) and National Metrology Institutes (NMIs). Many small countries have only a Secondary Standards Dosimetry Laboratory (SSDL) but not a PSDL. These smaller countries

rely on either the International Bureau of Weights and Measures (BIPM), the International Atomic Energy Agency (IAEA), or other PSDLs for traceable calibrations (see figure below). The role of the PSDLs and SSDLs is to calibrate secondary standards and/or user instruments using a primary standard dosimetry methodology. The primary standards housed in PSDLs and SSDLs are ultimately maintained in reference to the relevant radiation dosimetry quantities as defined by the BIPM, the IAEA, or an equivalent body.

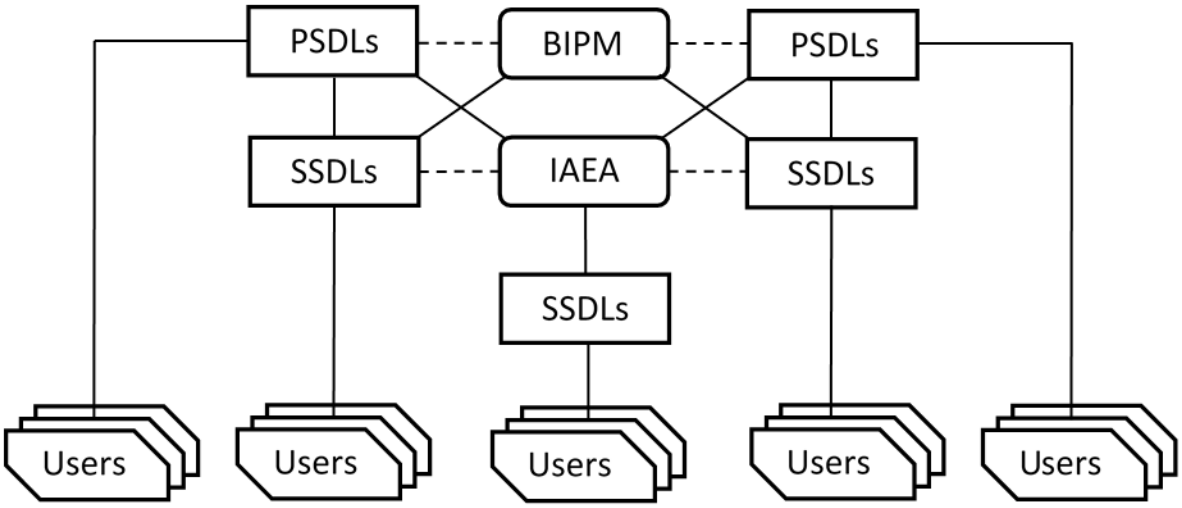

Figure 2.1. This figure represents the international arrangement for traceability in radiation dosimetry. Solid lines indicate calibration services, and dashed lines indicate references or comparisons. Diagram adapted from the IAEA TRS-398 report.[4]

Calorimetry is used as a primary standard in radiation dosimetry. The other two primary standards used are Fricke dosimetry and ionometry, briefly introduced in Chapter 1. Although absorbed dose-to-water can be accurately determined with any of these three methodologies, Fricke dosimetry and ionometry require the radiation field to be thoroughly characterized before obtaining measurements. In this sense, calorimetry can be regarded as the 'most absolute' of the primary standards, and it is the principal system used in most PSDLs.[1,2]

2.2. Materials for Absorbed Dose Calorimeters

As discussed in Chapter 1, absorbed dose calorimeters are usually composed of either water or graphite. Naturally, the quantity of absorbed dose-to-water is more directly obtained in water calorimeters. However, there are inherent and practical aspects of graphite calorimeters that can render them more advantageous depending on the intended application. Graphite has a specific heat capacity that is six times lower than that of water, which means that for an equal dose of ionizing radiation, the temperature rise observed in graphite is six times higher.[5-7] Low specific heat capacities are desirable in calorimeters because the signal is proportional to the temperature rise. For this reason, graphite calorimeters can obtain better signal-to-noise ratios than water calorimeters. Another reason why graphite calorimeters may be favourable is that the chemical heat defect, which refers to the consumption and release of energy in radiation-induced chemical reactions, is considered negligible in graphite.[8] Heat defect can have an additive or subtractive effect on the total temperature rise experienced by the medium as a consequence of irradiation, hence affecting the signal. By comparison, water is more susceptible to radiochemical effects, e.g. radiolysis, which makes heat defect an important consideration.[9]

In terms of practicality, there are more reasons why one might opt for a graphite calorimeter. Graphite has a very high thermal conductivity relative to water. This results in an absorbed dose value that is averaged over the whole graphite volume. Conversely, water calorimeters are better for point-measurements. Another difference is that graphite is solid at room temperature. This is important because it means that graphite bodies can be machined to specification to create custom-shaped, robust configurations. Combined, these considerations make graphite calorimeters arguably more convenient for use in the clinical environment.[3] The major disadvantage of graphite calorimetry is that one requires a dose-to-graphite to dose-to-water conversion factor to account for stopping power and fluence perturbation effects.

2.2.1. Graphite Calorimeters

Graphite calorimeters are generally built in a nested three-body configuration. The inner body is known as the ‘core’, and it is the sensitive volume of the detector. The intermediate bodies, which surround the core, are referred to as the ‘jackets’ (as is the case for this thesis) or ‘inner-shields’. The outer body, enclosing both the jacket and the core, is known as the ‘shield’. Photon beams with conventional radiotherapy dose rates produce temperature rises (that is, signals) on the order of millikelvin, whereas the temperatures in the outside environment can fluctuate by as much as one kelvin in a matter of minutes. Thus, the nested geometry and the extra graphite bodies (the jacket and shield) are used to thermally insulate the core.

To further aid the thermal insulation of the core, the graphite bodies are distanced through air, or more commonly, vacuum gaps. The assembly is held together with threads that are typically made of thermal insulators like silk or polystyrene. The temperature is sensed with thermistors, which are resistors that have temperature-dependent ohmic values, hence the resistance varies according to the temperature of the thermistor’s surrounding environment. Thermistors may also be used for heating by applying a sufficiently large current; this is known as Joule or Ohmic heating.[2,5]

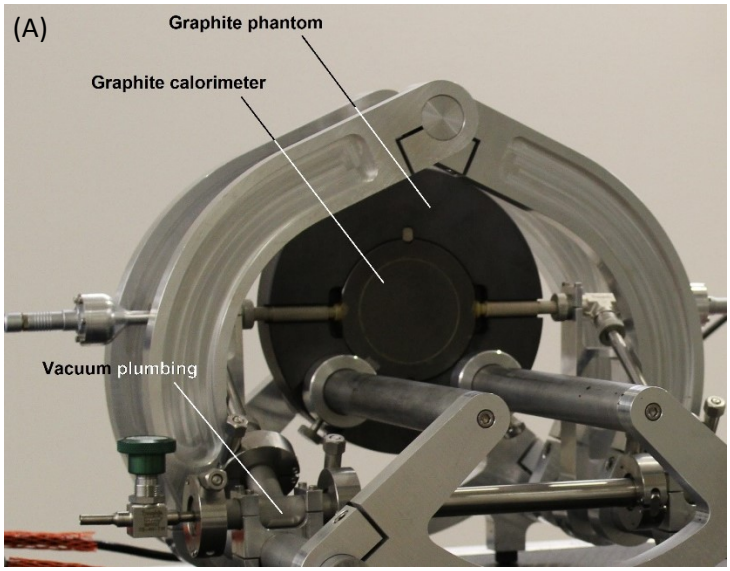

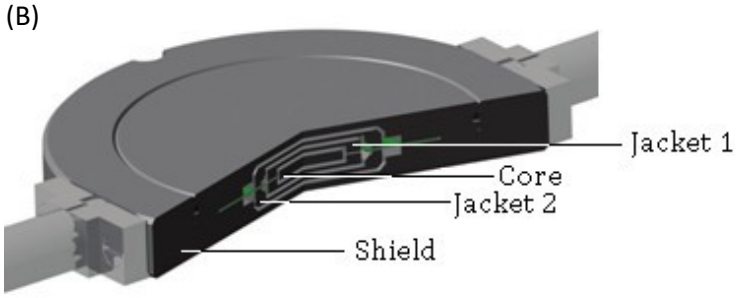

Figure 2.2. (A) General view and (B) cross-sectional diagram of the primary standard absorbed dose graphite calorimeter at the National Physical Laboratory (NPL), the United Kingdom's national PSDL. This calorimeter design is used for dosimetry of MV photons and MeV electrons. Images courtesy of Dr. James Renaud (NPL, 2014).

2.3. Quasi-Adiabatic Mode

To measure absorbed dose, a graphite calorimeter can operate in two modes; quasi-adiabatic mode and isothermal mode. Quasi-adiabatic mode may be thought of as the passive mode of operation. In this mode, the calorimeter is essentially operated as a thermometer. The absorbed dose to the sensitive volume is calculated using the fundamental principle of absorbed dose calorimetry, formulated in Eq. (2.1) and introduced in the last chapter. This principle is summarized here again: Energy from ionizing radiation will be expended primarily in the form of heat, which will produce a temperature rise in the irradiated medium that is proportional to the radiation dose absorbed.

$$D_{\text{med}} = c_P \cdot \Delta T \cdot \prod_i k_i \qquad (2.1)$$

Here, c_P is the specific heat capacity at constant pressure of the medium occupying the sensitive volume, ΔT is the measured temperature rise, and the k_i's represent various correction factors. These corrections and their meanings are explained at the end of this chapter.

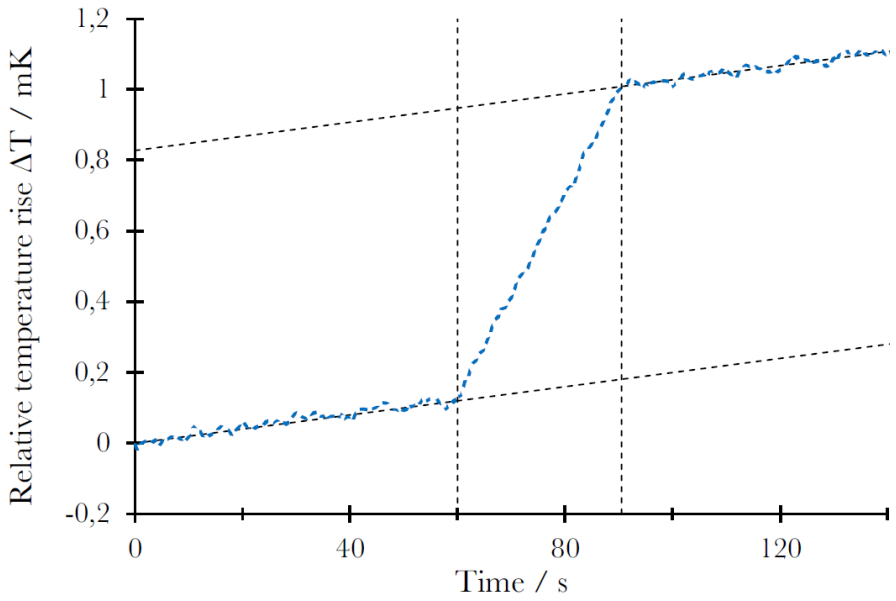

Figure 2.3. A typical quasi-adiabatic run. Initially, the sensitive volume is close to thermal equilibrium; it starts out with a relatively small temperature drift over time, known as 'pre-drift'. The temperature signal will exhibit a rise during irradiation and will arrive to a new, higher temperature after the beam is turned off, followed by a final 'post-drift'. The black vertical dashed lines represent the start and finish of the irradiation period, and the black horizontal dashed lines represent linear fits for the pre-drift and post-drift. Borrowed with permission from Côté (2019).[10]

The mode is termed quasi-adiabatic because, for graphite calorimeters, the temperature rise of the jacket matches (almost exactly) the temperature rise observed in the core. This means there is virtually no heat loss from the core to the jacket, i.e. there is an adiabatic boundary condition imposed on the core. Similarly, the sensitive volume of water calorimeters is also held in an adiabatic condition by placing the sensitive water volume within a large thermally insulated water tank in order to minimize heat dissipation to the outside environment.[2,3,5]

2.4. Isothermal Mode

Isothermal mode is the second mode in which an absorbed dose graphite calorimeter may be operated to measure absorbed dose.[11] This mode can be regarded as the active mode of operation, since it requires real-time adaptive heating to maintain the different graphite bodies at constant

temperatures. This can be achieved in different ways, but regardless of the method employed, it will invariably require some electrical power to dissipate heat into the core. When irradiated, the graphite will experience an absorption of energy, and the amount of electrical power required to keep the graphite bodies at the desired temperature setpoint will be reduced. The resultant drop in power can be quantified and the absorbed dose-to-medium can be inferred from it, as shown in Eq. (2.3). This equation states that the power drop, which is by definition the amount of energy absorbed per unit time, is equal to the absorbed dose rate divided by the total mass of the medium.

$$\dot{D}_{\mathrm{med}} = \frac{P(t)}{m_{\mathrm{med}}} \tag{2.2}$$

$$D_{\mathrm{med}} = \frac{1}{m_{\mathrm{med}}} \cdot \int P(t)\, dt \tag{2.3}$$

In practice, the power drop typically resembles a noisy step function, as seen in Figure 2.4, and the absorbed dose can simply be approximated with Eq. (2.4), where ΔP is the average power drop, Δt is the beam-on time, and the last term again represents all relevant corrections.

$$D_{\mathrm{med}} = \frac{\Delta P \Delta t}{m_{\mathrm{med}}} \cdot \prod_i k_i \tag{2.4}$$

Isothermal mode requires three separate sets of electronic circuits (one for each graphite body) in order to achieve constant temperatures in the core, jacket and shield. Hence, although conceptually not much more difficult than quasi-adiabatic mode, isothermal operation is inherently more complicated to implement practically.[2,5,11]

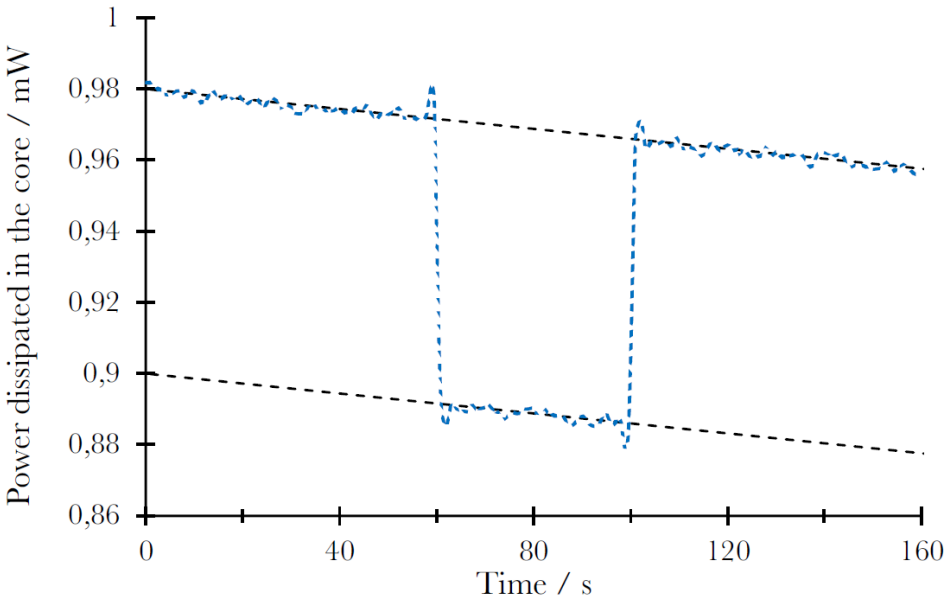

Figure 2.4. A typical isothermal run. Analogous to the quasi-adiabatic run, the power curve will begin with an initial 'pre-drift' and finish with a final 'post-drift'. The black dashed lines represent linear fits to the drifts (i.e. when the beam is off) and to the datapoints during beam-on time. During irradiation, the power required to heat the core to the desired setpoint will be reduced. If the dose rate is constant, the power drop will resemble a step function. Borrowed with permission from Côté (2019).[10]

2.5. Correction Factors for Absorbed Dose Graphite Calorimetry

As demonstrated in the last two sections, accurate absorbed dose determination with a graphite calorimeter requires correction factors k_i whether operating in quasi-adiabatic or isothermal mode. In a full and thorough consideration of inaccuracies, one can find many correction factors to apply to the measured quantity. Indeed, for PSDL calorimeters, correction factors can be numerous. They can include corrections for volume averaging effects that occur due to non-uniform dose distributions which are averaged over the core volume – this is a very important consideration when dealing with smaller radiation field sizes. There are also corrections to account for material impurities, since the core is never composed of pure graphite, and other corrections to account for different kinds of heat transfer effects. One of the inherent aspects of the calorimeter design discussed in this thesis (see Chapter 3) is that concessions are made in accuracy in order to obtain

gains in usability and practical function. For this reason, the discussion on corrections here will be limited to the two most important correction factors: k_{imp} and k_{ht}.

2.5.1. Mass Impurity Correction Factor

The mass impurity correction, k_{imp}, accounts for any non-graphite materials within the core. These can include thermistors, glass beads from the thermistor tips, copper wires, glue, etc. All these materials have physical properties different from those of graphite and hence will affect the absorbed dose to the core. Their composition will dictate whether they increase or decrease the total dose relative to a pure graphite core. Concretely, a non-graphite impurity i with mass m_i will absorb a dose D_i, which will be different from the dose absorbed by an equal mass of graphite, D_{gr}. The dosimetric effect of this impurity can be corrected by taking the ratio of the doses, $\frac{D_i}{D_{\mathrm{gr}}}$, and weighting it by the mass of the impurity relative to the mass of the core, as shown in Eq. (2.5). The summation over index i represents a correction for all impurities.[2,5]

$$k_{\mathrm{imp}} = \frac{m_{\mathrm{core}}}{m_{\mathrm{gr}} + \sum_i m_i \cdot \left(\frac{D_i}{D_{\mathrm{gr}}}\right)} = \frac{m_{\mathrm{gr}} + m_i}{m_{\mathrm{gr}} + \sum_i m_i \cdot \left(\frac{D_i}{D_{\mathrm{gr}}}\right)} \tag{2.5}$$

The above ratio of absorbed doses can be calculated using Monte Carlo methods, but due to the small volume of some of the components involved and their correspondingly small masses, this is computationally expensive and impractical. Instead, these contributions can be approximated using the mass energy absorption coefficient ratio, as seen in Eq. (2.6). This approximation assumes that the photon energy fluence is constant in both the graphite and the impurities, and that the range of the secondary electrons in the impurities is small compared to the dimensions of the impurities themselves. The data for the mass energy absorption coefficients as a function of energy for most materials used in radiotherapy is available at the United States' NIST (National Institute of Standards and Technology) website.[12]

$$k_{\text{imp}} \approx \frac{m_{\text{core}}}{m_{\text{gr}} + \sum_i m_i \cdot \left(\frac{\mu_{\text{en}}}{\rho}\right)_{\text{gr}}^{i}} \tag{2.6}$$

2.5.2. Heat Transfer Correction Factor

The heat transfer correction factor accounts for any non-radiation heat sources that can expose the core to additional heat, hence affecting the total temperature or power output observed. If unaccounted for, this additional heat can be interpreted as part of the energy transferred by the ionizing radiation, and this can lead to either an underestimation or an overestimation of the absorbed dose, depending on which mode the calorimeter is being operated in. Due to the difference in the nature of the measurements, the formulation of k_{ht} is different in the quasi-adiabatic and isothermal scenarios. The starting point is the same:[2]

$$\begin{aligned} \Delta E_{\text{tot,thermal}} &= \Delta E_{\text{rad}} + \Delta E_{\text{elec}} + \Delta E_{\text{transfer}} \\ &= m_{\text{core}} \cdot c_{p,\text{core}} \cdot \Delta T_{\text{core}} \end{aligned} \tag{2.7}$$

The total thermal energy (or the total heat), $\Delta E_{\text{tot,thermal}}$, consists of the sum of thermal energy introduced by the ionizing radiation in the absorbing medium, ΔE_{rad}, the electrical power used to heat the thermistors, ΔE_{elec}, and the heat from other sources, $\Delta E_{\text{transfer}}$. The aim of the heat transfer correction factor is to isolate the radiation term. In quasi-adiabatic mode, the signal is the total temperature rise and there is no adaptive heating, so ΔE_{elec} is zero. Thus:

$$\Delta E_{\text{rad}} = k_{\text{ht}} \cdot \Delta E_{\text{tot,thermal}} \tag{2.8}$$

$$k_{\text{ht}} = 1 - \frac{\Delta E_{\text{transfer}}}{m_{\text{core}} \cdot c_{p,\text{core}} \cdot \Delta T_{\text{core}}} \tag{2.9}$$

In isothermal mode, the signal is the amount of electrical power required to maintain the core at a set temperature. The heat transfer correction factor is therefore:

$$\Delta E_{\text{rad}} = k_{\text{ht}} \cdot \Delta E_{\text{elec}} \tag{2.10}$$

$$k_{\text{ht}} = \frac{m_{\text{core}} \cdot c_{p,\text{core}} \cdot \Delta T_{\text{core}}}{\Delta E_{\text{elec}}} - \frac{\Delta E_{\text{transfer}}}{\Delta E_{\text{elec}}} - 1 \tag{2.11}$$

In practice, k_{ht} is usually approximated by computing the ratio of the (numerically simulated) absorbed dose obtained from an ideal signal in a scenario without any heat transfer, and the absorbed dose obtained from a signal in a normal scenario with heat transfer.

$$k_{\mathrm{ht}} \approx \frac{D_{\text{without heat transfer}}}{D_{\text{with heat transfer}}} \tag{2.12}$$

This factor is calculated computationally. The physics, software specifications and the general procedure associated with these simulations are a major focus of this thesis and will be discussed in detail in Chapter 4 and onwards. For this discussion, it suffices to say that k_{ht} is determined by (a) simulating a detector signal with the usual heat transfer equations and with the relevant thermal properties of each material in the calorimeter design, (b) simulating another signal in which the heat transfer coefficients of all materials is set to zero, and then (c) taking the ratio between the two generated signals as shown in Eq. (2.12).

It is important to note that there will always be some heat transfer between the different bodies of the calorimeter, but most of this heat transfer is usually accounted for with the signal analysis method. Thus, k_{ht} represents any excess heat transfer in addition to that already accounted for implicitly in the signal analysis step.

Before concluding the discussion on heat transfer corrections, it is worthwhile to clarify that the correction factor for heat defect, k_{hd}, is treated separately from k_{ht}. The heat defect correction factor accounts for any discrepancies in the measured absorbed dose resulting from the fact that not all the energy from the incoming radiation is transformed into heat. Once irradiated, the absorbing material can experience radiation-induced exothermic or endothermic chemical reactions, as mentioned at the beginning of this chapter. These will add or remove heat from the medium and affect the total signal and hence affect the final absorbed dose value obtained. This is easily seen in the case of quasi-adiabatic operation, since the signal is the temperature rise itself.

For graphite calorimetry, radiochemical considerations are not so important, and the heat defect is assumed to be negligible.[2,5] For the work in this thesis, k_{hd} will always be assumed to be unity.

2.5.3. Medium Conversion Factor

To get D_{w}, a multiplicative factor $f^{D_{\mathrm{core}} \rightarrow D_{\mathrm{w}}}$ is used to convert the absorbed dose-to-graphite value. Strictly speaking, $f^{D_{\mathrm{core}} \rightarrow D_{\mathrm{w}}}$ is not a correction factor but rather a conversion factor, as it does not adjust the dose value to account for potential inaccuracies, but rather fundamentally changes the meaning of the quantity itself. Nevertheless, the absorbed dose-to-graphite to absorbed dose-to-water conversion factor is one of the most prominent sources of uncertainty. The medium conversion factor can be obtained analytically. However, absorbed dose calorimeters possess increasingly complex geometries and material compositions, making it extremely difficult to obtain an accurate medium conversion factor analytically. Thus, the conversion factor is most frequently calculated using Monte Carlo radiation transport techniques where the beam sources and calorimeter designs are modelled numerically.[10] The details of the calorimeter design discussed in this thesis are found in the next chapter.

2.5.4. Sources and Types of Uncertainty

The uncertainty on a numerical quantity is expressed as a range of values, going from a minimum to a maximum bound, in which the true value of the quantity of interest can be said to exist with a given level of confidence. The uncertainty of a value is not the same as its error; an error is the difference between a measured quantity and its true or expected value.

Uncertainties can be classified as Type A or Type B. Type A uncertainties are determined from a series of repeatable measurements and are therefore estimated with statistical analysis methods. Type B uncertainties are simply those that are not like Type A and can be regarded as quantifying intrinsic biases in the measurements. Type B uncertainties are estimated using relevant available

information, e.g. estimates from literature or from previous empirical knowledge.[13] In graphite calorimetry, there are usually more Type B uncertainties than Type A, and they are also generally larger in magnitude. Type A uncertainties include uncertainties associated with measurement reproducibility and repeatability, accuracy in linac set-up and field geometry, etc. Type B uncertainties include uncertainties associated with thermistor calibration curves, specific heat capacity reference values, data fits, correction factors, etc. Combined relative standard uncertainty on absorbed dose-to-water values using primary standard graphite calorimeters can be as low as 0.4–0.5 % with careful operation.[14]

2.6. References

1. Andreo P, Burns DT, Nahum AT, Seuntjens JP, Attix FH. Fundamentals of Ionizing Radiation Dosimetry. Wiley. 2017.

2. Renaud J, Palmans H, Sarfehnia A, Seuntjens J. Absorbed dose calorimetry. *Phys Med Biol*. 2020;65(5):05TR02.

3. Seuntjens J, Duane S. Photon absorbed dose standards. Metrologia. 2009;46:39–58.

4. International Atomic Energy Agency (IAEA). Absorbed dose determination in external beam radiotherapy: An International Code of Practice for Dosimetry based on Standards of Absorbed Dose to Water. *Technical Report Series* 398. Vienna, AT: 2006.

5. Renaud, J. On the Development of Absorbed Dose Calorimeter Systems for Absolute Clinical Dosimetry. PhD Thesis, McGill U. Department of Physics. Montreal, QC: 2016.

6. Osborne NS, Stimson HF, Grinnings D. Measurements of heat capacity and heat of vaporization of water in the range 0 to 100 C. *J. Res. Natl. Bur. Stand*. 1939;23:197-259.

7. Picard S, Burns DT, Roger P. Determination of the specific heat capacity of a graphite sample using absolute and differential methods. *Metrologia*. 2007;44:294–302.

8. Domen SR, Lamperti PJ. A heat-loss compensated calorimeter: Theory, design and performance. *J. Res. Natl. Bur. Stand.* 1974;78:595-610.

9. Klassen NV, Ross CK. Water Calorimetry: The Heat Defect. *J. Res. Natl. Inst. Stand. Technol*. 1997;102:63-74.

10. Côté, B. Numerical design and development of a probe-format graphite absorbed dose calorimeter for use in small-field radiotherapy. MSc Thesis, McGill U. Department of Physics. Montreal, QC: 2019.

11. Daures J, Ostrowski A. New constant-temperature operating mode for graphite calorimeter at LNE-LNHB. *Phys. Med. Biol*. 2005;50(17):4035-52.

12. Hubbell JH, Seltzer SM. Tables of X-Ray Mass Attenuation Coefficients and Mass Energy-Absorption Coefficients 1 keV to 20 MeV for Elements Z = 1 to 92 and 48 Additional Substances of Dosimetric Interest. Ionizing Radiation Div. of the National Inst. of Standards and Technology. Gaithersburg, MD; 1995.

13. The NIST Reference on Constants, Units and Uncertainties. Uncertainty of Measurement Results. America Society of Clinical Oncology. Cancer Terms: Treatment. Gaithersburg, MD: 1998; physics.nist.gov/cuu/Uncertainty. Accessed July 2020.

14. Laboratoire national de métrologie et d'essais (LNE). Radiotherapy: Graphite calorimetry, a basis of the traceability chain. Paris, FR: 2010.

Chapter 3. The Aerrow MK7 Calorimetry System

3.1. Small-Field Dosimetry

Accurate dosimetry of small and nonstandard fields is relatively difficult. Ideally, absorbed dose measurements should be acquired at a specific point in space. Part of the challenge in obtaining accurate and reliable dosimetric measurements in small fields lies in the fact that dosimeters have finite sizes and therefore they have sensitive volumes rather than sensitive points-of-measurement. Owing to this reality, any dosimeter will produce, to some extent, a volume-averaging effect (VAE) of the dose profile over its sensitive volume.[1] An additional complication is that the fluence perturbation effects are much more prominent when a large fraction of the total area of the field is occupied by the detector, and hence larger corrections are needed for perturbations at small fields. When the size of the sensitive volume is large compared to the size of the field itself, the fluence perturbation effect and the VAE can create significantly large discrepancies between the real and measured absorbed dose, thereby producing larger uncertainties. This is problematic because it can lead to the miscalibration of a clinical instrument and potentially affect the quality of patient treatment. Thus, it is important to have appropriately sized detectors that have been specifically designed for small-field dosimetry.

3.1.1. Small-Field Physics

To begin the discussion on small-field dosimetry, it is necessary to define the term 'small field' in a more quantitative and precise manner.

3.1.1.1. Definition of a Small Field

Any field is considered 'small' when one or more of the following three conditions is observed:[1,2]

i. There is a loss of lateral charged particle equilibrium (LCPE) along the beam axis.

ii. There is partial occlusion of the primary photon source by the collimation devices.

iii. The detector size is large compared to the field size at the depth of measurement.

3.1.1.2. Loss of LCPE

The first condition, the loss of LCPE, is related to the range of the secondary electrons created by a photon beam. In broad photon fields, charged particles set in motion in regions outside the central beam axis may traverse into the central axis region and 'replace' charged particles that were set in motion inside the central axis region but have traversed outside of it. This fluence-compensating condition is known as transient charged particle equilibrium (TCPE). LCPE refers to the lateral component of TCPE, i.e. charged particles traversing into or out of regions to the side of (rather than along) the central beam axis. When the photon field is sufficiently narrow, the range of the secondary electrons becomes larger than the field's width and LCPE is lost. The loss of LCPE will affect the absorbed dose on the central axis.[3]

Figure 3.1. A schematic depiction of (A) LCPE in a broad photon field and (B) loss of LCPE in a narrow photon field. In (A), external secondary electrons (yellow) enter laterally into the central axis region (red) and compensate for the secondary electrons exiting the central axis region, but there is no such compensation in (B).

One can estimate the minimum field radius required for LCPE (r_{LCPE}) from empirical relations given in the literature.[1,2] This parameter is dependent on the beam quality because the range of the secondary electrons depends on the energy of the beam. For example, a 6 MV photon beam with a standard 10 × 10 cm^2 field size and an SSD of 100 cm will have a r_{LCPE} at a depth of 10 cm in water of approximately 1.1 cm.

3.1.1.3. Primary Source Occlusion

The second condition is related to the devices used for collimating the beam. Modern clinical accelerators employ many kinds of collimators. To exemplify with an isocentric linac, the collimation is typically done with collimation jaws and with multi-leaf collimators (MLCs). When the photon field is sufficiently narrow (i.e. when the collimators are brought very close together) the collimators begin to occlude part of the primary beam itself, as shown in Figure 3.2. This causes the photon penumbrae to overlap, resulting in a significant decrease in beam output. Primary source occlusion affects both the energy and the angular fluence profile of the beam, and hence affects the response of detectors in small fields.[3]

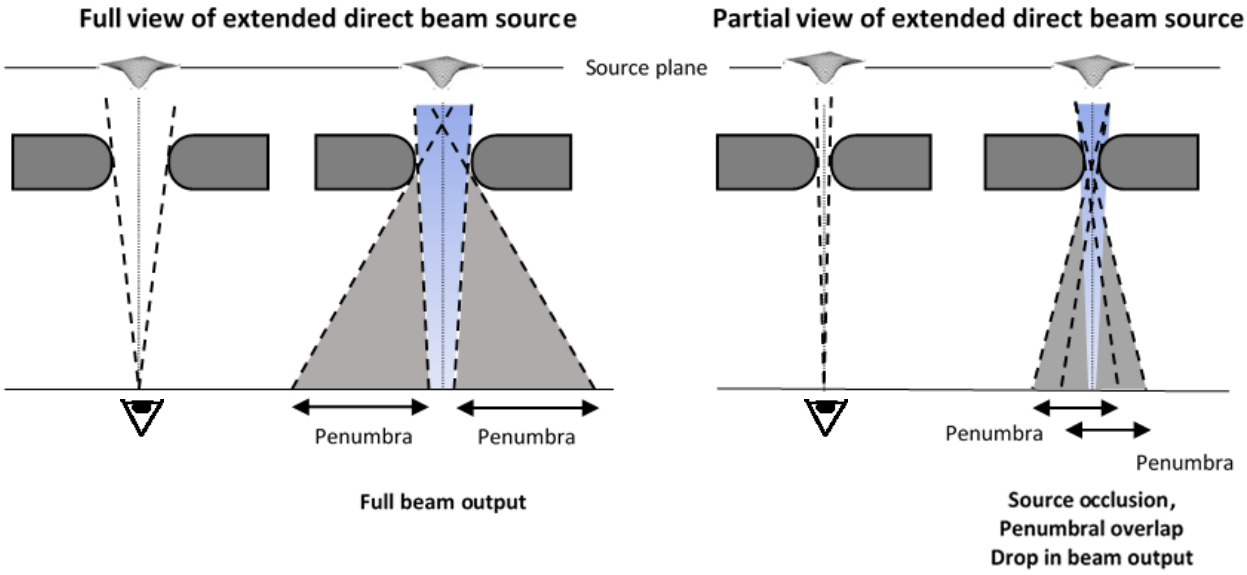

Figure 3.2. Schematic representation of the source occlusion effect. Left: A wide photon field with full beam output. Right: A narrowly collimated photon field with reduced beam output. Diagram adapted from Aspradakis et al. (2010).[4]

3.1.1.4. Detector Size

The third condition is detector-related rather than beam-related and is considerably more subjective; it states that if the physical dimensions of the detector are similar or larger than the dimensions of the field, then the field can also be considered ‘small’. When this is the case, dosimetric measurements will suffer from both the VAE and the fluence perturbation effect, as

explained in the introduction. As will become apparent throughout this chapter, these two effects are always interrelated and further complicate small-field dosimetry.

To specify this condition in a more quantitative manner, it is helpful to consider that most ionization chambers used in radiotherapy have sensitive volumes in the range of 0.2 – 0.6 cm^3.[5] The individual axis dimensions of these sensitive volumes can vary considerably between different chambers due to the broad range of shapes and intended functions, but most have collecting lengths of around 2.0 cm. Furthermore, the small-field range is considered to begin at a nominal square-field size of 2 × 2 cm^2. This seemingly arbitrary boundary is derived in part from the range of commonly used detector sizes, as well as the LCPE and collimator considerations discussed above.

This basic and introductory discussion might seem to suggest that the challenges of small-field dosimetry could be simply overcome by developing a sufficiently small ionization chamber. While this is theoretically correct, it is not so easy to accomplish in practice. Besides size considerations for engineering and manufacturing, there is an important limiting factor that presents itself when miniaturizing an ionization chamber (or any other radiation detector); namely, that the signal decreases proportionally to the decrease in mass and/or volume. This effect makes it difficult to produce ionization chambers with extremely small sensitive volumes. Current state-of-the-art small-field ionization chambers have volumes of around 0.02 cm^3.[6]

3.1.2. Small-Field Correction Factors

In light of these effects and challenges, correction factors are used to adjust absorbed dose readings in small-field dosimetry, and they are applied as multiplicative factors in the same way as those discussed in Chapter 2. To this end, the field output factor, defined in Eq. (3.1), is used to characterize the response of a given dosimeter in small fields.

$$\Omega_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}} = \frac{D_{\text{w},Q_{\text{clin}}}^{f_{\text{clin}}}}{D_{\text{w},Q_{\text{msr}}}^{f_{\text{msr}}}} = \frac{M_{Q_{\text{clin}}}^{f_{\text{clin}}}}{M_{Q_{\text{msr}}}^{f_{\text{msr}}}} \, k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}} \tag{3.1}$$

The subscripts and symbols in this equation are defined in Table 3.1 below and follow the formalism presented in small-field dosimetry protocols such as the IAEA-AAPM TRS-483 and the ICRU Report 91, referenced throughout this chapter.[1-4]

Table 3.1. Quantities and symbols used for describing the field output factor for small-field dosimetry.

Symbol	Description
f_{clin}	Clinical nonreference field.
Q_{clin}	Beam quality of a clinical nonreference beam.
f_{msr}	Machine-specific reference field.
Q_{msr}	Beam quality of a machine-specific reference beam.
$D_{\text{w},Q}^{f}$	Absorbed dose-to-water at a given reference depth (z_{ref}) in a water phantom irradiated with a beam of quality Q and field size f.
M_{Q}^{f}	Dosimeter reading at a given reference depth (z_{ref}) in a water phantom irradiated with a beam of quality Q and field size f, corrected for quantities other than beam quality.
$\Omega_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$	Field output factor.
$k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$	Small-field output correction factor.

Therefore, the field output factor $\Omega_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$ is the ratio of the absorbed dose-to-water in a small or nonreference (clinical) field to that in a reference field, at a specific depth in water (usually z_{ref}). From this definition, it becomes apparent that all the small-field corrections for a specific detector are contained within the small-field output correction factor, $k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$. One can interpret the small-field output correction factor as being the ratio of the detector medium dose conversion factors in the small (or nonreference) field and in the machine-specific reference field. An alternate way to state this is that $k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$ transforms the detector medium conversion factor from $f_{\text{msr}}^{D_{\text{det}}\rightarrow D_{\text{w}}}$ to $f_{\text{clin}}^{D_{\text{det}}\rightarrow D_{\text{w}}}$, as shown in Eq. (3.2).

$$k_{Q_{\mathrm{clin}},Q_{\mathrm{msr}}}^{f_{\mathrm{clin}},f_{\mathrm{msr}}} = \frac{D_{\mathrm{w},Q_{\mathrm{clin}}}^{f_{\mathrm{clin}}} / \overline{D}_{\mathrm{det},Q_{\mathrm{clin}}}^{f_{\mathrm{clin}}}}{D_{\mathrm{w},Q_{\mathrm{msr}}}^{f_{\mathrm{msr}}} / \overline{D}_{\mathrm{det},Q_{\mathrm{msr}}}^{f_{\mathrm{msr}}}} = \frac{f_{\mathrm{clin}}^{D_{\mathrm{det}} \to D_{\mathrm{w}}}}{f_{\mathrm{msr}}^{D_{\mathrm{det}} \to D_{\mathrm{w}}}} \tag{3.2}$$

Here, $\overline{D}_{\mathrm{det},Q}^{f}$ can refer to either the corrected detector reading (M_Q^f) as formulated before, or to an approximate substitute for the detector reading: the absorbed dose-to-detector as computed from Monte Carlo transport simulations. As this suggests, the small-field output correction factor can be calculated either experimentally or approximated via numerical methods.

To summarize, the two effects which (although present in any field size) are of special importance when dealing with small fields are (a) the fluence perturbation effect due to the presence of the detector and (b) the VAE due to the detector's finite size. The field output factor formalism was adopted to correct for the former.[2] The radial beam profile correction (P_{rp}) is used to correct for the VAE and was originally developed to correct for volume-averaging of non-uniformities in the beam profile, even in larger reference fields.[7]

3.2. The Aerrow: A Graphite Probe Calorimeter

Having established the necessary background in graphite calorimetry and small-field dosimetry, it is now possible to examine the novel absorbed dose calorimeter design that is central to this thesis.

3.2.1. Introduction to the Aerrow Calorimetry System

Until recently, absorbed dose calorimeters have been housed exclusively in PSDLs and NMIs.[8,9] There are many good reasons for this; in addition to the complex technicalities associated with maintaining a primary standard, the operation of an absorbed dose calorimeter requires:

i. Space: Absorbed dose calorimeters with their cooling/vacuum and readout systems are generally large and bulky.

ii. Cost: Absorbed dose calorimeters require laboratory-grade instrumentation and appropriate facility resources, in addition to custom fabricated components.

iii. Expertise: Absorbed dose calorimeters are extremely sensitive and their operation and setup is non-trivial; they require specialized technical expertise from the user.

For these reasons, absorbed dose calorimeters have not become a popular dosimeter choice in the clinical environment where time and cost efficiency are important considerations. Since 2012, the absorbed dose calorimetry group at McGill University has developed and reported on a calorimetry system, termed Aerrow, that has been designed in a familiar probe format in an attempt to translate this technology from the PSDLs to the radiotherapy clinic.[10,11] This compact and easily-portable calorimetry system has been successfully demonstrated to provide accurate absorbed dose measurements in high-energy photon fields.[12] Aerrow is usable in either water or solid phantoms. Moreover, due to the magnetic field independence of its response, Aerrow presents advantages in dosimetry of magnetic resonance image-guided radiotherapy machines. In short, it can provide absorbed dose measurements at least as accurate as reference-class ionization chambers, in absolute terms, and without the need for magnetic field correction factors.[13]

The usage of specialized EBRT accelerators that are no longer capable of producing reference fields (e.g. CyberKnife®, Gamma Knife®, the upcoming PET-linac by RefleXion™, etc.) continues to grow as the utility of small and irregular fields is increasingly exploited. This fact, combined with the relatively challenging nature of small-field dosimetry, has motivated the production of a new iteration of Aerrow. Aerrow MK7 is a half-scale model that occupies an eighth of the volume of the original. Detectors that can accurately determine absorbed dose in small fields, both in absolute terms and in the clinical environment, do not currently exist. There are distinct advantages to having such dosimeters available. In the context of QA, they could provide an immediate and independent secondary check of field output factors with minimal correction factors required.

It is worthwhile to remark that Aerrow is not currently meant to be operated as a primary standard as it does not match the accuracy of PSDL absorbed dose calorimetry systems. To achieve its goal of clinical implementation, Aerrow necessarily and inherently makes design concessions in favor of usability. This is true as well for Aerrow MK7. For the time being, the explicit intention of the miniaturized prototype is simply to extend the utility of Aerrow to small fields. Accordingly, it has been designed to mimic a thimble ionization chamber in both shape and size. The following sections contain a closer examination of the geometry and materials used in Aerrow MK7.

3.2.2. Aerrow MK7 Geometry

Aerrow MK7 is a cylindrically layered graphite calorimeter that is composed of three graphite bodies (from centermost out: core (C), jacket (J) and shield (S)) and two rigid thermal insulation layers intercalated between them, as shown in Figure 3.3 in the next section. Eight thermistors are embedded into the detector, three for sensing (1C, 1J, 1S) and five for heating (1C, 2J, 2S).[*] The physical dimensions of the Aerrow MK7 components are specified in Table 3.2 below.

The purpose of this configuration has been discussed in the previous chapter: to guard the sensitive volume — in the case of Aerrow MK7, the cylindrical graphite core — from the large temperature changes occurring in the outside environment. Besides attenuating temperature fluctuations in the core, the external graphite bodies are used in isothermal mode to establish a gradient so that heat flows outwards. This ensures that the power drop in the core can be attributed strictly to radiation energy. The thermal insulation layers also help in keeping the graphite bodies at constant temperatures by minimizing heat loss.

[*] Strictly speaking, the components used for heating the jacket and the shield were ‘heating contacts’, not thermistors. These heating contacts were simply copper wires that were heated via electrical currents, i.e. Joule heating. Thus, the temperatures of the jacket and the shield were controlled by placing the tips of these wires in direct contact with the graphite bodies.

Table 3.2. Physical dimensions of the different Aerrow MK7 components. The real dimensions of the components are marginally smaller than those listed here; the values in this table include within them small air gaps that exist between components and are therefore larger. All length measurements have an estimated uncertainty of ± 0.02 mm. The uncertainties on the approximate volumes were estimated by compounding measurement uncertainties and different air gap thickness estimates.[14] Note the nearly equal volumes for the core and the jacket.

Component	Diameter (mm)	Height (mm)	Approximate volume (mm^3)
Thermistor (×1)	0.32	1.25	0.10 ± 0.01
Core	3.05	5.00	36.35 ± 0.02
Aerogel insulation layer 1	3.75	5.70	26.5 ± 0.1
Jacket	4.45	6.40	36.58 ± 0.02
Aerogel insulation layer 2	5.45	7.30	70.7 ± 0.1
Shield	6.45	8.30	100.89 ± 0.04
Acrylic stem	7.80	> 50	—

3.2.3. Aerrow MK7 Materials

The benefits of using graphite have already been discussed in the last chapter; the principal one being a high thermal diffusivity which means that the absorbed dose can be averaged over the entire volume of the core. The thermal conductivity of the graphite used in this calorimeter (Isotropic Graphite Grade G458, Tokai Carbon Co., Ltd.) is approximately 139 $W \cdot m^{-1} \cdot K^{-1}$,[15] and its specific heat capacity at constant pressure is approximately 713 $J \cdot kg^{-1} \cdot K^{-1}$.[16-18] These values are estimates based on the literature. The density value specified by the manufacturer, 1.86 $g \cdot cm^{-3}$, has been experimentally verified.

One aspect in which Aerrow MK7 diverges significantly from the primary standards graphite calorimeters seen in Chapter 2 is the lack of evacuated gaps used to thermally insulate the graphite layers. Although not unreasonably sized, vacuum pumps are expensive and inconvenient to carry in the clinical setting. More importantly, due to its porous nature and its ability to trap air bubbles, evacuating graphite is a time-consuming process and produces an overall fragile setup. These impracticalities defeat the purpose of having a clinically suitable dosimeter.

Instead, Aerrow MK7 incorporates rigid silica aerogel layers for thermal insulation. The implementation of silica aerogel has been the enabling idea behind this miniature probe design.[19] Aerogel offers the greatest thermal insulation of any class of solids.[20] Furthermore, the formulation of aerogel used in this work (Airloy® X103 Strong Aerogel Tiles, Aerogel Technologies, LLC.) is machinable, and therefore provides structural support and keeps the internal assembly stationary and concentric. The thermal properties of this particular aerogel formulation, as estimated from literature or manufacturer specifications, are: c_P = 1050 J·kg^{-1}·K^{-1}, k = 43 mW·m^{-1}·K^{-1}, and ρ = 0.4 g·cm^{-3},[21,22] giving a thermal diffusivity that is ~ 10^3 times lower than graphite. The dosimetric effects of the density and type of aerogel have been extensively investigated.[23]

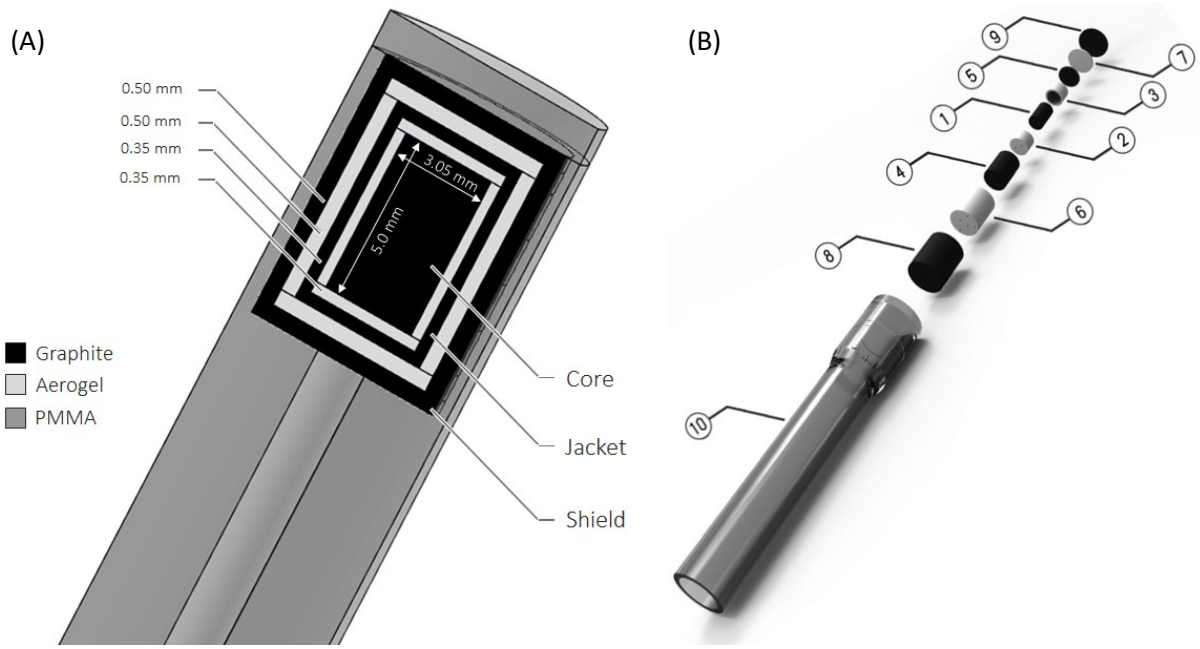

Figure 3.3. (A) Cross-sectional view of a software-rendered 3D model of Aerrow MK7. The numbers on the top left indicate the thicknesses of the layers external to the core. The embedded thermistors are not depicted. (B) Assembly of the larger predecessor prototype, Aerrow, with numbered components. In order: (1) core; (2), (3) first aerogel layer; (4), (5) jacket and jacket cap; (6), (7) second aerogel layer; (8), (9) shield and shield cap; (10) acrylic stem. The small holes at the bottom of the components are where the thermistors are wired through. Images adapted from Renaud et al. (2019).[13]

3.3. Numerical Characterization of Aerrow MK7

Extensive work has been done to characterize Aerrow MK7 numerically.[14] Monte Carlo (MC) transport simulations have been used for a variety of different studies in order to quantitatively analyze the dosimetric properties of this probe. Its thermal properties have been studied with heat transfer simulations using finite element analysis software like COMSOL Multiphysics® (in this thesis) as well as compartmental models in MATLAB® (in previous work). The MC studies and their results are presented below; heat transfer modelling will be introduced in Chapter 5. The bulk of the numerical work revolves around computing the necessary correction factors to the highest possible accuracy. The two major corrections needed for absorbed dose determination with Aerrow MK7, namely k_{ht} and k_{imp}, have already been determined and studied at length in previous work.[14] Thus, their values will be presented here with only a brief overview of the methods employed to attain them.

3.3.1. Mass Impurity Correction Factor

The mass impurity correction factor for Aerrow MK7 has been determined to have a maximum value of 1.002 ± 0.3 % using Monte Carlo techniques to estimate it.[14] A rougher estimate of k_{imp} using the approximate method mentioned in Eq. (2.6) gives an upper bound of 1.003 ± 0.3%. In any case, k_{imp} makes for a minor adjustment of at most 0.3% to the raw absorbed dose value.

3.3.2. Heat Transfer Correction Factor

A simple compartment model was built in MATLAB® by Côté (2019) to simulate the temperature trace of the core when operating in quasi-adiabatic mode.[14] This approximate model consists of a five-body geometry: core, jacket, shield, acrylic, and a thin water layer external to the probe. Each body experiences a single uniform temperature and exchanges heat predominantly with the adjacent bodies, but the model allows for heat exchange to occur between any two components. The rate of heat transfer permitted between two bodies i and j is governed by a user-controlled

coefficient, k_{ij}. The model possesses an adiabatic boundary condition that serves as a heat sink. This boundary represents the large volume of water (beyond the thin water layer represented by the fifth body) contained within the phantom surrounding the probe. The heat transfer rate coefficients (k_{ij}) are manually adjusted until the simulated temperature traces match the experimental ones. Once this is achieved, the new coefficients, $k_{ij_{\mathrm{eff}}}$, become representative of effective heat transfer rates that account for the presence of the aerogel layers.

The procedure outlined in Section 2.5.2 can then be followed: the absorbed dose is extracted from the trace, the simulation is repeated with all $k_{ij} = 0$, and the absorbed dose is extracted again. The k_{ht} factor is then calculated using Eq. (2.12). Using this technique, the heat transfer correction factor above and beyond what is accounted for with the dose analysis procedure was found to be 1.001 ± 0.2 %, making it a relatively unimportant correction.[14]

3.3.3. Medium Conversion Factors

The absorbed dose-to-graphite core (D_{core}) to absorbed dose-to-water (D_{w}) conversion factors were calculated using the EGSnrc code system by computing the ratio of the dose scored in a small volume of water to the dose scored in the sensitive volume of the calorimeter centered at the same location. Table 3.3 below details the Monte Carlo parameters used for these simulations.[24]

Table 3.3. Monte Carlo study specifications for the evaluation of Aerrow MK7's medium conversion factors.

Item name	Description	References
Code version	The modeling of the probe was carried out using the egs_chamber user code with the egs++ geometry package V2017. The model geometry was verified using the egs_view program.	25,26
Validation	Four in-house Monte Carlo beam models were commissioned for Varian TrueBeam® golden data for 6 MV FFF, 6 MV, 10 MV FFF, and 10 MV beams.	27,28
Timing	Approximately 70,000 CPU hours (combined time for all Monte Carlo studies in this paper) distributed over 123 cores. Individual simulations ranged between 10-100 CPU hours.	
Source description	Using BEAMnrc, phase-space files were created for jaw-defined fields of 10 × 10 cm^2 at the surface of a water phantom with an SSD of 100 cm.	29
Cross-sections	NIST ESTAR density corrections were used to generate PEGS4 datasets for all materials. XCOM photon cross sections were used throughout.	
Transport parameters	Parameter: Value ECUT, AE: 512 keV PCUT, AP: 1 keV SMAX: 10^{10} ESTEPE: 0.25 XIMAX: 0.5 EM ESTEPE: 0.02 (water phantom), 0.01 (Aerrow MK7) Boundary crossing algorithm: Exact Electron-step algorithm: PRESTA-II	
VRT parameters	Parameter: Value XCSE: 16 Russian Roulette survival factor: 256 ESAVE: 512 keV	
Number of histories	Between 2–4 × 10^{10} histories to achieve a type A uncertainty better than 0.1 % on each simulation.	
Statistical methods	Combined type B systematic uncertainty of 0.35 % assigned to be consistent with previous MC-simulated conversion factors for Aerrow.	13

D_{core} and D_{w} were calculated for four different photon beam qualities. Using the photon component of the PDD under reference conditions (a parameter known as $\%dd(10)_{\mathrm{X}}$),[5] the beam qualities can be specified as 63.5, 66.4, 71.1 and 73.5 %.

The model recreated reference dosimetry conditions — the centre of the sensitive volume was positioned at a depth of 10 cm in a 30 × 30 × 30 cm^3 water phantom, the SSD was set to 100 cm

and the equivalent square field size at depth (S_{clin}) was set to 11.0 cm. The detector was placed in the horizontal direction, i.e. perpendicular to the beam's central axis. The D_{core} values were then compared to the absorbed dose values scored in a water disk of thickness 0.03 cm and radius of either 0.5 cm for FFF beams or 1 cm for WFF beams. Ideally, D_{w} should be calculated at a point; these dimensions were chosen as the balance between computation time and volume-averaging effect.[30] The resultant $f^{D_{\mathrm{core}} \rightarrow D_{\mathrm{w}}}$ factors are shown below.

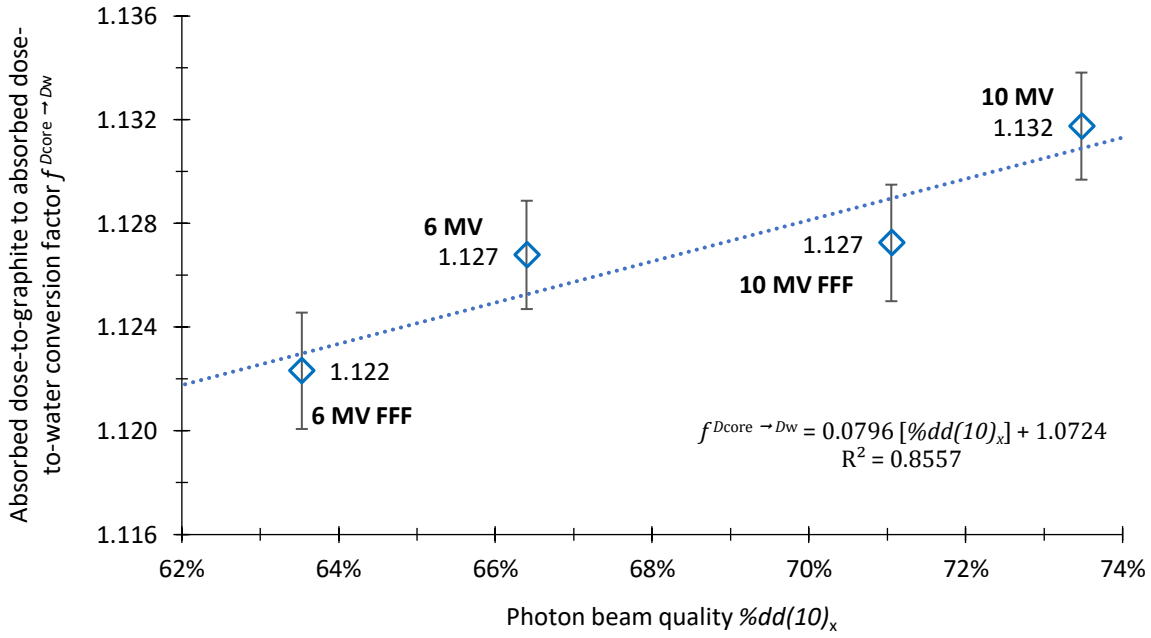

Figure 3.4. Absorbed dose-to-graphite to absorbed dose-to-water conversion factors for Aerrow MK7 for four different beam qualities on the Varian TrueBeam® accelerator under reference conditions. The error bars represent the type A uncertainty (i.e. the Monte Carlo-calculated statistical uncertainty) with a coverage factor of $k = 2$.

The absorbed dose-to-graphite to absorbed dose-to-water conversion factors exhibit a small but steady incline. This increase is a direct consequence of the slight energy dependence seen in the stopping power ratio of water to graphite in this energy range. Nevertheless, the maximum variation on $f^{D_{\mathrm{core}} \rightarrow D_{\mathrm{w}}}$ is only 1.5 %. These values are consistent with the conversion factors determined for the Aerrow MK7's predecessor model.[12,19]

3.3.4. Small-Field Output Correction Factors

Previous work has been conducted to determine the manner in which different silica aerogel densities in the probe can affect Aerrow's performance.[23] However, a new Monte Carlo study was performed to examine the consequences of varying the aerogel density on the small-field output correction factors of Aerrow MK7.

To do this, a BEAMnrc model of a 6 MV photon beam from a Varian Novalis SRS® linear accelerator, which has been extensively validated down to a nominal field size of 0.5 × 0.5 cm^2, was used as the radiation source.[31] The geometrical set-up remained as in the previous study (i.e. recreating reference conditions) with one exception: the detector was simulated in the vertical orientation to minimize the VAE due to the radial non-uniformity of the beam. To determine D_w, a water disk of thickness 0.1 cm and radius 0.05 cm was used as the scoring region for all field sizes. The jaw-defined geometrical field size was varied in the range of 10 × 10 cm^2 to 0.45 × 0.45 cm^2, corresponding to S_{clin} values between 11.0 cm and 0.54 cm. For each field, the aerogel density was varied from 0.01 $g \cdot cm^{-3}$ to 0.55 $g \cdot cm^{-3}$, as seen in Figure 3.5. This range reflects the commercially available densities of the Airloy® X103 aerogel tiles.[22] The small-field output correction factors were calculated for every simulated field size and aerogel density using Eq. (3.2). All the Monte Carlo transport, physics and VRT parameters were kept as in the last study (see Table 3.3).

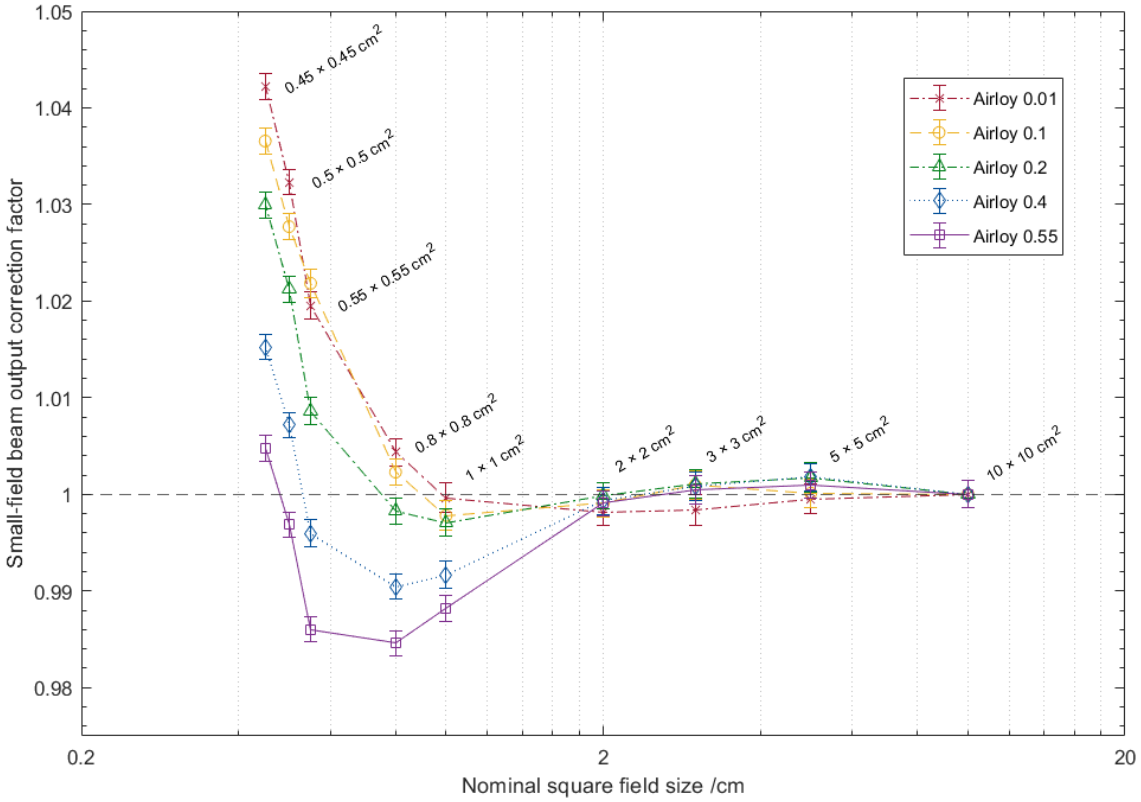

Figure 3.5. MC-calculated small-field output correction factors for Aerrow MK7, $k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$, as a function of nominal field size for a 6 MV photon beam in reference conditions (SSD = 100 cm, z = 10 cm) for aerogel densities ranging from 0.01 g·cm^{-3} to 0.55 g·cm^{-3}. The error bars represent the type A uncertainty with a coverage factor of k = 2.

Since an aerogel density of 0.4 g·cm^{-3} produced the field output correction factors that were closest to unity on average for all field sizes, this was the density that was selected for the construction of the prototype. An isolated representation of the $k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$ values obtained for an aerogel density of 0.4 g·cm^{-3} can be found in Figure 3.6 below.

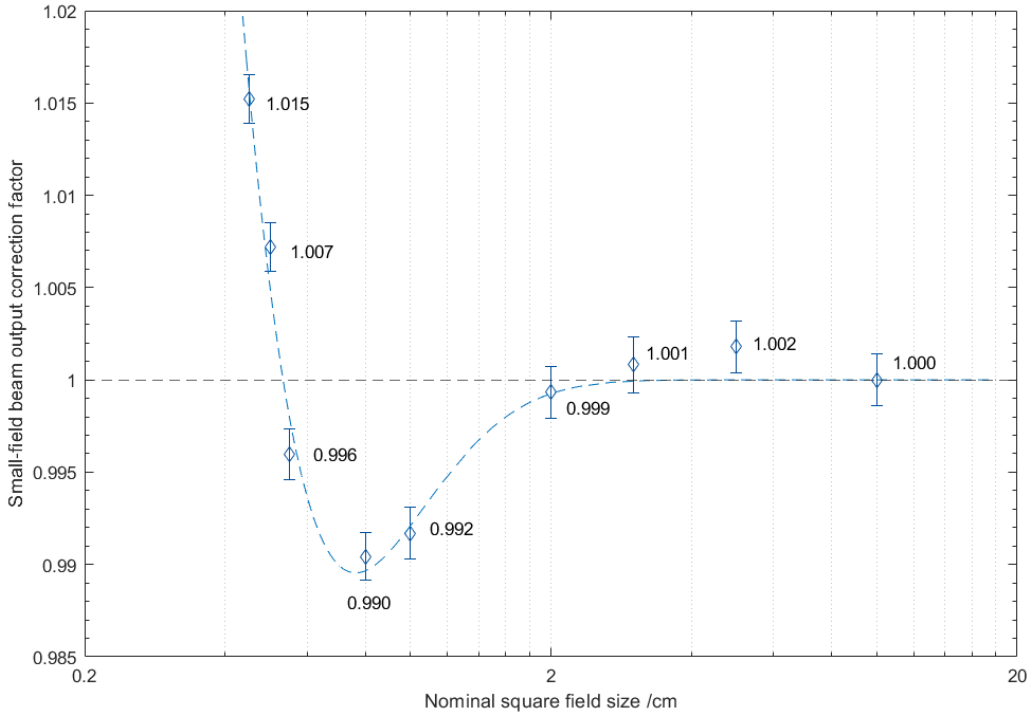

Figure 3.6. MC-calculated small-field output correction factors for Aerrow MK7, $k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$, as a function of nominal field size for a 6 MV photon beam in reference conditions (SSD = 100 cm, z = 10 cm) for the prototype with an aerogel density of 0.4 g·cm^{-3}. The error bars represent the type A uncertainty with a coverage factor of k = 2.

The dashed line in Figure 3.6 represents a fit applied to the data, described by the general function $f(x) = (1 + e^{-ax+b})(1 - e^{-cx+d})$ with coefficients $\{a, b, c, d\}$ = {7.073, 0.2413, 2.480, -2.235}. This function was chosen to reflect the general shape of $k_{Q_{\text{clin}},Q_{\text{msr}}}^{f_{\text{clin}},f_{\text{msr}}}$ as a function of field size that is observed in many other detectors.[2] This shape is driven by two competing factors: the density-dependent fluence perturbation effect due to the presence of the detector and the VAE due to the detector's finite volume. As can be appreciated from this graph, the small-field output correction factors for Aerrow MK7 in a 6 MV beam is unity to within ± 1 % down to an S_{clin} dimension of 0.54 cm. It is worth noting that these corrections are similar to or smaller than the corrections required for other popular small-field dosimeters.[2] Its minor small-field output correction factors and its capability of determining absorbed dose in absolute terms suggest that Aerrow MK7 is a suitable candidate for accurate dosimetry of small and nonstandard fields.

3.4. References

1. Palmans H, Andreo P, Huq MS, et al. Dosimetry of small static fields used in external photon beam radiotherapy: Summary of TRS-483. *Med Phys*. 2018;45: e1123-e1145.

2. International Energy Agency (IAEA). Dosimetry of small static fields used in external beam radiotherapy: an international code of practice for reference and relative dose determination. *Technical Report Series* 483. Vienna, AT: 2017.

3. ICRU Report 91: Prescribing, recording, and reporting of stereotactic treatments with small photon beams. *J. Int. Comm. Radiat. Units Meas*. 2014;14:1–160.

4. Aspradakis M, Byrne JP, Palmans H, et al. IPEM Report 103: Small field MV photon dosimetry. International Atomic Energy Agency (IAEA). Vienna, AT: 2010; 42(14).

5. Almond PR, Biggs PJ, Coursey BM, et al. AAPM's TG-51 protocol for clinical reference dosimetry of high-energy photon and electron beams. *Med Phys*. 1999;26:1847–1870.

6. PTW: The Dosimetry Company. PinPoint Ion Chambers 31014/31015. Freiburg, DE: 2020; ptwdosimetry.com/en/products. Accessed September 2020.

7. McEwen MR, DeWerd LA, Ibbott GS, et al. Addendum to the AAPMs TG-51 protocol for clinical reference dosimetry of high-energy photon beams. *Med Phys*. 2014;41:041501.

8. Seuntjens J, Duane S. Photon absorbed dose standards. *Metrologia*. 2009;46:S39–S58.

9. Renaud J, Palmans H, Sarfehnia A, Seuntjens J. Absorbed dose calorimetry. *Phys. Med. Biol*. 2020;65:05TR02.

10. Renaud, J. Development of a graphite probe calorimeter for absolute clinical dosimetry: Numerical design optimization, prototyping and experimental proof-of-concept. MSc Thesis, McGill U. Department of Physics. Montreal, QC: 2012.

11. Renaud, J. On the Development of Absorbed Dose Calorimeter Systems for Absolute Clinical Dosimetry. PhD Thesis, McGill U. Department of Physics. Montreal, QC: 2016.

12. Renaud J, Sarfehnia A, Bancheri J, Seuntjens J. Aerrow: a probe-format graphite calorimeter for absolute dosimetry of high-energy photon beams in the clinic. *Med Phys*. 2018;45:414–428.

13. Renaud J, Sarfehnia A, Bancheri J, Seuntjens J. Absolute dosimetry of a 1.5 T MR-guided accelerator-based high-energy photon beam in water and solid phantoms using Aerrow. *Med Phys*. 2020;47(3):1291-1304.

14. Côté, B. Numerical design and development of a probe graphite absorbed dose calorimeter for use in small-field radiotherapy. MSc Thesis, McGill U. Department of Physics. Montreal, QC: 2019.

15. Tokai Carbon Co., Ltd. Tokai Carbon - Isotropic Graphite Grade G458: Typical Properties. Oldbury, UK: 2020; tokaicarbon.co.jp. Accessed September 2020.

16. McEligot D, Swank DW, Cottle DL, Valentin, FI. Thermal Properties of G-348 Graphite. Idaho Falls, ID: 2016.

17. Picard S, Burns DT, Roger P. Determination of the specific heat capacity of a graphite sample using absolute and differential methods. *Metrologia*. 2007;44:294–302.

18. Williams AJ, Burns DT, McEwen MR. Measurement of the specific heat capacity of the electron-beam graphite calorimeter. National Physical Laboratory. Teddington, UK: 1993.

19. Renaud J, Marchington D, Seuntjens J, Sarfehnia A. Development of a graphite probe calorimeter for absolute clinical dosimetry. *Med Phys*. 2013;40:020701.

20. Gurav JL, Jung IK, Park HH, Kang ES, Nadargi DY. Silica aerogel: synthesis and applications. *J. Nanomater*. 2010;2010:11 pages.

21. Scheuerpflug P, Hauck M, Fricke J. Thermal properties of silica aerogels between 1.4 and 330 K. *J. Non-Cryst. Solids*. 1992;145:196–201.

22. Aerogel Technologies, LLC. Airloy® X103 Strong Aerogel Tiles. Boston, MA:2020; buyaerogel.com. Accessed September 2020.

23. Bancheri J, Seuntjens J, Sarfehnia A, Renaud J. Density effects of silica aerogel insulation on the performance of a graphite probe calorimeter. *Med Phys*. 2019;46:1874–1882.

24. Sechopoulos I, Rogers DWO, Bazalova-Carter M, et al. RECORDS: improved reporting of Monte Carlo Radiation transport studies: report of the AAPPM Research Committee Task Group 268. *Med Phys*. 2017;45:e1–e5.

25. Wulff J, Zink K, Kawrakow I. Efficiency improvements for ion chamber calculations in high energy photon beams. *Med Phys*. 2008;35:1328 –1336.

26. Kawrakow I. egspp: The EGSnrc C++ class library. NRC Technical Report No. PIRS-899. National Research Council of Canada. Ottawa, ON: 2005.

27. Beyer GP. Commissioning measurements for photon beam data on three TrueBeam linear accelerators, and comparison with Trilogy and Clinac 2100 linear accelerators. *Journal of Applied Clinical Medical Physics*. 2013;14:273-288.

28. Berryman A, Renaud MA, Seuntjens J. Auto-commissioning for Varian Accelerator for the 6 MV, 6 MV FFF, 10 MV, and 10 MV FFF beam qualities. McGill U. Medical Physics Internal Report.

29. Rogers DWO, Walters B, Kawrakow I. BEAMnrc users manual. NRC Technical Report No. PIRS-12:509. National Research Council of Canada. Ottawa, ON: 2009.

30. Muir BR, Rogers DWO. Monte Carlo calculations of kQ, the beam quality conversion factor. *Med Phys*. 2010;37:5939–5950.

31. Papaconstadopoulos P, Tessier F, Seuntjens J. On the correction, perturbation and modification of small field detectors in relative dosimetry. *Phys Med Biol*. 2014;59:5937.

Chapter 4. Absorbed Dose Determination with Aerrow MK7

4.1. Overview of Aerrow MK7's Dosimetric Performance

After numerically characterizing the dosimetric properties of the probe and determining the correction and conversion factors necessary for absorbed dose determination, Aerrow MK7's performance was tested experimentally. Calorimetry runs were performed in both quasi-adiabatic and isothermal mode. To evaluate the accuracy of the results, the absorbed dose measurements of Aerrow MK7 were compared with those of an Exradin A1SL ionization chamber, which is apt for small-field dosimetry and has been calibrated at the National Research Council of Canada (NRC). As will be demonstrated in the following section, it is possible to perform accurate absorbed dose determination with Aerrow MK7 when operating in quasi-adiabatic mode. However, when operating in isothermal mode, the detector suffered from a significant and irregular overresponse; to be precise, absorbed dose-to-water measurements were consistently overestimated by 17–24 %. The cause of this overresponse was initially unknown. Identifying the source of a response problem such as this is a laborious task since there are numerous steps involved in determining D_{w}. Resultantly, many hypotheses can be postulated as to where this overresponse originates from. In fact, the principal objectives of this thesis can be stated in the context of this problem:

i. **To examine and scrutinize all the likely sources in the chain of absorbed dose determination** (going backwards, from computational methods to electronics and construction) **that can potentially produce said overresponse.**
ii. **To identify a problematic step in the chain of absorbed dose determination and establish** (based on these findings) **a working theory as to why the overresponse was observed.**
iii. **To ascertain possible solutions and materialize them into potential amendments for the design of a new prototype, Aerrow MK8.**

In order to achieve these objectives, the COMSOL Multiphysics® modelling software was used to carefully study the heat transfer processes occurring inside the detector. The following sections include a detailed overview of the dosimetric performance of Aerrow MK7, and an introduction to heat transfer modelling in COMSOL will be provided in the next chapter.

4.2. Dosimetric Performance in Quasi-Adiabatic Mode

Before presenting the absorbed dose values, it is worth clarifying the way in which Aerrow MK7 realizes a quasi-adiabatic mode of operation. As explained in Chapter 2, quasi-adiabatic operation receives its name from the fact that the temperature rise of the jacket matches the temperature rise of the core as closely as possible, thereby minimizing heat loss from the core to the jacket. In primary standards graphite calorimeters, this is usually achieved via active-controlled heating in the jacket.[1] Unlike active-controlled heating in isothermal mode, the jacket temperature is made to match the freely rising temperature of the core during irradiation rather than a predefined setpoint. In Aerrow MK7, the quasi-adiabatic condition is realized passively by setting the volume of the jacket (almost) equal to the volume of the core, as seen in Table 3.2. Since both components are made of the same graphite, the mass of the two components is roughly the same, and c_P is the same for both. This results in a nearly identical temperature rise.

4.2.1. Measurement Methodology

The quasi-adiabatic measurements were taken with a clinical Varian TrueBeam® linac inside an in-house built temperature-controlled water phantom. This phantom is routinely used for water calorimetry measurements and can achieve a high level of thermal stability — typically better than 1 mK/min.[2] The SSD was set to 100 cm and the center of the detector core was positioned at a depth of 10 cm within the phantom. The detector was placed in the horizontal direction (i.e. with its longitudinal axis perpendicular to the beam's central axis) and the sensitive volume was aligned with the crosshairs of the linac's light field. Proceeding an overnight stabilization period,

calorimetric measurements were acquired for water temperatures ranging from 4.32 °C to 4.46 °C over a 3-hour period. As mentioned in Chapter 1, it is known that calorimetric measurements suffer from convection effects at room temperature.[3] Performing measurements at a nominal temperature of 4 °C combined with regular stirring of the water phantom eliminates large convective motions within the phantom and hence reduces the amount of heat lost to the water environment.[4] Three sets of measurements were taken with beam parameters as listed below.

Table 4.1. Beam parameters used to obtain absorbed dose measurements from Aerrow MK7 in quasi-adiabatic mode and from Exradin A1SL following reference dosimetry methodology.[5,6] The number of monitor units delivered, the beam quality (Q), the repetition rate, the total time of irradiation (Δt) and the equivalent square field size at depth (S_{clin}) are given for each set of measurements.

Measurement Set	Monitor units delivered (MU)	Q (MV)	Repetition rate (MU·min^{-1})	Δt (s)	S_{clin} (cm)
Set I	600	10 FFF	2400	15.0	11.0
Set II	350	6 FFF	1400	15.0	11.0
Set III	600	10 FFF	2400	15.0	2.16

The subsequent day, the measurement sets were repeated with an NRC-calibrated Exradin A1SL ionization chamber. The irradiation conditions remained the same with the exception that the water phantom was brought up to room temperature. For each run, the irradiation conditions were equalized for the two dosimeters not only in terms of geometric setup and total amount of MU delivered, but also with respect to the repetition rates and total irradiation times. In doing this, any potential bias arising from different beam-on times or dose rates can be discarded when comparing the dose values obtained from each detector.

4.2.2. Absorbed Dose Analysis Method

Figure 2.3. illustrated that the temperature trace, or signal, from a typical quasi-adiabatic run resembles a linear piecewise function composed of a pre-drift, a beam-on period in which the slope is positive and noticeably increased, and a post-drift. When dealing with this type of signal, the

conventional mid-run extrapolation technique can be used to extract the absorbed dose. In the mid-run extrapolation technique, the pre-irradiation and the post-irradiation temperatures are extrapolated to the mid-run timepoint, and the temperature rise is calculated at that point.[7] The underlying assumption in this technique is that the time scale of heat transfer between bodies is large compared to the irradiation period. Thus, the amount of heat dissipated away from the core is not significant throughout the duration of the run, and the beam-on and post-irradiation drift curves are therefore approximately linear.

However, this assumption did not hold true for Aerrow MK7. Instead, the signals from Aerrow MK7 were highly nonlinear and showed prominent heat loss effects, as seen in Figure 4.1 below. The temperature traces exhibited a fall-off similar to that of an exponential decay function. This shape is strongly indicative that the sensitive volume was losing heat at a rapid rate.

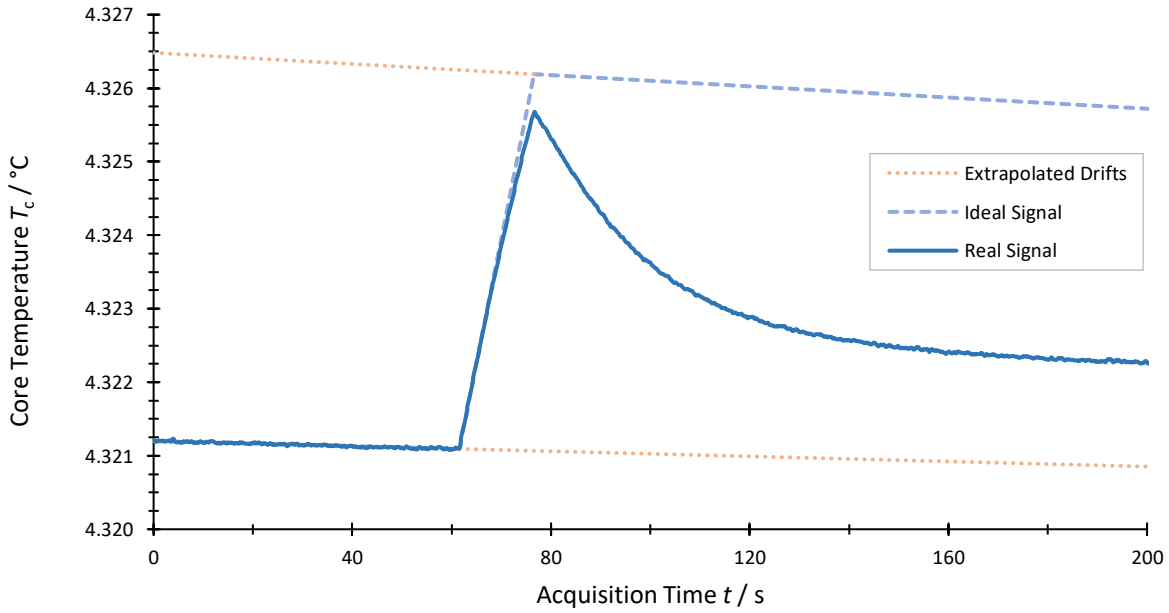

Figure 4.1. Raw signal of a quasi-adiabatic calorimetry run with Aerrow MK7, uncorrected for pre-drift. The irradiation conditions correspond to Measurement Set I, detailed in Table 4.1. The dashed blue line represents an idealized version of the signal with no heat loss effects. The dotted orange lines are extrapolations of the pre-irradiation and post-irradiation drifts, which would be used for the mid-run extrapolation technique.

This pronounced behaviour can be attributed primarily to the large amount of heat being dissipated from the core to the cooler surrounding water phantom, which acts as a heat sink and is many times more massive than the graphite core. Additionally, the water phantom has a rate of heating that is six times lower than that of the graphite within the probe, and this effect further increases heat transfer. The results suggest that Aerrow MK7 contains insufficient thermal insulation and that future measurements may have to be performed in a graphite phantom.

Since the post-irradiation drift curves are not linear in the case of Aerrow MK7, the mid-run extrapolation technique cannot be properly applied to analyze the signals. An alternative methodology was devised in order to extract the absorbed dose. This methodology has already been discussed at length in previous work.[8] In brief, the proposed methodology involves extracting the absorbed dose rate from the initial slope of the beam-on portion of the signal, $\dot{D}(t_0)$, and combining it with the total measured irradiation time Δt to arrive at the absorbed dose. In using this method, it is assumed that the dose rate is constant throughout the irradiation period. Determining the dose rate from the derivative of the signal immediately after beam-on is justified because the probe is essentially in thermal equilibrium with the surrounding water at that instant. Thus, the effects of heat loss seen above are minimal at the very beginning of the irradiation.

4.2.3. Absorbed Dose Measurements in Quasi-Adiabatic Mode

Having introduced the experimental set-up and the signal analysis methodology, it is now possible to present the absorbed dose measurements obtained with Aerrow MK7 when operating quasi-adiabatically. The results are tabulated below. The measurements from the Exradin A1SL ionization chamber are also included for comparison.

Table 4.2. Dose comparison between Aerrow MK7 used in quasi-adiabatic mode and Exradin A1SL used according to the AAPM TG-51 protocol.[5] All reported absorbed dose-to-water values (D_{w}) correspond to 100 MU delivered at $z_{\max}$. The uncertainties ($\mu_{D_{\mathrm{w}}}$) correspond to a coverage factor of k = 1.

Q	10 MV FFF		6 MV FFF		10 MV FFF	
S_{clin}	11.0 cm		11.0 cm		2.16 cm	
	D_{w} (cGy)	$\mu_{D_{\mathrm{w}}}$ (%)	D_{w} (cGy)	$\mu_{D_{\mathrm{w}}}$ (%)	D_{w} (cGy)	$\mu_{D_{\mathrm{w}}}$ (%)
Aerrow MK7	100.8	1.1	100.8	1.2	84.4	1.4
Exradin A1SL	100.4	1.1	100.3	1.1	84.5	1.3
Difference (%)	0.6		0.7		0.3	

Absorbed dose readings acquired with Aerrow MK7 in quasi-adiabatic mode were corrected for heat transfer (k_{ht}). The readings were also corrected with the small-field output correction factor $k_{Q_{\mathrm{clin}},Q_{\mathrm{msr}}}^{f_{\mathrm{clin}},f_{\mathrm{msr}}}$ for the 2 × 2 cm^2 (S_{clin} = 2.16 cm) field.

For the Exradin A1SL dose measurements, corrections were applied to account for voltage polarity (P_{pol}), ion recombination (P_{ion}), and pressure-temperature (P_{TP}) dependencies. These corrections are common and well-documented for ionization chambers.[5,6] According to the TG-51 addendum, collecting volumes with lengths under 5 mm present volume averaging effects below 0.05 % in FFF beams compared to flat beams.[6] Hence, no FFF corrections were applied to the Exradin A1SL absorbed dose readings. An uncertainty of 0.2 % was added to the measurement in the 2 × 2 cm^2 field to account for the small-field perturbation effect.[6]

When calculating the relative difference between the dose values of the two detectors, an additional relative uncertainty of 0.2 % is included to account for any potential drift in the beam output of the linac from Day A (in which Aerrow MK7 measurements were done) to Day B (in which Exradin A1SL measurements were done). This uncertainty was estimated by calculating the average daily drift of the linac's output from a log of QA measurements for the week in which measurements were acquired. With this additional uncertainty, the relative difference in D_{w} is at most 0.7 %, which is well within the combined standard uncertainty $\mu_{D_{\mathrm{w}}}$ for any given

measurement. This level of agreement is noteworthy given that the two dosimeters operate on fundamentally different principles. Furthermore, these results serve as validation for the Monte Carlo model, which seems to hold well even at smaller fields. It can be concluded that Aerrow MK7 is able to accurately determine absorbed dose-to-water in quasi-adiabatic mode when compared to a calibrated reference-class ionization chamber.

4.2.4. Absorbed Dose Uncertainty Estimation

The uncertainties associated with the Exradin A1SL measurements were estimated and calculated using the recommended best practices of the AAPM TG-51 addendum.[6] An estimated uncertainty budget for Aerrow MK7 is presented in Table 4.3 for the 10 MV FFF beam with a nominal field size of 10 × 10 cm^2 (S_{clin} = 11.0 cm). The thermistor calibration procedure listed in Table 4.3 involves benchmarking the thermistor resistance readings (Ω) at given known temperatures (°C) in order to later interpolate temperature values for any measured resistance. The measured irradiation time uncertainty is estimated by taking the temporal resolution of the linac (0.1 s) as a fraction of the entire irradiation period (15.0 s) and treating it as a uniform square distribution. This means that if the beam is set to turn on at some timepoint t = 0 seconds, then it is assumed that the beam is equally likely to turn on at any time between 0 and 0.1 seconds. The resulting number is added to itself in quadrature to account for the uncertainty in the beam-off timepoint as well. The uncertainty on k_{ht} was estimated from a sensitivity analysis of the absorbed dose extraction method.[8] The effects of the linac output ramp-up are considered negligible since the fitting procedure discussed in Section 4.2.2 excludes the datapoints immediately after the beam is turned on, thus justifying the constant dose rate approximation.

Table 4.3. Estimated uncertainty budget for the determination of D_w using Aerrow MK7 in quasi-adiabatic mode for a 10 MV FFF, 10 × 10 cm^2 photon field. All presented uncertainties correspond to a coverage factor of k = 1. All type A uncertainties were estimated using the TG-51 addendum.[6] The thermistor calibration uncertainty reflects the estimated uncertainty associated with fitting the temperature data. The medium conversion factor uncertainty reflects the estimated uncertainty associated with the Type A and Type B (e.g. cross sections) nature of the Monte Carlo calculation. The uncertainty on the specific heat capacity of graphite was estimated based on data in the literature.

Source of uncertainty	Type A (%)	Type B (%)
SSD setting	0.2	
Depth setting	0.2	
Field-size setting	0.2	
Linac stability	0.2	
Reproducibility	0.1	
Thermistor calibration		0.2
Medium conversion factor		0.4
Specific heat capacity value of graphite		0.7
Measured irradiation time		0.5
Heat transfer correction factor		0.2
Quadratic summation	0.4	1.0
Combined standard uncertainty on D_w (k = 1)	1.1	

Table 4.3 shows that the largest sources of uncertainty come from: (a) the graphite-to-water dose conversion factor calculated with Monte Carlo (0.4 %), (b) the specific heat capacity value of graphite (0.7 %), and (c) the measured irradiation time, which is used for calculating the absorbed dose (0.5 %). As explained in previous chapters, the medium conversion factor is expected to be one of the largest contributors to the combined standard uncertainty. The specific heat capacity value of graphite, which is estimated from the literature, should ideally not introduce such a large uncertainty. This uncertainty can be most effectively reduced by designing an experimental set-up for in-house determination of c_P values.[9] Because the dose-extraction methodology adopted in this work involves the measured irradiation time, longer beam-on periods are desirable in order to reduce the uncertainty associated with beam timing. Alternatively, beam timing accuracy can be improved by embedding a fast-sampling diode in the detector stem. As it is, the combined standard uncertainty on the D_w measurements obtained with Aerrow MK7 in quasi-adiabatic mode is

already relatively low (1.1–1.4 %) — this is on par with that of an ionization chamber under reference conditions.

4.3. Dosimetric Performance in Isothermal Mode

Aerrow MK7 was not capable of producing accurate dose measurements in isothermal mode. This conclusion is not altogether unexpected given that isothermal mode is inherently more complex, as explained in Chapter 2. The difficulty is compounded by the miniaturization of the probe, which brings with it many challenges. The absorbed dose measurements obtained in isothermal mode are presented below. In keeping with the principal objectives stated in Section 4.1, it is necessary to discredit as many sources of error as possible in order to narrow the possibilities of the most likely cause(s) of the overresponse. The rest of this chapter includes a series of 'refutations' of initial hypotheses that can be made regarding the origin of the malfunction in isothermal mode.

4.3.1. Absorbed Dose Measurements in Isothermal Mode

The following measurements were taken with a Varian TrueBeam® linac under reference conditions: the SSD was set to 100 cm, a nominal field size of 10 × 10 cm^2 was used, and the detector was placed (in the horizontal direction) at a depth of 10 cm in a 30 × 30 × 30 cm^3 water tank held at room temperature. Two sets of measurements were taken, exactly one week apart. Many consecutive isothermal runs were completed in each measurement set, totalling 37 isothermal runs with varying beam parameters. For the sake of brevity, a selection of only 5 runs from each set will be presented below. The beam parameters corresponding to each run are listed in Table 4.4, along with the final absorbed dose-to-water measurement obtained for that run.

The results and information contained in Table 4.4 will be discussed further after having analyzed and reduced the number of possible causes for the overresponse.

Table 4.4. Beam parameters used to obtain absorbed dose measurements from Aerrow MK7 in isothermal mode under reference conditions. The beam quality (Q), the repetition rate, the total time of irradiation (Δt) and the measured power drop (ΔP) are given for each run. The uncertainties on the measured power drops ($\mu_{\Delta P}$) correspond to a coverage factor of k = 1. The corresponding absorbed dose-to-water (D_{w}) values are reported per MU delivered at $z_{\max}$.

Measurement Set - Run Number	Q (MV)	Repetition Rate (MU·min^{-1})	Δt (s)	ΔP (μW)	$\mu_{\Delta P}$ (%)	D_{w} (cGy)
I-1	6	600	60	5.998	0.07	1.211
I-2	6	600	60	6.098	0.13	1.231
I-3	6	600	60	6.120	0.15	1.236
I-4	6	600	60	6.056	0.15	1.223
I-5	6	600	180	6.540	0.07	1.321
II-1	10 FFF	1200	60	12.93	0.12	1.225
II-2	10 FFF	1200	60	12.34	0.11	1.169
II-3	10 FFF	2400	120	25.34	0.10	1.200
II-4	10 FFF	2400	120	25.20	0.09	1.193
II-5	10 FFF	2400	120	23.83	0.14	1.128

An estimated uncertainty budget for D_{w} was never finalized for isothermal mode because the magnitude of the overresponse far exceeds any possible uncertainty. In any case, the uncertainty budget estimation would look almost identical to the one shown in Table 4.3 for quasi-adiabatic mode. In fact, the combined standard uncertainty on isothermal measurements is likely to be smaller since the largest source of uncertainty in Table 4.3, namely the specific heat capacity of graphite, is not directly used for calculating the absorbed dose.

4.3.2. Absorbed Dose Analysis Method

The absorbed dose is obtained via a fitting algorithm that works based on linear regression analysis. When the beam is turned on, the radiation will act as a nearly uniform heat source on the core and lead to a drop in the power output of the heating thermistor. The power curve of the core's heating thermistor will therefore consist of a pre-drift, a drop, and a post-drift. To extract the absorbed dose-to-graphite from this curve, the datapoints of the isothermal curve are first separated into two datasets: the first is the 'baseline' and contains all the datapoints in the pre-drift and post-drift, the second is the 'drop' and contains all the datapoints during beam-on time. A linear fit is then applied

to the first dataset, such that the baseline can be characterized by two parameters $\{a_0, a_1\}$ which are the polynomial coefficients for the line of best fit. A line with the same slope but different intercept, characterized by $\{b_0, a_1\}$, is fitted to the datapoints in the drop. The difference in intercepts $b_0 - a_0 < 0$ is representative of the power drop. The dose D_{core} can be calculated from the (absolute) power drop, the measured irradiation time, and the mass of the sensitive volume, as shown in Eq. (2.4). This method is only accurate under the assumption that there is no significant difference in the heat transfer processes occurring during beam-on as compared to beam-off, i.e. the slopes of the baseline and the drop are assumed to be comparable.

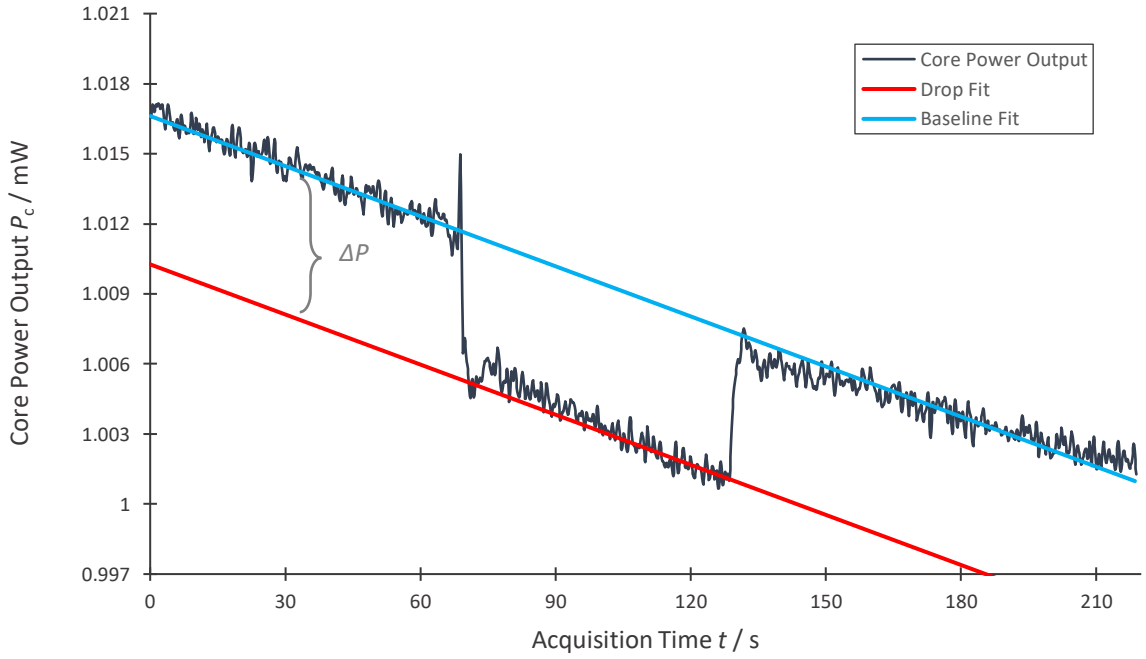

Figure 4.2. Power output of the heating thermistor in the core for run I-1 of the isothermal dataset (see Table 4.4). The blue line is the fit associated with the baseline, described by polynomial coefficients $\{a_0, a_1\}$. The red line, $\{b_0, a_1\}$, represents the fit applied to the drop. The difference in intercepts is approximately equal to the power drop.

To reject the hypothesis that the dose analysis method introduces bias, it is necessary to delve a little deeper into the methodology. The first step is to construct the baseline dataset, which includes the datapoints before beam-on and after beam-off. Then, a fit of first order is applied to the baseline. This is done with the method of least squares regression. Given a dataset:

$$S = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\} \tag{4.1}$$

Then using the method of least squares, one can find a line of best fit $y = a_0 + a_1 x$ where:

$$a_0 = \bar{Y} - a_1 \bar{X} \tag{4.2}$$

$$\bar{X} = \text{avg}(\vec{x}) = \frac{\sum_{i=1}^{n} x_i}{n} \tag{4.3}$$

$$\bar{Y} = \text{avg}(\vec{y}) = \frac{\sum_{i=1}^{n} y_i}{n} \tag{4.4}$$

$$a_1 = \frac{\sum_{i=1}^{n} (x_i - \bar{X})(y_i - \bar{Y})}{\sum_{i=1}^{n} (x_i - \bar{X})^2} \tag{4.5}$$

Once this line has been calculated, a line of best fit is calculated also for the drop — the dataset which includes the points after beam-on and before beam-off. However, the new line of best fit is constrained to have the same slope (a_1) as the first one. In other words, the intercept (b_0) must be optimized such that the average distance between the drop datapoints and the new line are minimized. The problem can be simplified further if the baseline is subtracted from the entire run to produce a 'drift-corrected' signal. In this new dataset, the lines of best fit have slope zero. Thus, finding the new intercept simply reduces to:

$$b_0 = \bar{Y} - b_1 \bar{X} = \bar{Y} - a_1 \bar{X} = \bar{Y} \tag{4.6}$$

Where $\bar{Y}$ now represents the average power of the drop datapoints. The difference between the two intercepts is therefore a good representation of the power drop itself.

$$|\Delta P| \approx |b_0 - a_0| = a_0 - b_0 \tag{4.7}$$

The b_0 parameter was calculated in two ways. The first involved calculating a straight-forward average of the power values of the datapoints in the drop, as in Eq. (4.6). The second method involved creating an array of lines of slope zero with intercepts b_0 ranging from the lowest to the highest power value in the drop. For each line, the coefficient of determination was calculated:

$$R^2 = 1 - \frac{\sum_{i=1}^{n} (y_i - b_0)^2}{\sum_{i=1}^{n} (y_i - \bar{Y})^2} \tag{4.8}$$

The line that maximized R^2 was chosen as the line of best fit for the drop, and its corresponding intercept was used as b_0. Although these two methods of calculating b_0 are arithmetically identical, they employ different computational functions. Thus, they are useful for determining whether the calculation method introduces bias or not. No difference was observed in the final value of the power drop (and hence the absorbed dose) between the two methods for any of the runs.

Finally, a third independent method was developed to extract the absorbed dose from the isothermal curves. Recall from Eq. (2.3) that D_{core} in isothermal mode can be defined exactly as an integral. The third method aims to calculate the dose directly from this definition, eliminating the need to assume that the baseline and drop slopes are similar. First, the baseline is subtracted from the entire run to produce a drift-corrected signal. The area under the signal is then calculated for the time window corresponding to the drop using trapezoidal numerical integration, as shown in Figure 4.3. The total area divided by the mass of the sensitive volume is representative of the absorbed dose-to-graphite. The final absorbed dose-to-water values obtained with the integration method were compared to the corresponding values obtained with the least squares fitting method; the two methods showed a maximum relative difference in D_w of 0.5 % for all isothermal runs.

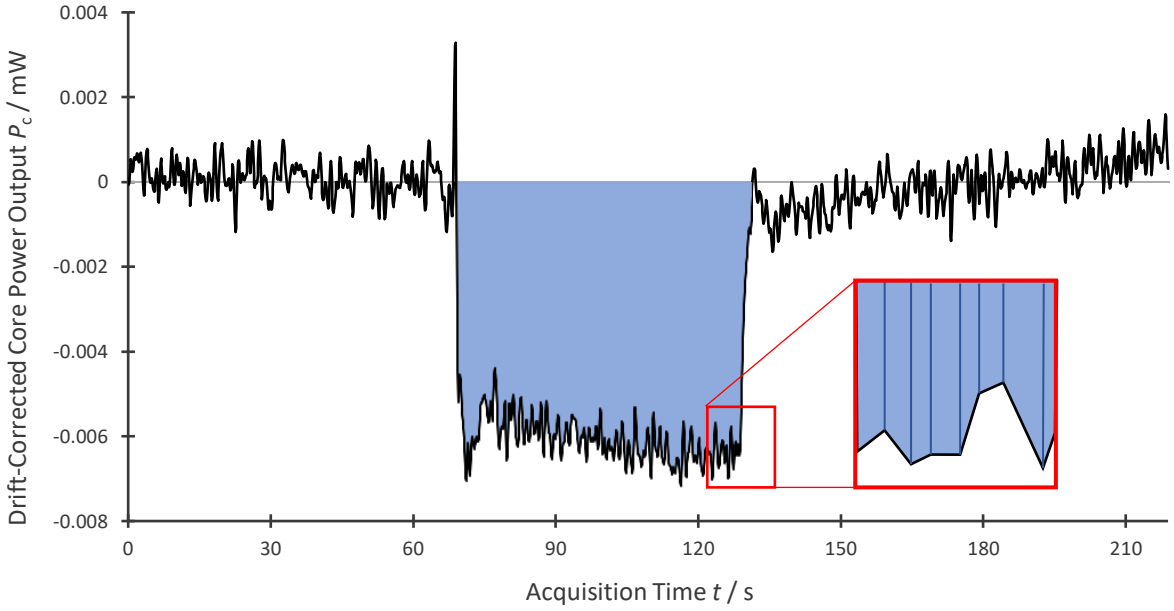

Figure 4.3. A visualization of the trapezoidal numerical integration technique on the drift-corrected signal of run I-1 of the isothermal dataset. A thin trapezoid is constructed for every two data points, with its vertices located at these data points. The width of the trapezoid is the difference in x-axis values (in this case, time) between the two points. The area of each trapezoid is calculated and added together. In the case of these isothermal runs, the trapezoidal integration technique yields an exact answer for the total area under the curve.

One final comment should be made regarding which datapoints are included in the baseline and drop. As can be appreciated from Figure 4.2 and 4.3, experimental isothermal curves contain a considerable amount of noise, and the beam-on and beam-off timepoints are not sharp steps. Furthermore, they are followed by stabilization periods, in which the output of the heating thermistor fluctuates as it readjusts to the sudden introduction/removal of the beam. Some of these fluctuations can be quite large and last for several seconds. Therefore, it is improper to include these datapoints in the analysis. All datapoints that fall in the 'exclusion zones' are discarded and do not form part of the baseline or drop datasets. The exclusion zones are defined as starting 5 seconds before the nominal beam-on or beam-off timepoints and ending 20 seconds after, as shown in Figure 4.4. Exclusion zones were used for all the above-mentioned dose analysis methods.

Reducing the total length of the exclusion zone from 25 to 20 seconds can produce a significant difference (up to 0.9 %) in the final value of D_w, depending on the amount of drift and noise present in the given run. It is clear however that variations in the inclusion or exclusion of certain datapoints cannot account for an overresponse in the range of 17–24 %.

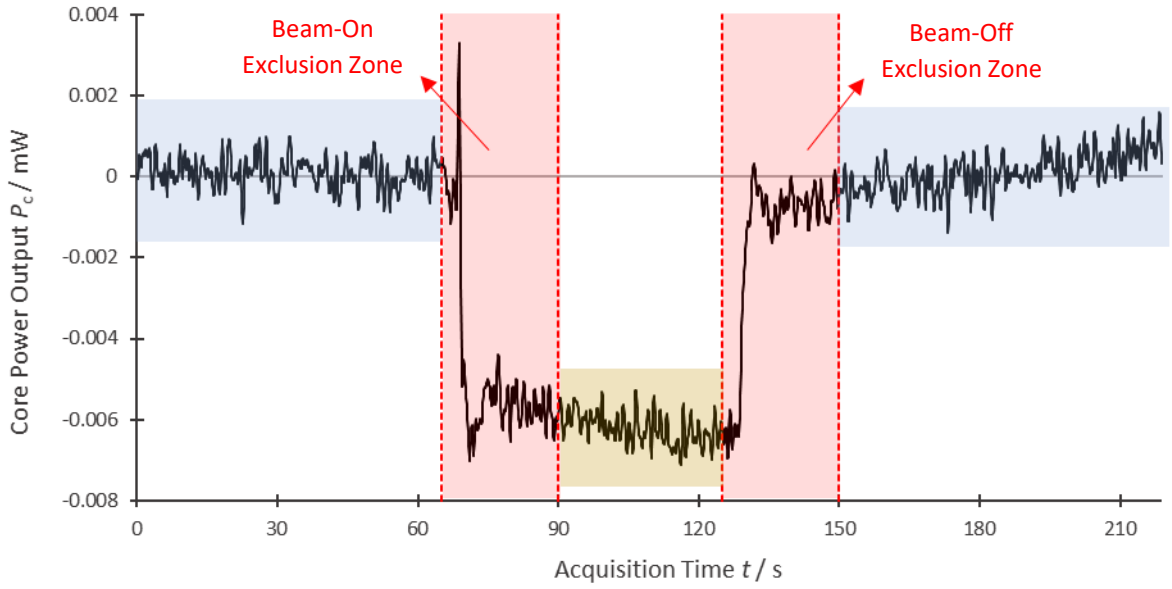

Figure 4.4. Illustration of the exclusion zones (red) for run I-1 of the isothermal dataset. The datapoints in these zones are discarded and unused for absorbed dose extraction. The exclusion zones start 5 seconds before the nominal beam-on or beam-off timepoint (70 s or 130 s, respectively) and end 20 seconds after. The blue-shaded regions represent the baseline. The yellow region in the middle represents the drop, to which either the fitting or the integration algorithm can be applied in order to find the absorbed dose.

The results of this study suggest that the dose analysis method can be discarded as a major source of error with a high degree of confidence.

4.3.3. Correction and Conversion Factors

For isothermal measurements with Aerrow MK7, five multiplicative factors are required in order to adjust and convert the raw dose reading extracted from the power signal into absorbed dose-to-water at depth. These are the mass impurity correction factor k_{imp}, the heat transfer correction

factor k_{ht}, the medium conversion factor $f^{D_{\mathrm{core}}\rightarrow D_{\mathrm{w}}}$, the dose-at-depth conversion factor $PDD(z)$ and the small-field output correction factor $k_{Q_{\mathrm{clin}},Q_{\mathrm{msr}}}^{f_{\mathrm{clin}},f_{\mathrm{msr}}}$. These can be rejected with a high degree of confidence as the possible cause(s) of the overresponse, as will now be demonstrated.

The magnitude of the mass impurity correction factor has already been discussed and can be safely rejected as the possible cause of the overresponse. In fact, the upper bound estimate of 1.003 for k_{imp} (a 0.3 % correction) mentioned in Section 3.3.1 is quite conservative. For this calculation, all the impurities were treated as being the same material: epoxy. This was the material which produced the mass energy absorption coefficient ratio $\left(\frac{\mu_{\mathrm{en}}}{\rho}\right)_{\mathrm{gr}}^{i}$ that was most different from unity from all the materials (glass, copper, epoxy, etc.) present in the impurities. In addition, the masses of the impurities themselves (thermistor encapsulants, thermistor beads, copper wires, etc.) were deliberately overestimated by taking only the upper bounds of the weight measurements.

As shown in the preceding chapter, the small-field output correction factor is at most 1.5 % for Aerrow MK7, and this maximum only occurs for a small field size of S_{clin} = 0.54 cm. The fields used in the isothermal runs have an S_{clin} dimension of 11.0 cm, where the field output correction factor is practically unity. An additional and important validation is the dosimetric performance of Aerrow MK7 in quasi-adiabatic mode. Since it was demonstrated that accurate absorbed dose determination (which necessitates the application of $k_{Q_{\mathrm{clin}},Q_{\mathrm{msr}}}^{f_{\mathrm{clin}},f_{\mathrm{msr}}}$) is possible in quasi-adiabatic mode, the magnitude of the small-field output correction factors is effectively validated down to the smallest experimental field size (S_{clin} = 2.16 cm). Thus, $k_{Q_{\mathrm{clin}},Q_{\mathrm{msr}}}^{f_{\mathrm{clin}},f_{\mathrm{msr}}}$ can be confidently discarded as the source of the overresponse.

The medium conversion factor offers a seemingly bigger possibility of error, since it represents an adjustment to the dose reading of many percent — approximately 12.7 % (revisit Figure 3.4) for

the energy ranges tested experimentally. It could be argued that, if miscalculated, $f^{D_{\mathrm{core}}\rightarrow D_{\mathrm{w}}}$ could account for an overresponse in the range of 17–24 %. There are three factors that make this scenario exceedingly unlikely: First, the medium conversion factor was calculated with Monte Carlo (using the same user code system that was used for calculating the small-field output correction factors) such that the type B uncertainty on the final value was 0.4 %. Since the maximum calculated value of $f^{D_{\mathrm{core}}\rightarrow D_{\mathrm{w}}}$ was 13.2 % for a 10 MV FFF beam, this puts it well outside the range of the overresponse. Secondly, once again, the accuracy of the absorbed dose results in quasi-adiabatic mode serve as a validation that this factor is accurate, since it was applied to those measurements as well. Lastly, the medium conversion factor values of Aerrow MK7 are consistent with those obtained for the previous Aerrow detector, which uses the same formulation of graphite. They are also of similar magnitude to the conversion factors required for PSDL graphite calorimeters, despite the considerable geometric differences with Aerrow MK7.[10,11]

The PDD value at depth is the easiest factor to discard as a possible source of error since this value is not specific to dosimetry with Aerrow MK7. The PDDs (of all beams) of the Varian TrueBeam® linac used for this experiment are routinely checked and verified as part of the QA program at McGill University.[12] Certainly, the uncertainty on the PDD value is negligible compared to the magnitude of the overresponse.

The heat transfer correction factor k_{ht} will be discussed in Chapter 5 as part of a more comprehensive exploration of the heat transfer processes in Aerrow MK7.

4.3.4. Mass of the Sensitive Volume

To evaluate the raw absorbed dose-to-graphite reading from an isothermal signal, one needs to have the power drop ΔP, the measured irradiation time Δt, and the mass of the sensitive volume m_{core}. The first two items and their associated uncertainties have already been examined in this

chapter. However, the latter still presents a possible source of error for the following reason: the sensitive volume is ill-defined for an absorbed dose calorimeter with non-vacuum thermal insulation layers. Typically, the graphite bodies of a graphite absorbed dose calorimeter are separated by evacuated gaps. The sensitive volume for such a calorimeter is easy to define; it is simply the volume occupied by the core. In Aerrow MK7 however, the boundary of the sensitive volume is blurred because of the presence of the first aerogel layer between the jacket and the core. There is no clear gap or discontinuity to delineate the sensitive volume. It can be argued that, even in this case, the sensitive volume is simply the innermost body that is maintained in an isothermal condition. However, there is significant heat transfer between the first aerogel layer and the core (especially in Aerrow MK7, as will be seen in Chapter 5) so the volume of the body that is kept in an isothermal condition is slightly larger than just the core, even if marginally so.

Traditionally, the sensitive volume of the previous Aerrow detector has been taken to be the core together with the entirety of the first aerogel layer. This definition has been adopted primarily based on the results of numerical models and it has led to accurate absorbed dose measurements in isothermal mode. This definition was adopted also for Aerrow MK7 and the value of m_{core} used incorporates the mass of the graphite core (and its inner contents, including thermistors) as well as the mass of the first aerogel layer. The question arises as to whether this definition is indeed applicable to Aerrow MK7. If it is not, then the mass of the sensitive volume would be different, and this would directly affect the absorbed dose-to-graphite calculations. A difference of this type could theoretically account for the 'overresponse' that is observed in the isothermal measurements.

Although it is true that the sensitive volume is an ill-defined quantity, the value of m_{core} can be confidently discarded as the source of the overresponse due to the following three reasons: (a) It cannot be argued that the uncertainty on m_{core} could potentially account for something on the order of magnitude of the overresponse (~ 20 %) because all of the probe's components (including

individual thermistors and copper wires) were repeatedly weighted in laboratory-grade scales, and the measured mass of the sensitive volume (m_{core} = 0.0836 g) possesses a type A uncertainty of 0.25 %. (b) Upon closer examination of the results in Table 4.4, it becomes evident that the overresponse is irregular and imprecise; its magnitude can vary by as much as 10 %. If the problem was the mass of the sensitive volume, one would observe a constant offset from the expected dose value, but this is not the case, even for like isothermal runs. (c) Most importantly, the variable m_{core} comes into the denominator of the absorbed dose equation for isothermal mode. This means that the mass of the sensitive volume would need to be larger, not smaller, for it to correct a dose overestimation. It is illogical for this to be the case. The question of the sensitive volume asks what fraction or percentage of the first aerogel layer should be included in its definition. If the mass of the first aerogel layer is already wholly included in m_{core}, then a significant portion of the jacket would need to be included as well in order to fix a dose overestimation of (on average) approximately 20 %. This is assumed to not be true because the jacket is its own isothermal body and is clearly outside the possible boundaries of the sensitive volume.

4.3.5. Electronics

The circuitry and the electronics behind Aerrow MK7 are quite complex and its discussion could easily form a chapter of its own. A lot of the details have been covered in previous literature.[8,10,11,13] In this thesis, an exposition of the electronics has been omitted because the majority of the work is primarily computational and numerical in nature. Nevertheless, to determine the cause of the overresponse, it is necessary to evaluate the electronics as a possible source of error. In short, the electronics can be separated into two parts: the readout circuitry and the thermistor circuitry.

The readout circuitry contains all the devices that are necessary for recording output as well as for regulating input, e.g. voltmeters and power supplies. The readout circuitry is an unlikely source of error because it is not detector-specific and has been extensively vetted since earlier Aerrow

prototypes. Moreover, malfunctions within the readout electronics can be confidently eliminated as the possible cause of the overresponse, because control measurements were performed to ensure that the readout circuitry had not changed or deteriorated during the time that the isothermal measurements were taken. Specifically, isothermal measurements with the previous Aerrow detector were taken on the same week and in the same irradiation conditions as the isothermal measurements taken with Aerrow MK7. The absorbed dose measurements obtained with Aerrow showed no overresponse; instead, they were accurate to within 1.1 % and comparable runs showed a maximum relative difference of 1.0 %.

The thermistor circuitry is detector-specific — each detector has its own thermistor circuitry because the same thermistors cannot be placed in two probes. The two functions that are completed through the thermistor circuitry are heating and sensing. The sensing aspect was effectively validated by the accuracy of the quasi-adiabatic results. It is not possible for the sensing circuitry to be problematic if accurate absorbed dose determination is possible in quasi-adiabatic mode, which relies solely on the sensing capacity of the thermistors.

The heating aspect can be verified by looking at two different components: the power curves and the temperature fluctuations of the target body. The power output curves verify that (at least from the power supply's point of view) sufficient amperage is being circulated through the thermistors and heating contacts in order for them to heat up. No discontinuities or jumps were observed in any of the power curves. This is significant because it entails a total of 111 power curves, three power curves for each run, corresponding to the three graphite bodies. Moreover, all the power curves were in the expected order of magnitude for each body: ~ 1 mW for the core, ~ 10 mW for the jacket, and ~ 300 mW for the shield. The second check is the temperature fluctuation of the target body. If the sensing function of a given thermistor is reliable and the temperature of the target body does not show large variations, then it can be concluded that the amount of power (and

therefore, heat) being dissipated by a heating thermistor or heating contact is indeed controlled and regulated. Stated another way, this verifies that the objective of heating, which is to maintain the isothermal setpoint, is being successfully met. For the core of Aerrow MK7, the average temperature deviation from the desired setpoint was around ± 10 µK, as can be seen in Figure 4.5 below. This fluctuation is low relative to the typical rise in temperature generated by a 6 MV photon beam, which is on the order of 1 $mK \cdot Gy^{-1}$. Hence, it can be concluded that the heating was stable and the isothermal condition was satisfied for the core. The same can be said for the jacket and the shield. The results suggest that the thermistor circuitry fulfilled its role in both sensing and heating with a high degree of accuracy and fidelity. Therefore, the electronics can be confidently discarded as the possible cause for the overresponse.

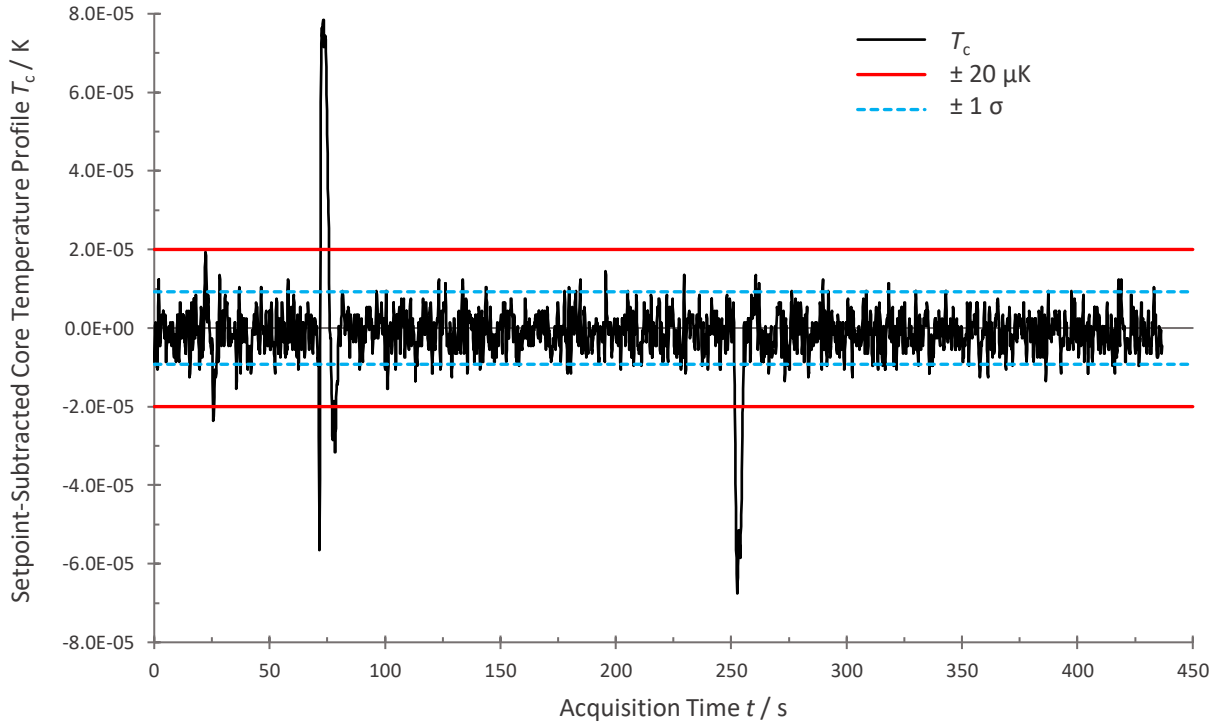

Figure 4.5. Temperature profile of the core's sensing thermistor for run I-5 of the isothermal dataset (see Table 4.4). The setpoint has been shifted down to 0 K for better visualization. Most of the datapoints lie within 20 µK (indicated by the red lines) from the setpoint. The standard deviation for this run, 9.23 µK, is represented by the dashed lines. The large peaks that appear at approximately 73 and 253 seconds are a manifestation of the rapid temperature changes that occur after turning the beam on or off. These should not be considered part of the typical temperature fluctuations.

4.3.6. Construction

During the second set of isothermal measurements, a critical observation was made regarding an error in the construction of the probe: one of the two heating contacts in the jacket was not glued properly and had come loose. Because of this, the jacket temperature was no longer able to be properly regulated in Aerrow MK7. This is a partial explanation as to why the irregularity of the overresponse is much greater in measurement set II. The inability to properly regulate the temperature setpoint of the jacket translates into a heat transfer problem, which can be more accurately studied with numerical models. As such, the evaluation of this effect as the possible cause for the overresponse will be concluded in Chapter 5.

4.3.7. Heat Transfer Processes

The conclusions established in this chapter suggest that the cause of the overresponse is likely internal and inherent to the design of Aerrow MK7. With this in consideration, it is now apt to return to the dose results reported in Table 4.4 and more closely examine their features. There are three critical findings that can be gathered from looking at the power curves, the temperature profiles, and the final absorbed dose values from the isothermal runs:

i. The length of the irradiation period aggravates the overresponse. This effect is less defined in the second measurement set, but the dose values in set II cannot be closely relied on because of the reasons described in the last section.
ii. The overestimation of the absorbed dose is caused by a steep downward drift in the slope of the beam-on portion of the power curve. This drift increases the total magnitude of the drop, leading to a larger absorbed dose value.
iii. The post-drift slopes of most isothermal runs also have pronounced drifts but in a direction inverse to that of the beam-on drifts, as seen in the power curves shown in

this chapter. This is indicative of a heat transfer effect that gradually undergoes stabilization after the beam is turned off.

These three aspects of the data indicate that there is a significant change in the heat transfer processes of the detector as a result of the energy that is introduced by the radiation. This would explain why the effect is present only during and after irradiation and why it is exacerbated as the length of the irradiation period increases. One possibility is the faulty heating in the jacket mentioned before, but this is not the only possibility. For now, it can only be stated that the isothermal malfunction in Aerrow MK7 is most likely related to thermal transport effects. To better assess this hypothesis, it is necessary to turn to numerical simulations. The next chapter introduces heat transfer modelling in the context of isothermal operation with Aerrow MK7.

4.4. References

1. Renaud J, Palmans H, Sarfehnia A, Seuntjens J. Absorbed dose calorimetry. *Phys. Med. Biol.* 2020;65(5):05TR02.

2. Sarfehnia A et al. Direct absorbed dose to water determination based on water calorimetry in scanning proton beam delivery. *Med Phys*. 2010;37:3541–3550.

3. Ross CK, Seuntjens J, Klassen NV, Shortt KR. The NRC sealed water calorimeter: correction factors and performance. Proc. NPL Workshop on Recent Advances in Calorimetric Absorbed Dose Standards National Physical Laboratory. Teddington, UK: 1999;90-102.

4. Seuntjens J, Palmans H. Correction factors and performance of a 4 °C sealed water calorimeter. *Phys Med Biol*. 1999;44:627–646.

5. Almond PR, Biggs PJ, Coursey BM, et al. AAPM's TG-51 protocol for clinical reference dosimetry of high-energy photon and electron beams. *Med Phys*. 1999;26:1847–1870.

6. McEwen MR, DeWerd LA, Ibbott GS, et al. Addendum to the AAPMs TG-51 protocol for clinical reference dosimetry of high-energy photon beams. *Med Phys*. 2014;41:041501.

7. Seuntjens J, Duane S. Photon absorbed dose standards. *Metrologia*. 2009;46:S39–S58.

8. Côté, B. Numerical design and development of a probe-format graphite absorbed dose calorimeter for use in small-field radiotherapy. MSc Thesis, McGill U. Department of Physics. Montreal, QC: 2019.

9. Picard S, Burns DT, Roger P. Determination of the specific heat capacity of a graphite sample using absolute and differential methods. *Metrologia*. 2007;44:294–302.

10. Renaud J, Marchington D, Seuntjens J, Sarfehnia A. Development of a graphite probe calorimeter for absolute clinical dosimetry. *Med Phys*. 2013;40:020701.

11. Renaud J, Sarfehnia A, Bancheri J, Seuntjens J. Absolute dosimetry of a 1.5 T MR-guided accelerator-based high-energy photon beam in water and solid phantoms using Aerrow. *Med Phys.* 2020;47(3):1291-1304.

12. Berryman A, Renaud MA, Seuntjens J. Auto-commissioning for Varian Accelerator for the 6 MV, 6 MV FFF, 10 MV, and 10 MV FFF beam qualities. McGill U. Medical Physics Internal Report.

13. Renaud, J. On the Development of Absorbed Dose Calorimeter Systems for Absolute Clinical Dosimetry. PhD Thesis, McGill U. Department of Physics. Montreal, QC: 2016.

Chapter 5. Heat Transfer Modelling

5.1. Introduction to Finite Element Analysis

Physical phenomena can be described with partial differential equations (PDEs). These PDEs govern the behaviour of functions $u(x, y, z, t)$ that represent quantities of interest (e.g. flow velocity, pressure, temperature) and that evolve according to the physics of the specific problem at hand (e.g. fluid dynamics, structural behaviour, or thermal transport). For most problems, the PDEs cannot be solved analytically, and solutions are approximated using numerical methods.

The finite element method (FEM) is a numerical method that approximates solutions to PDEs by dividing the spatial variables into smaller subparts, a process known as spatial discretization. In FEM, points known as 'nodes' are created over the two- or three-dimensional spatial domain. These nodes are joined with lines to create a mesh-like structure, and the subregions of the 'mesh' are called 'finite elements'. The PDEs are solved internally at each finite element, and the boundary conditions are solved in accordance with the internal solutions of neighboring elements. The net result is a coherent solution of the PDEs over the entire domain. Dividing the space into smaller finite elements is advantageous because (a) local effects can be accurately captured, (b) complex geometries can be easily represented, (c) dissimilar material properties between finite elements can be accounted for, and (d) finite elements converge to a solution relatively fast.[1,2]

FEM is particularly useful for problems that can be described in terms of stressors. The PDEs across the mesh are solved simultaneously and describe how each node responds when external stimuli (stressors) are applied to it or to its neighboring elements. If the PDEs are well behaved, the system will gradually resolve all stresses and the functions will begin to settle, converging to a solution. Thermal transport is often studied using FEM because it can be conceptualized as a problem in which the temperature profile of a given domain responds to the presence of different thermal stressors, i.e. heat sources and heat sinks. This thesis employs COMSOL Multiphysics®,

a FEM-based simulation platform, for modelling heat transfer in Aerrow MK7. Studying or analyzing a phenomenon with FEM is referred to as finite element analysis (FEA).[1-3]

5.2. Modelling Heat Transfer in Aerrow MK7 with COMSOL

The general procedure for designing a study in COMSOL can be separated into four major steps, which shall be referred to here as 'meshing', 'materials', 'physics' and 'simulation' to simplify the terminology. The meshing step involves the creation of the geometry, where the dimensions and shapes of the model objects are defined. In the materials step, the different domains of the geometry are assigned the necessary material properties for conducting the study. For example, some of the properties required for a heat transfer study include the mass density, thermal conductivity, specific heat capacity, and emissivity of a given material. The physics step is where the desired COMSOL module is selected and the physics of the problem are specified. COMSOL includes a list of modules which have presets for different types of problems such as wave optics, acoustics, electrochemistry and heat transfer. The modules have built-in options for initial conditions, boundary conditions, allowed interactions, treatment of variables, and many other parameters which facilitate the specification of the governing physics. Modules may be selected individually or combined with other modules for coupled multi-physics problems. The simulations in this thesis were conducted using COMSOL's Heat Transfer Module.[3] The fourth and final step involves tuning the simulation parameters (e.g. total time of simulation, length of time steps, maximum allowed uncertainties, etc.) before running the simulation.

The following sections provide detailed descriptions, in terms of these four general steps, of the COMSOL model used for analyzing isothermal operation with Aerrow MK7.

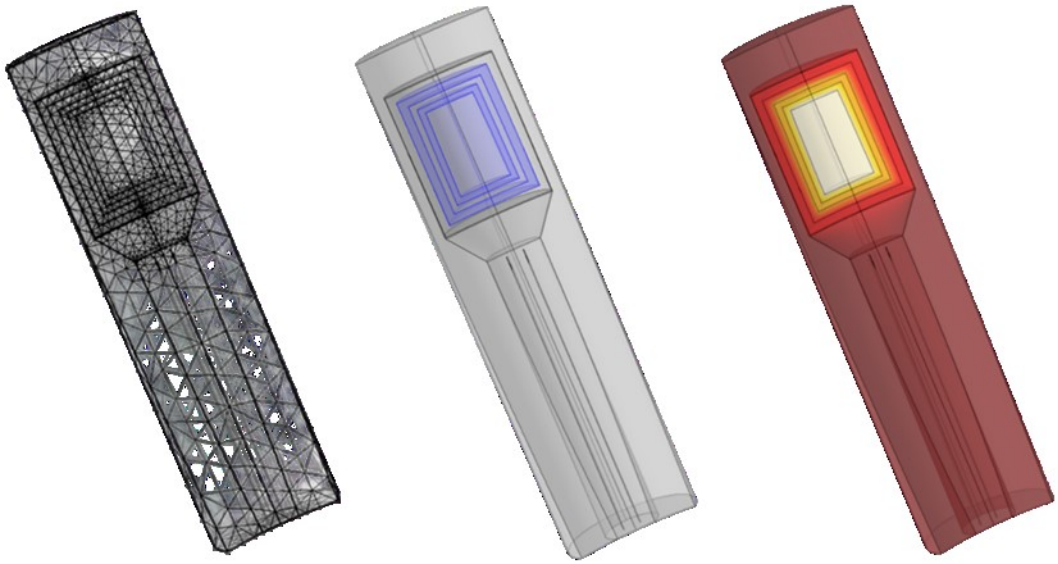

Figure 5.1. These cross-sectional profiles of Aerrow MK7 depict the different stages of designing a heat transfer study in COMSOL. Left: the meshing stage, in which the geometry is defined. Middle: The materials stage, in which the different domains are assigned material properties. The regions highlighted in blue correspond to the insulating layers, which were assigned thermal properties corresponding to those of aerogel. Right: The physics and simulation stages, in which the temperature profile is solved. The profile shown here ascends from dark red (colder) to white (hotter).

5.2.1. Geometry Design & Mesh Settings

In COMSOL, models may be defined spatially as three-dimensional, two-dimensional with axial symmetry, or two-dimensional. To be more specific, a 2D axisymmetric model is a cross-sectional 2D model that can be revolved around an axis to produce a 3D visualization of an axisymmetric object. The model is only solved in the two-dimensional profile, which makes it considerably faster to compute.[4] The design of Aerrow MK7 is cylindrically symmetric except for the thermistors and heating contacts, which are placed asymmetrically throughout the probe. Thus, the Aerrow MK7 COMSOL model was originally constructed three dimensionally to accurately capture the asymmetrical positioning of the embedded electrical components. These initial 3D simulations proved to be very time-consuming to process, taking on average 40 hours, so the model was approximated as an axisymmetric 2D geometry. This change shortened the computation time by an order of magnitude.

Since Aerrow MK7 is not strictly axisymmetric, the reduction to a 2D axisymmetric model meant that a simplification had to be made with regards to the placement of the thermistors and heating contacts embedded in the graphite bodies. To produce axial symmetry, there was only one option available: placing these components directly on the axis of revolution (the longitudinal axis). A visualization of the geometry reduction process is presented in Figure 5.2 below.

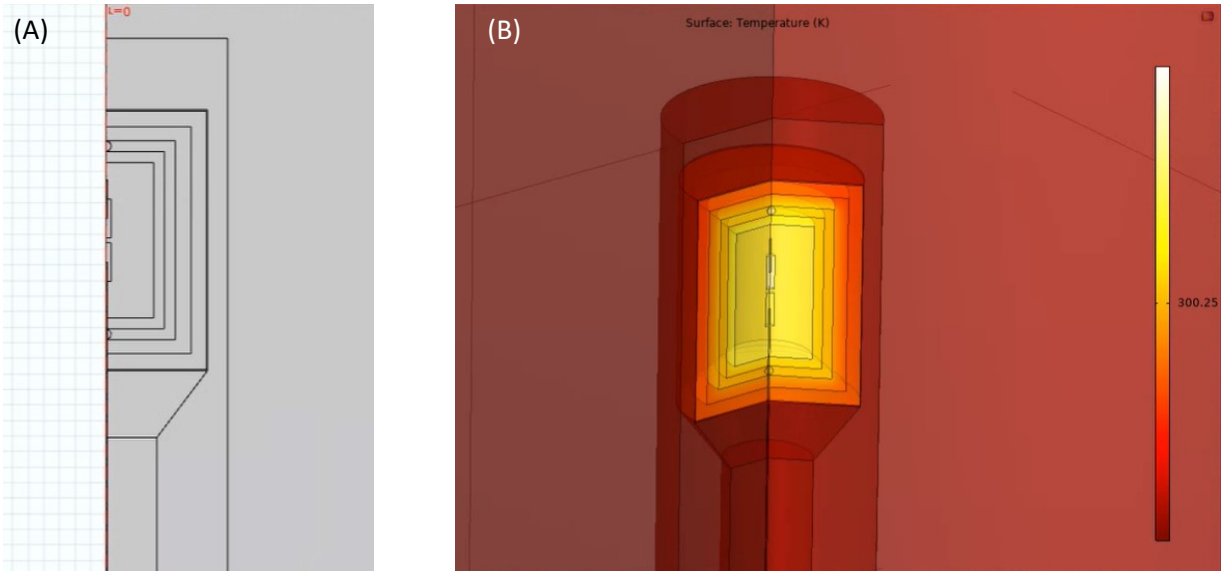

Figure 5.2. Two screenshots of the 2D axisymmetric model of Aerrow MK7 in COMSOL. (A) The 2D model, corresponding to half of a cross-section of Aerrow MK7. This 2D plane is the domain in which solutions are generated. (B) A sample solution revolved around the longitudinal axis to produce a 3D representation of the temperature profile. Note that the thermistors, wires and spherical heating contacts are placed on-axis to conserve axial symmetry.

The dimensions of the different model components were specified according to the real measured dimensions listed in Table 3.2. Two cylindrical thermistors were placed in the core, one for sensing and another for heating, as in the real detector. A long, thin cylinder (r = 0.035 mm) extends from the lower thermistor all the way down the stem; its purpose is to simulate the amount of heat transferred down the copper wires. In a 3D geometry, a copper wire would extend downwards from each thermistor. Because the thermistors had to be placed on the axis of revolution, the wire can only connect to the lower thermistor. The upper thermistor is connected to a shorter, thin cylinder of equal radius that serves as a mock copper wire (see Figure 5.2). Since the real jacket

has no heating thermistors but was rather heated directly through copper wire contacts, two spheres of radius 0.159 mm were placed at the top and bottom of the jacket to act as heating points.

The major disadvantage of moving to a 2D axisymmetric configuration now becomes apparent: the sphere at the bottom of the jacket is in direct contact with the copper wire that connects to the lower thermistor in the core. If used for heating, the lower thermistor can transfer heat down to the bottom spherical contact and vice versa. This means that the heating functions of these two bodies are coupled. Although the amount of heat transferred down the copper wires is negligible, this heat coupling effect must be accounted for in certain thermistor configurations (see Sections 5.3.2.1–5.3.2.2). In the physical probe, the jacket heating contacts are placed to the sides of the jacket and their heating functions are independent from that of the heating thermistor in the core.

No heating contacts or thermistors were placed in the shield. Instead, the shield in this model was heated as a whole and the volume used for temperature sensing in the shield was the shield itself. The reason for this design is that, unlike the jacket, the real shield had two functional heating points and the data (revisit Chapter 4) indicate that the shield was uniformly and effectively heated. An air gap of thickness 0.025 mm was placed between the shield and the acrylic stem to capture the real mechanical tolerances of the assembly step. Air gaps between the aerogel and graphite components were not included due to their narrowness; the inclusion of thin objects like air gaps and copper wires increases computing time drastically since it increases the density of small finite elements in the mesh.

The acrylic stem was given a length of 3.0 cm. The length of the real probe is considerably longer and connects to a cable of approximately equal diameter that extends much further. Including the whole probe and cable would have greatly increased the total volume and resolution of the mesh. Hence, the stem was truncated since the region of interest is only that surrounding the core. The

model concludes with a water cylinder (r = 4.0 cm, h = 8.0 cm) that surrounds the entire probe. As opposed to customizable user-controlled meshes, the mesh for this model is an automatically generated physics-controlled mesh. The resolution setting for the mesh was set to 'extra-fine', which has maximum and minimum element size limits of 1.6 mm and 0.006 mm, respectively. This is the second finest out of nine possible settings for element size; this level of resolution was needed to properly model the thin objects mentioned in the last paragraph.

5.2.2. Material Properties

The physical and thermal properties of aerogel and graphite, described in Section 3.2.3, were assigned to the corresponding bodies in the model. The other components were assigned a combination of user-defined material properties as well as some predefined properties from the COMSOL materials library, which holds data for commonly used media. Table 5.1 below contains a summary of all the relevant material properties for each component in the model.

Table 5.1. Material properties assigned to the different components of the Aerrow MK7 model in COMSOL for a study of conductive heat transfer. The mass density (ρ), specific heat capacity at constant pressure (c_P), and thermal conductivity (k) values are listed for each medium at a temperature of 300 K and under 1 atm of pressure. User-defined values are indicated with a number sign ($^{\#}$).[5-9]

Material	Component(s)	ρ [kg·m^{-3}]	c_P [J·kg^{-1}·K^{-1}]	k [W·m^{-1}·K^{-1}]
Copper	Wires	8940	384	393 $^{\#}$
Polyimide	Thermistors	1350 $^{\#}$	1100 $^{\#}$	0.140 $^{\#}$
Graphite	Core, Jacket, Shield	1860 $^{\#}$	713 $^{\#}$	139 $^{\#}$
Aerogel	Insulating layers	400 $^{\#}$	1050 $^{\#}$	0.043 $^{\#}$
Air	Air gaps, Stem cavity	1.204	1006	0.026
Polymethyl methacrylate (PMMA)	Stem	1190	1420 $^{\#}$	0.190 $^{\#}$
Water	Water environment	1000	4185	0.601

5.2.3. Physics of Heat Transfer

Thermal energy (or 'heat') can be transferred through conductive, convective or radiative processes. Thermal conduction refers to heat transfer realized by the collisions of higher energy

particles/molecules with lower energy ones. Convection is a mechanism of heat transfer caused by the (macroscopic) motion of fluids. Radiative heat transfer takes place through the transport of photons. The COMSOL Heat Transfer Module contains multiple interfaces for modelling different types of heat transfer mechanisms in different types of media. The model discussed in this chapter utilizes the 'Heat Transfer in Solids' interface, which allows for radiative and conductive heat transfer in solid media. Mathematically, heat transfer in solids is governed by:

$$\rho c_P \left(\frac{\partial T}{\partial t} + \mathbf{u}_{\text{trans}} \cdot \nabla T\right) + \nabla \cdot (\mathbf{q} + \mathbf{q}_{\text{r}}) = Q + Q_{\text{ted}} \tag{5.1}$$

Where T is the temperature of the medium, $\frac{\partial T}{\partial t}$ is the partial derivative of the temperature with respect to time, $\mathbf{u}_{\text{trans}}$ is the velocity vector of translational motion, ∇ is the del operator, $\mathbf{q}$ is the heat flux vector by conduction, $\mathbf{q}_{\text{r}}$ is the heat flux vector by radiation, Q is the heat term that accounts for additional heat sources (e.g. exothermic chemical reactions, Joule heating, etc.) and lastly Q_{ted} is the thermoelastic dampening term, which accounts for thermoelastic effects in solids.[10] Eq. (5.1) is derived from the heat balance equation, shown below, which is an alternate way of expressing the first law of thermodynamics:

$$\frac{dE_\Omega}{dt} = Q_{\text{exch}} - P_{\text{str}} \tag{5.2}$$

In the above equation, $\frac{dE_\Omega}{dt}$ represents the time derivative of the total internal energy E_Ω of a domain Ω, Q_{exch} represents the total exchanged heat rate and P_{str} is the stress power. P_{str} is the sum of the pressure-volume work and the viscous stress dissipation of a fluid, as shown in Eq. (5.3). For a fluid with velocity vector $\mathbf{u}$, pressure p, and viscous stress tensor τ:

$$P_{\text{str}} = -\int_\Omega p(\nabla \cdot \mathbf{u})dv + \int_\Omega (\tau : \nabla \mathbf{u})dv \tag{5.3}$$

Here, dv is the infinitesimal volume element of Ω, '$\cdot$' represents the dot product and ':' represents a contraction operator.[10] In the Aerrow MK7 model, the study of heat transfer by conduction is of primary interest and fluid dynamics are not considered. The water and air components in the model

are effectively treated as stationary solids, meaning that their fluid velocity vectors are assigned a value of zero, so $P_{\mathrm{str}} = 0$ for the entire model. These approximations are acceptable for this model because the volumes of fluid components are not big enough to produce large convective motions that could affect the detector's response. In experimental measurements, the real Aerrow MK7 is submerged in a $(30)^3$ cm^3 cubic water phantom, which can set off considerable convective motions at room temperature, especially during irradiation. In the future, the model will be expanded to account for convection by adding the 'Heat Transfer in Fluids' interface to control the water and air domains. More on future work will be discussed in the concluding remarks of Chapter 6.

Thus, the only mechanisms for heat transfer in this model are contained within the total exchanged heat rate variable, Q_{exch}, which accounts for thermal conduction, radiation and other heat sources:

$$Q_{\mathrm{exch}} = -\int_{\partial\Omega} (\mathbf{q} \cdot \mathbf{n}) ds - \int_{\partial\Omega} (\mathbf{q}_{\mathrm{r}} \cdot \mathbf{n}) ds + \int_{\Omega} Q dv \tag{5.4}$$

The radiative and conductive heat flux vectors are dotted with the external normal vector ($\mathbf{n}$) to $\partial\Omega$, the boundary of Ω, to give the total heat flux coming out of Ω. This quantity is integrated over $\partial\Omega$ through the infinitesimal line element ds. With these terms, the heat balance equation becomes:

$$\int_{\Omega} \rho \left(\frac{dE}{dt}\right) dv + \int_{\partial\Omega} (\mathbf{q} \cdot \mathbf{n}) ds + \int_{\partial\Omega} (\mathbf{q}_{\mathrm{r}} \cdot \mathbf{n}) ds = \int_{\Omega} Q dv \tag{5.5}$$

In the 'observer' or 'laboratory' frame of reference (as opposed to the material's frame of reference), Eq. (5.5) can be expressed in localized form as:[10]

$$\rho \left(\frac{dE}{dt}\right) + (\nabla \cdot \mathbf{q}) + (\nabla \cdot \mathbf{q}_{\mathrm{r}}) = Q_{\mathrm{tot}} \tag{5.6}$$

$$\rho \left(\frac{\partial E}{\partial t} + \mathbf{u}_{\mathrm{trans}} \cdot \nabla E\right) + (\nabla \cdot \mathbf{q}) + (\nabla \cdot \mathbf{q}_{\mathrm{r}}) = Q_{\mathrm{tot}} \tag{5.7}$$

Eq. (5.7) is equivalent to Eq. (5.1), and it can be simplified further for the model at hand. Since the detector-phantom assembly remains motionless during experiments and the configuration of Aerrow MK7 is stationary, the relative translational velocity $\mathbf{u}_{\mathrm{trans}}$ of all components is zero.

Furthermore, due to the low absolute temperatures at which Aerrow MK7 is operated, radiative heat transfer is not important and can be safely ignored. This can be more rigorously demonstrated by applying the Stefan-Boltzmann law. From this law, the total power emitted by a body that does not absorb all incident radiation (i.e. a grey body) is given by:

$$P = A\varepsilon\sigma T^4 \tag{5.8}$$

In Eq. (5.8), A is the surface area of the grey body, ε is its emissivity, σ is the Stefan-Boltzmann constant (5.6704×10^{-8} W·m^{-2}·K^{-4}) and T is the thermodynamic temperature of the grey body.[11,12] During isothermal operation, the internal components of Aerrow MK7 are held at approximately 300 K, around 10 K hotter than the outside environment, which is at room temperature. In this case, the maximum thermal power that can be radiatively emitted from a graphite cylinder with a radius of 1.5 mm and a height of 5 mm (approximate dimensions of the core) is ~ 10^{-12} W. This is much less than the total thermal power introduced by the beam, which is on the order of microwatts. In fact, the power dissipated by thermal radiation is immeasurable as it is three orders of magnitude smaller than the average noise seen on the Aerrow MK7 power curves, which exhibit nanowatt fluctuations. For these reasons, conduction is the only important heat transfer mechanism for this model. Hence, the only relevant material properties are those that affect conduction (see Table 5.1) and other properties like the emissivity are not needed. Thus, Eq. (5.1) simply becomes:

$$\rho c_P \left(\frac{\partial T}{\partial t}\right) + \nabla(\mathbf{q}) = Q + Q_{\text{ted}} = Q_{\text{beam}} + Q_{\text{therm}} + Q_{\text{ted}} \tag{5.9}$$

The conductive heat flux vector can be replaced by the product of the thermal conductivity and the temperature gradient, $\mathbf{q} = -k\nabla T$, to recast Eq. (5.9) into the more familiar and well-known heat conduction equation. The model contains two user-defined heat sources, so Q can be separated into two parts. The first, Q_{beam}, simulates the action of the beam and dissipates power such that every object receives the same thermal energy per unit mass. For each object, the power dissipated by Q_{beam} per unit volume is defined as $\frac{dP}{dV} = \rho_{\text{obj}}\dot{D}\Delta t$. To exemplify, for a 10 Gy·min^{-1} dose rate and an irradiation period of 60 seconds, the entire geometry is made to ‘absorb’ 10 Gy uniformly.

Although the absorbed dose varies spatially for real beam sources, this approximation greatly simplifies the analysis; the core power curves are expected to exhibit clear drops of exactly 10 Gy. The second user-defined heat source, Q_{therm}, represents the adaptive Joule heating from the thermistors. The way in which this heat source dissipates power will be explained in Section 5.3.1.

The discussion of the physics can be finalized by specifying the initial and boundary conditions. The model boundaries were held to an adiabatic condition, $\mathbf{q} \cdot \mathbf{n} = 0$. This condition allows the temperature at the boundaries to change freely, but heat cannot escape beyond the water phantom. Since the model is used to simulate isothermal operation, the three graphite bodies must be held at constant temperatures. Like the real setpoints, the model setpoints were chosen to have relative differences of 0.1 K: 300.36 K for the core, 300.26 K for the jacket, and 300.16 K for the shield. To facilitate a more rapid stabilization of the temperature profile and hence to reduce computation time, these setpoints were defined as the initial temperatures for these bodies. The aerogel layers were set at intermediary initial values, i.e. they had a difference of ± 0.05 K with the adjacent graphite bodies. All other components were set initially to room temperature (293.15 K).

5.2.4. Temporal Parameters & Tolerance Settings

The simulations were given total durations of either 350 or 590 seconds. These times were chosen to fit a 50 s equilibration period (determined to be the time required for the temperature profile to stabilize such that the maximum fluctuation seen on the power curves was ± 1.5 μW) followed by either two or four beams of 60 s durations, with each beam having a 60 s pre-drift and post-drift. The time steps taken by the solver were ‘free’. This means that the time elapsed between solutions is determined by the physics of the problem and can vary depending on the status of convergence, the geometrical complexity, and the number of PDEs to be solved. Usually, initial steps are as small as half a nanosecond, and the length of the time steps increases gradually as the number of unsolved quantities decreases. However, a maximum time step of 0.1 seconds was imposed.

The solver utilized the backwards differentiation formula (BDF) method of second order (s = 2) and the backward Euler method (s = 1) to solve the equations.[13,14] The relative tolerance, which controls the relative error of a single step, was set to 10^{-6}. The absolute tolerance produced satisfactory results when set to 10^{-8} and smaller. The tight tolerance margins are chiefly due to the small magnitude of the power drops observed for the core, which are on the order of microwatts and cannot be properly discerned with lower tolerance settings.

5.3. Modelling Isothermal Operation with Aerrow MK7 in COMSOL

To the best of the author's knowledge, no numerical models have been developed to simulate isothermal operation in a graphite (or any other solid body) absorbed dose calorimeter. Therefore, isothermal modelling represents one of the most important original contributions from this thesis. The following sections discuss the findings from a selection of studies performed with this model.

5.3.1. PID Controllers

Up to now, the heating function has been described as 'adaptive' or 'active-controlled' without covering the exact method in which this real-time adaptiveness is achieved. As adaptive heating is an integral part of isothermal operation, it is necessary to understand the underlying mechanism that controls the heating power and the way it reacts. This mechanism is the proportional-integral-derivative (PID) controller. PID controllers continuously calculate an error value, $e(t)$, between a process variable and its desired setpoint or target value. In the case of Aerrow MK7, the variable of interest is the measured temperature (T) of the core/jacket/shield, and the error is calculated through $e(t) = |T - T_{\text{set}}|$, where T_{set} is the temperature setpoint for that body. In order to match the setpoint, the PID determines how to adjust the heating power by looking at the 'control function', formulated below, which contains three terms: a term proportional to the error, a term with the integral of the error, and a term with the derivative of the error.[15]

$$u(t) = K_{\mathrm{p}} e(t) + K_{\mathrm{i}} \int_0^t e(t')dt' + K_{\mathrm{d}} \frac{de(t)}{dt} \tag{5.10}$$

The purpose of these terms is to introduce different components of information on how the error is behaving and where the error is likely to head in the immediate future, so that the PID can adjust the power output accordingly. The different terms are given coefficients, $K_{\mathrm{p,i,d}}$, which act as weighting factors for each means of assessing the error. The goal of the PID controller, which is to settle to the setpoint quickly and maintain the error as low as possible thereafter, is attained by tweaking these coefficients, a procedure known as 'PID tuning'.[†] PI controllers ($K_{\mathrm{d}} = 0$) have typically been sufficient for fast and accurate adaptive heating in Aerrow MK7.

It should be noted that there are two types of PID controllers being discussed here: one is 'real' and controls the physical temperature of the graphite bodies through the electronic circuitry, the other is 'virtual' and controls the isothermal COMSOL model through computer-defined functions. The working principle behind them — the control function in Eq. (5.10) — is the same. The rest of this section contains a series of graphs which illustrate the process of applying and tuning the three virtual PID systems in the core, jacket and shield.

PID tuning is a lengthy process for two major reasons. The first is that it involves the manual tweaking of 9 parameters (3 PID coefficients for 3 graphite bodies) primarily through trial-and-error. It would not be possible to present the results from all PID tuning studies here, so the general procedure is illustrated in Figure 5.3 below, which represents the tuning of $K_{\mathrm{d,C}}$ for the core PID after tuning the $K_{\mathrm{p,C}}$ and $K_{\mathrm{i,C}}$ coefficients. This preliminary study was completed when first implementing the PID in the core and is not conditioned by a maximum limit on the power output, hence the fast settling times of less than a second, even when starting from room temperature.

[†] A simple yet excellent illustration of PID tuning and the effect of the different terms on the process variable as a function of time can be accessed here, borrowed from the Wikipedia page on PID controllers.

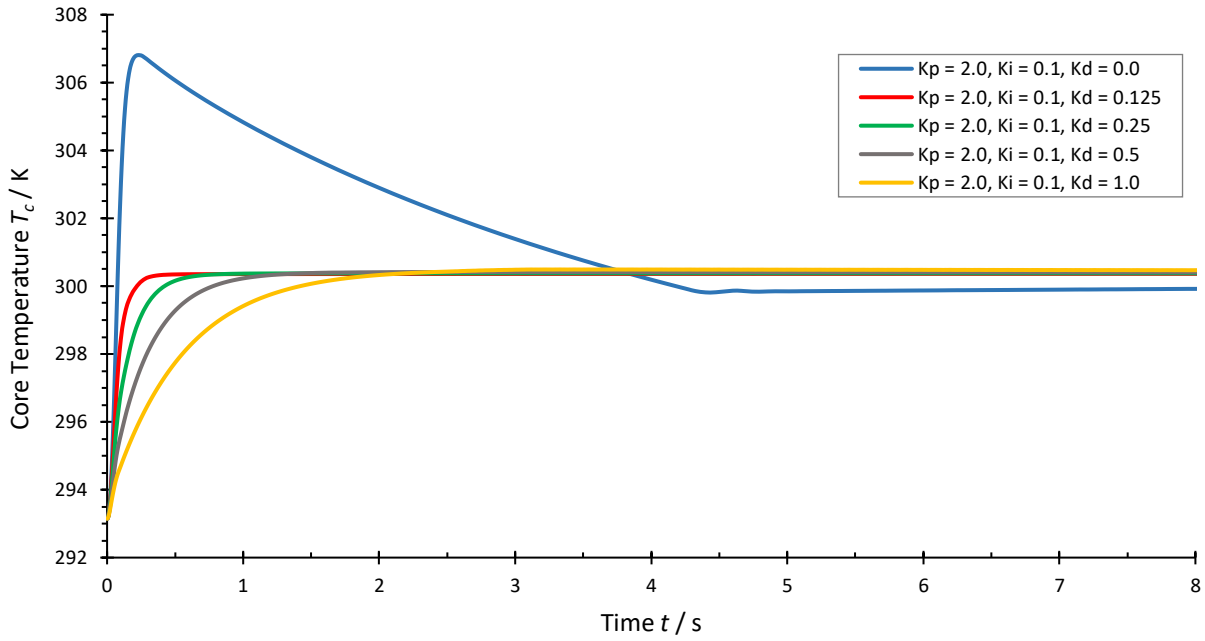

Figure 5.3. Temperature of the core's sensing thermistor in the Aerrow MK7 isothermal COMSOL model as a function of time for different values of the derivative coefficient, $K_{d,C}$, in the core's control function. The temperature setpoint was 300.36 K. In this PID tuning study, no maximum limit was set on the possible power output of the PID controller.

In reality, the PID controllers have an upper limit on the amount of power that can be dissipated to each body (core: 50 mW, jacket: 100 mW, shield: 500 mW) and these drive the settling time to be much longer. The PID controllers in COMSOL were adjusted to reflect these realistic constraints. To compensate, the initial temperatures of the graphite bodies were set to their respective setpoints, as mentioned before, in order to reduce computation time.

One additional limitation had to be considered in the tuning of PIDs, not related to any physical constraint but out of practicality: the computation time itself seemed to be (almost) directly proportional to the magnitude of the chosen coefficients, as shown in Figure 5.4 below. Due to the black-box nature of COMSOL, the reason for this effect remains to be understood. Although the time-to-stabilization can (minorly) affect the total computation time, any non-zero PID coefficient, $K_{p,i,d}$, should theoretically give a computation time comparable to that of any other non-zero value.

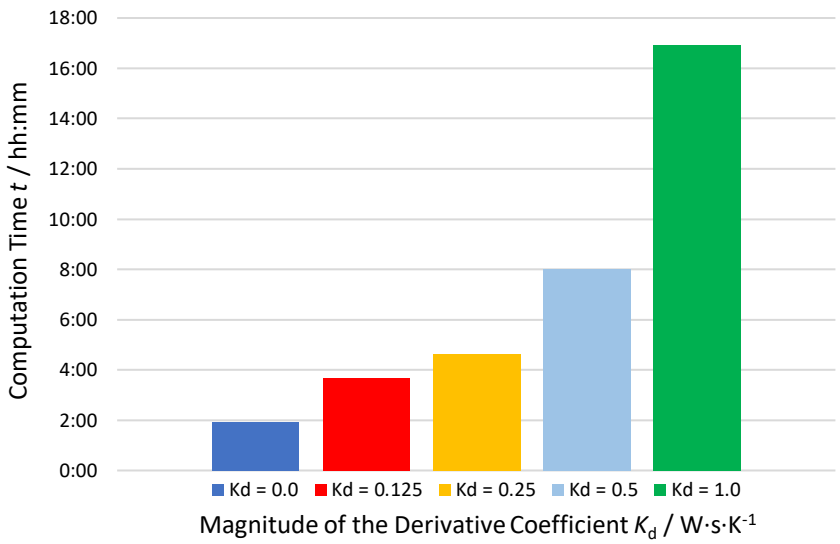

Figure 5.4. Total computation time for a PID tuning study using the Aerrow MK7 isothermal COMSOL model as a function of the magnitude of the derivative coefficient, $K_{d,C}$, of the PID controller in the core. All other parameters were kept identical for all five simulations.

Because of this limitation, the value of $K_{d,C}$ was not set to 0.125 W·s·K^{-1} as suggested by the study in Figure 5.3, but was rather kept to zero because of the already-large value of $K_{p,C}$. The introduction of the jacket and shield PIDs, which also help the core temperature to stabilize faster, was sufficient to compensate for the slower settling time and to correct the final setpoint offset seen in Figure 5.3 for $K_{d,C} = 0$.

As hinted by this last statement, the second reason why PID tuning is time-consuming is because the action of one PID can strongly affect the action of another, just as with the real PIDs. One of the main challenges in building this isothermal model was to figure out the interdependencies between all three PIDs and choose the coefficients that would produce acceptable and realistic results in each graphite body.

The PIDs were implemented from the inside out, starting with the core. In hindsight, introducing the PIDs in the bigger bodies first might have proved easier. However, introducing the PIDs in this

order shed some insight on the overresponse problem, which was the motivation for developing the isothermal model in the first place. The following two figures show the power curves for the core obtained with (a) a PID in the core only, and (b) PIDs in both the core and the jacket.

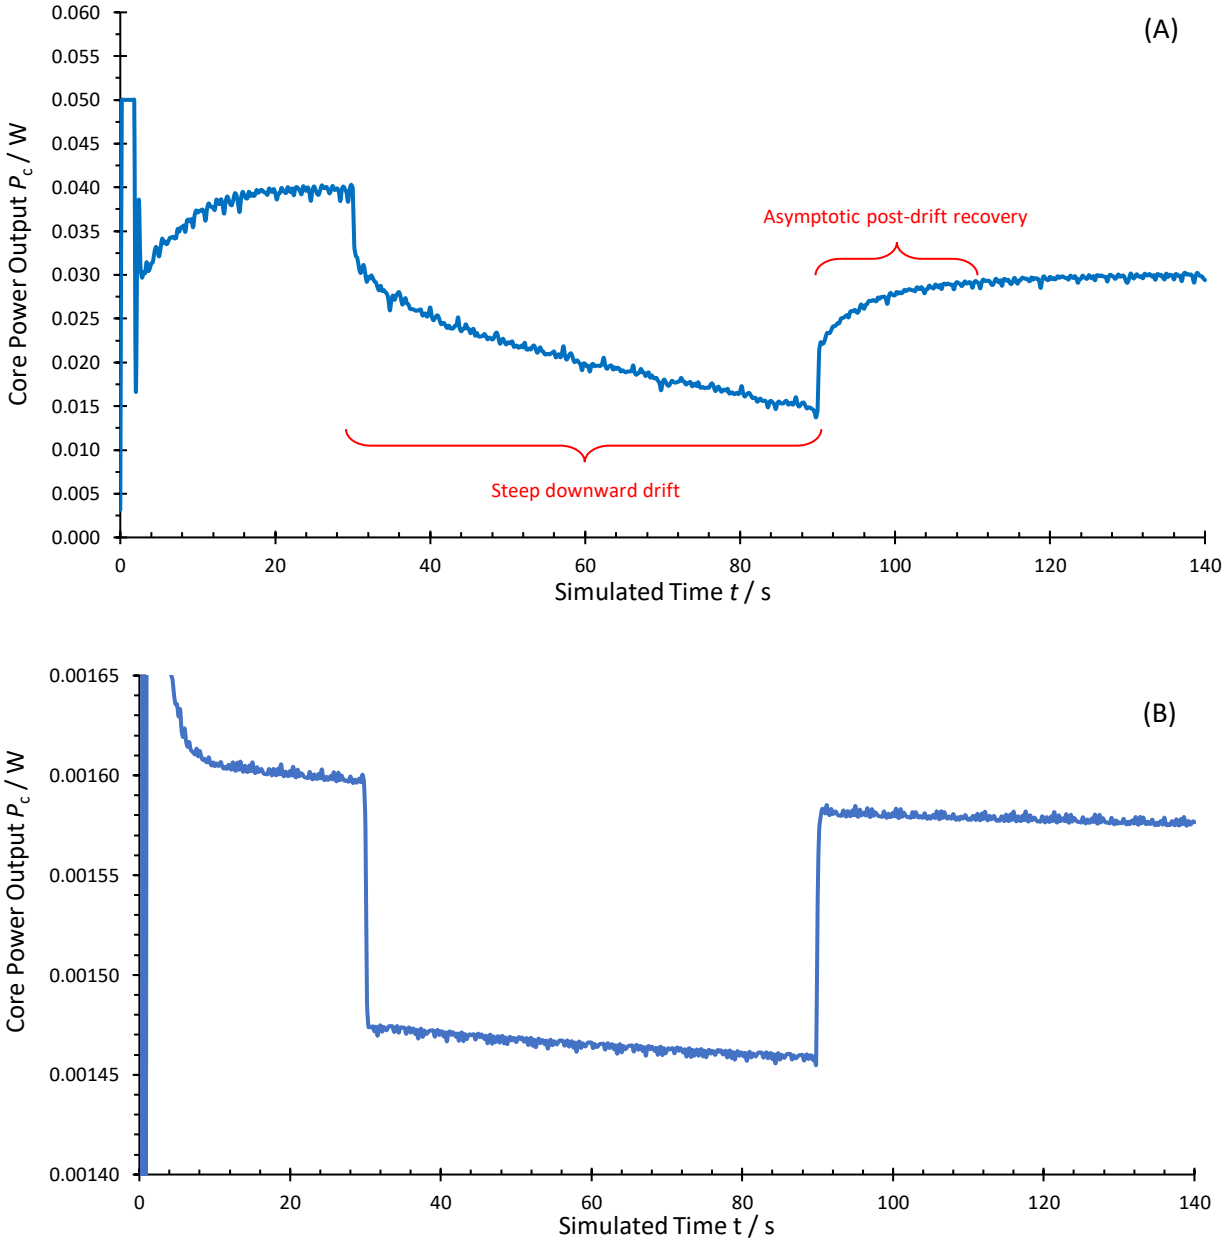

Figure 5.5. Power output for the core's heating thermistor in the Aerrow MK7 isothermal COMSOL model for (A) a 1000 Gy·min^{-1} beam heat source of 60 s with only one PID controller; (B) a 100 Gy·min^{-1} beam heat source of 60 s with PID controllers in both the jacket and the core. The temperature setpoints were 300.36 and 300.26 K for the core and the jacket, respectively. The large fluctuations seen in both plots during the first several seconds reflect the initial adjustment of the PIDs. Note the difference in the scale of the resting power output between (A) and (B).

As is apparent from these two figures, the power output values for the core approach realistic experimental values much more closely when a (functional) PID is placed in the jacket. When there are no jacket and shield PIDs, the heating thermistor in the core is forced to output a lot more power. Additionally, the temperature variations from the outside environment cause the power curve to show large fluctuations. Due to these large fluctuations, the beam source must be on the order of 1000 Gy in order to discern a power drop in the core. Once the jacket PID is placed, a beam of 100 Gy is easily discernible, as seen in Figure 5.5B. Finally, with the jacket and shield PIDs, beam heat sources on the order of clinical irradiations (1-10 Gy) are discernible, producing power curves like those seen in Chapter 4, except with no large drifts or overresponses.

Two conclusions can therefore be drawn from this study. The first is that no overresponse is seen in the COMSOL model when all PIDs are working and properly tuned. Secondly, although the power curves are not strictly comparable due to the large difference in total dose delivered, there is an important qualitative observation to be made from the power curve in Figure 5.5A. Namely, two distinguishing features (indicated in red on the graph) from the real isothermal curves were able to be replicated when running the isothermal model with no jacket or shield PID: the large downward drifts that lead to a dose overestimation, and the asymptotic-like recovery of the post-drift. This suggests that the overresponse is potentially linked to a problem with one of these PIDs.

For the sake of completeness, the final PID coefficients that produced satisfactory power curves with tolerable computation times were: $K_{p,C} = 2.0\ W \cdot K^{-1}$, $K_{i,C} = 0.1\ W \cdot s^{-1} \cdot K^{-1}$, $K_{p,J} = 1.5\ W \cdot K^{-1}$, $K_{i,J} = 1.0\ W \cdot s^{-1} \cdot K^{-1}$, $K_{p,S} = 1.0\ W \cdot K^{-1}$, $K_{i,S} = 0.5\ W \cdot s^{-1} \cdot K^{-1}$, and $K_{d,all} = 0$.

5.3.2. Thermistor Placement Tests

In order to further the discussion on the relation between the overresponse and PID functionality, it is necessary to first explore a (seemingly) unrelated set of experiments: the thermistor placement

tests. There are two possible thermistor-related sources of error that can be addressed within the context of heat transfer modelling: the placement of the thermistors within the core, and the amount of heat transferred down the copper wires. As will be demonstrated in the remainder of this chapter, the thermistor placement tests proved critical for establishing the following three conclusions:

i. Neither the placement of the thermistors nor the amount of heat transferred down the copper wires significantly affects the measured absorbed dose.
ii. The heat transfer correction factor, k_{ht}, left pending from the discussion in Chapter 4, is substantial for isothermal mode — it represents a correction on the order of 6 %.
iii. The functionality of the PID in the jacket is directly tied to the overresponse.

5.3.2.1. Effect of Heat Leakage from the Copper Wires on the Absorbed Dose

The first concern is that of heat transfer in the wires. One can postulate that if enough heat is dissipated away from the core through the copper wires, the rate of heat loss from the core would be accelerated, and more power would be needed to keep the core isothermal. This would result in an apparent overestimation of the absorbed dose. The following considerations make this a very improbable cause for the overresponse seen in Aerrow MK7.

The first is that this effect has already been studied in the literature; Choi et al. (2019) describes a similar COMSOL model for a three-body graphite calorimeter in a 2D axisymmetric geometry.[16] The conductive heat leakage through the copper wire (almost equal in dimensions to the wire used in the Aerrow MK7 model) was reported to be 0.6–0.7 % of the power input to the core. This range is consistent with results obtained from physical experiments conducted by Radu et al. (2010).[17]

Moreover, isothermal simulation runs were performed with a simplified version of the Aerrow MK7 model with no copper wire. After a 50 s stabilization period, the detector was subjected to

two consecutive beam heat sources, delivered 60 s apart from each other, that uniformly dissipated 10 Gy in 60 s. The runs were repeated with the copper wire present. The maximum relative difference in the final value of absorbed dose between any of the four beams was 1.1 %. Therefore, it can be concluded with a high degree of confidence that the conductive heat leakage from the copper wires cannot account for an overresponse in the range of 17–24 %.

5.3.2.2. Effect of Thermistor Placement on the Absorbed Dose

Because of its millimetric dimensions, the core of Aerrow MK7 can only contain one heating and one sensing thermistor, unlike previous Aerrow prototypes. For this reason, the placement of the thermistors was initially suspected to have some relation to the overresponse. With only one heating thermistor, it was thought, the temperature gradients within the core might be large enough that the sensing thermistor could be in a 'cold spot'. Then, the perceived temperature of the core would have been less than if the core had been uniformly heated or if the thermistors had been closer together, and the amount of power required to keep the core isothermal would be overestimated. The thermistor placement tests were designed to address this hypothesis.

These tests consisted of exposing the model to four consecutive beam heat sources for eight different model geometries. Once again, the beams dissipated 10 Gy in 60 s uniformly and were separated in time by 60 s, with the first beam activating 50 s into the simulation. A sequence with a high number of beams was chosen to evaluate whether the temperature gradients in the core changed over time, and how this affected the dose. The eight different geometries reflected different spatial distributions of the thermistors within the core.

Since the axial symmetry of the model had to be conserved, the only parameter that could effectively be changed was the longitudinal separation between the two thermistors. Nevertheless, the horizontal displacement between the thermistors in the core can never be more than 2.73 mm,

so a separation in the longitudinal axis only is still a good approximation. Additionally, the heating and sensing thermistors can be inversed for each separation setting. The geometries consist of three different separation distances: 2.35, 1.25, and 0.15 mm. Placing the thermistors directly against the core boundaries or directly against one another proved to be problematic for meshing. For this reason, the longitudinal separation between the thermistors was never fully maximized (2.50 mm) or minimized (0 mm). The tests also include one scenario with the heating and sensing functions of the jacket inversed, and one scenario with a non-working jacket PID, i.e. neither of the spherical contacts was given a heating function. A summary of the geometry settings for each test and the corresponding absorbed dose-to-graphite results obtained for that test are tabulated below.

Table 5.2. Results from the thermistor placement tests in the Aerrow MK7 isothermal COMSOL model. For each geometry, the separation between thermistors ($\Delta d_{\mathrm{therm}}$), the heating and sensing configurations in the jacket and core, and the absorbed dose-to-graphite for each beam (D_{core}) are given. For the positioning of the heating and sensing points given in the PID columns, 'H' denotes a heating thermistor/contact and 'S' denotes a sensing thermistor/contact. The uncertainties on the dose values were estimated by changing the power curve analysis method and by varying the total length of the exclusion zones (revisit Chapter 4).

Geometry	$\Delta d_{\mathrm{therm}}$ [mm]	Core PID	Jacket PID	D_{core} [Gy]			
				Beam 1	Beam 2	Beam 3	Beam 4
1	1.25	H: Lower S: Upper	H: Upper S: Lower	9.44 ± 0.09	9.23 ± 0.07	9.20 ± 0.07	9.19 ± 0.08
2	1.25	H: Upper S: Lower	”	9.38 ± 0.08	9.18 ± 0.08	9.14 ± 0.07	9.14 ± 0.08
3	0.15	H: Lower S: Upper	”	9.37 ± 0.09	9.16 ± 0.08	9.13 ± 0.07	9.13 ± 0.09
4	0.15	H: Upper S: Lower	”	9.34 ± 0.09	9.14 ± 0.08	9.11 ± 0.08	9.09 ± 0.07
5	2.35	H: Lower S: Upper	”	9.65 ± 0.07	9.43 ± 0.06	9.39 ± 0.06	9.39 ± 0.06
6	2.35	H: Upper S: Lower	”	9.41 ± 0.08	9.20 ± 0.06	9.17 ± 0.06	9.16 ± 0.06
7	1.25	H: Lower S: Upper	H: Lower S: Upper	8.78 ± 0.07	8.72 ± 0.06	8.7 ± 0.1	8.33 ± 0.06
8	1.25	H: Lower S: Upper	H: NONE S: Lower	13.4 ± 0.4	14.2 ± 0.5	14.2 ± 0.3	14.2 ± 0.3

A striking observation from the results in Table 5.2 is that the dose values are consistently underestimated. Inadvertently, these discrepancies shed valuable insight on the heat transfer correction factor, k_{ht}, in isothermal mode. Ultimately, the underresponse is related to the effect of the insulating layers on the heat transfer processes of the core. The discussion on the source of this underresponse will be resumed in the next section.

Ignoring the underresponse momentarily, one can notice some other features in the data of Table 5.2. The first is that the absorbed dose values show a slight trend downwards as the simulation progresses, i.e. the response from the later beams is underestimated by a slightly larger amount. Contrary to the initial hypothesis, the analysis demonstrated that this was not an indication of a change in the temperature gradients of the core over time. Although the absolute difference between the coldest and hottest points within the core could be as large as 0.001 K, the temperature profile as a function of time for any fixed point in the core was static over the entire simulation to within ± 0.00003 K. In other words, the gradual decrease in the absorbed dose values is not attributable to thermal gradients within the core.

Rather, the decrease is a manifestation of the small total volume of the model and the adiabatic boundary condition. As heat is not allowed to escape the model, the energy from all beam heat sources is permanently trapped. The net result is a (small) temperature increase for all domains over time, meaning that the power required to keep the graphite bodies at their desired setpoints is gradually decreased. The effect on the total power curve is illustrated in Figure 5.6 below. This effect can be ameliorated by changing to an isothermal boundary condition or by increasing the total volume of the water cylinder to resemble the larger water phantom used in physical experiments (at the cost of computation time).

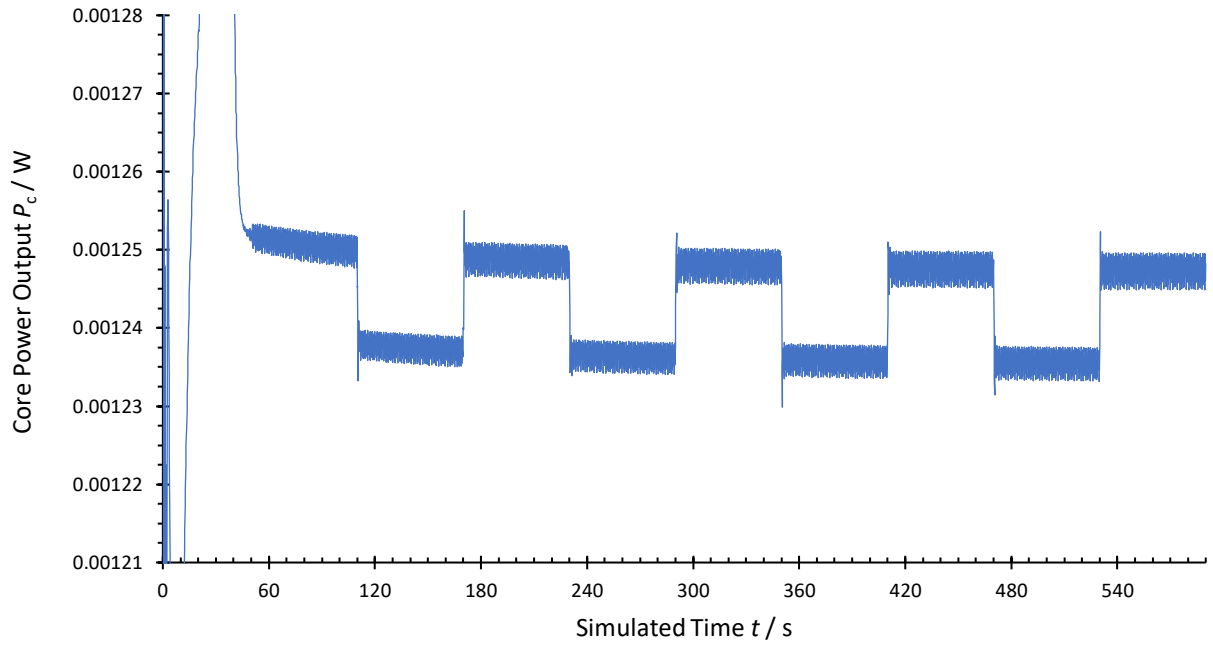

Figure 5.6. Power output curve for the core in the Aerrow MK7 isothermal COMSOL model for geometry 1 of the thermistor placement tests. The large fluctuations seen in the first 50 seconds correspond to the initial stabilization of the temperature profile. The nearly constant variations seen after the stabilization period are primarily dictated by the tolerance constraints of the model on the PIDs. Note the slight downward trend in the resting power output.

Another major observation one can extract form the data in Table 5.2 is that the placement of the thermistors does not seem to play a large role in the determination of absorbed dose. The results from the first seven geometries are directly comparable. For any given beam, the dose results from five of the first seven geometries agree to within combined uncertainties. There is one minor and one major deviation in these results; geometry 5, which consistently produced larger dose values than the other geometries by a small but significant amount, and geometry 7, which produced smaller dose values by as much as 1.06 Gy. As will be explained in the next section, these two exceptions are opposite manifestations of the same underlying physical effect, and they are more closely related to the source of the underresponse than to the placement of the thermistors. In conclusion, the data indicate that the placement of the thermistors does not have a significant effect on the absorbed dose provided that (a) the heating thermistor of the core and the heating contact

in the jacket are not too close together, as in geometry 7, and (b) the heating and sensing thermistors in the core are relatively close ($\Delta d_{\mathrm{therm}} \leq 2.35$ mm) to one another. Furthermore, the placement of the thermistors can be confidently discarded as the source of the overresponse.

A final and important piece of information provided by Table 5.2 is the large overresponse seen in geometry 8. Strictly speaking, this is not a thermistor placement test but rather a PID functionality test in which the jacket PID was disactivated. Therefore, the dose values obtained in this test should not be evaluated directly against the values obtained in other tests. However, the results from geometry 8 prompted further investigation on how the overresponse seen in the real isothermal runs might relate to PID functionality. This discussion will be concluded in Section 5.3.4.

5.3.3. Insulating Material Tests

The unexpected underresponse observed in the thermistor placement tests was initially thought to be caused by a lack of numerical accuracy. To remediate this, nearly all the settings in the model (time step size, mesh settings, tolerance settings, treatment of decimal points, etc.) were refined, and their effects on the dose were investigated. In short, none of these produced a significant change in the final absorbed dose values extracted from the power curves. Furthermore, it was conclusively verified that the model was indeed dissipating 10 Gy uniformly across all domains to within 0.2 %.[‡] Since the numerical accuracy of the model could not account for the difference, the underresponse had to be a manifestation of some real physical effect.

Eventually, it was found that this effect was caused by the presence of the aerogel layers; the only model change that returned the D_{core} values to 10 Gy was the replacement of aerogel with vacuum in the thermal insulation layers. With all other simulation parameters equalized, the reintroduction

[‡] The exact power output of a heat source can be retrieved directly from COMSOL. The uncertainty value listed here corresponds to the standard deviation on the power outputs obtained from multiple consecutive beam heat sources in a single simulation, so it is a type A uncertainty with a coverage factor of $k = 1$.

of aerogel caused the dose values to be underestimated again. The simulations were repeated with air as the insulation material. Air produced a small but noticeable underestimation in the absorbed dose values. Table 5.3 below contains the dose results from the insulating material tests, in which the model was subjected to two 60 s-long 10 Gy beam heat sources after a 50 s stabilization period. These tests were done with the same set-up as geometry 1 of the thermistor placement tests, i.e. $\Delta d_{\mathrm{therm}}$ = 1.25 mm and the heating was done by the lower thermistor in the core, and by the upper contact in the jacket.

Table 5.3. Results from the insulating material tests in the Aerrow MK7 isothermal COMSOL model. For each insulating material, the user-assigned mass density (ρ), specific heat capacity at constant pressure (c_P), and thermal conductivity (k) values are given for T = 300 K and P = 1 atm. The absorbed dose-to-graphite (D_{core}) is also given for each beam. As before, the uncertainties on the dose values were estimated by changing the power curve analysis method and by varying the total length of the exclusion zones.

Insulating Material	ρ [kg·m^{-3}]	c_P [J·kg^{-1}·K^{-1}]	k [W·m^{-1}·K^{-1}]	D_{core} [Gy]	
				Beam 1	Beam 2
Aerogel	400	1050	0.043	9.5 ± 0.1	9.4 ± 0.1
Air	1.204	1006	0.026	10.00 ± 0.04	9.96 ± 0.04
Vacuum	0.001	1006	0.050	10.02 ± 0.09	10.05 ± 0.09

As evidenced by Table 5.3, the ‘vacuum’ material used is not truly vacuum but rather a user-defined material with very low density. This is because COMSOL cannot complete computations with non-material (zero density, no material properties) domains. The thermal conductivity and the specific heat capacity of ‘vacuum’ were deliberately kept similar to those of the other materials to show that the major contributing factor is the density. In fact, the results show that the underestimation is a function of the material density. Although they are slightly larger than 10 Gy, the dose values reported for vacuum are the most accurate since the real values are themselves slightly over 10 Gy. This is because the beam heat sources do not turn off sharply at the desired timepoints. Due to the finite size of the time steps, the heat sources can only turn off gradually over the lapse of 0.1–0.3 s, leading to a total delivery of 10.02–10.05 Gy. Having a low but non-

zero density, the dose values for air are almost accurate but they do not make up the difference completely. The dose values for aerogel, the densest of the three materials, have the highest underestimation. The following paragraph proposes a working theory to explain this behaviour.

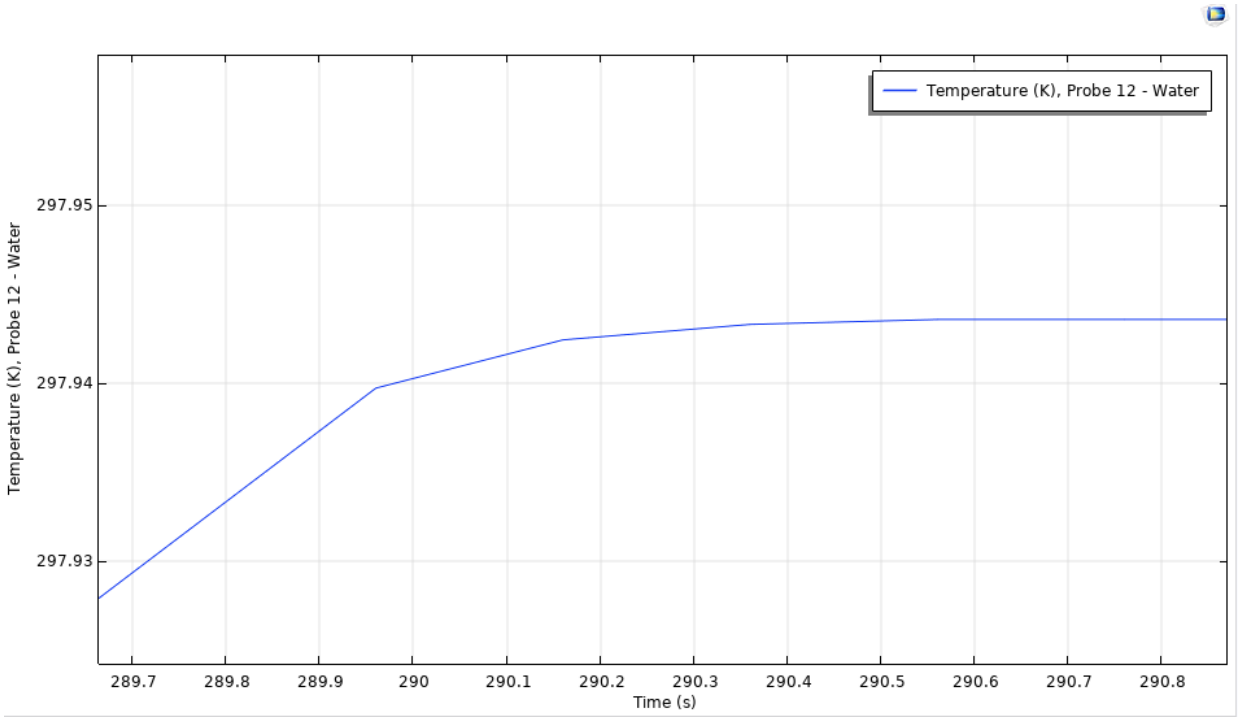

Figure 5.7. Screenshot of the temperature profile as a function of time for the sensing thermistor in the core depicting the gradual and smooth turn-off of the beam heat sources. This temperature profile was obtained with a different COMSOL model that simulates quasi-adiabatic mode of operation. By setting the thermal conductivities of all materials in the model to zero, heat loss is eliminated, and the action of the beam is reflected directly onto the temperature curve.

During isothermal operation, the graphite bodies are kept at their predefined temperature setpoints. However, there is no active-controlled heating in the insulation layers. During irradiation, the aerogel will experience a temperature rise of approximately 0.95 $mK \cdot Gy^{-1}$. This increase in temperature will reduce the rate of heat loss from the core to the first aerogel layer, thereby reducing the amount of power required to keep the core isothermal during and after irradiation. Therefore, the absorbed dose-to-graphite extracted from the power curves will be underestimated.

This effect can account for the underresponse seen in the isothermal runs presented in Table 5.2. Furthermore, this theory can explain the exacerbated underresponse seen in geometry 7, as well as the marginally increased (and more accurate) dose value obtained for geometry 5. In geometry 7, the lower thermistor and the lower spherical contact are both simultaneously used for heating. Because of their proximity and the coupling of their heat functions through the wire, the region of aerogel between them will experience a high temperature rise relative to the other regions of the aerogel layer. The excess heat from this 'hot spot' will further decrease the rate of heat loss from the core to the first aerogel layer, hence producing a more pronounced underresponse. Conversely, geometry 5 has the heating thermistor located at the bottom of the core and the heating contact located at the top of the jacket. Additionally, the thermistor separation is maximal, meaning that the separation between heating points is also maximal. Thus, the heat is more uniformly distributed in these inner layers, so no 'hot spot' is generated in the aerogel. As a result, the rate of heat loss from the core is not reduced so greatly, and the power required to keep it isothermal is closer to the real power introduced by the beam heat source.

Finally, the insulating material tests can be used to get an estimate of the heat transfer correction factor for Aerrow MK7 in isothermal mode. The reason for this is that the low-density 'vacuum' used in this model is almost equivalent to a material with no thermal conductivity, so the heat transfer from the core and the jacket is effectively zero. Recall that the heat transfer correction factor can be calculated as the ratio of a signal obtained in a scenario with no heat transfer to that obtained in a scenario with normal heat transfer processes. In this case, the heat transfer correction factor for Aerrow MK7 can be approximated by taking the ratio of the dose values obtained with vacuum to those obtained with aerogel. Using the dose values for the first beam in Table 5.3:

$$k_{\mathrm{ht,iso}} \approx \frac{\overline{D}_{\text{without heat transfer}}}{\overline{D}_{\text{with heat transfer}}} \approx \frac{\overline{D}_{\text{vacuum}}}{\overline{D}_{\text{aerogel}}} = \frac{10.02}{9.5} = 1.055 \qquad (5.11)$$

The dose values from the second beam could also be used but these may be skewed by the heat trapped in the model from the first beam. The heat transfer correction factor remains to be more accurately estimated by turning all thermal conductivities properly to zero. The estimation would also benefit from a bigger model volume or an isothermal boundary condition (to reduce the effect of heat entrapment) and an optimal thermistor configuration (to reduce the effect of aerogel heating from the heating points in adjacent bodies). This estimate is not expected to change dramatically.

5.3.4. Jacket PID Functionality Tests

After establishing a likely source for the underresponse, one conspicuous question remains: If the isothermal COMSOL model predicts that Aerrow MK7 will show an underresponse, why did the real isothermal runs show an overresponse? How can these results be reconciled? The starting point to this question has already been suggested by the data, which will be briefly reviewed here:

i. Only one simulation in the thermistor placement tests (Table 5.2) did produce an overresponse. This was the simulation with no working PID in the jacket (geometry 8).
ii. The qualitative behaviour of Figure 5.5A, which depicts the function of Aerrow MK7 with no PID in the jacket (or shield), matches that of the real isothermal runs.
iii. The construction error that led to an unattached heating contact in the jacket, mentioned in Chapter 4, is still left to be considered as the possible source of the overresponse.

Remarkably, the results in (i) and (ii) represent the only times in which an overresponse on the correct order of magnitude (in fact, larger still) was able to be replicated. All other sources of error considered thus far could not account for an overresponse of ~ 20 %. The contact problem mentioned in (iii) is also important because it hampered the active-controlled heating function in the real jacket. These three observations strongly suggest that the overresponse seen in the real isothermal runs could be accounted for with a non-functional or partially functional jacket PID.

To test this theory, jacket PID functionality tests were conducted with the isothermal COMSOL model. The simulations with a non-functioning jacket PID (i.e. no heating function) were repeated. Additionally, three more simulations were conducted with partially working jacket PIDs. The second and third tests possess a half-functional PID to closely model the effect of one (out of two) heating contacts coming loose. In the real probe, the heating contacts are glued to the side of the jacket. Ideally, these tests should divide the jacket into lateral halves. However, in an axisymmetric model, the best approximation is to divide the jacket into two halves longitudinally, i.e. top and bottom. The fourth test models a jacket PID that does not act on the jacket caps; it only provides active-controlled heating to the lateral sides of the cylinder. This test is especially interesting because the caps are machined separately and they are assembled to the rest of the jacket with drops of cyanoacrylate, which may inhibit the proper conduction of heat throughout the jacket.

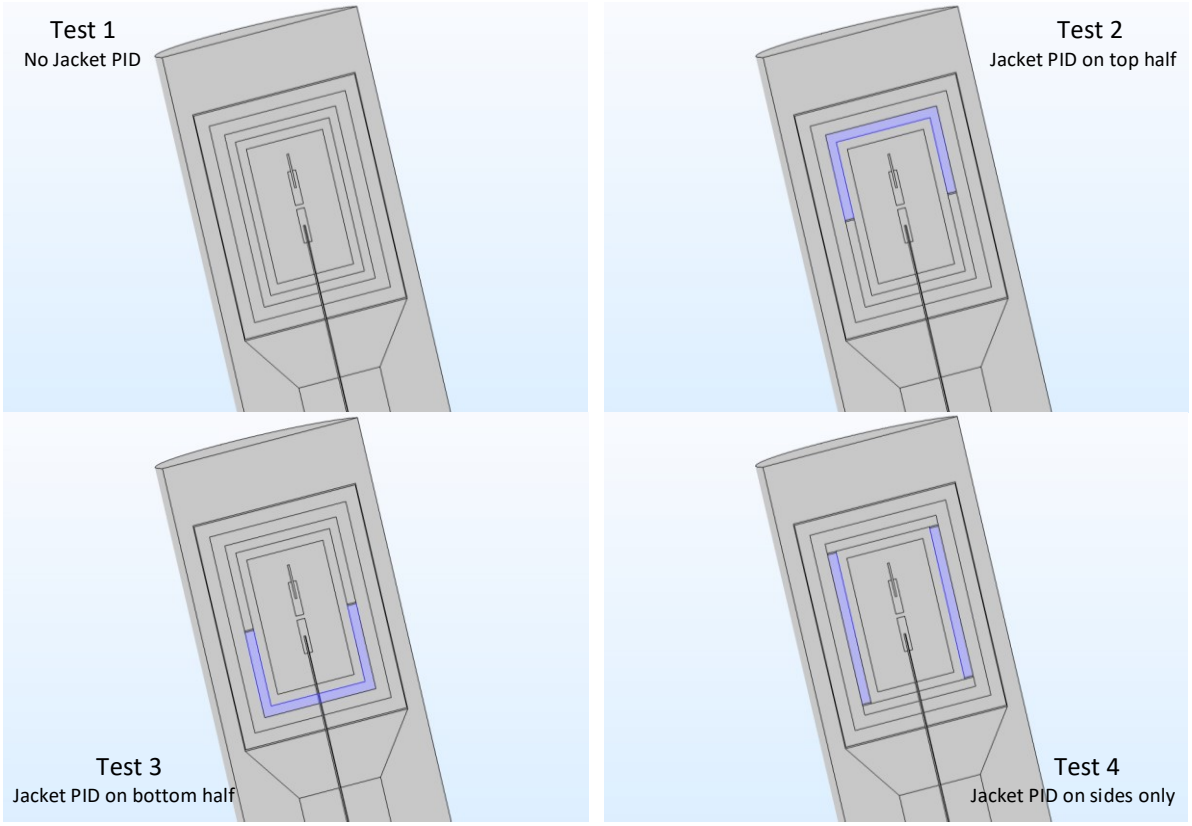

Figure 5.8. Four screenshots of the Aerrow MK7 isothermal COMSOL model depicting the four different jacket PID functionality tests. Note the absence of the spherical contacts. The regions highlighted in blue show the domains acted upon by the jacket PID; these are the volumes that acted as the heating and sensing components for each simulation.

For the jacket PID functionality tests, the model was subjected to two 60 s-long 10 Gy beam heat sources after a 50 s stabilization period, the separation between thermistors was set to 1.25 mm, the material assigned to the thermal insulation layers was aerogel, and the heating in the core was done by the lower thermistor. The results are presented below. The uncertainties for these tests are relatively bigger than for previous tests because the drifts and variations in the power output curves of the core are more pronounced due to the loss of a fully functional jacket PID.

Table 5.4. Results from the jacket PID functionality tests in the Aerrow MK7 isothermal COMSOL model. For each test, the heating function in the jacket is described, and the absorbed dose-to-graphite (D_{core}) value obtained for each beam is given. As with previous tests, the uncertainties on the dose values were estimated by changing the power curve analysis method and by varying the total length of the exclusion zones.

Test	Jacket PID Description	D_{core} [Gy]	
		Beam 1	Beam 2
1	No PID in the jacket	15.1 ± 0.3	13.9 ± 0.4
2	PID acts only on the top half of the jacket	14.2 ± 0.4	11.0 ± 0.2
3	PID acts only on the bottom half of the jacket	13.2 ± 0.4	11.0 ± 0.2
4	PID acts only on the lateral sides of the jacket	9.9 ± 0.1	9.2 ± 0.1

Noticeably, the dose values for a given beam do not agree within combined uncertainties between tests. This is to be expected, since the probe is deliberately made to work in suboptimal conditions. The results indicate that a malfunctioning PID in the jacket is indeed capable of producing an overresponse. These results are logical; if the PID is not working, the temperature setpoint cannot be maintained. The jacket will therefore naturally tend to cool down to room temperature, and the increased temperature gap with the core will accelerate the rate of heat loss from the core. In turn, the core will need to receive an increased amount of power to remain isothermal, resulting in an overestimation of the absorbed dose extracted from the power curves.

In the interest of transparency, it should be mentioned that this increasing temperature gap between the core and the jacket was not observed experimentally. This is part of what made the source of

the overresponse so difficult to identify. The experimental temperature curves from the sensing thermistor in the jacket were very stable, showing typical fluctuations of less than ± 0.04 mK. However, part of what is demonstrated by the jacket PID functionality tests is that the jacket PID can operate accurately in some regions of the jacket and still leave other regions without effective active-controlled heating. In other words, although the temperature of the region surrounding the jacket's sensing thermistor was locally stable, it can be supposed that other parts of the jacket were drifting away from the setpoint. This is tantamount to a jacket PID that is not completely functional, even if indicated otherwise by the sensing thermistor. In such a case, an overresponse could still be observed.

Not only can a malfunctioning PID produce an overresponse, but in addition, the overresponse itself seems to be a function of the degree of functionality in the PID. Table 5.4 demonstrates that as the functionality of the jacket PID is recovered, the magnitude of the overresponse lessens. The results from tests 2 and 3 are particularly informative, as they represent an overresponse in the range observed for the real isothermal runs. This study suggests that the overresponse is a consequence of a partially functional jacket PID, specifically one that is approximately half-functional, which could be explained by the faulty heating caused by the loose heating contact in the jacket.

The results from this test will be recapitulated in the next chapter, which will discuss key findings from this and other chapters in an overview of all the experimental and computational tests presented in this thesis.

5.4. References

1. COMSOL Multiphysics Cyclopedia. Detailed Explanation of Finite Element Method (FEM). Burlington, MA: 2016; comsol.com. Accessed October 2020.

2. Manor Tool & Manufacturing Company. The Benefits of Finite Element Analysis in Manufacturing. Schiller Park, IL: 2020; manortool.com. Accessed October 2020.

3. COMSOL Inc. Products. Heat Transfer Modelling Software for Analyzing Thermal Effects. Burlington, MA: 2016; comsol.com. Accessed October 2020.

4. COMSOL Multiphysics Video Gallery. Creating 2D Models from 3D Geometries in COMSOL Multiphysics. Burlington, MA: 2016; comsol.com. Accessed October 2020.

5. Engineering ToolBox. Thermal Conductivity of Metals, Metallic Elements and Alloys. 2005; engineeringtoolbox.com. Accessed October 2020.

6. Karditsas PJ, Baptiste MJ. Thermal and Structural Properties of Fusion related Materials: Pure Copper. Oxfordshire, UK: 1995.

7. Brandrup J, Immergut EH, Grulke EA. Physical Constants of Poly(methyl methacrylate). In Brandrup's *Polymer Handbook*, Second Edition. Wiley. New York, NY: 1975.

8. Goodfellow Online Catalogue Source. Polymethylmethacrylate (PMMA, Acrylic) Material Information. Cambridge, UK: 2008; goodfellow.com. Accessed October 2020.

9. Rule DL, Smith DR, Sparks LL. Thermal Conductivity of a Polyimide Film Between 4.2 and 300K, With and Without Alumina Particles as Filler. Gaithersburg, MD: National Institute of Standards and Technology Internal Report (NISTIR) 3948. 1990.

10. COMSOL Documentation. Foundations of the General Heat Transfer Equation. Heat Transfer Module User's Guide. Burlington, MA: 2015;89-100.

11. Spakovszky ZS, Greitzer EM, Waitz IA. Applications of the Second Law: The Thermodynamic Temperature Scale. Retrieved from Massachusetts Institute of Technology, Thermodynamics and Propulsion 16.01-16.02. Delivered Fall 2008.

12. Quinn TJ, Martin JE. A radiometric determination of the Stefan-Boltzmann constant and thermodynamic temperatures between -40 °C and +100 °C. *Philos. Trans. R. Soc. London*. 1985;316(1536),85-189.

13. Ascher UM, Petzold LR. Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM). 1998.

14. Iserles, A. A First Course in Numerical Analysis of Differential Equations. Cambridge, UK: Cambridge University Press. 1996.

15. Control Tutorials for MATLAB and Simulink (CTMS). Introduction: PID Controller Design. Ann Arbor, MI: 2019; ctms.engin.umich.edu. Accessed October 2020.

16. Choi Y, Jeon KJ, Park Y, Hyun S. Numerical simulation of a heat-loss compensated calorimeter. *Int. J. Comp. Meth. and Exp. Meas*. 2019;7(3),285–296.

17. Radu D, Guerra AS, Ionita C, Astefanoaei I. Heat loss through connecting thermistor wires in a three-body graphite calorimeter. *Metrologia*. 2010;47(3),179–191.

Chapter 6. Conclusion

6.1. Aerrow MK7: An Overview

The aim of this work has been to report on the development of the Aerrow MK7 dosimetry system. Aerrow MK7 is a miniaturized version of its predecessor, Aerrow, a probe-format graphite calorimeter intended to bring the advantages and accuracy of absorbed dose calorimetry to the clinical setting.[1,2,3] Increasingly, new radiation therapy modalities and delivery machines exploit the use of small and nonstandard photon fields to treat small and irregularly shaped tumours. The small dimensions of these fields mandate the use of appropriately sized detectors.[4] By reducing the volume of the detector by a factor of eight, the design of Aerrow MK7 seeks to extend the functionality of Aerrow to meet the growing need for accurate small-field clinical dosimetry. The availability of Aerrow MK7 as a commercial dosimeter would be of great significance, as it represents the first dosimeter capable of producing accurate absorbed dose measurements for small high-energy photon fields, in the clinical environment, and in absolute terms.

Figure 6.1. Photograph of a side-by-side comparison of Aerrow MK7 (red), an Exradin A1SL ionization chamber (black), and a coin for size reference. The ionization chamber pictured is the same one used for comparing the dosimetric performance of Aerrow MK7 in quasi-adiabatic mode. The red and porous external appearance of Aerrow MK7 is given by the waterproof plastic coating that surrounds the entire acrylic stem and part of the cable.

However, the miniaturization process brought considerable challenges that prevented the detector from operating at full functionality. To gain a better understanding of these challenges and how to overcome them, the Aerrow MK7 system has been subjected to various studies and tests throughout this thesis. The following sections provide a review of the most important results obtained from these investigations.

6.2. Quasi-Adiabatic Mode of Operation

It was found that when operating in quasi-adiabatic mode, Aerrow MK7 is capable of accurate absorbed dose-to-water determination down to nominal field sizes of 2 × 2 cm^2, corresponding to an equivalent square field size at depth (S_{clin}) of 2.16 cm. The absorbed dose-to-water measurements obtained with Aerrow MK7 were compared to those obtained with a small field-suitable reference-class ionization chamber, the Exradin A1SL.[4] The maximum percentage difference between the dose measurements of the two detectors was 0.7 % for any given (equivalent) irradiation. Furthermore, the measurements from these two dosimeters agreed within combined standard uncertainties — the D_w measurements from Aerrow MK7 had uncertainties of 1.1–1.4 %, which is almost the same accuracy that was achieved with the ionization chamber. The uncertainties can be further improved with some relatively minor work, listed in Table 6.1 at the end of this chapter. This level of accuracy is noteworthy given that the quasi-adiabatic signals displayed a pronounced heat loss effect, causing them to decay relatively rapidly. This problem was efficiently circumvented through an alternative dose analysis method, which has been described at length in previous work.[5]

6.3. Isothermal Mode of Operation

When operating in isothermal mode, Aerrow MK7 was not capable of producing accurate absorbed dose-to-water measurements. The probe suffered from an irregular overresponse that saw absorbed dose measurements overestimated by 17–24 % in most cases, with some isothermal runs producing

as much as a 31 % overresponse and others producing as little as a 12 % overresponse. The cause of this overresponse was initially unknown. Due to the numerous factors involved in absorbed dose determination, finding the source of a response problem in Aerrow MK7 is a non-trivial task. The principal objectives of this thesis can be framed in the context of this problem: (a) to identify the likely source(s) of the overresponse and (b) to propose implementable changes to the detector's design to mitigate the problem of the overresponse. The following section contains an overview of the studies conducted to achieve the former objective. The remainder of the chapter will discuss design revisions associated with the latter objective.

6.3.1. Review of Results

To identify the source of the overresponse, it was necessary to meticulously analyze all steps in the absorbed dose determination chain, from physical to computational aspects. The results from these analyses demonstrated that none of the following could potentially account for errors large enough to explain the overresponse: the readout circuitry, the thermistor circuitry, the conductive heat leakage from the copper wires, the placement of the thermistors within the core, the change of the thermal gradients within the core over time, the measured mass of the sensitive volume, the absorbed dose analysis method, the inclusion/exclusion of datapoints close to the beam-on and beam-off timepoints, the mass impurity correction factor, the heat transfer correction factor for isothermal mode, the small-field output correction factor, the percentage dose at depth conversion factor, or the medium conversion factor.

One may ask whether all these potential sources of error could, in combination, produce an overresponse on the correct order of magnitude. Albeit next to impossible that all these factors add constructively, this is a possibility. However, the maximum possible error that could be attained in this manner is calculated to be around 10 %. Therefore, all factors mentioned above can be confidently discarded as the possible source(s) of the overresponse.

There is one caveat regarding the placement of thermistors and the heat transfer correction factor for isothermal mode. Although they cannot account for the overresponse, they can still affect the final absorbed dose value by several percent, as seen in the last chapter. The effect of the placement of the thermistors and heating contacts on the absorbed dose can be ignored only when (a) the heating and sensing thermistors in the core are within 2.35 mm of each other, and (b) the heating components from two adjacent graphite bodies are not close together. As for $k_{\mathrm{ht,iso}}$, a rough early estimate puts it at 1.055, representing an upwards correction of 5.5 % on the absorbed dose. The real value is expected to be anywhere between 5 to 10 %. The magnitude of this correction factor is driven primarily by the presence of the aerogel layers, which experience a temperature rise during irradiation, thereby reducing the rate of heat loss from the core to the jacket.

A FEM-based model of Aerrow MK7 was constructed with COMSOL Multiphysics® to simulate thermal transport between the different components of the detector and their surroundings. These heat transfer studies revealed that the overresponse is most likely caused by a malfunctioning PID controller in the jacket. Specifically, an overresponse of 10–30 % can be reproduced in simulations with a jacket PID that is approximately half-functional. It is possible for these conditions to hold true in the physical probe since a construction error led to the detachment of one of two contacts used for active-controlled heating in the jacket.

6.3.2. Discussion

This thesis presented an avenue for modelling the isothermal mode of operation in a three-body absorbed dose calorimeter, something that has not been previously accomplished. This model is of potential interest to all researchers involved in absorbed dose calorimetry and even to those involved with primary absorbed dose standards at large. The model could easily be modified to suit other applications, like modelling isothermal operation in a primary standards aluminum calorimeter for ultra-high dose rate radiotherapy.[6]

The model predicted that for any non-vacuum material used for the thermal insulation layers, an underresponse would be observed. A working theory was proposed to explain this effect. This theory was able to explain the cause of the underresponse as well as the variations seen in the magnitude of the underresponse for the different thermistor placement tests. Qualitatively, this theory is also able to explain the nature of the asymptotic-like recovery of the post-drift exhibited by the experimental isothermal runs. As the aerogel layers cool down, the rate of heat loss from the core to the jacket accelerates, and the power required to keep the core isothermal is slowly restored to the level required prior to irradiation.

The model also helped to establish a working theory for the overresponse seen in the real isothermal runs. The theory is capable of explaining both the source of the overresponse as well as its irregular nature. Since one of the heating contacts in the jacket is not properly attached, it is natural that the total effective functionality of the jacket varies from one run to the next. This is supported by the experimental measurements — the dose values obtained during the second week, during which the jacket heating contact was noticeably less functional, possessed a larger degree of variation.

Two questions arise regarding the consistency between model predictions and physical results. One is, if the aerogel layers can themselves heat up during irradiation, why was an overresponse not observed for the quasi-adiabatic results? The answer to this question is two-fold. First, to produce an overresponse, the aerogel would necessarily have to act as a heat source and dump thermal energy into the core. Reducing the rate of heat loss from the core would not affect the response negatively, unlike in isothermal mode. The aerogel cannot act as a heat source because the specific heat capacity of aerogel is larger than that of graphite, so the rise in temperature per unit absorbed energy is less. Second, even if the first point was not true, the absorbed dose in the quasi-adiabatic runs was calculated from the dose rate evaluated at the very beginning of the

irradiation period. At this timepoint, the heat transfer effects from the aerogel have not had enough time to change the signal, so an overresponse would not be registered.

Another question one can ask regarding the model predictions is: why was such an underresponse not seen in the isothermal measurements performed with the predecessor prototype, Aerrow? The following two considerations may explain why. The first is that the formulation of aerogel used in Aerrow is half as dense, having a mass density of 0.2 $g \cdot cm^{-3}$. The results from the insulating material tests established that the underresponse is primarily dictated by the density of the material employed for thermal insulation. The less dense the insulating layers, the lesser the underresponse. The second consideration is that, even though the relative scaling between all components is equal for both detectors, the surface area-to-volume ratio is not. For any cylindrical core with height h and radius r, the surface area-to-volume ratio is $\frac{SA}{V} = 2\left(\frac{1}{h} + \frac{1}{r}\right)$. Since Aerrow MK7 is a half-scale of Aerrow, its surface area-to-volume ratio is twice as large. This is important because the bigger the surface area-to-volume ratio of the core, the more contact there is per unit mass with the adjacent aerogel layer. Because the presence of aerogel is responsible for reducing the rate of heat loss, leading to an underresponse, a smaller surface area-to-volume ratio is preferable.

To this end, a mathematical function is being developed to reoptimize the geometry of the probe based on these findings. The function provides optimal dimensions for all the aerogel and graphite components. Additionally, the function incorporates practical constraints relating to small-field dosimetry, typical diameter of solid phantom inserts, thickness of electrical components, etc. The design of the next prototype iteration, Aerrow MK8, will be justified based on these considerations.

6.4. Future Work

Table 6.1 below contains a list of possible future projects that could expand and/or improve the operation of Aerrow MK7 in terms of functionality, practicality, accuracy and reliability.

Table 6.1. A list of possible future experiments and design revisions to integrate into Aerrow MK7 and/or its setup.

Item	Description	Purpose
	Quasi-Adiabatic Mode of Operation	
Experimental determination of c_P	Perform an in-house determination of the specific heat capacity of the graphite used in Aerrow MK7.	To reduce the uncertainty associated with the value of c_P, which is currently estimated from literature.
Fast-sampling diode	Embed a fast-sampling diode into the acrylic stem of Aerrow MK7.	To reduce the uncertainty associated with the measured time of irradiation, which is currently estimated from the sampling frequency of the linac itself (0.1 s).
Longer irradiations	Subject Aerrow MK7 to longer-duration beams when operating in quasi-adiabatic mode.	To reduce the relative uncertainty associated with the measured time of irradiation.
Active-controlled heating in the jacket	Implement an active-controlled heating system in the jacket to match the temperature rise measured in the core.	To reduce the rate of heat loss from the core to the jacket, hence satisfying the adiabatic condition more accurately.
Graphite phantom	Place the Aerrow MK7 in a graphite phantom rather than a water phantom.	To reduce the rate of heat loss from the probe to the large surrounding water phantom, which acts as a heat sink.
Higher thermal insulation	Increase the quality of thermal insulation in the probe. This can be done either by adding another aerogel layer beyond the shield or by decreasing the density of the aerogel. The latter option also has benefits for the performance of isothermal mode, as discussed throughout this chapter.	To reduce the rate of heat loss from the probe to the large surrounding water phantom, which acts as a heat sink.
	Isothermal Mode of Operation	
Thicker jacket	Design a prototype, in accordance with practical constraints, that includes a jacket of larger thickness.	To be able to embed the heating contacts and sensing thermistors of the jacket PID directly into the jacket, diminishing the chance of detachment.
3D COMSOL model	Upgrade the existent Aerrow MK7 model in COMSOL from a 2D axisymmetric to a 3D geometry. Tolerable computation times are feasible using parallel processing (PP) on a high-performance computing cluster.	To more accurately represent the asymmetrical placement of the thermistors within the probe; to better understand the effect of thermistor placement on the absorbed dose.
Convection in COMSOL model	Add the 'Heat Transfer in Fluids' interface in the COMSOL model of Aerrow MK7.	To more accurately model convective motions within the air and water domains.
Isothermal boundary condition in COMSOL model	Impose an isothermal boundary condition on the COMSOL model of Aerrow MK7.	To more accurately model the amount of heat dissipated from the probe into the large water phantom.
Larger and more realistic volume of COMSOL model	Modify the Aerrow MK7 model in COMSOL to include a larger water phantom (also feasible with PP).	To more accurately model the amount of heat dissipated from the probe into the large water phantom.
Accurate determination of $k_{\mathrm{ht,iso}}$	Conduct a more thorough and systematic investigation of the heat transfer correction factor for isothermal mode using the Aerrow MK7 model in COMSOL.	To reduce the uncertainty associated with the value of $k_{\mathrm{ht,iso}}$; to determine the dependencies (if any) of $k_{\mathrm{ht,iso}}$; e.g. dependencies on dose rate, thermistor placement, insulating material, etc.

6.5. References

1. Renaud J, Marchington D, Seuntjens J, Sarfehnia A. Development of a graphite probe calorimeter for absolute clinical dosimetry. *Med Phys*. 2013;40:020701.

2. Renaud J, Sarfehnia A, Bancheri J, Seuntjens J. Aerrow: a probe-format graphite calorimeter for absolute dosimetry of high-energy photon beams in the clinic. *Med Phys*. 2018;45:414–428.

3. Renaud, J. On the Development of Absorbed Dose Calorimeter Systems for Absolute Clinical Dosimetry. PhD Thesis, McGill U. Department of Physics. Montreal, QC: 2016.

4. International Energy Agency (IAEA). Dosimetry of small static fields used in external beam radiotherapy: an international code of practice for reference and relative dose determination. *Technical Report Series* 483. Vienna, AT: 2017.

5. Côté, B. Numerical design and development of a probe-format graphite absorbed dose calorimeter for use in small-field radiotherapy. MSc Thesis, McGill U. Department of Physics. Montreal, QC: 2019.

6. Bourgouin A, Schüller A, Hackel T, Kranzer R, Poppinga D, Kapsch RP, McEwen M. Calorimeter for real-time dosimetry of pulsed ultra-high dose rate electron beams. *Frontiers in Physics*. 2020;8:400-410.