ABSTRACT

Chord Length Methods for Proton Computed Tomography

John P. Stenson, M.S.

Mentor: Keith E. Schubert, Ph.D.

Proton Computed Tomography (also known as pCT) is an emerging medical imaging method. The three dimensional images reconstructed from pCT scans provide information about the Relative Stopping Power (RSP) of the object scanned. Proton and ion therapies can administer radiation with higher conformity than traditional photon therapies. The RSP values provided by pCT scans are particularly useful when planning proton and ion related therapies. pCT reconstructions depend on solving the linear system of equations represented by $Ax = b$. It is necessary to accurately construct the matrix A by determining its elements. Which elements in A are non-zero is determined by Most Likely Path methods. The values of these non-zero elements are known as chord lengths. This work investigates different methods for determining chord length, and the effect they have on pCT reconstructions.

Chord Length Methods for Proton Computed Tomography

by

John P. Stenson, B.S.

A Thesis

Approved by the Department of Electrical and Computer Engineering

_____

Kwang Lee, Ph.D., Chairperson

Submitted to the Graduate Faculty of
Baylor University in Partial Fulfillment of the
Requirements for the Degree
of

Master of Science in Biomedical Engineering

Approved by the Thesis Committee

_____

Keith Schubert, Ph.D., Chairperson

_____

Robert J. Marks, Ph.D.

_____

Jonathan Rylander, Ph.D.

Accepted by the Graduate School

December 2017

_____

J. Larry Lyon, Ph.D., Dean

Page bearing signatures is kept on file in the Graduate School.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER ONE

Introduction

Modern medical treatment depends heavily on internal imaging methods, particularly for the treatment of cancer and other structural abnormalities. This is partially because internal imaging methods largely remove the necessity of performing exploratory surgery to diagnose pathologies. However, a major benefit of such imaging methods is their applicability to planning radiation therapy for cancer treatment. Proton Computed Tomography, referred to as pCT, is an emerging medical imaging modality. As a computed tomography method, pCT reconstructs three dimensional images of the internal structure of the scan subject by evaluating data gathered from particles which passed through the subject. Images resulting from pCT appear similar to those created by x-ray Computed Tomography, but actually represent different data. X-ray CT methods produce density maps of tissues, whereas pCT scans yield maps of the Relative Stopping Power of tissues (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015). The Relative Stopping Power of a material is a measurement of the energy change induced by passing through an object. Proton CT reconstruction methods find these RSP values by solving a linear system of equations representing the paths of protons through the subject (Schulte, Penfold, Tafas, and Schubert 2008; Penfold, Rosenfeld, Schulte, and Schubert 2009; Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015; Karbasi, Schultze, Giacometti, Plautz, Schubert, Schulte, and Bashkirov 2015). In this system, it is essential to estimate the lengths of voxel-path intersections (Schulte, Penfold, Tafas, and Schubert 2008; Penfold, Rosenfeld, Schulte, and Schubert 2009). This work compares different methods for determining the lengths of these intersections. The accuracy of the reconstructions produced are evaluated using traditional methods and a new, therapy-focused technique.

1.1 Proton-Matter Interactions

An understanding of the behavior of protons is useful when considering both proton computed tomography and proton therapy. Therefore, a brief summary is provided here. Protons do not interact with matter in the same way as photons. These interactions are more difficult to predict, but provide some unique properties which can be used advantageously. The three types of interactions which can occur while a proton is traveling through matter include: stopping, Multiple Coulomb Scattering, and nuclear interactions (Paganetti 2016).

1.1.1 Stopping

The first kind of interaction which a proton will undergo is that of stopping. The fact that protons do stop in matter is very important (Paganetti 2016). The deceleration induced is what gives rise to the term "stopping power," the standard measurement of matter produced by pCT reconstructions (Paganetti 2016). This deceleration is due to the collisions of the proton with the electrons of atoms. The proton loses more energy per time as it traverses farther because the energy lost is a function of the time it spends near the electron with which it collides (Paganetti 2016). Therefore a slower proton will be slowed more in the same distance of the same element than would a faster proton. During a pCT scan, protons actually coming to rest in the body is undesirable, as their behavior cannot be tracked in a detailed manner. Additionally, this would cause increased radiation dose to the subject scanned. Therefore, scans for reconstructing pCT usually employ higher energy protons than would be used for a treatment. The stopping theory of proton-matter interaction is particularly important for proton therapies (Paganetti 2016). The point immediately before which a proton stops is known as the Bragg peak, and is the area of largest radiation dose, while areas after the stopping point receive little

to no radiation dose (Paganetti 2016; Patyal 2007). This allows proton therapies to have a very good conformity, a subject discussed in detail in Chapter Two.

1.1.2 Multiple Coulomb Scattering

Another interaction which protons may undergo is that of multiple Coulomb scattering, or MCS. MCS causes small changes in the path of the proton due to the electromagnetic forces acting between the protons in the nucleus of atoms and the free proton (Paganetti 2016). As both are positive charges, they repel each other. This causes the proton to undergo many small changes in its direction, leading to the term scattering (Paganetti 2016). The angles produced by this kind of interaction are minimal, but they do prevent protons from traveling in simple to predict straight lines. Trying to find the path which such scattering caused a proton to follow is the focus of Most Likely Path procedures, an important element of pCT reconstruction techniques.

1.1.3 Nuclear

The final kind of interaction which can happen is a nuclear interaction between the proton and the nucleus of an atom. Low interest interactions occur when a proton elastically collides with a nucleus (Paganetti 2016). This may simply scatter the proton slightly, or it may slightly excite the nucleus with which it collides (Paganetti 2016). The more important interaction occurs when nonelastic collisions occur, causing a small burst of radiation, and preventing the proton from being tracked further, as it is indistinguishable from other protons set loose by the collision (Paganetti 2016). These kinds of interaction are unavoidable and produce negligible harmful effects, but are of no particular use to pCT imaging or therapy techniques.

1.2 Proton Computed Tomography

Proton Computed Tomography was first proposed in 1963 by A.M. Cormack, and has since been examined by many researchers (Cormack 1963; Hanson 1979; Li, Liang, Singanallur, Satogata, Williams, and Schulte 2006; Schulte, Penfold, Tafas, and Schubert 2008). Initial interest in pCT was due to the lower radiation dose required to image internal tissues when compared to x-ray Computed Tomography (Hanson 1979; Hanson, Bradbury, Koeppe, Macek, Machen, Morgado, Paciotti, Sandford, and Steward 1982; Li, Liang, Singanallur, Satogata, Williams, and Schulte 2006). However, this was not enough to offset the difficulties of pCT at that time. Limits to spatial resolution, reconstruction time and applicability of the RSP values generated limited the effort put into pCT development. More recent investigation into proton and ion therapies has increased interest the Relative Stopping Power data provided by pCT images (Li, Liang, Singanallur, Satogata, Williams, and Schulte 2006; Schulte, Penfold, Tafas, and Schubert 2008; Bashkirov, Schulte, Coutrakon, Erdelyi, Wong, Sadrozinski, Penfold, Rosenfeld, McAllister, and Schubert 2009; Penfold, Rosenfeld, Schulte, and Schubert 2009).

1.2.1 Resolution

Ever since the development of pCT, the primary obstacle to its adoption has been the lack of spatial resolution due to Coulomb scattering (Hanson 1979; Hanson, Bradbury, Koeppe, Macek, Machen, Morgado, Paciotti, Sandford, and Steward 1982). In early pCT scans, there were deficiencies in spatial resolution as compared to x-ray CT scans(Hanson 1979). However,as far back as 1982, pCT was employed to generate images comparable with those created by x-ray CT scans, though they did not achieve identical performance(Hanson, Bradbury, Koeppe, Macek, Machen, Morgado, Paciotti, Sandford, and Steward 1982). A study in 1994 proved there were limits to the spatial resolution of proton CT scans (Schneider and Pedroni 1994).

However, as reconstruction and scanning techniques have been developed and improved, it has been demonstrated that the locations of protons inside the scan subject can be tracked to within less than 0.5 mm, matching the resolution of common CT and MRI machines (Bruzzi, Blumenkrantz, Feldt, Heimann, Sadrozinski, Seiden, Williams, Bashkirov, Schulte, Menichelli, Scaringella, Cirrone, Cuttone, Randazzo, Sipala, and Presti 2007; Bashkirov, Schulte, Coutrakon, Erdelyi, Wong, Sadrozinski, Penfold, Rosenfeld, McAllister, and Schubert 2009; Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016). This level of precision has eliminated one major point of concern about the utility of pCT for medical purposes. In the case of this experiment, the images reconstructed employ voxels with dimensions: $1mm \times 1mm \times 2.5mm$ in the x, y, and z, dimensions, respectively.

1.2.2 Reconstruction

Other obstacles to the common usage of pCT were previously the accuracy of RSP values and the time spent on the reconstruction process. The reconstruction of proton Computed Tomography images depends on finding an approximate solution to a linear system of equations (Schulte, Penfold, Tafas, and Schubert 2008; Penfold, Rosenfeld, Schulte, and Schubert 2009; Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015; Karbasi, Schultze, Giacometti, Plautz, Schubert, Schulte, and Bashkirov 2015). These are represented as a matrix equation of the form $Ax = b$. The matrix A represents the path taken by the protons during the scan, the matrix x represents the RSP values of each voxel in the scan volume, and the matrix b represents the total stopping power along the entire path of the corresponding proton (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015). In pCT, the elements of the b matrix are derived from measured data, and the values within the A matrix are determined using a chosen path-estimation technique. However, most

reconstructions, including all those used in relation to this work, make use of the Most Likely Path (MLP) technique outlined by Schulte *et al.* in a 2008 paper, with some improvements (Schulte, Penfold, Tafas, and Schubert 2008). The MLP method only determines which elements of A are non-zero, and the values of these elements must be calculated to achieve accurate pCT reconstructions (Penfold, Rosenfeld, Schulte, and Schubert 2009). The values of the A matrix correspond to the length of intersection between a given voxel and proton path (Penfold, Rosenfeld, Schulte, and Schubert 2009). This work investigates the different possible methods by which these intersection lengths can be calculated, and the methods under consideration are described in detail in Chapter Four. These methods have clear effects on both timing of reconstructions and on accuracy (Penfold, Rosenfeld, Schulte, and Schubert 2009). Once the values of A are calculated, the system $Ax = b$ can be solved via a variety of algorithms (Penfold, Schulte, Censor, Bashkirov, McAllister, Schubert, and Rosenfeld 2010). The methods used to determine the solution to this equation have great effect on the time required to reconstruct images (Penfold, Schulte, Censor, Bashkirov, McAllister, Schubert, and Rosenfeld 2010). The method used in all reconstructions for this work is the Robust Diagonally Relaxed Orthogonal Projections method, described by Karbasi *et al.* in the 2015 article (Karbasi, Schultze, Giacometti, Plautz, Schubert, Schulte, and Bashkirov 2015). This method has proven accurate, robust to errors in the system, and efficient (Karbasi, Schultze, Giacometti, Plautz, Schubert, Schulte, and Bashkirov 2015). Once the system is solved for the RSP values of each voxel in x, images can be constructed demonstrating the differences in RSP values for the entire scan volume.

CHAPTER TWO

Proton and Ion Therapy

A resurgence of interest in proton and ion radiation therapies led to a coinciding increase in efforts investigating clinical-grade proton Computed Tomography systems (Schulte, Penfold, Tafas, and Schubert 2008). Proton and ion radiation therapies involve bombarding an area of interest with protons or heavier ions in order to dose a tumor or other structure in the target area (Patyal 2007; Paganetti 2016). The major points of interest for proton and ion therapies when compared to more traditional photon radiation therapies are dosimetry and precision (Patyal 2007). Treatments employing protons or other ions can deliver the same dose to a tumor as those using photons, but at a much lower collateral dosage (Patyal 2007). However, this can only be achieved when treatment plans can be created with precision. The desire for such accurate treatment plans is the driving force behind the interest in the Relative Stopping Power maps proton Computed Tomography can provide (Penfold, Rosenfeld, Schulte, and Schubert 2009; Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016).

2.1 Advantages of Proton Therapy vs. Photon Therapy

The goal of radiation therapy treatments is to maximize the dose in the target area and minimize dose in other areas. The measure of how much of the total administered dose is in the target area is known as conformity. The primary advantage of proton therapy over photon therapy is the greater conformity of proton therapy (Patyal 2007). This allows proton therapy techniques to be as effective at treating tumors or other abnormalities without causing as much collateral damage. This reduction in damage to surrounding tissues is one of the primary motivators for the interest in proton therapy techniques (Patyal 2007). An examination of dose defini-

tions, the behavior which causes the conformity advantage, and the methods which can take advantage of it follows.

2.1.1 Doses and Conformity

To understand the advantages of the conformity difference, it is important to understand the terms involved. A dose is the amount of energy deposited in a volume. The energy deposited in the target area is known as the effective dose. This is the part of the total dose that actually affects the tumor or other targeted tissue. The portion of total energy deposited at a lesser depth than the target area is known as the entrance dose, while the portion of energy deposited at a greater depth than the target area is known as the exit dose. Together, these values make up the collateral dose, the total energy deposited outside the target area. The total dose is the value of all energy deposited in the entire volume. In other words:

$$D_{Total} = D_{Effective} + D_{Collateral} = D_{Effective} + D_{Entrance} + D_{Exit} \tag{2.1}$$

Conformity is a measurement of how much a particular treatment conforms to the dosage ideal (Petkovska, Tolevska, Kraleva, and Petreska 2010). Conformity Index, or CI is defined to be:

$$CI = \frac{V_{Reference}}{V_{Target}} \tag{2.2}$$

(Petkovska, Tolevska, Kraleva, and Petreska 2010). A value which also describes conformity in terms of dose may be easier to understand and is given in Equation 2.3.

$$DCR = \frac{D_{Effective}}{D_{Total}} \tag{2.3}$$

This value is referred to henceforth as the Dose Conformity Ratio, or DCR. From the DCR, a related equation for the Dose Non-conformity Ratio (DNR) is apparent, and outlined in Equation 2.4.

$$DNR = \frac{D_{Collateral}}{D_{Total}} = 1.0 - DCR \tag{2.4}$$

In an ideal treatment plan, $DCR = 1.0, DNR = 0.0$. This would mean that all of the energy was deposited in the target area, and no energy was deposited elsewhere. This is an impossibility, but generally, the closer to this ideal a treatment is, the less collateral damage will be inflicted while maintaining the necessary dose to be effective in the target area. Proton-based therapy has a better potential to approach this ideal than do photon-based therapies, and this is the major motivation behind research into proton therapies (Sadrozinski, Bashkirov, Keeney, Johnson, Peggs, Ross, Satogata, Schulte, Seiden, Shanazi, *et al.* 2004; Patyal 2007). The difference in potential DCR values is due to the differences in the dose values of the therapies.

2.1.2 Doses of Protons vs. Photons

The difference in the interactions of protons and photons leads to a difference in dose behavior. Recall that the definition of dose is the amount of energy deposited per volume. Because of the behavior which creates the Bragg peak, protons deposit much more of their energy towards the end of their path than the beginning, and eventually stop, depositing no more energy past their endpoint (Sadrozinski, Bashkirov, Keeney, Johnson, Peggs, Ross, Satogata, Schulte, Seiden, Shanazi, *et al.* 2004; Patyal 2007). Conversely, photons deposit most of their energy towards the beginning of their path, and do not stop (Sadrozinski, Bashkirov, Keeney, Johnson, Peggs, Ross, Satogata, Schulte, Seiden, Shanazi, *et al.* 2004). The differences in dose at different depths can be seen in Figure 2.1. The difference between the native proton beam (red) and the modified proton beam (blue) is one of design. The modified beam represents the results of changing the properties of the protons in the beam to create a Spread-out Bragg peak, a method used to ensure the whole target area receives sufficient energy (Patyal 2007). It can be clearly seen that the photon beam deposits a much larger portion of its energy outside the marked target area than does the modified proton beam. This behavior is why proton-based therapies possess a much higher

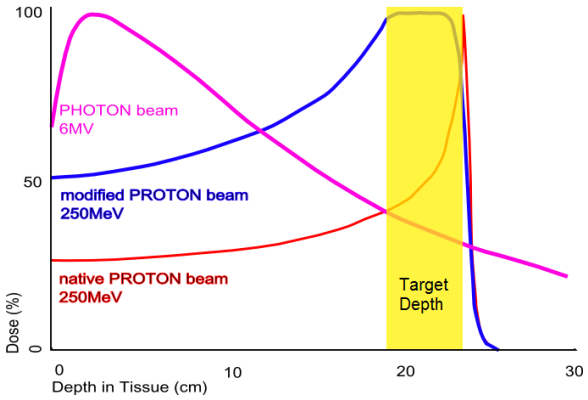

Figure 2.1: Dose as a function of depth for Photons vs. Protons. Note the extended Bragg peak on the modified Proton beam (Blue)

conformity than photon-based therapies (Patyal 2007). The almost non-existent exit dose of protons is one reason for the greater conformity of proton therapy (Patyal 2007). Another reason is the difference ratios of entrance dose and effective dose (Patyal 2007). Although in Figure 2.1 the apparent absolute entrance doses of the photon and modified proton methods are near identical, note the dose of each in the target depth. The photon method has a much lower dose in the target depth than the modified proton method. This means that even if the absolute entrance dosages are the same, the conformity of the modified proton beam is better than that of the photon beam. When the difference in exit dosage is considered, it becomes clear that proton methods are more conformal than photon methods.

2.1.3 Implications of Conformity

The conformity of a treatment method can be viewed as a measure of its safety efficiency. The greater the conformity, the less unintended damage inflicted while maintaining treatment efficacy, and therefore, the safer the treatment is for the patient (Patyal 2007). This can be easily derived from Equation 2.1. The inverse linear relationship between the Dose Conformity Ratio and the Dose Non-conformity Ratio

is outlined in Equation 2.5, which is obtained by dividing Equation 2.1 by D_{Total}, the total dose administered.

$$1.0 = DCR + DNR \tag{2.5}$$

Therefore, if the DCR increases, the DNR must decrease to maintain this relationship. As explained, proton therapy has a higher DCR than photon therapy. The greater conformity of the proton method can be exploited in two ways: maintaining effective dose while reducing collateral dose, or maintaining total dose while increasing effective dose (Paganetti 2016). Too low an effective dose is just as undesirable as too high a collateral dose, as each can have undesired effects (Patyal 2007). Treatment planning is therefore an optimization problem. Consider first the case where a target area must receive at least a certain dose to be effective, (i.e. $D_{Effective} \geq k$). By the definition of DCR (Equation 2.3), it is seen that for a fixed DCR, as required dose (k) increases, total dose must increase proportionally.This relationship is shown in Equation 2.6.

$$DCR \times D_{Total} = D_{Effective} \geq k \rightarrow D_{Total} \geq \frac{k}{DCR} \tag{2.6}$$

However, if the DCR is increased, the minimum D_{Total} for any given value of k decreases. Proton therapy's higher DCR allows the total dose administered to a patient to be lowered, while maintaining the effective dose administered. This in turn lowers the collateral dose without any loss of efficacy, and with it, the chance of complications. Conversely, consider the case where a patient must have a collateral dose no more than a certain amount (i.e. $D_{Collateral} \leq l$). When this required relationship is presented mathematically, Equation 2.7 is the result.

$$DNR \times D_{Total} = D_{Collateral} \leq l \rightarrow D_{Effective} \leq l \times \frac{DCR}{1 - DCR} \tag{2.7}$$

This implies that for a given DCR, any necessary reduction in l has a corresponding reduction in effective dose. For a greater DCR value, the effective dose for a given value of l is greater. This means that a greater amount of energy can be applied to the

target area without any increase in the collateral dose, which prevents a corresponding increase in likelihood of complications (Patyal 2007). These two relationships are the primary advantages of proton therapy over photon therapy. While the physics of protons make these relationships possible, accurate prediction models must be made to fully exploit them.

2.2 Relative Stopping Power

Accurate proton therapy treatment plans rest on the ability to predict and manipulate the location of the Bragg peak (Paganetti 2016). This can be done with knowledge of the stopping power of each portion of the target area and surrounding tissues (Penfold, Rosenfeld, Schulte, and Schubert 2009). As opposed to the density images yielded by x-ray CT, pCT produces images of the Relative Stopping Power (RSP), (Sadrozinski, Bashkirov, Keeney, Johnson, Peggs, Ross, Satogata, Schulte, Seiden, Shanazi, *et al.* 2004). This measurement is more useful in the planning of proton and ion therapy radiation treatments than the results of xCT scans (Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016). Previous proton therapy planning methods involved converting the Hounsfield value of a voxel in an xCT scan to the stopping power using a calibration curve (Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016). However, this led to errors of up to a few percent, due to the lack of an unique relationship between the Hounsfield value and stopping power (Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016). This imprecision demands a more direct method of measuring stopping power. pCT is that method.

2.2.1 *Definition*

Stopping power is a measurement of the energy loss of a proton caused by passage through a substance (Sadrozinski, Bashkirov, Keeney, Johnson, Peggs, Ross, Satogata, Schulte, Seiden, Shanazi, *et al.* 2004). This is expressed in terms of energy lost per distance, such as $\frac{MeV}{cm}$ (Paganetti 2016).

Relative Stopping Power compares the Stopping Power of a substance to that of water, represented as a ratio. For example, $\text{RSP}_{air} \approx 0.0013$, indicating that passage through air causes very little energy loss compared to traveling through an equal distance of water (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015). Conversely, $\text{RSP}_{Teflon} \approx 1.8280$; thus a proton loses nearly twice as much energy per distance traveled in Teflon as in water (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015). The RSP values produced by pCT are superior to the conversion of density readings of xCT when constructing treatment plans for proton and ion radiation therapies (Li, Liang, Singanallur, Satogata, Williams, and Schulte 2006; Schulte, Penfold, Tafas, and Schubert 2008; Penfold, Rosenfeld, Schulte, and Schubert 2009).

2.2.2 *Use in Treatment Planning*

When drafting a treatment plan for proton or ion therapy, it is necessary to accurately predict the location of the Bragg peak (Penfold, Rosenfeld, Schulte, and Schubert 2009). As the Bragg peak is the area of most dose, accurate predictions allow dose to be administered in a conformal manner, limiting the dose to non-targeted tissues (Patyal 2007). Determining a Bragg peak's location depends on finding the depth at which a proton will reach a critical energy, (and therefore velocity), depositing much more energy per volume passed through than at earlier points on its path, before coming to rest (Paganetti 2016). To find this depth, the RSP of the entire path followed by the proton must be found. This is accomplished by summing the

RSP of each voxel along the path, weighted by the intersection length of the path with that voxel. This method requires the RSP value of each voxel to be as accurate as possible.

Originally, RSP values of individual voxels were determined by conversion from the Hounsfield values from xCT systems (Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016). However, this resulted in uncertainties up to a few percent, leading to uncertainty in dosage values of treatment plans (Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016). As stated previously, this lack of precision is due to the lack of a unique relationship between the Hounsfield values of a substance and the RSP of that substance, and lead to reductions in achievable dose conformity (Plautz, Bashkirov, Giacometti, Hurley, Johnson, Piersimoni, Sadrozinski, Schulte, and Zatserklyaniy 2016). With the advent of pCT, it is possible to directly determine the RSP values, eliminating the uncertainty due to the conversion. The resulting improvement in RSP values directly enables better prediction of path RSP, and proton behavior within the subject. In turn, this improves the conformity of dose achievable during treatment, with benefits as outlined above.

2.3 Total RSP On-Path Evaluation

With the importance of accurate Relative Stopping Power values defined, the need for a method to determine the accuracy of RSP values is clear. Current methods focus on the comparison of mean and standard deviation for specific regions of interest to those of theoretical ideals (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015). However, the intended use of pCT reconstructed images to predict the RSP of proton paths inside the body for use in proton therapy treatment planning has lead to another method for error checking, though both systems complement each other, and will be most effective when used in tandem.

2.3.1 Measuring RSP Accuracy

One technique, used previously and also in this work, is to measure the mean of certain regions of interest, and compare the mean to a known value for the average RSP of that region (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015). This is a good benchmark for the general accuracy of the image reconstructed, because it will clearly show large discrepancies in value between reconstructed and actual. The standard deviation of each region of interest is also measured in current work as an indicator of consistency. If the mean of an area is near the true average, but standard deviation is large, any path prediction terminating inside that area will suffer in a way that paths passing through the entire region will not. The mean and standard deviation of regions of interest are important values for determining accuracy of pCT reconstructions. However, upon consideration of the methods employed when planning proton therapy treatment, another method has been employed.

2.3.2 TROPE Method

This method is referred to as Total RSP On-Path Error (TROPE). The principal purpose of TROPE is to find the total RSP of several paths leading to a voxel, and compare them to the values produced by theoretical ideal data. This process is then repeated for each voxel of interest. The version of TROPE used during this work tested the RSP of paths along each axis in three dimensions, and along all paths exactly 45° between axes. This was done for ease of implementation and visualization. Additionally, by searching a variety of paths in all 3 dimensions, error tendencies can be determined. Unfortunately, no ideal data exists at the time of this experiment, and the process of creating ideal data could not be fit into the time frame of the work. The direct relationship between the TROPE method and treatment planning allows the effects of error on planning to be more readily observed. It also will determine what kinds of errors have greater impact than others. Solutions for errors which strongly

affect treatment planning should be prioritized. A further algorithmic description of this method can be found in Appendix E.

CHAPTER THREE

Previous Work

As stated, the primary purpose of this work is to investigate the impact of different chord-length calculation methods on proton Computed Tomography reconstructions. One such study was conducted by Penfold *et al.* in 2009, which stated that effective mean chord length approach was the best method (Penfold, Rosenfeld, Schulte, and Schubert 2009). This led to effective mean chord length being adopted as the standard, and reconstructions using this technique have been proven to be sufficiently accurate for medical purposes (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015). In this work, a new method, referred to as mean chord length, is proposed. The results of reconstructions using mean chord length calculations are compared to those using the effective mean chord length approach.

3.1 Importance of Chord Length Values

When reconstructing images using data gathered for proton computed tomography, it is necessary to calculate both the A matrix, or system matrix, and the b vector, or total Relative Stopping Power vector. A is $n \times m$ where n is the number of proton histories considered, and m is the number of voxels in the volume. Thus, any element in A, a_j^i, represents the length of the intersection between the path of proton i and voxel j (Penfold, Rosenfeld, Schulte, and Schubert 2009). This value is referred to as a chord length. Similarly, b is a n-length vector, where each b_i corresponds to the total RSP of the path of the i^{th} proton. x is not known at the beginning of reconstruction, and the goal of reconstruction is to accurately predict the values of x. x is a m length vector containing the RSP values of each voxel. This relationship, though simple to describe, is outlined in Equation 3.1.

$$Ax = b \tag{3.1}$$

Because Equation 3.1 must be solved for x accurately, it can be seen that errors in the A matrix will lead to errors in the estimation of the x vector. Errors in b will also cause inaccuracies in x, but the errors in b are a function of equipment sensitivity, not reconstruction methods, and are therefore not considered in this work.

3.1.1 Impact of Errors

The method by which error in the system matrix A causes error in the image vector x can be seen by example. Consider the (trivial) case where $n = 1$, that is, only one proton history is considered when reconstructing the image. In this case, Equation 3.1 reduces to Equation 3.2.

$$\sum_{j=1}^{m} a_j x_j = b \rightarrow x_k = b\frac{a_k}{\sum_{j=1}^{m} a_j}, k = 1, 2, 3, ..., m-1, m \tag{3.2}$$

It is clear that even in this simple case, error in a particular a_j affects all values in x which are on the path, not only the case where $k = j$. Error only affects values of x on the path because if $a_k = 0$, $x_k = 0$ regardless of other values. For example, underestimation of a given a_k will cause the corresponding x_k to be underestimated, and all other values of x_j to be overestimated. If, however, all values of a_j were systematically underestimated, all values of x_j would be overestimated. The small reduction in x_k due to the underestimation of a_k would be overshadowed by the increase in x_k due to all other values of $a_{j\neq k}$ being underestimated. The inverse effects occur for overestimation of a_j values. It can be extrapolated that in the much more complex case of pCT reconstruction, where n is on the order of hundreds of millions, and m is on the order of two million, errors in A will affect a large magnitude of values.

3.1.2 Systemic Nature of Errors

Because of the large number of values involved in pCT reconstruction, an isolated error in a single a_j^i would have a negligible effect on the x vector, especially if the

error is small. However, the values in A are not determined experimentally, but rather by a combination of the Most Likely Path technique, and the chosen Chord Length calculation method. Because both MLP and Chord Length calculation are estimation methods, they will always produce some error, and the errors produced have patterns. For example, if a constant chord length is used, such as the width of a voxel, errors in a_j^i will be largest when the angle of the path i relative to the reconstruction grid is closest to 45°. A Chord Length calculation technique may attempt to correct for this error, but in that case, the technique must depend on the angle of the path, causing any errors at a given angle to be perpetuated for all paths with that same angle. This means that any error in the Chord Length calculation will affect a large number of paths, and thus a large amount of a_j^i values. This consistency means that errors will not occur only in isolated a_j^i values, and therefore will not be inconsequential. Errors in MLP produce errors similar to those produced by inaccuracies in Chord Length calculations, but with a potential for a greater magnitude of error, as MLP determines if a particular a_j^i is zero or non-zero. If an MLP method determines that an a_j^i is zero when it should posses a value, the effect is as though a Chord Length calculation technique had an error equal to the entire expected value of a_j^i. The same holds for elements determined by MLP techniques to be non-zero when they should be zero. A simple case of a systemic error in Chord Length calculation would be a Chord Length method that always underestimated the length of intersection between voxel j and path i. Any reconstruction using such a method would consistently overestimate the RSP values of the reconstruction volume. This is because it would appear that all protons lost their energy over a smaller distance than actually traveled.

3.2 Previous Chord Length Calculation Methods

A variety of Chord Length Calculation methods have been employed for pCT. In 2009, the values of different methods were investigated, and Effective Mean Chord

Length was selected as the best available with respect to efficiency of accuracy and computation time (Penfold, Rosenfeld, Schulte, and Schubert 2009). Although some of these methods have been dismissed as less useful than Effective Mean Chord Length, and are not compared in the experiment for this work, an overview of the different possible methods is given here, to provide context. Those wishing a more thorough review of methods not under consideration in this work are encouraged to read the 2009 article by Penfold *et al.* (Penfold, Rosenfeld, Schulte, and Schubert 2009).

3.2.1 Constant Chord Length

The simplest possible chord length calculation, Constant Chord Length determines which voxels are intersected by the given proton path using Most Likely Path methods. Each of these voxels would then be assigned a constant value, selected at the beginning of reconstruction, while all other voxels would retain a value of 0 (Penfold, Rosenfeld, Schulte, and Schubert 2009). In the case examined in the 2009 article, this length was taken to be the width of a voxel in the XY plane (Penfold, Rosenfeld, Schulte, and Schubert 2009). This is outlined in Equation 3.3, where l is taken to be the width of a voxel in the x or y directions. This assumes a square construction grid in the XY plane.

$$a_j^i = \begin{cases} l & \text{if voxel } j \text{ is on path } i, \\ 0 & else. \end{cases} \tag{3.3}$$

Constant chord length is very fast, as no additional calculations are necessary beyond those required for MLP determination. However, this method was dismissed because of consistent underestimations of RSP, up to 13% when using Algebraic Reconstruction Technique (Penfold, Rosenfeld, Schulte, and Schubert 2009).

3.2.2 Exact Chord Length

The Exact Chord Length method was proposed to eliminate the error inherent in the Constant Chord length approach. Exact Chord Length involved determining the length of intersection between a path and voxel at each step of the Most Likely Path process (Penfold, Rosenfeld, Schulte, and Schubert 2009). This led to the possibility of updating three a_j^i values simultaneously, for every step of the MLP process for every proton path (Penfold, Rosenfeld, Schulte, and Schubert 2009). The details of the Exact Chord Length method can be found in Penfold *et al.*, but a simple description is provided (Penfold, Rosenfeld, Schulte, and Schubert 2009). In essence, for each MLP step, the Exact Chord Length divided the length of an MLP step among up to three voxels, adding those values to the a_j^i of each of those voxels for the given path (Penfold, Rosenfeld, Schulte, and Schubert 2009). If the step remained in the same voxel as it started, the entire step length is added to that voxel's chord length. If the step went to an orthogonally adjacent voxel, the path was split proportionally between those two voxels. If the step went to a diagonally adjacent voxel, the portion of the step in starting and ending voxels was added to those voxels, with the remainder of the step length added to the intermediate voxel which the path must have passed through, but was not indexed by the step. Ultimately, though the Exact Chord Length method provided accurate RSP values for voxels, it was dismissed because of the approximately 20% increase in construction time per iteration when compared to either Constant Chord Length or Effective Chord Length approaches (Penfold, Rosenfeld, Schulte, and Schubert 2009).

3.2.3 Effective Mean Chord Length

Effective Mean Chord Length, as outlined by Penfold *et al.*, was designed to address the shortcomings of Constant Chord Length approaches, without being as computationally demanding as the Exact Chord Length method (Penfold, Rosenfeld,

Schulte, and Schubert 2009). The principle of Effective Chord Length is that for a given path, all voxels can be assigned a single chord length, which will represent the mean chord length along the path, including any errors due to overstepping (Penfold, Rosenfeld, Schulte, and Schubert 2009). The formula for Effective Mean Chord Length is presented in Equation 3.4, though the full derivation of this formula is not presented in this paper. It can be found in the 2009 article by Penfold *et al.* (Penfold, Rosenfeld, Schulte, and Schubert 2009).

$$\bar{\Delta}(\theta)_{\text{eff}} = \frac{l}{3}\left(\frac{(s/l)^2 \sin 2\theta - 6}{(s/l)\sin 2\theta - 2(\cos\theta + \sin\theta)} + \frac{(s/l)^2 \sin 2\theta}{2(\cos\theta + \sin\theta)}\right), \theta \in [0, \frac{\pi}{2}] \qquad (3.4)$$

In this case, l is the width of a voxel in the x and y directions, s is the step size of the MLP procedure, and θ is the average angle of the path relative to the reconstruction grid in the XY plane (Penfold, Rosenfeld, Schulte, and Schubert 2009). The values of l and s are constant throughout a given reconstruction. This version of the Effective Mean Chord Length formula assumes X and Y voxel dimensions are identical. This squareness allows the value of Effective Mean Chord Length to vary only based on θ, and to restrict the values of θ input to the range $\theta = [0°, 90°]$. Effective Mean Chord Length was found to be an efficient means for estimating values of a_j^i, as it was nearly as accurate as the Exact Chord Length method, but took 20% less time per iteration (Penfold, Rosenfeld, Schulte, and Schubert 2009).

3.2.4 Updated Effective Mean Chord Length

Further investigation into the Effective Mean Chord Length method led to an updated version of Effective Mean Chord Length, which accounts for variation in the path angle relative to the xy plane. Using a slight variation of the formula outlined in Equation 3.4, the change in the distance due to variation in the z direction is calculated. This factor is then multiplied by the xy plane Effective Mean Chord Length to produce a value accounting for variation in the z dimension. This updated version was derived by Scott Penfold, and has been in consistent use in reconstructions

since 2010 (Penfold 2010). As such, it is examined alongside other chord length methods during the experimental portion of this work. The formula for Updated Effective Mean Chord Length is given in Equation 3.5.

$$\bar{\Delta}(\theta,\phi)_{\text{u-eff}} = \frac{\bar{\Delta}(\theta)_{\text{eff}}}{3}\left(\frac{(s/l_m)^2\sin 2\phi - 6\frac{h}{l_m}}{(s/l_m)\sin 2\theta - 2(\frac{h}{l_m}\cos\phi + \sin\phi)} + \frac{(s/l_m)^2\sin 2\phi}{2(\frac{h}{l_m}\cos\phi + \sin\phi)}\right) \tag{3.5}$$

The value of l_m is given by Equation 3.6. It is referred to as the "Mean Pixel Width"

$$l = \frac{1}{\cos\theta + \sin\theta} \tag{3.6}$$

3.3 State of Chord Length Methods

The findings of the 2009 article by Penfold led to the adoption of Effective Mean Chord Length as the standard method used during pCT reconstructions. Its accuracy is comparable to that of the Exact Chord Length method, but its computation time is similar to that of the Constant Chord Length method, making it a good compromise between efficiency and accuracy (Penfold, Rosenfeld, Schulte, and Schubert 2009; Penfold and Censor 2015). Minimal investigation into alternative methods has been performed since 2009 *et al.* (Penfold and Censor 2015). As a result, Effective Mean Chord Length has remained the default method for pCT reconstructions since that time. This work aims to determine the best chord length method for proton computed tomography reconstructions.

CHAPTER FOUR

Design and Implementation

The purpose of this work is to evaluate different methods for determining chord lengths. As such, there must be new methods to test that were not tested previously. The primary objective of any new method is to reduce the number and complexity of calculations required, without any loss of accuracy. Efforts in this direction lead to the derivation of True Mean Chord Length, Simplified Mean Chord Length, and Approximate Mean Chord Length. These three methods were used for pCT reconstructions. The reconstructions were tested for accuracy and compared to results of reconstructions employing Effective Mean Chord Length.

4.1 True Mean Chord Length

Though Effective Mean Chord Length has been found to be a sufficient method for determining the chord length values required for pCT reconstruction, the possibility for a more accurate method exists. Because of this, it is possible to derive a new method for determining chord length using definitions of expected value. This new formula is referred to as the True Mean Chord Length.

4.1.1 Definition and Origin

The True Mean Chord Length endeavors to find a formula to find the expected value of a chord through a voxel at given traversal angles. Starting with the Pythagorean Formula, and employing trigonometric identities, an expression for True Mean Chord Length can be determined. The final formula for True Mean Chord Length can be found in Equation 4.1.

$$\bar{\Delta}(\theta,\phi)_{mean} = l\frac{\sec\phi\csc(\theta+\frac{\pi}{4})}{\sqrt{2}}, \theta \in [0,\frac{\pi}{2}] \tag{4.1}$$

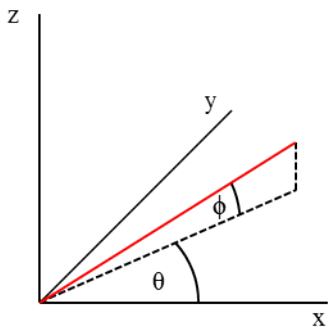

Figure 4.1: Definitions of θ and ϕ

In this case, θ is determined in the same way as for Effective Chord Length (Equation 4.2), while ϕ is the angle of the proton path with the xy plane. This can be found by taking the inverse tangent of the z distance over the distance in the xy plane (Equation 4.3). See Figure 4.1 for clarity. l is the voxel width set for the reconstruction.

$$\theta = \tan^{-1}\left(\frac{|y_{out} - y_{in}|}{|x_{out} - x_{in}|}\right) \tag{4.2}$$

$$\phi = \tan^{-1}\left(\frac{|z_{out} - z_{in}|}{\sqrt{(y_{out} - y_{in})^2 + (x_{out} - x_{in})^2}}\right) \tag{4.3}$$

This value is an expression of the expected value of chords traversing a voxel in the xy plane, and treating θ as a given value. An as-yet unpublished study by Paniz Karbasi demonstrated that greater than 96% of proton paths deviate two voxels or less in the z-direction over the entire path (Data Pending Publication). This means that in greater than 96% of cases, a proton will enter and exit a voxel in the same xy plane as the voxel. Therefore, because the z component of distance will almost never cause the proton to exit the voxel in the z direction, the distance in the xy plane need only be augmented by a factor of $\sec\phi$ to yield the entire distance in that voxel. This assumption is key, as it prevents the math required for determining True Mean Chord Length from becoming prohibitive. As l is a constant for the entirety

of any given reconstruction, $\bar{\Delta}_{mean}$ only varies with the θ and ϕ of a given path. See Appendix A for a full derivation of the True Mean Chord Length formula.

4.1.2 Implementation

The True Mean Chord Length formula was coded to be compatible with the reconstruction code currently in use at Baylor University. Like Effective Mean Chord Length, the True Mean Chord Length is computed only once for each path. pCT reconstruction methods at Baylor University make use of Graphics Processing Units in order to perform large sections of the reconstruction in parallel. This parallelism leads to large speed increases over serial reconstruction methods. Previous chord length methods, such as Constant Chord Length and Effective Mean Chord Length were programmed in such a way that they could be computed on Graphics Processing Units in order to allow for parallelization. The function which computes True Mean Chord Length was designed to be compatible with the parallel GPU setup. This allows the reconstruction process to retain its parallelism, and the speed advantages which come with it. The full C++ code for the True Mean Chord Length may be found in Appendix A.

4.2 Simplified Mean Chord Length

Although the True Mean Chord Length method is mathematically accurate, reductions in the complexity of the formula employed were desired. This could potentially lead to reductions in computation time, without a meaningful loss of accuracy, a consistent goal in pCT reconstruction research. Known factors about the proton paths present in pCT scans allow for an assumption to be made which can reduce the complexity of the True Mean Chord Length formula. This leads to the Simplified Mean Chord Length.

4.2.1 *Definition and Derivation*

In the case of True Mean Chord Length, the term $\sec\phi \approx 1$ for greater than 96% of paths. This can be deduced from the same observation that in greater than 96% of cases, proton paths deviate less than 2 voxels in the z direction. Because the z distance is much smaller than the xy distance, it can be deduced that $\phi \cong 0$. When ϕ is taken to be zero, Equation 4.1 reduces to Equation 4.4. Appendix A outlines the exact math behind this reduction in complexity.

$$\bar{\Delta}(\theta)_{s-mean} = l\frac{\csc(\theta + \frac{\pi}{4})}{\sqrt{2}}, \theta = [0, \frac{\pi}{2}] \tag{4.4}$$

In pCT scans, θ varies greatly on the interval $[0, \frac{\pi}{2}]$, and therefore cannot be ignored, meaning that Equation 4.4 is the most simple form possible for Simplified Mean Chord Length. In this case, as in Effective Mean Chord Length, for any given reconstruction, the value of $\bar{\Delta}_{s-mean}$ only varies based on the value of θ. This is unlike the True Mean Chord Length, which depends on both θ and ϕ.

4.2.2 *Implementation*

Like True Mean Chord Length, Simplified Mean Chord Length is calculated only once per path, and is coded such that it may be calculated on a GPU. This allows pCT reconstructions using this method to retain all the parallelism currently employed in the methods used at Baylor University. Although the final form of Simplified Mean Chord Length appears minimally different from that of True Mean Chord Length, there are significant differences in implementation. Equation 4.4 differs from Equation 4.1 only by the absence of the $\sec\phi$ term. However, there are a large number of operations present in the determination of this term. Although $\sec\phi$ is a single calculation, it is seen in Equation 4.3 that determining ϕ itself involves many operations. By assuming $\phi \approx 0$, all these calculations can be bypassed. Calculating θ requires a slightly fewer operations as ϕ. Therefore, removing the calculations

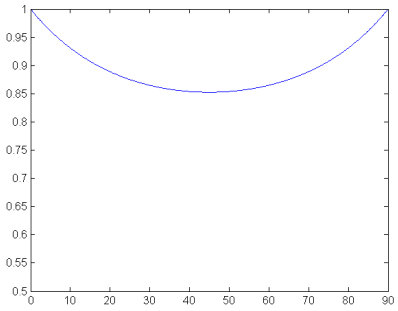

(a) Effective Mean Chord Length

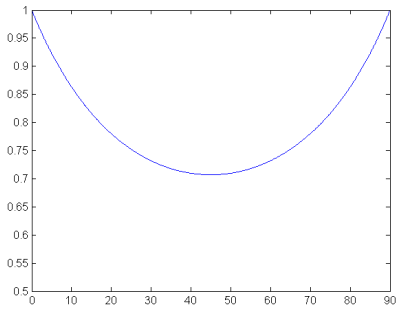

(b) Simplified Mean Chord Length

Figure 4.2: Effective and Simplified Mean Chord Length vs. θ, where $l = 1.0$

related to ϕ reduces the total operations required by roughly one-half. The C++ code implementing Simplified Mean Chord Length can be found in Appendix A.

4.3 Approximate Mean Chord Length

Trigonometric functions are computationally expensive compared to other operations such as multiplication and addition. As such, methods which employ more simple operations have the potential to run significantly faster than those using trigonometric operations. Therefore, it was desired to find a method which accurately produced chord lengths with a minimum of trigonometric functions. These efforts led to the Approximate Mean Chord Length.

4.3.1 Definition and Derivation

Plots of both Simplified Mean Chord Length and Effective Chord Length show a shape very similar to a parabola. This can be seen in Figures 4.2. In both cases, the maximum value of $\Delta_1$occurs at $\theta = 0, \frac{\pi}{2}$, and the minimum value occurs at $\theta = \frac{\pi}{4}$. Therefore, the equation for Approximate Mean Chord Length may be assigned the form shown in Equation 4.5.

$$\bar{\Delta}(\theta)_{approx} = A\left(\theta - \frac{\pi}{4}\right)^2 + b \tag{4.5}$$

A and b are functions of both the dimensions of the voxels employed in a reconstruction, and the step size used in the MLP procedure for that reconstruction. However, A and b are fixed values for a given reconstruction, and can therefore be calculated once at the beginning of the reconstruction process. Equations 4.6 and 4.7 give definitions for b and A respectively.

$$b = \frac{l}{3}\frac{24 - \sqrt{2}r^3}{8\sqrt{2} - 4r}, r = \frac{s}{l} \tag{4.6}$$

$$A = l\frac{16(1-b)}{\pi^2} \tag{4.7}$$

As a reminder, l is the width of a voxel in the xy plane, and s is the step size chosen for the MLP procedure used for this reconstruction. After the initial determination of A and b, the number of calculations required to determine $\bar{\Delta}(\theta)_{approx}$ for any given path is minimal.

4.3.2 Implementation

As with the other methods, Approximate Mean Chord Length is implemented such that it can be calculated on a GPU. This preserves parallelism in the reconstruction process. The values of A and b are calculated once at the beginning of the reconstruction process, while the value of $\bar{\Delta}(\theta)_{approx}$ is calculated once per path. The full code for the Approximate Mean Chord Length method can be found in Appendix D.

4.4 Experimental Setup

First, code for True Mean Chord Length and Approximate Mean Chord Length was implemented. Then, reconstructions employing each of the different chord length methods were performed for simulated data. These methods include the previously determined Effective method, the updated Effective method, the new True Mean method, Simplified Mean, and two versions of the Approximate Chord Length method. One Approximate method estimated the Effective Mean Chord Length

method, and the other targeted the Simplified Mean Chord Length method. Selected Regions of Interest were evaluated for the reconstructions, and the mean Relative Stopping Power of those regions compared to the ideal. Changes in RSP value error due to size of ROI or true RSP value were also investigated. Finally, the difference between total stopping power along paths to selected voxels were compared using the TROPE method.

4.4.1 Reconstruction Methods

Other than variations in the formulas used for chord length, all reconstructions in the experiment were performed identically. The steps are listed here in brief, but the details of the reconstruction method can be found in the 2015 article by Schultze *et al.*, and the steps listed here are drawn from that description (Schultze, Karbasi, Giacometti, Plautz, Schubert, and Schulte 2015).

First, proton histories are split into two groups: those possessing a Water Equivalent Path Length (WEPL) equal to or less than 0, and those which do not. The protons with WEPL at or less than 0 are used to build a hull describing the object. Voxels inside the hull will be updated by the reconstruction, while those outside will not. In parallel with this process, proton histories that were statistical outliers are removed from further consideration. Using the remaining proton histories, a simple Filtered Back Projection is performed to provide the initial image which will be modified via the iterative method selected.

Then, the system matrix A must be defined. Recall that the image reconstruction depends on finding an approximate solution to the system $Ax = b$. It is during this step that the various chord length calculation methods will change the values in A. A given a_j^i will be either zero or non-zero for all reconstructions, although the specific non-zero value will differ based on the chord length method used. More

details about the creation of the A matrix can be found in the paper by Penfold *et al.* (Penfold, Rosenfeld, Schulte, and Schubert 2009).

Once the A matrix is generated, the A and b terms in the reconstruction equation are defined. The initial iterate provides a starting point for the x vector. Once this equation is set up, iterative techniques are employed to find a solution for x, the image which produced the values observed. In this experiment, all reconstructions employed the Robust Diagonally Relaxed Orthogonal Projections (DROP) method, which has been shown to provide images of equal or greater quality to those produced by the standard DROP technique (Karbasi, Schultze, Giacometti, Plautz, Schubert, Schulte, and Bashkirov 2015). Detailed descriptions of this Robust DROP method can be found in the 2015 article by Karbasi (Karbasi, Schultze, Giacometti, Plautz, Schubert, Schulte, and Bashkirov 2015). For the reconstructions in this experiment, the block size was set to 3200 proton histories, and the relaxation parameter was set to $\lambda = 0.000150$. The images reconstructed have dimensions of $240 \times 240 \times 32$ voxels, in the x, y, and z dimensions, respectively. The voxel dimensions are 1mm×1mm×2.5mm. Selected images from these reconstructions are presented in Chapter Five.

4.4.2 Reconstruction Data

Reconstructions were performed using simulated data. In all cases, the subject of the scans was the Catphan ® CTP404 Phantom, henceforth "CTP404". This phantom was selected because it is well suited to test the validity of pCT reconstructions. Additionally, proton scan data of this phantom had previously been collected for a simulated situation. The CTP404 phantom is cylindrical, with a diameter of 150mm, with sections of known RSP values inserted throughout. A diagram of the CTP404 can be found in Figure 4.3. The inner circles are Acrylic spheres with diameters of 2, 4, 6, 8, and 10mm. The outer circles represent cylindrical inserts with

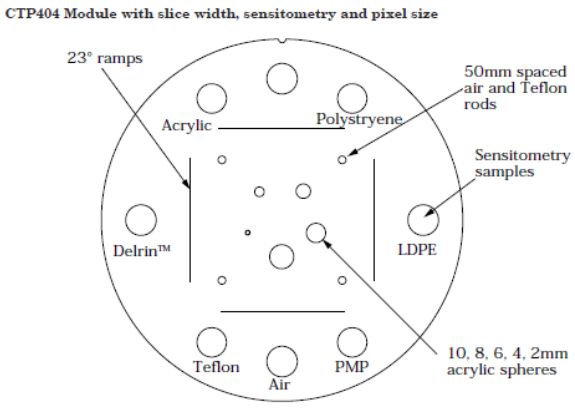

Figure 4.3: A diagram of the Catphan ® CTP404 Sensitom Phantom.

Table 4.1: Expected RSP values of materials

Material Insert	RSP Value
Air	0.0013
Polystyrene	1.0386
LDPE	0.9973
PMP	0.8770
Teflon	1.8280
Delrin	1.3560
Acrylic	1.1600

diameters of 12mm. Each insert is composed of a known material, which has a known Relative Stopping Power. The RSP values of each insert are listed in Table 4.1. The materials are in order, starting with the topmost air insert, and rotating clockwise around the outer circle.

The data used for this reconstruction was generated by a simulated scan of the CTP404 phantom. The simulation was performed in Geant4, using 4° changes in projection angle. The scan resulted in approximately 381.5 million proton histories which were used in the reconstruction.

CHAPTER FIVE

Experimental Results and Discussion

Image analysis was performed for reconstructions produced by the four new chord length methods and the existing Effective Mean Chord Length method. Mean RSP values and standard deviations were compared for selected Regions of Interest in reconstructed images. The changes were compared across chord length methods, RSP values, and ROI size. Additionally, path comparisons were performed for selected voxels in the reconstructed volume using the TROPE method. Results indicate that Updated Effective Mean Chord Length is the most accurate reconstruction method, though Approximate Effective Mean Chord Length closely approximates Effective Mean Chord Length, and appears worthy of further investigation.

5.1 Image Analysis

Analysis of the reconstructed images was performed using both ImageJ and the TROPE method as implemented in MATLAB®. Different Regions of Interest were defined, and their mean and standard deviations compared across reconstruction methods. Each of the eight material inserts were analyzed, including those with air. For each insert, four concentric circular Regions of Interest (ROI) were defined, with diameters of 12, 10, 8, and 6mm. Additionally, ROIs were defined targeting the four largest acrylic spheres. The mean RSP and standard deviation of each of these regions was determined. The slice selected for the image analysis was slice 16 out of 32. This slice was selected because it is near to the center of the Acrylic spheres inside the phantom, allowing them to be examined, and adding complexity to the image to be reconstructed. The resulting images from all reconstructions can be found in Figure 5.1.

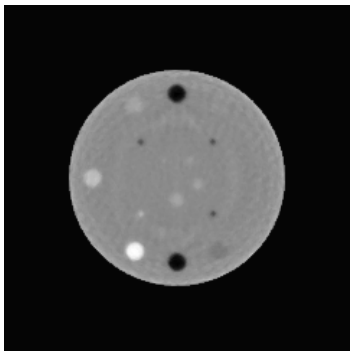

(a) Effective Mean Chord Length

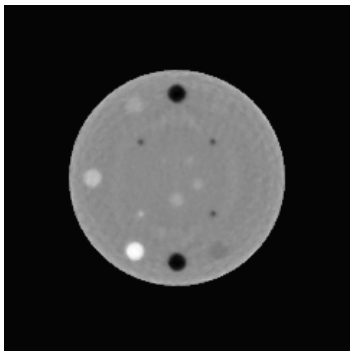

(b) Updated Effective Mean Chord Length

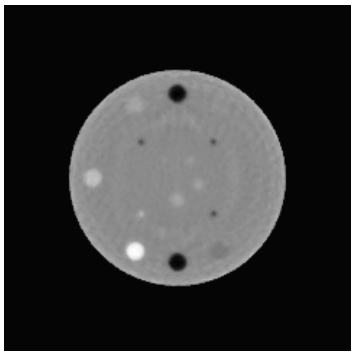

(c) Approximate Effective Mean Chord Length

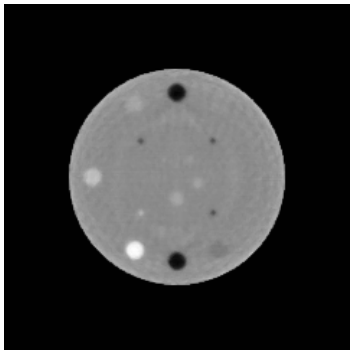

(d) True Mean Chord Length

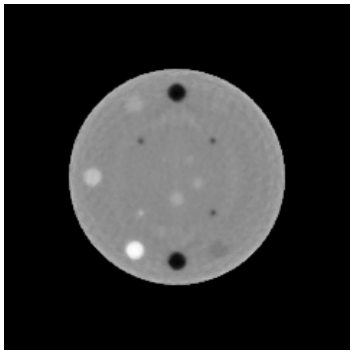

(e) Simplified Mean Chord Length

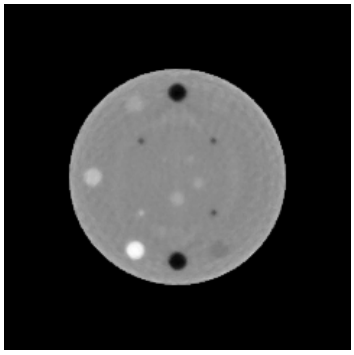

(f) Approximate Simplified Mean Chord Length

Figure 5.1: Reconstructions Produced by various Chord Length Methods

Table 5.1: Expected and Reconstructed RSP Values by Method

Method	Air	PMP	LDPE	Polystyrene	Acrylic	Delrin	Teflon
Expected	0.0013	0.8770	0.9973	1.0386	1.1600	1.3560	1.8280
Effective	0.1800	0.9020	0.9970	1.0340	1.1390	1.3120	1.7040
Updated Eff.	0.1800	0.9030	0.9980	1.0350	1.1400	1.3130	1.7070
App. Eff.	0.1760	0.8960	0.9900	1.0280	1.1320	1.3040	1.6960
True	0.2600	1.0170	1.1270	1.1640	1.2780	1.4610	1.8790
Simplified	0.2570	1.0300	1.1350	1.1740	1.2860	1.4750	1.8970
Appr. Simple	0.2480	1.0160	1.1200	1.1590	1.2710	1.4560	1.8770

5.1.1 Relative Stopping Power Values

After examination, it was clear that the Updated Effective Chord Length was the best method tested for acquiring accurate RSP values. Traditional Effective Chord Length and Approximated Effective Chord Length were also very accurate, and performed at very nearly the level of Updated Effective Chord Length. True Mean Chord Length, Simplified Mean Chord Length, and Approximated Simplified Mean Chord Length all had distinctly less accuracy. The RSP values generated by these three methods were very near to each other. Table 5.1 presents the mean RSP values obtained over the area of the entire insert for each material, while Table 5.2 presents the percent discrepancy of each case, organized similarly. This data corresponds to the data collected in the ROI with a 12mm diameter. The values for only one air insert are shown.

The continuing tendency of iterative construction techniques to underestimate values above 1.0 and overestimate values below 1.0 can be seen easily from the data in Table 5.1. Methods derived from the True Mean Chord Length method tend to give higher overall estimates than those derived from the Effective Chord Length method. However, impact of this overestimation on the accuracies of these two groups is more easily seen when presented with a percent discrepancy method. The results from the air inserts are not included in percent discrepancy because the small values

Table 5.2: Percent Discrepancy RSP Values by Method

Method	PMP	LDPE	Polystyrene	Acrylic	Delrin	Teflon
Effective	2.850%	-0.030%	-0.440%	-1.810%	-3.240%	-6.780%
Updated Eff.	2.960%	0.070%	-0.350%	-1.724%	-3.170%	-6.620%
App. Eff.	2.170%	-0.730%	-1.020%	-2.414%	-3.830%	-7.220%
True	15.960%	13.010%	12.070%	10.172%	7.740%	2.790%
Simplified	17.450%	13.810%	13.040%	10.862%	8.780%	3.770%
App. Simple	15.850%	12.300%	11.590%	9.569%	7.370%	2.680%

involved led to percent discrepancies in the thousands of percent. The accuracy of methods can be easily compared without the values for air. The estimation methods clearly fall into two groups: those employing or approximating the Effective Mean Chord Length methods, and those employing or approximating the True Mean Chord Length methods. Estimations using methods based on Effective Mean Chord Length are significantly more accurate than those based on True Mean Chord Length.

Accuracy of RSP values reconstructed is important, but also of note is the standard deviation of a ROI. A region may have an accurate mean, but be very noisy. Areas with smaller standard deviation are considered to be less noisy. Thus, Table 5.3 compares the standard deviation of the methods, by region. Methods derived from Effective Mean Chord Length again outperform those based on True Mean Chord Length, having equal or smaller standard deviations in every case. One trend which may be of importance is that the Approximate methods have equal or lesser standard deviation than the corresponding methods they estimate. Therefore, the images produced by Approximate methods are less noisy than their non-estimated counterparts.

5.1.2 Variation within Region of Interest

After analysis of entire inserts, analysis was also performed on concentric subsections of the inserts. This was done to gain information about how RSP estimations

Table 5.3: Standard Deviation of Reconstructed RSP Values by Method

Method	Air	PMP	LDPE	Polystyrene	Acrylic	Delrin	Teflon
Effective	0.163	0.030	0.031	0.022	0.037	0.069	0.121
Updated Eff.	0.164	0.030	0.030	0.022	0.037	0.069	0.121
App. Eff.	0.163	0.030	0.030	0.022	0.037	0.068	0.121
True	0.170	0.032	0.033	0.023	0.038	0.074	0.126
Simplified	0.173	0.030	0.031	0.023	0.038	0.073	0.129
App. Simple	0.172	0.030	0.031	0.023	0.038	0.072	0.129

vary within the same ROI. Figure 5.2 plots the percent discrepancies between expected and reconstructed RSP values for each of the six investigated methods. It is seen that for most materials, as the portion of the area examined increases towards total, the magnitude of percent discrepancy also increases. This means that values are more accurate towards the center of a given region. Additionally, examination of the data shows that the larger the portion of the ROI considered, the more that the entire area will tend towards 1.0. It can be seen from the graph that no particular chord length method produces results with better behavior in this regard, as the shapes of all lines are nearly identical, despite the offset of the True Mean Chord Length based methods.

5.1.3 Variation with Relative Stopping Power

Investigation was performed into how accuracy varies with the expected Relative Stopping Power. Figure 5.2 shows an oft-noted trend in pCT reconstructions: percent discrepancy increases as expected RSP values grow further from 1.0. Also, the effects of changes in proportional region size increase as RSP values grow further from 1.0. Again, no particular chord length method performs better or worse in this regard.

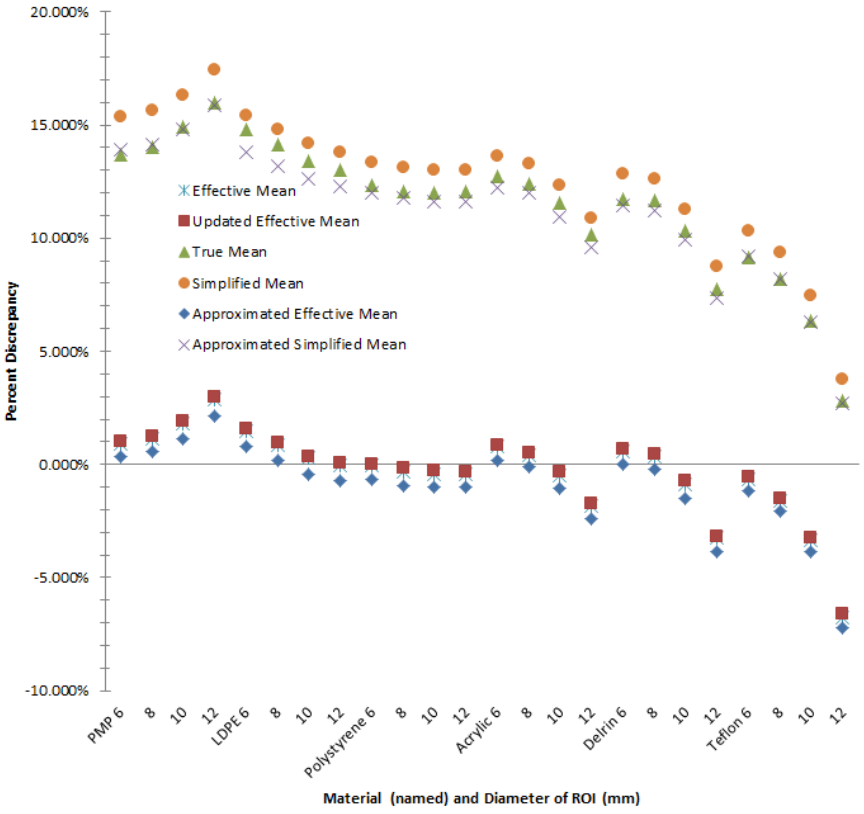

Figure 5.2: Percent Discrepancies. Sorted by true RSP value and size of region.

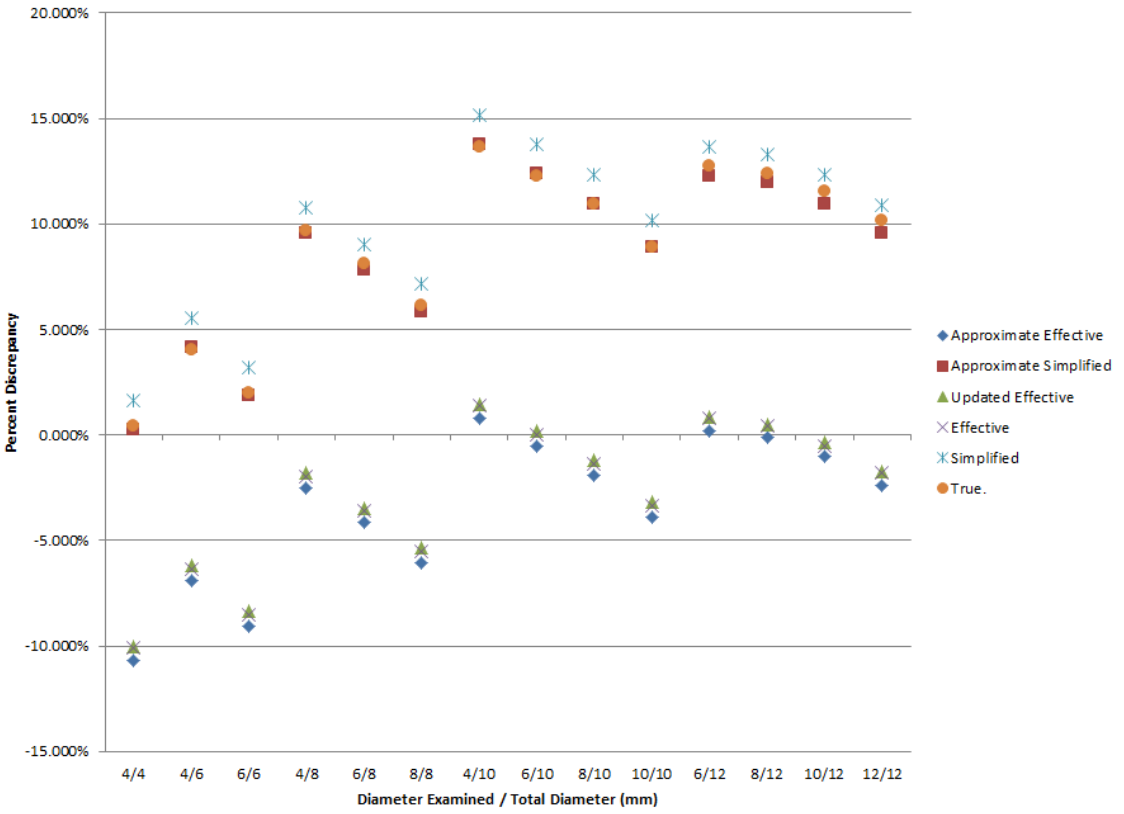

Figure 5.3: Variation of Percent Discrepancy by Size of Region of Interest

5.1.4 Variation with Size of Region of Interest

Another characteristic of the reconstructions examined was how accuracy and consistency changed with respect to the toal size of the ROI under consideration. To accomplish this, measurements of the cross-sections of the acrylic spheres in the CTP404 were taken, in addition to the acrylic cylinder. The measurements varied in size from 4mm up to the full diameter of the insert measured. Then, percent discrepancy was calculated for each of those measurements. The results are shown in Figure 5.3. First, note that smaller regions are estimated to have lower RSP values. It can also be seen that smaller regions have less internal consistency than larger regions. This can be inferred from the slope of each group of measurements corresponding to a given insert. Note that although there is a smaller percent discrepancy indicated for the True Mean Chord Length family of methods when considering smaller ROIs,

this is not an indicator of any particular accuracy in small regions inherent in these methods. It is simply due to the consistent overestimation of those methods as a whole, and the previously noted general tendency for smaller regions to have lower reconstructed RSP values.

5.2 Path Analysis

For each of the Regions of Interest investigated, the RSP of paths leading to voxels near the center of the region was calculated. This is similar to the calculations which would be performed when planning ion therapy treatments, as discussed in Chapter Two. An example of the values obtained can be found in Table 5.4. After this, a modified version of the TROPE method was run, which tests only the voxels selected, rather than testing every voxel in the volume. Additionally, for this specific implementation, an ideal for comparison was not available. Therefore the path RSP is found for the reconstruction, but there is no ideal against which to compare the path RSPs. The voxel under consideration in this example is located in the middle of the Teflon insert, in slice 16. Again the results are in two distinct groups: those methods based on the Effective Mean Chord Length, and those methods based on the True Mean Chord Length. The tendency to overestimate individual voxel RSP values present in the True Mean family of methods causes distinct overestimation of path RSP. Note that the values produced by Effective Chord Length, Updated Effective Chord Length, and Approximate Effective Chord Length never differ by more than 0.01, usually less than 0.005. The tendency of Approximate methods to underestimate voxel RSP values compared to the methods they approximate is reflected in the consistent underestimation of path RSP values when compared to the same methods.

Table 5.4: Path Relative Stopping Powers to Center of Teflon Insert

Vector	Effective	Updated Eff.	App. Eff.	True	Simplified	App. Simple
(-1,-1,-1)	1.3876	1.3893	1.3807	1.5340	1.5423	1.5257
(-1,-1,0)	1.0415	1.0429	1.0353	1.1714	1.1839	1.1690
(-1,-1,1)	1.3225	1.3241	1.3156	1.4580	1.4724	1.4560
(-1,0,-1)	1.4593	1.4610	1.4521	1.6095	1.6223	1.6051
(-1,0,0)	1.0465	1.0477	1.0402	1.1656	1.1827	1.1678
(-1,0,1)	1.5185	1.5203	1.5111	1.6717	1.6858	1.6683
(-1,1,-1)	1.4605	1.4622	1.4532	1.6091	1.6236	1.6063
(-1,1,0)	1.2000	1.2011	1.1931	1.3498	1.3626	1.3461
(-1,1,1)	1.4836	1.4852	1.4763	1.6392	1.6504	1.6331
(0,-1,-1)	1.4226	1.4243	1.4153	1.5698	1.5801	1.5629
(0,-1,0)	1.0462	1.0476	1.0397	1.1708	1.1825	1.1670
(0,-1,1)	1.4250	1.4267	1.4178	1.5667	1.5821	1.5651
(0,0,-1)	1.4897	1.4915	1.4822	1.6446	1.6566	1.6387
(0,0,1)	1.5518	1.5536	1.5442	1.7072	1.7219	1.7038
(0,1-1)	1.4915	1.4932	1.4842	1.6451	1.6606	1.6433
(0,1,0)	1.3026	1.3040	1.2953	1.4537	1.4636	1.4462
(0,1,1)	1.5528	1.5546	1.5452	1.7169	1.7275	1.7094
(1,-1,-1)	1.3448	1.3463	1.3378	1.4888	1.4996	1.4829
(1,-1,0)	1.0755	1.0766	1.0689	1.2113	1.2241	1.2084
(1,-1,1)	1.2889	1.2904	1.2818	1.4290	1.4421	1.4253
(1,0,-1)	1.4484	1.4501	1.4411	1.5974	1.6109	1.5934
(1,0,0)	1.1863	1.1876	1.1790	1.3184	1.3396	1.3223
(1,0,1)	1.5162	1.5179	1.5087	1.6699	1.6841	1.6661
(1,1,-1)	1.4286	1.4303	1.4213	1.5771	1.5916	1.5740
(1,1,0)	1.2494	1.2507	1.2421	1.4053	1.4115	1.3941
(1,1,1)	1.4200	1.4216	1.4126	1.5729	1.5822	1.5646

Table 5.5: Average Time of DROP Iterations for Various Chord Length Methods

Method	Average Time (s)
Effective	26.999
Updated Eff.	27.138
Approx Eff.	27.004
True	26.914
Simplified	27.022
Approx. Simplified	27.008

5.3 Timing Analysis

A timing analysis was also performed to examine the tradeoff between accuracy and reconstruction times for each method. First, the average time of each iteration of the DROP algorithm was found for each chord method. It was found that variations in chord length method had minimal effect on the overall timing of a DROP iteration, as all executions had average times within one quarter of a second. This means that chord length calculations are neither a bottleneck, nor a substantial portion of the computations performed in an iteration of the DROP method. Table 5.5 lists these average times per iteration. It is worth noting that the overall timing of current methods of reconstruction is not affected by changes in chord length methods. However, the difference in timing between methods should still be evaluated. To this end, a test was run of each of the chord length methods. All six methods were tested on the same set of randomly generated angle pairs, and their timing recorded. 180 million angle pairs were tested, in 180 batches of 1 million each. 180 million is approximately the number of histories used in reconstructions of CTP404 phantoms, and the number was chosen to reflect that relationship. Times were recorded per batch, and the times reported in Table 5.6 are the average times per batch across all batches.

Clearly, the Effective and Updated Effective methods take substantially more time than any others, with Updated Effective Mean Chord Length taking approxi-

Table 5.6: Average Time in ms of Chord Length Calculation for 1M History Batches

Method	Average Time (ms)
Effective	235.1
Updated Eff.	491.6
Approx. Eff.	1.1
True	58.6
Simplified	34.6
Approx. Simplified	1.2

mately twice the time of Effective Mean Chord Length. Though True and Simplified Mean Chord Length methods take a fraction of the time of the Effective methods, they too show a large difference between the True method, which depends on two angles, and the Simplified method, which only depends on only one. The timing of the two Approximate methods was nearly identical, which is to be expected, as the calculations performed are identical, and the only difference between them are the values of A and b calculated at the beginning of runtime. In the case of Approximated Simplified Chord Length, the Approximate method is nearly 49 times faster than the Simplified Method. In the case of Effective Chord Length, the Approximate method is almost 218 times faster than the Effective Chord Length method.

5.4 Discussion

Given the data, it appears that Updated Effective Chord Length is the best method currently available for pCT Reconstruction. It possesses the highest level of accuracy with respect to RSP values, does not cause a significant slowdown in current reconstruction methods, and its variation with RSP values, within Regions of Interest, and with size of ROI is nearly identical to those possessed by other methods. Approximate Effective Chord Length very nearly approaches the values of Effective Chord Length, and should be considered a valid approximation.

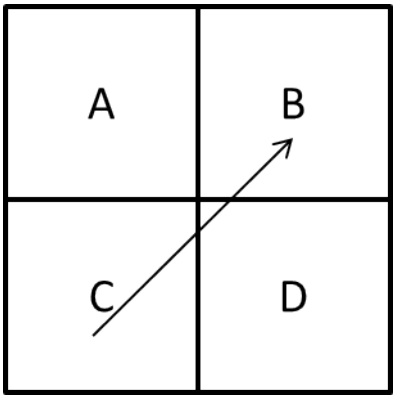

Figure 5.4: An MLP Step Resulting in an Undetected Voxel

5.4.1 Effective versus True

As seen, the family of Effective methods outperforms the family of True methods across all measurements of accuracy. The only advantage True methods appear to have is greatly reduced computation time. The hypothesized reason for this discrepancy in accuracy is the presence of the "undetected" term. It is possible for a Most Likely Path step to "skip" a voxel. This occurs when the points along the MLP are calculated such that they occur in voxels only diagonally adjacent. An example of this case is shown in Figure 5.4. The voxels in which the points reside are set to the chord length established for that path, but the skipped voxel remains set at zero. Effective Chord Length methods attempt to account for this possibility by adding the "undetected" term to the voxels which are detected. This increases the chord length of certain voxels to be greater than the actual length of path intersection, but helps to make up for errors caused by the undetected voxels. True and Simplified methods have no such term, and the shorter total path lengths produced by these methods lead to the systematic overestimation of RSP values. However, the reduced number of calculations produces clear speed advantages for the family of True methods. However, as these speed advantages currently have negligible effect on the overall speed of the reconstruction, Effective methods are the clear choice. This is even more clear when comparing the Approximated versions of both methods, as they have identi-

cal speed performance, but Approximated Effective Mean Chord Length outperforms Approximated Simplified Mean Chord Length in all other categories. An interesting relationship between Effective Methods and True methods is that as the step size approaches zero, ($s \to 0$), Effective methods converge to the True methods. This can be seen in Equations 5.1 and 5.2, which demonstrates the convergence of Effective Mean Chord Length to Simplified Mean Chord Length.

$$\bar{\Delta}_{\text{eff}} = \frac{l}{3}\left(\frac{(s/l)^2 \sin 2\theta - 6}{(s/l)\sin 2\theta - 2(\cos\theta + \sin\theta)} + \frac{(s/l)^2 \sin 2\theta}{2(\cos\theta + \sin\theta)}\right), s \to 0 \tag{5.1}$$

$$\bar{\Delta}(\theta)_{\text{eff}} = \frac{l}{3}\left(\frac{-6}{-2(\cos\theta + \sin\theta)}\right) = \frac{1}{\sin\theta + \cos\theta} = \bar{\Delta}(\theta)_{\text{s-mean}} \tag{5.2}$$

Although that is not the form of Simplified Mean Chord Length implemented in the code, or outlined in Equation 4.4, this value is equal to Simplified Mean Chord length across all $\theta \in [0, \frac{\pi}{2}]$, as seen in Appendix A. This proves to confirm the validity of the Effective Chord Length method, for if MLP step size were zero, no voxel could ever be missed, and thus there would be no need for undetected terms. However, given that current MLP methods do not have a 100% accuracy rate, the Effective Mean Chord Length group of methods should be used over the True Mean Chord Length Methods, until and unless MLP methods advance sufficiently that the differences in accuracy become negligible, or the speed advantages of True Mean Chord Length methods outweigh any concerns over accuracy.

5.4.2 Actual versus Approximate

In the case of Actual methods compared to Approximate methods, there is no clear best option. Approximate methods and their corresponding Actual method perform nearly the same for all tests of accuracy, and display no difference in behavior with variation in ROI size, portion of ROI considered, or actual RSP values. There was consistent underestimation of RSP values for all Approximate methods,

when compared to Actual methods, though this underestimation was small. The hypothesized reason for this underestimation is the consistent tendency of Approximate methods to produced larger chord lengths than Actual methods. It is possible that adjustments to the formulas used to calculate the constants employed in Approximate methods could close this gap, by causing chord lengths for some angles to be overestimated, and others underestimated. This may produce more accurate images, when using Approximate Chord Length methods, though this was not investigated in this work. Approximate methods are vastly faster than their corresponding Actual methods, but this difference does not currently lead to any distinct change in overall reconstruction time. Therefore, it is concluded that Actual methods should be used in current systems, until chord length calculations have more impact on overall reconstruction time, or more accurate constant selection renders the accuracy difference between Actual and Approximate methods negligible.

CHAPTER SIX

Conclusion

Proton Computed Tomography is an emerging medical imaging modality, with particular applications in treatment planning for proton and ion radiation therapies. These therapies are superior to traditional photon therapy, but require accurate Relative Stopping Power Prediction to accurately plan treatments. Reconstructions of images using pCT depend on solving a system of equations in the form $Ax = b$. In this system, the value a_j^i represents the length of a intersection between the i^th proton path and the j^th voxel. This value is referred to as a chord length. This work investigates the impact and validity of various methods for calculating chord lengths.

6.1 Determinations

After examination of available methods, it was determined that the best available method for pCT reconstructions is currently the Updated Effective Mean Chord Length method, also referred to as the Three-Dimensional Effective Mean Chord Length method. This method has the best accuracy out of those examined. Although it is significantly slower per chord length calculated than all other methods tested, the current variation in overall reconstruction times due to different chord length calculation methods is negligible. Therefore Three-Dimensional Effective Mean Chord Length should be held as the current standard for pCT reconstruction. Throughout this work, approximate methods were compared to more exact methods. It was found that approximate methods are reasonably accurate when compared to the method they approximate, but current editions have a tendency to overestimate chord lengths and underestimate Relative Stopping Powers. It is possible that the difference in chord length values produced by approximate and exact methods could

be reduced to a degree that approximate methods may be used because of their considerable timing advantages.

This work also tested the validity of the TROPE technique for evaluating pCT reconstruction usefulness for planning proton and ion therapy treatments. The results obtained via the modified TROPE method employed correlate with those garnered from more traditional analysis techniques. This adds credence to the TROPE method's value as a comparison method, especially as used to determine usefulness of a reconstruction technique for proton and ion therapy treatment planning, where the path-dependent values presented are more applicable.

6.2 Future Work

Ideally, a next step would involve further attempts to optimize the selection of the A and b parameters for the Approximate method, so that it more accurately estimates the values produced by the Two-Dimensional Effective Chord Length method. It may be possible to select A and b such that the results from the Approximate method are indistinguishable from those obtained using the exact method. If this is possible, there would remain no reason to use Two-Dimensional Effective Chord Length when accuracy remains the same, and the Approximate method obtains values many times faster. A perhaps more important step would be add another multiplicative term to the Approximate method to account for the variations in Three-Dimensional Effective Chord Length with respect to ϕ. This potentially would increase the overall accuracy of the Approximate method, and could eventually allow it to replace Three-Dimensional Effective Chord Length in the same manner as described for the Two-Dimensional version.

Additionally, efforts which increase the accuracy of the Most Likely Path method should be encouraged. As MLP methods become more accurate, the difference between Effective Chord Length and True Chord Length decreases. It may eventually

be that MLP methods will advance such that the differences in reconstructions performed with Effective Chord Length methods and True Mean Chord Length methods are negligible. A simple way to do this involves reducing the step size of the MLP procedure. However, this takes a heavy toll on reconstruction times, and should not necessarily be the first step.

An intriguing idea is one which compromises between Exact Chord Length and other chord length methods. Previous studies indicated that there were small differences in accuracy between Effective Chord Length and Exact Chord Length methods, but a large discrepancy in timing (Penfold, Rosenfeld, Schulte, and Schubert 2009). However, it may be possible to combine these two methods such that an Effective Chord Length is not assigned for an entire proton path, but only a portion thereof. This has the potential to increase accuracy without costing much additional computation time. This would be especially true if the Approximate method can be made highly accurate without adding computations. For example, rather than calculating the Effective Chord length for an entire path, the first 50 voxels on that path could be identified using MLP methods. Then, using the changes in coordinates from the entry voxel, angles θ and ϕ can be calculated. These angles could be used to assign an Effective Chord Length over this portion of the path. The procedure would then repeat, with the previous endpoint becoming the new starting point for comparison. This procedure could increase accuracy with a minimum of increased computation time, especially if the Approximate method may be used.

Regarding the TROPE method, the first step would be to create "ideals" of phantoms commonly used for testing. Then the theoretical ideal RSP of path could be compared to that generated by the reconstruction technique, and the full version of TROPE executed. Once this was complete, TROPE could be used to determine what types of errors cause errors in treatment planning, as well as which treatment paths are most error prone.

6.3 Closing Remarks

The validity and usefulness of several chord length calculation methods has been assessed, with Three-Dimensional Effective Chord Length emerging as the obvious choice in the current state of pCT reconstruction methods. Advancements in Most Likely Path techniques or more accurate calibrations of the Approximate method may eventually allow Approximate or True Mean methods to become valid options. Currently, however, the reduction in accuracy for negligible improvements in reconstruction time argue against their use.

APPENDICES

APPENDIX A

Derivation of Mean Chord Length

The formula for True Mean Chord Length is derived as follows. First, chord length is a function of θ, ϕ, x, y, and z. In this case, θ and ϕ represent the angle in the xy plane, and the angle with the xy plane, respectively. x, y, and z are the entry points of the path into the voxel. The total intersection length is determined by the minimum of three values. If all inputs to the function are known, the three values and true chord length can be defined as outlined in Equation A.1. Because of the symmetry, it can be assumed that θ and ϕ are between 0 and $\frac{\pi}{2}$ inclusive. For the duration of the derivation, it is assumed

$$\Delta(x, y, z, \theta, \phi) = \begin{cases} l_1 = (1-y)\csc\theta\csc\phi, l_1 < l_2 and l_1 < l_3 \\ l_2 = (1-x)\sec\theta\csc\phi, l_2 < l_1 and l_2 < l_3 \\ l_3 = (1-z)\csc\phi, l_3 < l_1 and l_3 < l_2 \end{cases} \tag{A.1}$$

The first simplifying assumption which can be made is based on an observation of proton behaviors. As noted in a study by Paniz Karbasi pending publication, over 96% of protons vary less than two voxels in the z direction over their entire path. This means that greater than 96% of voxel intersections will enter and exit the voxel through the side, rather than through the top or bottom. Because of this, quantity l_3 can be removed from consideration, as it will almost never be the limiting factor. This changes the equation for chord length to Equation A.2.

$$\Delta(x, y, z, \theta, \phi) = \begin{cases} l_1 = (1-y)\csc\theta\csc\phi, l_1 < l_2 \\ l_2 = (1-x)\sec\theta\csc\phi, else \end{cases} \tag{A.2}$$

With the simplifying assumption that the path will enter from the sides combined with the rotational symmetry, it can be seen that if $x \neq 0$, $y = 0$, and vice versa. Therefore expressions can be found for the case where $\theta \in [0, \frac{\pi}{4}]$ (Equation A.3.)

$$\Delta(x,y,z,\theta,\phi) = \begin{cases} l_1, 1-\tan\theta < y \leq 1 \\ \\ l_2, 0 \leq y \leq 1-\tan\theta \end{cases} \tag{A.3}$$

Next recall the definition of Expected Value for a function of a random variable.

$$E[U] = \int_0^\infty f(u)p(u)du \tag{A.4}$$

In this case, only x or y will be changing at a given time, therefore the expected value of $\Delta(x,y|\theta,\phi)$ can be computed as sums of functions of only x or y, by combining the definition of expected value and Equation A.3.

$$\Delta(x,\bar{y}|\theta,\phi) = \int_0^{1-\tan\theta} l_2 p(y)dy + \int_{1-\tan\theta}^1 l_1 p(y)dy \int_0^1 l_2 p(x)dx \tag{A.5}$$

It is necessary to find the probability density function for x and y. This is a function of θ, the likelihood of a chord intersecting when $x = 0$ increases with $\sin\theta$, and the likelihood of that chord intersecting when $y = 0$ increases with $\cos\theta$. An assumption is made that any point on a given side of the voxel is equally likely as any other, leading to Equation A.6.

$$\begin{aligned} \bar{\Delta}(x,y|\theta,\phi) = \frac{(1-\tan\theta)\sin\theta}{\sin\theta+\cos\theta}\int_0^{1-\tan\theta} l_2 dy + \\ \frac{\sin\theta\tan\theta}{\sin\theta+\cos\theta}\int_{1-\tan\theta}^1 l_1 dy + \frac{\sin\theta}{\sin\theta+\cos\theta}\int_0^1 l_2 dx \end{aligned} \tag{A.6}$$

Substituting in expressions for l_1 and l_2.

$$\begin{aligned} \bar{\Delta}(x,y|\theta,\phi) = \frac{(1-\tan\theta)\sin\theta}{\sin\theta+\cos\theta}\int_0^{1-\tan\theta} (1-x)\sec\theta\csc\phi dy + \\ \frac{\sin\theta\tan\theta}{\sin\theta+\cos\theta}\int_{1-\tan\theta}^1 (1-y)\csc\theta\csc\phi dy + \frac{\sin\theta}{\sin\theta+\cos\theta}\int_0^1 (1-x)\sec\theta\csc\phi dx \end{aligned} \tag{A.7}$$

After evaluating the definite integral, the expression evaluates to the expression in Equation A.8.

$$\bar{\Delta}(x,y|\theta,\phi) = \frac{\sec\phi}{\sin\theta+\cos\theta} \tag{A.8}$$

It was found that for values of $\theta \in [0, \frac{\pi}{2}]$, Equation A.1 is equivalent to Equation A.9.

$$\bar{\Delta}(x,y|\theta,\phi) = \frac{\sec\phi\csc(\theta+\frac{\pi}{4})}{\sqrt{2}} \tag{A.9}$$

This equation assumes a voxel width of 1.0. To account for the possible other voxel lengths, the simple multiplicative factor of l is added, where l is the path length. This results in the expression used to in the code to calculate the True Mean Chord Length.

$$\bar{\Delta}(x, y|\theta, \phi) = l\frac{\sec\phi\csc(\theta + \frac{\pi}{4})}{\sqrt{2}} \tag{A.10}$$

To change from True Mean Chord Length to Simplified Mean Chord Length, the assumption that $\phi \approx 0$ is made. This leads to the expression in Equation A.11.

$$bar\Delta(x, y|\theta, \phi) = l\frac{\csc(\theta + \frac{\pi}{4})}{\sqrt{2}} \tag{A.11}$$

This is the equation used for calculating Simplified Mean Chord Length. The code implementing both True Mean Chord Length and Simplified Mean Chord Length is shown in Appendix B.

APPENDIX B

Code for Mean Chord Length

The code used to compute True Mean Chord Length and Simplified Mean Chord Length is shown in this appendix, with some commentary on implementation.

For True and Simplified Mean Chord Length, one small change to the configurations method was required. A value storing the $l/\sqrt{2}$ term in the formula was calculated, and stored in the configurations object which is carried through all functions above the True and Simple Mean Chord length methods. The next change required was in the Device Functions header file, where the function prototypes were added. These changes were minor, and therefore the code implementing them is not included. However, the primary change was in the Device Functions source file, where the functions were implemented.

File: pCT_Reconstruction_D.cpp

```
__device__ double TrueMeanChordLength_GPU(double abs_theta,
double x_dist, double y_dist, double z_dist,configurations * config_d)
{
double eff_theta; \\To store theta between 0 and 90
eff_theta = abs_theta - ((int)(abs_theta/(PI/2)))*(PI/2) + (PI/4);
double phi = atan(z_dist/sqrt(x_dist*x_dist + y_dist*y_dist));
return config_d->mean_frac/(cos(phi)*sin(eff_theta));
}
```

```
__device__ double SimplifiedMeanChordLength_GPU(double abs_theta,
configurations * config_d)
{
double eff_theta;
eff_theta =abs_theta-((int)(abs_theta/(PI/2)))*(PI/2)+(PI/4);
return config_d->mean_frac /(sin(eff_theta));
}
```

It can be seen that the calculation of ϕ itself requires both more input arguments, and substantially more calculations. This is one reason for the speed differential observed between the two methods.

APPENDIX C

Derivation of Approximate Mean Chord Length

The formula for Approximate Mean Chord Length is derived as follows. First, the shapes of both Simplified Mean Chord Length and Effective Mean Chord Length were graphed. Then, different functions which could approximate the shape were graphed for comparison. Estimations considered included cosine and sine functions, and a quadratic curve. It was observed that the shape was most closely approximated by a quadratic curve. The form of the curve is given in Equation C.1.

$$\bar{\Delta}_{Approx}(\theta) = A(\theta - \theta_0)^2 + b \tag{C.1}$$

It was obvious that $\theta_0 = \frac{\pi}{4}$, the point of lowest mean chord length, and the parabola must be shifted accordingly. To find b, the value of the method approximated was found for $\theta = \frac{\pi}{4}$, as this would be the minimum value of the parabola. This value is outlined by Equation C.2.

$$b = \frac{l}{3}\frac{24 - \sqrt{2}r^3}{8\sqrt{2} - 4r}, r = \frac{s}{l} \tag{C.2}$$

Note that because Effective Mean Chord Length reduces to True Mean Chord Length as $s \to 0$, the value of b for Approximate Simplified Mean Chord Length can be found by setting $s = 0$, despite the lack of an s term in Simplified Mean Chord Length. A is calculated in such a way that $\bar{\Delta}_{Approx}(0) = \bar{\Delta}_{Approx}(\frac{\pi}{2}) = 0$. This leads to Equation C.3.

$$A = l\frac{16(1-b)}{\pi^2} \tag{C.3}$$

The values of A and b need only be calculated once per reconstruction, as they will not vary for fixed voxel and step sizes, which do not vary within a reconstruction. Combined with minimal additional operations, and a lack of trigonometric functions makes the Approximate Method much simpler to implement in code, as shown in Appendix D.

APPENDIX D

Code for Approximate Mean Chord Length

The code implementing the Approximate Mean Chord Length method is shown.

The addition of two new terms to the configurations method was required to streamline the Approximate Mean Chord Length method. The additions include terms for the values of A and b, as required by the Approximate Mean Chord Length. In the implementation for this experiment, two different functions were used to calculate Approximate Mean Chord Length, though they are effectively identical. Therefore, there are four values added in this implementation.

File: Configurations.h

```
float approx_eff_chord_base;
// [cm] distance for base of Approx Effective Chord Length Method
float approx_eff_chord_range;
// [cm] distance for range of Approx Effective Chord Length Method
float approx_simple_chord_base;
// [cm] distance for base of Approx Simple Chord Length Method
float approx_simple_chord_range;
// [cm] distance for range of Approx Simple  Chord Length Method
```

These values are initialized at the same time as other values in the configurations method.

File: pCT_Reconstruction_H.cpp

```
void initialize_parameters(){
...
parameter_container.step_frac = MLP_U_STEP/VOXEL_WIDTH;
    // fraction of mlp step to voxel width in xy plane
    parameter_container.mean_frac = VOXEL_WIDTH / sqrt(2.0);
    parameter_container.approx_eff_chord_base =
    1/3.0 * (24-sqrt(2.0) * parameter_container.step_frac *
     parameter_container.step_frac * parameter_container.step_frac)
    /(8.0*sqrt(2.0) - 4.0*parameter_container.step_frac);
    // [cm] distance for base of approx chord length method
    parameter_container.approx_eff_chord_range = 16.0*
    (1.0- parameter_container.approx_eff_chord_base)/(PI*PI);
    // [cm] distance for base of approx chord length method
    parameter_container.approx_simple_chord_base = 1/(sqrt(2.0));
    // [cm] distance for base of approx chord length method
    parameter_container.approx_simple_chord_range = 16.0*
    (1- parameter_container.approx_simple_chord_base)/(PI*PI);
    // [cm] distance for base of approx chord length method
}
```

As with the True and Simplified Mean Chord Length methods, the primary change for the Approximate Chord Length methods comes in the Device Functions source file.

```
__device__ double ApproxEffectiveChordLength_GPU(double abs_theta,
configurations * config_d)
{
double eff_theta;
eff_theta = abs_theta - ((int)(abs_theta/(PI/2)))*(PI/2)-(PI/4);
return ((config_d->approx_eff_chord_range * eff_theta * eff_theta +
config_d->approx_eff_chord_base)*config_d->voxel_width);
}
__device__ double ApproxSimplifiedChordLength_GPU(double abs_theta,
configurations * config_d)
{
double eff_theta;
eff_theta = abs_theta - ((int)(abs_theta/(PI/2)))*(PI/2)-(PI/4);
return ((config_d->approx_simple_chord_range * eff_theta * eff_theta +
config_d->approx_simple_chord_base) * config_d->voxel_width);
}
```

As seen, the only difference between the methods implemented is which constants are referred to. With some small changes, it would be possible to create a generalized method that selects its constant values at runtime, rather than having two separate methods for what is essentially the same function. Because these methods are nearly identical, it is unsurprising that they perform extremely similarly during timing trials.

APPENDIX E

TROPE Method

The Total RSP Of Path Error (TROPE) method is designed to test a reconstruction's effectiveness accuracy. It is tailored to evaluating the behavior of a reconstruction when planning treatments for proton therapy, but can be used for any reconstruction. It is particularly useful for noting error tendencies which could negatively affect treatment planning. The premise of TROPE is to test a variety of paths leading to voxels in the reconstruction and compare the results to ideal paths of the same object.

The purpose of the TROPE method is to analyze the efficacy of treatment plans using different reconstruction methods, rather than the direct representative accuracy. It can also be used to note error tendencies related to path direction.

The version of the TROPE method implemented for this work tests twenty-six different paths to each voxel selected for testing. Every path on an axis or at a 45° from an axis is tested. The paths are labeled by the vector leading outward from the voxel of interest, not inward. Therefore, path (-1,0,1) leaves the voxel of interest, progressing one voxel in the negative x direction and one voxel in the positive z-direction until it reaches the edge of the reconstructed volume.

As the algorithm walks along the path requested, it sums the Relative Stopping Power of each voxel encountered, and eventually reports the average for the path. This average value represents the ratio of the Water Equivalent Path Length of the path to the actual length of the path. Once this process has been performed for both the reconstructed image and the ideal image, the differences can between the reconstruction and the idea can be found for each path and voxel of interest. The difference between the reconstructed value and the ideal value will indicate the error tendency of proton treatments planned using the reconstruction data. For example,

the reconstruction gives a value of 1.5 for a given path, but the ideal gives a value of 1.0 for the same path. In this case, a proton launched using the energy derived from the reconstruction data would travel significantly longer than necessary. This difference would be an extreme case, but serves to illustrate information conveyed by the two values.

The TROPE method was implemented in MATLAB®. The code is split between two functions, which are presented below.

The first function calls the individual path-wise RSP evaluation for each path under consideration. pixel_path_RSP.m

```
function RSP = pixel_path_RSP(Image,xp,yp,zp)
    RSP = zeros(1,8);

    ind = 1;
    for xstep = -1:1
        for ystep = -1:1
            for zstep = 0%-1:1

                if((zstep ~=0) || (ystep ~=0) || (xstep ~= 0))

                    RSP(ind) =
                    direction_RSP(xp,yp,zp,Image,xstep,ystep,zstep);
                    ind = ind+1;
                end
            end
        end
    end
end
```

The next function finds the path RSP of a given path for the selected pixel. direction_RSP.m

```
function rspTotal = direction_RSP(xp,yp,zp,I,xstep,ystep,zstep)


    [x,y,z] = size(I);
    rspTotal = 0.5* I(yp,xp,zp);
    xi = xp+xstep;
    yi = yp+ystep;
    zi = zp+zstep;
    numsteps = 0.5;
    in = true;
    while((xi >=1) && (xi <=x) && (yi >=1)
    && (yi <=y) && (zi >=1) && (zi <=z) && in)
        vox = I(yi,xi,zi);
        if(vox ==0)
            in = false;
        else
            rspTotal = rspTotal + I(yi,xi,zi);
            xi = xi+xstep;
            yi = yi+ystep;
            zi = zi+zstep;
            numsteps = numsteps +1;
        end
    end
    rspTotal = rspTotal/numsteps;
end
```

Note that because MATLAB® stores images in a row-column format, rather than x,y format, the indexes are not in the traditional x,y,z order.

BIBLIOGRAPHY

Bashkirov, V., R. Schulte, G. Coutrakon, B. Erdelyi, K. Wong, H. Sadrozinski, S. Penfold, A. Rosenfeld, S. McAllister, and K. Schubert (2009). Development of proton computed tomography for applications in proton therapy. In *AIP Conference Proceedings*, Volume 1099, pp. 460–463. AIP.

Bruzzi, M., N. Blumenkrantz, J. Feldt, J. Heimann, H. F. W. Sadrozinski, A. Seiden, D. C. Williams, V. Bashkirov, R. Schulte, D. Menichelli, M. Scaringella, G. A. P. Cirrone, G. Cuttone, N. Randazzo, V. Sipala, and D. L. Presti (2007, Feb). Prototype tracking studies for proton ct. *IEEE Transactions on Nuclear Science 54*(1), 140–145.

Cormack, A. M. (1963). Representation of a function by its line integrals, with some radiological applications. *Journal of applied physics 34*(9), 2722–2727.

Hanson, K., J. Bradbury, R. Koeppe, R. Macek, D. Machen, R. Morgado, M. Paciotti, S. Sandford, and V. Steward (1982). Proton computed tomography of human specimens. *Physics in medicine and biology 27*(1), 25.

Hanson, K. M. (1979). Proton computed tomography. *IEEE Transactions on Nuclear Science 26*(1), 1635–1640.

Karbasi, P., B. Schultze, V. Giacometti, T. E. Plautz, K. Schubert, R. W. Schulte, and V. Bashkirov (2015). Incorporating robustness in diagonally-relaxed orthogonal projections method for proton computed tomography.

Li, T., Z. Liang, J. V. Singanallur, T. J. Satogata, D. C. Williams, and R. W. Schulte (2006). Reconstruction for proton computed tomography by tracing proton trajectories: A monte carlo study. *Medical Physics 33*(3), 699–706.

Paganetti, H. (2016). *Proton therapy physics*. CRC Press.

Patyal, B. (2007). Dosimetry aspects of proton therapy. *Technology in cancer research & treatment 6*(4_suppl), 17–23.

Penfold, S. and Y. Censor (2015). Techniques in iterative proton ct image reconstruction. *Sensing and Imaging 16*(1), 19.

Penfold, S., R. W. Schulte, Y. Censor, V. Bashkirov, S. McAllister, K. Schubert, and A. B. Rosenfeld (2010). Block-iterative and string-averaging projection algorithms in proton computed tomography image reconstruction.

Penfold, S. N. (2010). Image reconstruction and monte carlo simulations in the development of proton computed tomography for applications in proton radiation therapy.

Penfold, S. N., A. B. Rosenfeld, R. W. Schulte, and K. E. Schubert (2009). A more accurate reconstruction system matrix for quantitative proton computed tomography. *Medical Physics 36*(10), 4511–4518.

Petkovska, S., C. Tolevska, S. Kraleva, and E. Petreska (2010). Conformity index for brain cancer patients. In *CONFERENCE ON MEDICAL PHYSICS AND BIOMEDICAL ENGINEERING*, pp. 56.

Plautz, T. E., V. Bashkirov, V. Giacometti, R. Hurley, R. Johnson, P. Piersimoni, H.-W. Sadrozinski, R. W. Schulte, and A. Zatserklyaniy (2016). An evaluation of spatial resolution of a prototype proton ct scanner. *Medical physics 43*(12), 6291–6300.

Sadrozinski, H.-W., V. Bashkirov, B. Keeney, L. R. Johnson, S. G. Peggs, G. Ross, T. Satogata, R. W. Schulte, A. Seiden, K. Shanazi, *et al.* (2004). Toward proton computed tomography. *IEEE Transactions on Nuclear Science 51*(1), 3–9.

Schneider, U. and E. Pedroni (1994, November). Multiple Coulomb scattering and spatial resolution in proton radiography. *Medical Physics 21*, 1657–1663.

Schulte, R. W., S. N. Penfold, J. T. Tafas, and K. E. Schubert (2008). A maximum likelihood proton path formalism for application in proton computed tomography. *Medical Physics 35*(11), 4849–4856.

Schultze, B., P. Karbasi, V. Giacometti, T. E. Plautz, K. Schubert, and R. W. Schulte (2015). Reconstructing highly accurate relative stopping powers in proton computed tomography.