IOP Publishing | Institute of Physics and Engineering in Medicine Physics in Medicine & Biology

Phys. Med. Biol. **61** (2016) 5868–5882 doi:10.1088/0031-9155/61/15/5868

Fast reconstruction of low dose proton CT by sinogram interpolation

David C Hansen[1], **Thomas Sangild Sørensen**[2] **and Simon Rit**[3]

[1] Department of Oncology, Aarhus University Hospital, Nørrebrogade 44, building 5, DK-8000 Aarhus C, Denmark
[2] Department of Clinical Medicine, Aarhus University, Palle Juul-Jensens Boulevard 82, DK-8200 Aarhus N, Denmark
[3] Université de Lyon, CREATIS, CNRS UMR5220, Inserm U1044, INSA-Lyon, Université Lyon 1, Centre Léon Bérard, 69008 Lyon, France

E-mail: daviha@rm.dk

Received 27 July 2015, revised 10 May 2016
Accepted for publication 20 June 2016
Published 18 July 2016

Abstract

Proton computed tomography (CT) has been demonstrated as a promising image modality in particle therapy planning. It can reduce errors in particle range calculations and consequently improve dose calculations. Obtaining a high imaging resolution has traditionally required computationally expensive iterative reconstruction techniques to account for the multiple scattering of the protons. Recently, techniques for direct reconstruction have been developed, but these require a higher imaging dose than the iterative methods. No previous work has compared the image quality of the direct and the iterative methods. In this article, we extend the methodology for direct reconstruction to be applicable for low imaging doses and compare the obtained results with three state-of-the-art iterative algorithms. We find that the direct method yields comparable resolution and image quality to the iterative methods, even at 1 mSv dose levels, while yielding a twentyfold speedup in reconstruction time over previously published iterative algorithms.

Keywords: proton CT, iterative reconstruction, stopping power

(Some figures may appear in colour only in the online journal)

1. Introduction

The primary advantage of particle therapy in cancer treatment is the sparing of healthy tissues thanks to the finite range of charged particles in matter. However, a very accurate prediction of the particle range is required in order to exploit the full benefit of the treatment, and range

© 2016 Institute of Physics and Engineering in Medicine Printed in the UK

uncertainties are currently the largest contributor to the dose margins employed (Paganetti 2012). Several methods have been proposed for improving particle range prediction and verification (Knopf and Lomax 2013), including prompt-gamma imaging (Min *et al* 2006), positron-emission tomography (PET) (Bennett *et al* 1978, Knopf *et al* 2011), dual energy CT (Yang *et al* 2010a, Hünemohr *et al* 2014) and ionoacoustics (Assmann *et al* 2015). With the exception of dual energy CT, all of these methods are mainly concerned with *in vivo* range verification and cannot be used for dose calculation and planning. While dual-energy CT yields better range estimates than traditional CT, there are still significant errors, particularly for metals but also in tissue (Hünemohr *et al* 2014, Hansen *et al* 2015). An alternative is to image the patient using proton CT, where an imaging proton beam is shot through the patient at low intensity, the energy-loss measured, and a map of the stopping-power throughout the patient reconstructed. This approach was originally proposed by Cormack (1963) but suffered from low spatial resolution due to the effects of multiple scattering. Huesman *et al* (1975) proposed to compensate for the scattering by measuring the angles and positions of every individual proton, before entering and after leaving the patient. Only within the last decade, detectors have become fast enough to perform such measurements and computers fast enough for the corresponding image reconstruction.

Two prototype proton CT scanners that allow for path modelling are currently in operation (Hurley *et al* 2012, Sadrozinski *et al* 2013). Several papers have investigated fast iterative methods for proton CT (Penfold *et al* 2010a, 2010b, Hansen *et al* 2014a). Nonetheless, a medium sized computing cluster is required to achieve clinically useful reconstruction times even when applying such accelerated methods (Karonis *et al* 2013). Two faster, direct reconstruction methods have been proposed: a filtered backprojection approach (the path-FBP) (Rit *et al* 2013) and a backprojection-then-filtering (BPF) approach (Poludniowski *et al* 2014). Barring interpolation and a truncation approximation in the BPF, the two methods are mathematically equivalent, but it is more straight forward to use the path-FBP method for fan and cone-beam type scans. Additionally, the path-FBP method was recently extended to work with two-step Hilbert transform-type reconstructions (Noo *et al* 2004), allowing for reconstructions from truncated projections (Rit *et al* 2015). In this work, we demonstrate that the path-FBP method breaks down at low imaging doses. We consequently extend the method to become applicable at low imaging doses using a sinogram interpolation technique, and compare the proposed technique with state-of-the-art iterative methods.

2. Proton CT reconstruction

For proton CT, the basic problem of tomographic reconstruction corresponds to solving

$$\mathcal{A}\mathbf{x} = \mathbf{f} + \epsilon \tag{1}$$

where $\mathbf{x}$ is the desired reconstruction of water equivalent stopping power (WESP), $\mathbf{f}$ contains the measured water equivalent path length (WEPL) of each individual proton and ϵ is the noise. $\mathcal{A}$ is a linear operator modeling the path integral of each individual proton path through the image. Several options exist for estimating the individual proton paths from the measured angles and positions, including the most-likely-path (MLP) (Williams 2004). Another option is to simply use a cubic spline to approximate the path (Li *et al* 2006), which was used in our previous works (Hansen *et al* 2014a, 2014b). The latter approach results in a slightly poorer spatial resolution than the MLP (Li *et al* 2006, Wang *et al* 2011). It does however have the advantage that it allows for fast GPU calculation of $\mathcal{A}\mathbf{x}$ without explicitly storing $\mathcal{A}$ as a matrix, as the coefficients are easily calculated on the fly. This allows for processing of much larger images and proton datasets, as the matrix easily contains on the order of 10^{10} non-zero elements. For this reason, the cubic spline approach was chosen for this work.

2.1. Path-FBP

In parallel x-ray CT, the standard filtered backprojection is given by

$$\mathbf{u} = \mathcal{A}^{\dagger}(\mathbf{k}_{\mathrm{1D}} * \mathbf{p}) \tag{2}$$

where $\mathbf{k}_{\mathrm{1D}} * \mathbf{p}$ represents the convolution of the projection data $\mathbf{p}$ with the 1D kernel $\mathbf{k}_{\mathrm{1D}}$. $\mathcal{A}^{\dagger}$ indicates the adjoint of $\mathcal{A}$, which in this case is the backprojection operator. Rit *et al* (2013) generalized this formula for proton paths. For each scanned angle θ, a 3D projection is made via the equation

$$\mathbf{p}_{\theta} = \mathcal{A}^{\dagger}_{\theta}(\mathbf{f}_{\theta}) \oslash \mathcal{A}^{\dagger}_{\theta}(\mathbf{1}). \tag{3}$$

$\mathbf{f}_{\theta}$ refers to the protons measured at angle θ, $\oslash$ denotes elementwise division, and $\mathbf{1}$ is the vector with 1 in every entry and the same length as $\mathbf{f}$. The second term can be seen as a normalization factor and the entire equation simply as the average WEPL through each voxel. If the protons move along straight line trajectories, no variation would be observed along the beam direction and equation (3) would simplify to the backprojection of a standard 2D tomographic projection. The final reconstruction from t projections is computed as

$$\mathbf{x} = \frac{1}{t}\sum_{\theta=0}^{\theta_t} \mathbf{k}_{\theta} * \mathbf{p}_{\theta} \tag{4}$$

where $\mathbf{k}_{\theta} * \mathbf{p}_{\theta}$ is the convolution of the filter orthogonal to the beam direction and the axis of rotation. In this work we use the Hann weighted ramp filter, with a cutoff frequency of 80% at CTDEI = 1 mSv and 50% and CTDEI = 0.5 mSv (Hansen *et al* 2014a).

2.1.1. Low dose. One problem of the path-FBP is that for low doses and high image resolution, there is no guarantee that protons will pass through every image voxel from every scanning angle. This leaves holes in the backprojection, causing the filtering to break down. In fact, for the doses used in this study, such holes are almost guaranteed to appear in every projection (99.96% probability, assuming straight lines and uniform distribution). We propose a very simple approach to fix this: missing pixels are interpolated based on an average of the four neighbours in the projection, excluding any that might be themselves be missing. For a single slice scan, this only works when no three holes are adjacent. However, at doses of CTDEI = 1 mSv and 0.3 mm voxel width, the probability of this occurring is about 0.02% per scan. This increases sharply with decreasing doses, but below 1 mSv, iterative methods are also known to start showing artefacts (Hansen *et al* 2014a). This is however also dependent on things such as voxel size, proton energy and size of the scanned phantom, so the results may not be directly comparable. For even lower doses, the interpolation could be done iteratively, until all holes are filled.

2.2. Statistical reconstruction

Statistical image reconstruction is a widely studied topic in x-ray CT (Beister *et al* 2012) and was recently applied to proton CT as well (Hansen *et al* 2014a). In statistical reconstruction, the goal is to reconstruct the image that matches the measured data with the highest probability given some statistical model of the noise and the image to be reconstructed. Accurately modelling this for proton CT is inherently complicated and varies heavily with the energy loss (ICRU Report 49 and International Commision on Radiation Units and Measurements 1993). In this work, two different statistical models were used. In both, variance was assumed to be normally distributed (Gaussian); in the first the noise was assumed to be equal for all detected particles while in the second a more complicated model was employed. The variance of the ith proton was assumed to be the root mean square of the variance from the detectors (σ_{det}) and the variance coming from the energy straggling of the protons (σ_{stragl}):

$$\sigma_i^2 = \sigma_{\text{det}}^2 + \sigma_{\text{stragl}}^2. \tag{5}$$

Under the assumption that the straggling happened in a uniform medium σ_{stragl} can be calculated via the theory of Tschalär (1968). The corresponding penalized weighted least squares (PWLS) reconstruction problem becomes

$$\min_{\mathbf{x}} \|\mathcal{W}(\mathcal{A}\mathbf{x} - \mathbf{f})\|_2^2 + \lambda \mathcal{R}(\mathbf{x}) \text{ s.t. } \mathbf{x}_{i,j,k} \geqslant 0 \tag{6}$$

where $\mathcal{W}$ is a diagonal matrix containing the inverse of the variance squared $\mathcal{W}_{ii} = \sigma_i^{-2}$. In the case of uniform variance, the problem reduces to a penalized least squares problem (PLS). $\mathcal{R}(\mathbf{x})$ is a noise suppressing roughness penalty and λ the corresponding weight. In this work the Huber norm is applied, defined by

$$\mathcal{R}(\mathbf{x}) = \sum_{i,j,k} \begin{cases} \dfrac{\text{TV}(\mathbf{x})_{i,j,k}^2}{2\alpha} & \text{if } \text{TV}(\mathbf{x})_{i,j,k} \leqslant \alpha \\ \text{TV}(\mathbf{x})_{i,j,k} - \dfrac{\alpha}{2} & \text{otherwise} \end{cases} \tag{7}$$

where $\text{TV}(\mathbf{x})_{i,j,k}$ indicates the well-known total variation at voxel i, j, k

$$\text{TV}(\mathbf{x})_{i,j,k} = \sqrt{(\mathbf{x}_{i,j,k} - \mathbf{x}_{i,j,k-1})^2 + (\mathbf{x}_{i,j,k} - \mathbf{x}_{i,j-1,k})^2 + (\mathbf{x}_{i,j,k} - \mathbf{x}_{i-1,j,k})^2} \tag{8}$$

Equation (6) was solved using an ordered subset momentum accelerated augmented Lagrangian algorithm (Kim *et al* 2015) with 5 iterations of the Arrow-Hurwicz primal-dual algorithm for solving the inner denoising problem (Arrow *et al* 1958, Chambolle and Pock 2011). The regularization weight λ was set to 1 for the PLS algorithm and 0.001 for PWLS for the Gammex phantom and 10 and 0.01 for the ICRP phantom. In both cases $\alpha = 0.02$. Details are provided in the appendix.

2.3. Total variation superiorisation

Another iterative reconstruction method used successfully for proton CT is total variation superiorization (TV-SUP) (Penfold *et al* 2010a, Censor *et al* 2010). In this approach, an iterative algorithm for the solution of the reconstruction problem $\mathcal{A}\mathbf{x} = \mathbf{f}$ is modified such that after every iteration, an additional denoising step is taken. This relies on the iterative algorithm to be robust to small pertubations and similarly relies on the denoising step to be relatively small. The performance of superiorization in the presence of noise (where $\mathcal{A}\mathbf{x} = \mathbf{f}$ does not have a solution) has not been studied, but the algorithm performs well for a number of practical applications (Censor *et al* 2010). In this work, an implementation of the diagonally relaxed orthogonal projection (DROP) algorithm (Penfold *et al* 2010a) was used. Again, further details are provided in the appendix.

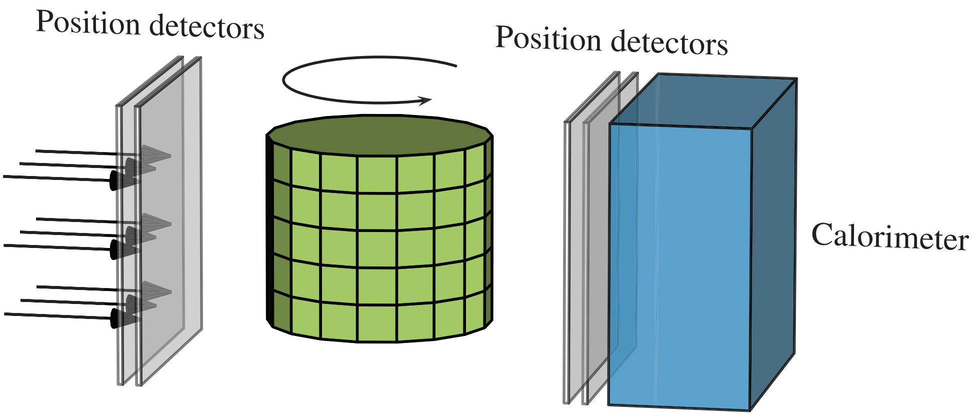

Figure 1. Geometry of the simulated proton CT scan, consisting of two position detectors in front of the phantom and two after the phantom, followed by a calorimeter.

3. Monte Carlo simulation

In order to evaluate the performance of each algorithm, a simulated proton CT scan was performed using the Monte Carlo transport code Geant4 (Allison *et al* 2006). The setup was modelled after the proton CT scanners in Sipala *et al* (2011a) and Sadrozinski *et al* (2013) (shown on figure 1). It consists of two pairs of two silicon strip detectors, each 200 μm thick, placed on each side of the phantom, allowing for detection of position and direction of individual particles. The distance between each silicon strip detector was 5 cm, and the detectors were given a spatial accuracy of $\sigma_{xy} = 200$ μm. A YAG:Ce calorimeter was placed after the position detectors, modelled after the PRIMA detector Sipala *et al* (2011b). Based on Sipala *et al* (2011b), an energy resolution of $\sigma_E(E) = \max(3\% \cdot E, \frac{24.15\ \mathrm{MeV}^2}{E} + 1.76\ \mathrm{MeV})$ was assigned. Only protons were considered in the scoring. This setup was also used in a previous publication on proton CT (Hansen *et al* 2015). First, a phantom was scanned using 360 angles at an energy of 250 MeV, at a dose of CTDEI = 1 mSv (Hansen *et al* 2014a), which corresponds to about 7500 protons per slice per angle. The phantom used was a replica of the Gammex RMI 467 electron density phantom (Gammex, Middleton, WI, USA), which has a diameter of 33 cm. In addition, a patient scan was also simulated, again using 360 angles but at an energy of 330 MeV. Here we used the female ICPR computational phantom (ICRP 2009) (see figure 2), and did a scan of a 10 cm slice of the lungs and a dose of CTDEI = 0.5 mSv. The voxelised phantom has a voxel size of $1.775 \times 1.775 \times 4.84$ mm^3 and consists of 141 different materials. After both simulations, all events in which less than 0.5% of the initial energy was lost and the angle between entrance and exit were more than 3σ were removed, as has been done in several other publications (Schulte *et al* 2008, Hansen *et al* 2014a).

3.1. Reconstruction

All reconstructions were performed on an NVIDIA Titan X GPU using the CUDA programming language and the Gadgetron image reconstruction framework (Hansen and Sørensen 2013). The reconstruction code is matrix-less, meaning that the matrix-vector products were implemented without explicitly forming the matrix, significantly reducing memory requirements. For the Gammex phantom, images were reconstructed on a 1024×1024 grid

Figure 2. Slice of the ICRP phantom used in this study.

with a 34 cm × 34 cm size. For the Path-FBP, the generated projections had a resolution of 1536 × 64 before interpolation. For the iterative methods, 200 iterations were used to ensure convergence.

For the patient simulation, a voxel size of $0.5 \times 0.5 \times 0.5$ mm^3 was used for reconstruction, on a 1024 × 1024 × 200 grid. For the path-FBP, the projections then had a resolution of 1536 × 64 × 1536 before interpolation. For PLS and WPLS, 10 iterations were used, while 60 iterations were used for the TV-SUP method. This was based on the results obtained with the Gammex phantom.

Further details of the code can be found in previous publications (Hansen *et al* 2014a, 2014b) and the code is freely available at https://github.com/dchansen/protonCT. For all the iterative algorithms, 12 subsets were used as suggested by Penfold *et al* (2010a).

4. Evaluation

For the Gammex phantom, image resolution was evaluated by taking the task specific modulation transfer function (MTF) (Richard *et al* 2012) of the cortical bone insert of the phantom. The MTF was assumed Gaussian and the MTF = 10% (MTF$_{10\%}$) point is reported as has been done in previous studies (Hansen *et al* 2014b, Seco *et al* 2013). For stopping power accuracy, the mean error was evaluated for each of the phantom inserts, and the largest one is reported for each reconstruction method.

For the patient simulation, projections of the containing the integral stopping power was calculated via the Radon transform for 30 different angles in a 180° span, in a grid of 1450 × 198. Points for which the ground truth projection was less than 10 cm WEPL were excluded to avoid edge effects. The data was summarized in an error-volume histogram (similar to a dose-volume histogram). In this analysis, the background, as well as the first and the last slice of the images were excluded.

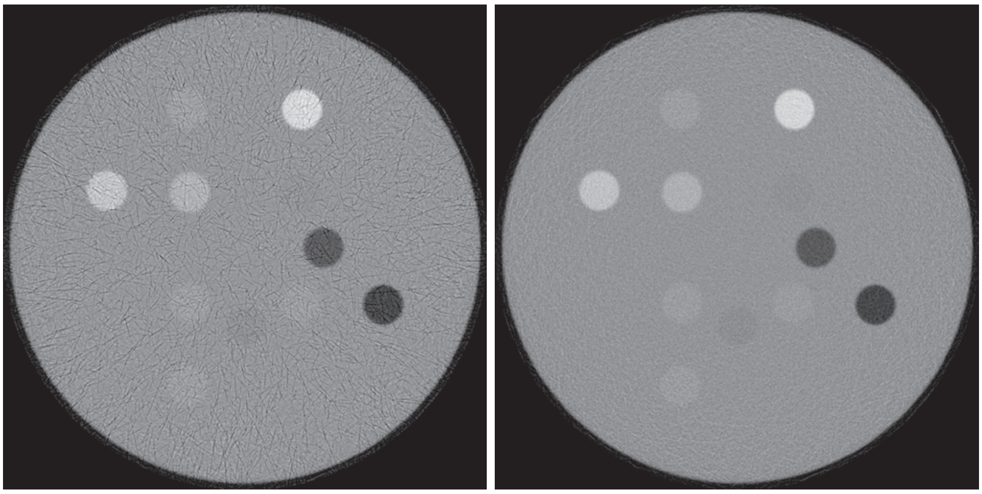

Figure 3. Left: 1 mSv proton CT of the Gammex phantom, reconstructed using the standard path-FBP. Right: proton CT reconstructed using path-FBP with projection interpolation. Window level is 0 to 2 WESP.

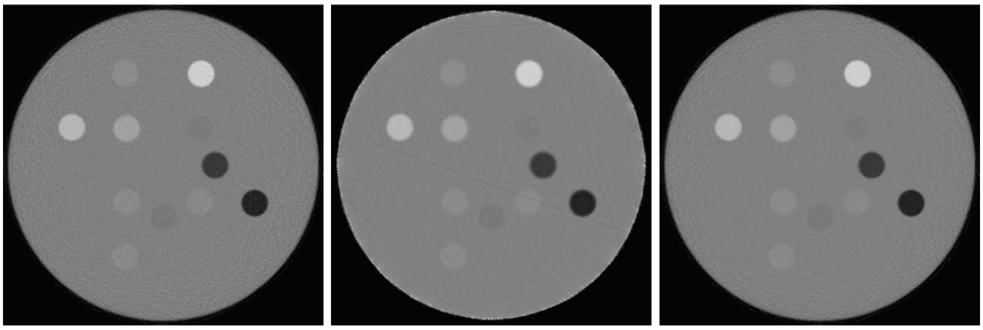

Figure 4. Iterative reconstruction algorithms after 200 iterations. Left: penalised least squares (PLS), middle: weighted penalised least squares (WPLS), right: total variation superiorisation.

5. Results

5.1. Gammex phantom

The standard path-FBP is displaying a high degree of artefacts, reducing the visibility of the inserts as shown on figure 3. In comparison, the version using projection interpolation is artefact free. For the iteratively reconstructed images (shown on figure 4), all three methods show clear and sharp images but also small noise artefacts near the edge of the phantom. Additionally, a few streak artefacts can be observed across the PWLS reconstruction. The resolution study is presented in figure 5. The PLS and PWLS methods both show a resolution of 4.0 lp cm^{-1}, while the TV-SUP method has a slightly higher resolution of 4.0 lp cm^{-1} and the path-FBP method a slightly lower resolution of 3.8 lp cm^{-1}. For the PLS and PWLS methods, the maximal resolution is reached after 9 iterations, while the TV-SUP method needs about 60 iterations to reach the level of the other methods, and more than 150 iterations to reach the final resolution. For the error study (figure 6), the PLS method reaches a stopping power error

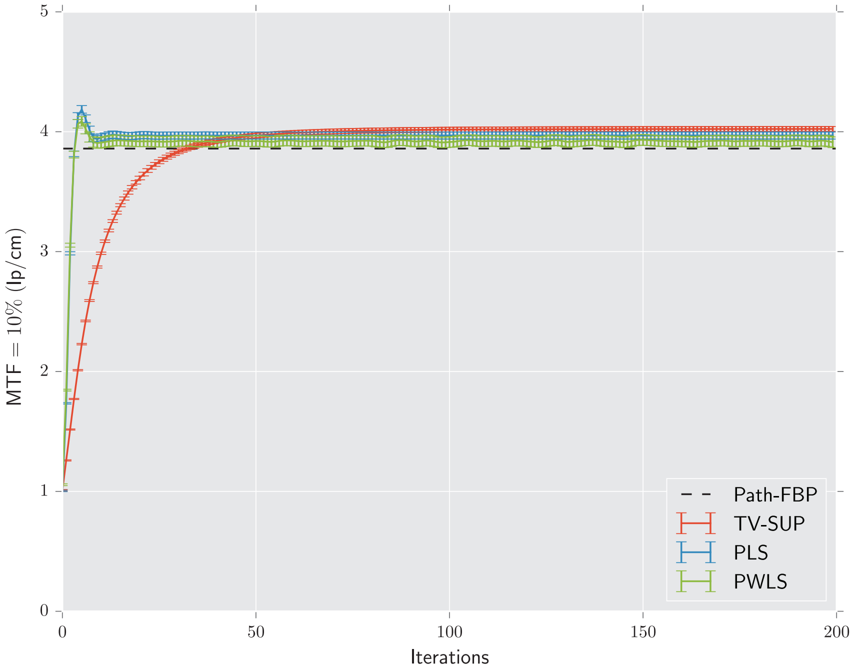

Figure 5. Plot of resolution (as calculated from the MTF = 10% point (Richard *et al* 2012)), as a function of the number of iterations for four different reconstruction algorithms. Note that the Path-FBP is not an iterative algorithm.

of 0.56 ± 0.11%, the PWLS method 0.89 ± 0.38%, TV-SUP 0.83 ± 0.04% and the Path-FBP 0.44 ± 0.10%. Both the PLS and PWLS methods have converged to their final value at 10 iterations, whereas TV-SUP is not converged until after 60 iterations. The total time for the Path-FBP reconstruction was 8 s. For the iterative algorithms, 200 iterations took 480 s for the PLS and WPLS, and 450 s for the TV-SUP method. Assuming the PLS and WPLS algorithms were converged after 10 iterations and the TV-SUP algorithm after 60, the reconstruction time was 25 s, 26 s and 138 s respectively.

5.2. *ICRP patient simulation*

Looking at the error projection histogram on figure 7, 91% of all beampaths of the PLS image are within 3.5% of the true WESP value. For WPLS, this is 0% and 69% for total variation superiorisation. The path-FBP method falls inbetween with 28% and 34%, with and without interpolation respectively. Looking at the reconstructions themselves as well as the reconstruction error (figures 8 and 9), a high degree of streak artefacts can be seen in the WPLS image, while both TV-SUP and path-FBP methods show a moderate degree of noise. In the PLS image, the noise is more suppressed, and it is possible to distinguish different tissue types visually.

Reconstruction time was 240 min for both the PLS and WPLS methods, 1130 min for TV-SUP and 12 min for the Path-FBP.

6. Discussion

One surprising result of this study is the low spatial resolution obtained as reported by the MTF measurements. While our findings are consistent with the values reported by Li *et al*

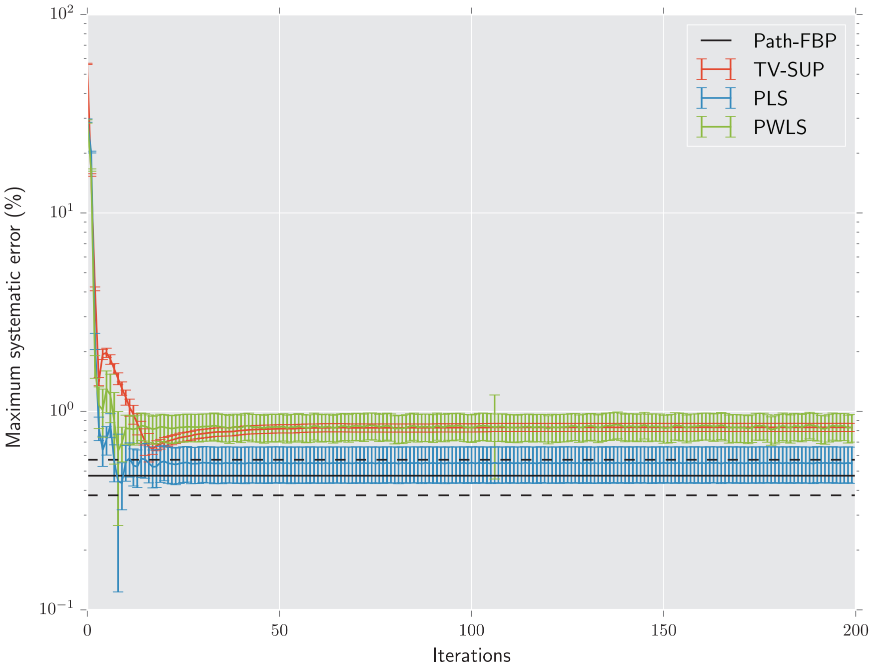

Figure 6. Plot of the largest mean stopping power error of the inserts in the phantom as a function of iterations for the four different reconstruction algorithms. Note that the Path-FBP is not an iterative algorithm.

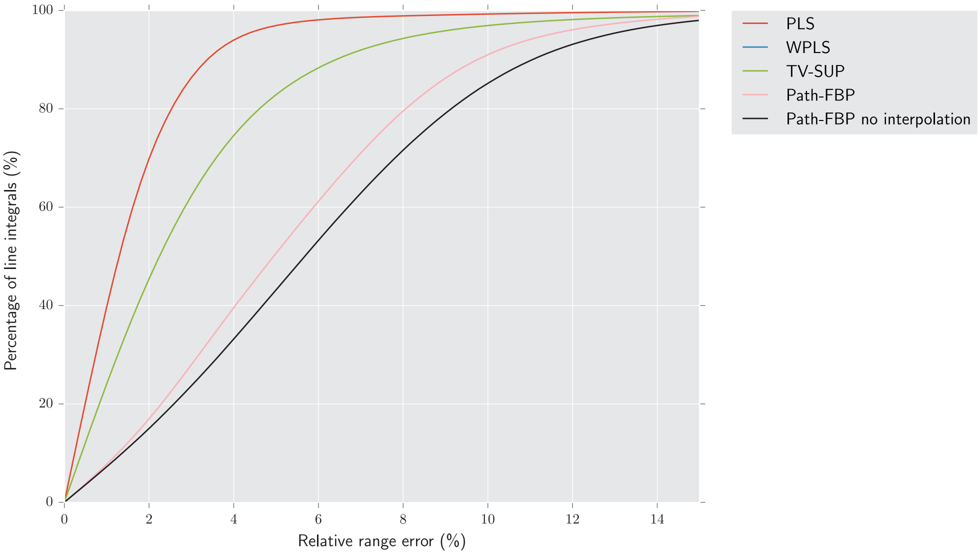

Figure 7. Plot of the relative error projection histogram of the patient simulation for the four different reconstruction algorithms. The plot denotes the percentage of beampaths below a specific error threshold.

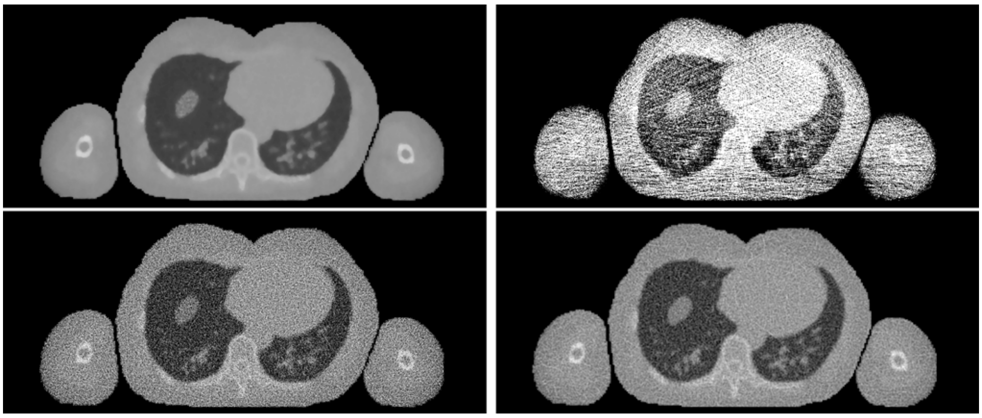

Figure 8. A mid slice of the simulated ICRP patient scan, reconstructed with the four different reconstruction algorithms. Top row, left: penalised least squares (PLS), right: weighted penalised least squares (WPLS). Bottom row, left: total variation superiorisation, right: path-FBP with interpolation. Window level was set from 0 to 2 WESP and the background was removed.

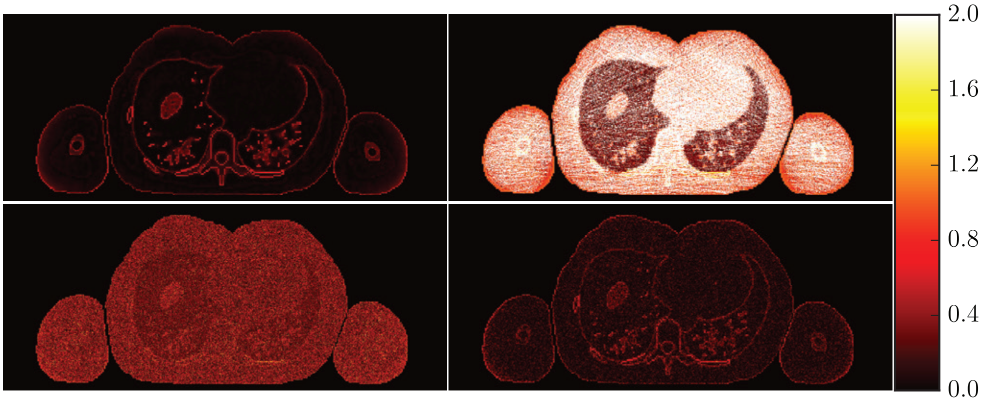

Figure 9. Images showing absolute error for a single slice of the simulated ICRP patient scan. Top row, left: penalised least squares (PLS), right: weighted penalised least squares (WPLS). Bottom row, left: total variation superiorisation, right: path-FBP with interpolation. Window level was set from 0 to 2 WESP and the background was removed.

(2006), later studies by Hansen *et al* (2014a), reported an $\mathrm{MTF}_{10\%} = 9\ \mathrm{lp\ cm}^{-1}$ and in Penfold *et al* (2010a) the $\mathrm{MTF}_{10\%}$ was reported to be over $60\ \mathrm{lp\ cm}^{-1}$. However, in both studies smaller phantoms were used, and in Hansen *et al* (2014a), detector resolution was not considered. In this study, while the TV-SUP algorithm does reach slightly higher resolution, it does so at the cost of accuracy. The convergence of the algorithm is also slower than what has been reported previously. We do however note that the data in this study is significantly more noisy, both due to the detector noise but also due to the larger phantom. This is well known to affect the performance of iterative optimization algorithms (Kelley 1999). For the patient phantom, TV-sup reached a much higher noise level than for the Gammex phantom, presumably due to the lower dose per voxel. It is possible that reducing the step-size and the number of subsets could have improved this, however as the authors were using previously recommended values, such modifications was considered beyond the scope of the present work.

Another result worth emphasizing is that the PWLS method performs worse than the simpler PLS method. There can be two explanations for this: first of all, the basic assumption of Gaussian distribution and uniform medium might not be sufficiently accurate, in particular for complex geometries such as a patient. Secondly, greater weight is given to protons that have lost little energy, meaning that the method is less robust if the estimated path of such events is not accurately predicted. This could for instance be the result of nuclear collisions. This explanation is consistent with the streak artefacts seen across the reconstructed Gammex phantom.

For the path-FBP, the resolution is slightly lower than the other methods due to the frequency cutoff in the Hann filter. Harder filters were considered but resulted in a much higher noise level and visually inferior image quality. As it is, the path-FBP has a significantly lower error than both the WPLS and TV-SUP methods and a much faster reconstruction time. While it is only three times faster than the PLS and WPLS method, it is 22 times faster than the TV-SUP method in 2D. It is expected that this will also hold at higher dose levels, as the forward and backprojection dominate the computational cost for both the path-FBP and iterative methods. Additionally, at higher dose levels the noise levels for the path-FBP are expected to drop significantly. In 3D, for the patient phantom, the speed gain was much more significant, with the path-FBP method being 20 times faster than PLS and WPLS, and almost 100 times faster than the TV-SUP method. This is due to the larger data size, which means that not all data can fit on the GPU at once, and there is a significant overhead in transferring data to and from the GPU. If the order of magnitude of the data was fixed (as would be the case in a clinical setting), the algorithms could be optimized to hide some of this overhead, reducing the speed difference.

For the ICRP phantom, the errors are generally larger than the often used 3.5% Yang *et al* (2012) for x-ray CT. The error for proton CT is however random, while x-ray CT yields systematic, tissue dependent, errors which may be larger for certain patients Yang *et al* (2012). Other authors have also reported stopping power errors as high as 8% in a clinical setting Moyers (2014). Even so, at these dose levels the error in the path-FBP images is larger than what would be used for treatment planning. They would still make a useful tool for identifying systematic deviations or changes in patient stopping power however.

While the PLS method yielded the best results for ICRP phantom, it is important to stress that the original phantom is piecewise constant. The Huber norm used for regularization was thus highly effective, as the phantom had many large and uniform areas. While it would have been possible to do simulations based on x-ray CT scans, these would have suffered from the noise and beam-hardening artefacts in the original x-ray images. The presented images are thus a good compromise, until real proton CT scans become feasible.

Further studies should also investigate the impact of each reconstruction technique on proton dose distributions as done with the path-FBP method by Arbor *et al* (2015).

7. Conclusion

In this study, we have compared three different iterative methods and a single direct method (path-FBP) for the reconstruction of low dose proton CT. The direct method, together with the iterative penalised least squares method yielded the lowest reconstruction error, while the difference in resolution of the investigated methods was less than 10%. We have also demonstrated that proton CT of clinically relevant size and resolution can be reconstructed in less than 15 min using a desktop using a single GPU. In short, the path-FBP algorithm provides high quality reconstructed images, while being an order of magnitude faster than previously published iterative algorithms and three times as fast as the fastest iterative algorithm presented in this paper.

Acknowledgments

This work was sponsored by the Danish Cancer Society (grant no. A5992) and the Lundbeck Centre for Interventional Research in Radiation Oncology (CIRRO).

Appendix. Algorithms

Algorithm 1. Momentum accelerated ordered subsets method with augmented Lagrangian multipliers for penalized weighted least squares (PWLS) (Kim *et al* 2015).

```
function OS-MOM($\mathbf{f}$)
        $\mathbf{x}^0 \leftarrow \mathbf{0}$
        $\mathbf{g}\mathcal{A}^\dagger\mathcal{W}^\dagger\mathcal{W}\mathcal{A}\mathbf{1}$                    $\triangleright$ SQS diagonal preconditioner
        $t^0 \leftarrow 1$
        for $k \leftarrow 1$ to iterations do
            for $l \leftarrow 1$ to subsets do                              $\triangleright$ Loop over subsets
                $m \leftarrow l + k \times$ subsets
                $\mathbf{y}^{m*} \leftarrow \text{subsets} \cdot \mathcal{A}_\mathbf{l}{}^\dagger(\mathbf{f}_\mathbf{l} - \mathcal{A}_\mathbf{l}\mathbf{x}^{m-1}) \oslash \mathbf{g}$      $\triangleright$ $\oslash$: elementwise-division
                $\mathbf{y}^m \leftarrow \text{DENOISE}(\mathbf{y}^{m*},\mathbf{g})$                $\triangleright$ Denoise using algorithm 2
                $t^m \leftarrow \frac{1}{2}\left(1 + \sqrt{1 + 4t^{m-1} \cdot t^{m-1}}\right)$
                $\mathbf{x}^m \leftarrow \mathbf{y}^m + \frac{t^{m-1} - 1}{t^m}(\mathbf{y}^m - \mathbf{y}^{m-1})$      Momentum acceleration step
    return $\mathbf{x}^{\text{final}}$
```

Algorithm 2. Arrow Hurwicz denoising (Arrow *et al* 1958, Chambolle and Pock 2011).

```
function DENOISE($\mathbf{x}^0, \mathbf{g}$)
    $\tau^1 \leftarrow 10^{-3}$
    $\gamma \leftarrow 0.35$
    $\sigma^1 \leftarrow \frac{1}{16\tau^1}$
    for $k \leftarrow 1$ to iterations do
        $\mathbf{p}^k \leftarrow y^{k-1} + \sigma^k\mathcal{K}\bar{\mathbf{x}}^{k-1}$                   $\triangleright$ $\mathcal{K}$ calculates the finite difference
        $y_i^k \leftarrow \frac{\frac{\mathbf{p}_i^k}{1+\sigma^k\alpha}}{\max\left(1, \left|\frac{\mathbf{p}_i^k}{1+\sigma^k\alpha}\right|\right)}$
        $\mathbf{x}^{k*} \leftarrow \mathbf{x}^{k-1} - \tau^k\mathcal{K}^\dagger\mathbf{y}^k + \frac{\tau^k}{\lambda}\mathbf{x}^0 \oslash \mathbf{g}$
        $\mathbf{x}^k \leftarrow \mathbf{x}^{k*} \oslash \mathbf{g}$
         $\theta^k \leftarrow \frac{1}{\sqrt{1 + 2\gamma\tau}}$
         $\tau^{k+1} \leftarrow \theta^k\tau^k$
         $\sigma^{k+1} \leftarrow \frac{\sigma^k}{\theta^k}$
      return $\mathcal{P}(\mathbf{x}^{\text{final}})$                             $\triangleright$ $\mathcal{P}$ sets all negative elements to 0
```

Algorithm 3. Diagonally relaxed ordered projections (DROP) with total variation superiorisation (TV-SUP) (Penfold *et al* 2010a).

```
function TV-DROP($\mathbf{f}$)
    $\gamma \leftarrow 1.9$
    $\mathbf{x}^0 \leftarrow \mathbf{0}$
    for $k \leftarrow 1$ to iterations do
        for $l \leftarrow 1$ to subsets do
            $m \leftarrow l + k \times$ subsets
            $\mathbf{b}_i^l \leftarrow \sum_{j \in \mathbf{l}} \|\mathcal{A}_{i,j}\|^2$    ▷ Calculate the squared norm of the rows of $\mathcal{A}_\mathbf{l}$
            $\mathbf{u}_i^l \leftarrow \max(1, \sum_{j \in \mathbf{l}} H(\mathcal{A}_{i,j}))$    ▷ $H$ is the Heaviside function.
            $\mathbf{x}^{m*} \leftarrow \gamma \mathcal{A}_\mathbf{l}^\dagger((\mathbf{f}_\mathbf{l} - \mathcal{A}_\mathbf{l}\mathbf{x}^{m-1}) \oslash \mathbf{b}^\mathbf{l}) \oslash \mathbf{u}^l$    ▷ $\oslash$: elementwise-division
            $\mathbf{x}^m, \beta^{m+1} \leftarrow$ SUPERIORISE($\mathbf{x}^{k*}, \beta^m$)    ▷ Using algorithm 4
    return $\mathbf{x}^{\text{final}}$
```

Algorithm 4. Total variation superiorization.

```
function SUPERIORISE($\mathbf{x}$, $\beta^0$)
    $\mathbf{y} \leftarrow \mathcal{P}(\mathbf{x})$    ▷ $\mathcal{P}$ sets all negative values to 0
    $\mathbf{s} \leftarrow \nabla \mathrm{TV}(\mathbf{y})$    ▷ Calculate total variation gradient
    $\mathbf{v} \leftarrow -\mathbf{s}/\|\mathbf{s}\|$
    $k \leftarrow 0$
    while $\mathrm{TV}(\mathbf{y} + \beta^k \mathbf{v}) > \mathrm{TV}(\mathbf{y})$ do
        $k \leftarrow k + 1$
        $\beta^k \leftarrow \beta^{k-1}/2$
    return $\mathbf{y} + \beta^k \mathbf{v}, \beta^k$
```

References

Allison J 2006 Geant4 developments and applications *IEEE Trans. Nucl. Sci.* **53** 270–8

Arbor N, Dauvergne D, Dedes G, Létang J, Parodi K, Quiñones C, Testa E and Rit S 2015 Monte Carlo comparison of x-ray and proton CT for range calculations of proton therapy beams *Phys. Med. Biol.* **60** 7585–99

Arrow K J, Hurwicz L and Uzawa H 1958 *Studies in Linear and Non-Linear Programming* ed H B Chenery *et al* (*Stanford Mathematical Studies in the Social Sciences* vol 2) (Stanford, CA: Stanford University Press)

Assmann W *et al* 2015 Ionoacoustic characterization of the proton Bragg peak with submillimeter accuracy *Med. Phys.* **42** 567–74

Beister M, Kolditz D and Kalender W A 2012 Iterative reconstruction methods in x-ray CT *Phys. Med.* **28** 94–108

Bennett G W, Archambeau J O, Archambeau B E, Meltzer J I and Wingate C L 1978 Visualization and transport of positron emission from proton activation *in vivo Science* **200** 1151–3

Censor Y, Davidi R and Herman G T 2010 Perturbation resilience and superiorization of iterative algorithms *Inverse Problems* **26** 65008

Chambolle A and Pock T 2011 A first-order primal-dual algorithm for convex problems with applications to imaging *J. Math. Imaging Vis.* **40** 120–45

Cormack A M 1963 Representation of a function by its line integrals, with some radiological applications *J. Appl. Phys.* **34** 2722

Hansen D C, Bassler N, Sørensen T S and Seco J 2014a The image quality of ion computed tomography at clinical imaging dose levels *Med. Phys.* **41** 111908

Hansen D C, Petersen B B Jr, Bassler N and Sørensen T S 2014b Improved proton computed tomography by dual modality image reconstruction *Med. Phys.* **41** 031904

Hansen D C, Seco J, Sørensen T S, Petersen J B B, Wildberger J E, Verhaegen F and Landry G 2015 A simulation study on proton computed tomography (CT) stopping power accuracy using dual energy CT scans as benchmark *Acta Oncol.* **54** 1638–42

Hansen M S and Sørensen T S 2013 Gadgetron: an open source framework for medical image reconstruction *Magn. Reson. Med.* **69** 1768–76

Huesman R, Rosenfeld A and Solmitz F 1975 Comparison of heavy charged particles and x-rays for axial tomograpic scanning *Technical Report* Technical Information Center US Department of Energy

Hünemohr N, Krauss B, Tremmel C, Ackermann B, Jäkel O and Greilich S 2014 Experimental verification of ion stopping power prediction from dual energy CT data in tissue surrogates *Phys. Med. Biol.* **59** 83–96

Hurley R F, Schulte R W, Bashkirov V A, Wroe A J, Ghebremedhin A, Sadrozinski H F W, Rykalin V, Coutrakon G, Koss P and Patyal B 2012 Water-equivalent path length calibration of a prototype proton CT scanner *Med. Phys.* **39** 2438

ICRP 2009 Adult reference computational phantoms *Ann. ICRP* 2

ICRU Report 49 and International Commision on Radiation Units and Measurements 1993 Stopping powers and ranges for protons and alpha particles *Technical Report* 49 ICRU Bethesda, MD

Karonis N T, Duffin K L, Ordoñez C E, Erdelyi B, Uram T D, Olson E C, Coutrakon G and Papka M E 2013 Distributed and hardware accelerated computing for clinical medical imaging using proton computed tomography (pCT) *J. Parallel Distrib. Comput.* **73** 1605–12

Kelley C 1999 *Iterative Methods for Optimization* (Philadelphia, PA: SIAM)

Kim D, Ramani S and Fessler J A 2015 Combining ordered subsets and momentum for accelerated x-ray CT image reconstruction *IEEE Trans. Med. Imaging* **34** 167–78

Knopf A C and Lomax A 2013 *In vivo* proton range verification: a review *Phys. Med. Biol.* **58** R131–60

Knopf A C, Parodi K, Paganetti H, Bortfeld T, Daartz J, Engelsman M, Liebsch N and Shih H 2011 Accuracy of proton beam range verification using post-treatment positron emission tomography/computed tomography as function of treatment site *Int. J. Radiat. Oncol.* **79** 297–304

Li T, Liang Z, Singanallur J V, Satogata T J, Williams D C and Schulte R W 2006 Reconstruction for proton computed tomography by tracing proton trajectories: a Monte Carlo study *Med. Phys.* **33** 699

Min C H, Kim C H, Youn M Y and Kim J W 2006 Prompt gamma measurements for locating the dose falloff region in the proton therapy *Appl. Phys. Lett.* **89** 183517

Moyers M F 2014 Comparison of x ray computed tomography number to proton relative linear stopping power conversion functions using a standard phantom *Med. Phys.* **41** 061705

Noo F, Clackdoyle R and Pack J 2004 A two-step Hilbert transform method for 2D image reconstruction *Phys. Med. Biol.* **49** 3903–23

Paganetti H 2012 Range uncertainties in proton therapy and the role of Monte Carlo simulations *Phys. Med. Biol.* **57** R99–117

Penfold S N, Schulte R, Censor Y, Bashkirov V, McAllister S, Schubert K E and Rosenfeld A B 2010a *Biomedical Mathematics: Promising Directions in Imaging, Therapy Planning and Inverse Problems* 1st edn, ed Y Censor *et al* (Madison, WI: Medical Physics Publishing) pp 347–367

Penfold S N, Schulte R W, Censor Y and Rosenfeld A B 2010b Total variation superiorization schemes in proton computed tomography image reconstruction *Med. Phys.* **37** 5887

Poludniowski G, Allinson N M and Evans P M 2014 Proton computed tomography reconstruction using a backprojection-then-filtering approach *Phys. Med. Biol.* **59** 7905–18

Richard S, Husarik D B, Yadava G, Murphy S N and Samei E 2012 Towards task-based assessment of CT performance: system and object MTF across different reconstruction algorithms *Med. Phys.* **39** 4115

Rit S, Clackdoyle R, Hoskovec J and Letang J M 2015 List-mode proton CT reconstruction using their most likely paths via the finite Hilbert transform of the derivative of the backprojection *Fully3D* (submitted)

Rit S, Dedes G, Freud N, Sarrut D and Létang J M 2013 Filtered backprojection proton CT reconstruction along most likely paths *Med. Phys.* **40** 031103

Sadrozinski H F W W *et al* 2013 Development of a head scanner for proton CT *Nucl. Instrum. Methods Phys. Res.* A **699** 205–10

Schulte R W, Penfold S N, Tafas J T and Schubert K E 2008 A maximum likelihood proton path formalism for application in proton computed tomography *Med. Phys.* **35** 4849–56

Seco J, Oumano M, Depauw N, Dias M F, Teixeira R P and Spadea M F 2013 Characterizing the modulation transfer function (MTF) of proton/carbon radiography using Monte Carlo simulations *Med. Phys.* **40** 091717

Sipala V *et al* 2011a PRIMA: an apparatus for medical application *Nucl. Instrum. Methods Phys. Res.* **658** 73–7

Sipala V *et al* 2011b YAG(Ce) crystal characterization with proton beams *Nucl. Instrum. Methods Phys. Res.* **654** 349–53

Tschalär C 1968 Straggling distributions of large energy losses *Nucl. Instrum. Methods* **61** 141–56

Wang D, Mackie T R and Tomé W A 2011 Bragg peak prediction from quantitative proton computed tomography using different path estimates *Phys. Med. Biol.* **56** 587–99

Williams D C 2004 The most likely path of an energetic charged particle through a uniform medium *Phys. Med. Biol.* **49** 2899

Yang M, Virshup G, Clayton J, Zhu X R, Mohan R and Dong L 2010a Theoretical variance analysis of single- and dual-energy computed tomography methods for calculating proton stopping power ratios of biological tissues *Phys. Med. Biol.* **55** 1343–62

Yang M, Zhu X R, Park P C, Titt U, Mohan R, Virshup G, Clayton J E and Dong L 2012 Comprehensive analysis of proton range uncertainties related to patient stopping-power-ratio estimation using the stoichiometric calibration *Phys. Med. Biol.* **57** 4095–115