Novel Optimization Strategies for Clinical FLASH Proton Therapy

Evaluation on stereotactic lung treatment plans

Rodrigo José Santo
rodrigo.j.santo@tecnico.ulisboa.pt

Instituto Superior Técnico, Lisboa, Portugal

December 2021

Abstract

The FLASH effect has gained increasing interest due to its higher healthy-tissue sparing than conventional radiotherapy. This has been observed for irradiation of high doses (> 8Gy) at ultra-high dose rates (> 40 Gy/s). Combined with the dosimetric advantages of proton therapy and its adequacy to treat deep-seated tumours, FLASH-PT has the potential to reduce toxicity and improve clinical outcome in some patients. However, current treatment planning software is unable to optimize FLASH.

In this project, the aim was at developing novel strategies for optimization of FLASH-PT within clinical dose requirements, that reduce side effects through FLASH. This is currently limited to cyclotron-accelerated beams and pencil-beam scanning. For the highest dose rates, the maximum commissioned cyclotron energy (244 MeV) is used, corresponding to shoot-through transmission beams, on stereotactic treatment of lung lesions. Dose rate is optimized via beam intensity, beam current and scanning-pattern optimization, based on the Dose-Averaged Dose Rate (DADR) and the Pencil-Beam Scanning Dose Rate (PBSDR), using iterative linearization, iterative convex relaxation and Genetic Algorithms.

A significant increase of dose rate is achieved with the proposed optimization strategies, through high beam currents and optimized snowflake-shaped scanning patterns. Based on current knowledge of FLASH, this may be of significant clinical benefit. Before clinical application, sensitivity to treatment machine parameters needs to be evaluated. Results can be further improved, weighting tissue sparing and trade-off, optimizing beam direction and partially irradiating with FLASH.

Keywords: FLASH, proton therapy, pencil-beam scanning, dose-rate optimization, scanning-pattern optimization

1 Introduction

For decades, the number of patients diagnosed with cancer has been increasing, in part due to a higher life expectancy. Mortality has been decreasing, however death rates are still high and treatments have unavoidable side effects.

Radiotherapy is one of the most important cancer treatment modalities, alongside surgery and systematic therapies (e.g. chemo and immunotherapies). Its efficacy is based on the fact that tumour cells are often more radiosensitive than healthy cells and are less capable of repairing damage. This differential effect is further increased through high-precision irradiation techniques, such as proton therapy, and through optimal targeting, using online imaging techniques. Unfortunately, patients may still experience side effects, with potentially severe impact on their quality of life.

Recent studies point out a fundamentally different way to optimize the differential damage: the FLASH effect. Healthy cells irradiated with high dose (>8 Gy) at very high dose rates (>40 Gy/s) have been seen to become up to 40% less damaged than in conventional irradiation, without compromising tumour cell kill [1][2]. Ultra-fast irradiation also enables mitigation of motion uncertainties.

The FLASH effect has been demonstrated for electrons [1] and proton beams [2], with the first human patient recently treated for a superficial lymphoma [3]. Electrons are suitable only for superficial tumours, whereas protons allow the treatment of deep-seated infiltrating tumours, for which side

effects may be much more severe and so may benefit significantly from FLASH. This has increased the interest in FLASH proton therapy (FLASH-PT).

FLASH-PT is currently limited to cyclotron-accelerated beams using pencil-beam scanning (PBS). Synchrotrons have a pulsed structure on the other of tenths to a few second, too slow for FLASH, while passive scattering significantly reduces dose rates, being only suitable for small targets. The required high dose rates are readily available with current clinical cyclotrons. The overall highest are achieved at the maximum cyclotron energy (250 MeV), with no energy modulation currently possible at FLASH-compatible times. This corresponds to transmission beams (TB), that shoot through the patient. Although relatively non-precise compared with Bragg peak beams, TBs mitigate range uncertainties and freeze anatomical motion. Despite analogous to photon beams, TBs have a flat dose-depth profile and a sharper lateral dose fall-off.

With PBS, the target volume is sequentially irradiated with local pencil beams, which makes quantifying the dose rate and evaluating FLASH difficult, since experimental studies have focused on simultaneous whole field irradiation. Two metrics have been proposed: the Dose-Averaged Dose Rate (DADR) and the Pencil-Beam Scanning Dose Rate (PBSDR). The DADR accounts for the local contribution of each pencil beam, weighting their instantaneous dose rate by their dose contribution. The PBSDR additionally accounts for the local time structure, through the time interval between the first and last pencil beams contributing >1 cG. It depends on the order each pencil beam is delivered.

In radiotherapy, patients are treated with personalized treatment plans with computer-optimized machine settings to deliver a curative dose to the tumour, while sparing normal tissue. Dose rate is not accounted and algorithms to optimize it are not available or have not been validated yet. This is challenging due to the non-convex and non-linear nature of the metrics, which compete with dose requirements. Indirect approaches on treatment configurations have been successful, although with suboptimal results and little control over trade-offs [4][5]. Recently, a simultaneous dose and DADR optimizer was proposed [6], which optimizes dose while satisfying dose rate requirements, although not necessarily maximizing FLASH.

In this project, the main objective is to reduce side effects in clinical radiotherapy through the FLASH effect, developing novel strategies for optimization and delivery of FLASH-PT plans within clinical constraints on dose, based on the present knowledge and clinical technology. Two direct FLASH optimization strategies are proposed using lung tumours as the first foreseen clinical application: based on the DADR, with full control over trade-offs, and on the PBSDR, a first-of-its-kind to the best of found knowledge. Evaluation was performed on analytical 2D models and clinical data, but focus is here set exclusively on the latter.

2 Background

Treatment planning software bridges the medical treatment requirements with the treatment delivery by the machine. Inputs include CT-scan delineations of important structures in the patient and a set of constraints, objectives and goals for the dose (wish-list). For PBS-PT, the optimized plans are Intensity Modulated Proton Therapy (IMPT) plans, the golden standard, and correspond to the optimal pencil beams, their location and intensity (weight).

The FLASH optimization strategies are implemented in the Erasmus MC in-house developed treatment planning software, the Erasmus-iCycle. It stands out due to its prioritized multi-criteria optimization [7], allowing to better specify clinical preferences, leading to treatment plans of exceptional quality. This is well-suited for clinical FLASH optimization, offering full control over the trade-offs.

Optimization is performed for 12 patients with one lung lesion, for 54 Gy/3 fractions with 3 equiangular coplanar beams, on a single beam per fraction regime, having median PTV of 6.4 cc and range 4.4-10.1 cc. PTV-based planning with a 5mm margin is applied. TBs with energy of 244 MeV, the maximum available in the in-house beam model, are used. For the PBSDR, 3 additional patients are considered, characterized by larger PTVs (23.6 cc, 52.9 cc, 83.9 cc).

3 Methods

3.1 Dose-Averaged Dose Rate

DADR optimization was introduced as two new objectives (mean and min) that can be specified at the wish-list, making it simple to adapt conventional treatments to account for the FLASH effect.

Direct optimization of the DADR is challenging because it is a ratio. Considering D_{ij} the dose-disposition matrix and dr_{ij} the instantaneous dose rate at voxel i by pencil beam j, with weight w_j, the DADR at voxel i is given by:

$$DADR_i = \sum_{j=1}^{n} \frac{D_{ij}w_j \times dr_{ij}}{\sum_{k=1}^{n} D_{ik}w_k} = \frac{N(w)}{D(w)} \quad (1)$$

However, the instantaneous dose rate dr_{ij} at voxel i by pencil beam j does not depend on the pencil-beam weight. It only depends on the beam current I_j and D_{ij}. Defining the delivery time of each pencil beam as $T_j = I_j w_j$, it follows:

$$dr_{ij} = \frac{d_{ij}}{T_i} = \frac{D_{ij}w_j}{I_j w_j} = I_j D_{ij} \quad (2)$$

Therefore, the DADR is a ratio of linear positive expressions. In addition, the hard constraints that must be satisfied are all linear since the dose constraints and objectives are linear (min, max, mean), while the DADR constraints can be linearized. To comply with clinical requirements, the DADR objectives are only applied on healthy tissue, after the dose objectives. Consequently, the search space is a bounded polyhedron, which is important for the convergence of the proposed strategies.

Linear Fractional Programming is a deeply studied topic, with extensive literature on optimization methods. The implementation is based on [8]: to transform the linear fractional program into a series of linear programs. This is illustrated on a simpler linear fraction program $f(x)$, with $x \in X$ where $X \in \mathbb{R}^n$ is a bounded polyhedron:

$$f(x) = \frac{ax+b}{cx+d} = \frac{N(x)}{D(x)} \quad (3)$$

If the objective function $f(x)$ is maximized, its optimum solution is $f(x^*) = q^*$ and the denominator is positive, it is true for any $x \in X$:

$$\frac{N(x)}{D(x)} \leq q^* \Leftrightarrow N(x) - q^* D(x) \leq 0 \quad (4)$$

The previous inequality is only 0 for x^*, meaning that if the optimum q^* is known, the optimization of ratio $f(x)$ is equivalent to:

$$\max_{x \in X} F(x, q^*) = N(x^*) - q^* D(x^*) = 0 \quad (5)$$

The difficulty is in finding q^* but it can be iteratively approximated by:

$$q_{k+1} = f(x_k) = \frac{N(x_k)}{D(x_k)} \quad (6)$$

$$x_{k+1} = \arg\max_{x \in X} \left[N(x) - q_{k+1} D(x)\right] \quad (7)$$

Which can be shown to converge to x^* and q^* in a finite number of steps, with $\arg\max$ the arguments of the maxima. This way, the linear fractional program is solved through a series of linear programs.

Maximization of the DADR objectives follows the same rationale but with several ratios $f_i(x)$, one for every voxel in the target volume. For the minimum DADR, the routine takes the form:

$$q_{k+1} = \min f(x_k) = \min \frac{N_i(x_k)}{D_i(x_k)} \quad (8)$$

$$x_{k+1} = \arg\max_{x \in X} \min \left[N_i(x) - q_{k+1} D_i(x)\right] \quad (9)$$

In Erasmus-iCycle's multi-criteria optimization, after each objective is optimized, it is transformed into a constraint. Considering the minimum constraint value q^-, the minimum DADR objective can be rewritten as a liner constraint:

$$\begin{aligned} \min f_i(x) \geq q^- \Rightarrow & f_i(x) = \frac{N_i(x)}{D_i(x)} \geq q^- \\ & \Leftrightarrow N_i(x) - q^- D_i(x) \geq 0 \\ & \text{for all } i \end{aligned} \quad (10)$$

For the mean DADR, the objective function is formulated as a multiobjective problem, since the goal is to maximize the DADR everywhere. The previous rationale can also be used:

$$q_{k+1} = \left[\ldots, f_i(x_k) = \frac{N_i(x_k)}{D_i(x_k)}, \ldots\right] \quad (11)$$

$$X_{k+1} = X \cap \{x \in \mathbb{R}^n | N_i(x) - q^i_{k+1} D_i(x) \geq 0\}$$

$$x_{k+1} = \arg\max_{x \in X_{k+1}} \sum_i \left[N_i(x) - q^i_{k+1} D_i(x)\right] \quad (12)$$

To transform the mean DADR in a linear constraint, the dose at the target volume is considered fixed because optimization is only recommended after the dose and the minimum DADR objectives.

The denominator can then be approximated by the dose for the starting solution x_0. For a minimum mean DADR constraint value q^-, this means:

$$\sum_i f_i(x) = \sum_i \frac{N_i(x)}{D_i(x)} \approx \sum_i \frac{N_i(x)}{D_i(x_0)} \geq q^- \quad (13)$$

Beam current also increases the dose rate, but it comes with a trade-off on dose since it constrains the minimum pencil-beam weight w_{min} due to treatment-delivery machine specific minimum pencil-beam delivery time t_{min}:

$$w_{min} = I \times t_{min} \Rightarrow d_{min} = \mathbf{D} \cdot w_{min} \quad (14)$$

Optimization of the beam current is challenging because it requires increasing the weights and selecting the most suitable pencil beams. The approach used is similar to [9]: alternate between optimizing the objectives and selecting the best beams. This is performed with algorithm 1.

Algorithm 1: Minimum-Weight Optimization

for N *iterations* **do**
  while *True* **do**
    $w \leftarrow \forall\, w'_i > w_{min}$
    $w' \leftarrow \arg\min\{f(w)|w \in W\}$
    if *Fails* **then**
      $w \leftarrow \forall\, w_0 > w_{min}$
      $w' \leftarrow \arg\min\{f(w)|w \in W \wedge w \geq w_{min}\}$
      if *Fails* **then** $w' \leftarrow w_0$ **and** Stop
    if $\forall\, w'_i > w_{min}$ **then** Stop
  $w_{min} \leftarrow \min w'$

The routine is applied to constrain and optimize the minimum weight. For the former case, $f(x)$ corresponds to the function of the objective being optimized, while for the latter it is set to 0, corresponding to finding a feasible solution. To help maximizing the minimum weight, a seed $w \times \delta$ is used, with the restart factor δ set to 1.05.

To transform DADR optimization in FLASH maximization, only the voxels at the target volume with dose higher than the FLASH dose threshold are optimized. After optimization, the transformed constraint is only applied to the voxels with dose and dose rate higher than the FLASH thresholds.

3.2 Pencil-Beam Scanning Dose Rate

PBSDR optimization is performed only at the scanning-pattern level, on IMPT treatment plans, with no changes on the dose distributions. This approached is used because the scanning pattern has a great impact on the dose-rate distribution.

Scanning-pattern optimization (SPO) is a combinatorial optimization problem, similar to the Traveling Salesman Problem (TSP). The solution for both is a pattern and the objective is a global metric of that pattern, comprised of local features: for the classic TSP this corresponds to the distance between pairs of cities, which adds up to the total path length, whereas for the SPO it corresponds to the local dose rate by contributing pencil beams, making up the FLASH coverage. Therefore, strategies for solving the TSP offer a starting point for SPO.

Genetic Algorithms (GA) are often discussed in literature to solve TSPs [10]. These are stochastic algorithms that work similarly to natural selection, using biologically inspired operators, such as mutation and crossover, to generate high-quality solutions. For TSPs, given an initial sample of possible routes, they mix, match and modify smaller paths connecting just a few cities and preserve the best, in search of the overall optimal routes. Local features get preserved, which is important considering that the PBSDR is a voxel-wise dose rate: short sets of nearby pencil beams should give higher dose rates on the locally-irradiated voxels.

An Island Model GA (IMGA) is implemented, to maximize FLASH coverage everywhere but the PTV. Pencil beams are assigned a unique number, so patterns can be represented as an ordered array. Population is randomly initialized. Mutation is performed using swap, insertion, simple inversion and shift operators. Partially-mapped ordered crossover (PMOX) is used, as in figure 1.

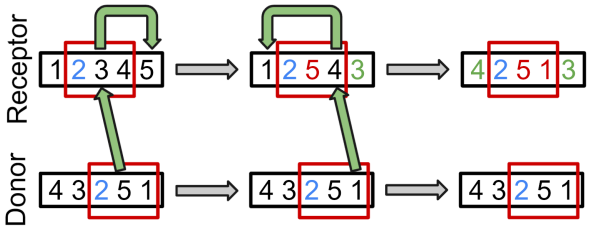

Figure 1: Partially-Mapped Ordered Crossover.

The IMGA should be tuned, to ensure the best solution in the least time, within the available resources. This is not critical though, since SPO has

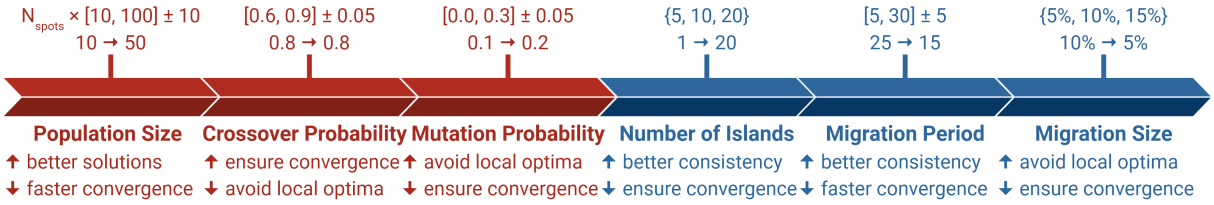

Figure 2: IMGA fine-tune routine, with the parameters and explanation on the bottom, alongside the range, default and optimized values on top.

no trade-off on dose and only aims to further increase the gains by FLASH. The fine-tune routine in figure 2 aimed at consistently finding the best solutions, starting with values used in TSPs [10] and testing for an arbitrary small number of iterations.

3.3 Results Evaluation

The results are evaluated on the FLASH-enhanced mean dose to the ipsilateral lung except the GTV. This is a good proxy for damage to healthy lung, since lung has a parallel organization scheme (it maintains function even if a fraction is sacrificed). It also depends linearly on dose, which makes it easy to calculate an effective FLASH enhancement ratio (FER). Comparison is performed against IMPT plans, the golden standard.

To model the FLASH gains, a FER is considered, by taking the effect as binary based on the dose and dose rate thresholds. FLASH-enhanced dose distributions are obtained by dividing the dose at healthy tissue by this ratio, modeling a reduced toxicity over IMPT plans. For lung tissue, a FER of at least 1.8 has been observed [1].

Using different thresholds for the FLASH effect, both on dose (6 Gy, 8 Gy, 10 Gy and 12 Gy) and dose rate (30 Gy/s, 40 Gy/s, 50 Gy/s, 60 Gy/s), enhanced dose distributions are calculated to evaluate the sensitivity of the solutions regarding FLASH uncertainties on the trigger conditions.

4 Results

4.1 Dose-Averaged Dose Rate

Fully-optimized treatment plans (dose followed by DADR optimization at every structure) are generated. In figure 3, the FLASH coverage is reported, for the minimum-required and maximum-compatible beam currents obtained by the optimizer, against IMPT plans, for which the beam current was calculated with a 5 ms minimum pencil-beam delivery time and the respective minimum weight. All optimized plans have full FLASH coverage, with some IMPT plans partially compatible.

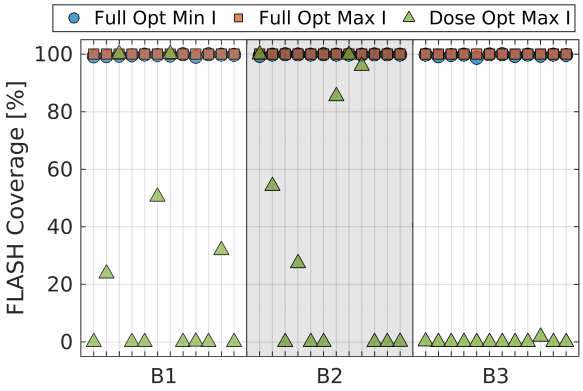

Figure 3: FLASH coverage for fully-optimized and conventional IMPT treatment plans.

To evaluate at what cost the full FLASH coverage is achieved for the fully-optimized plans, the mean dose at the ipsilateral lung is compared against IMPT plans in figure 4. Fully-optimized plans have a non-FLASH mean dose identical to IMPT plans and a substantially lower FLASH-enhanced dose. On direction *B3*, the non-FLASH mean dose for fully-optimized plans is lower than for IMPT plans.

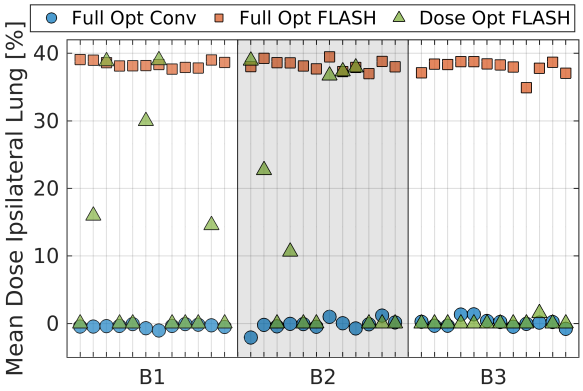

Figure 4: Mean dose differences at the ipsilateral lung for fully-optimized plans relative to IMPT plans.

[1] Obtained from thoracic irradiation of mice[1], for which 30 Gy at FLASH conditions was seen to be less fibrogentic than 17 Gy conventional (FER = 30/17≈1.8).

The mean dose for the fully-optimized plans has a median deterioration of 0.37%, 0.11% and -0.18% for each of the three directions, with an interquartile range of 0.28%-point, 0.56%-point and 0.64%-point. On improvements, results show medians of 38.2%, 38.1% and 38.3% for directions *B1*, *B2* and *B3*, respectively, alongside an interquartile range of 0.8%-point, 0.9%-point and 1.1%-point.

Full optimization calculates the minimum beam current required to achieve optimal FLASH coverage, but it also attempts to make the treatment plans compatible with higher currents by maximizing the minimum weight at the last step. This translates into windows of beam currents compatible with the plans, as reported in figure 5.

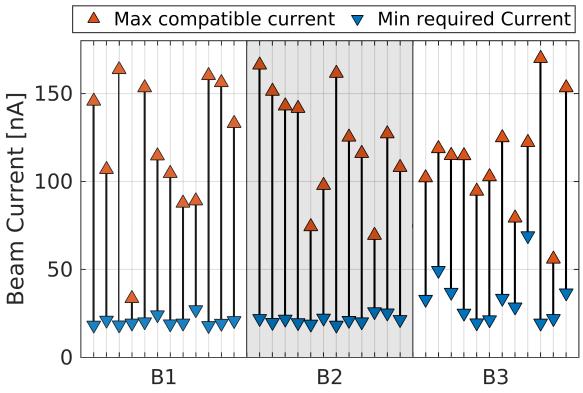

Figure 5: Beam current windows for fully-optimized treatment plans that guarantee optimal FLASH coverage within the wish-list requirements.

The majority of the plans is compatible with a wide window of currents. For directions *B1*, *B2* and *B3*, the minimum-required beam current has a median of 19 nA, 21 nA and 31 nA, alongside an interquartile range of 2 nA, 2 nA and 15 nA respectively. For the beam-current window of [33,56] nA, 80% of the plans are fully compatible.

To evaluate the sensitivity of the fully-optimized treatment plans to FLASH uncertainties, FLASH-enhanced dose distributions are calculated using various thresholds and the minimum beam currents, considering the overall treatment of the patient by summing up the doses from each direction.

Treatment plans are evaluated on the mean dose to the ipsilateral lung in figure 6, with the median improvements reported in table 1. The results show an overall consistency, with an expected degradation for higher dose and dose-rate thresholds.

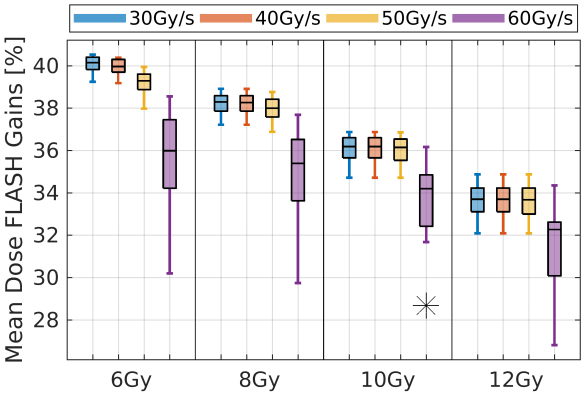

Figure 6: Sensitivity to different FLASH thresholds of the FLASH improvements on the mean dose over IMPT plans for fully-optimized plans.

Dose\Rate	30 Gy/s	40 Gy/s	50 Gy/s	60 Gy/s
6 Gy	40.1%	30.0%	39.3%	36.0%
8 Gy	38.3%	38.3%	38.0%	35.4%
10 Gy	36.2%	36.2%	36.1%	34.2%
12 Gy	33.7%	33.7%	33.7%	32.3%

Table 1: Median mean dose FLASH improvements for fully-optimized plans on FLASH thresholds.

4.2 Pencil-Beam Scanning Dose Rate

Scanning-pattern optimization is performed with the version that gave the best results during validation: IMGA using 20 islands for 20 independent runs. A beam current of 40 nA is used. The optimized patterns are evaluated by comparing the FLASH improvements on the mean dose to the ipsilateral lung over non-FLASH dose, for the best and worst-reported patterns against predefined ones, including the clinical golden standard (*Snake*). These results are reproduced in figure 7.

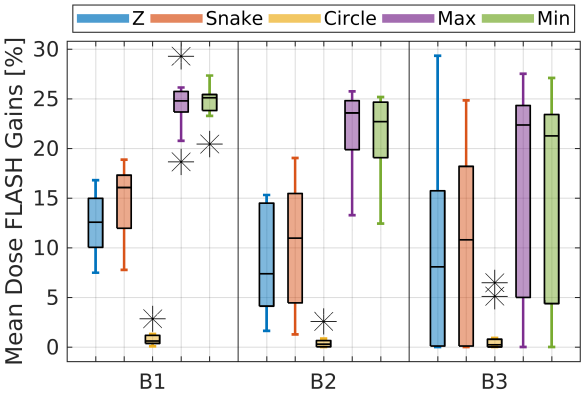

Figure 7: Mean dose improvements at the ipsilateral lung over non-FLASH for different patterns.

The optimized patterns are considerably better than the predefined ones, with higher FLASH gains. Relative median improvements over the best predefined patterns of 9.3%-point, 10.8%-point and 5.5%-point are obtained for the optimized patterns, on directions *B1*, *B2* and *B3*, with interquartile ranges of 4.8%-point, 4.8%-point and 6.2%-point. The median difference between the best and worst optimized patterns is 0.4%-point, 0.3%-point and 0.3%-point, with interquartile ranges of 0.8%-point, 0.9%-point and 1.0%-point. Average running times of 8 minutes are obtained, corresponding to a median of 34 pencil beams.

To evaluate the scalability of the optimizer, since it was fine-tuned for relatively small lung lesions, patterns were optimized for patients with larger PTVs of 23.6 cc, 52.9 cc and 83.9 cc. Despite the increased complexity, the optimized patterns have a consistently higher FLASH coverage than the predefined ones. For those volumes respectively, relative median improvements of 22.3%-point, 17.2%-point and 17.0%-point are obtained for the optimized patterns over the best predefined. The median difference between the best and worst optimized patterns is 3.6%-point, 7.2%-point and 1.7%-point. The average running time for a single run is 24 minutes, 3 hours and 4 hours, with 64, 107 and 159 median numbers of pencil beams.

To understand how the optimal FLASH coverage and consequently the highest potential tissue sparing is achieved, the best optimized patterns for the larger three PTVs plus another with value 8.2 cc are reproduced on top of the corresponding PBSDR distributions in figure 8, along beam direction *B1*, in an arbitrary slice.

The optimized patterns consistently show a similar snowflake shape: closed circular loop, following a radial inwards-outwards movement around a central region, closing the circle in the same initial radial direction. As the tumour volume increases, the optimized patterns have their start and end points more and more to the center. For the largest patterns, the closed loop more resembles a swirl, wrapping around itself. The corresponding PBSDR distributions have a *C* shape for the smaller patterns and a swirl for the larger ones.

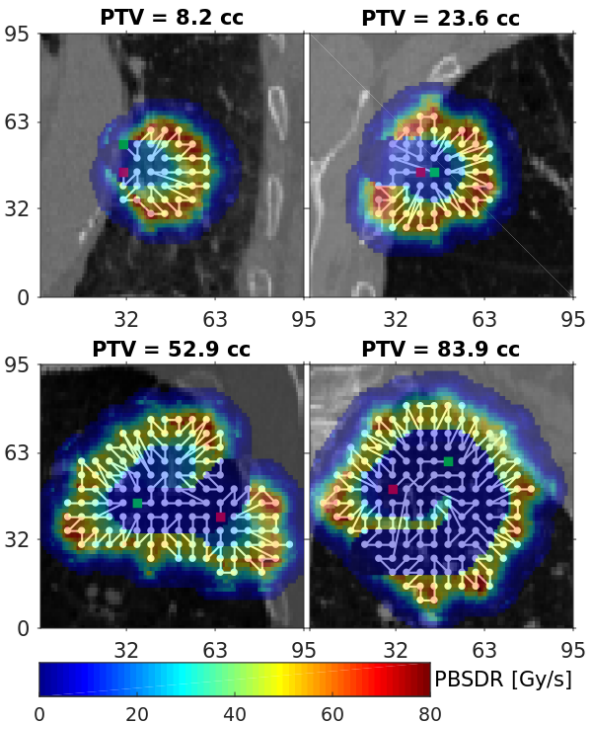

Figure 8: PBSDR distributions of the best patterns for varying lung tumour volumes, overlapped with the pattern, starting in green and ending in red.

The sensitivity of solutions to FLASH uncertainties is evaluated on the mean dose to the ipsilateral lung in figure 9. This is calculated from FLASH-enhanced dose distributions for each direction, considering various thresholds and for the best-performing scanning patterns, summed up to translate the overall treatment for every patient.

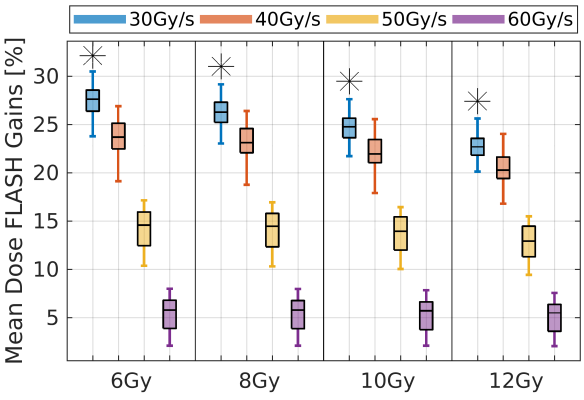

Figure 9: Sensitivity of the improvements on the mean dose at the ipsilateral lung for the best scanning patterns on different FLASH thresholds.

The results are relatively consistent on the dose threshold but change considerably on the dose rate threshold. The median improvements of the FLASH-enhanced mean dose at the ipsilateral lung

over the non-FLASH dose are reported in table 2.

Dose\Rate	30 Gy/s	40 Gy/s	50 Gy/s	60 Gy/s
6 Gy	26.4%	22.5%	13.3%	4.7%
8 Gy	25.1%	22.1%	13.1%	4.7%
10 Gy	23.6%	21.2%	12.8%	4.6%
12 Gy	21.8%	19.5%	12.0%	4.3%

Table 2: Median mean dose improvements for best optimized patterns on various FLASH thresholds.

5 Discussion

5.1 Dose-Averaged Dose Rate

The optimization strategies generate plans with improved FLASH coverage, as reported in figure 3. Some IMPT plans are already FLASH compatible, which for lung tumours is due to the high prescription dose of 54 Gy/3 fractions (hyperfractionation) and single beam per fraction using TBs. This translates in higher minimum weights and beam currents, leaving more room for optimization.

The mean dose at the ipsilateral lung in figure 4 shows that there is room for improvements on the dose rate, with minimum trade-off on dose. In some cases, better non-FLASH doses are obtained, due to the more thorough pencil-beam reduction performed for the DADR objectives. Whenever IMPT plans are already fully FLASH compatible, the FLASH gains might need to compensate for the dose trade-off. This is because optimization is performed per-structure, so it is possible for dose rate at one structure to limit the dose optimization at the following structures.

Figure 5 shows compliance to different setups, as the fully-optimized plans are compatible with a wide range of beam currents. The optimization strategies iteratively increase the current in small steps before evaluating compatibility with much higher values, explaining the wide windows. The maximum-compatible current is tied to the minimum weight, so plans requiring more pencil-beam modulation are expected to be only compatible with smaller maximum currents, since the best dose distributions depend on more pencil beams, each delivering relatively less dose. The need for more modulation comes from organs at risk (OARs) shot-through by TBs.

The existence of a window of beam currents compatible with most plans reveals that when optimized, beam current can shift from a patient-specific parameter to a generic treatment parameter. Fixing the beam current within that window in the treatment planning software is guaranteed to deliver optimal FLASH-compatible plans, with the least trade-offs. This is an important result because it might not be feasible to change the beam current on a patient basis within the clinical workflow, at least with current clinical technology.

Direction *B3* has less consistent results, with IMPT plans having no FLASH coverage, while fully-optimized plans have higher median beam current and wider interquartile range compared with other directions. This illustrates how FLASH can be optimized indirectly, in this case through the beam direction. This is due to a larger path length from the beam entrance to the tumour and can be critical. As the beams go through the patient, scattering increases and so does pencil beam overlap, translating into lower dose rates. This can be amplified by dense structures in the beam path. Nonetheless, the optimization strategies demonstrate compliance to non-optimal setups, as full FLASH coverage is always reported.

Despite specifically optimized for 8 Gy and 40 Gy/s, the treatment plans have consistently high FLASH gains for different thresholds as reported in figure 6 and table 1. To some extent, the thresholds are not critical, even when considering the minimum-required beam current. This should follow from the fact that the dose rate at the dose compatible regions of the ipsilateral lung should naturally be relatively higher, as scattering is still not pronounced. Therefore, the current required for full FLASH coverage there should be smaller than the minimum reported, since this is optimized for FLASH at every dose-compatible region. The decreasing improvements for higher dose thresholds were expected, since the higher it is, smaller fractions of the same volume can benefit from FLASH.

By being compatible with a wide window of beam currents, treatment plans can be made more or less sensitive to FLASH uncertainties *a posteriori* on

the dose rate threshold, depending on the setup of the delivery settings, without generating new plans. The higher the beam current, the higher the dose rate and so the less sensitive the solutions are. Optimization can be made further consistent by specifying lower dose and higher dose-rate thresholds.

5.2 Pencil-Beam Scanning Dose Rate

The fine-tuned optimizer successfully finds scanning patterns with FLASH gains consistently higher than the predefined ones, without any trade-off on dose, as illustrated in figure 7. Despite the stochastic nature of the optimizer, the difference on sparing between the best and worst-performing patterns is small, suggesting that running the optimizer just once guarantees high-quality solutions, speeding up optimization for clinical applications.

Snowflake-shaped patterns are obtained for the optimized patterns, as illustrated in figure 8. This follows from a clever sacrifice of dose rate at some regions to improve it at others. Priority is given to maximization of the dose rate wherever it is easier to achieve, which corresponds to regions with less pencil-beam overlap. The center of the pattern should be sacrificed, since there is more overlap due to more neighbor pencil beams, each with a high dose contribution. The outwards regions have less dose and can be fully irradiated quickly. This is enabled by the branches of the snowflakes.

The swirl patching for larger target volumes is a logical consequence of the higher number of pencil beams. If the snowflake shape was applied, branches would have to be considerably longer and so the time to reach opposite ends of one branch would not be FLASH compatible, translating into very low dose rates, similarly to the predefined patterns. Instead, it is more advantageous to adopt patches of branches, connected in a swirl formation, despite some dose rate sacrifice between extremes of different branches, as seen in figure 8.

Some optimized patterns have slightly less FLASH gains than predefined ones. This can be explained by the lower-resolution dose-deposition matrix used by the optimizer, which despite sufficient for dose, is not accurate. For evaluation, the full accurate dose-deposition matrix is used.

Optimization is performed outside the PTV, although it would be more logical to also include the GTV-PTV margin, as it is healthy tissue irradiated with high dose. This was not implemented to speed up optimization, since no coherent data was available for that margin besides the full dose-deposition matrix, which would make optimization infeasible. This is not an issue though, as high FLASH gains are obtained, because healthy tissue exists in front and behind the PTV, so the GTV-PTV margin is indirectly optimized. However, including it can be critical for larger tumour volumes, near the beam entrance and for relatively low beam currents.

For some patients, the optimized patterns for direction *B3* have low FLASH gains, despite being consistent between runs, which suggests an issue with the beam direction. These underperforming patters have empty regions in their interior, where there could be pencil beams. The inner voxels of those empty regions require the irradiation of most of the beams around, translating into low dose rates. Inverting the direction can increase FLASH compatibility, while preserving plan quality.

For the underperforming treatment plans, direction *B3* is associated with larger pathlengths from the beam entrance to the tumour and through all the patient. These pencil beams, chosen by the treatment planning software and not the pattern optimizer, should be preferred because of lateral scattering. For deep-seated tumours, lateral scattering should be already significant, therefore, to keep the high dose conformal and localized there, less but heavier pencil beams are required with empty spaces around them. This guarantees the dose is still high at the middle but minimal at the healthy region, which is unavoidable due to scattering. For tumours closer to the beam entrance, scattering is minimal, so it is easier to keep the dose conformal and localized by using more and lighter beams.

On FLASH sensitivity, figure 9 and table 2 show that the dose-rate threshold has great impact on the FLASH gains for the optimized patterns. In the regime where the pencil-beam delivery time dominates over the switching time, sensitivity to the dose-rate threshold can be guaranteed *a posteriori*

by adjusting the beam current accordingly. Since the relation between current and dose rate is linear, the optimized patterns for 40 nA and 40 Gy/s would also be the best-performing for 60 nA and 60 Gy/s. This is the case here, as the median minimum delivery time for 40 nA is 1.8 ms and the switching time is ~0.2 ms. Regarding the dose threshold, the results are consistent, due to the symmetric dose-rate distributions.

6 Conclusion

A significant increase of dose rate may be achieved with the proposed optimization strategies for stereotactic PT of lung lesions using TBs, through higher beam currents and optimized snowflake-shaped scanning patterns. The optimized treatment plans have a quality in terms of dose similar or even identical to conventional IMPT plans. Nonetheless, this quality is constrained by a set of *a priori* parameters: beam current, minimum delivery time, beam direction and FLASH model.

With current and near-future PT, combined with the proposed optimization approaches, a significant enhancement in dose rates is feasible. Based on the current knowledge of FLASH, this may be of significant clinical benefit for stereotactic treatment of selected patients with lung lesions. For the best balance between dose and dose rate, it is essential to determine the parameters of the treatment-delivery machine and the thresholds at which FLASH occurs. The effect of beam current fluctuations, the pencil-beam scanning and switching times should be evaluated. Beam-direction optimization is important, manually performed by minimizing distances and avoiding dense structures.

FLASH optimization could be improved, weighting tissue sparing and the trade-off to achieve it, to guarantee optimal FLASH-enhanced dose distributions. Snowflake-like patterns could be used as a starting point to reduce the optimization time. Integrated plan and pattern optimization on both metrics could yield better and more consistent results. Hybrid approaches could improve results and allow the application in other tumour sites, with multi-beam setups for which only a fraction of radiation dose is delivered under strict FLASH conditions.

References

[1] Vincent Favaudon et al. “Ultrahigh dose-rate (FLASH) irradiation increases the differential response between normal and tumor tissue in mice”. In: *Science Translational Medicine* 6.245 (2014), 245ra93.

[2] Manuela Buonanno, Veljko Grilj, and David J. Brenner. “Biological effects in normal cells exposed to FLASH dose rate protons”. In: *Radiotherapy and Oncology* 139 (2019). FLASH radiotherapy International Workshop, pp. 51 –55.

[3] Jean Bourhis et al. “Treatment of a first patient with FLASH-radiotherapy”. In: *Radiotherapy and Oncology* 139 (2019). FLASH radiotherapy International Workshop, pp. 18 –22.

[4] Steven van de Water et al. “Towards FLASH proton therapy: the impact of treatment planning and machine characteristics on achievable dose rates”. In: *Acta Oncologica* 58.10 (2019). PMID: 31241377, pp. 1463–1469.

[5] Daniel Tsang. “Investigating the influence of pencil beam scanning on the dose rate and the optimal pencil beam scanning patterns”. In: *Internship Report* (2021).

[6] Hao Gao et al. “Simultaneous dose and dose rate optimization (SDDRO) for FLASH proton therapy”. In: *Medical Physics* n/a.n/a (2020).

[7] Sebastiaan Breedveld et al. “iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans”. In: *Medical Physics* 39.2 (2012), pp. 951–963.

[8] J. R. Isbell and W. H. Marlow. “Attrition games”. In: *Naval Research Logistics Quarterly* 3.1-2 (1956), pp. 71–94.

[9] Hao Gao et al. “Minimum MU optimization (MMO): an inverse optimization approach for the PBS minimum MU constraint”. In: *Physics in Medicine & Biology* 64.12 (2019), p. 125022.

[10] P. Larrañaga et al. “Genetic Algorithms for the Travelling Salesman Problem: A Review of Representations and Operators”. In: *Artificial Intelligence Review* 13.2 (1999), pp. 129–170.