



**TÉCNICO**  
LISBOA

# **Novel Optimization Strategies for Clinical FLASH Proton Therapy**

Evaluation on stereotactic lung treatment plans

**Rodrigo José Santo**

Thesis to obtain the Master of Science Degree in

**Engineering Physics**

Supervisor(s): Dr. Steven Johannes Martinus Habraken  
Prof. Patrícia Carla Serrano Gonçalves

## **Examination Committee**

Chairperson: Prof. Maria Teresa Haderer de la Peña Stadler

Supervisor: Dr. Steven Johannes Martinus Habraken

Members of the Committee: Prof. Mischa Sebastiaan Hoogeman

Prof. Ruben Maurício da Silva Conceição

**December 2021**



I declare that this document is an original work of my own authorship and that it fulfills all the requirements of the Code of Conduct and Good Practices of the Universidade de Lisboa.



## Acknowledgements

Nine months and one thesis: few words may describe a dissertation but fail to tell its journey, with all the struggles and challenges, the enthusiasm and passion. I would have not reached this far and surely not this enthusiastically without all the support and help of many kindhearted people, to whom I would like to express my sincere appreciation.

First, I would like to thank my supervisor, Steven Habraken, for being a true mentor, with all the care and support on both the thesis project and the experience abroad, always kindly making sure I was fully passionate and motivated about the work and headed in the right direction. I am also grateful to the whole Radiotherapy department at Erasmus Medical Center for eagerly welcoming me, making me feel part of their family. A special thanks to the Master's students I worked alongside with on a daily basis. In particular, a warm thank you to Axel and Sofia, for the extra incentives and comfort, making every day truly unpredictable, exciting and fun.

I am deeply in debt to all my family for constant help and inspiration. A special thanks to my parents and sister for the never ending support and motivation, constantly encouraging me to follow my dreams, teaching me to always challenge myself and to never give up. Despite far away, I never stop feeling the warmth of everyone back home, instead always feeling so close, consistently spoiled with love. My gratitude extends to my grandmothers for always checking on me and wishing me luck, even coping with modern technology to ease some of the distance. Also to my uncles for never stopping to believe in me and for always being curious on my latest findings, truly interested and excited to learn all the details about my work.

I am also extremely grateful for all the sympathy and kindness of my friends, especially those that followed my progress up close, from the beginning until the very end. A big thank you to Daniel, David, Pedro and Tiago for their interest and enthusiasm on my work. Even though online, I had great fun and true joy on every moment together, genuinely reassuring me our friendship knows no borders and no distance. My sincere gratitude goes also to Hakyong for being part of this important journey, contributing with endless love and support, indifferent to all the time and space between us. I am thankful for the contagious happiness and positivity boosts that never failed to cheer me up during hard times.

I humbly extend my thanks to all kind people that contributed to this thesis, to everyone I crossed paths with on this journey.



## Collaboration Framework

This research project was fully developed by the author at the Erasmus University Medical Center (Erasmus MC), based in Rotterdam, The Netherlands, as a guest of the Radiotherapy Medical Physics and Technology research group of the Radiotherapy department, over a duration of 9 months (February 2021 to October 2021), with full supervision and collaboration of Dr. Steven Habraken, medical physicist and principal investigator at Erasmus MC. The project also counted on input from Prof. Dr. Mischa Hoogeman, head of Medical Physics at Erasmus MC and at the Holland Proton Therapy Center (HollandPTC), and Prof. Dr. Sebastiaan Breedveld, assistant professor at Erasmus MC.

This thesis was developed to obtain the Master of Science Degree in Engineering Physics at Instituto Superior Técnico (IST), Universidade de Lisboa, based in Lisbon, Portugal. The project was also officially supervised by Prof. Dr. Patrícia Gonçalves, associate professor at the Physics department and at the Engineering and Nuclear Sciences department of IST, the degree-granting institution. Within the collaboration framework, the author was awarded a scholarship under the Erasmus+ program.



The Erasmus MC is the largest and most versatile of all eight university medical centres in the Netherlands and one of the most authoritative scientific university medical centers in Europe [a]. It is committed to a healthy population and excellence in health care through research and education.

The Department of Radiotherapy within the Erasmus MC Cancer Institute is one of the largest clinics in the Netherlands with more than 5,500 patients treated each year [b]. The Radiotherapy Medical Physics and Technology research group has been at the forefront in developing automated treatment planning and high-precision adaptive radiotherapy approaches. This includes online and real-time adaptive radiotherapy, robotic stereotactic body radiotherapy and proton therapy. Other focus areas are biology image-guided radiotherapy, and technology for interventional radiotherapy such as brachytherapy and hyperthermia. An important goal of the research program is to implement completed research and to evaluate the benefit.





## Resumo

O efeito FLASH tem despertado interesse devido à sua maior preservação de tecido saudável que radioterapia convencional. Este efeito tem sido observado para irradiação de elevadas doses (>8 Gy) a taxas ultra-elevadas (>40 Gy/s). Combinado com as vantagens dosimétricas da terapia com prótons e sua adequação para tratar tumores profundos, FLASH-PT tem o potencial de reduzir a toxicidade e melhorar tratamentos. Contudo, os atuais software de planejamento de tratamento não estão aptos para otimizar o efeito.

Neste projeto são desenvolvidas novas estratégias para otimização de FLASH-PT, que reduzam efeitos colaterais através de FLASH dentro das restrições clínicas. FLASH-PT está atualmente limitada a feixes acelerados por ciclotrões, usando pencil-beam scanning. Para maximizar as taxas de dose, a energia máxima do ciclotrão é considerada, correspondendo a transmission beams, aplicados no tratamento estereotáxico de lesões pulmonares. A taxa de dose é otimizada por otimização da intensidade, da corrente e do padrão de scanning dos feixes, com base na taxa de dose ponderada à dose (DADR) e taxa de dose PBS (PBSDR), usando linearização iterativa, relaxamento convexo iterativo e Algoritmos Genéticos.

Aumentos significativos da taxa de dose são alcançado com as estratégias propostas, através de correntes de feixe elevadas e padrões em floco-de-neve, que, com base no conhecimento atual de FLASH, podem ser de grande benefício clínico. Antes da aplicação clínica, a sensibilidade aos parâmetros da máquina de tratamento deve ser avaliada. Os resultados podem ser melhorados, ponderando os ganhos e danos nos tecidos, otimizando a direção do feixe e irradiando parcialmente com FLASH.

**Palavras-chave:** FLASH, radioterapia com prótons, pencil-beam scanning, otimização da taxa de dose, otimização do padrão de scanning



## 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 ( $>8$  Gy) 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 does not account for dose rates, not optimizing FLASH.

In this project, the aim was at developing novel strategies for optimization and delivery of FLASH-PT within clinical constraints on dose, that reduce side effects through FLASH. This is currently limited to cyclotron-accelerated beams and using pencil-beam scanning. To achieve the overall highest dose rates, protons with the maximum commissioned cyclotron energy (244 MeV) are used, corresponding to shoot-through transmission beams (TB), applied to stereotactic treatment of lung lesions. Dose rate is optimized via beam intensity, beam current and scanning-pattern optimization, based on two metrics - 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 higher 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



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Radiotherapy	1
1.2	Proton Therapy	2
1.3	FLASH Radiotherapy	3
1.4	FLASH Proton Therapy	3
1.4.1	FLASH Evaluation	5
1.4.2	Dose-Averaged Dose Rate	5
1.4.3	Pencil-Beam Scanning Dose Rate	6
1.5	FLASH Proton Therapy Optimization	7
1.5.1	Indirect Optimization	7
1.5.2	Direct Optimization	8
1.6	Problem Definition	9
1.7	Research Goal	10
1.8	Thesis Outline	10
<b>2</b>	<b>Background</b>	<b>11</b>
2.1	2D dose optimization	11
2.2	Erasmus-iCycle	11
2.2.1	Erasmus-iCycle Workflow	12
2.2.2	Erasmus-iCycle Pencil-Beam Weights	13
<b>3</b>	<b>Methods</b>	<b>15</b>
3.1	Dose-Averaged Dose Rate (DADR)	15
3.1.1	Simple Linear Fractional Programming	16
3.1.2	Maximization of the minimum DADR	17
3.1.3	Maximization of the mean DADR	17
3.1.4	Beam-Current Optimization	19
3.1.5	FLASH Thresholds	22
3.2	Pencil-Beam Scanning Dose Rate (PBSDR)	22
3.2.1	Genetic Algorithms	23
3.2.2	Island Model Genetic Algorithm	24
3.2.3	Scanning-Pattern Optimization	25
3.2.4	Fine-tune routine	30
3.3	Optimization Setup	31
3.3.1	Analytical 2D models	31
3.3.2	Clinical Scenario	32
3.4	Evaluation	33
3.4.1	Plan comparison	33
3.4.2	FLASH Sensitivity	34
<b>4</b>	<b>Results</b>	<b>35</b>
4.1	Dose-Averaged Dose Rate (DADR)	35
4.1.1	Analytical 2D Models	35
4.1.2	Clinical Scenario	38
4.2	Pencil-Beam Scanning Dose Rate (PBSDR)	44
4.2.1	Analytical 2D Models	44

4.2.2	Clinical Scenario . . . . .	49
<b>5</b>	<b>Discussion</b>	<b>60</b>
5.1	Dose-Averaged Dose Rate Optimization . . . . .	60
5.1.1	Analytical 2D Models . . . . .	60
5.1.1.1	Tumour field . . . . .	60
5.1.1.2	Square field . . . . .	61
5.1.2	Clinical Scenario . . . . .	62
5.1.2.1	Comparison with IMPT treatment plans . . . . .	62
5.1.2.2	Beam Current Analysis . . . . .	63
5.1.2.3	Direction Analysis . . . . .	64
5.1.2.4	Component analysis . . . . .	64
5.1.2.5	FLASH Sensitivity Evaluation . . . . .	65
5.2	Pencil-Beam Scanning Dose Rate Optimization . . . . .	66
5.2.1	Analytical 2D Models . . . . .	66
5.2.1.1	Version Comparison . . . . .	66
5.2.1.2	Pattern Comparison . . . . .	66
5.2.1.3	Target Comparison . . . . .	67
5.2.2	Clinical Scenario . . . . .	68
5.2.2.1	Version Comparison . . . . .	68
5.2.2.2	Convergence Analysis . . . . .	68
5.2.2.3	Suboptimal Evaluation . . . . .	68
5.2.2.4	Pattern Comparison . . . . .	69
5.2.2.5	Scalability Evaluation . . . . .	70
5.2.2.6	Direction Analysis . . . . .	71
5.2.2.7	FLASH Sensitivity Evaluation . . . . .	71
<b>6</b>	<b>Conclusion</b>	<b>73</b>
<b>7</b>	<b>Appendix</b>	<b>78</b>
7.1	Proof of DADR optimization routines . . . . .	78
7.1.1	Simple Linear Fractional Programming . . . . .	78
7.1.2	Maximization of the minimum DADR . . . . .	78
7.1.3	Maximization of the mean DADR . . . . .	80



# List of Figures

1.4.1	Dose deposition as a function of depth for beams of different particles: electrons, photons and protons. . . . .	4
1.4.2	Tissue revisits between pencil beams. . . . .	5
1.4.3	Scanning-pattern contribution for the PBSDR on a region of interest. . . . .	6
1.5.1	Predefined scanning patterns. . . . .	8
3.2.1	Example of two different scanning patterns and their corresponding <i>path representation</i> for a $3 \times 2$ regular grid. . . . .	25
3.2.2	Partially-Mapped Crossover (PMX) operator. . . . .	26
3.2.3	Partially-Mapped Ordered Crossover (PMOX) operator. . . . .	27
3.2.4	Swap mutation operator. . . . .	28
3.2.5	Insert mutation operator. . . . .	28
3.2.6	Simple inversion mutation operator. . . . .	28
3.2.7	Shift mutation operator. . . . .	28
3.2.8	Migration routes for a ring-like organization of 4 islands. . . . .	29
3.2.9	IMGA fine-tune routine. . . . .	30
4.1.1	FLASH coverage for fully-optimized solutions on analytical 2D models. . . . .	35
4.1.2	Relative solution quality for fully-optimized solutions on analytical 2D models. . . . .	36
4.1.3	Global solution quality for fully-optimized solutions on analytical 2D models. . . . .	36
4.1.4	Spatial distribution of the normalized pencil-beam weights for DADR-only-optimized solutions on analytical 2D square fields. . . . .	37
4.1.5	Dose distribution of DADR-only-optimized solutions for analytical 2D square fields. . . . .	37
4.1.6	Dose-rate distribution of DADR-only-optimized solutions for analytical 2D square fields. . . . .	38
4.1.7	FLASH coverage for fully-optimized and conventional IMPT treatment plans. . . . .	38
4.1.8	FLASH coverage for fully-optimized treatment plans with the the minimum-required and the maximum-compatible beam currents. . . . .	39
4.1.9	Relative mean dose at the ipsilateral lung for fully-optimized treatment plans. . . . .	39
4.1.10	Beam-current windows for fully-optimized treatment plans. . . . .	40
4.1.11	DADR distribution at different slices within a patient with a small lung lesion. . . . .	40
4.1.12	Relative mean dose at the ipsilateral lung for fully-optimized and current-optimized treatment plans. . . . .	41
4.1.13	Effective FLASH Enhancement Ratio at the ipsilateral lung for fully-optimized and current-optimized treatment plans. . . . .	41
4.1.14	Sensitivity of the improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the maximum-compatible beam current. . . . .	42
4.1.15	Sensitivity of the effective FLASH Enhancement Ratio at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the maximum-compatible beam current. . . . .	43
4.1.16	Sensitivity of the improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the minimum-required beam current. . . . .	43
4.1.17	Sensitivity of the effective FLASH Enhancement Ratio at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the minimum-required beam current. . . . .	44



4.2.1	Relative FLASH coverage for suboptimal solutions of different version of the scanning-pattern optimizer on analytical 2D models. . . . .	45
4.2.2	Relative FLASH coverage for different versions of the scanning-pattern optimizer on analytical 2D models as a function of the running time. . . . .	45
4.2.3	Relative FLASH coverage for different number of pencil beams on analytical 2D models as a function of the running time. . . . .	46
4.2.4	FLASH coverage for optimized and predefined scanning patterns on analytical 2D models. . . . .	46
4.2.5	Best-performing optimized patterns for analytical 2D models. . . . .	47
4.2.6	PBSDR distributions of the best-performing optimized patterns for analytical 2D models. . . . .	47
4.2.7	Best-performing optimized patterns for analytical 2D models optimized including GTV. . . . .	48
4.2.8	PBSDR distributions of the best-performing optimized patterns for analytical 2D models optimized including the GTV. . . . .	48
4.2.9	FLASH coverage for the patterns optimized excluding and including the GTV on analytical 2D models. . . . .	49
4.2.10	Relative FLASH coverage for suboptimal solutions of different version of the scanning-pattern optimizer on clinical data. . . . .	49
4.2.11	Relative FLASH coverage for different version of the scanning-pattern optimizer on clinical data. . . . .	50
4.2.12	Convergence of the scanning pattern optimizer for two tumour sizes. . . . .	50
4.2.13	Optimized scanning patterns for clinical data of small lung lesions. . . . .	51
4.2.14	Improvement on the mean dose at the ipsilateral lung for different scanning patterns. . . . .	52
4.2.15	Effective FLASH Enhancement Ratio at the ipsilateral lung for different scanning patterns. . . . .	52
4.2.16	Best-performing optimized patterns for different patients with small lung lesions. . . . .	53
4.2.17	Optimized scanning patterns on clinical data for various lung tumour volume. . . . .	53
4.2.18	Improvement on the mean dose at the ipsilateral lung for various lung tumour volumes. . . . .	54
4.2.19	Effective FLASH Enhancement Ratio at the ipsilateral lung for various lung tumour volumes. . . . .	55
4.2.20	Best-performing optimized patterns for various lung tumour volumes. . . . .	55
4.2.21	PBSDR distributions of the best-performing optimized patterns for various lung tumour volumes. . . . .	56
4.2.22	Best-performing optimized patterns for the original and the reversed direction <i>B3</i> . . . . .	56
4.2.23	Improvement on the mean dose at the ipsilateral lung for the original and reversed direction <i>B3</i> . . . . .	57
4.2.24	Effective FLASH Enhancement Ratio at the ipsilateral lung for the original and reversed direction <i>B3</i> . . . . .	57
4.2.25	Sensitivity of the improvement on the mean dose at the ipsilateral lung for the best-performing optimized scanning patterns on different FLASH thresholds. . . . .	58
4.2.26	Sensitivity of the effective FLASH Enhancement Ratio at the ipsilateral lung for the best-performing optimized scanning patterns on different FLASH thresholds. . . . .	59

**List of Tables**

- 4.1.1 Median improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the maximum-compatible beam current. 42
- 4.1.2 Median improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the minimum-required beam current. 44
- 4.2.1 Median improvements on the mean dose at the ipsilateral lung for the best-performing optimized scanning patterns on different FLASH thresholds with the maximum compatible beam current. 58
- 7.1.1 Treatment Planning wish-list. 82

## Glossary

**BP:** Bragg Peak

**DADR:** Dose-Averaged Dose Rate

**Erasmus MC:** Erasmus Medical Center

**FER:** FLASH Enhancement Ratio

**GA:** Genetic Algorithm

**GTV:** Gross Tumour Volume

**HollandPTC:** Holland Proton Therapy Center

**IMGA:** Island Model Genetic Algorithm

**IMPT:** Intensity Modulated Proton Therapy

**OAR:** Organ At Risk

**PBS:** Pencil-beam scanning

**PBSDR:** Pencil-Beam Scanning Dose Rate

**PT:** Proton Therapy

**PTV:** Planning Tumour Volume

**RT:** Radiotherapy

**TB:** Transmission Beam



# 1 Introduction

Over decades, the number of patients diagnosed with cancer has been increasing and is expected to continue rising in the future [1], in part due to higher life expectancy and lower incidence and mortality of cardiovascular diseases. This has motivated developments on both cancer prevention and treatment procedures, to decrease mortality and to improve future complications. Thanks to that, data shows that over the last 25 years cancer mortality has been decreasing but death rates are still high [2].

Cancer treatment is split in three main modalities, with different specificities, which are often used in combination to ensure the best results. Surgery is recommended for local tumours, systematic therapies (e.g. chemo and immunotherapies) for systemic diseases and radiotherapy for local-regional control. Radiotherapy has been consistently one of the most important treatment modalities, as around 50% of all cancer patients are indicated for it, at least once during the course of their disease, irrespective of the world region [3]. Treatments with radiotherapy may be associated with side effects though, arising from unavoidable damage to healthy tissue, which research on the field has been trying to reduce.

Sparing of healthy tissue in radiotherapy treatments has been successfully improved through high-precision irradiation techniques, such as proton therapy, and through optimal targeting, such as live treatment adaptation. However, a new way of enabling reduced healthy-tissue toxicity has been identified by biological studies: the FLASH effect. This has been observed for cells when irradiated with a high dose, for a very short time, at ultra-high dose rate. Combined with the dosimetric advantages of precision irradiation techniques, the potential for substantially improved clinical outcomes is great. FLASH-compatible proton therapy (FLASH-PT) has been the focus of several recent studies, which mostly cover the feasibility and optimization of treatment plans with the current clinical technology.

## 1.1 Radiotherapy

Radiotherapy is effective due to a differential response to ionising radiation by the DNA of cells. Tumour cells are often more radiosensitive than healthy cells and are less capable of repairing damage. This translates in reduced growing capabilities and death for tumours, whereas healthy tissue is partially spared. There are different ways of irradiating tissue but the focus here is on external beams, where one or more conformal beams, most commonly of photons, are targeted at the tumour from outside of the patient's body. These beams are modulated in intensity to deliver a treatment-specific prescribed dose. External beam radiotherapy is commonly applied either as a primary treatment modality or as a locoregional treatment, after surgery to kill remaining tumour cells left in the target area. Recent studies also suggest the alongside use of immunotherapy, to amplify the radiation damage to the DNA of cancer cells, while reducing side effects to healthy tissue.

The radiosensitivity differences between healthy and tumour cells is referred to as therapeutic bandwidth and is modeled by the TCP (Tumour Control Probability) and the NTCP (Normal Tissue Complication Probability). Treatment plans should increase this therapeutic bandwidth, by maximizing the TCP (tumour damage) while minimizing the NTCP (healthy-tissue damage). In addition, tumour cells are also less capable of self-repairing their DNA compared with healthy cells, which can be exploited through fractionation of the treatment into several sessions. By giving more time for healthy tissue to recover, the therapeutic bandwidth is increased, lessening possible future complications.

One of the main challenges in radiotherapy is to reduce the exposure of healthy tissue to radiation, which is unavoidable because tumours infiltrate normal tissue. Additionally, there are always uncertainties in the tumour location and shape, requiring safety margins that might contain healthy tissue. This unavoidable

damage to the healthy cells can lead to immediate or long term effects and complications, such as dry mouth, dysphagia, pneumonitis and fibrosis, with potentially severe impact on quality of life.

Research on radiotherapy aims at reducing side effects without sacrificing tumour cell kill. This can be improved through high-precision irradiation techniques, such as proton therapy, and through optimal targeting, using online imaging techniques combined with motion correction (such as online patient setup through tracking) and with live treatment plan adaptation. However, increasing the conformity of dose to the target has limits, motivating the search for alternative methods to reduce healthy tissue toxicity. Recent biological studies point out to a fundamentally new different way to achieved this: the FLASH effect.

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. Treatment planning software assures that, bridging the medical treatment requirements with the treatment delivery by the machine. These optimal plans are generated based on information about the beam properties, such as directions, positions, energies and intensities, alongside tissue characteristics, namely response to radiation. Inputs include CT-scan delineations of important structures in the patient and a set of constraints, objectives and goals for the dose restrictions and requirements in the patient.

## **1.2 Proton Therapy**

Proton therapy is a radiotherapy modality that uses high-energy protons, ranging from 70 to 250 MeV, to deliver dose. It allows for precise irradiation as protons have a peaked energy deposition curve, with most of their energy being deposited at a well defined depth (Bragg peak), which depends on their initial energy. This translates into substantially lower entry and exit doses and, consequently, in a better sparing of healthy tissue compared to conventional photon irradiation. The dosimetric advantages of proton therapy and the maturation of the technology has led to a growth in the number of proton therapy centers throughout all world, with more under construction or at a planning stage [4].

Despite the recent expanded interest in proton therapy, clinical implementations have been used over the past seventy years, with the first reported patient being treated back in the 1950s [5]. These first treatments used scatterers to shape and aim the radiation field to the target volume, through a cascade of interactions between the proton beam and the specific material of the modulation devices. Due to its simplicity, passive scattering is still used today for the majority of patients treated with proton therapy. However, it has been replaced at new facilities by a more versatile delivery system: the pencil-beam scanning (PBS). PBS works by adjusting the location of the Bragg peak for a proton beam (pencil beam) through energy degradation, incidence-angle adjustment and magnet steering, and modulating the intensity of each beam to accurately irradiate any tumour volume, irrespective of its shape and position.

The development of proton therapy was only possible due to the invention of cyclotrons in 1929, a type of particle accelerator that produces continuous charged particle beams, such as of protons, of high intensity at a fixed energy. This is achieved by accelerating charged particles using rapidly varying electric fields and confining them to a spiral trajectory within a flat cylindrical vacuum chamber, through a static magnetic field.

Proton treatments with more complex types of particle accelerators are also possible and carried out in some facilities. In synchrotron-based systems, particles are accelerated around a fixed closed-loop path through an increasing magnetic field, enabling the extraction of any energy directly from the accelerator. This is a major advantage since the use of degraders for energy selection in cyclotrons is associated with scattering, which increases beam spreading and consequently decreases dose conformity to the

target, specially for lower energies. However, synchrotrons produce beams in bunches, on at a time and with tenths to a few seconds apart. Cyclotrons are by far the most common accelerators used for clinical proton therapy due to their simplicity and compactness. Synchrotrons offer higher versatility for other radiotherapy modalities though, as they can be used with different types of ions, most notoriously carbon, and are able to produce micron-sized ultra-high dose rate X-rays.

Intensity Modulated Proton Therapy (IMPT) plans are the golden standard for proton therapy, corresponding to optimal dose distributions. These are personalized treatment plans, for which the beam parameters are modulated to maximize damage to tumour and sparing of healthy tissue, improving clinical outcomes.

### **1.3 FLASH Radiotherapy**

The radiosensitivity of cells has been suggested for decades to depend on the dose rate, both from observations that treatments at low dose rates are relatively ineffective and that extremely high dose rates are more deadly [6][7]. However, studies have only recently concluded that the effect of dose rate on cell survival is differential. Healthy tissue irradiated with a high dose at a ultra-high dose rate, for a very short time (a flash of dose), has been observed to be less damaged compared to conventional irradiation, without compromising tumour cell kill [8][9]. This differential toxicity is know as the FLASH effect and has the potential to substantially improve treatment quality. To date, this effect has only been demonstrated on culture cells and small mammals but recently the first clinical trial was conducted, with the first human patient being treated with a FLASH-compatible radiotherapy plan [10].

The FLASH effect is still mostly not understood but two explanations are often suggested, supported by preliminary evidences. It is suggested that fast irradiation induces hypoxic conditions (low concentration of oxygen), which consequently reduce DNA damage [11][12]. It explains why tumour cells are not affected by the FLASH effect as many tumours are already hypoxic due to their abnormal growth. By generating hypoxic conditions in healthy tissue, the differential effect of radiotherapy can be increased. It is also suggested that fast irradiation does less damage to circulating immune system cells in the blood [13].

Current knowledge on how to trigger the FLASH effect is limited, with a large set of irradiation parameters reported to potentially contribute to the effect. Consequently, it is still not possible to fully and accurately exploit the reduced toxicity enabled by FLASH. Nevertheless, a set of 3 thresholds have been consistently reported to trigger the effect in several pre-clinical studies: doses higher than 10 Gy, irradiated for times lower than 100 ms and at dose rates higher than 100 Gy/s. Although different experiments have reported different minimum thresholds, the previous set of parameters is compatible with all observations, therefore guaranteeing FLASH beyond doubt. More conservative approaches consider instead thresholds of 8 Gy, 40 Gy/s and 200 ms.


### **1.4 FLASH Proton Therapy**

Most studies on FLASH radiotherapy have been conducted using electrons for dose delivery [8][14], because electron-beam setups are more widely available and FLASH-compatible dose rates can be achieved with conventional linear accelerators, albeit with modifications, allowing instantaneous pulses with dose rates up to  $10^6$  Gy/s. However, because electrons have a small range within the patient, FLASH electron therapy is only suitable for superficial treatments. Even if electrons are given enough energy to reach deeper in the patients, because they are very light particles, scattering is pronounced, resulting into significantly broader and less well-defined beams.

Recent experiments have demonstrated the FLASH effect for protons too [9][15]. This is an important milestone for FLASH radiotherapy because protons allow for better healthy tissue sparing and higher treatment versatility, enabling conformal irradiation of deep-seated tumours. Furthermore, the required high dose rates are readily available with present-day clinical proton system. General-purpose clinical FLASH-RT is thus currently only possible with protons [16].

Despite the enormous potential of FLASH-PT, meeting the requirements to trigger the effect is currently challenging. Present-day clinical cyclotrons deliver continuous proton beams at a cyclotron frequency of  $10^3$  Hz and are readily capable of instantaneous dose rate of about 800-1000 Gy/s, which is compatible with the FLASH effect. However, FLASH-compatible electron therapy is delivered in very short pulses of ultra-high dose rate, at a frequency of  $10^2$  Hz and simultaneously to the whole target. This is not possible with current cyclotrons because scattering significantly reduces dose rates, which become incompatible with FLASH, being limited to very small targets on biological experiments. Synchrotrons are not suitable either because even though they allow for energy preselection, their pulsed time structure is in the order of tenths to a few second, too slow for FLASH [17]. Therefore, FLASH-PT is currently limited to cyclotron-accelerated pencil-beam scanning, meaning that the tumour volume is not irradiated completely all at once, but instead locally through a succession of pencil beams. In addition, it is also critical to account for the contributions of the treatment delivery machine, such as the dead time between pencil beams, for switching energy layers, for steering the pencil beam and for the adjusting the beam current required for each layer/beam. Nonetheless, the FLASH effect has also recently been observed for this treatment delivery scheme [18].

To guarantee the highest possible dose rates, the fixed extraction energy of the cyclotron is used, which for commercially-available cyclotrons corresponds to 250 MeV. Energy degraders are not considered due to the slow switching time between energy layers, currently incompatible with FLASH. These single high-energy proton beams shoot through the patient and have their Bragg peak outside, being designated by transmission beams (TB). TBs are relatively non-precise compared to Bragg peak beams but they mitigate range uncertainties and freeze anatomical motion. They are analogous to photon beams in the sense that they transverse all the patient, but TBs have a flat dose-depth profile and a sharper lateral dose fall-off.



**Figure 1.4.1:** Dose deposition as a function of depth for beams of different particles: electrons, photons and protons. The highly-energetic protons shoot through and have a Bragg peak outside the patient, only depositing dose with the tail.


The first studies on FLASH-PT have focused on tumour types that are suggested to benefit greatly with the effect. Small lung lesions have been getting the most attention, being considered a benchmark for treatment comparison. With lungs being comprised of mostly air and motion from breathing being inevitable during irradiation, conventional Bragg peak treatments are associated with significant range and motion uncertainties. Transmission beams can overcome these issues, since there is no associated range uncertainty and they are less sensitive to motion. Treatment quality with TBs is comparable to



that of using photons, which would conventionally be prescribed for lung lesions [19]. Additionally, lung lesions are commonly treated in stereotactic hypofractionation schedules, meaning fewer fractions, each with high doses. This is important to trigger the FLASH effect since it enables higher compatibility with the FLASH dose threshold, translating into a higher fraction of healthy tissue that may benefit from the reduced toxicity.

### 1.4.1 FLASH Evaluation

Evaluating the FLASH effect for proton therapy using PBS is not trivial since dose is delivered locally in a sequence of small fields, contrasting with the continuous whole volume irradiation used in experimental *in vivo* FLASH studies. PBS translates into tissue revisits and overlap between beams, which ultimately compromises the dose rate. This represents a new degree of freedom that conventional proton therapy does not take into account but is essential to trigger the FLASH effect.



**Figure 1.4.2:** Tissue revisits between pencil beams.

Current knowledge on FLASH is unable to give insight on how tissue revisits contribute to the effect. Consequently, new metrics were proposed based on different aspects of the dose delivery scheme for PBS, allowing to consistently and universally compare different treatment plans in clinical trials. Two of the most used FLASH dose-rate metrics are the Dose-Averaged Dose Rate (DADR) and the Pencil-Beam Scanning Dose Rate (PBSDR). Present-day treatment planning software does not take this new degree of freedom into account, requiring new optimization strategies to tune and find the best balance between the FLASH and the clinical dose requirements.

### 1.4.2 Dose-Averaged Dose Rate

Proposed by van de Water *et al.* [20], the Dose-Averaged Dose Rate (DADR) accounts for the local contribution of each pencil beam, weighting their instantaneous dose rate by their dose contribution for every voxel. For a given voxel  $i$ , this corresponds to the following expression:

$$DADR_i = \sum_{j=1}^n \frac{(D_{ij}w_j) \times dr_{ij}}{D_i} \quad (1.4.1)$$

with  $D_{ij}$  the dose deposition matrix and  $dr_{ij}$  the dose rate at voxel  $i$  by pencil beam  $j$  and  $w_j$  the weight of the pencil beam  $j$ .

This metric naturally incorporates the fact that small doses have little effect on the dose rate, by weighting the dose rate of each pencil beam by its dose contribution. This makes it unnecessary to include any arbitrary threshold, while still taking into account the fraction of dose that is given at FLASH-compatible doses at every voxel.


By offering a more mathematical approach to the problem and by not depending on the dose delivery time structure, direct optimization of the DADR is complex but still challenging. However, the simplicity

of the DADR gives it an incomplete representation of the problem, as it tries to summarize the effect of an intermittent delivery in one number. Furthermore, it does not account for the sequential delivery of the pencil beams, instead considering they are irradiated all at once.

### 1.4.3 Pencil-Beam Scanning Dose Rate

Pencil beams are delivered one at a time, following a patterns, without continuously irradiating the same voxels at the same dose rate. This translates into dead times that although expected to have some role on FLASH, as they contribute to the total time to accumulate dose, their effects are not yet understood.

Folkers *et al.* [21] tried to account for this contribution on a new metric, the Pencil-Beam Scanning Dose Rate (PBSDR). For each voxel  $i$ , this metric considers an effective irradiation time  $T_i$  that starts when the accumulated dose exceeds an arbitrary dose  $d_T$ , recommended in literature to be fixed at 1 cGy, and stops when the accumulated dose comes within  $d_T$  of the total dose  $D_i$ . This way, the effective accumulated dose is  $D_i - 2d_T$  and the PBSDR is just the ratio of this dose and the effective time. The PBSDR metric is better illustrated in figure 1.4.3, for a particular example of a scanning pattern and the corresponding dose time structure.



**Figure 1.4.3:** Scanning-pattern contribution for the PBSDR on a region of interest (ROI). The slope of the blue line is the Pencil-Beam Scanning Dose Rate  $dr^{PBS}$  at the ROI.

The effective time for a given voxel includes not only the irradiation times of contributing pencil beams, but also the dead time between them, to which contribute both scanning and irradiation times of non-contributing beams, as illustrated for the ROI by the dotted lines on figure 1.4.3. For that pattern, it is seen that the effective time at the ROI must take into account the irradiation time of the rightmost red and orange pencil beams, even though they do not contribute with dose there. A pattern that would irradiate the three leftmost beams first (the only contributing with dose there), would give a higher PBSDR at the ROI, reducing both the irradiation time of non-contributing beams and the scanning time.

Although the local time structure of dose delivery is suggested to be important, it does not mean that the PBSDR is necessarily better and that the other metrics are irrelevant. If the onset of the tissue's state to cause FLASH is fast but the relaxation to the normal state is slow, then the time window during which the effect can be considered triggered is wide. The dead time between pencil beams is not relevant, meaning the PBSDR is not a good approximation of the dose rate. However, biological studies have shown that under the oxygen depletion hypothesis, both the onset of hypoxic conditions and reoxygenation is fast. Therefore, the effect is optimized if the highest dose is delivered in the shortest period of time, at the highest dose rate and in short pulses, so that irradiation is over before reoxygenation. This supports the need to take into account the sequential irradiation of the pencil beams, reinforcing that the PBSDR metric is relevant to evaluate FLASH.

The PBSDR can be interpreted as a more general metric, with the DADR corresponding to a particular extreme case, where all the pencil beams are delivered at the same there, so there is no pattern to account for. Consequently, if a voxel is compatible with FLASH based on the PBSDR, it will also be compatible regarding the DADR, but the opposite is not valid. Nonetheless, the DADR distributions might offer some insight on the optimal PBSDR distributions, since conversely, if a voxel is not compatible with FLASH on the DADR, it will never be compatible on the PBSDR.

## 1.5 FLASH Proton Therapy Optimization

Current clinical treatment planning software is not able to generate proton therapy treatment plans that are fully compatible with FLASH, as several studies have concluded. However, an appropriate choice of planning strategies and beam delivery scenarios has been reported to improve dose rates, although with trade-off on the dose distribution in the patient [20][21][22][23]. These results are promising as they show that FLASH-compatible plans are possible to some extent, but they illustrate the need for adequate optimization strategies, not only capable of fully exploiting FLASH but also of keeping the plans clinically safe for the patients.

Because no FLASH model exists yet, optimization of the effect can only be performed by satisfying the FLASH thresholds in the most regions possible. Research has mostly been focused on dose-rate distributions, evaluating both compatibility with FLASH and the trade-off on the dose distributions, which arises from the fact that dose and dose rate objectives compete with one another. On the one hand, higher beam currents can greatly increase dose rates, but because of the minimum pencil-beam delivery time is a fixed machine parameter, plans become limited to a proportionally higher minimum pencil-beam weight. This should translate into higher minimum doses in the patient, potentially rendering plans clinically unsafe. On the other hand, even though a smaller number of pencil beams should reduce overlap and tissue revisits, each one has to contribute more. Consequently, these plans have higher doses in the patient, with less conformal distributions.

### 1.5.1 Indirect Optimization


The simpler approach to the problem is to indirectly optimize FLASH, by exploiting some of the dose rate dependencies. As noted by Groen [24], because the number and separation of pencil beams are strongly tied to the delivery time, they can play an important role on the dose rate. The higher the number of pencil beams on the same target, the closer together they are and consequently there is more overlap and tissue revisits. The dose rate is then expected to be lower, with this effect being stronger for larger field sizes. A higher number of pencil beams also translates into a more complex scanning pattern, which will take longer to deliver. This motivated an attempt to optimize the dose rate by reducing the pencil-beam overlap. When optimizing and evaluating stereotactic lung proton therapy plans, despite the overlap minimization routine reducing the number of pencil beams and increasing their separation, the treatment plans did not generally benefit from it. A reduction on the delivery time and increase of dose rate was observed for some beams, but because this was achieved by reducing their dose contribution, other beams had to compensate for it. Therefore, these beams have higher delivery times and lower dose rates at the regions they compensate. In addition, a decrease in the treatment plan quality was also observed, due to a lower conformity of high dose areas to the target volume.

Van de Water *et al.* [20] took a different approach, evaluating different planning strategies for different delivery scenarios, in search of the best settings for proton therapy treatment plans fully compatible with FLASH. Focus was also set on reducing the total number of pencil beams, with the aim of reducing irradiation times and consequently increasing FLASH compatibility. Based on the four head and neck

tumour patients evaluated, it was concluded that mean dose rates compatible with FLASH throughout all irradiated volume are only achievable when using transmission beams all around the patient (arc shoot-through beams), with higher doses per fraction (6 Gy as opposed to conventional 2 Gy), considering hypothetical variation of beam current for every pencil beam, with no safety limits (up to  $10^{12}$  protons/s as opposed to conventional  $10^9$  protons/s) and minimizing pencil-beam overlap. However, these plans came at the cost of substantially higher integral dose, which is still deemed unsafe based on the conventional clinical requirements. Additionally, some settings are still not achievable with current technology.

To overcome the current technological limitations, van Marlen *et al.* [23] later proposed splitting a plan into different subplans, each associated with a different beam current. Pencil beams are selected for each subplan based on their weight, allowing to deliver heavier beams at higher beam currents without sacrificing the global minimum pencil-beam weight, which is delivered in a different subplan at a lower beam intensity. On the ten evaluated head and neck tumour patients, plans generated with this strategy were seen to have an overall higher dose rate and reduced pencil-beam delivery time. However, because different subplans have different pencil beams in the same beam direction, the delivery of the global plan requires more gantry rotations, consequently resulting into a higher global scanning time. Although not evaluated, this increase of the total irradiation time might be enough to make plan splitting not compatible with FLASH, requiring each subplan to be evaluated individually on the effect. Some of these subplans might not be compatible with the FLASH dose threshold, depending on the contributing pencil beams.

On the scanning pattern, Tsang [25] focused on reducing the dead time between pencil beams, more specifically on the irradiation time of non-contributing beams. Different scanning patterns were evaluated on analytical 2D models with square homogeneous fields and one stereotactic lung proton therapy treatment plan. It was observed that minimization of the total path length lead to substantially better dose rate distributions in comparison with predefined patterns, such as the *Z*, *Snake* and *Circle*, illustrated in figure 1.5.1. Despite the results showing that the improvements greatly depend on the configured start and end positions, they illustrate that there is room for optimization and that the scanning pattern can be critical for FLASH.



**Figure 1.5.1:** Predefined scanning patterns, namely the *Z*, *Snake* and *Circle* patterns. Patterns start with a green square and end with a red square.

Despite encouraging, showing that there is room for dose-rate optimization without substantial degradation of the treatment plan quality, these indirect approaches do not guaranteed maximization of the FLASH coverage. In addition, it is impossible to quantify or control any dose trade-off *a priori* with these strategies, which only reinforces the need for direct optimization strategies.

### 1.5.2 Direct Optimization

Direct optimization of the FLASH effect is expected to give better results and to allow more control over the trade-offs between dose and dose rate. This is a challenging problem, due to the complex formulation

of the dose-rate metrics, which are associated with nonlinear and nonconvex functions.

The first reported direct optimizer was proposed by Gao *et al.* [26], based on the DADR formulation of the dose rate. It uses an Iterative Convex Relaxation method, splitting the complex problem into a series of simpler subproblems. For three lung cancers patients, the algorithm generated transmission-beam plans with substantially improved FLASH coverage on dose rate at the target volume and at other critical organs, with dose coverage at the tumour comparable to that of conventional IMPT plans. These substantially improved dose rates were achieved through beam current maximization, performed taking into account dose requirements but prioritizing dose rate goals in the healthy volume around the tumour. This region can potentially benefit the most from FLASH, because dose on healthy tissue is the highest there. With this approach, the generated plans are guaranteed to have dose rates compatible with FLASH and to satisfy the dose objectives as far as possible. This priority on the dose rate explains the observed slightly worse sparing of OARs compared to conventional IMPT plans.

A variation of this optimizer was later introduced by Lin *et al.* [27] to take into account both transmission beams and conventional Bragg peak beams. By sampling Bragg peaks so that they are placed inside the tumour, the target dose conformality and the dose to healthy tissue is seen to improve. Transmission beams primarily cover the tumour boundary, necessary to achieve both FLASH-compatible dose rates at the OARs close to the tumour target and plan robustness. With this approach, FLASH coverage is comparable to that of only using transmission beams but the target dose conformality is substantially improved and the dose at OARs is reduced. However, although not analysed in that study, potential gains from the FLASH effect might be lost because FLASH compatible Bragg peak hardware is still under development. Currently, only the dose delivered by transmission beams can be associated with FLASH. Even though they are placed inside the tumour target, these Bragg peak beams have a non-negligible contribution to the healthy tissue at the vicinity of the tumour boundary. Therefore, it is possible for the contribution of the transmission beams alone not to be enough to meet the FLASH dose threshold, reducing FLASH compatibility.

## 1.6 Problem Definition

A treatment plan is only considered compatible with FLASH if all or most of the regions receiving doses above the FLASH dose threshold are irradiated at a dose rate above the FLASH dose rate threshold, for an irradiation time below the FLASH irradiation time threshold. Although different treatment planning strategies and optimizers have been already evaluated on compatibility with the FLASH effect, direct FLASH optimization has not yet been performed.

The state of the art simultaneous dose and dose-rate optimization methods are reportedly dose-rate biased, first ensuring FLASH-compatible dose rates and later optimizing the dose. Consequently, dose will always be sacrificed whenever and wherever necessary. From a clinical perspective, this lack of control over the dose trade-off makes it impossible to generally guarantee both the effective treatment of the tumour and the safety of the patient. In addition, by only focusing on a FLASH dose-rate threshold, the current approaches proposed in literature might be overestimating the resulting FLASH coverage. They might unnecessarily sacrifice healthy tissue to improve the dose rate at regions that might not be compatible with FLASH on dose.

Despite being pointed out in literature as an important contributing factor, the irradiation time is often overlooked. FLASH evaluation has been focused on the DADR metric, which does not take into account the treatment delivery time. The PBSDR metric is often skipped, even though previous studies have shown that the dose-rate distributions greatly depend on the pencil-beam scanning pattern. In addition,

no direct scanning-pattern optimization strategies have been proposed for this metric yet, to the best of found knowledge, neither on FLASH nor dose rate. This is due to the non-convex and non-linear nature of the PBSDR metric and the inherent challenges of pattern optimization.

## 1.7 Research Goal

In this project, novel optimization strategies for clinical FLASH proton therapy are developed. The research goal is to reduce the side effects in clinical radiotherapy through the FLASH effect, developing strategies for safe optimization and delivery of FLASH-compatible proton therapy plans, based on the present knowledge of FLASH and taking into account the conventional clinical requirement on dose and the present-day technological challenges.

Two direct FLASH optimization strategies are proposed, based on the two different dose-rate metrics introduced: DADR and PBSDR, using lung tumours as the first foreseen clinical application. FLASH compatibility is here defined by irradiation of doses above 8 Gy, delivered at dose rates above 40 Gy/s and for irradiation times below 200 ms. Therefore, clinical optimization on FLASH aims to maximize the volume that is irradiated within these conditions, while satisfying the clinical requirements on dose for effective and safe treatments considered in conventional proton therapy.

First, the optimization strategies are evaluated on analytical 2D models. This enables to develop insight into the intrinsic pathways that maximize FLASH coverage, such as the pencil-beam weight distribution, the dose trade-offs and pattern shapes, while understanding the technical challenges and shortcomings of the proposed strategies. After this validation, the optimization strategies are implemented in a treatment planning software. Hereafter, transmission-beam treatment plans with and without FLASH optimization are generated and compared on dose, the corresponding dose-rate metric and FLASH coverage.

From the clinical point of view, the proposed optimization strategies are implemented in the the Erasmus Medical Center in-house developed treatment planning software - Erasmus-iCycle [28]. The proton therapy setup considered is that of the HollandPTC, which consists of a cyclotron particle accelerator, followed by a beam transportation system leading the proton beam to the gantry on the treatment room, where the patient is irradiated. A constant beam energy of 244 MeV is used throughout all this study, the maximum commissioned in the in-house beam model. Since the focus here is on proof of principle, robustness planning is not applied, neither a thorough clinical validation of the treatment plans.

## 1.8 Thesis Outline

This document is split into 5 additional sections. Section 2 starts with providing the theoretical background on dose optimization methods, both for analytical 2D models and the Erasmus MC in-house developed treatment planning software - Erasmus-iCycle. Section 3 describes the methods implemented to generate the results, alongside the data used and the evaluation approaches. Results are then reported in section 4. A discussion of the results and methods used, explanations and implications can be found in section 5. Finally, the conclusions are drawn in section 6, alongside an outlook.

## 2 Background

Some background knowledge is provided on how Erasmus-iCycle, the treatment planning software used, works. This is relevant for this project because the proposed optimizations strategies are built around it. Additionally, the dose optimization approach used for the analytical 2D models is also described.

### 2.1 2D dose optimization

The dose contribution of a unit weight of a given pencil beam at every region of the space is given by a dose-deposition matrix  $\mathbf{D}$ , which has the following form:

$$\mathbf{D} = \begin{bmatrix} D_{1,1} & D_{1,2} & \dots & D_{1,n} \\ D_{2,1} & D_{2,2} & \dots & D_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ D_{m,1} & D_{m,2} & \dots & D_{m,n} \end{bmatrix} \quad (2.1.1)$$

where  $D_{i,j}$  represents the dose deposited at voxel  $i$  by unit weight of the pencil beam  $j$ ,  $n$  is the total number of pencil beams and  $m$  is the total number of voxels. Considering the pencil-beam weights are represented by  $w$ , it comes that the dose  $d$  in the patient is then:

$$\mathbf{D} \cdot w = d \quad (2.1.2)$$

The pencil-beam weights need to be optimized so that the best dose distribution is obtained, one that balances complete irradiation of the tumour with sparing of the surrounding healthy tissue. For this, the quadratic different between the target dose  $d_T$  and the deposited dose  $d$  is minimized, by varying the pencil-beam weights  $w$  using a quadratic optimizer:

$$\min_w (\mathbf{D} \cdot w - d_T)^2 \quad (2.1.3)$$

The dose distribution obtained by the previous optimization problem is considered the best possible for the corresponding 2D model, so it is used to set dose constraints when optimizing the dose rate.

### 2.2 Erasmus-iCycle

Erasmus-iCycle was developed in-house at Erasmus MC and stands out from other treatment planning software due to its prioritized multi-criteria optimization [28]. This allows to better specify clinical preferences to input on the optimization routine, thereby leading to treatment plans of exceptional quality while maintaining high generalization over different patients.

The prioritized nature of Erasmus-iCycle is well-suited for clinical FLASH optimization because the dose and dose-rate objectives compete with one another. This way, there is a better control of the trade-off between dose and dose rate, enabling a finer management over which objectives can be compromised to achieve better dose rate distributions and which ones are critical. Therefore, the clinical requirements on dose can be ensured, guaranteeing the safe and efficient treatment of the patient, with the added bonus of FLASH.

### 2.2.1 Erasmus-iCycle Workflow

Erasmus-iCycle takes as input CT-scan delineations of important structures in the patient and a list of all optimization parameters, objectives and constraints, alongside their clinical preference, desired goal and corresponding target structure. This list is termed wish-list. The constraints have to be necessarily met for a plan to be feasible [29], while the objectives are just goals that the optimizer strives to achieve, according to the desired preferences. These wish-lists are not patient specific, as the clinical requirements are general for patient groups with the same tumour site. By reusing the same wish-list, it is possible to save time compared to conventional planning approaches, as those require a time-consuming adaptation for each patient. The output is an accurate optimized dose distribution to the patient and a set of pencil beams, with their location and their optimized weight.

The workflow of Erasmus-iCycle is therefore divided into three main phases: initialization, optimization and finalization. At the initialization phase, all the input data is processed, whereas at the finalization phase, the accurate dose distribution in the patient is calculated and outputted, alongside the optimized pencil beams. The optimization phase, being the most complex and where the modifications for FLASH optimization were made, deserves a more in-depth explanation.

#### Optimization Phase

One of the initial challenges of any proton therapy treatment planning software is to select an appropriate set of pencil beams. There are different strategies for pencil-beam selection but considering that transmission beams are to be used throughout all this project, a regular grid of such beams is preferred. This corresponds a 3D grid of pencil beams, generated from a given lateral spacing and an energy spacing, where all the beams that do not go through the tumour, according to the beam directions specified, are removed. In this particular case, because quasi mono-energetic transmission beams of 244 MeV are used, the resulting regular grid is only bidimensional.

From the eligible pencil beams, the dose-deposition matrix is calculated. This is similar to the matrix described for the analytical 2D models, with the difference that the type of tissue each pencil beam goes through is taken into account, by extracting information from the CT scans of the patient. In addition, it is infeasible to use the accurate dose-deposition matrix for optimization, because it is too large and consequently the running time would get impractical within the clinical workflow. Therefore, a lower resolution version is used, comprised of just the necessary voxels for an adequate dose calculation.

The weights of the pencil beams are optimized in a prioritized multi-criteria optimization routine, using the *2-phase  $\epsilon$ -constraint (2pec)* method [30]. This method splits optimization into two stages, based on the fact that when possible, it is more beneficial to first optimize an objective only up until a certain threshold - the goal - and then move to the following priority objective, instead of directly optimizing the higher priority objective to its fullest. This enables more flexibility in general, by leaving more room for lower priority objectives to meet the associated goals.

At the first phase of the *2pec* method, each objective  $f_i(w)$  is optimized, respecting all hard constraints  $\mathbf{g} \leq \mathbf{0}$ , in order of their priority  $i$ , only up until the goal  $b_i$  is satisfied, when possible. Otherwise, the closest value to the goal is taken. In the first iteration of this first phase, the objective having the highest priority is optimized:

$$\begin{aligned} \min_w \quad & f_1(w) \\ \text{s.t.} \quad & \mathbf{g}(w) \leq \mathbf{0} \end{aligned} \tag{2.2.1}$$



After each objective has been optimized, it is transformed into a hard constraint with a given bound  $\epsilon_i$ , depending on the result  $w^*$  and according to the following rule:

$$\epsilon_i = \begin{cases} b_i & , f_i(w^*) \times \delta < b_i \\ f_i(w^*) \times \delta & , f_i(w^*) \times \delta \geq b_i \end{cases} \quad (2.2.2)$$

where  $\delta$  is a relaxation constant, introduced to leave more room for the optimization of subsequent objectives, usually set to 1.03 (representing a relaxation of 3%). The next optimization step goes through objective  $f_2$ , now keeping  $f_1$  constrained:

$$\begin{aligned} \min_w \quad & f_2(w) \\ \text{s.t.} \quad & \mathbf{g}(w) \leq \mathbf{0} \\ & f_1(w) \leq \epsilon_1 \end{aligned} \quad (2.2.3)$$

In the second phase of the method, the objectives that had their goals met are optimized to their fullest, while keeping all other constrained. Therefore, for each  $f_k$  constrained to  $b_k$ , in order of priority:

$$\begin{aligned} \min_w \quad & f_i(w) \\ \text{s.t.} \quad & \mathbf{g}(w) \leq \mathbf{0} \\ & f_k(w) \leq \epsilon_k \quad , \quad k \in \{1, \dots, N\} \setminus i \end{aligned} \quad (2.2.4)$$

When this routine finishes, the pencil-beams weight have been fully optimized. However, the weights obtained for some of the pencil beams might not be achievable at the machine, due to the minimum pencil-beam weight constraint. All the pencil beams with weights bellow 1.05 times that minimum are removed and a feasibility check is performed based on the previously obtained results:

$$\begin{aligned} \min_w \quad & 0 \\ \text{s.t.} \quad & \mathbf{g}(w) \leq \mathbf{0} \\ & f_k(w) \leq \epsilon_k \quad , \quad k \in \{1, \dots, N\} \end{aligned} \quad (2.2.5)$$

This procedure is repeated till either all the pencil beams have weights higher than the minimum or the plans is infeasible, in which case it returns the previous best plan, despite having underweighted beams. In most cases, however, the minimum pencil-beam weight is low enough for this routine to successfully satisfy this constraint. By considering an extra 5% above the minimum weight, the pencil-beam removal routine is sped up, although it might lead to infeasible plans more easily.

### 2.2.2 Erasmus-iCycle Pencil-Beam Weights

The optimized pencil-beam weights represent the configuration that better satisfies the clinical requirements. These weights  $w$  can have arbitrary units, depending on how the dose-deposition matrix  $\mathbf{D}$  is constructed and on its corresponding units, usually differing for different beam models. Since in conventional planning all objectives are based on dose, the units of the pencil-beam weights are irrelevant for the optimizer and is only relevant for plan delivery, when interpreted by the proton treatment machine. However, understanding the physical units of the pencil-beam weights becomes essential when working with dose rates, as they depend on the beam current.

The proton beam current  $I$  is a measure of the amount of charge delivered by the proton beam per unit of time, being reported for proton therapy system in nA, which is equivalent to C/s. Knowing the charge

of the proton ( $e = 1.60217662 \times 10^{-19}$  C), the beam current can be expressed in units of protons per unit time, the most common unit being gigaprotons per second Gp/s. The relation between pencil-beam weights and protons is then required, so that the beam current can be used to calculate irradiation times and dose rates.

In this study, the HollandPTC beam data was used, for which the pencil-beam weights are normalized so that for each proton energy, for a unit pencil-beam weight, the dose deposited in a water-equivalent material at the Bragg peak  $D_{max}$  is 1 Gy. Experiments at this beam setup reported a transmission efficiency  $\eta$  of 43% for the beam transportation system, going from the cyclotron to the research bunker, at energies of 250 MeV. Therefore, a cyclotron current of  $I = 10$  nA corresponds to a beam current in the bunker, where the patient is to be irradiated, of  $I_{bunker} = 4.3$  nA which is equivalent to 26.9 Gp/s. At this beam current, a dose-rate of  $dr_5 = 20.8$  Gy/s at 5 cm in a water-equivalent material was measured. For the 244 MeV transmission beams used in this project, it is known from previously measured dose-depth curves that the dose at 5 cm in that same water-equivalent material is 28.5% of  $D_{max}$ , meaning that the dose-rate at the Bragg peak  $dr_{peak}$ , located outside the patient, is:

$$dr_{peak} = \frac{dr_5}{0.285} = 73.0 \text{ Gy/s} \quad (2.2.6)$$

Consequently, it comes that at  $D_{max}$  we have the following relation:

$$\frac{dr_{peak}}{I_{bunker}} = 0.368 \text{ Gp/Gy} \quad (2.2.7)$$

Considering that for the HollandPTC beam data a pencil-beam weight has  $D_{max} = 1$  Gy at the Bragg peak, it directly follows that a unit pencil-beam weight in Erasmus-iCycle corresponds to 0.368 Gp for these 244 MeV proton transmission beams.

### 3 Methods

In this section, the optimization strategies for the two dose-rate metrics are introduced and described. The section is split into two parts, the first covering the Dose-Averaged Dose Rate (DADR) and the latter the Pencil-Beam Scanning Dose Rate (PBSDR). Two additional subsections are included in the end, one describing the data and settings for optimization and the other covering the metrics to evaluate and compare the optimized solutions.

#### 3.1 Dose-Averaged Dose Rate (DADR)

Optimization of the DADR was implemented at the optimization phase of Erasmus-iCycle, so that pencil beam weights are directly adjusted to maximize DADR, while keeping compatibility with the prioritized nature of this treatment planning software. DADR optimization was introduced as two new types of objectives that can be specified at the wish-list, making it simple and easy for wish-lists previously used on conventional treatments to be adapted for the FLASH effect. These correspond to maximization of the minimum DADR and maximization of the mean DADR, which can be prioritized to any structure and volume in the patient, alongside conventional dose constraints and objectives.

The main challenge of directly optimizing the DADR comes from the fact that it corresponds to a ratio, depending on the optimization variable  $w$  on both numerator and denominator. Considering  $D_{ij}$  the dose-disposition matrix and  $dr_{ij}$  the instantaneous dose rate at voxel  $i$  by pencil beam  $j$  and  $w_j$  the weight of the pencil beam  $j$ , the DADR at voxel  $i$  is given by:

$$DADR_i = \frac{\sum_{j=1}^n D_{ij}w_j \times dr_{ij}}{\sum_{k=1}^n D_{ik}w_k} = \frac{N(w)}{D(w)} \quad (3.1.1)$$

The problem can be simplified though, by noting that the instantaneous dose rate  $dr_{ij}$  at voxel  $i$  by pencil beam  $j$  does not depend on the pencil-beam weight  $w_j$ . It only depends on the beam current  $I_j$  and the dose-deposition matrix  $D_{ij}$ . Defining the delivery time of each pencil beam as  $T_j = I_jw_j$ , it follows:

$$dr_{ij} = \frac{d_{ij}}{T_j} = \frac{D_{ij}w_j}{I_jw_j} = I_jD_{ij} \quad (3.1.2)$$

From the previous observation, it comes that the DADR is just a ratio of two linear expression - a linear ratio - which helps to simplify the problem further. Linear Fractional Programming is a deeply studied topic, as these types of functions appear extensively on different fields, such as stock cutting, resource allocation, ship scheduling and cargo-loading problems [31]. Therefore, there is extensive literature on specific methods and approaches to optimize these functions.

In addition, when optimizing these DADR objectives, the hard constraints that must be satisfied are all linear: all the dose constraints and objectives that can be specified in the wish-lists are linear (minimum, maximum and mean doses), while the DADR constraints can be transformed into linear constraints. This means that when optimizing each objective, the search space will always be polyhedral. To better comply with the clinical requirements, the DADR objectives are only applied on healthy tissue and after the dose objectives, meaning that the search space is bounded. Additionally, the numerator and denominator on the DADR expression are both always positive in the search space, meaning that the DADR is a continuous positive function. These properties make it possible to use simpler and more straightforward algorithms.

### 3.1.1 Simple Linear Fractional Programming

Maximization of the minimum DADR and maximization of the mean DADR is based on the same general idea proposed by Isbell and Marlow [32] in 1956 for a bounded polyhedral search space: to transform the linear fractional program into a series of linear programs.

To better illustrate how this iterative linearization works, a step by step example is first given on a simple linear fractional program, having objective function  $f(x)$ , with  $x \in X$  where  $X$  is a bounded polyhedron in  $\mathbb{R}^n$ :

$$f(x) = \frac{ax + b}{cx + d} = \frac{N(x)}{D(x)} \quad (3.1.3)$$

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

$$\frac{N(x)}{D(x)} \leq q^* \Leftrightarrow N(x) - q^*D(x) \leq 0 \quad (3.1.4)$$

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

$$\mathcal{F}(q^*) = \max_{x \in X} F(x, q^*) = N(x^*) - q^*D(x^*) = (ax^* + b) - q^*(cx^* + d) = 0 \quad (3.1.5)$$

The problem is that the value  $q^*$  is not known but it can be iteratively approximated and validated, as the solution to the problem is found when the maximum of the linear function  $F(x, q)$  is 0. Based on these observations, the routine summarized in algorithm 3.1.1 solves the maximization of the linear ratio  $f(x)$ , with  $\arg \max$  the argument of the maxima and parameter  $\delta$  the early-stopping criteria. The convergence of this routine to the optimum is proven in the Appendix section.

---

#### Algorithm 3.1.1: Optimization of the linear ratio function $f(x)$

---

**Result:**  $f(x^*) = q^*$

**Input:**  $x_0 \in X$

$k \leftarrow 1$

**do**

$q_k \leftarrow f(x_{k-1}) = N(x_{k-1})/D(x_{k-1})$
$x_k \leftarrow \arg \max F(x, q_k) \Leftrightarrow \arg \max \{N(x) - q_k D(x) \mid x \in X\}$ using any linear optimizer
$k \leftarrow k + 1$

**while**  $F(x_k, q_k) \geq \delta$

$x^* \leftarrow x_k$

$q^* \leftarrow q_k$

---

Despite the challenges on maximizing a linear ratio, it is easy and straightforward to set it as a constraint. Considering the objective function  $f(x)$  and the minimum constraint value  $q^-$ , it follows:

$$f(x) = \frac{N(x)}{D(x)} \geq q^- \Leftrightarrow N(x) - q^-D(x) \geq 0 \quad , \quad (3.1.6)$$

which is simply a linear constraint. For implementation in Erasmus-iCycle, this is particularly useful not only because it is easy to set linear fractional constraints, but also because after each DADR objectives

is optimized, it is transformed into a linear constraint, which will keep the search space as a bounded polyhedron.

### 3.1.2 Maximization of the minimum DADR

Maximization of the minimum DADR belongs to a subset of problems inside the fractional programming field, as it corresponds to a *maxmin* objective function. This optimization problem is solved with a routine similar to that for the simple linear ratio [33], as it may be view as a generalization. The only difference is that instead of a single linear ratio, there are several of them, one for each voxel of the structure targeted by the objective. Therefore, the optimization problem is now expressed as:

$$\max_{x \in X} \left[ \min_{1 \leq i \leq M} f_i(x) \right] = \max_{x \in X} \left[ \min_{1 \leq i \leq M} \frac{a_i x + b_i}{c_i x + d_i} \right] = \max_{x \in X} \left[ \min_{1 \leq i \leq M} \frac{N_i(x)}{D_i(x)} \right] = q^* \quad (3.1.7)$$

For this new objective function, the linearized functions will have a similar form as before, with the rest of the optimization routine staying identical, as reported by algorithm 3.1.2. In the Appendix section it is proven that this routine converges to the optimum of the *maxmin* linear fractional problem. Parameter  $\delta$  sets an early stopping criteria and  $\arg \max$  is the argument of the maxima.

---

**Algorithm 3.1.2:** Optimization of a *maxmin* linear fractional program with  $M$  functions  $f_i(x)$

---

**Result:**  $f(x^*) = q^*$

**Input:**  $x_0 \in X$

$k \leftarrow 1$

**do**

$q_k \leftarrow \min_{1 \leq i \leq M} f_i(x_{k-1}) = \min_{1 \leq i \leq M} [N_i(x_{k-1})/D_i(x_{k-1})]$   
 $x_k \leftarrow \arg \max F(x, q_k) \Leftrightarrow \arg \max \{ \min_{1 \leq i \leq M} [N_i(x) - q_k D_i(x)] \mid x_k \in X \}$  using any  
 linear optimizer  
 $k \leftarrow k + 1$

**while**  $F(x_k, q_k) \geq \delta$

$x^* \leftarrow x_k$

$q^* \leftarrow q_k$

---

Transforming a *maxmin* linear fractional objective into a constraint is also straightforward. It just requires applying the same transformation used for the simple linear fractional programming problem but to every ratio. Considering the objective function  $\min_{1 \leq i \leq M} f_i(x)$  and the minimum constraint value  $q^-$ , it follows:

$$\min_{1 \leq i \leq M} f_i(x) \geq q^- \quad \Rightarrow \quad f_i(x) = \frac{N_i(x)}{D_i(x)} \geq q^- \Leftrightarrow N_i(x) - q^- D_i(x) \geq 0 \quad \text{for all } i = 1, \dots, M \quad (3.1.8)$$

### 3.1.3 Maximization of the mean DADR

Maximization of the mean DADR is a substantially more difficult problem because the objective function corresponds to a sum of linear ratios. There have been extensive studies on optimization of these dunctions, with several methods proposed in literature [34][35][36][37], however, these are all quite limited to just a few ratios and even unfeasible within Erasmus-iCycle. Thankfully, it is possible to reformulate this objective in a different way, by noting that the goal of maximizing the mean DADR is simply that of maximizing the DADR everywhere. Optimization of this single objective can be interpreted as a multiobjective problem, with each subobjective having the same weight and corresponding to the DADR at every voxel

of the target structure. This greatly simplifies the problem as it is possible to solve multiobjective linear fractional programming (MOLFP) problems with the same straightforward approach used for the simple linear fractional programming problem [38].

The maximization of the mean of a set of linear ratios  $f_i(x)$ , with  $x \in X$  and  $X$  a bounded polyhedron, can then be reformulated as a multiobjective problem the following way, with solution  $q^* \in \mathbb{R}^M$ :

$$\max_{x \in X} \sum_{i=1}^M f_i(x) = \max_{x \in X} \sum_{i=1}^M \frac{a_i x + b_i}{c_i x + d_i} = \max_{x \in X} \sum_{i=1}^M \frac{N_i(x)}{D_i(x)} \Rightarrow \max_{x \in X} [f_1(x), \dots, f_M(x)] = q^* \quad (3.1.9)$$

Each of these subobjective functions  $f_i(x)$  is to be linearized similarly to the approach for the *maxmin* linear fractional programs. The key difference is that the new objective function is a weighted sum of the linearized objectives, with a  $q_k^i$  per subobjective per iteration. Furthermore, each subobjective is constrained to be at least equal to the value of the previous iteration, resulting into a different search space  $X_k$  at each iteration. This translates into the routine summarized by algorithm 3.1.3, which is proven to converge to the optimal solution of the multiobjective linear fractional program in the Appendix section. Parameter  $\delta$  corresponds to the early stopping criteria and  $\arg \max$  to the arguments of the maxima.

---

**Algorithm 3.1.3:** Maximization of a multiobjective linear fractional program with  $M$  functions  $f_i(x)$

---

**Result:**  $f(x^*) = q^*$

**Input:**  $x_0 \in X$ ,  $w \in \mathbb{R}^M$  with  $\sum_{i=1}^M w_i = 1$

$w^* \leftarrow 1 / \min_{1 \leq i \leq M} w_i$

$k \leftarrow 1$

**do**

$q_k \leftarrow [f_1(x_{k-1}), \dots, f_M(x_{k-1})] = [N_1(x_{k-1})/D_1(x_{k-1}), \dots, N_M(x_{k-1})/D_M(x_{k-1})]$

$X_k \leftarrow X \cap \{x \in \mathbb{R}^N \mid N_i(x) - q_k^i D_i(x) \geq 0, i = 1, \dots, M\}$

$x_k \leftarrow \arg \max F(x, q_k) \Leftrightarrow \arg \max \{\sum_{i=1}^M w_i [N_i(x) - q_k^i D_i(x)] \mid x \in X_k\}$  using any linear optimizer

$k \leftarrow k + 1$

**while**  $F(x_k, q_k) w^* \geq \delta$

$x^* \leftarrow x_k$

$q^* \leftarrow q_k$

---

Although the routine described in algorithm 3.1.3 is the one proposed in literature, the actual implementation is slightly different to prevent possible divergence errors by the linear optimizer when solving the linearized problems. Possible issues might arise from constraining each subobjective to be at least equal to the value of the previous iteration when a lot of subobjectives are to be considered. To overcome this, the constraints are relaxed by a small enough constant  $\gamma$ , which is updated every iteration  $k$  to be even smaller, so that the routine still converges to an optimal solution:

$$\gamma_{k+1} = (1 - 0.01 \times 5^{-k}) \quad (3.1.10)$$

It should be noted that maximization of the mean DADR is not totally equivalent to the multiobjective problem, since the latter does not take into account the overall effects of increasing a given objective through decreasing another one. Therefore, this multiobjective interpretation does not maximize the mean DADR. However, the mean DADR will still increase from the initial solution because the DADR at every voxel of the target structure is guaranteed to increase. As it is interpreted here, the maximization

of the mean DADR is still very important and should be performed after maximization of the minimum DADR at the same target structure. This is because the *maxmin* objectives only target the minimum value, leaving the remaining voxels with a suboptimal dose rates.

Transforming the maximization of the mean DADR from an objective to a constraint is neither simple nor straightforward, because it corresponds to a sum of linear ratios. Even though the mean DADR is not maximized, it is increased and so it would be more logical to constrain the objective as a mean. This can be simplified by considering that optimization of the mean DADR is only performed after the dose and the *maxmin* DADR objectives on the same target structure. Therefore, the dose there should not change further, meaning the denominator at the DADR for every voxel, which corresponds to the dose at that voxel, can be approximated by the dose for the starting solution  $w_0$ . For a minimum mean DADR constraint value  $q^-$ , it follows:

$$\frac{1}{M} \sum_{i=1}^M DADR_i(w) = \frac{1}{M} \sum_{i=1}^M \left[ \frac{\sum_{j=1}^n D_{ij} w_j \times dr_{ij}}{\sum_{k=1}^n D_{ik} w_k} \right] = \frac{1}{M} \sum_{i=1}^M \frac{N_i(w)}{D_i(w)} \approx \sum_{i=1}^M \frac{N_i(w)}{D_i(w_0)} \geq Mq^- \quad (3.1.11)$$

Through this approximation, the mean DADR constraint is transformed from a sum of linear ratios into a sum of linear functions, which itself is also linear, greatly simplifying optimization of other objectives.

### 3.1.4 Beam-Current Optimization

The previous optimization routines focus on maximizing the DADR by readjusting the pencil-beam weights but dose rates can also be increased by increasing the beam current. However, the beam current cannot be blindly increased due to treatment delivery machine specific limitations, that translate into dose trade-offs. This trade-off arises from a fixed minimum pencil-beam delivery time  $t_{min}$ , independent of the beam current. The higher the beam current  $I$ , the higher the minimum possible pencil-beam weight  $w_{min}$  and consequently the higher the minimum dose delivered to the patient  $d_{min}$ :

$$w_{min} = I \times t_{min} \quad \Rightarrow \quad d_{min} = \mathbf{D} \cdot w_{min} \quad (3.1.12)$$

This shows that the beam current should be also optimized alongside the dose and the DADR constraints, so that optimal DADR distributions could be obtained within the clinical dose requirements for a safe and effective treatment. Considering the previous relation, it comes that maximizing the beam current is equivalent to maximizing the minimum pencil-beam weight. This is challenging because it requires both increase of the weights and selection of the most suitable pencil beams. In other words, the pencil-beam weights have to either be maximized or set to 0, which makes the problem substantially more difficult.

In conventional proton therapy treatment planning the minimum pencil-beam weight also needs to be constrained due to other intrinsic machine limitations, such as the noise level on the monitor chamber and the stability of the beam current. These limitations translate into a fixed minimum beam weight, below which the delivery of pencil beams cannot be assured without uncertainties [39]. In literature, this is known as the minimum Monitor Unit (MU) constraint and it has been the focus of several studies.

The approach used here is the same proposed by Gao *et al.* on both conventional proton treatment plans [40] and on simultaneous dose and dose-rate optimization [26]: iterative convex relaxation. The idea is to optimize and constrain the minimum pencil-beam weights by alternating between optimizing the objectives and selecting the most appropriate beams from the set initially sampled. Minimum beam

weight optimization should only be performed for the DADR objectives because the beam current is not relevant for the dose objectives. This routine follows the procedure reported by algorithm 3.1.4.

---

**Algorithm 3.1.4:** Optimization with a minimum pencil-beam weight  $w_{min}$  constraint

---

**Result:**  $f(w^*)$

**Input:** Objective  $\min f(w)$ , search space  $W$ , restart factor  $\delta$ , initial pencil-beam weights  $w_0$  and minimum pencil-beam weight  $w_{min}$

$w' \leftarrow w_0$

**while**  $w' < w_{min}$  **do**

$w \leftarrow w'$  with all the pencil beams with weights bellow  $w_{min}$  removed

$w \leftarrow w \times \delta$

$w' \leftarrow \arg \min \{f(w) \mid w \in W\}$

**if Fails then**

$w \leftarrow w_0$  with all the pencil beams with weights bellow  $w_{min}$  removed

$w \leftarrow w \times \delta$

$w' \leftarrow \arg \min \{f(w) \mid w \in W \wedge w \geq w_{min}\}$  using any optimizer

**if Fails then**  $w' \leftarrow w_0$

        Stop

**end**

**end**

$w^* \leftarrow w_0$

**for every pencil beam  $j$  in  $w_0$  do**

**if pencil beam  $j$  removed from  $w'$  then**  $w_j^* \leftarrow 0$

**else**  $w_j^* \leftarrow w'_j$

**end**

---

Optimization with a minimum pencil-beam weight constraint is split into two stages: pencil-beam removal and weight maximization. First, the whole initial set of sampled beams is taken into account, optimized and removed iteratively. Pencil-beam removal is more likely not to be able to satisfy the minimum beam weight constraint  $w_{min}$ , so weight maximization follows. This step should not fail because the routine is applied to every objective function, meaning that the constraint is always ensured. Therefore, it should be always possible to optimize the objective with a hard constraint on the pencil-beam weights, defined by the minimum weight of the initial solution  $w_0$ . However, weight maximization is expected to leave less room for optimization of the objective. In case both optimization steps fail, the initial solution  $w_0$  is returned, which does satisfy the minimum pencil-beam weight but it is not optimal for the current objective.

This routine is applied to both constrain and optimize the minimum beam weight because if some pencil beams are removed, their contribution has to be compensated by the other beams. The weight of the pencil beams should then increase beyond their previous value, which was already higher or equal to the minimum beam weight. However, when optimizing the minimum weight, there is no objective function and so the optimizer just performs a feasibility check ( $f(x) = 0$ ). The routine is performed several times, with the minimum weight being updated after each successful step, for a given number of iterations, until it fails, converges or it is high enough to increase the beam current to a value that satisfies the DADR objective. To further help the optimizer find a solution with higher minimum weight, after removing the unsuitable pencil beams, the weights are multiplied by a restart factor  $\delta$  higher than 1, in this case set to



1.05. Optimization of a DADR objective is summarized by algorithm 3.1.5.

---

**Algorithm 3.1.5:** Optimization of a DADR objective with a minimum pencil-beam weight  $w_{min}$  constraint

---

**Result:**  $f(w)$ ,  $w$ ,  $w_{min}$  and  $I$

**Input:** DADR objective and goal  $f(w) \geq q$ , number of iterations  $N$ , initial beam current  $I_0$ , minimum pencil-beam delivery time  $t_{min}$ , initial pencil-beam weights  $w_0$  and minimum pencil-beam weights  $w_{min}^0$

$w'_{min} \leftarrow w_{min}^0$

$w' \leftarrow w_0$

**for**  $N$  iterations **do**

$w' \leftarrow$  run minimum pencil-beam weight  $w'_{min}$  constraint routine with objective 0 and initial solution  $w'$

**if** *Fails*  $\vee \min w' = w'_{min}$  **then**

        | Stop

**end**

$w'_{min} \leftarrow \min w'$

$I_{req} \leftarrow$  minimum beam current  $I$  required to satisfy DADR goal  $q$  with  $w'$

$w_{req} \leftarrow I_{req} \times t_{min}$

**if**  $w'_{min} \geq w_{req}$  **then**

        |  $I \leftarrow I_{req}$

        |  $w_{min} \leftarrow w_{req}$

        | Stop

**end**

**end**

$w \leftarrow \arg \max \{f(w) \mid w \in W \wedge w \geq w_{min}\}$  using iterative linearization

**if** goal  $f(w) \geq q$  was reached **then**

    |  $I \leftarrow$  minimum beam current  $I$  required to satisfy DADR goal  $q$  with  $w$

    |  $w_{min} \leftarrow I \times t_{min}$

**else**

    |  $I \leftarrow I_0$

    |  $w_{min} \leftarrow w_{min}^0$

**end**

---

After successfully optimizing the minimum weight and consequently the beam current, the DADR objectives are optimized separately with a hard constraint on the minimum pencil-beam weight. This is because the optimization strategies used for the DADR objectives are iterative, meaning that combining them with the iterative algorithm 3.1.4 to constrain the minimum weight would be too time consuming, perhaps even unfeasible in clinic.

Even if the beam current is already high enough to satisfy DADR goals, optimization is still performed because it may allow to afterwards relax the beam current and consequently the minimum pencil-beam weight, leaving more room for other objectives. If the dose rate goal is not met, the minimum weight and the beam current are reverted back to their initial values before optimization. This is because the dose rate goal should be associated with a FLASH threshold, meaning that if the goal is not met for that structure, the conditions to trigger FLASH cannot be satisfied completely and so any trade-off on dose is considered unnecessary.

After all objectives have been optimized, the minimum weight is maximized, through the same beam-current optimization routine used for the DADR objectives, with the difference that it only stops when the minimum weight cannot be increased any further. This makes the plan compatible with higher beam currents, as high as possible without sacrificing any of the objectives. The optimizer then returns a window of beam currents compatible with that treatment plan and that guarantee optimal FLASH coverage.

### 3.1.5 FLASH Thresholds

FLASH is here defined by doses above 8 Gy, delivered at dose-rates above 40 Gy/s and for irradiation times below 200 ms. This means that purely optimizing the dose rate does not necessarily translate into an optimization of the FLASH effect. For that, it is important to consider the other thresholds too, which must be specified in the wish-list.

Considering that the clinical requirements on dose are prioritized and that FLASH is still mostly not understood, it is currently not acceptable to directly optimize the dose to make it compatible with the FLASH effect. The dose objectives should therefore be prioritized over the dose rate objectives, at least when evaluating the same structure. The dose threshold should be only used *a posteriori*, to select the regions compatible with FLASH on dose and where optimization of the dose rate can be beneficial. By limiting the dose-rate optimization to just a set of selected voxels, the search space should get less constrained afterwards, better enabling the goals of other objectives to be met. This is implemented in the DADR optimization routines by selecting and optimizing only at the voxels with a dose higher than the FLASH dose threshold.

Although the dose rate is directly optimized, its threshold to trigger FLASH is also important because it represents a sufficient goal for the optimizer: higher dose rates will not increase the FLASH gains, at least given the current knowledge on the effect. This leaves more room to the optimization of the other dose and dose rate objectives, consequently translating into overall better and more compatible FLASH treatment plans. Erasmus-iCycle is already compatible with goals for the objectives, meaning that no additional implementations is required.

The dose rate threshold should also be taken into account when transforming the optimized DADR objective into a constraint. Because FLASH compatibility is to be ensured after optimization at each structure, dose rate constraints should only be set on the regions that have both a dose and a dose rate higher than the thresholds. The other objectives on different structures are then expected to be left with more room to get further improved, increasing the potential for better dose distributions and higher FLASH compatibility at other structures.

In summary, the FLASH dose and the dose-rate thresholds are essential to transform dose-rate optimization into FLASH optimization. By performing optimization and analysis on just these two thresholds, the irradiation time is not required to be taken into account because it gets automatically defined.

## 3.2 Pencil-Beam Scanning Dose Rate (PBSDR)

Optimization of the PBSDR is carried only at the scanning-pattern level, over IMPT dose-only-optimized treatment plans, meaning that the pencil-beam weights are not modified. Although weight readjustment may maximize further the FLASH coverage, different patterns have been observed to significantly modify the PBSDR distribution, without any trade-offs on the dose distribution to the patient. Scanning-pattern optimization (SPO) is a combinatorial optimization problem, which grows on difficulty the larger the target is, as a consequence of a higher number of pencil beams: the number of possible patterns increases factorially.

One of the main challenges of SPO comes from the non-linear and non-convex nature of the PBSDR metric. Furthermore, this metric depends on a dose threshold to determine when the effective irradiation time starts and when it stops. This makes the problem is more complex, extremely difficult for determinism optimizers. However, the study performed by Tsang [25] offers some hints, as it concludes the PBSDR is a local metric, as patterns with minimum path length were observed to have higher dose-rates. Therefore, short sequences of nearby pencil beams should always translate into higher PBSDR on the locally-irradiated voxels.

The previous observations suggest that SPO is 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 [41] (GA) are often discussed in literature to solve TSPs [42]. This subclass of Evolutionary Algorithms is inspired by Biology and works similarly to natural selection, using biologically inspired operators such as mutation, crossover and selection, to generate high-quality solutions to optimization problems [43]. 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: small sets of nearby pencil beams should give higher dose rates on the locally-irradiated voxels.

### 3.2.1 Genetic Algorithms

First introduced by Fraser [41] in 1957, Genetic Algorithms have become a hot research topic, with various applications in very diverse fields, such as biology, chemistry, computer aided design, cryptanalysis, microelectronics, production planning, robotics, etc. In these algorithms, the search space of a problem is represented as a collection of *individuals*, each represented by a string of characters, often referred to as *chromosomes*. The goal of the Genetic Algorithm is to find the individuals with the best genetic characteristics, which are compared on a *fitness function*. Because GAs are used in problems with very large search spaces, only a portion is analyzed at a time, making up what is referred to as *population*.

GAs work routinely in a three step procedure, similarly to how populations evolve in Nature: crossover, mutation and selection. First, the initial population is generated and the fitness of each individual is determined. Next, in every iteration of the algorithm, parents are selected from the population in pairs and produce children through *crossover*, which are then added to the population. These newly generated individuals correspond to a mix of the two parents and include genetic material from the two. On top of that, there is a probability, usually very small, for *mutation* of those new individuals, changing their hereditary qualities. Finally, the population is reduced to its initial size by a *selection* process. Each iteration of the algorithm is referred to as a *generation* and the whole routine represents the *evolution*. The algorithms stops according to a criteria that normally take into account the uniformity of the population, the convergence rate and the solution quality [42].

The convergence of GAs to solutions better than those of the initial population is assured by a balance between the crossover and the mutation operators. Crossover is responsible for sharing and preserving genetic material, enabling an increase of the average quality of the population. The best sets of chromosomes, referred to as the best *genes*, should be more represented by being part of more individuals, which is only possible through crossover. Mutation is needed to explore new states, helping the algorithm to avoid local optima by modifying the current genes. An adequate choice of crossover and

mutation operators and probabilities is important to have convergence to an optimum in a reasonable number of generations. In conclusion, crossover is important for convergence but it can lead to local optima, whereas mutation is important to avoid local optima but it can prevent convergence altogether.

### 3.2.2 Island Model Genetic Algorithm

Genetic Algorithms are stochastic optimization methods, meaning that they are only able to find sufficiently good solutions to the optimization problem, not guaranteeing that in independent runs, those solutions will be always the same and equally near the global optimum. This is a major concern on clinical applications since it is obviously not acceptable that the quality of the treatment depends on the random nature of the optimizer. However, this is not a concern here, because there is no trade-off on the quality of the treatment plan by optimizing the scanning pattern. Nonetheless, it is still considered important to improve the structure of the GA for higher consistency, guaranteeing higher FLASH gains than for the predefined patterns.

Distributed Genetic Algorithms (DGA) are a class of GAs that take into account information from different independent runs of simple GA models by running them in parallel at the same time and periodically exchanging information between them. These algorithms often lead to faster convergence and more consistent solutions for different runs, alongside superior numerical performance, when compared with the simple GA formulation [44].

There are different types of DGAs depending on how the simpler GAs are structured and on how the information is shared between them. In this implementation, an Island Model Genetic Algorithms (IMGA) was used. The idea behind the IMGA is to have several simple GA models focusing at different regions of the search space, by having them evolving independently from one another, like islands, each with its own unique population. From time to time, based on a *migration policy*, some individuals of each island migrate to a different one. This introduces new genetic material, similarly to mutation, with the key difference that the new genes should be specific of a different region in the search space. This should give more insight at the global scale of the problem and consequently, the population of all islands, as a whole, should more consistently converge to a global optimum.

Despite the different possible configurations, Genetic Algorithms have a common general structure, which is summarized in algorithm 3.2.1 for the specific Island Model version.

---

**Algorithm 3.2.1:** General structure of an Island Model Genetic Algorithm

---

*Initialize* population at random for every island

**while not** *Stop* **do**

**for each island do**

*Produce children* from the parents through crossover

*Mutate* the individuals

*Extend* the population adding the children and mutants to it

*Evaluate* the quality of the individuals based on the fitness function

*Reduce* the extended population through selection

**end**

**if conditions apply then** *migrate* individuals from one island to another

**end**

---

### 3.2.3 Scanning-Pattern Optimization

The choice of operators and parameters should allow to exploit the properties and characteristic of the PBSDR metric. The extensive studies already performed for the similar Traveling Salesman Problem [42] are a good starting point.


**Fitness Function** The goal of SPO is to maximize FLASH coverage, by maximizing the percentage of regions that having a dose above 8 Gy outside the tumour, also have a PBSDR above 40 Gy/s. Internally, this is also carried out by maximizing the PBSDR at the dose compatible regions, but only up until the FLASH dose rate threshold.

Setting the FLASH coverage as the fitness function might not be ideal, as its abstract meaning does not directly relate to any property of the patterns. Therefore, it might be useful to hint at the pattern shape, to speed up optimization and to help focusing on better directions in the search space. Based on the observations by Tsang [25] that patterns with a minimized total path length have higher dose rates, adding that contribution to the fitness function might be useful. Analyzing the total path length of several patterns and considering that the FLASH coverage is always between 0% and 100%, it was concluded that subtracting 10% of the average distance between pencil beams (the ratio of the total path length by the number of pencil beams) to the FLASH coverage should improve convergence and solution quality.

Adding information on the total path length of the pattern should also lead to better overall solutions, that minimize the scanning time between different pencil beams. Even though this is not taken into account in the PBSDR definition, it should have a non-negligible effect on the dose rate and consequently on FLASH.

**Problem Representation** The problem should be computationally implemented in such a way that the properties of the solution are naturally represented by the chromosomes and genes in the GA. In this case, the PBSDR is a very local metric that should give preference to sets of nearby beams, which should always translate into high dose rates on the locally irradiated voxels. Therefore, the scanning pattern can be intuitively represented by the ordered sequence of the pencil beam's index: if pencil beam  $i$  is the  $j$ -th element of the list, pencil beam  $i$  is the  $j$ -th beam to be irradiated.

This natural representation is denominated by *path representation* and is the one mostly used on GA applications to TSPs, both because of its intuitive nature as well as the good results obtained with it [42]. In figure 3.2.1 a simple example of two different path representations are given, corresponding to two different patterns for a  $3 \times 2$  regular grid:



**Figure 3.2.1:** Example of two different scanning patterns and their corresponding *path representation* for a  $3 \times 2$  regular grid.

**Population Size** The population is a critical part of the GA because it works as a natural memory, storing information regarding both the results of previous generations and the best genes. Therefore, it is important to have a large enough population so that more information can be preserved and consequently, better and more reliable solutions can be obtained [45]. However, the larger the population, the more individuals and, therefore, more memory is required. This also translates into more function evaluations, resulting into longer optimizations. Additionally, the convergence is also slower for larger population because it requires more generations to spread the best genes over the whole population. Therefore, a balanced population size is required.


For IMGAs, choosing the appropriate population size also needs to take into account the number of islands, because the population is to be equally spread over them. It is desirable to have a smaller population at each island to achieve local convergence on different regions of the search space [46]. It is also important to take into account the pattern size because the higher the number of pencil beams, the more complex the problem is, requiring larger populations.

In addition, the initial population must be initialized. Hints based on predefined scanning patterns can be provided, however, these might restraint the optimizer from exploring the whole search space, leading to suboptimal solutions. The simplest initialization is to consider random individuals, which despite translating into a slower convergence rate, guarantees a steadily and improving evolution. The slow convergence for the early generations can be overcome by sampling a higher number of random individuals, leading to an overall better but still distinct random population.

The initial population was randomly initialized, by sampling random permutations of the pencil beams making up 10 times the population size.

**Crossover Operator** A suitable crossover operator is, in general, one that mixes the genetic material of both parents in such a way that the resulting children inherit the best genes. This is a major problem for the path representation, as simply swapping sets of pencil beams between the parents results in unfeasible solutions, with several beams showing up more than once in the sequence. Therefore, crossover operators require a mechanism to repair the new solutions while keeping them as close to the original parents as possible, in order not to introduce mutations.

Different crossover operators have been suggested for the path representation and applied to the TSP problem [42]. One of the most commonly used is the *Partially-Mapped Crossover* (PMX) [47]. This crossover operator passes to the children information on the ordering of the parents, by mapping a portion of one parent's sequence into a portion of the other parent's sequence, with the remaining information being exchanged. A practical example of the operator generating one child is represented in figure 3.2.2.




**Figure 3.2.2:** Partially-Mapped Crossover (PMX) operator.

The PMX operator first selects uniformly at random two cut points along the sequences. As in figure 3.2.2, suppose that the first cut point is selected between the second and third sequence element, and

the second one at the last sequence element. The set between the cut points is called the *mapping section*. Considering the generation of just one child, the parent giving the information on the mapping section is the *donor* while the other parent is the *receptor*. The mapping process is then carried element by element: the *i*-th element on the donor's mapping section is placed at the same position in the receptor. The element that was previously in that position in the receptor moves to the initial position of the element that replaced it. In this example, the first element in the donor's mapping section is 2 at position 5 in the pattern. This is then placed at position 5 in the receptor, while beam 3, which was previously in that position, is moved to the previous position of beam 2 in the receptor, namely to position 2. The routine is repeated but inverting the role of donor and receptor, generating two distinct children.

The PMX operator allows to exploit the fact that the genes passed during crossover correspond to small sequences of beams. However, given the strong dependence of the PBSDR on the pattern, the repair mechanism of this operator might be too disturbing and significantly reduce the quality of the solutions. Moreover, while the mapping sequence of the donor is easily identified in one of the children, depending on the size of the pattern and the size of the mapping sequence, it is not always straightforward to identify the contribution of the receptor in that same children. This can be seen by comparing the initial parents and the final child in figure 3.2.2.

In order to allow more information of the receptor parent to be passed to the child, a new crossover operator based on the PMX is proposed: *Partially-Mapped Ordered Crossover* (PMOX). This new operator is also motivated by the idea that it should be the small sequence of beams that matter and not its location in the pattern. Therefore, in the PMX routine, instead of placing the first element of the mapping sequence in the receptor at the same position it appears on the donor, it is only used to identify the *mapping start* in the receptor. The other elements in the mapping sequence are then placed after the mapping start, following the same procedure as in the PMX routine. A practical example of the operator generating one child is represented in figure 3.2.3.




**Figure 3.2.3:** Partially-Mapped Ordered Crossover (PMOX) operator.

In this implementation, crossover was allowed for all individuals but it was associated with a probability. However, if the two parents are equal, the children are automatically excluded from the population, in order to avoid early dominance and suboptimal convergence. To further preserve information of the receptor parent and to focus on more local genes, the size of the mapping sequence was limited to a third of the total length of the pattern.


**Mutation Operator** A suitable mutation operator should be able to both fine tune the selected individual, so that the quality of the solution gets improved, and introduce substantial changes in the solution, so that the algorithm jumps out of local optima. A set of four mutation operators were used: swap, insertion, simple inversion and shift. Each mutation operator was associated with the same mutation probability but considered independently.

The *swap mutation operator* [48] randomly selects two pencil beams in the sequence and exchanges them. Considering the second and the fifth pencil beams are randomly selected, the effect of this operator is illustrated in figure 3.2.4. This operator is chosen because it allows for a fine-tune of the patterns, making only very small changes, although the quality might get substantially affected.




**Figure 3.2.4:** Swap mutation operator.

The *insertion mutation operator* [49] randomly chooses a pencil beam in the sequence, removes it and inserts it again but in a randomly selected position. This is represented in figure 3.2.5. Although translating into substantial changes in the absolute position of the pencil beams, this operator allows for a fine tune on the relative position, which should be useful.




**Figure 3.2.5:** Insert mutation operator.

The *simple inversion mutation operator* [50] selects randomly two cut points in the sequence and reverses the subsequence between them. Considering that the cut points are randomly chosen between the first and second elements and between the fourth and fifth elements, the result is illustrated in figure 3.2.6. The simple inversion mutation operator is chosen to possibly fine-tune the genes, which in this problem correspond to subsequences of pencil beams.



**Figure 3.2.6:** Simple inversion mutation operator.

The *shift mutation operator* [51] shifts the whole sequence to the right or to the left by a random number, making all the elements to be relocated to a new absolute position but preserving their relative order, as represented in figure 3.2.7. This operator is considered useful because it allows to experiment with different start and end beams, without substantially changing the PBSDR distribution, as the relative positions are preserved.



**Figure 3.2.7:** Shift mutation operator.

**Selection Mechanism** After new solutions have been generated, the population is reduced to its initial size by selecting the most suitable individuals. Because the selected individuals will be used to




generate new children in the next generation, the best should be selected so that the GA is driven to improve the population fitness over succeeding generations. The degree to which the better individuals are favoured is the *selection pressure*. It is critical to choose a balanced pressure because if it is too high, it is more probable for the GA to prematurely converge to a suboptimal solution, as diversity is killed off. If it is too low though, the convergence rate will be slow and the GA will take substantially longer to find the optimal solution, perhaps never finding it in a feasible time.

Tournament Selection [52] is a very common selection mechanism which applies selection pressure by holding a tournament among several individuals, as many as the *tournament size*. The tournaments are held up by randomly picking solutions from the population, which can be either replaced afterwards or removed entirely from the pool. The winner of the tournament is the individual with the highest fitness of those competitors and it is added to the next generation's population. This new population should have a higher average fitness than previous generations, as it will be only made up of tournament winners. Increased selection pressure can be provided by increasing the tournament size and replacing the competitors in the tournament pool. The winner of a larger tournament will, on average, have a higher fitness than the winner of a smaller tournament.

In this application, tournament selection was performed without replacement, which makes the tournament size fixed and given by the initial population size. The tournament pool is comprised of the best individuals from the previous generation, the new children and new mutants. Diversity is considered essential for SPO because the fitness function is a misleading indicator of how close a pattern is to the optimal sequence, regarding the relative order of the elements. Given an optimal scanning pattern, if the positions of two beams are swapped, the fitness function changes substantially but the patterns still remain similar. This justifies the focus on diversity within the population.

**Migration Policy** Migration is a unique concept to IMGAs. In this implementation, the islands are arranged in a ring and migration happens around the ring in an incremental way, simultaneously for all islands [46]. In the first migration, a fixed percentage of the population is randomly chosen and moved from their current island to their immediate neighbor to the left. In the second migration, new migrants are randomly picked and moved to the island which is two moves to the left in the ring. In general, the migration destination is incremented by 1 and moves around the ring, as illustrated in figure 3.2.8. Migrations occur every fixed number of generations until each island has sent one set of individuals to every other island (not including itself), at which point it is repeated from the beginning.



**Figure 3.2.8:** Migration routes for a ring-like organization of 4 islands.

The number of generations between migration must be high, so that each island is given room to locally converge. It cannot be too high though, or the newly added genetic material will either be irrelevant or dominate over all the population. A balance should be found so that the islands do not focus on completely different regions of the search space but instead get steered by each other towards a global


better search direction. As for the portion of migrating individuals, it should be high enough to introduce a considerable amount of new genes but not too high to kill the local diversity at the migration address.

The number of islands is also an important parameter as the more island, the better the search space is explored. However, if there are too many islands, depending on the migration policy, it might get more difficult to have a global convergence among all the islands. This is because they might converge to local optima before the information is shared between all of them.

Migration is here performed at the end of the generation, after selection. This was chosen so that the migrants are guaranteed to perform crossover with the natives at least once, enabling the introduction of new genetic material in the local population.

### 3.2.4 Fine-tune routine

The numerical parameters of the GA that can and should be fine-tuned to obtain the best solutions in the least amount of generations, with the highest consistency possible between independent runs and within the available resources. More specifically, there are six parameters that can be configured: population size, crossover probability, mutation probability, number of islands, migration period and migration size. The routine followed is reported in figure 3.2.9, with the configuration used being just one of many possibilities, aimed at obtaining the best solutions possible. This is not critical, as the goal of SPO is to improve the FLASH gains and it never introduces trade-offs on the dose.



**Figure 3.2.9:** IMGA fine-tune routine, with the parameters and explanation on the bottom, alongside the range, default and optimized values on top.

Fine-tuning should start without the island configuration (equivalent to one island only), in order to remove all IMGA related parameters, such as the number of islands and the migration policy. This way, the population size, the crossover probability and the mutation probability are the only parameters to assign. Different configurations should then be evaluated by running the algorithm on the same setup for an arbitrary small number of generations and comparing the fitness function, enabling quick testing of different parameters, one at a time. Overall, population size and crossover probability should be high, while mutation probability should be low. Tuning started with the population size since there is no range of values associated with it, fixing the crossover probability on 0.8 and mutation probability on 0.1. Since the parameters depend on each other during evolution, choosing a different tuning order might lead to different optimal configurations. Some parameters might compensate others, further highlighting that there isn't just one correct setup.

The problem gets substantially more complex and difficult for larger patterns, which should be taken into account in the population size. Considering that studies of GA applications on TSP fix the population size on values one order of magnitude higher than the number of cities, sizes equal to multiples of 10 times the number of pencil beams were tested. After trial and error on several test runs, it was concluded that a global population size of 50 times the number of pencil beams represented a good balance between convergence speed, solution quality and consistency.

The crossover probability was tuned next, varying it in steps of 0.05 around 0.8. It was ultimately concluded that 0.8 was a suitable value. Mutation probability followed, using the same tuning scheme but around 0.1. Several test runs showed that a mutation probability of 0.2 lead to the best results in the shortest number of generations.

The island configuration was then introduced, adding three extra parameters to be tuned, namely the number of island, the migration period and the migration size. Focus was then shifted to the consistency of the optimizer on different independent runs, requiring several runs for a small number of generations to test the parameters. Since the global population size is fixed, the number of islands determines the local population size of each island. The number of islands was first tuned, setting migration to happen every 25 generations with 10% of the local population.

Considering that the goal of the island model is for each island to focus on different regions of the search space, the number of islands should be set based on the population size per island and the number of pencil beams. The number of islands was then varied between 5, 10 and 20, corresponding to population sizes per island of 10, 5 and 2.5 times the number of pencil beams. These trials concluded that 20 islands lead to more reliable and better results between independent runs.

Migration policy followed, first the period and then the size. Several tests were performed, varying the number of generations between migration in steps of 5 around 25. It was concluded that 15 generations ensured local convergence without allowing the islands to drift too far apart. Migration sizes of 5%, 10% and 15% were tested, from which it was verified that 5% allowed the introduction of enough genetic material without dominating over the destination island.

For full evaluation, the stopping criteria was defined as 25 generations without improvement on any of the islands, allowing for each of them to fully converge to its optimum. No early stopping is performed, in order to achieve the best solutions possible with this optimizer, at the expense of the running time.

### 3.3 Optimization Setup

#### 3.3.1 Analytical 2D models

On the 2D approach, the analytical models are composed of a circular tumour with a annulus margin around it, in a square field, irradiated using pencil beams sampled from a regular square grid, selecting the ones falling inside the tumour. These pencil beams are modelled using standard Gaussian distributions with normalized amplitudes, so that the pencil-beam weight directly corresponds to the dose deposited at the peak.

The models are irradiated with variable number of pencil beams, fixing the beam width  $\sigma$  at 3 unit lengths while varying the beam separation  $\Delta\mu$ , through ratios  $\Delta\mu/\sigma = \{0.7, 0.8, 0.9, 1.0, 1.4, 2.0\}$ . The tumour radius is fixed at  $2.5\sigma$ , with a annulus margin of  $2.5\sigma$  defined around it, and the field size is set to  $20\sigma \times 20\sigma$ . To simulate a stereotactic treatment of lung lesions, the target dose at the tumour and its margin is set to 18 dose units and 0 outside.

**DADR** For DADR optimization, constraints are set based on the dose-only-optimized solutions. These are comprised of minimum constraints on the mean and minimum doses at the tumour and margin, with maximum constraints on the mean and maximum doses at the margin and surrounding area. To allow trade-offs when optimizing the dose rate, these constraints are relaxed through a relaxation factor  $\delta$ , equal to 3%. The beam current is given by the minimum pencil-beam weight, based on a pencil-beam

delivery time of 15 time units. This is chosen so that the dose-only-optimized plans require dose-rate optimization to be compatible with FLASH.

**PBSDR** For pattern optimization, the pencil-beam weights from the dose-only-optimized solutions are used and the optimizer is run 10 times for each ratio, in order to build up statistics on the overall performance. The beam current is fixed at 400 beam weight units per time unit, selected so that the predefined patterns translate in low dose rates but there is still room for optimization.

### 3.3.2 Clinical Scenario

Optimization on clinical data is performed on anonymized planning-CT scans and delineations of 12 patients with one single lung lesion, that were clinically treated with photons on a Cyberknife in 3 fractions of 18 Gy, prescribed to isodose levels ranging from 72%-80%, with a PTV not exceeding 10 cc (median 6.4 cc, range 4.4-10.1 cc). GTV delineations are available and a 5 mm CTV-PTV margin is applied. For scanning-pattern optimization, 3 additional patients are considered, characterized by larger PTVs, specifically 23.6 cc, 52.9 cc and 83.9 cc.

To make the analysis consistent and general, the same wish-list is used to generate the treatment plans for all the patients, which is based on previous studies on FLASH for the same clinical setup. This is only valid because the tumours considered are all lung lesions, so the same set of clinical requirements and constraints applies to all the patient population. The same regular grid was used to sample pencil beams for all treatment plans, with 5 mm of lateral spacing and energy of 244 MeV.

Due to the FLASH dose threshold, it is possible for plans to never be compatible with the effect if the target dose is too small. Therefore, to guarantee the highest FLASH compatibility, hypofractionation is considered, with a prescription dose of 18 Gy per fraction, on a total of 3 fractions, similarly to the conventional treatments delivered to these patients. Dose is delivered with 3 equiangular coplanar beams, one at each fraction: 40°, 160° and 280°. Conventionally, several beam directions would be delivered within the same fraction but current clinical technology does not allow to rotate the gantry within FLASH-compatible times.

The constraints and objectives specified in the wish-list depend on the structure, as reproduced in the appendix, in table 7.1.1. For the GTV and PTV, the dose is constrained to be higher than 98% of the target dose. Upper limits are not constrained there but set as objectives, to be optimized. At the regions on the vicinity of the PTV, objectives are set to minimize the minimum and the mean dose in a gradient of decreasing doses at shells around the PTV, to ensure better conformality. At the other healthy regions in the patient, the dose goals are set to 0. None of the constraints nor objectives is optimized robustly, as the focus of this project is on proof of principle.

**DADR** For the DADR optimization at the treatment planning level, the minimum pencil-beam delivery time is fixed to 5 ms. This is chosen so that the relation between the beam current and the minimum pencil-beam weight is more restrictive, consequently making optimization more complex and requiring a more careful balance between dose and dose rate. For this value, IMPT plans are only compatible with smaller beam currents. The DADR objectives are optimized after the dose objectives on a per-structure basis, except at the GTV. A final DADR objective is included, corresponding to optimization of the dose rate everywhere in the patient.

**PBSDR** Scanning-pattern optimization is performed on IMPT treatment plans, with a fixed beam current. If this value is too low, no plan will be FLASH compatible and scanning-pattern optimization

is ineffective. If the beam current is too high, every pattern is compatible with FLASH, making pattern optimization unnecessary. Therefore, a suitable beam current should be used, one that is clinically feasible, for which the predefined patterns show a low FLASH coverage, while leaving enough room for optimization. A beam current of 40 nA represents the right balance for this setup.

### 3.4 Evaluation

Despite the differences between the two dose-rate metrics, FLASH optimization aims to reduce the toxicity at the healthy tissue in both scenarios, meaning that evaluation and comparison of the results should be similar, using the same set of approaches.

#### 3.4.1 Plan comparison

The results of the proposed optimization strategies were compared on the FLASH coverage. This is the percentage of voxels outside the GTV that have FLASH compatible dose rates among the voxels irradiated with a FLASH compatible dose. The mean dose at the ipsilateral lung except the GTV is also used for comparison. This is a good proxy for damage to the 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 is useful to evaluate the overall gains by FLASH. The optimized results are compared against the golden standard for the mean ipsilateral lung dose, given by the conventional IMPT plans.

The healthy tissue sparing from FLASH is modeled by a FLASH enhancement ratio (FER), by considering the effect as binary based on the dose and dose-rate thresholds. This is a ratio by which dose at healthy tissue irradiated at FLASH-compatible conditions is divided, to model a reduced toxicity relative to IMPT-delivered dose. For lung tissue, studies have observed an FER of at least 1.8<sup>1</sup>, so it is the value considered here. An effective FER is also calculated, which represents the global FLASH gain at the ipsilateral lung except the GTV, if the whole volume was compatible with FLASH regardless of dose and dose rate. It is defined as the ratio between the conventional non-FLASH and the FLASH-enhanced doses, so it can easily be calculated with the mean dose.

On FLASH optimization with the DADR metric, since new treatment plans are generated, it is important to evaluate how do they compare with the conventional IMPT plans regarding dose. Therefore, the non-FLASH mean dose at the ipsilateral lung except the GTV sets the benchmark. The FLASH-enhanced dose is calculated for both the FLASH-optimized plans and the corresponding conventional IMPT plans, for which the dose rate is calculated based on the minimum pencil-beam weight and the minimum pencil-beam delivery time.

For scanning-pattern optimization, the FLASH-enhanced doses are compared for different predefined patterns: the *Z*, *Snake* and *Circle*. The best and worst-performing patterns found by the optimizer on the various independent runs are both evaluated. The current clinical standard for proton therapy using pencil-beam scanning is the *Snake* pattern, performed energy layer by energy layer.

On the analytical 2D models, the comparison with the dose-only-optimized solutions is different. For the DADR-optimized weights, since dose optimization is performed based on homogeneity to a target dose, the trade-offs are evaluated and compared on that objective function (quadratic difference). As for the FLASH gains, a similar objective function is used but taking the target as the prescribed dose at the GTV, zero outside the PTV and the FLASH-enhanced prescription dose at the PTV margin.

---

<sup>1</sup>This value is obtained from a study on thoracic irradiation of mice<sup>[8]</sup>, for which irradiation with 30 Gy at FLASH conditions was seen to be less fibrogenic than 17 Gy conventional (FER = 30/17 $\approx$ 1.8).

### **3.4.2 FLASH Sensitivity**

In order to evaluate how sensitive the FLASH-optimized solutions are, different dose and dose rate thresholds are considered. To evaluate how the thresholds reflect on the reduced toxicity achieved with FLASH, new FLASH enhanced dose distributions are calculated for the solutions optimized for FLASH with 8 Gy and 40 Gy/s, using dose values of {6, 8, 10, 12} Gy, alongside dose rate values of {30, 40, 50, 60} Gy/s. Both improvements on the mean dose at the ipsilateral lung and the effective FER are calculated, relative to conventional IMPT treatment plans.

## 4 Results

In this section, the results of the proposed optimization strategies are presented. First, the results for FLASH optimization at the treatment planning level with the DADR metric are presented, followed by scanning-pattern optimization for FLASH coverage maximization according to the PBSDR metric. The optimizers are first validated for analytical 2D models before being applied to clinical patient data.


### 4.1 Dose-Averaged Dose Rate (DADR)

The fully-optimized - dose followed by DADR - solutions are evaluated on the dose trade-offs and gains over dose-only-optimized solutions, the optimum weight distribution and the sensitivity of the optimized solutions to different FLASH thresholds.

For the early stopping criteria on optimization of both DADR objectives, defined by parameter  $\delta$ , the reported values for the objective function and for the corresponding linearized function values were analyzed at each step of the iterative linearization routines. It was concluded that for the minimum DADR,  $\delta = 0.01$  guarantees a difference on the order of  $10^{-2}$  Gy/s relative to the possible results of running the algorithm for longer. For the mean DADR,  $\delta = 100$  ensures a difference on the order of  $10^{-1}$  Gy/s relative to the results of letting the routine run for more steps.

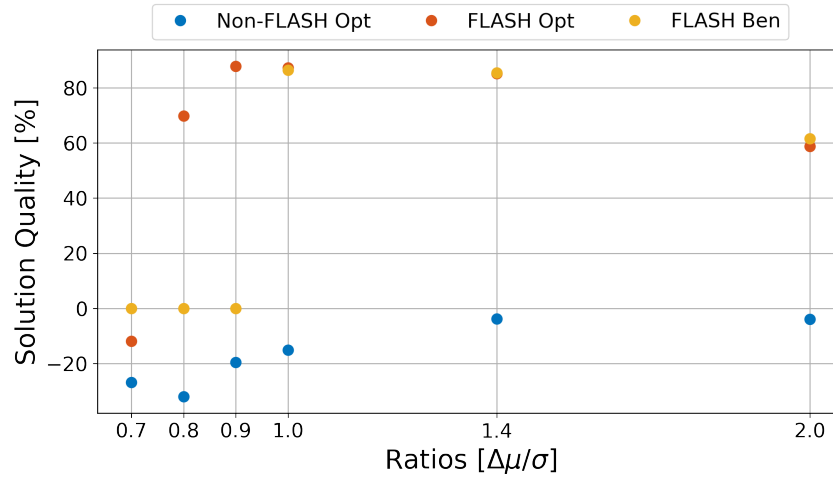
#### 4.1.1 Analytical 2D Models

The FLASH coverage for the fully-optimized solutions on the analytical 2D models with different ratios  $\Delta\mu/\sigma$  of beam separation and beam width is reported in figure 4.1.1. FLASH coverage is increased for all ratios but full coverage is only achieved for the highest ratios. The lowest ratio shows an improvement of only 2.53%.




**Figure 4.1.1:** FLASH coverage for fully-optimized solution on analytical 2D models with different ratios  $\Delta\mu/\sigma$ . The beam current is given by the minimum weight of each solution, considering a minimum delivery time of 15 time units.

For each ratio  $\Delta\mu/\sigma$ , the quality of the fully-optimized solutions is compared against the respective dose-only-optimized solutions. The results are reported in figure 4.1.2, with and without FLASH enhancement. Solution quality is evaluated using the dose homogeneity objective function, corresponding to the quadratic difference between the dose and a target. For the conventional dose, the target is set as the prescribed dose when inside the PTV, and zero when outside. FLASH gains are instead evaluated considering the prescribed dose at the GTV, zero outside the PTV and the FLASH enhanced prescription dose at the PTV margin. The solution quality degrades for all ratios when compared with the dose-only-optimized solutions, more considerably for the lower ratios.



**Figure 4.1.2:** Relative solution quality for fully-optimized solutions on analytical 2D models with different ratios  $\Delta\mu/\sigma$ , compared against the respective non-FLASH dose distribution of benchmark dose-only-optimized solution.

The ratio itself is also a measure of the quality of the plans since the solutions get better for lower ratios. To globally compare quality of different plans, since the setup is the same, the solution quality of each fully-optimized scenario is evaluated against the best solution, corresponding to the smallest ratio  $\Delta\mu/\sigma = 0.7$ . This is represented in figure 4.1.3. As expected, there is a degradation of the dose-only-optimized solutions as the ratio is increased.




**Figure 4.1.3:** Global solution quality for fully-optimized solutions on analytical 2D models with different ratios  $\Delta\mu/\sigma$ , compared against the best non-FLASH dose distribution, which corresponds to the solution for the smallest ratio 0.7.

Relating the global solution quality of figure 4.1.3 with the FLASH coverage in figure 4.1.1, it is observed that the ratio 1.0, although not fully FLASH compatible, has the same FLASH but better non-FLASH quality than its fully-optimized solution. This same solution has a better non-FLASH quality than the only FLASH compatible solution for ratio 0.9, which is the one obtained from full optimization. Similarly, for the smallest ratio 0.7, FLASH coverage is slightly increased through optimization but the corresponding FLASH quality is worse than the non-FLASH quality of the corresponding dose-only-optimized solution.

To understand how DADR optimization steers the spatial distributions of pencil-beam weights, optimization of the minimum and mean DADR is performed on a square field of size  $15\sigma \times 15\sigma$ , for the set of ratios  $\Delta\mu/\sigma = \{0.7, 1.0, 2.0\}$ . Only one constraint is applied, namely on the maximum dose, set to




18 dose units. In figure 4.1.4, the normalized spatial weight distribution is reproduced for these optimized solutions. These pencil-beam weight distributions are seen to be well-defined, symmetric and geometrically-organized, varying concentrically for different ratios.




**Figure 4.1.4:** Spatial distribution of the normalized pencil-beam weights for DADR-only-optimized solutions on analytical 2D square fields with different ratios  $\Delta\mu/\sigma$ . The normalized weight of each pencil beam is given by its color.

The dose distribution corresponding to the DADR-only-optimized solutions on the square fields is illustrated in figure 4.1.5. The dose distribution have shapes similar to the corresponding spatial pencil-beam weight distribution. The very concentric and smooth varying distribution for the lowest ratio gets broader and broader so that for the intermediate ratio it is very uniform. It broadens further for higher ratios until overlap between pencil beams is totally avoidable, resulting into very local peaks of high dose.



**Figure 4.1.5:** Dose distribution of DADR-only-optimized solutions for analytical 2D square fields with different ratios  $\Delta\mu/\sigma$ .

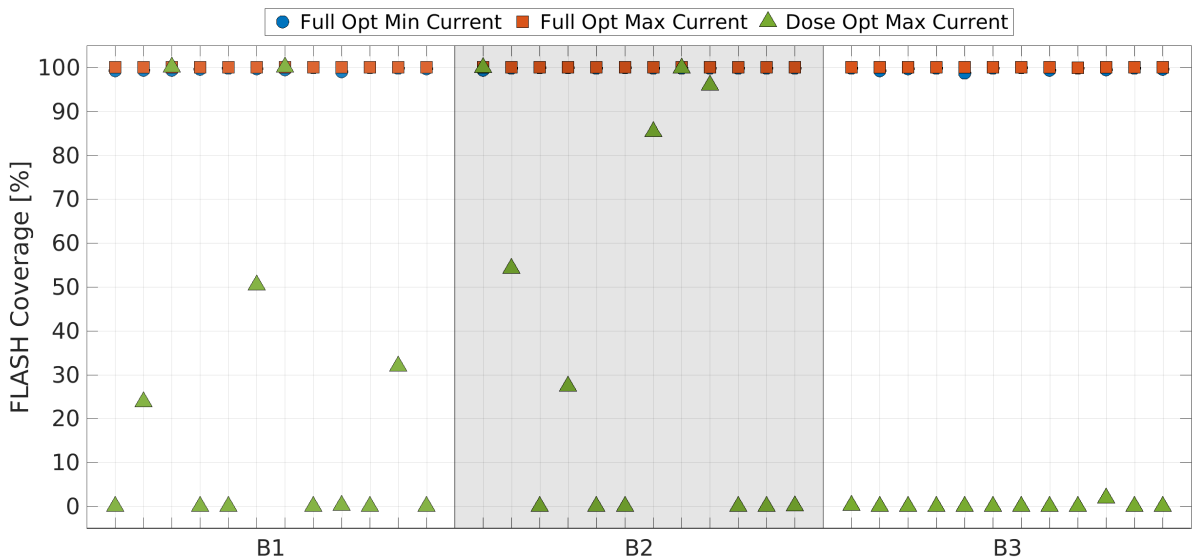
The corresponding dose rate distributions are reproduced in figure 4.1.6. Distributions gets less and less uniform the higher the ratio, broadening the range of values observed. For the lowest ratio, dose rate is uniformly constant everywhere with an intermediate value, whereas for the highest ratio the distribution is very peaked.



**Figure 4.1.6:** Dose-rate distribution of DADR-only-optimized solutions for analytical 2D square fields with different ratios  $\Delta\mu/\sigma$ . The beam current is fixed at 60 pencil-beam weight units per time unit.


#### 4.1.2 Clinical Scenario

For the clinical dataset of 12 patients with small lung lesions, fully-optimized - dose followed by DADR - treatment plans are generated. In figures 4.1.7, the FLASH coverage is reported, comparing the results for the minimum-required and maximum-compatible beam currents reported by the optimizer, against conventional IMPT plans, corresponding to dose-only optimization. All the optimized treatment plans have full FLASH coverage, with some IMPT plans already compatible with the FLASH effect, not requiring further optimization.




**Figure 4.1.7:** FLASH coverage for fully-optimized treatment plans against IMPT dose-only-optimized plans. Points on the horizontal correspond to the same patient, on the same direction. The beam current considered for conventional IMPT plans was calculated with a 5 ms minimum pencil-beam delivery time and the respective minimum weight.

However, focusing on the minimum required beam current reported by the optimized, zoomed in for figure 4.1.8, 100% coverage is never actually achieved. This is minor though, with very little impact on the results, as the median FLASH coverage is still 99.79%, with an interquartile range of 0.35%.



**Figure 4.1.8:** FLASH coverage for fully-optimized treatment plans, using the minimum-required and the maximum-compatible beam currents reported by the optimizer. Points on the horizontal correspond to the same patient, on the same direction.

Despite the substantial improvement of FLASH coverage in the fully-optimized plans, it is important to understand at what cost this is achieved. In figure 4.1.9, the mean dose at the ipsilateral lung is compared against the conventional IMPT treatment plans. The trade-off on dose is very small, while the FLASH-enhanced dose is always substantially better than the non-FLASH IMPT dose. On direction B3, the mean dose to the ipsilateral lung for some of the fully-optimized treatment plans is reportedly lower than for the corresponding IMPT plans.




**Figure 4.1.9:** Mean dose difference at the ipsilateral lung for fully-optimized treatment plans relative to conventional IMPT plans. FLASH enhanced dose considers a reduction of dose by 1.8 at the FLASH compatible regions. Points on the horizontal correspond to the same patient, on the same direction.

The optimized treatment plans have a median deterioration on the mean dose at the ipsilateral 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, respectively. Regarding improvements, results show medians of 38.27%, 38.08%

and 38.30% for directions  $B1$ ,  $B2$  and  $B3$ , respectively, alongside an interquartile range of 0.79%-point, 0.91%-point and 1.10%-point.


With the proposed optimization strategies for full optimization, the minimum-required beam current to achieve optimal FLASH compatibility is calculated, according to the wish-list configuration. In addition, the treatment plan are also made compatible with higher beam currents by maximizing the minimum weight at the last step. This translates into windows of beam currents compatible with each optimized treatment plan and that guarantee optimal FLASH coverage. The different ranges for the fully-optimized treatment plans are reported in figure 4.1.10.



**Figure 4.1.10:** Beam-current windows for fully-optimized treatment plans, consisting of the minimum-required beam current and the maximum compatible beam current for the highest FLASH coverage within the specifications of the wish-list.


The majority of the plans is compatible with a wide window of beam currents. For directions  $B1$ ,  $B2$  and  $B3$ , the minimum required beam current has a median of 19.41 nA, 21.20 nA and 30.77 nA, alongside an interquartile range of 2.33 nA, 2.44 nA and 15.02 nA respectively. For the beam current window of [33,56] nA, 80% of the plans are fully compatible.

Understanding how the dose rate changes within the patient can give insight about what directions should be chosen for higher FLASH compatibility. In figure 4.1.11, the DADR distribution is reproduced along the beam direction for different slices corresponding to different depths in the patient. The deeper into the patient, the lower the overall dose rate.



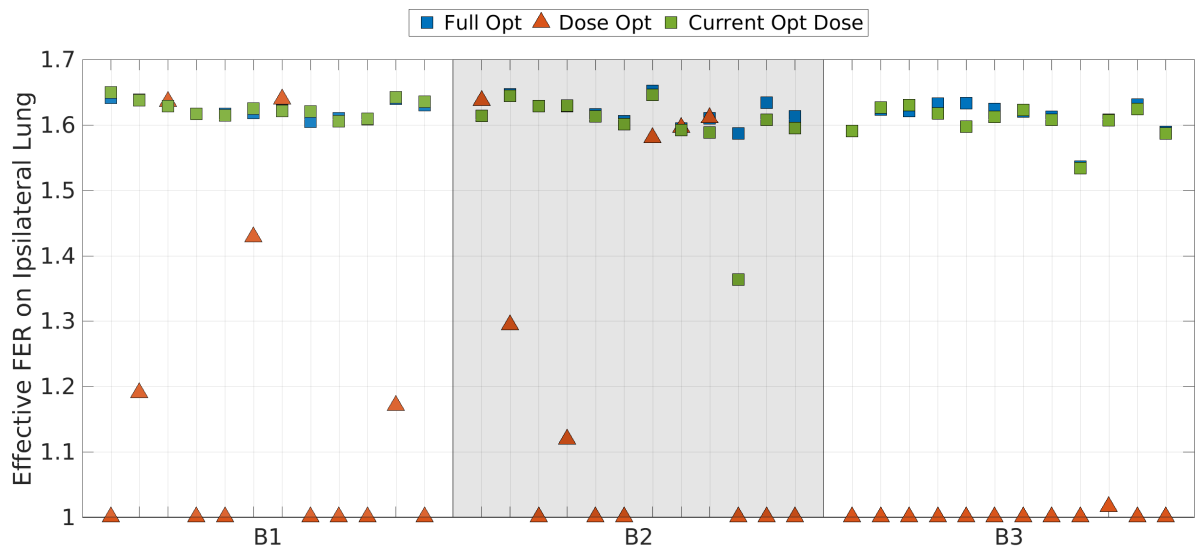
**Figure 4.1.11:** DADR distribution within a patient with a small lung lesion, along the beam direction for direction  $B1$ , at different slices corresponding to different depths in the patient. The title indicates the depth relative to the first slice, in mm.

To compare the contribution and importance of the two steps in DADR optimization, namely beam-current optimization and weight readjustment, new treatment plans are generated for the same patient data but performing current optimization. This is defined as dose-only optimization followed by a last step of beam-current maximization. In figure 4.1.12, the dose at the mean ipsilateral lung is reported against the conventional IMPT plans, whereas the effective FER is reported in figure 4.1.13.




**Figure 4.1.12:** Mean dose difference at the ipsilateral lung for fully-optimized and current-optimized treatment plans relative to conventional IMPT plans, considering both non-FLASH and FLASH-enhanced dose. Points on the horizontal correspond to the same patient, on the same direction.

The mean dose improvements at the ipsilateral lung for the current-optimized plans is increased by a median of 0.02%-point, -0.14%-point and -0.08%-point over the fully-optimized treatment plans, with an interquartile range of 0.26%-point, 0.74%-point and 0.41%-point for directions *B1*, *B2* and *B3*, respectively. Regarding dose degradation, it is reduced by a median of 0.06%-point, -0.07%-point and -0.22%-point for the same directions, with an interquartile range of 0.46%-point, 0.75%-point and 0.81%-point.



**Figure 4.1.13:** Effective FLASH Enhancement Ratio at the ipsilateral lung for fully-optimized and current-optimized treatment plans against conventional IMPT plans. The effective FER corresponds to the ratio between the conventional and the FLASH-enhanced (mean) dose. Points on the horizontal correspond to the same patient, on the same direction.

To evaluate how sensitive the treatment plans are to different FLASH thresholds, FLASH-enhanced dose distributions are calculated for each plan using various thresholds and the maximum-compatible beam current. The contribution of each direction on the treatment of the patient is summed up to report full treatment quality. The results are evaluated on the mean dose to the ipsilateral lung in figure 4.1.14 and on the effective FLASH Enhancement Ratio in figure 4.1.15.

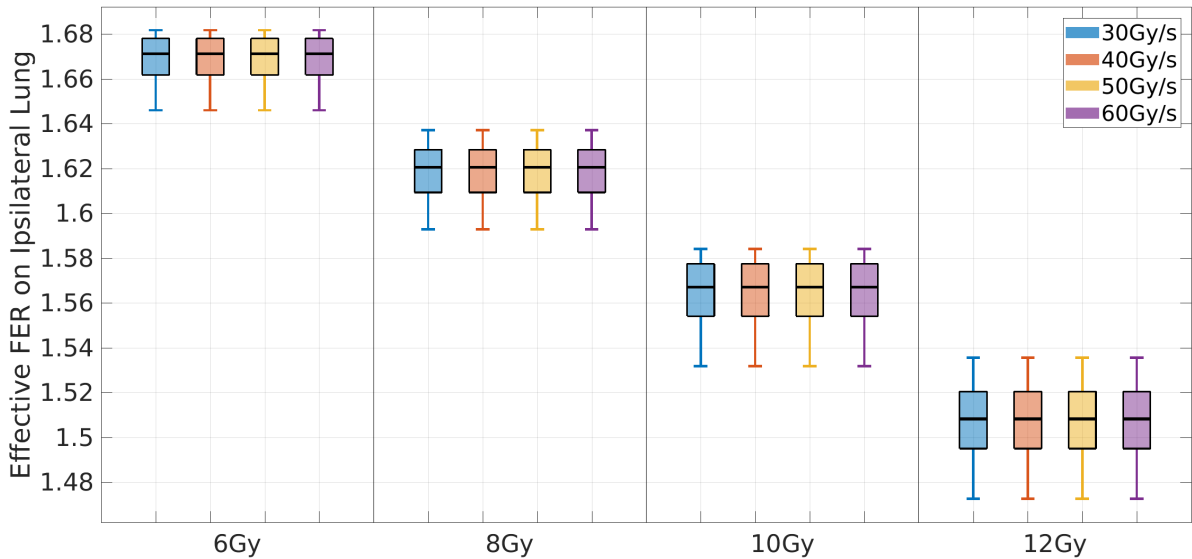


**Figure 4.1.14:** Sensitivity of the improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the maximum-compatible beam current, relative to conventional IMPT treatment plans.

The results are very consistent, showing the same performance for different dose-rate thresholds and an expected degradation for higher dose thresholds. The median improvements are reported in table 4.1.1, considering the overall treatment of the patient, summing up the doses of each direction.


Dose\Dose rate	30 Gy/s	40 Gy/s	50 Gy/s	60 Gy/s
6 Gy	40.16%	40.16%	40.16%	40.16%
8 Gy	38.29%	38.29%	38.29%	38.29%
10 Gy	36.18%	36.18%	36.18%	36.18%
12 Gy	33.69%	33.69%	33.69%	33.69%

**Table 4.1.1:** Median improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the maximum-compatible beam current, relative to conventional IMPT treatment plans.



**Figure 4.1.15:** Sensitivity of the effective FLASH Enhancement Ratio at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds, calculated with the maximum-compatible beam current. The effective FER corresponds to the ratio between the conventional and the FLASH-enhanced (mean) dose.

The same sensitivity analysis was performed but instead considering the minimum-required beam current for each plan. The results are reported on the FLASH improvements in figure 4.1.16 and on the effective FLASH Enhancement Ratio in figure 4.1.17.




**Figure 4.1.16:** Sensitivity of the improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the minimum-required beam current, relative to conventional IMPT treatment plans.

The results are still consistent, similar to those for the maximum possible beam current, although the higher dose-rate thresholds now translate into some degradation. The median improvements are reported in table 4.1.2, considering the overall treatment of the patient, summing up the doses of each direction.

Dose\Dose rate	30 Gy/s	40 Gy/s	50 Gy/s	60 Gy/s
6 Gy	40.14%	39.96%	39.28%	35.99%
8 Gy	38.29%	38.26%	37.99%	35.39%
10 Gy	36.18%	36.18%	36.14%	34.20%
12 Gy	33.69%	33.69%	33.67%	32.27%

**Table 4.1.2:** Median improvements on the mean dose at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the minimum-required beam current, relative to conventional IMPT treatment plans.



**Figure 4.1.17:** Sensitivity of the effective FLASH Enhancement Ratio at the ipsilateral lung for fully-optimized treatment plans on different FLASH thresholds with the minimum-required beam current. The effective FER corresponds to the ratio between the conventional and the FLASH-enhanced (mean) dose.

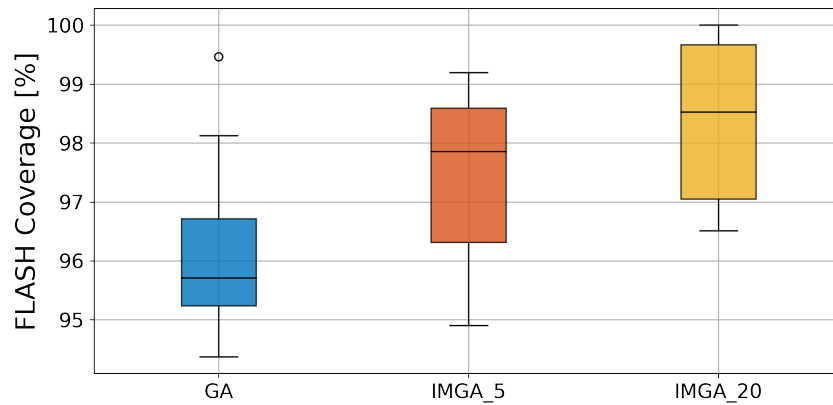
## 4.2 Pencil-Beam Scanning Dose Rate (PBSDR)

Scanning-pattern optimization for FLASH coverage maximization according to the PBSDR metric is evaluated on the performance and convergence of the optimizer, comparing different configurations and different predefined patterns. The sensitivity of the FLASH gains for the optimized scanning-patterns to different FLASH thresholds is also evaluated.

### 4.2.1 Analytical 2D Models

The quality of the different scanning patterns obtained for 10 independent runs of different versions of the scanning-pattern optimizer is reported in figure 4.2.1, for the ratio  $\Delta\mu/\sigma = 1$ . All the versions evaluated have consistent results, with the Island Model Genetic Algorithm finding the best patterns, within a shorter window of possible solutions.






**Figure 4.2.1:** Relative FLASH coverage for the suboptimal patterns generated with 10 independent runs of different versions of the scanning-pattern optimizer on analytical 2D models with ratio  $\Delta\mu/\sigma = 1$ . The results are normalized to the maximum found for the total 30 runs.

For ratio  $\Delta\mu/\sigma = 1$ , the solutions for the three tested versions of the scanning-pattern optimizer have a median FLASH coverage relative to the best found pattern of 95.71%, 97.86% and 98.53% for the simple GA, IMGGA with 5 islands and with 20 islands, respectively, with a full range, i.e. the distance between the overall minimum and maximum values, of 5.09%-point, 4.29%-point and 3.48%-point.


The performance of different versions of the scanning-pattern optimizer is reported in figure 4.2.2 as a function of the running time for the ratio  $\Delta\mu/\sigma = 1$ . Bootstrapping is performed with data from 10 runs, drawing data samples of size up to 10 and calculating the maximum FLASH coverage for each sample. For every number of runs, bootstrapping is performed 1000 times and the median relative FLASH coverage is calculated. If enough runs are performed, the optimizer will be more likely to find a better solution, at the expense of longer computational times.



**Figure 4.2.2:** Relative FLASH coverage for different versions of the scanning-pattern optimizer on analytical 2D models with ratio  $\Delta\mu/\sigma = 1$  as a function of the running time. This corresponds to as many runs as the order of the point in the x axis and determines the sample size for bootstrapping. The reported FLASH coverage is the median of the maximum coverage for 1000 samples obtained through bootstrapping, normalized to the maximum value found on all runs for all versions.


The Island Model Genetic Algorithm with 20 islands always performs better but in the time to run it just once, the simple GA can be run for almost three times. The IMGGA with 5 islands has better relative FLASH coverage in less runs than the simple GA, but because each one takes substantially longer, at some point it becomes more advantageous to run more versions of the simpler faster GA model.

In order to understand how the complexity of the algorithm can increase the running time and how the algorithm performs for configurations different from the one used to fine-tune it, the relative FLASH coverage is evaluated as a function of the running time for different ratios. These results are reported in figure 4.2.3, using the final version of the scanning-pattern optimizer: IMGGA with 20 islands. The complexity of the problem increases with the number of pencil beams, translating into substantially longer running times.



**Figure 4.2.3:** Relative FLASH coverage on analytical 2D models for different numbers of pencil beams, given by the different ratios  $\Delta\mu/\sigma$ , as a function of the running time. The running time corresponds to as many runs as the order of the point in the x axis, which determines the sample size for bootstrapping. The reported FLASH coverage is the median of the maximum coverage for 1000 samples obtained through bootstrapping, normalized to the maximum FLASH coverage. The higher the number of pencil beams, the lower the ratio.


In figure 4.2.4, the FLASH coverage is reported for the best and worst-performing optimized scanning patterns, against predefined patterns. FLASH coverage is always higher for the optimized scanning patterns, independently of the ratio  $\Delta\mu/\sigma$ , with higher improvements for the lower ratios. The predefined patterns significantly underperform for a large number of pencil beams.



**Figure 4.2.4:** FLASH coverage for optimized and predefined scanning patterns on analytical 2D models with different ratios  $\Delta\mu/\sigma$ . The *Max* and *Min* correspond to the patterns with maximum and the minimum FLASH coverage reported by 10 independent runs of the optimizer.


The FLASH coverage for the best-performing scanning patterns found by the optimizer is similar and consistent for the different ratios, with a median of 76.83% and a full range of 5.86%-point. For ratios  $\Delta\mu/\sigma = \{0.7, 1.0, 2.0\}$ , there's an improvement over the best predefined patterns of 544.10%, 316.57% and 48.46% respectively.

In order to better understand how the optimizer improves FLASH coverage, the best optimized scanning patterns are reproduced in figure 4.2.5, found for 10 independent runs. The overall shape of patterns is consistently similar: closed circular loops, following a radial inwards-outwards movement around the GTV, resembling snowflakes. In addition, they all have one end of the pattern in the center, within the GTV, while closing the circular loop in the same initial radial direction.




**Figure 4.2.5:** Best-performing optimized patterns for analytical 2D models with different ratios  $\Delta\mu/\sigma$ , out of 10 independent runs. Patterns start with a green square and end with a red square. The GTV is identified by the light circle.

In figure 4.2.6, the PBSDR distributions of the patterns in figure 4.2.5 for  $\Delta\mu/\sigma = \{0.7, 1.0, 2.0\}$  are reproduced. These dose-rate distributions have similar shapes for the different ratios, all resembling the letter C, with sacrifices at the interior core region, corresponding to the GTV, and along the direction of the start and end branches.




**Figure 4.2.6:** PBSDR distributions of the best-performing optimized patterns for analytical 2D models with different ratios  $\Delta\mu/\sigma$ , out of 10 independent runs. The GTV is identified by the light circle, with the corresponding pattern reproduced on top, starting with a green square and ending with a red one.

To analyze the impact of objective's targeted region scanning-pattern optimization is performed including the GTV. The best-performing scanning patterns found by the optimizer and their dose rate distributions are reproduced in figures 4.2.7 and 4.2.7, found for 10 independent runs for ratios  $\Delta\mu/\sigma = \{0.7, 1.0, 2.0\}$ .




**Figure 4.2.7:** Best-performing optimized patterns for analytical 2D models with different ratios  $\Delta\mu/\sigma$ , optimized including the GTV, out of 10 independent runs. Patterns start with a green square and end with red square. The GTV is identified by the light circle.

Similar snowflake-shaped patterns are obtained, consequently translating into identical C-shaped PB-SDR distributions. Some differences are seen at the GTV region, with the newly optimized patterns reporting higher dose rates, specially at the margin regions near the GTV boarder.



**Figure 4.2.8:** PBSDR distributions of the best-performing optimized patterns for analytical 2D models with different ratios  $\Delta\mu/\sigma$  optimized including the GTV, out of 10 independent runs. The GTV is identified by the light circle, with the corresponding pattern reproduced on top, starting with a green square and ending with a red one.


The FLASH coverage of the patterns optimized excluding and including the GTV is compared in figure 4.2.9, only considering voxels outside the GTV for evaluation. FLASH coverage is generally better for the scanning patterns optimized taking into account the GTV. The median FLASH coverage increases by 0.21%-point, 1.22%-point and -0.20%-point over the previous results, the maximum values increase by 0.83%-point, 0.70%-point and -0.20%-point and minimum values increase by 1.10%-point, 2.49%-point and -0.20%-point, for ratios  $\Delta\mu/\sigma = \{0.7, 1.0, 2.0\}$ , respectively.



**Figure 4.2.9:** FLASH coverage for patterns optimized with and without the GTV for different ratios  $\Delta\mu/\sigma$  for 10 independent runs on analytical 2D models.

#### 4.2.2 Clinical Scenario


For one patient, with a PTV of 8.2 cc, different versions of the scanning-pattern optimizer are tested in order to validate the configuration proposed. Independent runs are performed for each version and for each treatment plan, corresponding to a different direction. The results are compared on the FLASH coverage, for which bootstrapping is performed with data from 20 runs, drawing bootstrapped data samples of size up to 20 and calculating the maximum FLASH coverage for each sample. For every number of runs, bootstrapping is performed 2000 times and the median is computed. The results are reported in figure 4.2.10 and 4.2.11.



**Figure 4.2.10:** Relative FLASH coverage for suboptimal patterns generated with 20 independent runs of different versions of the scanning-pattern optimizer on clinical data. The reported FLASH coverage is normalized to the maximum found for each direction.


The three proposed versions of the scanning-pattern optimizer consistently find patterns near the optimum, with the best performance being obtained with the Island Model Genetic Algorithm using 20 islands. The GA, IMGA with 5 islands and with 20 islands achieve a median relative FLASH coverage of 99.45%, 99.45% and 99.63% respectively, when grouping the directions together, with a full range of 12.88%-point, 2.33%-point and 2.33%-point. Regarding computational time, the IMGA with 20 islands

is the best for direction *B2*, whereas direction *B3* is seen to be fully optimized for just one run, with any of the versions. Direction *B1* shows the most consistent performance for IMGA with 5 islands in short running times, but this version never finds the results achieved by the 20 islands or even the simple GA.



**Figure 4.2.11:** Relative FLASH coverage for 20 independent runs of different versions of the scanning-pattern optimizer, on the same patient data with a PTV of 8.2 cc. The running time corresponds to as many runs as the order of the point in the x axis, which determines the sample size for bootstrapping. The reported FLASH coverage is the median of the maximum coverage for 2000 samples obtained through bootstrapping, normalized to the maximum FLASH coverage.


The performance of the final configuration of the optimizer - IMGA with 20 islands - is also evaluated on convergence, by plotting the median relative coverage out of all the independent runs for every generation of the Genetic Algorithm. Because larger tumours require the irradiation of more pencil beams, the patterns get larger and the problem gets more complex. To evaluate the general performance of the optimizer, a patient with a PTV of 83.9 cc is also considered. In figure 4.2.12, the convergence of the algorithm is reproduced.



**Figure 4.2.12:** Convergence of the scanning pattern optimizer for a patient with a PTV of 8.2 cc and another with 83.9 cc, on three beam directions. The number of generations and the metric are normalized to the values reported by the best run.

For the smaller PTV, the convergence of the algorithm is relatively steep in the beginning and almost linear up until very near the optimal solution, with 80% of the performance obtained in 23% of the total number of generations, for the direction with the worst median results. For the larger PTV, despite a slower convergence, the algorithm is still relatively quick and solutions improve almost linearly, taking up to 48% of the generations to reach 80% of the best performance. The running time increases on average from 6 minutes to 4 hours for a single run, from the small lung lesion to the larger one.


To evaluate the overall consistency of the optimizer, the results from 20 independent runs for each treatment plan are compared, on a patient and direction basis. In order to better visualize the window of possible values, the results are normalized to the best metric found for the respective treatment plan. This is reported in figure 4.2.13.



**Figure 4.2.13:** Optimized scanning patterns for clinical data of 12 patients with small lung lesions, for three different directions. The optimizer is independently run 20 times for each treatment plan and the result are normalized to the maximum FLASH coverage found.


Evaluating the results on a direction basis, the optimizer achieves a median relative FLASH coverage of 99.83%, 100% and 100%, with an interquartile range of 1.16%-point, 1.01%-point and 0.24%-point for directions *B1*, *B2* and *B3*, respectively. Average running times of 8 minutes are obtained, corresponding to a median number of pencil beams of 34.

The improvements achieved by scanning-pattern optimization are evaluated by comparing the FLASH gains for the best and worst-performing optimized patterns against predefined ones. These results are compared on the improvements over the mean dose to the ipsilateral lung for conventional IMPT plans and on the effective FLASH Enhancement Ratio, reproduced in figures 4.2.14 and 4.2.15.



**Figure 4.2.14:** Improvement on the mean dose at the ipsilateral lung over IMPT plans for different scanning patterns, including predefined ones such as the *Z*, *Snake* and *Circle* from left to right, and the *Max* and *Min*, corresponding to the patterns with maximum and the minimum FLASH coverage found for 20 runs of the optimizer.


The optimized patterns are considerably better than any of the predefined, with higher FLASH gains. Relative median improvements of 9.34%-point, 10.79%-point and 5.45%-point over those for the best predefined patterns are obtained for the optimized patterns, on directions *B1*, *B2* and *B3*, respectively, with interquartile ranges of 4.76%-point, 4.84%-point and 6.24%-point. The median difference between the FLASH improvements of the best and worst-performing optimized patterns is 0.44%-point, 0.33%-point and 0.33%-point, with interquartile ranges of 0.83%-point, 0.91%-point and 0.96%-point, respectively for each direction.



**Figure 4.2.15:** Effective FLASH Enhancement Ratio at the ipsilateral lung for different scanning patterns, including predefined ones and the *Max* and *Min*, corresponding to the patterns with maximum and the minimum FLASH coverage found for 20 runs of the optimizer. The effective FER corresponds to the ratio between the conventional and the FLASH-enhanced (mean) dose.




To understand how the best FLASH coverage is achieved, the best-performing patterns out of 20 independent runs for four different patients are reproduced in figure 4.2.16, for direction *B1*. These optimized patterns show the same snowflake shape: closed circular loop, following a radial inwards-outwards movement around a central region.



**Figure 4.2.16:** Best-performing optimized patterns for different patients with small lung lesions, visualized along the beam direction *B1*. Patterns start with a green square and end with a red square.


To evaluate the scalability of the optimizer, scanning patterns were optimized for 3 additional patients, characterized by a larger PTVs of 23.6 cc, 52.9 cc and 83.9 cc. The window of values obtained for 10 independent runs of the optimizer for each patient and direction are reported in figure 4.2.17, also including the results for the patient with PTV of 8.2 cc.



**Figure 4.2.17:** Optimized scanning patterns for four patients with different lung tumour volume, for three different directions. The optimizer is independently run 10 times and the result are normalized to the highest metric found.


Despite the increase in complexity, the optimizer finds scanning patterns that are better than the predefined ones for all volumes, with consistently higher FLASH coverage within narrow windows of possible values. For volume sizes of 8.2 cc, 23.6 cc, 52.9 cc and 83.9 cc, median relative FLASH coverage of 99.64%, 99.69%, 98.75% and 97.98% are reported, respectively, with interquartile ranges of 1.22%-point, 0.62%-point, 1.80%-point and 2.88%-point. The average running time for a single run is 6 minutes, 24 minutes, 3 hours and 4 hours, with 36, 64, 107 and 159 median numbers of pencil beams, respectively.

In figure 4.2.18 is reported the FLASH improvements over the mean dose to the ipsilateral lung for conventional IMPT plans, while in figure 4.2.19 is reported the effective FLASH Enhancement Ratio on the same structure. In all scenarios, the optimized patterns are considerably better than the best predefined, enabling more FLASH healthy-tissue sparing.




**Figure 4.2.18:** Improvement on the mean dose at the ipsilateral lung over IMPT plans for various lung tumour volumes using predefined patterns such as the Z, Snake, Circle from left to right, and the Max and Min, corresponding to the patterns with the maximum and minimum FLASH coverage found for 10 runs of the optimizer.

Comparing the FLASH-enhanced mean dose at the ipsilateral lung for the optimized scanning patterns against the best predefined patterns, the gains get higher as the tumour volume increases. The improvements on the mean dose at the ipsilateral lung for the optimized patterns have a median increase of 8.38%-point, 22.25%-point, 17.21%-point and 16.98%-point, over the best predefined patterns, for volumes 8.2 cc, 23.6 cc, 52.9 cc and 83.9 cc, respectively. The median difference between the FLASH improvement of the best and worst-performing optimized patterns is 5.55%-point, 3.61%-point, 7.19%-point and 1.68%-point, respectively for each tumour volume.




**Figure 4.2.19:** Effective FLASH Enhancement Ratio at the ipsilateral lung for various lung tumour volumes using predefined patterns and the *Max* and *Min*, corresponding to the patterns with maximum and the minimum FLASH coverage reported by 10 runs of the optimizer. The effective FER corresponds to the ratio between the conventional and the FLASH-enhanced (mean) dose.

The best-performing optimized pattern for each of the patients with different tumor volumes is reproduced in figure 4.2.20, for direction *B1*, with the corresponding the PBSDR distributions reported in figure 4.2.21, along the beam direction, in an arbitrary slice and overlapped with the scanning pattern.



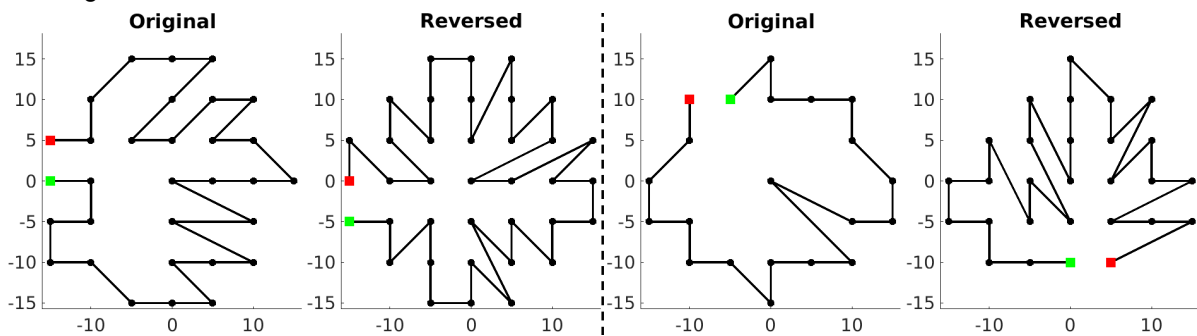
**Figure 4.2.20:** Best-performing optimized patterns for various lung tumour volumes, visualized along the beam direction for direction *B1*. Patterns start with a green square and end with a red square.

The optimized patterns show the same snowflake shape, with movement around a closed circular loop, following a radial inwards-outwards around a central region. 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 4.2.21:** PBSDR distributions of the best-performing optimized patterns for various lung tumour volumes, along the beam direction for direction  $B1$ , overlapped with the corresponding scanning pattern. Patterns start with a green square and end with a red square.

For one of the directions of the treatment plans, namely  $B3$ , the optimized patterns show a wide range of values, suggesting a lower compatibility with the FLASH effect. To understand the source of this inconsistent performance, new treatment plans are generated inverting the direction, which should translate into equally good treatment plans from a dose perspective. The scanning patterns for the original and the reversed directions are reproduced in figure 4.2.22, for two patients where reversing increases FLASH coverage.



**Figure 4.2.22:** Best-performing optimized patterns for the original and the reversed direction  $B3$ , on two patients with improved FLASH coverage on reversal. Patterns start with a green square and end with a red square.

The results for the new optimized scanning patterns are compared to the original beam direction on the improvements of the mean dose to the ipsilateral lung over conventional IMPT plans in figure 4.2.23, and on the effective FER in figure 4.2.24.




**Figure 4.2.23:** Improvement on the mean dose at the ipsilateral lung for the original and reversed direction *B3*, including the best predefined and best-performing optimized patterns, found for 20 runs of the optimizer.

Whenever beneficial, reversing the beam direction leads to a median increase of the FLASH mean dose improvements at the ipsilateral lung of 13.81%-point, over the optimized results for the original direction, with an interquartile range of 13.16%-point.



**Figure 4.2.24:** Effective FLASH Enhancement Ratio at the ipsilateral lung for the original and reversed direction *B3*, including the best predefined and best-performing patterns, reported by 20 runs of the optimizer. The effective FER corresponds to the ratio between the conventional and the FLASH-enhanced (mean) dose.

The sensitivity of the optimized solutions to different FLASH thresholds is evaluated on the mean dose improvements over the conventional IMPT plans and on the effective FLASH Enhancement Ratio, in figures 4.1.14 and 4.1.15, respectively. This is calculated from FLASH-enhanced dose distributions for each plan on each beam direction, summed up together to translate the full treatment, considering various thresholds and the best-performing optimized scanning pattern.

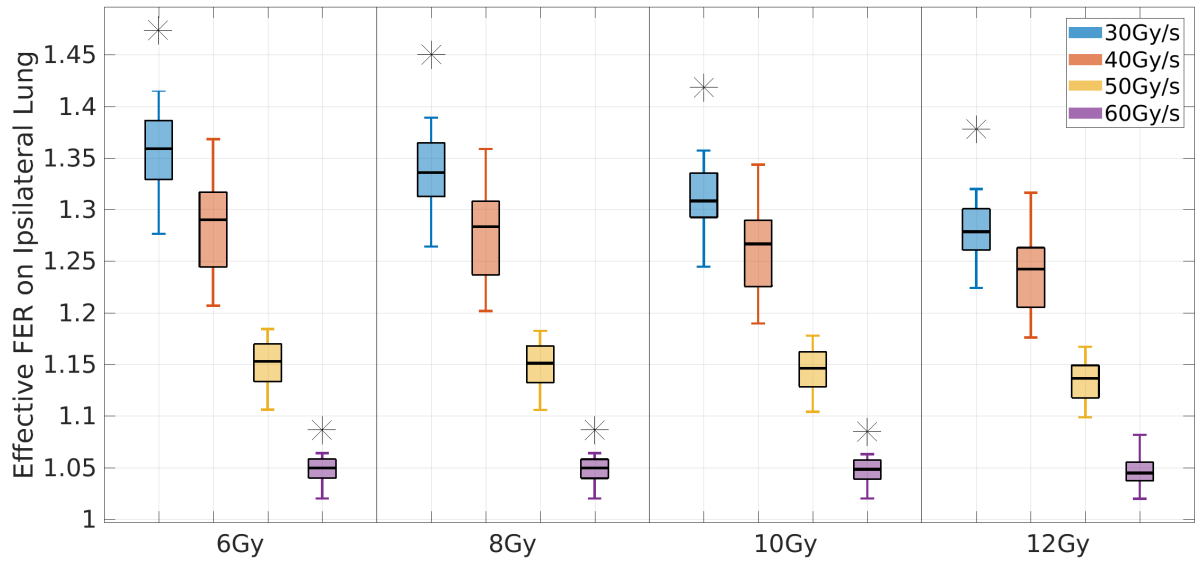


**Figure 4.2.25:** Sensitivity of the improvement on the mean dose at the ipsilateral lung for the best-performing optimized scanning patterns found for 20 independent runs, on different FLASH thresholds.

The results are relatively consistent on the dose threshold but degrade considerably on the dose-rate threshold. The median improvements over the mean dose for conventional IMPT plans are reported in table 4.2.1, considering the overall treatment of the patient, summing up the doses of each direction.

Dose\Dose rate	30 Gy/s	40 Gy/s	50 Gy/s	60 Gy/s
6 Gy	26.40%	22.47%	13.26%	4.72%
8 Gy	25.14%	22.07%	13.12%	4.70%
10 Gy	23.56%	21.05%	12.76%	4.60%
12 Gy	21.78%	19.50%	12.01%	4.28%

**Table 4.2.1:** Median improvements on the mean dose at the ipsilateral lung for the best-performing optimized scanning patterns on different FLASH thresholds, relative to conventional IMPT treatment plans.



**Figure 4.2.26:** Sensitivity of the effective FLASH Enhancement Ratio for the best-performing optimized scanning patterns found for 20 independent runs, on different FLASH thresholds. The effective FER corresponds to the ratio between the conventional and the FLASH-enhanced (mean) dose.

## 5 Discussion

This section is split into two subsections, one for the DADR based treatment plan optimization and the other for the PBSDR based scanning-pattern optimization, each divided into 2 subsections encompassing the analytical 2D models and the clinical application.

### 5.1 Dose-Averaged Dose Rate Optimization

In this subsection, the results from the optimized solutions for 2D models and the optimized clinical treatment plans are discussed on the preferred pathway for DADR maximization, its implications on the pencil-beam weights and trade-off on dose. On clinical data, the contributions of beam-current optimization and weight readjustment for the DADR objectives are analysed.

#### 5.1.1 Analytical 2D Models

Solutions with higher dose rate are successfully generated with the proposed optimization strategies, based on the DADR metric, for different ratios of beam separation and beam width  $\Delta\mu/\sigma$ , in different setups.

##### 5.1.1.1 Tumour field

The optimization strategies successfully increase the dose rate of the dose only optimized solutions in a way that maximizes FLASH coverage, as illustrated by figure 4.1.1. There's a trade-off between the dose and dose rate objectives, with priority given to the former, which explains the low FLASH coverage for the lowest ratio. That solution is characterized by a higher number of pencil beams, each contributing less than for higher ratios. Therefore, the minimum weight is lower and the maximum-compatible beam current is low too. In addition, the higher number of pencil beams enables better dose distributions, which are harder to satisfy when optimizing the dose rate, sacrificing the FLASH coverage. Weight readjustment is also shown not to be able to substantially modify the dose rate distribution, so beam-current optimization is essential.

FLASH coverage, is not, however, increased without any trade-off, as figure 4.1.2 shows. For smaller ratios, the beam current has to be substantially increased since FLASH compatibility is low. Therefore, the minimum weight has to increase considerably, which can translate into a significant degradation of the dose distributions. However, these solutions are also more constrained due to the higher quality enabled by their higher degree of freedom, so full FLASH coverage is not always achieved. Plans get less degraded relative to the initial dose-only-optimized solution as the ratio increases, because the required jump in beam current shortens, due to a higher initial minimum weight.

Figure 4.1.3 illustrates that although the relative degradation is smaller for the high ratios, the solution quality is substantially worse. However, if FLASH is considered, the dose sparing might still be advantageous. This requires further validation though, as the sparing observed for higher ratios can also be due to less conformity and less dose to the target structures, perhaps too low to effectively treat the tumour. This shows a need for a controlled trade-off on selecting the pencil beams, so that the required dose conformity is achievable, while leaving enough room for dose-rate optimization.

When relating the global solution quality in figure 4.1.3 with the FLASH coverage in figure 4.1.1, the inevitable trade-offs from DADR optimization are observed. Although FLASH coverage is increased to its fullest whenever possible, it might not be justifiable when coverage is already very high. This demonstrates that FLASH coverage alone is not an accurate indicator of the gains enabled by this effect.



It becomes essential to balance the gains of FLASH with the sacrifices to achieve it in terms of dose, by taking into account both non-FLASH and FLASH-enhanced dose distributions and opting for partial FLASH coverage whenever beneficial. The proposed optimization strategies must be modified, as they currently only aim for FLASH coverage maximization.

Furthermore, it is observed that the trade-off introduced by the DADR optimization can be higher than the trade-off introduced by increasing the ratio  $\Delta\mu/\sigma$ . Because increasing the ratio means less pencil beams and so a higher contribution by each, the minimum weight and consequently the maximum-compatible beam current should be higher, making the plan more FLASH compatible. Therefore, the ratio can also be an important parameter for optimal FLASH compatibility, which can be taken into account in a fully integrated optimization by selecting an appropriate grid from where to sample the pencil beams. This assures beam current is a free parameter, that can be tuned for each treatment plan.

The two previous points on the implementation of the optimization strategies might not be a problem at the clinical treatment planning level. Since for these 2D models the dose constraints are set based on the respective dose-only-optimized solution, the lower ratios have little room to optimize the dose rate due to better dose distributions. If the constraints were the same between all ratios, the pencil-beam removal performed by the beam-current maximization step of the DADR optimization routine would ensure all unnecessary pencil beams were removed, possibly leading to higher FLASH coverage. The ratio would just have to be low enough to be compatible with the constraints and to offer a high degree of freedom. This is what happens in the treatment planning software, since the clinical objectives and goals are defined *a priori*, setting a baseline that every plan has to meet.

In the clinical scenario, the acceptable trade-off is specified by specifying how important preserving the dose is and where it should be prioritized. Although this is also applied for the analytical 2D models, the constraints are not directly tied with the objective function. Dose was optimized by maximizing the homogeneity of the field, but the constraints are set on the mean, maximum and minimum dose at the different structures, chosen to make optimization simpler. If the objective function was instead constrained after being optimized, it would be equivalent to set a minimum acceptable solution quality. Furthermore, in a treatment planning software, partial FLASH coverage can be exploited to some extent, for optimal dose distributions. If structures are optimized one at a time, first in dose and then dose rate, it might be easier to achieve full FLASH coverage at some structure, while prioritizing dose at others.

#### 5.1.1.2 Square field

In order to develop deeper insight on how pencil-beam weights are distributed and assigned to optimize FLASH, dose rates objectives are optimized everywhere for a square field but without beam-current optimization. The spatial pencil-beam weight distributions of figure 4.1.4 are explained by the fact that in the middle, the pencil beams have the highest overlap possible. At the corners it is minimal, since the number of neighbor pencil beams is substantially lower. To achieve a higher dose rate at a certain position, the closest pencil beams have to dominate over all the tails of the contributing Gaussian distributed beams.

To maximize DADR, the relative weight of local pencil beams must be higher over neighbor ones, increasingly higher for more beam overlap. For the lowest ratio, the beams are the closest together so overlap is maximal, explaining the very concentric distribution of weights. As the ratio increases, the pencil beams get further apart and overlap becomes less and less an issue. This explains why weights become more and more uniform, as seen for ratio  $\Delta\mu/\sigma = 1.0$ . However, at some point, pencil beams are so far apart that intermediate regions get irradiated at very low dose rates. This is counterbalanced

by instead increasing the weights, higher at the borders than at the center, as it minimizes possible overlap, which the results for ratio  $\Delta\mu/\sigma = 2.0$  support.

The dose rate distributions demonstrate that the optimization strategies are capable of exploiting pencil-beam density. Since increasing the dose rate substantially at one region would decrease it at nearby regions, the strategies find a middle ground through balancing the distribution everywhere whenever the pencil-beam separation allows. Furthermore, the dose and dose rate distributions for the lowest and intermediate ratios suggest that the dose-rate objectives do not necessarily compete with the dose ones. For the lowest ratio, the uniform dose rate is achieved with a very conformal dose distribution, high at the center and decreasing all around. Similarly, the intermediate ratio allows for near uniform dose and dose rate distributions. Therefore, it might be possible to have full FLASH coverage optimization with little to no sacrifice on dose in some setups, by exploiting beam separation through pencil-beam removal.

### 5.1.2 Clinical Scenario

Application of the DADR optimization methods to clinical patient data is more complex. However, the results in figure 4.1.7 show that the proposed strategies generate plans with improved FLASH coverage. Therefore, dose rate optimization is successfully converted in true direct FLASH optimization, achieved through the selection of the FLASH-compatible voxels to optimize and constrain within each objective's target structure.

This clinical implementation of the DADR objectives should calculate the minimum-required beam current for full FLASH coverage. When focusing on the reported coverage for these current values, 100% coverage is never actually achieved though, which is specially visible in figure 4.1.8. This is surprising considering that the treatment plans are configured to sacrifice the dose to achieve full FLASH coverage and noting that the same plans are compatible with higher beam currents. This is suggested to be an artifact of how the optimization strategies were implemented.

Erasmus-iCycle uses a dose-deposition matrix with a lower resolution when optimizing the different objectives, which encompasses less voxels than the final outputted accurate dose distribution in the patient. This is performed to speed up optimization, since the full accurate dose-deposition matrix would be too large. As a consequence, it is possible that some voxels get excluded from optimization simply because they are not well represented by the low-resolution matrix, falsely reporting a dose lower than the FLASH threshold for selection. In addition, it is also possible that the DADR calculations yield higher values with the simplified dose-deposition matrix, leading to a false early stop. This can be solved by increasing the resolution of the dose-deposition matrix used in optimization, at the expense of a higher running time.

The fact that full FLASH coverage is never obtained for the minimum-required beam current can also be explained by the voxel selection routine. For each objective, the voxels compatible with FLASH on dose are selected in the beginning of the routine and never updated afterwards. If the dose distribution changes, which it inevitably does, the selected voxels are no longer guaranteed to include all the FLASH compatible regions throughout the optimization routine. This is purposely implemented though, since the optimization strategies do not guarantee convergence otherwise. A workaround is to relax the dose threshold, giving a margin for those compatibility changes, although potentially introducing more trade-offs.

#### 5.1.2.1 Comparison with IMPT treatment plans

Figure 4.1.7 shows that some IMPT plans are already compatible with the FLASH effect, illustrating that there is room for FLASH optimization. For stereotactic lung tumours, this is due to the high prescription

dose of 18 Gy (hyperfractionation) and by full delivery of a single direction per fraction, using shoot-through pencil beams. Higher minimum weights are possible, translating into higher beam currents, in some cases sufficiently high to allow near full FLASH coverage. However, the conventional treatment planning approaches cannot guarantee any FLASH compatibility, requiring instead specific optimization strategies.

The fact that IMPT treatment plans have little to no FLASH coverage for direction *B3*, while showing some coverage for the other two directions, is an example of how the FLASH effect can be optimized indirectly, through appropriate choice of settings for each setup. Although direction *B3* is not ideal, choosing a direction *a priori* compatible with higher dose rates is not obvious and requires a manual patient-specific configuration. Nonetheless, the optimization strategies are able to compensate for the smaller compatibility, as full FLASH coverage is always achieved, demonstrating versatility to non-optimal setups. Therefore, no manual patient-specific tuning is required to guarantee optimal FLASH compatibility with the proposed strategies.

The mean dose at the ipsilateral lung in figure 4.1.9 validates that there is considerable room for improvement on the dose rate of IMPT treatment plans with minimum dose trade-off, for stereotactic treatment of lung tumours using transmission beams. In some cases, better non-FLASH dose may be achieved, which is explained by the pencil-beam reduction performed for every DADR objective in the full optimization routine. This is a more aggressive but also a more consistent routine, performed several times throughout the optimization. The conventional pencil-beam reduction routine for IMPT treatment planning is only performed at the end and does not push the solutions to their limit, seen by the fact that only a few plans have weights compatible with beam currents that enable FLASH.

The FLASH-enhanced mean dose at the ipsilateral lung for most IMPT treatment plans is significantly low, as expected from their low FLASH coverage. However, in some cases where the IMPT treatment plan is already fully FLASH compatible, sparing for full optimization is slightly lower, since the FLASH gains might need to compensate for the dose trade-off. Optimization is performed per-structure, first focusing on dose and then on dose rate, therefore it is possible for dose rate optimization in one structure to limit the dose optimization in the others that follow. This balance of FLASH gains and sacrifices on dose is not taken into account by the proposed optimization strategies and the wish-list used. Instead they focus only on maximization of the FLASH coverage.

The healthy lung fraction is critically important on clinical treatments, specially for patients with multiple lesions. However, this is only one of several structures where dose and dose rate is optimized. Sacrifices and gains should be spread throughout the whole patient, so it is possible for one region to compensate another. Nonetheless, the wish-list used prioritizes first the ipsilateral lung before other healthy structures, guaranteeing optimal FLASH coverage with the least dose sacrifice there. If dose is considered more important and no trade-off is desired, optimization should focus first on dose everywhere and only later on dose rate for FLASH coverage.

#### **5.1.2.2 Beam Current Analysis**

The implemented optimization routine calculates the minimum-required and the maximum-compatible beam currents for full FLASH coverage. Figure 4.1.10 shows that the optimization strategies are versatile, by making plans on different setups all compatible with a wide range of beam currents. This is due to the fact that the implemented strategies start with a conservative approach, iteratively increasing the beam current in small steps, before evaluating full compatibility with considerably higher values. The higher median beam current and wider interquartile range for direction *B3* follows from the already discussed small FLASH compatibility for this direction.

Regarding the maximum beam current, compatibility greatly depends on the objective values achieved by the treatment planning software, and it is not necessarily related to the dose sacrifices. In general, because this beam current is tied to the minimum weight, it is expected that plans requiring more pencil-beam modulation will only be compatible with smaller maximum currents. This is because the best dose distributions depend on more pencil beams, each delivering relatively less dose. In the set of patients analyzed, this is observed for patients where the transmission beams shoot through organs at risk (OARs).

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. In addition, this shows that FLASH coverage is not sensitive to beam-current fluctuations. Fixing the beam current within the window of possible values in the treatment planning software guarantees that fully FLASH-compatible plans are generated, with the least trade-off possible. Such result may be critical for clinical applications, because it might not be feasible to change the beam current on a patient basis within the clinical workflow. It might also be difficult to guarantee a constant beam current without fluctuations, at least with current clinical technology.

### 5.1.2.3 Direction Analysis

For the suboptimal beam direction  $B3$ , the shoot-through pencil beams have a longer path in the patients, with a longer distance from the beam entrance to the tumour. This can play a major role because as the beams go through the patient, scattering increases and so does pencil-beam overlap, translating into lower dose rates. This is seen when analyzing the data along the beam direction in figure 4.1.11. For some patients, the density of the structures in the beam path are also critical, because they can also substantially increase scattering, specifically when going through the spinal chord and the thoracic wall.

Despite direction  $B3$  requiring overall higher minimum beam currents for full FLASH compatibility, the maximum-compatible currents are still relatively high for most patients. This suggest that the low FLASH compatibility for this direction is not a limitation imposed by the dose objectives and constraints, but instead a consequence of the spatial setup. Therefore, the problem is just on clever pencil-beam selection and reduction, which the implemented optimization strategies are able to achieve through the more thorough beam-current optimization routine. Furthermore, the relatively wide window of beam currents for all treatment plans shows that there is enough room for the optimizer to overcome the difficulties introduced by non-optimal beam directions, with minimal dose sacrifice.

An appropriate choice of the beam direction is important for optimal FLASH coverage, but not relevant when performing full dose and dose rate optimization, at least for small lung lesions. Nonetheless, the possibility of the required beam current to be so high in other setups that the corresponding minimum pencil-beam weight translates into substantial trade-offs cannot be discarded. Beam direction becomes important to guarantee the best FLASH coverage with the least possible global trade-off, supporting the need for beam direction optimization.

### 5.1.2.4 Component analysis

Although weight readjustment is an important part of DADR optimization, because dose rate varies linearly with the beam current, it is expected for beam-current optimization to be more important for FLASH optimization. Therefore, the more complex nature of the weight readjustment might lead to dose degradation more easily. This is further justified by the wide window of beam currents compatible with the plans and the already FLASH-compatible IMPT plans, which demonstrate that there is enough room to optimize FLASH by only increasing the beam current.

The results for dose-only-optimized treatment plans with maximized beam current, reported in figures 4.1.12 and 4.1.13, show no improvements over the previous fully-optimized plans, with similar dose degradation. This was not expected since the beam current is maximized only after all dose objectives, constrained not to sacrifice any of the objective values obtained. This suggests that the sacrifice on dose by these newly generated plans is allowed by the optimizer, as a consequence of the relaxation step required for the *2pec* algorithm to make room for all objectives.

Event though this same relaxation is performed when generating the conventional IMPT treatment plans, improvements on one dose objective might not substantially degrade the others. The plans never get limited by all constraints and so the reported objective values are better than the relaxed values, at least in this setup with the suggested wish-list. However, when performing beam-current maximization, some trade-off is inevitably obtained because that thorough routine pushes the constraints to the limit. Since the same trade-off is observed for both current and fully-optimized plans, it is proven that the latter guarantees optimal FLASH coverage with minimal dose trade-off. This also suggests that dose-rate optimization at one structure has little to no impact on dose for the surrounding structures, as already discussed for the analytical 2D models.

In some cases, the fully-optimized plans show higher healthy-tissue sparing in the ipsilateral lung than the current-optimized ones, which could follow from the beam-current optimization and weight readjustment at each structure, after dose is optimized. This enables a more careful pencil-beam reduction throughout all optimization. Furthermore, the fact that pencil beams removed for one objective are still considered for optimization of the following objectives translates into a higher degree of freedom to better optimize all objectives. This could also explain the fact that some of current-optimized plans are not fully FLASH compatible but the fully-optimized ones are, despite similarly small dose trade-offs.

#### 5.1.2.5 FLASH Sensitivity Evaluation

The treatment plans optimized for 8 Gy and 40 Gy/s are evaluated for different FLASH dose and dose-rate thresholds, first on the maximum-compatible beam current, in figure 4.1.14 and 4.1.17. These thresholds are critical for the FLASH optimization performed here, so the results were expected to strongly depend on the thresholds specified. However, the maximum beam current for each plan is high enough to guarantee very high dose rates everywhere in the patient, and so the uncertainties on the FLASH dose-rate thresholds are not relevant, at least until 60 Gy/s, the maximum considered here. The decreasing improvements for higher dose thresholds were expected, since for higher thresholds, only smaller fractions of the same volume can benefit from the potential reduced toxicity enabled by FLASH.

For the minimum-required beam current, reported in figures 4.1.16 and 4.1.17, the minimal differences seen for the 30 Gy/s and 50Gy/s thresholds could follow from the fact that the ipsilateral lung is just one of the many optimized structures. In addition, the regions compatible with the FLASH dose threshold are very near the tumour and halfway through the beam path, so the dose rate of each pencil beam should naturally be relatively higher, as scattering is still not pronounced there. Therefore, the beam current required for full FLASH coverage at the ipsilateral lung should be smaller than near the exit point of the transmission beams. Because the optimizer aims for full coverage everywhere, the reported minimum-required current should translate into a dose rate higher than the 40 Gy/s at the region of the ipsilateral lung compatible with the FLASH dose threshold.

Comparing the sensitivity of the results for the two extreme beam currents, it comes that by being compatible with a wide window of beam currents, treatment plans can be made more or less sensitive to FLASH uncertainties *a posteriori*, depending on the delivery settings, without requiring new plans to

be generated. Solutions can be made further consistent though, by specifying lower dose and higher dose-rate thresholds for the FLASH effect.

## 5.2 Pencil-Beam Scanning Dose Rate Optimization

In this subsection, the results for scanning-pattern optimization on both the analytical 2D models and on patient data are discussed, focusing on analysing the convergence and consistency for various settings of the optimizer, comparing the results with predefined patterns on distinct setups, such as various tumour volumes and beam directions.

### 5.2.1 Analytical 2D Models

The proposed scanning-pattern optimizer successfully optimizes the scanning pattern, increasing considerably the FLASH coverage over the predefined patterns, such as the *Z*, *Snake* and *Circle*, for all ratios of beam separation and beam width  $\Delta\mu/\sigma$ .

#### 5.2.1.1 Version Comparison

Despite the stochastic nature of the scanning-pattern optimizer, the results are consistent for 10 independent runs on ratio  $\Delta\mu/\sigma = 1$ , as seen in figure 4.2.1. This is the outcome of the proposed fine-tuning routine, which allows to find a set of parameters that thoroughly configures the optimizer, always guaranteeing high quality. The better solutions and narrower window of possible values for the Island Model shows that by independently running smaller populations and sharing information among them, more regions of the solution space are explored.

If the optimizer is run several times and the best solution out of those runs is selected, a better performance may be obtained, at the expense of computation time, as illustrated in figure 4.2.2. Although the increased consistency of the IMGGA versions guarantees a better performance for a lower number of runs, the trade-off on running times can make it less useful than running the simpler and faster GA more times. Therefore, choosing and fine-tuning a version of the optimizer should balance the consistency of the solutions with the available resources, such as time and computer power.

The performance and the running time greatly depend on the size of the pattern, with optimization for larger patterns taking substantially longer to converge and requiring more runs to guarantee that the best possible solution is found. This is illustrated in figure 4.2.3. Nevertheless, the performance of the algorithm is still consistent, very near the optimum, despite it being tuned only for ratio  $\Delta\mu/\sigma = 1$ , supporting the versatility of the IMGGA and the general effectiveness of the proposed fine-tune routine.

#### 5.2.1.2 Pattern Comparison

Figure 4.2.4 shows that the gains of scanning-pattern optimization are substantially higher for patterns with more overlap. As the pencil-beam separation increases, the pattern becomes less relevant since less pencil beams contribute for the same nearby voxels. The similar FLASH coverage obtained for the different ratios  $\Delta\mu/\sigma$  suggests that this ratio is not a relevant parameter for optimization, despite the different complexities. Therefore, a smaller beam separation can be used with little trade-off on FLASH coverage, but guaranteeing a substantially better dose distribution due to the higher degree of freedom for dose optimization. A balance is needed though, as optimization takes significantly longer for larger patterns, as previously seen in figure 4.2.3. Furthermore, if the ratio is too low, the solution might be incompatible with the beam current initially considered, due to the machine-constrained minimum pencil-beam delivery time.

The optimal FLASH coverage is obtained with snowflake-shaped patterns, as reproduced in figure 4.2.5. 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. Therefore, the center of the pattern should be sacrificed, since there is more overlap due to more neighbor pencil beams, each with a higher dose contribution. In contrast, the outwards regions have less dose and can be fully irradiated quickly. The branches of the snowflakes enable the sacrifice of the interior regions and maximization at the exterior.

The C-shaped PBSDR distributions of figure 4.2.6 are a result of the previous rationale to maximize FLASH coverage. The lower dose rates in the interior come with no sacrifice on FLASH coverage, since that tissue is not considered FLASH compatible. The opening of the circular loop corresponds to the start and end branches of the pattern, which are inevitably sacrificed to close the loop and so to complete the dose delivery. These patterns and distributions show the versatility of the optimizer, demonstrating that it is capable of understanding through the evolutionary process that FLASH coverage can only be maximized if some regions are sacrificed on dose rate.

### 5.2.1.3 Target Comparison

Differences at the objective's target are observed to result in very similar patterns, with the same snowflake shapes, as reproduced in figure 4.2.7. Despite including in optimization the interior region, corresponding to the GTV, the optimized dose rate distributions of figure 4.2.8 show the same sacrifice in the center. Therefore, the snowflake shape is demonstrated not to be a consequence of the geometry of the problem, but instead to represent the best balance of dose rate everywhere. These results also show that the optimizer is compliant with different geometrical setups, different from the one used to configure it. This further supports that fine-tuning the algorithm for simpler 2D models may ensure equally fine-tuned performance in more complex 3D clinical data.

The slightly better results for optimization including the GTV were not expected, because the objective function corresponds exactly to the FLASH coverage when excluding the GTV from optimization, so it gets directly optimized. Including the GTV should steer the focus of the optimizer away from the annulus around the GTV, which is the region where dose rate needs to be optimized to maximize the FLASH coverage. However, excluding the GTV makes the optimizer blind to the central region and consequently, the effect of pencil beams directly delivered in the GTV is less direct. This could make it more difficult for the optimizer to arrange the central pencil beams in a better-performing configuration. This effect is smaller for a reduced number of pencil beams, since the number of pencil beams directly delivered in the GTV is smaller and so the optimizer is less blind. For ratio 2.0, including the GTV is observed to lead to slightly worse results.

Although including the GTV can improve the performance of the algorithm, this approach is not recommended because it is impossible to guarantee that the FLASH coverage will always improve. This is not an issue when optimizing scanning patterns for clinical data though, since transmission beams are used. Even if the GTV or the whole PTV are excluded, there isn't a substantial lack of information like in the 2D case. When looking through the beam direction, there is still healthy tissue compatible with FLASH on dose right in front and behind those structures, meaning that they will also get indirectly optimized. In the 2D case, the optimizer is blind to a central core and focus on the annulus around it, whereas in 3D the optimizer is instead just blind to a sphere inside a relatively larger cylinder.

## 5.2.2 Clinical Scenario

Scanning-pattern optimization on clinical patient data with Genetic Algorithms is successful in FLASH coverage maximization, independent of the tumour size and shape, the number of pencil beams and the beam direction. There is no dose trade-off when optimizing the patterns, which further solidifies the relevance of this optimizer.

### 5.2.2.1 Version Comparison

The consistent results for the three proposed versions of the scanning-pattern optimizer in figure 4.2.10 show that the parameters that fine-tune optimization on the analytical 2D models guarantee equally good results when applied to the more complex 3D clinical scenario. This shows that per treatment plan fine-tuning, even if feasible, is not necessary for consistently high-quality results. Although the IMGGA with 20 islands is observed to always find the best-performing patterns in the evaluated set of 20 runs, the different versions are not necessarily mutually exclusive.

When comparing the running time, figure 4.2.11 shows that consistently guaranteeing the best solution possible is difficult, requiring a steep trade-off on running time. The most suitable version depends on the available resources and should be chosen balancing the solution quality with the computational time. In a clinical setting, because running time is bound to be a limited resource, it might be only possible to run the optimizer once. The Island Model using 20 islands is observed to give the best overall results in that scenario. In the end, this is just a matter of acquiring more and better computers because the problem is parallelizable, as each run is independent and doesn't need to be performed sequentially.

### 5.2.2.2 Convergence Analysis

In general, figure 4.2.12 shows that the convergence rate of the IMGGA is relatively steep in the beginning and almost linear up until very near the optimal solution, at which point it slows down considerably and only improves slightly. This suggests that each island converges at different speeds, so waiting for all of them to converge is time consuming and leads to little improvement. Optimization can then be sped up with little trade-off on the solution quality, by implementing better stopping criteria that takes into account the convergence rate.

Comparing the patients with PTVs of 8.2 cc and 83.9 cc, from the smaller to the larger volume, the number of pencil beams in the worst scenario increases from 31 to 159 pencil beams, a factor of roughly 5. However, convergence only takes a factor of 2 longer, even though the search space increases factorially with the number of pencil beams. The computation time increases substantially though, which might make early stopping critical for clinical pattern optimization of large tumour volumes. Since time is a limited resource in the clinical workflow, these differences in running time suggest that the tumour size can be a fundamental limitation for clinical applications. Nevertheless, the results demonstrate the versatility of the optimizer to different problems and setups, completely different from the ones used to fine-tune it, and the effectiveness of the proposed fine-tuning routine.

### 5.2.2.3 Suboptimal Evaluation

The narrow windows of relative FLASH coverage for the different scanning patterns found by the optimizer, as reported in figure 4.2.13, demonstrate again the consistency of the optimizer and of the fine-tuning routine. The differences between the results for various patients on each directions support the idea that there should be a treatment-plan-specific setup, consisting of every parameter of the optimizer,



that guarantees the best performance in the least amount of time. However, that it is not necessary since the optimized scanning-patterns translate in a FLASH coverage already very near the optimum.

The high median values and narrow interquartile ranges obtained for each direction show that running the optimizer just once is a viable option. Therefore, the running time can be reduced substantially, with very little trade-off on the solution quality, which is still guaranteed to be consistently very near optimum. This too can be helpful and even essential for integration of scanning-pattern optimization in the clinical workflow.

#### 5.2.2.4 Pattern Comparison

The results of figures 4.2.14 and 4.2.15 show that there is always a substantial gain by optimizing the scanning pattern for the beam current value considered here, namely 40 nA. The differences on sparing the ipsilateral lung between the best and the worst-performing patterns on the set of 20 runs for each treatment plant are virtually negligible. This shows that the small differences between the minimum and maximum obtained FLASH coverage translate into even smaller differences on the FLASH-enhanced dose, further supporting the previously-proposed approaches that sacrifice FLASH coverage to speed up optimization at the clinical workflow.

Snowflake-shaped patterns are reproduced in figure 4.2.16, further validating this shape as the best balance between sacrificing and maximizing dose rate. Optimization for 3D clinical data is more complex, since scattering increases along the beam direction and the deeper into the patient, translating into more pencil-beam overlap. Snowflake-shaped patterns are still optimal though, as they are consistently obtained on the simpler 2D models for different levels of beam overlap, given by different ratios of beam width and beam separation. The PBSDR distributions are very similar along the beam direction, with slices normal to this direction having the same previously-discussed C shape.

Despite the improvements on the mean dose, suboptimal results might be obtained from the fact that the optimizer only takes into account a simpler dose-deposition matrix with less resolution, the same used by the treatment planning software. Although this matrix should offer a sufficient representation of all pencil-beam dose contributions, it is not accurate. For the evaluation of the FLASH-enhanced dose, the full accurate dose-deposition matrix is instead used. This could explain the fact that for direction B3, there are optimized patterns with slightly less potential healthy-tissue sparing than for Z patterns.

It should also be noted that optimization is performed taking into account all tissue outside the PTV and compatible with FLASH on dose. Conceptually, it would be more logical to instead optimize FLASH excluding GTV, as the GTV-PTV margin is healthy tissue and is irradiated to a high dose. However, the data available from treatment plan optimization does not naturally allow this, as the dose-deposition matrices used by Erasmus-iCycle for each structure have different sizes, not consistent with the actual size of the structures. In this case, the matrix for the GTV-PTV margin is comprised of as many voxels as the matrix for the whole patient outside the PTV. Therefore, merging the two structures would make FLASH coverage biased to the GTV-PTV margin, steering optimization to focus there.

Optimizing directly on the accurate dose-deposition matrix is an obvious solution but infeasible, since this matrix is substantially larger, making optimization impossible in a reasonable time. Compressing this large accurate matrix in a consistent way for all structures is another possibility. However, this would require a lengthy data preparation routine, since each pencil beam is associated with an accurate dose-deposition matrix, which would take as much as time as the actual optimization. It would also introduce another parameter, namely the compressing ratio, requiring fine-tuning to achieve a balance between accuracy and running time.

The results show that this is not necessary though, at least for small lung lesions, as the current approach still translates into substantially improved FLASH gains. This follows from the fact that healthy tissue exists right in front and right in back of the PTV, meaning that the dose rate will indirectly be optimized at the GTV-PTV margin too. In addition, the size of this margin in comparison with the remaining volume compatible with FLASH on dose is relatively small, accounting for less than 5%. Therefore, although accounting for the GTV-PTV margin is expected to improve the results, the overhead necessary does not compensate for small lung lesions.

For large tumours and low beam currents though, optimization of FLASH coverage may not maximize healthy-tissue sparing. Large tumours account for a larger fraction of the total volume compatible with FLASH on dose, meaning the importance of the PTV-GTV margin for optimization increases. Furthermore, because dose rate decreases in depth, it may never be sufficiently high to become FLASH compatible at the regions deep in the patient, but it may be high enough to enable FLASH near the beam entrance, if optimized. Therefore, it becomes necessary to choose where to increase the dose rate, through prioritized FLASH optimization on different volumes. The discussion about whether or not to include the PTV-GTV margin is a first example of this: it may be critical for large tumours, elongated along the beam direction.

#### 5.2.2.5 Scalability Evaluation

The consistent results for the various tumour volumes in figure 4.2.17 can be explained by the setup used to fine-tune the optimizer, consisting of a relatively large number of pencil beams, namely 88. Even though that setup corresponds to a simple 2D model, the higher number of pencil beams makes the problem more complex. Therefore, fine-tuning should have made the optimizer more compliant with more complex problems, with some trade-off on smaller instances.

Figures 4.2.18 and 4.2.19 show that mean dose improvements with the optimized patterns are substantially higher than those for the predefined patterns. The improvements get higher the larger the tumour sizes, since the predefined patterns sweep the volume between opposite extremes and do not focus on full local irradiation of nearby voxels. These larger tumours might require some patching of patterns to locally increase the dose rate and so FLASH coverage. This is observed for the optimized scanning patterns in figure 4.2.20, with the swirling of the snowflake shape.

The swirl patching of the optimized patterns for the larger tumour volumes is a logical consequence of the higher number of pencil beams. If the normal snowflake shape was applied, branches would have to be considerably longer, so long that the time to reach opposite ends of one branch would no longer 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 swirling formation, despite some unavoidable sacrifice of dose rate in the regions between extremes of different branches. This is observed in the PBSDR distributions of figure 4.2.21. This swirl patching is expected to continue for even larger patterns.

The previous predictions on the PBSDR distribution are confirmed in figure 4.2.21, which illustrates the dose rate distributions along the beam direction, overlapped with the best scanning patterns for the same set of four different tumour volumes. The distributions are very similar to those obtained for the analytical 2D models, which was already expected considering the similar scanning patterns. The only difference are for the larger patterns, where the swirl patching translates into a sacrifice of dose rate at the regions between extremes of different branches. The results follow this same shape along the beam direction, although the dose rate values are observed to decrease the deeper into the patient, a consequence of the increased scattering.

### 5.2.2.6 Direction Analysis

For some patients, the optimized patterns for direction *B3* have a low FLASH coverage but the results are still consistent. Therefore, this should not be an issue of the optimizer but of the beam direction itself. Figure 4.2.22 shows that these underperforming patterns have empty regions in their interior, where there should be some pencil beams. This explains the reduced FLASH coverage because the inner voxels of those empty regions, specially the ones at the center, will always be sacrificed on the dose rate. Those voxels require the irradiation of most of the pencil beams around them, with each contributing only with a small dose, at a small instantaneous dose rate. This would not be a problem for a higher beam currents, since the instantaneous dose rate of every pencil beam would be higher. However, predicting *a priori* a suitable beam current is challenging and changing it for every treatment plan might not be feasible with current clinical technology, within the clinical workflow.

The holed patterns explain the underwhelming results and show that despite everything, the optimizer still actively tries to push the FLASH coverage as high as possible. It suggests that pencil-beam reduction might hurt FLASH compatibility. It might be possible to obtain identical patterns but with the holes filled with the missing pencil beams and by splitting the contribution of the other nearby beams among them, which should then have higher FLASH compatibility. For the underperforming treatment plans, it is seen that direction *B3* is, compared with the other directions, associated with larger paths from the entrance location to the tumour and even through all the patient.

When reversing the direction *B3*, it is observed that whenever FLASH coverage increases in figures 4.2.23 and 4.2.24, the distance from the beam entrance to the tumour is reduced. Figure 4.2.22 shows that the newly optimized patterns do not have empty regions in their inside. Holed patterns should be preferred for larger distances because of lateral scattering. For deep tumours, there should be already noticeable lateral scattering of the pencil beams. Therefore, to keep the high dose conformal and very localized there, less but heavier pencil beams are required with empty spaces around them. This guarantees that the dose is still high at the middle but the inevitable dose contribution from scattering at the healthy region is much lower. When the tumour is located near the beam entrance, scattering is still not pronounced, so it is easier to keep the dose conformal and localized by using more and lighter pencil beams.

### 5.2.2.7 FLASH Sensitivity Evaluation

Figures 4.2.25 and 4.2.26 show that for the beam current used, the dose rate threshold is very important and greatly determines the possible healthy tissue sparing in the ipsilateral lung. This demonstrates that the optimizer focuses only on getting the dose rate above the specified threshold, which leaves more room for optimization at other regions, but not on increasing the dose rate as much as possible, which would make results less sensitive.

From a dose perspective, the results are still consistent, which should follow from the symmetric dose rate distributions obtained with the optimized patterns. Increasing the dose threshold correspond to shrinks the compatible volume, so the distribution is still very similar, translating in minimal differences in FLASH improvements for small lung lesions. In this setup, this should correspond to shrinking the radius of the cylinder that goes through the tumour and is normal to the beam direction. When too close to the GTV, FLASH gains are very small because dose rate is sacrificed there by the snowflake-shaped patterns.

Considering that the beam current is tied to the dose-rate threshold, there is some room to make solutions less sensitive to this threshold *a posteriori* by adjusting the current accordingly. This ultimately depends on the switching time between pencil beams, which is not here taken into account. For relatively low

currents, the pencil-beam delivery time is much higher than the switching time and there's approximately a linear relation between dose rate and beam current. Therefore, the optimized patterns are the best performing for all configurations that preserve that proportion: the optimized patterns for a current of 40 nA and dose rate threshold of 40 Gy/s would also be the best performing for 60 nA and 60 Gy/s.

As the beam current is increased though, the pencil-beam delivery time decreases and consequently the switching time becomes more important. However, higher currents make FLASH coverage less pattern-dependent, so the switching time is not expected to substantially change the results. In this application to small lung lesions, considering that the median minimum delivery time for a beam current of 40 nA is 1.8 ms and the switching time is usually around 0.2 ms, not taking into account this parameter is not expected to significantly change the best-performing patterns.

## 6 Conclusion

A significant increase of dose rate may be achieved with the proposed optimization strategies for stereotactic proton therapy of lung lesions using transmission beams, through higher beam currents and optimized snowflake-shaped scanning patterns. The optimized treatment plans have a quality in terms of dose very similar or even identical to conventional IMPT treatment plans. Nonetheless, dose, dose rate and trade-offs between them are constrained by a set of *a priori* parameters, namely the beam current, the minimum delivery time, the beam direction and the FLASH model.

With current and near-future proton therapy, combined with the proposed optimization approaches, a significant enhancement in dose rates is feasible. Based on current knowledge of FLASH, this may be of significant clinical benefit for stereotactic treatment of selected patients with lung lesions. Even though solutions are observed not to critically depend on the setup, the optimization strategies need first to be thoroughly tested before clinical application. To ensure 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, specially on the dose rate. The effect of beam current fluctuations, the scanning and the switching time between pencil beams should be also evaluated. Beam-direction optimization is also important, manually performed by minimizing distances and avoiding dense structures.

FLASH optimization could go one step further, weighting the potential 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 for scanning-pattern optimization, to substantially reduce the running time of the optimizer. Furthermore, integrated optimization of the treatment plan and the scanning pattern on both metrics could yield better and more consistent results. Hybrid approaches could also improve results further and open the door to application in other tumour sites, considering multi-beam setups for which only a fraction of radiation dose is delivered under strict FLASH conditions, for example using a mix of Bragg peaks and shoot-through transmission beams.

## References

- [1] Sophie Pilleron et al. “Global cancer incidence in older adults, 2012 and 2035: A population-based study”. In: *International Journal of Cancer* 144.1 (2019), pp. 49–58. DOI: [10.1002/ijc.31664](https://doi.org/10.1002/ijc.31664).
- [2] Claudia Santucci et al. “Progress in cancer mortality, incidence, and survival: a global overview”. In: *European Journal of Cancer Prevention* 29.5 (2020), pp. 367–381. ISSN: 0959-8278. DOI: [10.1097/CEJ.0000000000000594](https://doi.org/10.1097/CEJ.0000000000000594).
- [3] Yolande Lievens, Josep M. Borrás, and Cai Grau. “Provision and use of radiotherapy in Europe”. In: *Molecular Oncology* 14.7 (2020), pp. 1461–1469. DOI: [10.1002/1878-0261.12690](https://doi.org/10.1002/1878-0261.12690).
- [4] PTCOG - Particle Therapy Co-Operative Group. Accessed: 28 December 2020. URL: <https://www.ptcog.ch/>.
- [5] Melody Ju, Abigail T. Berman, and Neha Vapiwala. “The Evolution of Proton Beam Therapy: Insights From Early Trials and Tribulations”. In: *International Journal of Radiation Oncology\*Biophysics* 90.4 (2014), pp. 733–735. ISSN: 0360-3016. DOI: <https://doi.org/10.1016/j.ijrobp.2014.07.035>.
- [6] F. G. Spear and Frederick William Andrewes. “The delayed lethal effect of radium on tissue cultures *in vitro*; Comparison of continuous and spaced radiation”. In: *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character* 108.756 (1931), pp. 190–195. DOI: [10.1098/rspb.1931.0031](https://doi.org/10.1098/rspb.1931.0031).
- [7] SHIRLEY HORNSEY and TIKVAH ALPER. “Unexpected Dose-rate Effect in the Killing of Mice by Radiation”. In: *Nature* 210.5032 (Apr. 1966), pp. 212–213. DOI: [10.1038/210212a0](https://doi.org/10.1038/210212a0).
- [8] 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 (July 2014), 245ra93–245ra93. DOI: [10.1126/scitranslmed.3008973](https://doi.org/10.1126/scitranslmed.3008973).
- [9] 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. ISSN: 0167-8140. DOI: [10.1016/j.radonc.2019.02.009](https://doi.org/10.1016/j.radonc.2019.02.009).
- [10] 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. ISSN: 0167-8140. DOI: [10.1016/j.radonc.2019.06.019](https://doi.org/10.1016/j.radonc.2019.06.019).
- [11] H. Weiss et al. “Oxygen Depletion in Cells Irradiated at Ultra-high Dose-rates and at Conventional Dose-rates”. In: *International Journal of Radiation Biology and Related Studies in Physics, Chemistry and Medicine* 26.1 (1974), pp. 17–29. DOI: [10.1080/09553007414550901](https://doi.org/10.1080/09553007414550901).
- [12] Guillem Pratx and Daniel S Kapp. “A computational model of radiolytic oxygen depletion during FLASH irradiation and its effect on the oxygen enhancement ratio”. In: *Physics in Medicine & Biology* 64.18 (2019), p. 185005. DOI: [10.1088/1361-6560/ab3769](https://doi.org/10.1088/1361-6560/ab3769).
- [13] Joseph D. Wilson et al. “Ultra-High Dose Rate (FLASH) Radiotherapy: Silver Bullet or Fool’s Gold?” In: *Frontiers in Oncology* 9 (2020), p. 1563. ISSN: 2234-943X. DOI: [10.3389/fonc.2019.01563](https://doi.org/10.3389/fonc.2019.01563).
- [14] Marie-Catherine Vozenin et al. “The Advantage of FLASH Radiotherapy Confirmed in Mini-pig and Cat-cancer Patients”. In: *Clinical Cancer Research* 25.1 (2019), pp. 35–42. ISSN: 1078-0432. DOI: [10.1158/1078-0432.CCR-17-3375](https://doi.org/10.1158/1078-0432.CCR-17-3375).
- [15] Eric S. Diffenderfer et al. “Design, Implementation, and *In Vivo* Validation of a Novel Proton FLASH Radiation Therapy System”. In: *International Journal of Radiation Oncology\*Biophysics* 106.2 (2020), pp. 440–448. ISSN: 0360-3016. DOI: [10.1016/j.ijrobp.2019.10.049](https://doi.org/10.1016/j.ijrobp.2019.10.049).

- [16] Joseph D. Wilson et al. “Ultra-High Dose Rate (FLASH) Radiotherapy: Silver Bullet or Fool’s Gold?” In: *Frontiers in Oncology* 9 (2020), p. 1563. ISSN: 2234-943X. DOI: [10.3389/fonc.2019.01563](https://doi.org/10.3389/fonc.2019.01563).
- [17] Simon Jolly et al. “Technical challenges for FLASH proton therapy”. In: *Physica Medica* 78 (2020), pp. 71–82. ISSN: 1120-1797. DOI: <https://doi.org/10.1016/j.ejmp.2020.08.005>.
- [18] Shannon Cunningham et al. “FLASH Proton Pencil Beam Scanning Irradiation Minimizes Radiation-Induced Leg Contracture and Skin Toxicity in Mice”. In: *Cancers* 13.5 (2021). ISSN: 2072-6694. DOI: [10.3390/cancers13051012](https://doi.org/10.3390/cancers13051012).
- [19] Benjamin Mou et al. “Feasibility of Proton Transmission-Beam Stereotactic Ablative Radiotherapy versus Photon Stereotactic Ablative Radiotherapy for Lung Tumors: A Dosimetric and Feasibility Study”. In: 9.6 (June 2014), e98621. DOI: [10.1371/journal.pone.0098621](https://doi.org/10.1371/journal.pone.0098621).
- [20] 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. DOI: [10.1080/0284186X.2019.1627416](https://doi.org/10.1080/0284186X.2019.1627416).
- [21] Michael M. Folkerts et al. “A framework for defining FLASH dose rate for pencil beam scanning”. In: *Medical Physics* n/a.n/a (2020). DOI: [10.1002/mp.14456](https://doi.org/10.1002/mp.14456).
- [22] Patricia van Marlen et al. “Bringing FLASH to the Clinic: Treatment Planning Considerations for Ultrahigh Dose-Rate Proton Beams”. In: *International Journal of Radiation Oncology\*Biophysics* 106.3 (2020), pp. 621–629. ISSN: 0360-3016. DOI: [10.1016/j.ijrobp.2019.11.011](https://doi.org/10.1016/j.ijrobp.2019.11.011).
- [23] Patricia van Marlen et al. “Ultra-High Dose Rate Transmission Beam Proton Therapy for Conventionally Fractionated Head and Neck Cancer: Treatment Planning and Dose Rate Distributions”. In: *Cancers* 13.8 (2021), p. 1859. ISSN: 2072-6694. DOI: [10.3390/cancers13081859](https://doi.org/10.3390/cancers13081859).
- [24] JA Groen. “FLASH optimisation in clinical IMPT treatment planning: optimisation and evaluation of stereotactic lung treatment plans with proton transmission beams”. In: *Thesis* (2020). URL: <https://repository.tudelft.nl//islandora/object/uuid:23b4e83e-87f8-41ac-8b05-6298aba2548e>.
- [25] Daniel Tsang. “Investigating the influence of pencil beam scanning on the dose rate and the optimal pencil beam scanning patterns”. In: *Internship Report* (2021).
- [26] Hao Gao et al. “Simultaneous dose and dose rate optimization (SDDRO) for FLASH proton therapy”. In: *Medical Physics* n/a.n/a (2020). DOI: [10.1002/mp.14531](https://doi.org/10.1002/mp.14531).
- [27] Yuting Lin et al. “SDDRO-joint: simultaneous dose and dose rate optimization with the joint use of transmission beams and Bragg peaks for FLASH proton therapy”. In: *Physics in Medicine & Biology* 66.12 (2021), p. 125011. DOI: [10.1088/1361-6560/ac02d8](https://doi.org/10.1088/1361-6560/ac02d8).
- [28] 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. DOI: [10.1118/1.3676689](https://doi.org/10.1118/1.3676689).
- [29] Sebastiaan Breedveld et al. “A novel approach to multi-criteria inverse planning for IMRT”. In: *Physics in Medicine and Biology* 52.20 (2007), pp. 6339–6353. DOI: [10.1088/0031-9155/52/20/016](https://doi.org/10.1088/0031-9155/52/20/016).
- [30] Sebastiaan Breedveld, Pascal R M Storchi, and Ben J M Heijmen. “The equivalence of multi-criteria methods for radiotherapy plan optimization”. In: *Physics in Medicine and Biology* 54.23 (2009), pp. 7199–7209. DOI: [10.1088/0031-9155/54/23/011](https://doi.org/10.1088/0031-9155/54/23/011).
- [31] Siegfried Schaible. “Fractional programming: Applications and algorithms”. In: *European Journal of Operational Research* 7.2 (1981). Fourth EURO III Special Issue, pp. 111–120. ISSN: 0377-2217. DOI: [https://doi.org/10.1016/0377-2217\(81\)90272-1](https://doi.org/10.1016/0377-2217(81)90272-1).
- [32] J. R. Isbell and W. H. Marlow. “Attrition games”. In: *Naval Research Logistics Quarterly* 3.1-2 (1956), pp. 71–94. DOI: <https://doi.org/10.1002/nav.3800030108>.

- [33] JP Crouzeix, JA Ferland, and S Schaible. "An algorithm for generalized fractional programs". In: *Journal of Optimization Theory and Applications* 47.1 (1985), pp. 35–49.
- [34] James E. Falk and Susan W. Palocsay. "Optimizing the Sum of Linear Fractional Functions". In: *Recent Advances in Global Optimization*. Princeton University Press, Dec. 1991, pp. 221–258. DOI: [10.1515/9781400862528.221](https://doi.org/10.1515/9781400862528.221).
- [35] Mirjam. Dur, Reiner. Horst, and Nguyen. Van Thoai. "Solving sum-of-ratios fractional programs using efficient points". In: *Optimization* 49.5-6 (Jan. 2001), pp. 447–466. DOI: [10.1080/02331930108844543](https://doi.org/10.1080/02331930108844543).
- [36] Chun-Feng Wang and Pei-Ping Shen. "A global optimization algorithm for linear fractional programming". In: *Applied Mathematics and Computation* 204.1 (Oct. 2008), pp. 281–287. DOI: [10.1016/j.amc.2008.06.045](https://doi.org/10.1016/j.amc.2008.06.045).
- [37] Hongwei Jiao et al. "An Effective Algorithm for Globally Solving Sum of Linear Ratios Problems". In: *Journal of Control Science and Engineering* 2017 (2017), pp. 1–7. DOI: [10.1155/2017/8138975](https://doi.org/10.1155/2017/8138975).
- [38] E. Valipour, M. Yaghoobi, and M. Mashinchi. "An iterative approach to solve multiobjective linear fractional programming problems". In: *Applied Mathematical Modelling* 38 (2014), pp. 38–49.
- [39] Benjamin M Clasié, Hanne M Kooy, and Jacob B Flanz. "PBS machine interlocks using EWMA". In: *Physics in Medicine and Biology* 61.1 (2015), pp. 400–412. DOI: [10.1088/0031-9155/61/1/400](https://doi.org/10.1088/0031-9155/61/1/400).
- [40] 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.
- [41] Alex S Fraser. "Simulation of genetic systems by automatic digital computers ii. effects of linkage on rates of advance under selection". In: *Australian Journal of Biological Sciences* 10.4 (1957), pp. 492–500.
- [42] 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. DOI: [10.1023/a:1006529012972](https://doi.org/10.1023/a:1006529012972).
- [43] Melanie Mitchell. *An Introduction to Genetic Algorithms*. Cambridge, MA, USA: MIT Press, 1998. ISBN: 0262631857.
- [44] Enrique Alba and Marco Tomassini. "Parallelism and evolutionary algorithms". In: *IEEE transactions on evolutionary computation* 6.5 (2002), pp. 443–462.
- [45] Kenneth Alan De Jong. *An analysis of the behavior of a class of genetic adaptive systems*. University of Michigan, 1975.
- [46] Darrell Whitley, Soraya Rana, and Robert B Heckendorn. "The island model genetic algorithm: On separability, population size and convergence". In: *Journal of computing and information technology* 7.1 (1999), pp. 33–47.
- [47] David E. Goldberg and Robert Lingle. "Alleles, Loci and the Traveling Salesman Problem". In: *Proceedings of the 1st International Conference on Genetic Algorithms*. USA: L. Erlbaum Associates Inc., 1985, pp. 154–159. ISBN: 0805804269.
- [48] I. M. Oliver, D. J. Smith, and J. R. C. Holland. "A Study of Permutation Crossover Operators on the Traveling Salesman Problem". In: *Proceedings of the Second International Conference on Genetic Algorithms on Genetic Algorithms and Their Application*. Cambridge, Massachusetts, USA: L. Erlbaum Associates Inc., 1987, pp. 224–230. ISBN: 0805801588.
- [49] D. B. Fogel. "An evolutionary approach to the traveling salesman problem". In: *Biological Cybernetics* 60.2 (1988), pp. 139–144. DOI: [10.1007/bf00202901](https://doi.org/10.1007/bf00202901).
- [50] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. Cambridge, MA, USA: MIT Press, 1992. ISBN: 0262082136.



- [51] Holger Mauch. "Closest Substring Problem – Results from an Evolutionary Algorithm". In: *Neural Information Processing*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 205–211. ISBN: 978-3-540-30499-9.
- [52] B. L. Miller and D. Goldberg. "Genetic Algorithms, Tournament Selection, and the Effects of Noise". In: *Complex Syst.* 9 (1995).
- [53] Werner Dinkelbach. "On Nonlinear Fractional Programming". In: *Management Science* 13.7 (1967), pp. 492–498. ISSN: 00251909, 15265501.

## 7 Appendix

### 7.1 Proof of DADR optimization routines

#### 7.1.1 Simple Linear Fractional Programming

In order to understand how the routine 3.1.1 converges to a solution, it is important to split analysis in a few steps [53]. First, it should be proved that  $q_{k+1} > q_k$  for all  $k$  with  $\mathcal{F}(q_k) \geq \delta$ . For that, it should be noted that  $\mathcal{F}(q_k) > 0$ , which is implied by the following observation, with  $x^+ \in X$  and  $q^+ = N(x^+)/D(x^+)$ :

$$\mathcal{F}(q^+) = \max\{N(x) - q^+D(x) \mid x \in X\} \geq N(x^+) - q^+D(x^+) = 0 \quad (7.1.1)$$

By definition, it then follows:

$$\begin{aligned} N(x_k) &= q_{k+1}D(x_k), \quad \text{hence} \quad \mathcal{F}(q_k) = N(x_k) - q_kD(x_k) \\ &= q_{k+1}D(x_k) - q_kD(x_k) > 0 \end{aligned} \quad (7.1.2)$$

Since  $D(x_k) > 0$ , it then automatically follows that  $q_{k+1} > q_k$ .

The second step is to verify that the routine converges to the optimal solution  $\lim_{k \rightarrow \infty} q_k = q^*$ . If this is not true, it must come that  $\lim_{k \rightarrow \infty} q_k = q^+ < q^*$  and by how the algorithm was constructed, that there is a sequence  $x_k^+$  with  $q_k^+$ , such that  $\lim_{k \rightarrow \infty} \mathcal{F}(q_k^+) = \mathcal{F}(q^+) = 0$ . Now, it should be noted that  $\mathcal{F}(q)$  is strictly monotonic decreasing, which can be proven by considering  $q' < q''$  and  $x''$  to maximise  $\mathcal{F}(q'')$ :

$$\begin{aligned} \mathcal{F}(q'') &= \max\{N(x) - q''D(x) \mid x \in X\} = N(x'') - q''D(x'') \\ &< N(x'') - q'D(x'') \\ &\leq \max\{N(x) - q'D(x) \mid x \in X\} = \mathcal{F}(q') \end{aligned} \quad (7.1.3)$$

Since  $\mathcal{F}(q)$  is strictly monotonic decreasing and  $q^+ < q^*$ , it comes that:

$$0 = \mathcal{F}(q^+) > \mathcal{F}(q^*) = 0 \quad (7.1.4)$$

Which is a contradiction. Hence, it follows that  $\lim_{k \rightarrow \infty} \mathcal{F}(q_k) = \mathcal{F}(q^*)$  and because  $\mathcal{F}(q)$  is continuous in  $X$ , as it is simply the maximization of a linear function, it is concluded that the iterative linearization converges to the optimal solution. However, it is not guaranteed that the algorithm will converge in a finite number of steps, neither that in case it does, the algorithm is fast enough to converge in a feasible amount of time. Because of this, the parameter  $\delta$  is introduced, which works as a tolerance, enabling the routine to be stopped earlier according to how precise the minimum solution is desired to be.

#### 7.1.2 Maximization of the minimum DADR

The proof of convergence of algorithm 3.1.2 is similar to what was shown for the simpler version. First, it should be proved that  $q_{k+1} > q_k$  for all  $k$  with  $\mathcal{F}(q_k) \geq \delta$ . Similarly, it should be noted that  $\mathcal{F}(q_k) > 0$ , which is implied by the following observation, with  $x^+ \in X$  and  $q^+ = \min_{1 \leq i \leq M} [N_i(x^+)/D_i(x^+)]$ :

$$\mathcal{F}(q^+) = \max\left\{ \min_{1 \leq i \leq M} [N_i(x) - q^+D_i(x)] \mid x \in X \right\} \geq \min_{1 \leq i \leq M} [N_i(x^+) - q^+D_i(x^+)] = 0 \quad (7.1.5)$$

Denoting  $J(x) = \{j \mid N_j(x)/D_j(x) = \min_{1 \leq i \leq M} N_i(x)/D_i(x)\}$  and considering that  $j \in J(x_k)$ , it follows:

$$\begin{aligned} q_{k+1} &= \min_{1 \leq i \leq M} \frac{N_i(x_k)}{D_i(x_k)} = \frac{N_j(x_k)}{D_j(x_k)} \quad \text{hence} \quad \mathcal{F}(q_k) = \min_{1 \leq i \leq M} [N_i(x_k) - q_k D_i(x_k)] \\ &\leq N_j(x_k) - q_k D_j(x_k) \\ &= D_j(x_k)(q_{k+1} - q_k) \end{aligned} \quad (7.1.6)$$

Rewriting the previous inequality and considering that  $\mathcal{F}(q_k) > 0$  and  $D_i(x) > 0$ , it follows:

$$q_{k+1} - q_k \geq \frac{\mathcal{F}(q_k)}{D_j(x_k)} > 0 \quad (7.1.7)$$

Hence, it is concluded that  $q_{k+1} > q_k$ .

Finally, it is just necessary to show that the routine converges to the optimal solution. First, the previous inequality should be rewritten, by introducing  $\bar{D}(x) = \max_{1 \leq i \leq M} D_i(x)$  and considering  $q^*$  the solution to the problem:

$$q_{k+1} - q_k \geq \frac{\mathcal{F}(q_k)}{\bar{D}(x_k)} \Leftrightarrow q^* - q_{k+1} \leq q^* - q_k - \frac{\mathcal{F}(q_k)}{\bar{D}(x_k)} \quad (7.1.8)$$

Because the ratios are all linear and the search space  $X$  is a bounded polyhedron space, it is guaranteed to exist a solution  $x^+$  for every  $\mathcal{F}(q^+)$  if  $q^+ \leq q^*$ . It is then possible to write for all  $i$ :

$$\mathcal{F}(q^+) \leq N_i(x^+) - q^+ D_i(x^+) = -\mu D_i(x^+) + N_i(x^+) - (q^+ - \mu) D_i(x^+) \quad (7.1.9)$$

Assuming  $\mu < q^+$  and introducing  $\underline{D}(x) = \min_{1 \leq i \leq M} D_i(x)$ , it follows:

$$\begin{aligned} \mathcal{F}(q^+) + (q^+ - \mu) \underline{D}(x^+) &\leq \mathcal{F}(q^+) + (q^+ - \mu) D_i(x^+) \leq N_i(x^+) - \mu D_i(x^+) \quad \text{for all } i = 1, \dots, M \\ &\leq \min_{1 \leq i \leq M} [N_i(x^+) - \mu D_i(x^+)] \leq \mathcal{F}(\mu) \end{aligned} \quad (7.1.10)$$

Considering that  $q_k < q^*$ ,  $\mathcal{F}(q^*) = 0$  and according to the previous relation, it is possible to write:

$$\mathcal{F}(q_k) \geq \mathcal{F}(q^*) + (q^* - q_k) \underline{D}(x^*) = (q^* - q_k) \underline{D}(x^*) \quad (7.1.11)$$

And consequently plugging the previous relation in inequality 7.1.8:

$$q^* - q_{k+1} \leq (q^* - q_k) \left(1 - \frac{\underline{D}(x^*)}{\bar{D}(x_k)}\right) \quad (7.1.12)$$

Since  $D_i(x)$  are all linear positive and the search space  $X$  is a bounded polyhedron space, it is guaranteed that  $\sup_k \bar{D}(x_k) < \infty$ . In addition, because  $q_k \leq q^*$ , it follows that  $0 < \underline{D}(x^*)/\bar{D}(x_k) \leq 1$  and consequently that  $q_k$  converges to  $q^*$ . Similar to the simpler case, a parameter  $\delta$  is introduced to allow early stopping of the algorithm when the solution is considered already good enough.

### 7.1.3 Maximization of the mean DADR

Similarly to the previous routines, it is necessary to prove that algorithm 3.1.3 also converges and to show that the weighted sum of the linearized objectives is a solution to the initial multiobjective problem. It is then necessary to show first that  $x^* \in X$  is an efficient solution to the multiobjective problem, if and only if  $\mathcal{F}(q^*) = 0$ , with  $q^* = [f_1(x^*), \dots, f_M(x^*)]$ . A solution  $x' \in X$  is efficient if there is no other  $x \in X$  such that  $f_i(x) \geq f_i(x') + \delta_i$  for all  $i = 1, \dots, M$  and  $f_j(x) > f_j(x') + \delta_j$  for at least one  $j$  or, in other words,  $x'$  is a solution that could only be improved on some objectives by sacrificing others. First, considering  $x^+ \in X^+$  a feasible solution of  $\mathcal{F}(q^+)$ :

$$\begin{aligned} f(x^+) = \left[ \frac{N_1(x^+)}{D_1(x^+)}, \dots, \frac{N_M(x^+)}{D_M(x^+)} \right] \geq q^+ &\Rightarrow N_i(x^+) - q_i^+ D_i(x^+) \geq 0 \quad \text{for } i = 1, \dots, M \\ &\Leftrightarrow w_i(N_i(x^+) - q_i^+ D_i(x^+)) \geq 0 \\ &\Rightarrow \sum_{i=1}^M w_i(N_i(x^+) - q_i^+ D_i(x^+)) = \mathcal{F}(q^+) \geq 0 \end{aligned} \quad (7.1.13)$$

However, if  $x^*$  is an efficient solution and the optimal value  $\mathcal{F}(q^*)$  is assumed nonzero, it comes from the previous relation that  $\exists x^+ \in X$  such that  $f(x^+) \geq q^* = f(x^*)$  and  $\sum_{i=1}^M w_i[N_i(x^+) - q_i^* D_i(x^+)] > 0$ . Consequently, there must exist at least one  $j$  such that

$$w_j(N_j(x^+) - q_j^* D_j(x^+)) > 0 \quad \Leftrightarrow \quad N_j(x^+) - q_j^* D_j(x^+) > 0 \quad \Leftrightarrow \quad f_j(x^+) > q_j^* = f_j(x^*) \quad (7.1.14)$$

So it is concluded that  $\exists j$  such that  $f_j(x^+) > f_j(x^*)$  but at the same time that  $f_i(x^+) \geq q_i^* = f_i(x^*)$  for all  $i = 1, \dots, M$ , which is in contradiction with the efficiency of the solution  $x^*$ . Therefore, if the solution is efficient, then  $\mathcal{F}(q^*) = 0$ .

Conversely, considering that  $\mathcal{F}(q^*) = 0$  but that  $x^* \in X$  is not an efficient solution, by definition there must exist a feasible solution  $x^+ \in X$  to the multiobjective problem such that  $f_i(x^+) \geq f_i(x^*)$  for all  $i = 1, \dots, M$  and  $f_j(x^+) > f_j(x^*) = q_j^*$  for some  $j$ , meaning that  $x^+$  is also a feasible solution of  $\mathcal{F}(q^*)$ . Consequently, since each term of  $\mathcal{F}(q^*)$  is non-negative, it comes for the same  $j$ :

$$\mathcal{F}(q^*) = \sum_{i=1}^M w_i[N_i(x^+) - q_i^* D_i(x^+)] \geq w_j(N_j(x^+) - q_j^* D_j(x^+)) > 0 \quad (7.1.15)$$

A contradiction is found, meaning that if  $\mathcal{F}(q^*) = 0$  then  $x^* \in X$  has to be an efficient solution. This concludes the proof of the if and only if relation in both directions for the solution of the multiobjective problem.

Next, it should be noted that the search space at each iteration is a subset of the previous search space  $X_{k+1} \subset X_k$ . For that, consider  $x_k$  the solution for  $\mathcal{F}(q_k)$  at iteration  $k$ :

$$\begin{aligned} x_k \in X_k &\Rightarrow f_i(x_k) \geq q_k^i = f_i(x_{k-1}) \quad \text{for } i = 1, \dots, M \\ &\Rightarrow \{x \in X : f_i(x) \geq q_{k+1}^i = f_i(x_k) \geq q_k^i = f_i(x_{k-1})\} \subset \{x \in X : f_i(x) \geq q_k^i = f_i(x_{k-1})\} \\ &\Rightarrow X_{k+1} \subset X_k \end{aligned} \quad (7.1.16)$$

It can be also shown that each sequence  $\{q_k^i = f_i(x_{k-1})\}$  is nondecreasing and convergent. This is because  $\forall k, x_{k+1} \in X^{k+1} \Rightarrow f_i(x_{k+1}) \geq f_i(x_k)$  for  $i = 1, \dots, M$ . This then implies that  $\{q_k^i =$

$f_i(x_{k-1})\}$  are nondecreasing series. Since  $f(x)$  is a bounded set, a consequence of the bounded search space  $X$  and the continuous objective functions, it comes that each  $\{q_k^i\}$  is a bounded sequence and consequently convergent too.

Finally, it is just left showing that the routine does converge to the solution, which can be proved by showing that the sequence of  $\{\mathcal{F}(q_k)\}$  converges to zero, since that only happens if and only if  $x_k$  is an efficient solution to the problem, with  $q_{k+1} = f(x_k)$ :

$$\begin{aligned}
N_i(x_k) = q_{k+1}^i D_i(x_k), \quad \text{hence} \quad \lim_{k \rightarrow \infty} \mathcal{F}(q_k) &= \lim_{k \rightarrow \infty} \sum_{i=1}^M w_i (N_i(x_k) - q_k^i D_i(x_k)) \\
&= \lim_{k \rightarrow \infty} \sum_{i=1}^M w_i (q_{k+1}^i D_i(x_k) - q_k^i D_i(x_k)) \\
&= \lim_{k \rightarrow \infty} \sum_{i=1}^M w_i (q_{k+1}^i - q_k^i) D_i(x_k) \quad (7.1.17)
\end{aligned}$$

It should be first noted that each  $\{D_i(x_k)\}$  is a bounded positive sequence, since  $D_i(x)$  is bounded positive on  $X$ . This plus the fact that each sequence  $\{q_k^i\}$  is nondecreasing convergent, as it was shown previously, allows to conclude that  $\lim_{k \rightarrow \infty} \mathcal{F}(q_k) = 0$ , so that the routine converges to a solution  $x^* \in X$  of the multiobjective linear fractional program. Nonetheless, a parameter  $\delta$  is also introduced to stop the algorithm earlier, when a good enough solution has already been found.

Structure	Priority	Min/Max	Type	Goal	Sufficient	Parameters
GTV	Constraint	Maximize (minimum)	linear	A		
GTV	2	Minimize (maximum)	linear	A*1.24		
PTV	Constraint	Maximize (minimum)	linear	A*0.98		
PTV without GTV	Constraint	Maximize (minimum)	linear	A*0.98		
PTV without GTV	1	Minimize (maximum)	linear	A*1.28		
Shell around PTV 3mm	3	Minimize (maximum)	linear	A*0.62		
Shell around PTV 3mm	4	Minimize (maximum)	mean	A*0.62		
Shell around PTV 3mm	5	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Shell around PTV 3mm	6	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Shell around PTV 6mm	7	Minimize (maximum)	linear	A*0.31		
Shell around PTV 6mm	8	Minimize (maximum)	mean	A*0.31		
Shell around PTV 6mm	9	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Shell around PTV 6mm	10	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Shell around PTV 9mm	11	Minimize (maximum)	linear	0		
Shell around PTV 9mm	12	Minimize (maximum)	mean	0		
Shell around PTV 9mm	13	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Shell around PTV 9mm	14	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Shell around PTV 20mm	15	Minimize (maximum)	linear	0		
Shell around PTV 20mm	16	Minimize (maximum)	linear	0		
Shell around PTV 20mm	17	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Shell around PTV 20mm	18	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Outside PTV	19	Minimize (maximum)	mean	0		
Outside PTV	20	Minimize (maximum)	linear	0		
Lung ipsilateral without PTV	21	Minimize (maximum)	linear	0		
Lung ipsilateral without PTV	22	Minimize (maximum)	mean	0		
Lung ipsilateral without PTV	23	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Lung ipsilateral without PTV	24	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Long other side	25	Minimize (maximum)	linear	0		
Long other side	26	Minimize (maximum)	mean	0		
Long other side	27	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Long other side	28	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Spinal chord	29	Minimize (maximum)	linear	0		
Spinal chord	30	Minimize (maximum)	mean	0		
Spinal chord	31	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Spinal chord	32	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Esophagus	33	Minimize (maximum)	linear	0		
Esophagus	34	Minimize (maximum)	mean	0		
Esophagus	35	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Esophagus	36	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Trachea	37	Minimize (maximum)	linear	0		
Trachea	38	Minimize (maximum)	mean	0		
Trachea	39	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Trachea	40	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Bronchus ipsilateral	41	Minimize (maximum)	linear	0		
Bronchus ipsilateral	42	Minimize (maximum)	mean	0		
Bronchus ipsilateral	43	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Bronchus ipsilateral	44	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
Outside PTV	45	Maximize (minimum)	DADR	B	B	C   D   $\delta_1$
Outside PTV	46	Maximize (minimum)	mean DADR	B	B	C   D   $\delta_2$
MU	47	Minimize (maximum)	linear	0		

**Table 7.1.1:** Treatment Planning wish-list used for all the Erasmus-iCycle optimizations performed in this study: *A* corresponds to the prescribed dose 18 Gy, *B* is the dose-rate goal, set to the FLASH dose rate threshold 40 Gy/s, *C* and *D* are the FLASH dose and the dose-rate thresholds, so set to 8 Gy and 40 Gy/s respectively,  $\delta_1$  and  $\delta_2$  are tolerances on the DADR optimization routines for the maximization of the minimum and mean, set to 0.1 and 100 respectively.