



Low Energy Linacs for 3D X-ray Scanning Applications

Samuel Smith, BEng (Hons), M.S

School of Engineering

The University of Lancaster

A thesis submitted for the degree of

Doctor of Philosophy

July, 2024

Low Energy Linacs for 3D X-ray Scanning Applications

Samuel Smith, BEng (Hons), M.S.

School of Engineering, Lancaster University

A thesis submitted for the degree of *Doctor of Philosophy*. July, 2024.

Abstract

Typical cargo scanning Linacs used to scan shipping containers are usually designed at S-band (2-4 GHz) frequencies and have energies of 3-6 MeV in order to obtain sufficient contrast during inspections. To scan smaller and thinner containers and to reduce the footprint of scanning systems, there is an interest in lower energy (1-3 MeV) and higher frequency (5-6 GHz) devices. In this thesis, a design of such a device is presented, with a compact (11.9 cm), five cell, 2 MeV, bi-periodic, C-band linac selected as the final design. Multi-objective optimisation methods and spline modelling techniques are used to optimise the cells with a Pareto analysis used to select a final design, also allowing for rapid design adjustment. A novel coupling method using nose cone slants is developed, giving an improvement of 28 % in the coupling factor between cells (0.7 % - 0.9 %), with minimal effect on the peak fields or shunt impedance (1.3 %). MO methods are also employed to optimise the cell lengths and RF amplitudes to increase the capture efficiency of the linac to over 90 % using re-capture methods. A complete thermal analysis is presented showing that the linac can handle up to 1.2 kW average power with less than 2 % error in the electric field. The study includes CFD simulations and an improved method for estimating the heat transfer coefficient by including bends when performing calculations which agrees with the CFD analysis. The design is then integrated into a full RF system that allows for three linacs to be fired and synchronised, using three frequencies with a 3 MHz separation (5.712 ± 3 MHz) on three sections of one 10 μ s RF pulse. It is then shown that this system is capable of generating quasi-3D images in a CT-like setup using 3D image reconstruction techniques.

Declaration

I declare that the work presented in this thesis is, to the best of my knowledge and belief, original and my own work. The material has not been submitted, either in whole or in part, for a degree at this, or any other university. This thesis does not exceed the maximum permitted word length of 80,000 words including appendices and footnotes, but excluding the bibliography.

Samuel Smith

Acknowledgements

This thesis would not have been possible without the guidance and support of many individuals who have contributed in numerous ways to my research, professional development, and personal growth.

First and foremost, I would like to express my deepest gratitude to my supervisor, Professor Graeme Burt, for his steadfast support, invaluable guidance, and insightful feedback throughout this process. His expertise, encouragement, and kindness have been instrumental in shaping this research and my development over the past decade. I also extend my sincere thanks to Dr Amos Dexter for his excellent supervision during the initial stages of my studies. Additionally, a big thank you to Dr Sadiq Setiniyaz for his assistance in all matters relating to low-energy electron beams. A special mention goes to my fellow PhD student and friend, Matthew Southerby, with whom I have shared the joys and challenges of this journey in parallel. Our collaborations and discussions have been enriching and memorable. I would also like to thank Dr Alick Macpherson and the rest of Team Rocket whose support has been greatly appreciated during the final months of completing this thesis, alongside starting a new position.

To the team at Rapiscan — James, Jasmin, Dave, and Dean — I thoroughly enjoyed my short time with you. Thank you for hosting me and for sharing your knowledge and expertise. I extend particular thanks to Dr Mike Jenkins for his encouragement, for organising the internship at Rapiscan, and for providing invaluable feedback and guidance as the research and design progressed.

To my parents and family, I thank you for your unconditional love and unwavering support, regardless of the path I chose.

Finally, to my partner and best friend, Rae, who has been there every step of the way. Thank you for walking this challenging road with me.

To everyone mentioned—and to all the wonderful friends and colleagues I have made along the way, who have shaped this process in ways big and small—thank you for being a part of this journey. Your support, kindness, and presence mean more than words can express.

*I sit beside the fire and think
of all that I have seen,
of meadow-flowers and butterflies
in summers that have been;*

*Of yellow leaves and gossamer
in autumns that there were,
with morning mist and silver sun
and wind upon my hair.*

*I sit beside the fire and think
of how the world will be
when winter comes without a spring
that I shall ever see.*

*For still there are so many things
that I have never seen:
in every wood in every spring
there is a different green.*

*I sit beside the fire and think
of people long ago
and people who will see a world
that I shall never know.*

*But all the while I sit and think
of times there were before,
I listen for returning feet
and voices at the door.*

J.R.R. Tolkien

Nomenclature

β	Relativistic beta
η	Dynamic viscosity
Γ	Reflection coefficient
λ	Wavelength
μ_a	Attenuation coefficient
ω	Angular frequency
B_{pk}	Peak magnetic flux density
c	Speed of light
C_p	Specific heat capacity at constant pressure
E_z	Axial electric field
E_{acc}	Accelerating gradient
E_{pk}	Peak electric field
eV	Electron volts
f	Frequency
I_0	X-ray intensity
k_c	Coupling constant

P_d	Dissipated power
Q	Quality factor
Q_0	Intrinsic quality factor
Q_e	External quality factor
Q_L	Loaded quality factor
Q_t	Transmitted quality factor
R_e	Reynold's number
R_s	Shunt impedance
R_{surf}	Surface resistance
S_c	Modified Poynting vector
T_m	Melting temperature
t_p	Pulse length
U	Stored energy
V_0	Axial voltage
V_c	Accelerating voltage
Z	Shunt impedance per unit length
Z_0	Impedance
Z_a	Atomic number
BDR	Breakdown rate
bpp	Breakdowns per pulse

CFD Computational fluid dynamics

CT Computed tomography

CVI Cargo and vehicle inspection

dB Decibel

DUT Device under test

EM Electromagnetic

FBP Filtered Back Projection

FBP Filtered back projection

FDK Feldkamp-Davis-Kress

FDM Frequency-Division Multiplexing

FOM Figure of merit

HEP High-energy physics

HFSS High-frequency structure simulator

HTC Heat transfer coefficient

KE Kinetic energy

LLRF Low level radio frequency

MLEM Maximum Likelihood Expectation Maximisation

MO Multi-objective optimisation

mT Milli-tesla

NDT Non-destructive testing

NURBS Non-uniform rational basis splines

OFHC Oxygen-free high thermal conductivity

PFN Pulse-forming network

PIC Particle-in-cell

RF Radio frequency

RRT Real-time tomography

SART Simultaneous Algebraic Reconstruction Technique

SS Steady-state

STL Stereolithography

SW Standing wave

TE Transverse electric

TIG Tungsten inert gas

TM Transverse magnetic

TTF Transit time factor

TW Traveling wave

ULD Unit load device

WG Waveguide

Contents

1	Industrial Accelerators and X-ray Imaging Techniques	1
1.1	Radio-Frequency Linacs	1
1.2	X-rays for Security	3
1.3	Cargo Scanning Linacs	4
1.3.1	Shipping Containers and Heavy Cargo	5
1.3.2	Types of CVI Systems	7
1.4	Cargo Scanning Linac System Components	7
1.5	Generating Images	10
1.5.1	Collimators	10
1.5.2	Detectors	10
1.5.3	Movement	11
1.5.4	Material Discrimination	13
1.5.5	CT Applications and Image Reconstruction	14
1.6	Design Criteria	15
1.6.1	Selection of RF Frequency	16
1.6.2	Commercial Low-Energy Linacs	17
1.7	Thesis Outline	19
2	RF Acceleration and Theory	20
2.1	RF Acceleration	20
2.2	Figures of Merit	22
2.2.1	Accelerating Gradient	23

2.2.2	Peak Fields	23
2.2.3	Shunt Impedance and Cavity Trade-offs	24
2.2.4	Field Limits	26
2.2.4.1	Electric Field	27
2.2.4.2	Breakdown Rate Scaling	27
2.2.4.3	Magnetic Field and Pulsed Heating	28
2.2.5	Cell-Cell Coupling	30
2.2.6	Group Velocity	32
2.3	Tolerances	33
3	Single Cell Optimisation	36
3.1	Design Criteria	36
3.1.1	Standing vs. Travelling Wave Cavities	37
3.1.2	Mode Selection	39
3.1.3	Frequency Perturbations	42
3.1.4	Initial Design Ideas	45
3.2	Single Cell Design	46
3.2.1	Multi-Objective Optimisation	47
3.2.2	Algorithms	50
3.2.3	Optimization Results	52
3.3	Algorithm Studies	55
3.3.1	Frequency Handling	55
3.3.2	Convergence	57
3.3.3	Geometric Comparison	58
3.4	Cell-Cell Coupling	59
3.4.1	Coupling Improvement	61
3.4.2	Effect of Aperture Radius and Wall Thickness	65
3.5	Final Cell Selection	68
3.6	Power of Pareto Design	70

4	Beam Capture Optimisation and Linac Tuning	73
4.1	Longitudinal Dynamics	73
4.2	Phase Slippage	76
4.3	Electron Bunching	77
4.3.1	Maximum Capture	77
4.3.2	Tapered Buncher	79
4.4	Cell Length Optimisation	80
4.4.1	MO Optimisation	83
4.4.2	First Cell Design	87
4.4.3	Final Design Selection	89
4.5	Coupled Periodic Standing Wave Cavities	90
4.5.1	Lossy Bi-Periodic Equivalent Circuit	93
4.5.2	Amplitude Tuning	96
4.6	Coupler Design	98
4.7	Full Cavity Design	103
4.7.1	Tuned Model Comparison	103
4.8	Mode Tuning	107
4.8.1	Beam Capture Tolerances	108
4.9	Particle-in-Cell Design Verification	112
5	X-ray Target Design	116
5.1	Target Theory	116
5.1.1	X-ray Production Processes	117
5.2	G4beamline Simulation Setup	119
5.3	X-ray Target Optimisation	122
5.4	Copper Target	123
5.5	X-ray Target Heating	125
5.6	Energy Deposition Model	128
5.6.1	Optimal Target Thickness	130
5.7	Final Design	133

6	Linac Thermal and Mechanical Analysis	135
6.1	Basic Heat Transfer Theory	135
6.1.1	Conduction	135
6.1.2	Convection	136
6.1.3	Radiation	137
6.2	Thermal Effects During Manufacture	137
6.2.1	Ansys Results	141
6.2.2	Heating During Operation	144
6.3	Single Cell Frequency Analysis	146
6.4	Full Linac Thermal Analysis	147
6.4.1	Forced Convection	149
6.4.1.1	Mean values	149
6.4.1.2	Boundary Layers	150
6.4.2	Heat and Cooling for Linacs	150
6.4.3	HFSS Results	154
6.5	Steady State Thermal Analysis	155
6.6	SS Simulation Setup	156
6.7	Thermal Deformation	158
6.8	Frequency and Field Errors	160
6.8.1	Frequency Errors	162
6.9	CFD Analysis	163
6.9.1	Simulation Setup	164
6.9.2	Meshing	164
6.9.3	Results	165
6.10	Validity of HTC Approximations	167
6.10.1	Heat Transfer in Curved Geometries	169
6.11	Operational Considerations	174
6.11.1	Frequency Adjustment	175
6.11.2	Power Limits	176

7	RF System Design	178
7.1	Multi-View Linac System	178
7.1.1	Initial RF System Considerations	178
7.1.2	Delay Lines	180
7.2	Frequency-Division Multiplexing	183
7.3	Frequency Switching	183
7.4	RF Pulse	185
7.5	Simulation Verification	186
7.5.1	Signal Generation	187
7.5.2	Validation	193
7.5.3	Power Requirements	197
7.6	Experimental Verification	197
7.6.1	Results	198
7.6.2	Validating Experimental Results	201
8	X-Ray Imaging and 3D Reconstruction Techniques	204
8.1	Image Reconstruction	204
8.2	2D Theory	205
8.2.1	Radon Transforms	206
8.2.2	Filtered Back Projection (FBP)	206
8.2.3	Number of Projections Required	207
8.3	Algorithms	208
8.3.1	Simultaneous Algebraic Reconstruction Technique	208
8.3.2	Maximum Likelihood Expectation Maximisation	209
8.4	Methods	209
8.5	Reconstruction Simulations	211
8.6	Threat Reconstruction	213
8.7	Density Variations	215
8.8	Distance Variations	216
8.9	Further Work	219

9	Conclusions	222
9.1	RF Design and Optimisation	224
9.2	Note on Multi-objective Optimisation	227
9.3	Beam Capture Optimisation	228
9.4	Linac Thermal Analyses	229
9.5	RF System Design	232
9.6	Quasi-3D Image Generation using Linacs	232
9.7	Future Work	233
9.8	Final Remarks	234
	References	236

List of Tables

1.1	Types of cargo scanning systems.	7
1.2	Penetration of steel and water vs. electron energy. [39]	15
1.3	C-band industrial linac examples.	16
1.4	Frequency scaling of cavity parameters.	17
2.1	RF single cell objectives.	26
2.2	Estimate of the frequency error caused by a 30 μm error on the radius.	35
3.1	Initial design specifications.	36
3.2	Initial gradient estimates.	46
3.3	Typical GA parameters used for optimisation.	51
3.4	Initial design parameters.	52
3.5	Effect of wall thickness on various FOMs and k_c	68
3.6	Final optimised $\beta = 1$ cell parameters.	70
4.1	Cell length optimisation objectives.	81
4.2	Astra MO input parameters and initial limits.	84
4.3	Matlab gamultiobj parameters used for optimisation of cell lengths.	84
4.4	Frequency change for outer radius change for the three first cell models.	89
4.5	Final design parameters selected from beam capture optimisation.	90
4.6	Single vs. tuned field comparison for $E = 45.13$ MV/m.	105

4.7	Tuned field 1D Astra simulation vs. full 3D PIC results for $E = 45.13$ MV/m.	114
5.1	G4beamline gaussian beam parameters.	120
5.2	Various characteristic X-ray peaks for a tungsten target.	122
5.3	Parameters used for composite target heating study.	126
6.1	Effect of parameters on deformation.	143
6.2	Initial analytical calculations of thermal properties.	154
6.3	Comparison of frequency shifts for three different fixed displacement options.	163
7.1	Initial delay line considerations.	180
7.2	EMMA cavity parameters used to compare simulations and experi- mental results.	203
9.1	Final $\beta = 1$ cell parameters.	225
9.2	Comparison with typical X-band linacs. X-band 1 [38], X-band 2 [199], X-band 3 [200]. Parameters are given as an average for the TW structure.	225

List of Figures

1.1	Plot showing reduction in successful hijackings after 1973. Image sourced from [15].	4
1.2	First Philip’s Saferay system, with manual placement of the object to be scanned. Image sourced from [15].	5
1.3	Example of an X-ray taken by the Philip’s Saferay system. Image sourced from [15].	5
1.4	Table showing the comparisons between X-ray vacuum tubes, radioactive sources and linac systems. Table sourced from [20].	6
1.5	Simple schematic of a CVI system.	8
1.6	Example of a CVI collimator. Image sourced from [9]	11
1.7	Example of a curved detector array. From left to right, pixel, module, curved array. Image sourced from [23]	12
1.8	(Left) Example of a gantry system, Rapiscan EAGLE G60, Image sourced from [25] (Right) Example of a portal system, Rapiscan EAGLE P60, image sourced from [24].	12
1.9	Mass attenuation vs. energy for different materials. Image sourced from [33].	14
1.10	Example image of a truck using color to represent material groups based on atomic number. Image sourced from [33]	14
1.11	Dual energy linac scanning example. Image sourced from [61]	18
1.12	Cargo scanning linac Examples. Image sources from [62]	18
2.1	Gap between cells (left), Electric field inside this gap (right).	21
2.2	Transit time factor as a function of gap length.	22
2.3	Simple single cell cavity vacuum space (left), cross-section (middle), copper outer walls (right).	23
2.4	Simple single cell cavity showing electric (left) and magnetic (right) fields in the volume.	24
2.5	Example of fields in a cavity cell, showing where the maximum occurs and the effect of adding nose cones. Fields are normalised to 1 J stored energy.	25
2.6	Equivalent unit cell circuit for electric coupling through an iris.	30
2.7	Plot of equation 2.28.	32

3.1	General dispersion curve for waveguide loaded with irises.	37
3.2	(Left) Travelling wave cavity filling vs (Right) Standing wave cavity filling. Image sourced from [81].	38
3.3	Dispersion curve showing 0, $\pi/2$ and π modes.	41
3.4	Amplitudes of π and $\pi/2$ modes.	41
3.5	Bi-periodic vs side coupled linacs. a.) Bi-periodic structure with no adjustments of coupling cell, b.) Reduced length of coupling cell to increase Z, c.) Coupling cells moved off axis, again increasing space for acceleration, d.) Example of coupling slots that couple the cavities together using the magnetic field, e.) Example of side coupled structure [63].	44
3.6	Linac example showing shielding to protect from radiation and the formula for calculating the area of the shielding shown in grey, highlighting the dependence on transverse size. A is the area, r is maximum radius of the cell and t is required shielding thickness. . . .	45
3.7	Optimum beam aperture vs shunt impedance used for C-band aperture selection, where shaded region represents the aperture radii where C-band performs best. Image sourced from [85].	46
3.8	NURBS spline geometry showing control points used to optimise cell.	47
3.9	Pareto optimal points in 2D. [87]	48
3.10	Weighted sum vs. multi-objective methods. [87]	49
3.11	Generic GA flow chart.	50
3.12	2D Pareto front results for $\beta = 1$ (left), top to bottom (Z vs. E_{pk}/E_{acc} , S_c/E_{acc} vs. E_{pk}/E_{acc} , B_{pk}/E_{acc} vs. E_{pk}/E_{acc} and $\beta = 0.5$ (right), top to bottom (Z vs. E_{pk}/E_{acc} , S_c/E_{acc} vs. E_{pk}/E_{acc} , B_{pk}/E_{acc} vs. E_{pk}/E_{acc}	53
3.13	Z vs. E_{pk}/E_{acc} for three values of β	54
3.14	Pareto optimal points in 3D.	54
3.15	Pareto front of Z vs. E_{pk}/E_{acc} , with the outer colour of each point showing the value of $\sqrt{S_c}/E_{acc}$ and inner colour the value of B_{pk}/E_{acc} . A cross section of the types of solutions that fall on the Pareto front are also shown, highlighting the interesting shapes at lower values of E_{pk}/E_{acc}	55
3.16	Frequency variance in combined Pareto optimal set.	56
3.17	Variation in 2D Pareto front for varying ϵ (defined in equation 3.20).	57
3.18	Convergence of the Z vs E_{pk}/E_a Pareto front.	58
3.19	Comparison of Z vs. E_{pk}/E_{acc} Pareto front for s NURBS geometry and standard shape model.	59
3.20	Schematic representation of cell-cell coupling calculation showing drive cell.	60
3.21	Coupling vs structure length for a 2 degree phase shift along the structure, a 1 % amplitude variation and various values of $d\omega/\omega_a$, which is the ratio of the frequency difference between coupling cells to the accelerating mode frequency.	61

3.22	Required coupling vs. acceptable tolerance for different values of the acceptable phase shift across the structure.	62
3.23	Two options for increasing the coupling between the coupling cells and accelerating cells. (Left) Chamfering (Right) Blending.	62
3.24	From left to right, blend option, chamfer option, baseline cavity for scans with $k_c = 0.771$ %.	63
3.25	Effect of changing the blend radius on coupling constant and shunt impedance.	63
3.26	Z and coupling constant vs. chamfer angle for three different chamfer lengths. (Top) 1 mm (Middle) 2 mm (Bottom) 2.7 mm.	64
3.27	Effectiveness of coupling proxy.	65
3.28	Z vs. E_{pk}/E_{acc} , S_c/E_{acc} Pareto front showing the coupling proxy value and the estimated limit for E_{pk}/E_{acc}	66
3.29	S_c/E_{acc} vs. E_{pk}/E_{acc} solutions with coupling proxy shown in colour. Pareto fronts are shown for binned values of Z. Potential final solution candidates circled in red.	66
3.30	Cells used for aperture radius investigation (top), plots showing the effect of the aperture radius on k_c (bottom left) Z and (bottom right).	67
3.31	Effect of wall thickness on the Z vs. E_{pk}/E_{acc} Pareto fronts (left), cavity candidate used for investigating the effect of wall thickness on k_c (right).	68
3.32	Final $\beta = 1$ cavity design showing fields from left to right (Cross-section, Electric field, modified Poynting vector, magnetic field). Normalised to 1J stored energy.	69
3.33	Surface fields along curve on final cell design.	69
3.34	Effect of adding chamfer on magnitude of electric field distribution. Original geometry (Left), after adding 2.8 mm chamfer with 25° angle (Right).	70
3.35	k_c vs. E_{pk}/E_{acc} showing Z in colour.	71
3.36	Coupling constant vs. linac energy for various cavity gradients. The two shaded regions represent the accessible coupling regions for different values of the beam pipe aperture radius, A.	72
3.37	Cross section of five $\beta = 1$ cells that satisfy the various E_{pk}/E_{acc} limits shown in Fig. 3.36.	72
4.1	Energy gain vs. Phase of a stable particle and particles arriving early or late.	74
4.2	Setup for accelerating cells to describe longitudinal motion. Image sourced from [63].	74
4.3	Relativistic β vs. energy for electrons.	75
4.4	$\phi_{initial}$ vs ϕ_{final} for $\beta_{inj} = 0.3$, $\beta_{final} = 1$ and $E_{acc} = 30$ MV/m.	76
4.5	Estimate of how the phase velocity, $\beta(z)$ (left) and amplitude, A(z) (right), should vary along the linac in a tapered velocity buncher.	80

4.6	(a) $\beta = 1$ 1D field (b) $\beta = 1$ electric field contour (c) $\beta = 0.5$ 1D field (d) $\beta = 0.5$ electric field contour.	82
4.7	(a) Scanning of the maximum electric field amplitude. (b) Calculated required power to obtain this field.	82
4.8	Scatter plots (last generation) for cell length optimization, showing how the average kinetic energy, % electrons returning to cathode and length/amplitude of the first cell interact.	85
4.9	Matrix for cell length Optimization, showing Pareto points from last generation at 40 MV/m and 45 MV/m.	86
4.10	First cell optimization results. Length vs. E_{pk}/E_{acc} (left), Length vs. B_{pk}/E_{acc} (right).	87
4.11	Progression in the design of the low amplitude first cell.	88
4.12	First cell optimization results showing a reduction in field length whilst maintaining the same value of β on the outer cell.	89
4.13	Second cell design.	90
4.14	Two numerical solutions	91
4.15	Chain of capacitively coupled resonators with nearest neighbour and next nearest neighbour coupling included.	92
4.16	(Top), Bi-periodic cavity chain showing nearest neighbour and next nearest neighbour coupling. (Bottom), 4 cell model with end cells and driving term V_d	94
4.17	Current amplitude in all 9 cells for a perfectly tuned bi-periodic capacitively coupled cavity.	95
4.18	Current amplitude in all 9 cells after tuning the coupling between the first few cells.	96
4.19	First three cells (top) and field amplitude of the $\pi/2$ mode for various values of the chamfer angle between the third and second cell (bottom).	97
4.20	First cell parameters used for tuning.	97
4.21	First cell length for different values of cg_{min} which increases the nose cone, pushing the mode field further up the cell. This shows how the length and amplitude can be tuned.	98
4.22	Effect of aperture radius and blend radius parameters on length and amplitude of the field in the first cell.	99
4.23	Circuit model showing waveguide and cavity connected by a coupler.	99
4.24	Taper to get from cell width to WR159 waveguide.	101
4.25	S_{11} vs. taper length for the taper design.	101
4.26	Coupler design with taper. Vacuum model (left), electric field through cell (centre), magnetic field strength through cell (right).	102
4.27	Peak magnetic field on coupling slot.	103
4.28	Final linac design. (a) Z-component of the electric field (b) Vacuum model (c) Absolute value of the electric field with copper walls.	104
4.29	Scan of β_c of a single cell model to find the critical coupling, leading $\beta_c = 1$ (critical coupling) for the full cavity.	105

4.30	Comparison of stitched 1D fields used for optimisation and final tuned fields.	106
4.32	Model used to initially tune coupling cell $\pi/2$ mode.	107
4.33	S11 vs. frequency and location of modes from frequency domain solver, before and after tuning.	108
4.34	Plot showing a comparison of the 9 modes in the structure with the analytical dispersion for a 9-cell structure.	109
4.35	Effect of changing the amplitude of the first cell on they key outputs of interest.	110
4.36	Effect of changing the length of the first cell on they key outputs of interest.	111
4.37	short	111
4.38	Scan of field on axis for a $\pm 5\%$ error on the first cell amplitude.	112
4.39	Imported fields from tuned CST model in PIC solver. (Top) Electric field, (Bottom) magnetic field.	113
4.40	Particle position monitor snapshot showing bunch formation and exiting bunch at 2 MeV.	114
4.41	Particle energy vs. time recorded at the end of the linac showing the bunches leaving the structure at 2 MeV.	115
4.42	Histogram of the beam at the end of the linac in the PIC simulation.	115
5.1	X-ray conversion efficiency for different energies and materials [107].	117
5.2	Description of the bremsstrahlung process. Image sourced from [108].	118
5.3	Typical X-ray spectrum showing characteristic X-rays. [109].	118
5.4	Histogram of the final beam size at the end of the linac.	119
5.5	Simple diagram of target simulations showing composite two layer target, incoming beam and X-ray detection area.	120
5.6	Images of simulation from G4beamline, red lines represent electron trajectories and green lines represent photon trajectories.	121
5.7	Low energy section of X-ray spectrum of a tungsten target showing characteristic X-rays.	121
5.8	Example of spectra for a tungsten (W) thickness scan with a 1 mm copper (Cu) backing. Np/Max represents the number of photons normalised to the maximum number from the scan.	123
5.10	Results for a copper only target. (Top) X-ray yield vs. copper thickness (Bottom) Electron leakage vs. copper thickness	125
5.11	Comparison of spectra from a 1.75 mm copper target and a copper-tungsten composite target.	126
5.12	Geometry of initial target heating simulations showing tungsten layer, copper layer, and tungsten section where the heat is generated.	127
5.13	Transient heating applied to element of tungsten to represent heating by the beam. (Top) A single pulse, (Bottom) multiple pulses at the 200 Hz repetition rate.	127

5.14	Heat distribution within the target. From top to bottom, start of the first pulse, after the decay of the first pulse, start of the last pulse, after the decay of the last pulse.	129
5.15	Maximum temperature in the target vs. time for two different boundary temperatures.	130
5.16	(a) Longitudinal energy deposition in the target. (b) Area independent energy deposition radially in target. (c) Longitudinal energy deposition in copper backing. (d) Longitudinal slices used to measure energy deposition. (e) Concentric segments used to measure radial energy deposition.	131
5.17	Gaussian heating model used for a more accurate assessment of target heating.	132
5.18	Gaussian heating thermal distributions. (Top) start of the last pulse, (Bottom) after the decay of the last pulse.	132
5.19	Comparison of target heating for a single cylinder and the gaussian model.	133
5.20	Comparison of maximum target temperature vs time for three copper backing thicknesses, using a gaussian deposition model.	133
5.21	Comparison of final X-ray spectrum 1 m away from the target for a Gaussian beam and the beam generated from the tuned CST Astra model.	134
6.1	Schematic of basic brazing technique.	138
6.2	Typical temperature vs. time brazing curve [120].	139
6.3	High temperature diffusion mechanisms [121].	140
6.4	Low pressure, high temperature diffusion mechanisms. Image sourced from [124].	140
6.5	Diffusion coefficients (left) and strain rates (right).	141
6.6	Piece-wise fit to strain rate.	142
6.7	Single cell thermal deformation model (left), Parameters in model (centre) deformation after heating (right).	142
6.8	Deformation vs. time Ansys result, maximum (green), minimum (red) and average (blue) deformation in the vertical direction.	143
6.9	Temperature dependent material properties of copper obtained from [124, 129]. Top left to bottom right: thermal conductivity, coefficient of thermal expansion, specific heat capacity, Young's modulus, density, pressure.	144
6.10	Heating during operation (single cell). Left to right: 3 mm wall thickness, 2 mm wall thickness, 1 mm wall thickness, cooling pipe arrangement.	146
6.11	Nose cone temperature rise and shunt impedance vs. wall thickness for the single cell geometry.	146

6.12	Steps required for single cell frequency shift analysis. (Top left) Single cell model in HFSS. (Top right) Imported heat flux from HFSS into the thermal solver. (Bottom left) Single cell thermal deformation. (Bottom right) Temperature distribution from thermal solver.	148
6.13	(Left) Displaced electric field after thermal deformation. (Right) Displacement due to deformation.	148
6.14	Boundary layer formation at the start of a pipe, also showing the variation of the HTC in this region. Image sourced from [113].	151
6.15	(Left) Comparison of S_{11} from the CST and HFSS models. (Right) Comparison of the the absolute value of E_z on axis for the CST and HFSS models.	155
6.16	Plot of the absolute value of the HFSS electric field in the x-y plane.	155
6.17	Simulation process flow for analysing the thermal deformation of RF cavities showing two options for obtaining the thermal distribution.	156
6.18	RF surface losses calculated in HFSS.	157
6.19	Setup of steady state thermal simulation showing imported flux faces, convection to the water pipes and air, and the inlet of the pipes.	158
6.20	Mapped and scaled RF surface losses from HFSS into the SS solver, maximum values are seen on the slots connecting to the coupler.	158
6.21	Steady state thermal temperature distribution.	159
6.22	Temperature rise along thermal elements in SS solver.	159
6.23	Total deformation for three different boundary cases. (Left, 800x magnification), Fixed at ends, (Centre, 500x magnification), Fixed top and bottom, (Right, 800x magnification), Fixed at all four locations.	160
6.24	Von Mises stress for the three boundary cases. Total deformation for three different boundary cases. (Left) Fixed at ends (Centre) Fixed top and bottom (Right) Fixed at all four locations.	160
6.25	Deformed mesh in HFSS, used to re-simulate structure.	161
6.26	Absolute value of the electric field after simulating the deformed mesh in HFSS.	161
6.27	Electric field error on axis after mechanical deformation.	162
6.28	Electric field error on axis after mechanical deformation for varying average power levels in the linac.	162
6.29	Absolute error on the electric field on axis after mechanical deformation.	163
6.30	CFD simulation setup in Ansys fluent. (Left) Inner structure for surface loss mapping and cooling pipes showing inlet and outlet. (Right) Outer wall and cooling pipe walls.	166
6.31	Ansys fluent mesh. (Left) Slice through the x-y plane at the start of the structure showing the fine mesh around the boundary layer of the pipes. (Right) Mesh on outer wall.	166
6.32	Water properties and fits used in Ansys fluent.	167
6.33	Slice of the y-z plane showing the resulting temperature distribution obtained using Ansys fluent.	167

6.34	Slice through the structure close to the inlet with two temperature scales. (Left) Slice through the structure showing the temperatures in the water pipes. (Right) Slice through the structure showing the temperature in the copper.	168
6.35	Slice through pipe outlet showing the boundary layer temperature rise in the water.	168
6.36	Comparison of Nu_c for curved and straight pipes, image from [149]. .	170
6.37	Surface Nusselt number, showing variation in the corners of the pipes compared with the straight sections.	170
6.38	HTC values along the pipes. (Left) Outer pipes visible (Right) Slice through all the pipes showing the values on the inner surface. . . .	171
6.39	HTC values at all of the nodes along the pipes with a moving average of 26 periods shown in red.	171
6.40	(Left) Water velocity along the pipes showing variation around the corners. (Right) Water temperature showing the difference between the SS linear temperature increase approximation and the real temperature distribution in the water.	171
6.41	Temperature on external wall of the linac with the area used for the average temperature measurement shown in red.	172
6.42	Temperature on outer wall vs value of the HTC for various different simulation configurations.	173
6.43	Comparison of the maximum temperature on the iris for different simulation setups.	174
6.44	Frequency vs. inlet water temperature showing the change required to bring the linac back on frequency.	175
6.45	Frequency shift vs. average power after reducing the inlet temperature to the value required to bring the linac back on resonance, down to 20 °C.	176
6.46	Maximum average input power to the structure before the outer linac wall reaches 60 °C vs. the water flow rate.	177
6.47	Temperature on outer shell vs. water inlet temperature for various average input powers and a flow rate of 5 L/min.	177
7.1	Initial design concept using delay lines and circulators to allow for the use of only one RF source.	179
7.2	Power loss vs. normalised frequency for various higher-order waveguide modes, operating 0.3 MHz from cut-off.	181
7.3	(Left) Percentage of power lost vs. delay in waveguide for 8 circular waveguide modes. (Right) Loss vs. v_g/c for 8 circular waveguide modes.	183
7.4	FDM schematic showing modulation, transmission, then demultiplexing.	184
7.5	Frequency switching RF system idea, using three linacs tuned to slightly different frequencies and an RF source with a large enough bandwidth to accommodate them.	185

7.6	RF pulse split into three sections with the frequency being changed every $3 \mu\text{s}$ by switching the klystron frequency using a LLRF system.	186
7.7	Circuit diagram of the modelled RF system showing the linac modelled as an RLC section and transformers representing the input coupler and probe.	186
7.8	Frequency domain spectrum of the input signal showing frequency content at 3 frequencies, separated by 3 MHz.	188
7.9	(Top) Simple $10\mu\text{s}$ square wave. (Bottom) FFT of the square wave showing sinc function.	190
7.10	Wigner-Ville distributions of the input signal, over a shorter length of $1 \mu\text{s}$, showing the three different frequencies in time. (Top) Wigner-Ville distribution. (Bottom) Smoothed pseudo Wigner-Ville distribution.	191
7.11	S_{11} curves for three linacs with 3 MHz separation between them, all within a 10 MHz bandwidth.	192
7.12	Simulation results of three linacs and three circulators, with $Q_0 = Q_e = 12000$ and $Q_t = 50000$. (Top) Reflected (left) and transmitted (right) voltage signals from the first linac. (Middle) Reflected (left) and transmitted (right) voltage signals from the second linac. (Bottom) Reflected (left) and transmitted (right) voltage signals from the third linac.	194
7.13	Simulation results of three linacs and three circulators, with varying Q_0 and Q_e and $Q_t = 50000$. (Top) Reflected (left) and transmitted (right) voltage signals from the first linac. (Middle) Reflected (left) and transmitted (right) voltage signals from the second linac. (Bottom) Reflected (left) and transmitted (right) voltage signals from the third linac.	195
7.14	Simulation results of three linacs and three circulators, with $Q_0 = Q_e = 12000$ and $Q_t = 50000$ for various values of the frequency difference between the resonances. (Top) Reflected (left) and transmitted (right) voltage signals from the first linac. (Middle) Reflected (left) and transmitted (right) voltage signals from the second linac. (Bottom) Reflected (left) and transmitted (right) voltage signals from the third linac.	196
7.15	(Left) Coaxial probe inserted into exit beampipe. (Right) Electric field showing coupling to probe.	198
7.16	S_{11} and S_{12} for coaxial probe inserted into cavity.	198
7.17	Realistic S-parameters for a C-band circulator.	199
7.18	Reflected (left) and Transmitted (right) signals from the first linac, including realistic circulator S-parameters.	199
7.19	Simulink Model for final RF system design to verify power requirements.	200
7.20	Signal reaching linac 3 after two circulators and reflection from two linacs showing $\approx 80 \%$ of the initial pulse reaching the third linac.	201

7.21	Measured signal showing the possibility of adjusting the frequency within the pulse. (Top) Time domain signal (Bottom) Frequency spectrum of the signal. Image reprinted with permission [170].	201
8.1	Geometry of object being scanned, it is rotated through θ , and the intensity of the transmission is stored for each rotation [175].	206
8.2	2D reconstruction using FBP method. (A) Original image (B) 1 projection (C) 3 projections (D) 4 projections (E) 16 projections (F) 32 projections (G) 64 projections [175].	208
8.3	CAD model of arbitrary object to test reconstruction.	210
8.4	Setup for simulating X-rays of an arbitrary object using ray-tracing.	210
8.5	Simulated X-rays of arbitrary object using ray-tracing (Left) Front view (Right) 45 degrees rotation.	211
8.6	(Left) 3D view of reconstruction, showing artefacts around the edges. (Right) Slice through the centre of 3D reconstruction showing reconstructed object in the centre.	211
8.7	3D image of reconstruction with points < 0 removed. Artefacts are still present, but the form of the object is visible.	212
8.8	Reconstruction of an object using 360 views in 1-degree increments with the FDK algorithm, after post-processing.	213
8.9	Simulated X-ray of realistic threat object with no internal parts.	213
8.10	Reconstruction of threat using 180 views split around 360 degrees.	214
8.11	Reconstruction of the threat using 5 views and FDK algorithm.	214
8.12	(Left) Reconstruction using 5 views and MLEM algorithm. (Right) Reconstruction using 3 views and MLEM algorithm.	215
8.13	CAD model of threat with density variations included.	216
8.14	Reconstruction of threat with density variations using 5 views and MLEM algorithm.	216
8.15	Reconstruction of threat with density variations using 3 views. (Left) MLEM (Right) SART	217
8.16	CAD model of boxes to test the ability of the algorithms to reconstruct positions accurately.	218
8.17	Reconstruction of boxes using 3 views and MLEM algorithm. Reconstruction of the smaller box is unsuccessful due to angles chosen for X-ray simulations.	218
8.18	CAD model of boxes that includes a wall and an object inside.	219
8.19	Simulated X-ray of boxes with objects inside.	219
8.20	Reconstruction of boxes using 3 views and MLEM algorithm.	220
8.21	Methods of employing CNNs for limited-view CT reconstructions.	221
9.1	Experimental setup for testing FDM technique using RF bunker at Daresbury laboratory.	234

- 9.2 (Left) Simulated reflected pulse from S-band linac showing first frequency (2.995 GHz) being absorbed into linac. (Right) Transmitted pulse into S-band linac showing linac accepting first frequency. . . . 235

Chapter 1

Industrial Accelerators and X-ray Imaging Techniques

1.1 Radio-Frequency Linacs

The first high energy (> 400 keV) particle accelerator was developed in 1932 by Cockcroft and Walton to facilitate the first man-controlled splitting of the atom. [1] It was a DC system that converted AC or pulsing DC to a high-voltage DC pulse used for electrostatic acceleration. Around the same time, resonant acceleration was also being developed, and the first linear accelerator (linac) that used time-varying radio frequency (RF) fields was first proposed by Ising in 1924. It was then eventually demonstrated by Wideröe in 1928 [1], where he put together a linac for his PhD Thesis, consisting of one drift tube between two accelerating gaps [2]. Since then, the development and usage of particle accelerators has grown rapidly, and there are now more than 40,000 operational accelerators in the world. Over 97 % of these accelerators are used for commercial purposes, and a large proportion of them are linacs. The current linear accelerator market is valued at US \$3.79B (2019) with projections to reach US \$6.8B by 2027. [3]

Aside from the more recognisable accelerators that are used for scientific research, such as the LHC at CERN [4], accelerators play a crucial role in a multitude of

sectors, including:

- Medical (particle and radiotherapy for the treatment of cancerous tumours)
- Security (generation of X-rays for cargo screening)
- Research (high-energy particle physics, Synchrotron radiation facilities for generating high brightness X-rays [5])

Linacs utilise electromagnetic power to accelerate charged particles, and they are capable of providing beams with energies ranging from keV to tens of GeV and currents from nA to kA, which are then used in a variety of ways depending on the application. One of the ways these charged particles can be used is to directly alter the properties of materials, such as surface hardening of metals, polymer cross-linking (for example, making shrink wrap) and improving semiconductor properties [6, 7]. The beams can also be used for sterilisation; as the energy from the electrons is deposited into cells, it can alter chemical bonds and damage DNA to destroy the reproductive capabilities of microorganisms. This method is used on medical instruments, pharmaceutical products, and food packaging. Beam sterilisation can offer improvements over traditional processing methods, such as better efficiency over thermal processing, a reduction in the use of toxic chemicals, and the ability to immediately use the product after sterilisation [8].

There are a plethora of other industrial applications of accelerators, with far too many to list here. They range in size from a few centimetres up to a few metres, depending on the required beam energy for the application and the type of accelerator design. However, in general, for industrial applications, the energies are within the range of 2-15 MeV. Some of these specific applications include [9–11]:

- Ion implantation for making semiconductor devices
- Production of radionuclides for diagnostic medical imaging
- Neutron production for material research

- Ion-beam analysis to probe the elemental composition of objects
- Waste water treatment
- Treatment of sludge and flue gases to remove pollutants
- Treatment of asphalt to improve wear resistance
- Generation of X-rays for non-destructive testing (NDT)
- Increasing the efficiency of electrical power transmission
- Leather tanning [12]

1.2 X-rays for Security

This thesis is concerned with the design of low-energy (sub 10 MeV) electron linacs, specifically for cargo scanning security applications where the charged particle beams generate X-rays when they are fired into a dense target. The history of X-ray scanning for security applications dates back to the 1960's and 1970's and was introduced primarily as a result of an increasing rate of aeroplane hijackings. Sixty-six aircraft of United States or foreign registration were hijacked from January 1, 1961 through to 1968 and 277 were hijacked from January 1, 1969 to 1972 [13, 14]. This led to an increase in security measures, mainly profiling, and then the use of metal detectors and X-ray machines on the profiled individuals [14]. In 1973, the FAA started screening all passengers and searching their bags; this and other measures quickly led to a rapid decrease in the number of successful hijackings, as shown in Fig. 1.1.

As there were already a number of X-ray tube systems being used for medical and NDT purposes, with energies ranging from a few keV up to 1 MeV, the X-ray sources were already available for a X-ray baggage scanning system. The methodologies for generating and processing X-ray images were also available, as they had been used for medical and NDT applications since the early 1900s [9]. One of the initial




Figure 1.1: Plot showing reduction in successful hijackings after 1973. Image sourced from [15].

systems used for baggage screening was the Philips Saferay, shown in Fig. 1.2 with an example of an X-ray image taken by the system, with a weapon inside a case shown in Fig. 1.3.

The system was very slow as it required the manual placement of an object, with each item having to be placed in front of the screen before each scan. Eventually, this system was phased out in favour of systems that used conveyor belts to increase scanning speed [15]. By the end of 1974, there were 260 X-ray luggage systems, 1100 walk-through metal detection systems, and 1200 hand-held scanning systems in use across U.S airports. [15].

1.3 Cargo Scanning Linacs

Following the introduction of baggage scanning for air passengers, the September 11 attacks (9/11) led to new protocols for scanning cargo worldwide, mostly due to concerns that radioactive material could be smuggled in through cargo routes [16]. In the U.S at the time of 9/11, none of the containers entering the U.S were being



Figure 1.2: First Philip's Saferay system, with manual placement of the object to be scanned. Image sourced from [15].




Figure 1.3: Example of an X-ray taken by the Philip's Saferay system. Image sourced from [15].

scanned. Now, the U.S mandates that 100 % of the cargo entering ports in the U.S should be scanned through the Container Security Initiative (CSI), though this is yet to be fully realised [17, 18].

1.3.1 Shipping Containers and Heavy Cargo

The requirements then changed from scanning small objects with thin walls, such as luggage to standard shipping containers with thicker steel walls capable of holding a

significant amount of cargo. Consequently, the X-ray sources required for scanning also changed. For a narrow beam of monoenergetic photons with incident energy I_0 penetrating a layer of material with mass thickness x and density ρ , it emerges with a new intensity I , given by

$$\frac{I}{I_0} = \exp [-(\mu_a/\rho)x] \quad (1.1)$$

Where μ_a is the attenuation coefficient [19]. This means that previously used sources, such as X-ray tubes and radioactive sources, did not provide the energies needed for sufficient material penetration. This led to the use of linacs instead of X-ray tubes, as they were able to provide energies greater than 1 MeV to meet this requirement. Fig. 1.4, shows a comparison between three types of X-ray sources, with the energies, penetration in steel and limitations of each system [20]. The linac systems used for these applications are generally referred to as Cargo and Vehicle Inspection (CVI) systems.

Technology	Energy Range	Approximate Steel Penetration/mm	Method	Application	Limitations
Vacuum Tubes	50 - 450 keV	38	Electron beam acceleration and rapid deceleration by an anode target to create Bremsstrahlung	Baggage and mail scanning	Poor penetration so not suitable for CVI
					Above 150 keV there are anode cooling issues
Cobalt-60 Source	1.17-1.33 MeV ^[4]	63.5 - 229	Radioactive source emits gamma rays of different energies. The photopeaks for a scintillation counter occur at 1.17 and 1.33 MeV.	Vehicle inspection and non-destructive testing (NDT)	Source lifetime and output dose variation with half life
Linac Systems	2 - 10 MeV	133 - 390	Electron beam acceleration using a linear accelerator (linac) and absorption by a high-Z target to produce X-rays.	Heavy container and vehicle scanning	Above 6 MeV there is neutron production and more shielding required
					Poor penetration below 3 MeV

Figure 1.4: Table showing the comparisons between X-ray vacuum tubes, radioactive sources and linac systems. Table sourced from [20].

1.3.2 Types of CVI Systems

The research discussed in this thesis is focused on the design of low-energy (2-6 MeV) linacs primarily for vehicle or light-cargo scanning applications. Generally, these types of systems can be split into four categories as shown in Table 1.1:

- Stationary
- Portal
- Gantry
- Mobile

Type of system	Description
Stationary	System cannot be easily relocated. Typically larger systems with higher energies and lower frequencies as shielding can easily be implemented.
Portal	System remains fixed while vehicle moves through the system. Some require a pulley system if the radiation levels are too high and some allow the vehicle to be driven.
Gantry	System moves along while vehicle remains stationary. Can operate in both directions to improve scanning speed.
Mobile	Compact system that can be mounted on other vehicles. All peripherals included with the device. Typically higher frequency for smaller dimensions (X-band).

Table 1.1: Types of cargo scanning systems.

The design work presented here is for a device that could potentially be used for either a mobile system type or a portal type.

1.4 Cargo Scanning Linac System Components

Cargo scanning systems that use linacs require several different components to operate, with an example of a generic CVI system set-up shown in Fig. 1.5. Each of the key components are discussed briefly below:

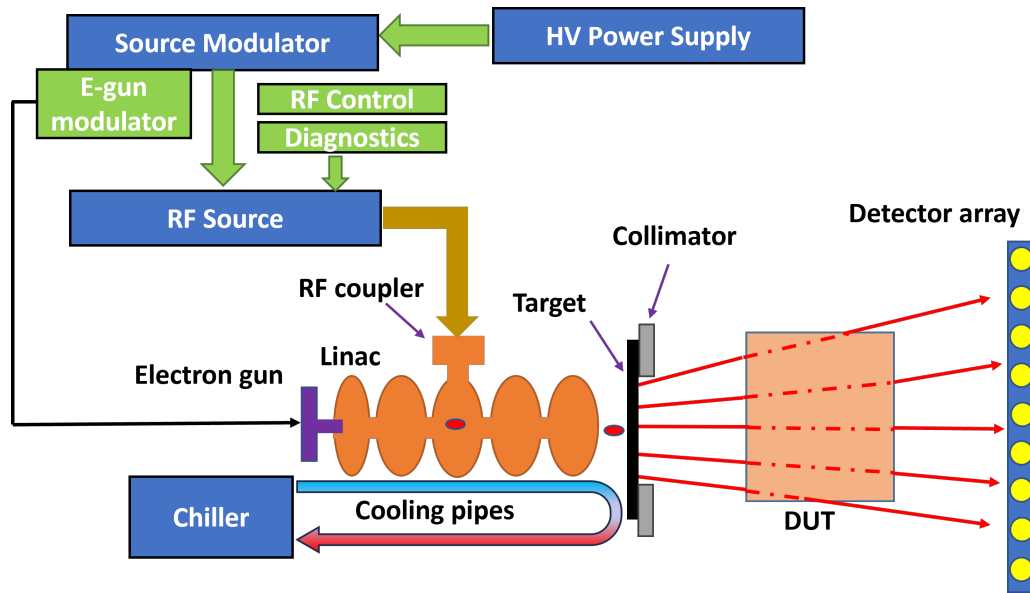


Figure 1.5: Simple schematic of a CVI system.

- Electron gun: The electrons are generated through a thermionic emission process on a cathode [21], with initial energies ranging from 15-50 keV and peak currents of 50-200 mA.
- High voltage supply: The high voltage supply provides the voltages for the modulator and the electron gun. They normally include regulators, filtering, and feedback to ensure that the pulses are stable. It also requires protection against overvoltage and overcurrent.
- Modulator: A modulator is used to provide low-current, high-voltage pulses from the low-voltage, high-current AC supply that goes to the HV supply to drive the electron gun and the linac. There are numerous types, with the most common one used for klystrons being the line-type modulator. The first part of the system is a pulse-forming network (PFN), which is a series of capacitors and inductors that charge up to create the shape of the output pulse so that it is compatible with the klystron. After this a switching device such as a thyrotron or Insulated Gate Bipolar Transistor is used to rapidly discharge the stored energy into a pulse transformer, which steps up the voltage. The

pulse widths for the linac pulse are usually a few μs , and the repetition rates can range from 50-800 Hz.

- RF source: The RF source is typically a magnetron [22], which is a type of electron tube widely used in microwave devices such as microwave ovens and radar systems. It operates on the principle of generating electromagnetic waves through the interaction of electrons with a magnetic field. They are generally cheaper than other sources but require frequency control to keep the frequency of the generated RF power the same as the frequency required by the linac. They also suffer from some drawbacks, such as pulse jitter and a probability of missing pulses [9]. The other option for the RF sources is Klystrons. These offer some benefits over magnetrons, as their frequency and power can be easily controlled and they do not miss pulses. They are generally larger and heavier but can also provide higher average powers compared to magnetrons. The lifetime of a klystron is also generally longer, but they are considerably more expensive than magnetrons, so they are rarely used for systems where cost is an important factor.
- RF cavity: The RF cavity takes the RF power from the RF source and accelerates the electron beam generated by the gun. They usually consist of a RF coupler to couple power into the structure, a number of accelerating cells, and sometimes coupling cells, depending on the mode of operation. Generally, the length is proportional to the required final energy. There are several different options for the RF cavity, and these are discussed in more detail in Chapter 3.
- Target: The target usually consists of a high atomic number (Z_a) material (typically tungsten) that generates bremsstrahlung when electrons are rapidly decelerated as they hit the target. The linac and target are typically water-cooled to avoid excessive heating, as only a few percent of the electron energy ends up as X-rays, with the rest being dissipated as heat.

- Control system and interface: The control system allows for the setting of various parameters, such as the linac repetition rate and the frequency of the magnetron. It also provides status information of all subsystems so that they can be monitored.
- Cooling system: A chiller is normally used to cool the linac, target, and RF source using a liquid coolant. The chiller provides the coolant at a constant temperature (typically between 30 °C and 35 °C).

1.5 Generating Images

There are a number of processing steps before the X-rays generated from the target allow for a quality X-ray image to be realised.

1.5.1 Collimators

In order to reduce the amount of X-ray scatter from other objects that are not of interest, various types of collimator are employed. The collimator is normally made from a high Z_a material such as lead. Some options for collimators include putting lead bricks in the beam path, placing shielding around the detectors just before the X-rays are detected, or placing internal shields in the linac. The materials (normally lead or tungsten) and thicknesses are chosen appropriately to match the beam energy, and the required X-ray beam profile. Fig. 1.6 shows an example of an external collimator where the distance between the plates can be adjusted using a motor, allowing a wide range of different shapes and dimensions to be scanned.

1.5.2 Detectors

The detectors are typically made from scintillating crystals which are materials that exhibit the phenomenon of scintillation - this refers to the emission of light when exposed to ionising radiation, such as X-rays. This property is exploited for




Figure 1.6: Example of a CVI collimator. Image sourced from [9]

cargo scanning, where the crystals convert incoming X-rays into detectable photons. The visible light is then sensed by photo-diodes and processed electronically. The crystals are selected based on trade-offs between radiation hardness, conversion efficiency, after-glow, response time, and cost [9]. Common materials used for crystals include sodium iodide (NaI), caesium iodide (CsI), and cadmium tungstate (CdWO₄). An example of a detector module is shown in Fig. 1.7. It consists of individual pixels made from the crystals and photodiodes, a module housing the electronics, and finally, the array is built up from multiple modules. A curved profile is sometimes utilised to ensure that the direct path of the X-rays from the source remains equidistant to all the detectors, thus maintaining a constant focal distance.

1.5.3 Movement

The detecting array is typically 1D, and a 2D image is generated either by moving the cargo across the array, for example, by driving a vehicle at a low speed, or by mounting the detectors onto rails which then move past the vehicle. In this way, a 2D image is constructed. Two examples of these types of systems are shown in Fig. 1.8. The G60 system on the left is on rails, which allows it to move past a stationary vehicle at 0.4 m/s, which enables 20 trucks per hour to be scanned. The P60 portal system on the right has scanning speeds from 0.8-2.2 m/s and a throughput of up



Figure 1.7: Example of a curved detector array. From left to right, pixel, module, curved array. Image sourced from [23]

to 100 vehicles per hour [24, 25]. The repetition frequency of the linac is normally adjusted to match the speed of the vehicle, since the image quality will depend on the number of pulses during the time window that the vehicle is passing the X-ray source and detectors. There are now systems capable of scanning rail cargo, with speeds up to 60 km/h, requiring repetition rates of up to 800 Hz [26].



Figure 1.8: (Left) Example of a gantry system, Rapiscan EAGLE G60, Image sourced from [25] (Right) Example of a portal system, Rapiscan EAGLE P60, image sourced from [24].

1.5.4 Material Discrimination

In recent years, the use of material discrimination methods has become popular as a way to identify the types of material within the cargo [27]. When the X-rays generated by the target propagate through the DUT, they are attenuated by different attenuation factors depending on the material. Analysis of the resulting spectrum allows for the generation of the image based on the different attenuation constants of the materials. Most material discrimination methods require two images to be taken at two different energies, usually 6/4 MeV or 8/4 MeV [28]. One of the methods that is used for material discrimination is the R-curve method [28], where two images are taken at two energies, E_1 and E_2 and then a ratio of their logs taken:

$$R(E_1, E_2, Z_a) = \frac{\log(T_1)}{\log(T_2)} = \frac{\mu(E_1, Z_0) D(E_1)}{\mu(E_2, Z_0) D(E_2)} \quad (1.2)$$

Here, Z_a is the atomic number, T is the transparency (normalised image to the number of photons) and $D(E)$ is the detector sensitivity [27]. This provides a rough estimate of the atomic number of the material. For a monochromatic energy, the material could be easily uniquely identified, but the spectra from these systems are generally a continuous distribution, meaning that, in reality, only groups of materials can be identified. Fig. 1.9 shows how the mass attenuation coefficient varies for different material types and X-ray energies. Most organic material and some explosives will have a composition similar to carbon, light machine and car parts similar to aluminium, containers and heavy machine parts similar to iron, and radioactive material/heavy elements similar to uranium. Using this difference in attenuation coefficient, different material groups can be differentiated using these discrimination techniques. An example of a generated image with colour coding to represent the material is shown in Fig. 1.10. Some examples of dual-energy interlaced system designs can be found in [29–32].




Figure 1.9: Mass attenuation vs. energy for different materials. Image sourced from [33].

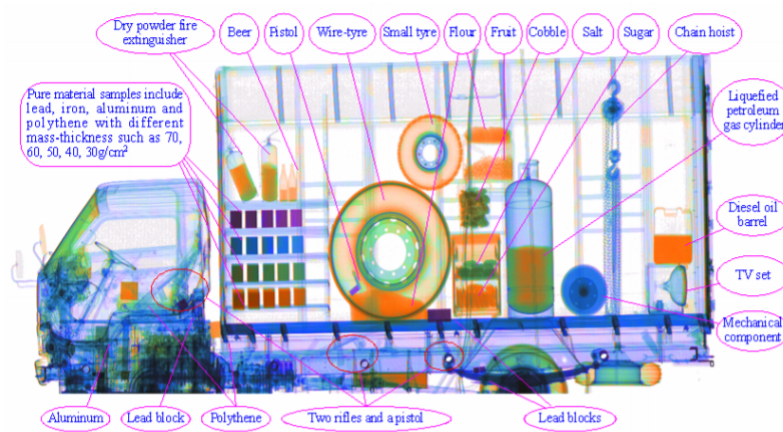


Figure 1.10: Example image of a truck using color to represent material groups based on atomic number. Image sourced from [33]

1.5.5 CT Applications and Image Reconstruction

Computed tomography (CT) methods have been used in the medical industry since the 1970s to provide images with much higher contrast compared with traditional radiography [34]. The basic principle involves measuring the attenuation of X-rays at hundreds of angles and then using various reconstruction algorithms to generate an image of the object's internal structure [35]. Lower energy CT systems are already used in many airports and for cabin and hold baggage, typically using dual-

view or multi-view machines [36]. There is now some interest in using CT methods with higher-energy linacs. However, this poses challenges due to the large cargo sizes and heavy linac equipment that must be rotated [37]. There are also concerns about whether the scans can be completed fast enough to make them feasible for the high-throughput requirements at most ports and border crossings.

1.6 Design Criteria

Standard shipping containers typically require beam energies of 3-9 MeV to obtain sufficient contrast in the images due to the thickness of the containers and the amount of cargo inside. There has been a push to scan smaller types of cargo, such as Aviation Cargo Unit load devices (ULDs), small trucks, and even passenger vehicles [38]. The energies currently used for cargo scanning (≈ 6 MeV) would not provide sufficient contrast when used on these smaller cargo types and therefore lower energy (1-3 MeV) devices are required for these applications. If the energy is too high, then most of the X-rays will pass through the container with little attenuation, meaning that the image will not have adequate contrast for discerning internal structure. Table. 1.2 shows the penetration depth of X-rays in steel and water for different energies. It shows that for thinner containers a lower energy is required. This thesis will present a design of a low-energy device (1-2 MeV) that could be integrated into a compact scanning system, to scan steel or metallic walls with thicknesses of up to 20 cm.

Energy (MeV)	1.25	2	3	4	6	9
Penetration Steel (mm)	133	205	297	352	406	430
Penetration Water (mm)	880	1370	2050	2530	3160	3640

Table 1.2: Penetration of steel and water vs. electron energy. [39]

1.6.1 Selection of RF Frequency

Many universities and commercial companies have developed linacs in the 3-9 MeV range for cargo scanning applications, with some at S-band (2-4 GHz) [29, 40, 41] and others at X-band (8-12 GHz) [38, 42]. Until recently, the majority of electron linacs used for cargo scanning have operated at one of these two frequencies. This is because there were many commercial RF power sources available at those frequencies as they were historically used for medical and radar applications.

ID, ref	1, [43]	2, [44]	3, [45]	4, [46]	5, [47]	6, [48]
Frequency (GHz)	5.52	5.712	5	5.712	5.712	5.712
Energy (MeV)	2	10	6	6	2 per turn	10
Use	X-ray, Industrial	Medical	X-ray, Industrial	Industrial + Medical	Microtron	Medical
Gradient MV/m	N/A	30.3	13.3	27.75	24	N/A
Z (MΩ/m)	141	98	110	130	108	140
P_{in} (MW)	1	2.66	1.5	2.16	0.6	11
Coupling type	On-axis (slot)	On-axis (slot)	On-axis (slot)	On-axis (slot)	On-axis (slot)	Hybrid
Capture %	30	35	N/A	40	N/A	1
E_{pk}/E_a	N/A	4.2	N/A	4.1	4	2.99
Q_0	11782	7100	8500	11500	11700	9991

Table 1.3: C-band industrial linac examples.

Typically, S-band has been used for stationary systems and X-band for portable ones. The higher frequencies at X-band mean that the structures are significantly smaller as they scale with the wavelength, with the downside being that the

manufacturing tolerances are tighter, which increases manufacturing costs or reduce the performance of the linacs if the tolerances are not met. Recently, several commercial sources at C-band (4-8 GHz) [49–51] have come on the market, sparking an interest in developing electron linacs at C-band for mobile cargo scanning and other applications. Some examples of recent C-band linac developments can be found in [43–47, 52, 53]. A list of the types of structure and key parameters from these sources is given in Table 1.3. Operating at C-band can reduce the cost of manufacturing compared with X-band linacs and provide a compact system compared to S-band that could potentially be portable if required [54]. Using C-band over S-band can also provide higher gradients due to an increased voltage breakdown threshold and a shorter filling time, both desirable properties. The shunt impedance per unit length also scales with \sqrt{f} , leading to a higher accelerating efficiency. C-band acquired its name as it was a compromise between S and X bands for radar systems; clearly, this is still holds true for some normal conducting accelerator systems [55]. The frequency scaling of some of the important parameters when discussing RF cavities are given in Table. 1.4. Some example SW Linacs are shown in 1.12 giving an idea of the dimensions of the components.

Cavity Parameter	Scaling with frequency
Quality factor, Q	$f^{-1/2}$
Power loss, P_d	$f^{-1/2}$
Surface resistance, R_{surf}	$f^{1/2}$
Shunt impedance per unit length, Z	$f^{1/2}$

Table 1.4: Frequency scaling of cavity parameters.

1.6.2 Commercial Low-Energy Linacs

A number of commercial suppliers also offer low-energy linacs (1-10 MeV). Some of these include the Linatron XP (0.95-1.3 MeV), the Linatron M1 (1-1.5 MeV, X-band) and the Linatron M3 (1-4.5 MeV, S-band), all from Varex imaging (spin-off

from Varian Medical Systems) [23]. MEVEX also offers low-energy linacs (1-2 MeV, 3 MeV) [56] and Radiabeam also offers some low-energy options (ARCIS 2-9 MeV [57], MIXI 2-6 MeV [58]) with all devices operating at either X-band or S-band. Some companies also offer fully integrated scanning systems such as Rapiscans Z Portal® [59] or the NUCTECH MT series (4-6 MeV) [60], both of which have dual energy technology for material discrimination. This thesis will present the design of a compact, low-energy, C-band accelerator for cargo scanning applications. It will also explore the use of this linac for MV CT imaging, using an RF system design to allow quasi-3D CT images to be generated quickly and with a low cost.




Figure 1.11: Dual energy linac scanning example. Image sourced from [61]




Figure 1.12: Cargo scanning linac Examples. Image sources from [62]

1.7 Thesis Outline

Chapter 2 of this thesis details some of the basic theory of RF acceleration and electron linacs required to design a linac system. It has been covered in great detail in many places, and the reader is referred to [63] for a more in-depth treatment. Chapter 3 is concerned with the optimisation process for the single cells of the linac. It details the multi-objective optimisation tools and algorithms used to design the cells, and the trade-offs made when selecting the final designs. Chapter 4 is concerned with the full linac design and optimisation. It begins by describing a beam optimization performed using particle tracking codes to minimize the number of lost electrons and improve cathode lifetimes. Then, it reports on the process of assembling the individual cells and tuning the linac, alongside details of particle-in-cell (PIC) simulations used to verify the final design. The optimised beam parameters from the linac design are then used to design a high atomic number X-ray target in Chapter 5. The target thickness is optimised to give the maximum X-ray yield from the electron beam-target interaction. A thermal analysis is also performed to ensure that the target does not overheat. Chapter 6 details a full thermal analysis of the linac, including computational fluid dynamics simulations (CFD). A comparison is drawn between analytical results for water flow in pipes and the results obtained from simulation packages. Chapter 7 incorporates the linac into a full RF system design that would allow multiple linacs to fire sequentially with a chosen delay between them. The system is simulated and compared with experimental results obtained by colleagues at Daresbury Laboratory, UK. Chapter 8 presents the results of image reconstruction using few (< 5) views from simulated X-rays of various objects, including potential threat items such as weapons. The objective of this was to test whether a 3 linac system would be adequate for generating useful quasi-3D images. Finally, Chapter 9 provides the conclusions of the thesis and gives recommendations on further work that should be completed to improve the design in the future.

Chapter 2

RF Acceleration and Theory

In this Chapter the basic theory essential for understanding the project is presented. Only key results are described, without delving into full details, as comprehensive treatments can be found elsewhere [63, 64].

2.1 RF Acceleration

The electric field along the axis of a standing wave in an RF cavity is given by

$$E_z(z, t) = E(z) \cos(\omega t + \phi) \quad (2.1)$$

where ω is the angular frequency and ϕ is the phase. If we imagine a particle arriving in the gap on the LHS of Fig. 2.1 and that there is an electric field on axis, then the field experienced by the particle is described as

$$E_z(r = 0, z, t) = E(0, z) \cos[\omega t(z) + \phi] \quad (2.2)$$

where

$$t(z) = \int_0^z v(z) dz \quad (2.3)$$

is the time the particle is at position z . At $t = 0$, the phase of the field relative to the crest is ϕ . The energy gain of the particle through the gap is then

$$\Delta W = q \int_{-L/2}^{L/2} E(0, z) \cos(\omega t(z) + \phi) dz \quad (2.4)$$

which can also be expressed as $\Delta W = qV_0 T \cdot \cos(\phi)$, where T is the *Transit Time Factor* (TTF) and V_0 is the axial voltage across the cavity defined by

$$V_0 = \int_{-L/2}^{L/2} E(0, z) dz \quad (2.5)$$

Assuming that $z = 0$ is at the centre of the cavity, the transit time factor is defined as

$$T = \frac{\int_{-L/2}^{L/2} E(0, z) \cos(\omega t(z)) dz}{\int_{-L/2}^{L/2} E(0, z) dz} - \tan(\phi) \frac{\int_{-L/2}^{L/2} E(0, z) \sin(\omega t(z)) dz}{\int_{-L/2}^{L/2} E(0, z) dz} \quad (2.6)$$




Figure 2.1: Gap between cells (left), Electric field inside this gap (right).

and for a particle arriving at the origin when the field is maximum ($\phi = 0$) and assuming a square electric field profile constant across the gap so that E_z falls immediately to 0 outside the gap, then Equation 2.6 reduces to

$$T = \frac{\sin(\pi g / \beta \lambda)}{\pi g / \beta \lambda} \quad (2.7)$$

[63]. Fig. 2.2 illustrates how the TTF varies with the gap size. It is clear from Equation 2.7 that the TTF should be equal to 1 for the maximum energy gain, meaning that the gap length should be 0. Of course, it is not practical to design a system with a gap length of 0, as the cavities need to have walls between them, but this fact must be balanced with other constraints such as peak fields and power requirements when designing a cavity with short gaps. In practice, a good approximation for the optimum cavity length for a standing wave accelerator is given by: [64]

$$L_{optimum} \approx \frac{\pi\beta c}{\omega}. \quad (2.8)$$




Figure 2.2: Transit time factor as a function of gap length.

2.2 Figures of Merit

This section describes the figures of merit (FOM) typically used to describe RF cavities similar to the cell shown in Fig. 2.3. These FOMs facilitate the optimisation of the cavity geometry and provide a way to quantify the cavity performance.




Figure 2.3: Simple single cell cavity vacuum space (left), cross-section (middle), copper outer walls (right).

2.2.1 Accelerating Gradient

The accelerating gradient E_{acc} will be used extensively throughout this thesis, as it defines how much energy is given to a particle per unit length on average. Taking Equation 2.4 and multiplying by the TTF, we get the accelerating voltage $V_c = V_0 T$. The accelerating gradient is then:

$$E_{acc} = \frac{|V_c|}{L_c} \text{ (MV/m)} \quad (2.9)$$

where L_c is the length of the cavity and the average voltage is used. The accelerating gradient allows for the normalisation of the other key field quantities so that they become field-independent.

2.2.2 Peak Fields

When designing a cavity, the maximum electric and magnetic fields would normally be minimised to reduce the probability of breakdown and heating effects [65]. Examples of the electric and magnetic field distributions in a simple cavity is shown in Fig. 2.4. For this reason, two useful quantities are the normalised peak electric field:

$$\frac{E_{pk}}{E_{acc}} \quad (2.10)$$

which is unitless, and the normalised peak magnetic flux density

$$\frac{B_{pk}}{E_{acc}} \quad (2.11)$$

which is normally given in units of (mT/MV/m). These will be described in more detail in the next section.




Figure 2.4: Simple single cell cavity showing electric (left) and magnetic (right) fields in the volume.

2.2.3 Shunt Impedance and Cavity Trade-offs

Another useful figure of merit is the shunt impedance defined as

$$R_s = \frac{|V_c|^2}{P_d} \quad (2.12)$$

where V_c is the cavity accelerating voltage including the TTF and P_d is the power dissipated in the cavity walls (RF power loss). It should be noted that in some texts the circuit definition of the shunt impedance is used where there is a factor of two difference:

$$R_{s,circuit} = \frac{|V_c|^2}{2P_d} \quad (2.13)$$

The difference arises from the choice of peak or RMS voltages where:

$$V_{RMS} = \frac{V_{pk}}{\sqrt{2}} \quad (2.14)$$

The shunt impedance measures the effectiveness of the cavity in producing

an axial voltage for a given power dissipation in the cavity walls. A quantity independent of the of the cavity length, useful for linear accelerators is the shunt impedance per unit length is given by:

$$Z \equiv \frac{R_s}{L_c} = \frac{|V_c|^2}{P_d/L_c} \quad (\Omega/\text{m}) \quad (2.15)$$

Z is used to refer to the shunt impedance per unit length including the TTF for the rest of the document. It should also be noted that other authors use different definitions for the shunt impedance for example the use of an effective shunt impedance when describing the shunt impedance over different particle energy ranges. The definition also may or may not include the transit-time factor.

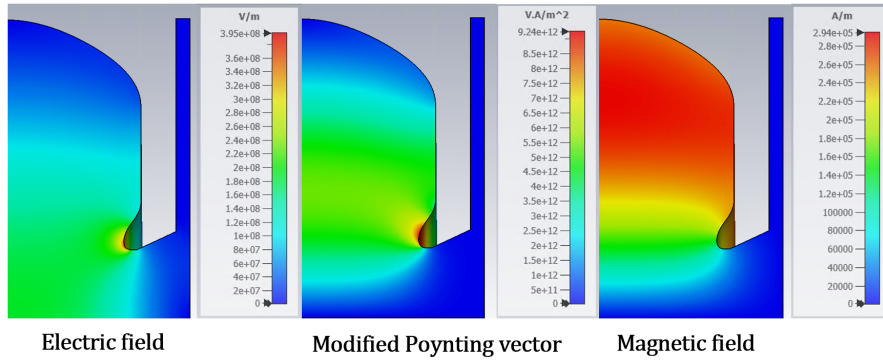


Figure 2.5: Example of fields in a cavity cell, showing where the maximum occurs and the effect of adding nose cones. Fields are normalised to 1 J stored energy.

The addition of nose cones, as depicted in Fig. 2.5, can increase Z by enhancing the axial electric field relative to the total stored energy, especially as the TTF increases with a reduction in gap length for the same voltage [63]. However, this enhancement comes at the cost of increased peak surface electric fields, which can significantly elevate the likelihood of RF breakdown, therefore the quantity E_{pk} must be minimised to avoid this, leading to a set of conflicting objectives. Another quantity that has recently been found to give a good indication of RF breakdown is the modified Poynting vector which also rises as the nose cones are made larger [66]. It is given by:

$$S_c = \Re(\bar{S}) + g_c \cdot \Im(\bar{S}) \quad (2.16)$$

Where \bar{S}_c is the complex Poynting vector given by:

$$\bar{S} = \bar{E} \times \bar{H} \quad (2.17)$$

and g_c is a weighting factor used to distinguish between the coupling of active and reactive power flow to field emission power flow. It is normally taken as 1/6 [66]. This is another quantity that must be minimised, although it typically becomes more important when operating at very high gradients. The key quantities of interest are then the three peak fields and the shunt impedance, shown in Table. 2.1 with the appropriate units.

Quantity	Maximise/Minimise	Units
E_{pk}/E_{acc}	Minimise	Unitless
B_{pk}/E_{acc}	Minimise	mT/MV/m
$\sqrt{S_c}/E_{acc}$	Minimise	$\sqrt{A/V}$
Z	Maximise	M Ω /m

Table 2.1: RF single cell objectives.

2.2.4 Field Limits

A more extensive treatment of the peak fields is now given along with some examples of numeric values for an assumed accelerating gradient of 25 MV/m. Recently, many studies have been made into the phenomenon of electrical breakdown in RF accelerating structures, and from this work practical working limits have been found for the fields in the cavities [67, 68], they are explored here.

2.2.4.1 Electric Field

It is generally believed that electric field enhancement caused by small defects on the cavity walls causes electrical breakdown in accelerating structures. The sites where the field is enhanced heat the cavity material and cause it to vaporise. The process causes a plasma discharge, which then reflects the incoming RF power. Although it is a statistical phenomenon, by minimising maximum surface fields, the probability that electrical breakdown will occur can be reduced to acceptable levels (a number normally used to compare structures is 10^{-6} breakdowns per pulse per metre (bpp/m)). This value comes from the CLIC project requirement of a maximum BDR (breakdown rate) of 3×10^{-7} bpp/m, set to limit the luminosity loss due to breakdown of a 3 TeV design to be less than 1 % [69]. This requirement is quite stringent, therefore 1×10^{-6} is considered a good value for short industrial or medical linac designs [70]. Before accelerating structures are used, they are normally 'conditioned', whereby the RF power is increased very slowly, ensuring that the breakdown rate is kept below some threshold. By doing this, the sites with a higher probability of breakdown can be vaporised in a controlled manner, burning off enhanced regions and eventually leading to a more uniform breakdown probability across the entire surface of the cavity.

2.2.4.2 Breakdown Rate Scaling

Recent tests at C-band suggest that the dependence of the breakdown rate on the maximum surface field and the pulse length is given as [70, 71]

$$\frac{E_{s,max}^{35} t_p^{4.9}}{BDR} = \text{constant} \quad (2.18)$$

Where BDR is the breakdown rate and t_p is the pulse length. Using the results presented in [70, 71] and scaling for a pulse length of $4 \mu\text{s}$ (typical for normal conducting superconducting linacs), for a breakdown rate of 10^{-6} , a maximum surface field of 230 MV/m should be achievable at C-band. Conservatively, an ultimate limit of

200 MV/m was used for this design. As mentioned previously, S_c can also be used as a predictor of electrical breakdown with good success [67]. The BDR has been shown to scale with S_c as

$$\text{BDR} \propto S_c^{15} \cdot t^5 \quad (2.19)$$

A generally accepted optimistic maximum value of S_c is $4 \text{ W}/\mu\text{m}^2$ for 10^{-6} bpp/m with a pulse length of 200 ns. If we assume a pulse length of $4 \mu\text{s}$, which is typical for NC cargo scanning linacs, then the optimistic limit for the cavity would be $1.5 \text{ W}/\mu\text{m}^2$.

2.2.4.3 Magnetic Field and Pulsed Heating

It has also been found that the magnetic field can also play a role in causing breakdowns, due to a cyclic heating effect [72–74]. As copper is not a perfect conductor, the fields penetrate a certain depth into the material and current flows through some resistance, resulting in a heating effect. This depth is called the skin depth, given by

$$\delta = \sqrt{\frac{2}{\sigma_c \mu \omega}} \quad (2.20)$$

Where σ_c is the conductivity, μ is the permeability of the conductor, given by $\mu_0 \mu_r$, where μ_r is the relative permeability of the material (normally taken as 1), and μ_0 is the permeability of free space, $\mu_0 = 1.25663706212(19) \times 10^{-6} \text{ NA}^{-2}$ [75] and ω is the operating angular frequency. The surface resistance at RF frequencies is $R_{surf} = 1/\sigma\delta$ so substituting equation 2.20 gives:

$$R_{surf} = \sqrt{\mu\omega/2\sigma_c} \quad (2.21)$$

The average power dissipated per unit area into a lossy conductor by a magnetic field tangential to the surface per cycle is given by:

$$P_{density} = \frac{dP}{dA} = \frac{R_{surf}}{2} |H_{||}|^2 \quad (2.22)$$

where H_{\parallel} is the parallel component of the local magnetic field strength on the surface, and dP/dA is the dissipated power per unit area. The surface temperature rise on the cavity surface depends on these properties and is given by:

$$\Delta T = \frac{2P_{density}\sqrt{t_p}}{\sqrt{\pi\rho\kappa c_\epsilon}} \quad (2.23)$$

where ρ is the material density, k is the thermal conductivity and c_ϵ is the specific heat capacity. Using these formulas, one can estimate a maximum value of B_{pk} in mT that keeps the temperature rise of the surface < 40 °C. This temperature is selected as at temperatures higher than 50 °C, noticeable changes can be observed in copper grain boundaries [76]. If the cavity uses short pulses, compared to the time it takes for the heat to diffuse into the bulk of the copper then this heating is exacerbated as it is localised in a small volume in the skin depth in the copper. When temperatures start to rise above this, thermal stresses and fatigue over many cycles can create cracks, voids, fatigue slip bands and other defects due to the yield strength of the copper being exceeded [77].

A convenient expression for the temperature rise in copper has been worked out to be: [78]

$$\Delta T = 127|H_{\parallel}|^2\sqrt{f_{rf}(\text{GHz}) t_p(\mu s)} \quad (2.24)$$

Where $|H_{\parallel}|$ is the magnitude of the parallel component of the magnetic field on the surface, f_{rf} is the cavity resonant frequency and t_p is the RF pulse length. As an example, for a C-band cavity at 5.712 GHz, to ensure the temperature rise of the surface remains < 20 K the value of B_{pk} must remain below 229 mT, for a pulse length of 4 μs . Industrial linacs require a robust and stable design; therefore, the limits on the maximum fields were reduced by ≈ 50 % in order to remain well below the maximum during operation. As an example, if the gradient of a cargo scanning linac is 25 MV/m, this leads to a corresponding surface electric field of 100 MV/m, and the other field limits shown in Table 3.4. Fig. 2.5 shows where the peaks of the

surface fields generally occur within a single cell of a cavity.

2.2.5 Cell-Cell Coupling

Another important parameter to consider for a periodic array of cells is the cell-to-cell coupling constant. This parameter is related to the frequency range of the cavity as will be shown and, therefore, gives an idea of the tolerances on the frequency for a given value of the coupling constant. An array of periodic cells that are electrically coupled through the irises can be modelled as an equivalent circuit, as shown in Fig. 2.6. This circuit represents one cell, and connecting them in series represents a cavity.




Figure 2.6: Equivalent unit cell circuit for electric coupling through an iris.

Finding the impedance and admittance of the circuit we have:

$$Z = j\omega L_1 + 1/j\omega C_1 = \frac{1}{j\omega C_1}(1 - \omega^2/\omega_0^2) \quad (2.25)$$

where $\omega_0^2 = 1/L_1 C_1$, and $Y = j\omega C_2$. Substituting these into

$$\cos \phi = 1 + ZY/2 \quad (2.26)$$

where ϕ is the phase advance per period and where equation 2.26 is obtained by applying the Floquet theorem and assuming a wave travelling in the + z direction having the form $I_n = I_{n-1}e^{-j\phi}$ and $I_{n+1} = I_n e^{-j\phi}$ where n represents the nth cell in an infinite array, we end up with

$$\cos \phi = 1 - \frac{C_2}{2C_1}(\omega^2/\omega_0^2 - 1) \quad (2.27)$$

then solving for ω^2 we get the dispersion relation

$$\omega = \omega_0 \sqrt{2(C_1/C_2)(1 - \cos \phi) + 1} \quad (2.28)$$

Which is plotted in Fig. 2.7. At $\phi = 0$,

$$\omega = \omega_0 \quad (2.29)$$

at $\phi = \pi$

$$\omega = \omega_0 \sqrt{4\frac{C_1}{C_2} + 1} \quad (2.30)$$

For $C_1/C_2 \ll 1$, Taylor expanding using

$$\sqrt{1+x} \approx 1 + \frac{x}{2} \quad (2.31)$$

gives

$$\omega = \omega_\pi \approx \omega_0 \left(1 + 2\frac{C_1}{C_2}\right) \quad (2.32)$$

Normalising the frequency differences to the $\phi = \pi/2$ mode at the centre, which is given by

$$\omega = \omega_{\pi/2} \approx \omega_0 \left(1 + \frac{C_1}{C_2}\right) \quad (2.33)$$

gives

$$\omega_{\pi/2}^2 = \omega_0^2 (1 + 2C_1/C_2) \quad (2.34)$$

Then

$$\omega^2 = \omega_{\pi/2}^2 (1 - k_c \cos \phi) \quad (2.35)$$

and

$$\omega \approx \omega_{\pi/2} \left(1 - \frac{k_c \cos \phi}{2}\right) \quad (2.36)$$




Figure 2.7: Plot of equation 2.28.

where $k_c = 2C_1/(C_2 + 2C_1)$ is the inter-cell coupling constant. If $k_c \ll 1$ the fractional bandwidth of the passband is $(\omega_\pi - \omega_0)/\omega_{\pi/2} \approx k_c$ which can be used to determine the coupling constant using finite element codes and the appropriate boundary conditions to determine the π and 0 modes.

2.2.6 Group Velocity

The group velocity is defined as:

$$v_g = \frac{d\omega}{dk} \quad (2.37)$$

where k is the wave number. Taking equation 2.36 and noting that the phase advance per period is $\phi = kL$, where L is the period, gives

$$\omega(k) \approx \omega_{\pi/2} \left(1 - \frac{k_c \cos(kL)}{2}\right) \quad (2.38)$$

Differentiating w.r.t k

$$v_g = \frac{d\omega(k)}{dk} = \omega_{\pi/2} \frac{k_c L}{2} \sin kL \quad (2.39)$$

This shows that as the magnitude of k_c increases, so too does the group velocity. v_g plays an important role in RF breakdown in cavities. The lower the group velocity the higher the ohmic losses, due to the increased stored energy per cell as the power flow is related to the group velocity as

$$P = v_g U \quad (2.40)$$

where P is the power flow and U is the stored energy. If the coupling constant and hence the group velocity are too low, the real part of the Poynting vector is increased as P is the integral of the Poynting vector over the cross-sectional area of the structure. This can increase the likelihood of RF breakdown as there is more power to be absorbed into arcs.

2.3 Tolerances

Manufacturing tolerances can have a significant impact on the performance of cavities used in S-band, C-band, and X-band applications. Cavities are designed to operate at specific resonant frequencies, and even slight deviations in physical dimensions or material properties can lead to shifts in these frequencies, affecting their performance. The sensitivity of these cavities to manufacturing tolerances is directly related to their operating frequency, with higher frequency cavities being more susceptible to errors.

The resonant frequency f_{mn0} of a cylindrical cell operating in the TM_{mn0} mode is given by:

$$f_{mn0} = \frac{c}{2\pi\sqrt{\epsilon_r}} \sqrt{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{p}{h}\right)^2} \quad (2.41)$$

Where, c is the speed of light, ϵ_r is the relative permittivity of the material inside the cavity, x_{mn} is the n -th root of the m -th order Bessel function, R is the radius of the cylindrical cavity, h is the height of the cylindrical cavity and p is the axial mode number (0 for TM accelerating cavity modes) [79]. Considering small changes

in the radius R and height h due to manufacturing tolerances, given by ΔR and Δh . These small variations will lead to a corresponding change in the resonant frequency, $\frac{\Delta f}{f}$.

Taking the total differential of the resonant frequency formula with respect to R and h

$$df = \frac{\partial f}{\partial R}dR + \frac{\partial f}{\partial h}dh \quad (2.42)$$

To get $\frac{\Delta f}{f}$, dividing this total differential by f , gives

$$\frac{\Delta f}{f} = \frac{1}{f} \left(\frac{\partial f}{\partial R} \Delta R + \frac{\partial f}{\partial h} \Delta h \right) \quad (2.43)$$

Taking the partial derivatives of f with respect to R and h . W.r.t R

$$\frac{\partial f}{\partial R} = \frac{\partial}{\partial R} \left(\frac{c}{2\pi\sqrt{\epsilon_r}} \sqrt{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{p}{h}\right)^2} \right) \quad (2.44)$$

Simplifying

$$\frac{\partial f}{\partial R} = -\frac{c}{2\pi\sqrt{\epsilon_r}} \cdot \frac{x_{mn}^2}{R^3} \cdot \frac{1}{\sqrt{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{p}{h}\right)^2}} \quad (2.45)$$

and w.r.t h

$$\frac{\partial f}{\partial h} = \frac{\partial}{\partial h} \left(\frac{c}{2\pi\sqrt{\epsilon_r}} \sqrt{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{p}{h}\right)^2} \right) \quad (2.46)$$

Simplifying

$$\frac{\partial f}{\partial h} = -\frac{c}{2\pi\sqrt{\epsilon_r}} \cdot \frac{p^2}{h^3} \cdot \frac{1}{\sqrt{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{p}{h}\right)^2}} \quad (2.47)$$

Now, combining the results to get $\frac{\Delta f}{f}$. Substituting the partial derivatives back into the equation for $\frac{\Delta f}{f}$

$$\frac{\Delta f}{f} = -\frac{x_{mn}^2}{R^2} \cdot \frac{\Delta R}{R} \cdot \frac{1}{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{p}{h}\right)^2} - \frac{p^2}{h^2} \cdot \frac{\Delta h}{h} \cdot \frac{1}{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{p}{h}\right)^2} \quad (2.48)$$

This expression provides the relative change in the resonant frequency due to

small changes in the radius and height of the cylindrical cavity cell. If it is assumed that the dimensional deviations ΔR and Δh are small relative to the original dimensions R and h , and keeping only first order terms, this expression can be approximated as:

$$\frac{\Delta f}{f} \approx -\frac{\Delta R}{R} \left(\frac{x_{mn}}{Rf} \right) - \frac{\Delta h}{h} \left(\frac{p}{hf} \right) \quad (2.49)$$

As the TM010 mode is usually used for accelerating cavities, this simplifies again to

$$\frac{\Delta f}{f} \approx -\frac{\Delta R}{R} \left(\frac{x_{mn}}{Rf} \right) \quad (2.50)$$

The second term $\frac{x_{mn}}{Rf}$ is a dimensionless number, with a small magnitude (on the order of 10^{-8}), leaving:

$$\frac{\Delta f}{f} \approx -\frac{\Delta R}{R} \quad (2.51)$$

Taking typical values of the radii of SW linacs at S, C and X band and calculating the frequency error for a $30 \mu\text{m}$ error on the radius, the values given in Table. 2.2 are obtained.

Frequency band	Radius (mm)	Δf for a $30 \mu\text{m}$ tolerance error (MHz)
S	33	2.7
C	21	8.2
X	10	36

Table 2.2: Estimate of the frequency error caused by a $30 \mu\text{m}$ error on the radius.

The S and C bands have manageable frequency shifts (< 10 MHz), which is on the order of typical RF source bandwidths meaning that they can be compensated for. X-band on the other hand has a frequency error an order of magnitude higher than S-band and would require a tolerance of $2.2 \mu\text{m}$ to achieve the same frequency error as S-band. Machining to this precision is costly, and requires advanced machining methods. This is part of the reason why X-band has been avoided for systems where cost is a significant factor.

Chapter 3

Single Cell Optimisation

3.1 Design Criteria

Before starting the initial design process, discussions were held with the industrial partner on the requirements for the linac design parameters for the project. The outcome of this discussion and the initial rough design specifications are listed in Table 3.1. These were chosen so that existing suppliers, infrastructure and auxiliary components could be reused.

Cavity Type	Standing wave bi-periodic $\pi/2$ structure
Frequency	C-band (5.712 GHz)
Number of cells	4-8
Beam energy	2-2.5 MeV
Pulse length	2-8 μs
e-gun voltage	~ 25 keV
Beam current	100 mA
Potential RF source	C-band multi-beam klystron (3.5 – 4 MW)

Table 3.1: Initial design specifications.

3.1.1 Standing vs. Travelling Wave Cavities

Fig. 3.1 shows a generic dispersion diagram of a periodic structure (circular waveguide with an iris placed periodically spaced by a distance, l , where n is the spatial harmonic number [63] and l is the distance between each iris. The irises are required to reduce the group velocity of the wave so that it can be synchronised with a beam. For a travelling wave structure where the ends are considered to be open, the wave flows from one end, and at every iris, some power is reflected, leading to power being transmitted to the end of the structure and some being reflected back towards the source. As the spacing between the irises becomes a multiple of half of the wavelength, the forward and reflected waves interfere and give rise to the sinusoidal shape of the dispersion curve. If, instead, closed walls are placed at the end of the structure, then the power that would leave the structure is now reflected back in, with this continuing until the cavity is filled with electromagnetic energy. The former case where the power leaves the structure is called a travelling wave (TW) linac and the latter a standing wave (SW) linac.




Figure 3.1: General dispersion curve for waveguide loaded with irises.

The two types of linac are shown in Fig. 3.2, where the difference between the filling of the linacs is shown. Choosing between the two options normally depends on the application and the constraints such as how much RF power is available, the particle energy, the length and the required pulse length or duty cycle. Travelling

wave structures are generally chosen when the structures are long ($> \approx 1.5$ m). Industrial linacs typically have a few cells that are shorter than $\beta = 1$ for bunching, and compared to TW cells, SW cells are able to achieve a higher shunt impedance for all cells, including those that have $\beta < 1$ [80]. This is due to the fact that TW structures are typically not operated in the π mode, as they require a non-zero group velocity, which reduces the acceleration efficiency. This is due to the increased number of iris' per unit length which adds extra losses. SW structures can also implement nose cones to increase the TTF, and although they can be included in TW structures they are generally less effective because of the shorter lengths, and can cause issues when trying to achieve adequate coupling between cells.

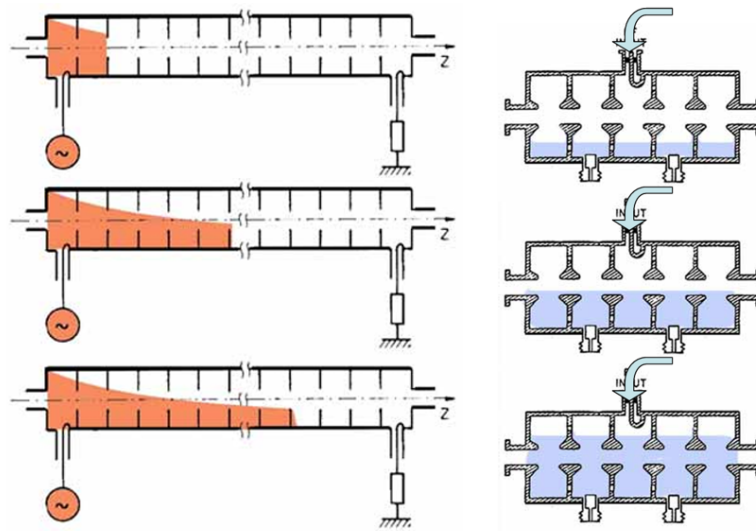


Figure 3.2: (Left) Travelling wave cavity filling vs (Right) Standing wave cavity filling. Image sourced from [81].

SW structures are, therefore, more efficient for short structures, but as they become longer, more coupling is needed to maintain field flatness, and this will lead to a drop in accelerating gradient. One of the disadvantages of SW structures is that the coupling is normally fixed which means that there can be a mismatch as the linac is filling, causing power to reflect back to the source, potentially causing damage. To mitigate this, circulators are required to allow the reflected power to be sent to a load, protecting the source but adding an extra costly component. TW structures

don't require circulators and there are examples of designs incorporating both SW and TW sections to benefit from both types (no circulator and high shunt impedance in the bunching section) [82]. TW linacs generally require a more complex design including an output coupler to absorb the remaining power and normally solenoids for RF focusing. Short SW structures used for electron acceleration generally do not require focusing as the RF defocusing term is negligible when electrons are relativistic and the natural divergence of the beam from the source has less of an effect over short distances. Despite the drawbacks of TW structures, if short filling times and higher bandwidths are required, they are the better choice. As one of the key requirements for the industrial partners was a low cost; a simpler and more compact design was preferred, leading to the selection of a SW structure.

3.1.2 Mode Selection

Typically SW structures used for high-energy physics are operated in the π -mode with field in each cell, and a 180 degree phase shift between each cell. An example of a cavity with this configuration is the proposed ILC cavity [83]. This configuration means that the beam sees the accelerating phase in each cell and the maximum efficiency is reached. Unfortunately, the π mode is located on the least steep part of the dispersion curve, as shown in Fig. 3.3, with a shorter distance to the next mode. If there are any frequency perturbations due to manufacturing defects, then this can cause problems with mode competition as multiple modes could be excited simultaneously. This is due to relationship between the bandwidth and separation between the 0 and π modes:

$$\Delta\omega = \omega_\pi - \omega_0 \approx k_c \omega_{\pi/2} \quad (3.1)$$

If the bandwidth is smaller then the likelihood of exciting unwanted modes is increased as they occupy a smaller range of frequencies. To combat this, the cavity could be made with tighter tolerances, but this drives up the cost. For industrial

and medical linacs, the cost is an important consideration in the design, so they are often designed to operate in the $\pi/2$ mode, which has the largest separation from the other modes, as it is in the centre of the dispersion curve. The nearest neighbour mode separation for the π and $\pi/2$ modes can be found as:

$$\frac{\Delta\omega_\pi}{\omega} = \frac{\omega_N - \omega_{N-1}}{\omega_{\pi/2}} \approx \frac{k_c}{2}(1 - \cos(\pi/N)) \approx \frac{k_c\pi^2}{4N^2} \quad (3.2)$$

$$\frac{\Delta\omega_{\pi/2}}{\omega} = \frac{\omega_{N/2} - \omega_{(N-1)/2}}{\omega_{\pi/2}} \approx \frac{k_c}{2}(1 - \cos(\frac{\pi}{2} - \frac{\pi}{2N})) \approx \frac{k_c\pi}{4N} \quad (3.3)$$

These equations show that the mode separation for the π mode is a factor of π/N smaller than the $\pi/2$ mode. The mode separation should be larger than the loss bandwidth (ω/Q_0) leading to [84]:


$$N_\pi < \frac{\pi}{2}\sqrt{k_c Q_0} \quad (3.4)$$

for the π mode and

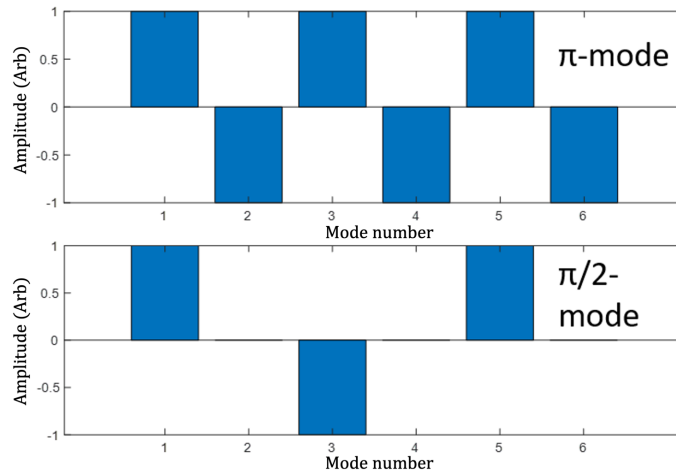
$$N_{\pi/2} < \frac{\pi}{4}k_c Q_0 \quad (3.5)$$

for the $\pi/2$ mode. This shows that the number of coupled cells is limited for structures operating in the π mode compared with the $\pi/2$ mode. Additionally, the other modes are symmetric about the $\pi/2$ mode, enhancing its resilience to frequency errors. As the phase and fields change smoothly by $\pi/2$ from cell to cell, compared with the sharp change of a π mode, this also contributes to increased stability.

When operating in the $\pi/2$ mode, every other cell is unfilled, which means that the accelerating efficiency is reduced by a factor of two as only half of the cells provide acceleration, as shown in Fig. 3.4. This is compensated for by either reducing the axial length of the unfilled cells, or by placing it off-axis on the side of the cell. Both methods allow for some of the lost efficiency to be regained. The $\pi/2$ mode was selected to provide this improved cavity stability and to reduce the effect of perturbations on the mode, allowing for looser manufacturing tolerances. For short

Figure 3.3: Dispersion curve showing 0, $\pi/2$ and π modes.

standing wave linacs there are generally two options available for coupling the cells together;

Figure 3.4: Amplitudes of π and $\pi/2$ modes.

- Coupling through cells on axis (electrically through the beam aperture or magnetically through dedicated slots).
- Coupling through cells located off-axis.

Examples of these types of cavity operating in the $\pi/2$ mode are shown in Fig. 3.5.

3.1.3 Frequency Perturbations

As well as the effects of manufacturing errors on a single cell shown in Chapter 2, frequency errors cause perturbations on the normal modes of coupled oscillators. Using the example of three magnetically coupled oscillators, the circuit equations take the form

$$LX_q = \frac{1}{\Omega_q^2}X_q \quad (3.6)$$

Where L is the matrix operator which is a function of the cavity frequency and the coupling constant, k_c , X_q is the eigenvector with elements representing the field amplitude (x_n) and $1/\Omega_q^2$ is the eigenvalue. L is given by [63]

$$L = \begin{bmatrix} 1/\omega_0^2 & k_c/\omega_0^2 & 0 \\ k_c/\omega_0^2 & 1/\omega_0^2 & k_c/2\omega_0^2 \\ 0 & k_c/\omega_0^2 & 1/\omega_0^2 \end{bmatrix} \quad (3.7)$$

and

$$X_q = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \quad (3.8)$$

The normal-mode eigenfrequencies and normalised eigenvectors for the $\pi/2$ and π modes are

$$q = 1, \pi/2 \text{ mode} : \Omega_1 = \omega_0, X_1 = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} \quad (3.9)$$

and:

$$q = 2, \pi \text{ mode} : \Omega_2 = \frac{\omega_0}{\sqrt{1 - k_c}}, X_2 = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix} \quad (3.10)$$

To perturb the solutions a perturbation matrix P is added to the unperturbed matrix L. The first order eigenvector corrections can be calculated using the series

$$\Delta X_q = \sum_{r \neq q} a_{qr} X_r \quad (3.11)$$

where X_r are the unperturbed eigenvectors, and the amplitudes a_{qr} are

$$a_{qr} = \frac{[X_q P X_r]}{\frac{1}{\Omega_r^2} - \frac{1}{\Omega_q^2}} \quad (3.12)$$

The perturbation matrix P is given by

$$L = \begin{bmatrix} -(2\delta\omega_0/\omega_0^3) & -k_c(2\delta\omega_0/\omega_0^3) & 0 \\ -k_c(2\delta\omega_1/\omega_0^3)/2 & -(2\delta\omega_1/\omega_0^3) & -k_c(2\delta\omega_1/\omega_0^3)/2 \\ 0 & k_c(2\delta\omega_0/\omega_0^3) & (2\delta\omega_0/\omega_0^3) \end{bmatrix} \quad (3.13)$$

Where k_c is the coupling constant, ω_0 is the resonant frequency of the uncoupled oscillators and $\delta_{1,2}$ are the error parameters. Choosing ω_0 as the unperturbed frequency averaged over the end-cell frequencies, $\pm\delta\omega_0$ as the frequency error of the end cells relative to ω_0 , $\delta\omega_1$ as the frequency error of the middle cell, then the results for the $\pi/2$ and π modes perturbed by frequency errors are

For the $\pi/2$ mode, valid through second order

$$\Omega_1 = \frac{\omega_0}{\sqrt{1 - 4(\delta\omega_0/\omega_0)^2}} \quad (3.14)$$

$$X_1 = \begin{bmatrix} 1 + \frac{4}{k_c^2} \frac{\delta\omega_1}{\omega_0} \frac{\delta\omega_0}{\omega_0} - \frac{2}{k_c^2} \left(\frac{\delta\omega_0}{\omega_0}\right)^2 \\ -\frac{2}{k} \frac{\delta\omega_0}{\omega_0} \\ -1 + \frac{4}{k_c^2} \frac{\delta\omega_1}{\omega_0} \frac{\delta\omega_0}{\omega_0} + \frac{2}{k_c^2} \left(\frac{\delta\omega_0}{\omega_0}\right)^2 \end{bmatrix} \quad (3.15)$$

and for the the π mode, valid through first order

$$\Omega_2 = \frac{\omega_0}{\sqrt{1 - k_c}} \sqrt{1 + \frac{\delta\omega_1}{\omega_0}} \quad (3.16)$$

$$X_1 = \begin{bmatrix} 1 - \frac{1-k_c}{2k_c} \left(\frac{\delta\omega_1}{\omega_0} - 4 \frac{\delta\omega_0}{\omega_0} \right) \\ -1 - \frac{1-k_c}{2k_c} \frac{\delta\omega_1}{\omega_0} \\ 1 - \frac{1-k_c}{2k_c} \left(\frac{\delta\omega_1}{\omega_0} + 4 \frac{\delta\omega_0}{\omega_0} \right) \end{bmatrix} \quad (3.17)$$

These show that frequency errors cause non-uniformity in the cells of the structure in proportion to the fractional error in the frequencies, but that the effect can be reduced by increasing k_c . First order corrections for the π mode show that the field distributions are sensitive to frequency errors of each oscillator whereas for the $\pi/2$ mode, only the centre cell, which is unexcited is affected to first order. The effects on the excited cell are only second order effects. This shows that the excited cells in the $\pi/2$ mode are insensitive to small frequency errors, and provides justification for the selection of this mode to reduce the required tolerances of the structure.

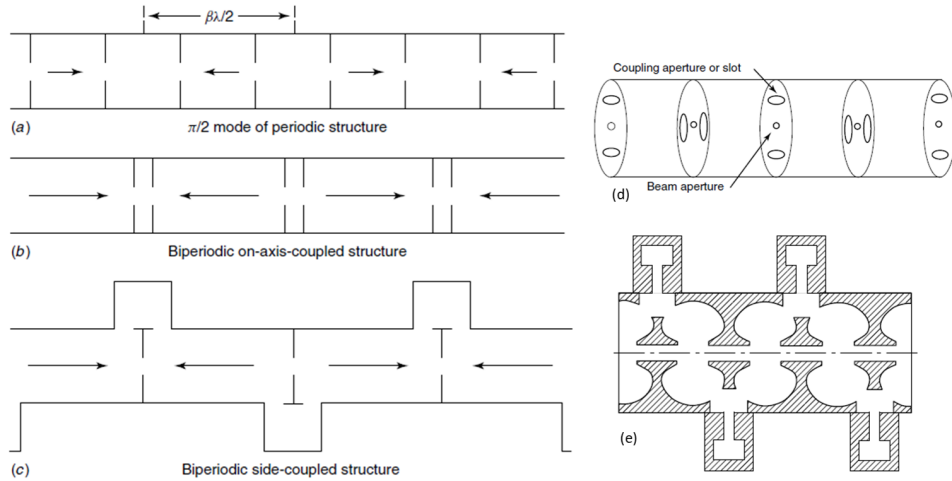


Figure 3.5: Bi-periodic vs side coupled linacs. a.) Bi-periodic structure with no adjustments of coupling cell, b.) Reduced length of coupling cell to increase Z , c.) Coupling cells moved off axis, again increasing space for acceleration, d.) Example of coupling slots that couple the cavities together using the magnetic field, e.) Example of side coupled structure [63].

One of the main drivers of the weight of the linacs used for security applications is the X-ray shielding that must cover the linac, as shown in Fig. 3.6. Utilising a bi-periodic linac design ensures that the transverse size of the linac is kept as small as possible, facilitating a decrease in shielding weight. This reduction in weight ultimately leads to a decrease in overall size and cost, thereby enhancing the potential for the device to be portable. For this reason, a bi-periodic cavity type was chosen for the design.




Figure 3.6: Linac example showing shielding to protect from radiation and the formula for calculating the area of the shielding shown in grey, highlighting the dependence on transverse size. A is the area, r is maximum radius of the cell and t is required shielding thickness.

3.1.4 Initial Design Ideas

The first step taken in the design process was to estimate the number of cells required. This initial estimate was made using the approximate parameters given in Table 3.2. It was known that at least one bunching cell with $\beta < 1$ would be required, as the electrons coming from the gun would not be relativistic. The assumed design energy was 2 MeV and an initial estimate of one $\beta = 0.5$ cell, one $\beta = 0.7$ cell, and three $\beta = 1$ cells were used based on the knowledge that it is challenging to design cells with $\beta < 0.5$, as shall be seen in later sections, and that the second cell should be somewhere in between $\beta = 0.5$ and $\beta = 1$. The total gain of 2 MeV was divided by an estimate of the TTF and then distributed among the cavities

proportionally to their length. This gave an estimate of the gradient required in $\beta = 1$ cells of ~ 26.3 MV/m. An initial estimate of the aperture radius (where the beam passes through) was chosen as 5 mm, based on research using spline geometries to create optimum standing wave cavity geometries that maximise shunt impedance for different frequency bands shown in Fig. 3.7 and discussed in [85].

Cell no.	β	$\beta\lambda/2$ (mm)	% Total L	Energy gain (MeV)	Gradient (MV/m)
1	0.5	13.13	12	0.24	13.1
2	0.7	18.38	16.7	0.33	18.4
3	1	26.26	23.8	0.47	26.3
4	1	26.26	23.8	0.47	26.3
5	1	26.26	23.8	0.47	26.3
Totals		110	100	1.972	

Table 3.2: Initial gradient estimates.




Figure 3.7: Optimum beam aperture vs shunt impedance used for C-band aperture selection, where shaded region represents the aperture radii where C-band performs best. Image sourced from [85].

3.2 Single Cell Design

The next step was to design a single C-band $\beta = 1$ cell, to evaluate the expected values of the shunt impedance and coupling constant. A single cell geometry was established using non-uniform rational basis splines (NURBS) based on recent work

on novel modelling methods for RF structures [86]. This method uses a number of control points and passes a spline through them to provide fine control of the cavity geometry, as shown in Fig. 3.8. These control points were related to parameters that can be easily identified by looking at cavity geometry, such as a minimum nose cone gap cg_{min} and a minimum nose cone radius that could be feasibly manufactured R_{min} . This method allows for much greater control over the cavity geometry as compared to using standard geometric shapes such as circular blends or straight sections with angles that allow for an improved shunt impedance for the same peak fields as described in [87]. Initial estimates of the thickness of the wall, t , and the length of the coupling cavity, L_c , were made, with both set at 2 mm.




Figure 3.8: NURBS spline geometry showing control points used to optimise cell.

3.2.1 Multi-Objective Optimisation

Historically, cavity geometries have been optimised using a single-objective optimisation approach, whereby one may try to maximise the shunt impedance while keeping the value of E_{pk}/E_{acc} , below a certain value, for example, < 2 . Although this method can work well, it requires prior knowledge of the maximum limits on the peak fields and will only provide a single solution and one optimised cavity geometry. If the

design requirements are changed slightly or new information is available regarding the maximum allowable field limits, then the optimisation process would need to be repeated with these new limits implemented. Instead, another option is to use multi-objective optimisation methods which are generally employed when the optimisation problem has two or more objectives that are in conflict with one another. These types of problem can be formulated as follows:

$$\begin{aligned}
 & \min/\max f_m(x), \quad m = 1, 2, \dots, M, \\
 & \text{subject to } g_j(x) \geq 0, \quad j = 1, 2, \dots, J, \\
 & x_{ilb} \leq x_i \leq x_{iub}, \quad i = 1, 2, \dots, n
 \end{aligned}
 \tag{3.18}$$

Where f_m are the objectives, g_j are the constraints, and x_i are the upper and lower bounds of the input parameters. Instead of obtaining a single solution to the problem, a set of solutions is found that defines the trade-offs between the objectives in an N-dimensional objective space. A 2D example is shown in Fig. 3.9 where the red curve shows the Pareto frontier that defines the trade-off between the objectives f_1 and f_2 when they are both to be minimised. This set of solutions are defined as Pareto optimal solutions, and they make up a boundary in the objective space where no objectives can be improved without sacrificing at least one of the other objectives.




Figure 3.9: Pareto optimal points in 2D. [87]

Another method that can be used is the weighted sum approach, as shown in Fig. 3.10. This method involves a linear combination of objectives with a weight applied to each of them. This method can also work well if the objective requirements for a problem are known, as the optimisation process for a single design will be faster. The weighted sum method also suffers though, when the trade-offs are required or the design needs to be changed, and again the whole optimisation must be repeated. MO methods were selected for the cavity optimisation in order to provide the ability to change the design and accurately assess the trade-offs for the cavity.

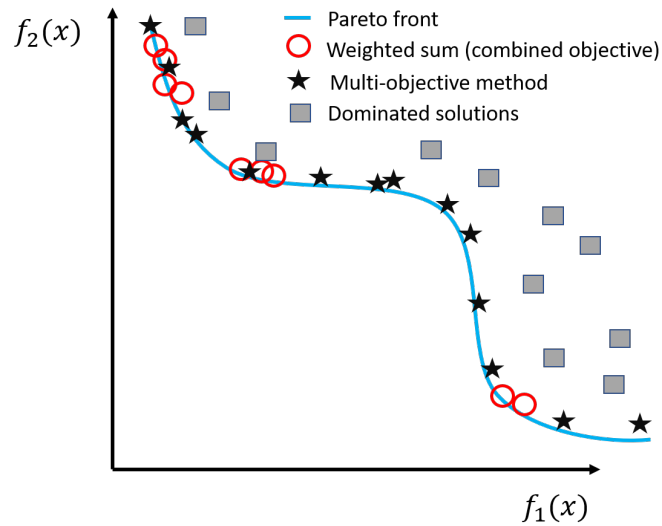


Figure 3.10: Weighted sum vs. multi-objective methods. [87]

There are several techniques that can be used to solve MO problems, but the ones selected for the cavity optimisation were genetic algorithms. These algorithms are inspired by biology and the process of evolution. The flow diagram shown in Fig. 3.11 details how they work for a general case. An initial random population of individuals is generated first, using a Latin hypercube sampling method. A fitness value is then calculated based on the objectives and weightings chosen. Individuals with the best fitness are used for operations such as crossover, where parts of each fit individual are combined to create new offspring, and mutation, where genetic diversity is introduced into the population by modifying some of the individuals' properties to ensure the algorithm does not get stuck in a local minimum. The new

offspring are used to replace less fit individuals, and the process is repeated until some objective criteria have been satisfied or a set number of generations have been reached.

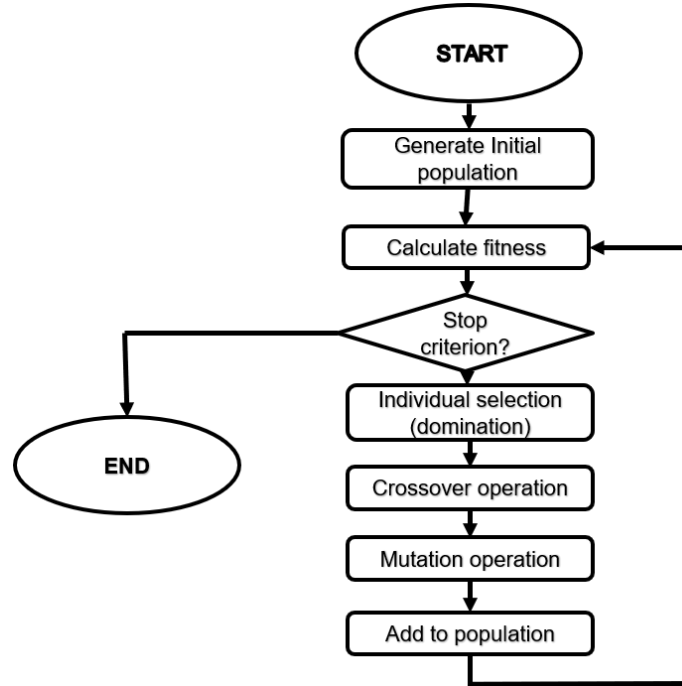


Figure 3.11: Generic GA flow chart.

The four objectives used for the optimization are shown below:

$$\text{Objectives} = \left\{ \frac{E_{pk}}{E_{acc}}, \frac{B_{pk}}{E_{acc}}, \frac{\sqrt{S_c}}{E_{acc}}, Z \right\} \quad (3.19)$$

Where the peak fields are to be minimized and the shunt impedance maximized. For the optimization, the eigenmode solver of CST Studio Suite [88] was used, along with Isight, a software used for design process flow [89].

3.2.2 Algorithms

The two main algorithms used for the optimisation were the well-known Fast Non-dominated Sorting Genetic Algorithm (NSGA-II) [90], and the Archive-Based Micro Genetic Algorithm (AMGA) which offers an improvement over NSGA-II by using

an archive to maintain its search history [91]. The parameters used for a typical optimisation run are shown in Table 3.3.

Parameter	NSGA-II	AMGA
Initial size	60	2000
Population size	80	40
No. function evaluations	20000	20000
Archive size limit	-	2000
Crossover probability	0.9	0.9
Mutation probability	-	1/n
Crossover distribution index (CDI)	9	9
Mutation distribution index (MDI)	18	18
No. generations	250	-

Table 3.3: Typical GA parameters used for optimisation.

The initial size is the size of the randomly generated population pool that samples the input space, and the population size is the population size for each subsequent generation. The crossover probability controls the likelihood that a crossover will occur between two parent solutions; it is generally set high so that a crossover will occur between most parent solutions, and a recommended initial value is 0.9, though this can be changed to 1 to always allow for crossover [89]. The mutation probability determines the likelihood that an individual in the set will mutate and is normally set low so that few individuals mutate in any generation. If it is set higher, then more of the search space will be explored, but the chances of the algorithm getting stuck and a sub-optimal solution are increased. The Crossover Distribution and Mutation Distribution Indices are inversely proportional to the changes in the design variables - a smaller value will reduce the chance of converging too quickly at the cost of a narrow search. The AMGA algorithm holds an archive of 2000 points, which is updated each generation using the following steps:

1. Begin optimisation
2. Generate initial population
3. Evaluate solutions of initial population
4. Update external archive population using initial population

5. Begin main loop
 - (a) Create parent population from the archive
 - (b) Create the mating pool from the parent population
 - (c) Create offspring population from the mating pool
 - (d) Evaluate solutions of the offspring population
 - (e) Update external archive
6. Repeat until either convergence criteria is met, or specified number of evaluations
7. End optimisation

3.2.3 Optimization Results

The 2D optimisation results are shown in Fig. 3.12 along with red lines showing the approximate field limits that were calculated previously. Two optimisations were performed, one for a $\beta = 1$ cell and one for a $\beta = 0.5$ cell. It can be easily seen that nearly all individuals were under the limit for B_{pk}/E_{acc} as there were no geometric features that would enhance the magnetic field, such as coupling slots. It was also apparent that the structure would probably be mainly limited by E_{pk}/E_{acc} as the majority of the points had values of $S_c/E_{acc} < 0.04 \sqrt{A/V}$.

Quantity	Limit	Limit at 25 MV/m
Peak E-field	100 MV/m	$E_{pk}/E_{acc} = 4$
Peak S_c	$1 \text{ W}/\mu\text{m}^2$	$\sqrt{S_c}/E_{acc} = 0.04$
Max surface temperature rise	20 K	-
Peak B field (from 20 K surface rise)	229 mT	$B_{pk}/E_{acc} = 9.16$

Table 3.4: Initial design parameters.

It should be noted that Z is reduced by almost a factor of 4 for the $\beta = 0.5$ cells, which means that they would contribute very little to the acceleration. Various other options for visualising the Pareto fronts were explored, with some examples shown in Fig. 3.14 and Fig. 3.15. The first was created by first binning the points into bins

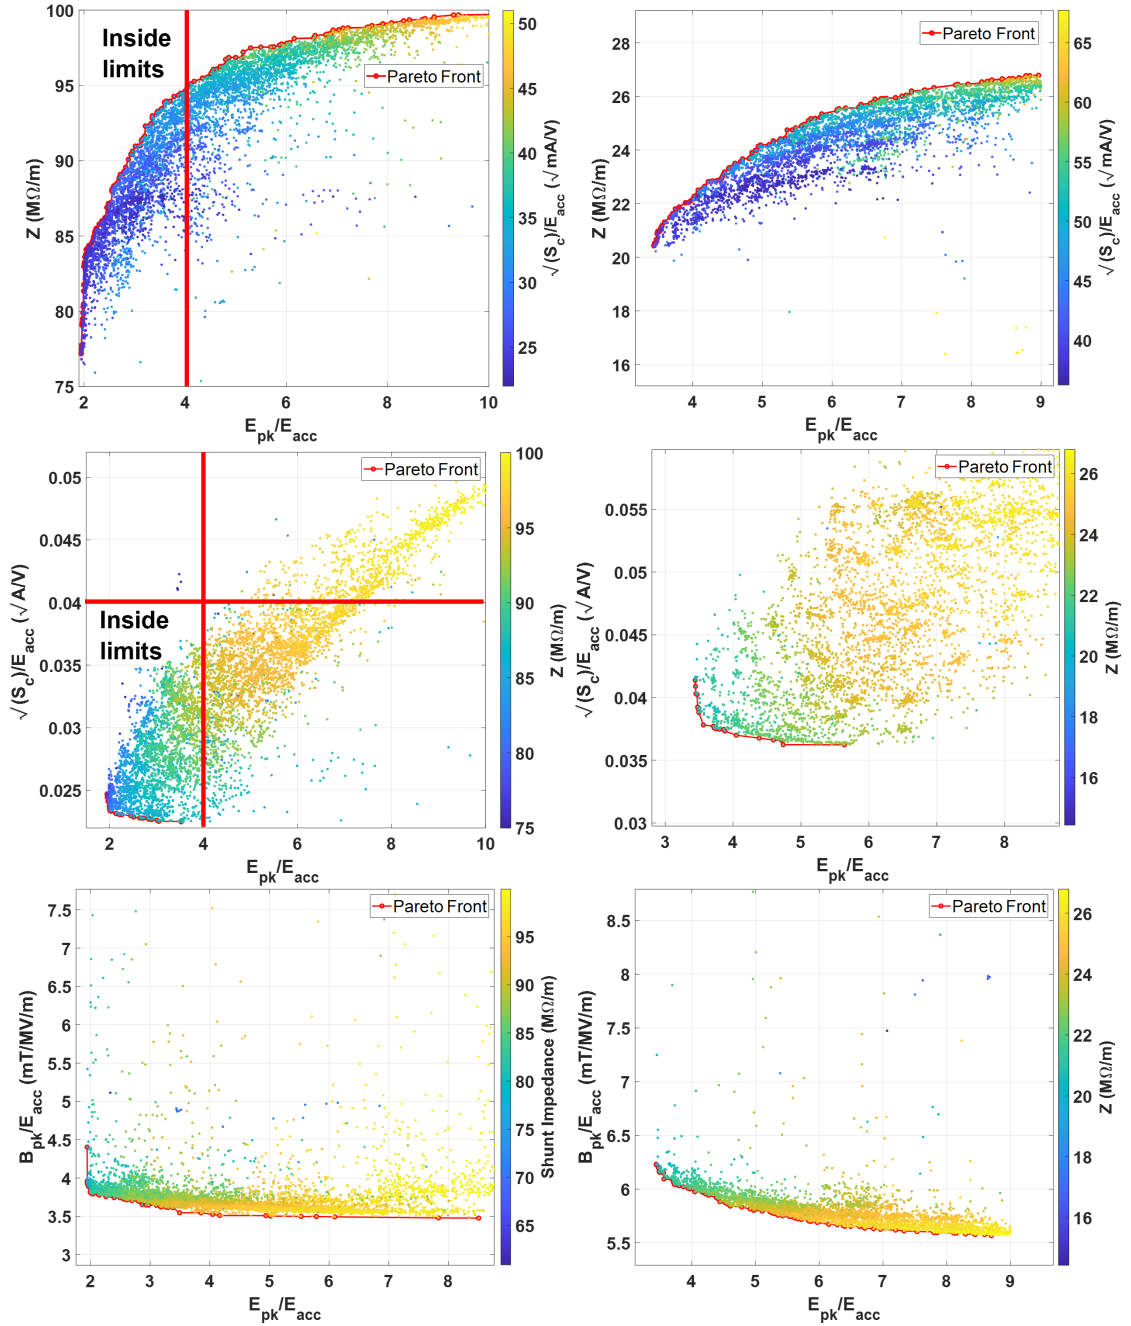


Figure 3.12: 2D Pareto front results for $\beta = 1$ (left), top to bottom (Z vs. E_{pk}/E_{acc} , S_c/E_{acc} vs. E_{pk}/E_{acc} , B_{pk}/E_{acc} vs. E_{pk}/E_{acc}) and $\beta = 0.5$ (right), top to bottom (Z vs. E_{pk}/E_{acc} , S_c/E_{acc} vs. E_{pk}/E_{acc} , B_{pk}/E_{acc} vs. E_{pk}/E_{acc}).

of $1 \text{ M}\Omega/m$ for the shunt impedance and then finding the 2D Pareto fronts for each bin. A representation of the 3D Pareto front can then be constructed showing all of the objective values, allowing one to easily see how the objectives interact and also

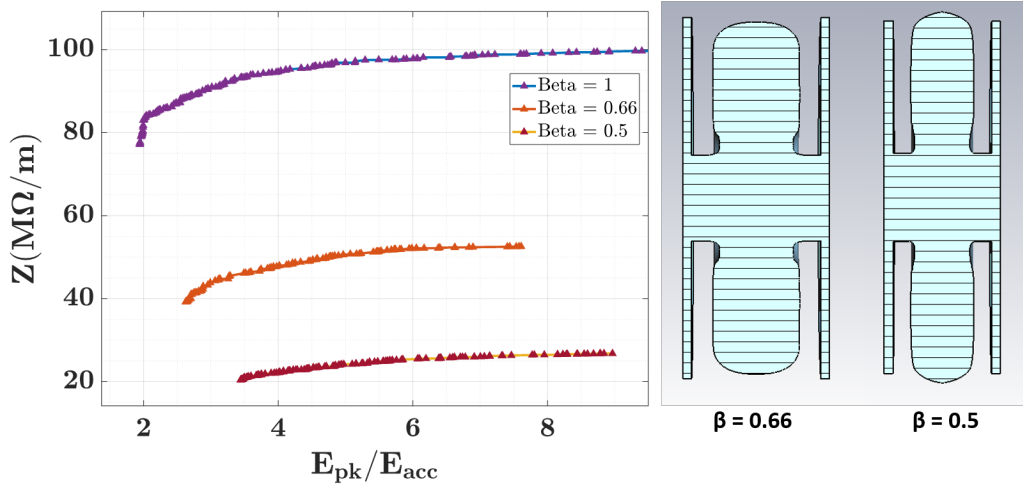


Figure 3.13: Z vs. E_{pk}/E_{acc} for three values of β .




Figure 3.14: Pareto optimal points in 3D.

allowing for potential final solutions to be selected. The fronts are also projected onto the corresponding axis, clearly showing the individual 2D Pareto fronts. The other shows the 2D Pareto front for Z vs. E_{pk}/E_{acc} whilst also showing the various cavity geometries that are generated. This allows for an understanding of what types of geometric features lead to which objective values. This led to some interesting




Figure 3.15: Pareto front of Z vs. E_{pk}/E_{acc} , with the outer colour of each point showing the value of $\sqrt{S_c}/E_{acc}$ and inner colour the value of B_{pk}/E_{acc} . A cross section of the types of solutions that fall on the Pareto front are also shown, highlighting the interesting shapes at lower values of E_{pk}/E_{acc} .

cavity geometries that may not have been considered otherwise, especially at low values of E_{pk}/E_{acc} (< 3). An extensive treatment of visualisation techniques for Pareto optimal sets specifically for RF cavities written by the author and colleagues is given in [87].

3.3 Algorithm Studies

A number of simulations were performed with different settings applied to the algorithms to ensure that the convergence and frequency error criteria were met. A comparison was also drawn between the two methods of modelling the structure (using splines and using standard shapes) to show the efficacy of the spline modelling approach.

3.3.1 Frequency Handling

For each simulation of the $\beta = 1$ cavity, the frequency of the TM_{010} accelerating mode had to be kept as constant as possible. Due to inherent numerical noise and meshing errors, it is difficult to keep exactly constant. This was dealt with by defining a value ϵ , such that for the simulations:

$$f_{TM_{010}} = f_{desired} \pm \epsilon \quad (3.20)$$

Typical values for ϵ of 5-30 MHz were used, with 20 MHz set for the final value. The frequency of the chosen solutions could then be tuned to the exact value, which was found to have a minimal effect on the final field values, as changing the outer radius does not affect the fields significantly. Figure 3.16 illustrates the viable data points within the Z vs. E_{pk}/E_{acc} 2D objective space, subject to a total frequency constraint of 60 MHz. To assess the influence of the frequency constraint on the estimated Pareto frontier, ϵ was systematically reduced from 30 MHz to 1 MHz, with the Pareto frontier found for each data set. The results of this analysis are shown in Fig. 3.17.




Figure 3.16: Frequency variance in combined Pareto optimal set.

Notably, at higher Z values, a more pronounced shift in the Pareto frontier was seen because most of the points there had a lower frequency and, therefore, a smaller radius, which increases the shunt impedance. As ϵ is decreased to 1 MHz, a larger shift was observed in the frontier, caused by increased noise due to the reduced number of feasible solutions from the total set. This could be

improved by performing a second set of simulations which are seeded with solutions with frequencies closer to 5.712 GHz, although as the cells could be tuned without effecting the objective parameters severely, this was not required. Since the time needed for convergence is directly tied to the ϵ value, opting for $\epsilon = 20$ MHz represented a compromise between simulation time and accuracy.

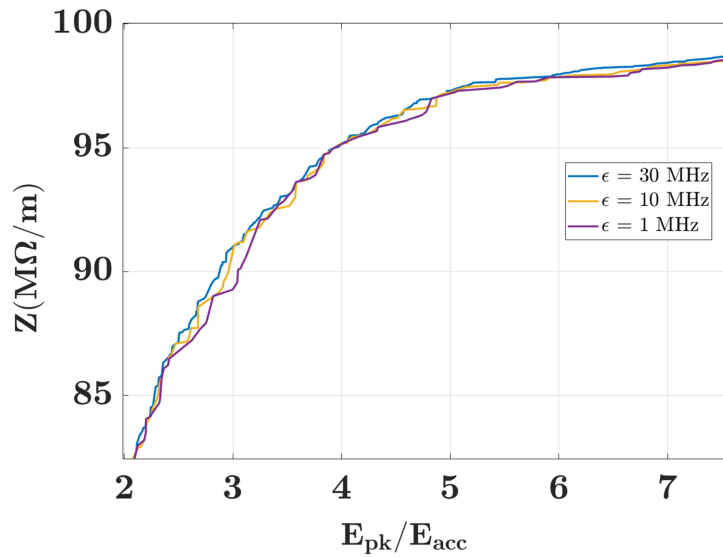


Figure 3.17: Variation in 2D Pareto front for varying ϵ (defined in equation 3.20).

3.3.2 Convergence

In Figure 3.18, a 2D Pareto front is shown, revealing the inherent trade-off relationship between Z and the ratio E_{pk}/E_{acc} . The true Pareto frontier is not known, but, nonetheless, insights can be gained into the convergence behaviour by plotting the current Pareto front against the generation number for a specific set of parameters and simulation run.

The figure shows that within the range of values E_{pk}/E_{acc} spanning 2.5 to 5, the Pareto front effectively converges after around 125 generations. Subsequent generations tend to increase the dispersion of the front, primarily at the extremities. In particular, the difference in Pareto fronts between 200 and 250 generations is

minimal, with the majority of data points aligning closely on the same curve. This suggests that, given the specific simulation parameters and constraints at hand, there is a limited advantage in extending the simulation for additional generations.




Figure 3.18: Convergence of the Z vs E_{pk}/E_a Pareto front.

3.3.3 Geometric Comparison

An optimisation was also performed using a cavity setup that used only circular blends and straight sections to form the geometry, which allowed a comparison to be made between the two modelling methods as shown in Fig. 3.19. The Z vs. E_{pk}/E_{acc} Pareto fronts are shown for both optimisations, and it was clear that the spline modelling method offers an improvement in the shunt impedance. This is especially pronounced at lower values of E_{pk}/E_{acc} where there is a drastic improvement in Z , due to the flexible nature of the geometry. This improvement is more modest at higher values of E_{pk}/E_{acc} , and at very high values there is little improvement, as the model with standard shapes is able to produce similar shapes to the spline when

the nose cone is large.




Figure 3.19: Comparison of Z vs. E_{pk}/E_{acc} Pareto front for s NURBS geometry and standard shape model.

3.4 Cell-Cell Coupling

Before a potential design could be selected, the cell-cell coupling had to be checked to ensure that there was adequate coupling between the cells. First, the required coupling was estimated analytically using the following formula, for which a full derivation can be found in [63]:

$$X_{2n} \cong (-1)^{n-m} X_{2m} \left[1 - \frac{2(m^2 - n^2)}{k_c^2 Q_a Q_c} + j \frac{4(m^2 - n^2) \delta\omega}{k_c^2 Q_a \omega_a} \right] \quad (3.21)$$

where X_{2n} is the field in accelerating cavity $2n$, X_{2m} is the field in the drive cavity, k_c is the coupling coefficient, Q_a and Q_c are the quality factors of the accelerating cavity and coupling cavity respectively, $\delta\omega$ is the frequency difference between the coupling and accelerating cavities, ω_a is the accelerating cavity frequency and m/n are used

to number the cavities. Equation 3.21 shows that there is an amplitude and a phase difference between the cells, both of which are affected by the coupling constant, while only the phase difference is affected by $\delta\omega$. An approximate frequency shift can be estimated by first estimating the manufacturing tolerance on the cavity radius, as this has a significant effect on the cell frequency. A typical manufacturer might be able to provide a $30 \mu m$ tolerance, so this was used as an initial estimate. The frequency shift due to this tolerance, as shown in Chapter 2, can be estimated by using:

$$\frac{\Delta f}{f} \approx \frac{\Delta R_d}{R_d} \quad (3.22)$$

Where R_d is the cavity radius dimension and ΔR_d is the error on this dimension. If the centre cell is assumed to be the 'drive' cell where power will be coupled into the cavity as shown in Fig. 3.20 and assume an acceptable phase difference of 2 degrees from the end cell to the centre, then equation 3.9 can be rearranged and solved for the required coupling constant. This can be plotted for an increasing number of cavities and therefore, length, as shown in Fig. 3.21, revealing how the coupling must be increased as the number of cavities increases.




Figure 3.20: Schematic representation of cell-cell coupling calculation showing drive cell.

The length of each cell was set to $\beta\lambda/2$, so 52.48 mm for $\beta = 1$. Fig. 3.22 shows the coupling requirements for various acceptable phase differences and tolerances of $30 \mu m$ and $50 \mu m$. Increasing the tolerance to $50 \mu m$ would provide an increased resilience to manufacturing errors and possibly cheaper manufacturing costs. Eventually, it was decided that a larger phase error of 3-4 degrees would be

acceptable due to the fact that no synchronisation with an existing beam is required, leading to an approximate coupling constant requirement of 0.9 %.




Figure 3.21: Coupling vs structure length for a 2 degree phase shift along the structure, a 1 % amplitude variation and various values of $d\omega/\omega_a$, which is the ratio of the frequency difference between coupling cells to the accelerating mode frequency.

3.4.1 Coupling Improvement

After the MO optimisation, a number of cavity candidates were selected, and the coupling constants found using the eigenmode solver in CST studio suite. First, the frequency of the $\pi/2$ mode in the accelerating cell is optimised to 5.712 GHz using magnetic boundaries. Periodic boundaries with a 180° phase shift are then used to optimise the second $\pi/2$ mode in the coupling cells to 5.712 GHz. Finally, the phase shift is set to 0 to find the 0-mode and the π -mode so that the coupling constant can be calculated. It was found that in the regions of interest on the Pareto front ($E_{pk}/E_{acc} \approx 4 - 4.5$), the coupling was only around 0.6 % - 0.8 % and therefore not sufficient. One option considered to increase the coupling was to increase the size of the aperture radius, although this would also reduce the shunt impedance, which was undesirable.

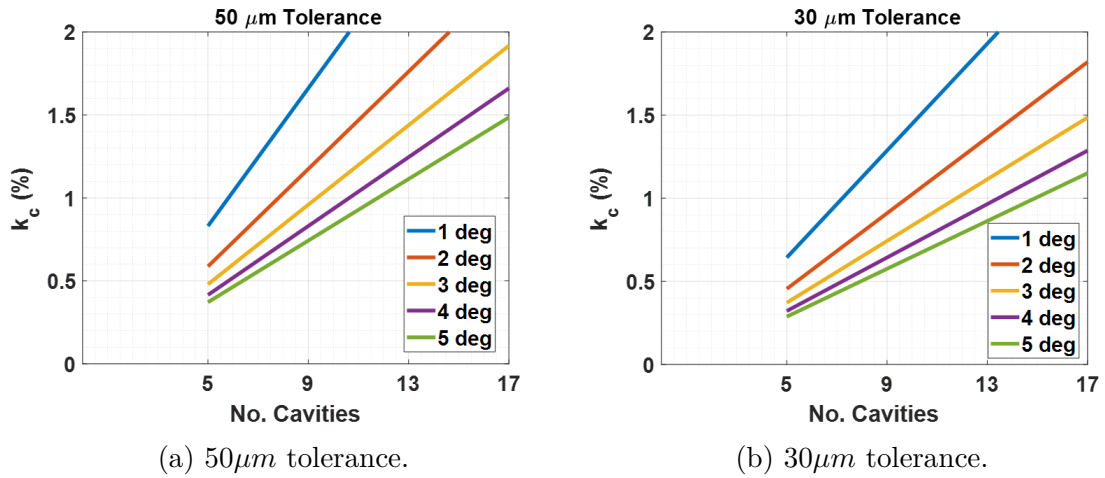


Figure 3.22: Required coupling vs. acceptable tolerance for different values of the acceptable phase shift across the structure.




Figure 3.23: Two options for increasing the coupling between the coupling cells and accelerating cells. (Left) Chamfering (Right) Blending.

Another option would be to choose a candidate with a shorter nose cone and, hence, a larger coupling, although again, the shunt impedance would be reduced. Some other methods were investigated to try and increase the coupling without affecting Z significantly; these are shown in Fig. 3.24. One approach was to include a chamfer, and the other was to add a rounding on the outer edge where the beam pipe connects to the coupling cavity. The two options are shown in more detail in Fig. 3.23. The results of parameter scans using these methods are shown in Fig. 3.25 and Fig. 3.26. It was found that the best option for improving the coupling is to have a chamfer that extends to the edge of the nose cone and has a moderately

shallow angle; for example, with a chamfer length of 2.7 mm and an angle of 20 degrees, the coupling can be improved by 20 % with only a 0.5 % reduction in Z and little effect on the peak fields.




Figure 3.24: From left to right, blend option, chamfer option, baseline cavity for scans with $k_c = 0.771$ %.






Figure 3.25: Effect of changing the blend radius on coupling constant and shunt impedance.


To accurately calculate the coupling coefficient, a number of steps and simulations are required, which made it unfeasible to calculate the coupling coefficient for each optimisation point. Instead, a proxy parameter was chosen to emulate the coupling: this was the furthest distance that one of the nose cone spline points stretches into the cavity, similar to cg_{min} in Fig. 3.8. For some of the candidates, simulations were performed to obtain the real values of the coupling, which were then plotted against the coupling proxy with the results shown in Fig. 3.27, showing the effectiveness of the coupling proxy. It was evident that the coupling proxy effectively



(a) 1 mm chamfer.



(b) 2 mm chamfer.



(c) 2.7 mm chamfer.

Figure 3.26: Z and coupling constant vs. chamfer angle for three different chamfer lengths. (Top) 1 mm (Middle) 2 mm (Bottom) 2.7 mm.

estimated the coupling constant with the chamfers around the region of interest, as most points closely aligned with a linear fit. This meant that in order to achieve the required 0.9 % coupling, the coupling proxy would need to be $> \approx 8.5$ mm.




Figure 3.27: Effectiveness of coupling proxy.

The optimisation was run again, with the chamfers included, and the coupling proxy recorded for each run. The 2D Pareto front for Z vs. E_{pk}/E_{acc} , is shown in Fig. 3.28. From this plot it can be seen that most solutions with $Z > 90 M\Omega/m$ have a coupling proxy that is too small, meaning that the coupling would not be sufficient. Fig. 3.29 shows the solutions in the S_c/E_{acc} vs. E_{pk}/E_{acc} 2D objective space, candidates for a final solution are circled in red, which are those with a coupling proxy value > 8.4 , below the E_{pk}/E_{acc} limit of 4 and with a high shunt impedance.

3.4.2 Effect of Aperture Radius and Wall Thickness

The next step was to investigate the effect of changing the aperture radius and wall thickness on the coupling coefficient and shunt impedance to see if further improvements could be made to the design, and to know how they change in case the manufacturing tolerances or peak field limits needed to be adjusted. Three potential cavity candidates were selected and the aperture radius was scanned, with the coupling coefficient calculated for each. One cavity candidate was selected for




Figure 3.28: Z vs. E_{pk}/E_{acc} , S_c/E_{acc} Pareto front showing the coupling proxy value and the estimated limit for E_{pk}/E_{acc} .




Figure 3.29: S_c/E_{acc} vs. E_{pk}/E_{acc} solutions with coupling proxy shown in colour. Pareto fronts are shown for binned values of Z . Potential final solution candidates circled in red.

high shunt impedance, one for low peak electric fields, and one a compromise between the two. The results are shown in Fig. 3.30. A linear relationship was found between the aperture radius and both the coupling and the shunt impedance for small changes in the aperture radius with a $\approx 0.35\%$ increase in coupling and $5\text{ M}\Omega/\text{m}$ decrease in Z for an increase of 0.5 mm in the aperture radius. For the wall thickness scans, the MO optimisation was repeated for three different wall thicknesses, along with checking the effects of the wall thickness on a cavity candidate that already had the chamfer included. These results are shown in Fig. 3.31 and Table 3.5. Increasing the wall thickness by 1 mm reduces the coupling by $\approx 0.2\%$ and reduces the shunt impedance by $\approx 5\text{--}10\text{ M}\Omega/\text{m}$ at values of $E_{pk}/E_{acc} = 4$, so a compromise had to be made between maximising Z and achieving the required coupling. It can also be seen that changing the wall thickness has little effect on the peak fields, meaning that the key trade-off was between the coupling constant and Z .

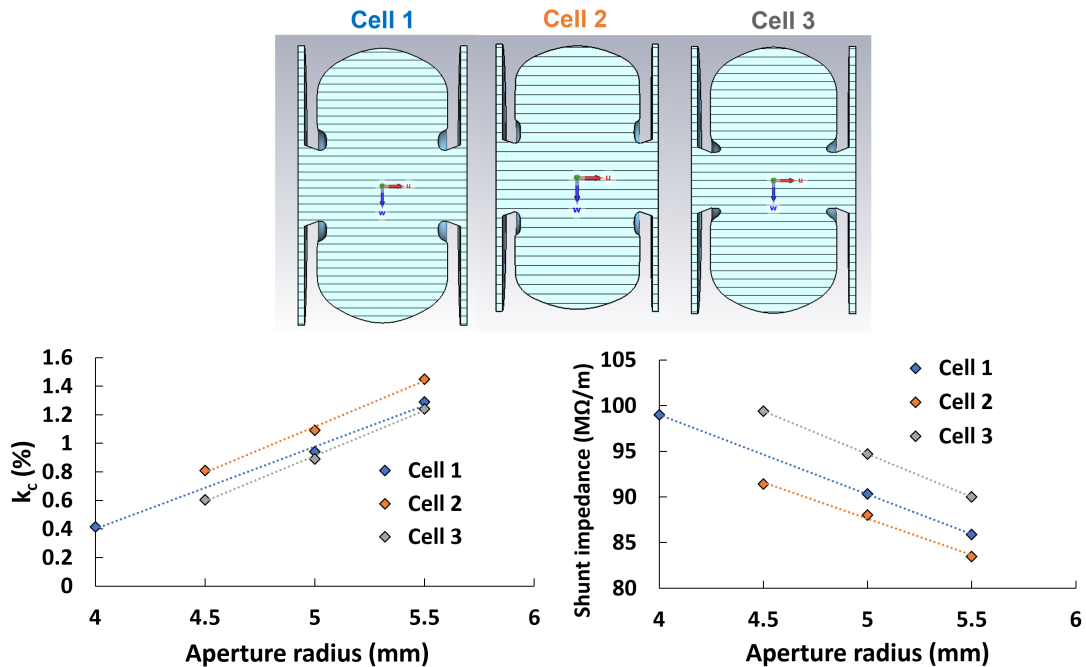


Figure 3.30: Cells used for aperture radius investigation (top), plots showing the effect of the aperture radius on k_c (bottom left) Z and (bottom right).

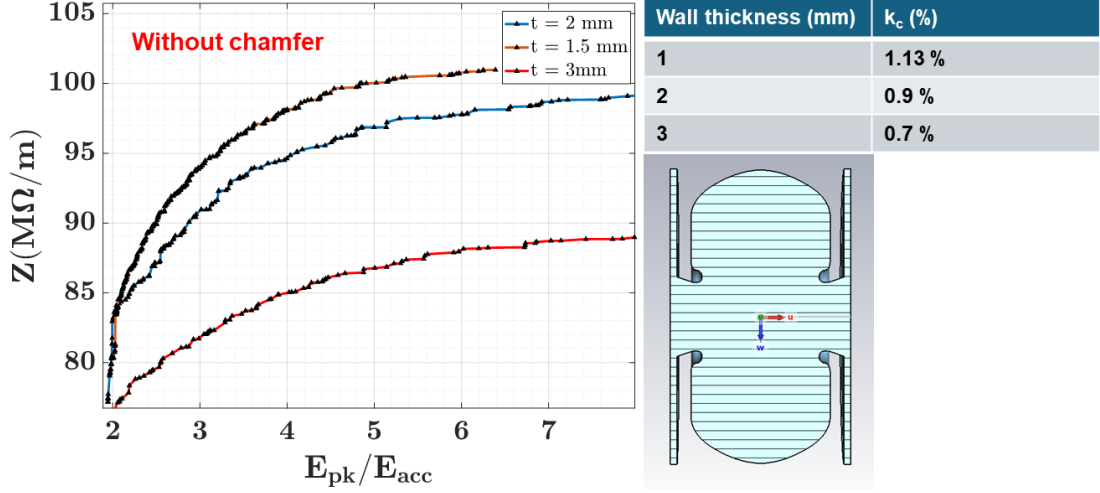


Figure 3.31: Effect of wall thickness on the Z vs. E_{pk}/E_{acc} Pareto fronts (left), cavity candidate used for investigating the effect of wall thickness on k_c (right).

Wall thickness (mm)	E_{pk}/E_{acc}	B_{pk}/E_{acc}	$\sqrt{S_c}/E_{acc}$	Z	k_c (%)
1	3.89	3.51	0.036	98.9	1.13
2	3.78	3.58	0.035	93.6	0.9
3	3.72	3.83	0.034	86.1	0.7

Table 3.5: Effect of wall thickness on various FOMs and k_c .

3.5 Final Cell Selection

The optimisation process was repeated for a wall thickness of 2.5 mm, for thermal dissipation reasons discussed in Chapter 6, while also recording the coupling proxy. The 2D Pareto front showing Z vs. E_{pk}/E_{acc} is shown in Fig. 3.28. The final $\beta = 1$ cell was selected by finding the solution with the highest value of Z , while still ensuring the coupling constant requirement ($k_c = 0.9$ %) and electric field limit ($E_{pk}/E_a < 4$) were met.

The final $\beta = 1$ cell choice shown in Fig. 3.32, showing the absolute value of the electric field, modified Poynting vector, and magnetic field. All three fields were recorded along a curve starting at the centre of the cavity and extending to the end of the nose cone shown in Fig. 3.33. The final design parameters for the single $\beta = 1$ cell are shown in Table 3.6. The final chamfer length was 2.8 mm with a 25°

chamfer angle, with a cell radius of 21.38 mm and a coupling cell radius of 21.48 mm. Without the chamfer the coupling was found to be 0.7 %, which meant a 28.6 % increase to 0.9 %. The shunt impedance was reduced by only 1.3 % from 90.58 $M\Omega/m$ to 89.4 $M\Omega/m$. An expanded view of the fields in the region of the coupling chamfer is shown in Fig. 3.34, where no field enhancement was observed, and there is little effect on the electric field distribution.




Figure 3.32: Final $\beta = 1$ cavity design showing fields from left to right (Cross-section, Electric field, modified Poynting vector, magnetic field). Normalised to 1J stored energy.




Figure 3.33: Surface fields along curve on final cell design.

Quantity	Value
E_{pk}/E_{acc}	3.99
B_{pk}/E_{acc}	3.71 (mT/MV/m)
$\sqrt{S_c}/E_{acc}$	0.0303 $\sqrt{A/V}$
Z	89.4 $M\Omega/m$
k_c	0.9 %
Q_0	12400
Frequency	5.712 GHz
Aperture radius	5 mm
Wall thickness	2.5 mm
Coupling cavity length	2 mm


 Table 3.6: Final optimised $\beta = 1$ cell parameters.


Figure 3.34: Effect of adding chamfer on magnitude of electric field distribution. Original geometry (Left), after adding 2.8 mm chamfer with 25° angle (Right).

3.6 Power of Pareto Design

This section describes how the power of Pareto optimisation can be fully exploited for the design of an RF cavity. Fig. 3.35 shows the coupling constant vs. E_{pk}/E_{acc} for all solutions of the $\beta = 1$ cell. The shunt impedance is represented by the colour, showing that a lower coupling constant generally gives a high shunt impedance, as expected. The wide range of values of the coupling constant that are accessible (0.8 % - 1.15 %) could be used for many different linac designs without having to repeat the optimisation. Fig. 3.36 shows how this would work in practice. The figure shows the required coupling constant vs. the average final particle energy at the end of the linac, along with lines of constant gradient from 15 MV/m to 40

MV/m. The E_{pk}/E_{acc} limit for a maximum field of 100 MV/m is also shown. The two shaded regions represent the available coupling that is possible for two values of the aperture radius, 5 mm and 5.5 mm. It is clear that a wide range of energies can be reached from 2 MeV to 9.5 MeV, with the number of cells ranging from 5-9.




Figure 3.35: k_c vs. E_{pk}/E_{acc} showing Z in colour.

Fig. 3.37 shows the optimal cavities for maximising the shunt impedance for each E_{pk}/E_{acc} limit. Depending on the energy and length requirements, one would select a gradient and then an optimal cavity from the Pareto front that satisfies both the coupling constant and peak electric field requirements. As an example, a 6 MeV design is considered, with the requirement that the linac should be as short as possible. Any gradients between 25 MV/m - 40 MV/m would be possible, but to minimise the length, 40 MV/m would be preferable. Therefore, a 6-cell design could be selected with a gradient between 35 MV/m and 40 MV/m, with an optimal cell selected based on the peak field limit. Any point is accessible within the shaded regions by adjusting the gradient to the desired value and then selecting a corresponding solution from the Pareto optimal set. It should be noted that this analysis does not include the bunching section which would be added afterwards, providing extra energy gain, allowing one to drop the gradient, or use fewer $\beta = 1$ cells.




Figure 3.36: Coupling constant vs. linac energy for various cavity gradients. The two shaded regions represent the accessible coupling regions for different values of the beam pipe aperture radius, A .




Figure 3.37: Cross section of five $\beta = 1$ cells that satisfy the various E_{pk}/E_{acc} limits shown in Fig. 3.36.

Chapter 4

Beam Capture Optimisation and Linac Tuning

4.1 Longitudinal Dynamics

Before the electron buncher design is presented, the key formulas for the longitudinal motion of particles in RF cavities are given here. Longitudinal focusing is provided by choosing the phase of the synchronous particle relative to the crest of the RF electric field. When this phase is selected to coincide with the rising field, as illustrated by the stable particle in Fig. 4.1, longitudinal restoring forces come into play. These forces cause early particles to encounter a smaller amplitude field while late particles encounter a larger one. As a result, particles centred around the stable phase remain in a stable bunch.

Considering the setup shown in Fig. 4.2. The synchronous phase is designated ϕ_s , synchronous energy W_s and synchronous velocity β_s . The particle phase in the n th cell is the phase of the RF field at the time when the particle is in the centre of the gap. As the energy changes from cell to cell, the length of the accelerating cells must change to keep the same phase. The approximate length of the $n+1$ cell is given by:

$$L_{n+1} = \frac{v_n \phi_a}{2\pi f} + \frac{dv_n}{dz} \frac{L_n \pi}{4\pi f} \quad (4.1)$$




Figure 4.1: Energy gain vs. Phase of a stable particle and particles arriving early or late.

where v_n is the velocity of the particle entering cell $n+1$ (leaving cell n) with phase advance ϕ_a . v_n can now be replaced with β_n (relativistic beta), which is v_n/c , then: [38, 64]:

$$L_{n+1} = \frac{\beta_n}{\frac{2f}{c} - \frac{1}{2} \frac{d\beta}{dz}} \quad (4.2)$$




Figure 4.2: Setup for accelerating cells to describe longitudinal motion. Image sourced from [63].

The rest mass of an electron, W_0 , is approximately 0.511 MeV and below this energy, electrons are considered not fully relativistic. At 1 MeV, they have reached close to 0.95c and at 3 MeV, 0.99c and are now considered fully relativistic; the β versus energy for electrons is shown in Fig. 4.3. The relativistic factors are given below:

$$\beta = \sqrt{1 - \frac{1}{\gamma^2}} = \sqrt{1 - \left(\frac{E_0}{E}\right)^2} = \sqrt{1 - \left(\frac{E_0}{E_0 + W}\right)^2} \quad (4.3)$$

As β approaches one rapidly, these formulas only apply to the first few cells, where the phase will be different for each cell. After that, the bunch will see the same phase in each gap throughout the rest of the linac.




Figure 4.3: Relativistic β vs. energy for electrons.

Defining the following variables for a particle with respect to the synchronous particle:

$$\Phi = \phi - \phi_s = \omega(t - t_s) \quad (4.4)$$

$$w = W - W_s \quad (4.5)$$

Where $t(\phi)$ is the arrival time (phase) of a particle at a gap and W is the kinetic energy of the particle at some position along the linac. The formulas that describe the motion of a non-synchronous particle with respect to the synchronous one are [92]:

$$\frac{dw}{dz} = qE_{acc}[\cos(\phi_s + \Phi) - \cos(\phi_s)] \quad (4.6)$$

and

$$\frac{d\Phi}{dz} = \frac{\omega}{cE_0\beta_s^3\gamma_s^3}w \quad (4.7)$$

These formulas show that as the electrons gain energy and $\gamma \gg 1$ the phase change from cell to cell quickly drops.

4.2 Phase Slippage

It is also useful to consider what happens when a DC beam with a $\beta_{inj} < 1$ is injected into a constant $\beta = 1$ structure. Before the electrons have reached relativistic speeds, and so have a constant phase, the relationship between β and ϕ from an initial set of conditions to a final set, is given by [92]:

$$\sin(\phi_{final}) = \sin(\phi_{initial}) + \frac{2\pi m_0 c^2}{\lambda_{RFQ} E_{acc}} \left(\sqrt{\frac{1 - \beta_{inj}}{1 + \beta_{inj}}} - \sqrt{\frac{1 - \beta_{final}}{1 + \beta_{final}}} \right) \quad (4.8)$$

Equation 4.8 tells us that for a given field and frequency, there are only some possible injection phases that can be captured. In order for a solution for ϕ_{final} to be found, the RHS must be within the interval $[-1, 1]$. This is shown schematically in Fig. 4.4 where only initial phases ranging from -135° to -45° can be captured when injecting a beam with $\beta_{inj} = 0.3$ into a $\beta_{final} = 1$, C-band structure with $E_{acc} = 30$ MV/m. There are several options for improving the capture, E_{acc} can be increased (though as shall be seen there is a limit to the maximum value), β_{final} can be reduced (although this will eventually be equal to 1) or the frequency can be reduced (in this case it is fixed at 5.712 GHz). The other option is to reduce the range of phases that the initial beam occupies by bunching them first. This is explored in the next section.




Figure 4.4: $\phi_{initial}$ vs ϕ_{final} for $\beta_{inj} = 0.3$, $\beta_{final} = 1$ and $E_{acc} = 30$ MV/m.

4.3 Electron Bunching

The thermionic guns used for typical industrial linacs produce electrons with energies ranging from 10-50 keV, and produce them over long pulses that are usually multiple RF periods long. This means that half of the electrons end up in the accelerating phase and half in the decelerating phase, which can cause them to travel back towards the cathode, if the amplitude of the field is too high (back-bombardment). This back-bombardment causes heating of the cathode and can reduce its lifetime. If the electrons are bunched beforehand using one of the various bunching methods to reduce the bunch length (and the phase region they occupy), then more can be captured. There has recently been a push to maximise the capture of the electrons to mitigate the effects of cathode bombardment [93]. This was achieved by creating a low-amplitude first cell, both to provide velocity bunching at an energy closer to the injection energy, and to allow electrons that start back-streaming to catch the next RF cycle before leaving the first cell and going back towards the cathode. These electrons are re-accelerated in the next cycle and come to the second cell with the correct phase to join the next bunch. They reported that the only losses were due to the back-streaming electrons from the second cell that have a higher energy due to the larger amplitude there and so cannot be recaptured. They were able to achieve a capture efficiency of over 90 % using this method, along with reporting a reduced energy spread of the final beam as a consequence of the improved bunching. It was decided that this method would be employed for the linac design to try and achieve the desired high capture efficiency.

4.3.1 Maximum Capture

A comprehensive review of electron bunching methods was recently written by Kutsaev, and the reader is referred there for a detailed treatment of the theory [94]. To capture 100 % of the electrons in a travelling wave linac, the following conditions must be met:

$$\beta_{inj} = \beta_{ph} \quad (4.9)$$

$$\left. \frac{d\beta_{ph}}{dz} \right|_{z=0} = 0 \quad (4.10)$$

$$A_m < \frac{\pi}{\beta_{ph}} \left(1 - \sqrt{1 - \beta_{ph}^2} \right) \quad (4.11)$$

Where β_{inj} is the injection relativistic beta, β_{ph} is the phase velocity in the structure, and A is the accelerating field amplitude given by

$$A_m = \frac{qE\lambda}{W_0} \quad (4.12)$$

where q is the charge of an electron, E is the gradient, λ is the wavelength, and W_0 is the initial particle energy. For a gradient of 25 MV/m, A has a value of ≈ 2.6 , which means that for the last condition to be met with a β_{inj} of 0.302, the gradient in the first cell should be < 5 MV/m. The first condition is difficult to realise as there are limitations on how short a cell can be as the irises have a finite thickness, and if the amplitude in the cell is high, it will be more prone to breakdown as it made excessively short. The last condition comes from requiring that the amplitude of the field is lower than the threshold which causes them to have a negative momentum. This means that the bunching cell should have a much smaller amplitude than the rest of the structure to meet this condition. The field should then increase to maximise the capture in the rest of the cavity, as described in Section 4.2. It is possible to achieve these conditions given enough space (such as using ballistic bunching with a drift space [95]), but meeting them in a compact design is difficult, and designers usually accept the losses when injecting electrons with $\beta_{inj} < \beta_{ph}$ [94].

As explained in [93], in order to achieve re-capture in a short design, the length of the first cell should be slightly longer than the value obtained by matching

the injection energy and the β of the cell. This allows electrons that do have a negative momentum to remain inside the cell before the next RF cycle. After the first bunching cell, the second cell should also be shorter than $\beta = 1$ to meet the conditions for high capture, whilst also providing acceleration so that the bunches can reach the third cell with the correct phase for acceleration. The bunches will then go with a constant phase into the remaining cells, which have $\beta = 1$ for acceleration to the final energy.

4.3.2 Tapered Buncher

An initial estimate of the lengths of the first cells was made based on analytical formulae derived using experimental data for tapered velocity bunchers, described in [94]. The maximum capture these formulae provide is limited, as there is only one parameter to vary, A_m . In addition, to achieve a compact design, the value of $\beta(z)$ will jump between cells, meaning a smooth function is not possible, reducing the capture. It was reported that typical values for these types of bunching sections are 40-60 %. Nonetheless, they are useful as a starting point for future optimisation. For a buncher design where the velocity and amplitude of the RF wave are tapered in the first cells, the phase velocity along the linac is estimated using the following formulae [94]:

$$\beta(z) = \frac{2}{\pi}(1 - \beta_{inj}) \arctan(k_q z^{k_2}) + \beta_{inj} \quad (4.13)$$

Where:

$$k_1 = 3.8 \cdot 10^{-3}(10.8^{A_m} - 1) \quad (4.14)$$

The amplitude variation along the bunching section was estimated using:

$$A(z) = k_3 - k_4 \cos\left(\frac{\pi}{k_5} z\right) \text{ for } 0 < z < k_5 \quad (4.15)$$

and:

$$A = A_M \text{ for } z \geq k_5 \quad (4.16)$$

and:

$$k_{3,4} = 0.5A_M \pm 0.15\sqrt{A_M} \quad (4.17)$$

$$k_5 = \frac{1}{1.25\sqrt{A_M}} \quad (4.18)$$

where β_{inj} is the injection β , A_M is given by equation 4.12 and k_{1-5} are fitting parameters. The results are shown in Fig. 4.5, which gives an estimate of how the cell lengths and amplitudes should vary along the linac. From these plots, it was estimated that reducing the length of the first two or three cells would be sufficient for transitioning from the low-energy gun to relativistic energies. The values calculated here would also be used to estimate the bounds of the amplitudes and lengths of the cells for the optimisation process described in the following sections.




Figure 4.5: Estimate of how the phase velocity, $\beta(z)$ (left) and amplitude, $A(z)$ (right), should vary along the linac in a tapered velocity buncher.

4.4 Cell Length Optimisation

The following section describes the optimisation process for the lengths of the cavity cells. As mentioned above, one of the main drivers of the total lifetime of industrial linacs is the cathode lifetime, which can be significantly reduced when

it is bombarded with excessive back-streaming electrons that do not end up getting accelerated [96]. For this reason, the number of electrons returning to the cathode during linac operation is an objective that should be minimised. The total capture efficiency of the linac is also of interest, as electrons can generate X-rays when they hit the walls of the linac, so ideally, all of the electrons would make it through the linac and generate useful X-rays at the target. After discussions with the industrial partner, it was decided that any electrons with energies > 1.5 MeV would be adequate for generating useful X-rays, so this was a second optimisation objective. Finally, the average kinetic energy should be close to 2 MeV so that, on average the X-ray penetration is at the design value. These are the objectives that were used for the cell length optimisation and they are summarised in Table. 4.1.

Optimisation parameter	Min/Max	Units
Average energy > 1.5 MeV	Max	MeV
Total capture efficiency	Max	%
Capture efficiency > 1.5 MeV	Max	%
No. electrons returning to cathode	Min	%

Table 4.1: Cell length optimisation objectives.

Simulations were performed in Astra [97] which is a space charge tracking algorithm that can be used to track electrons as they traverse external electromagnetic fields. A representative particle distribution of a standard thermionic cathode was used, with an average kinetic energy of 25 keV. 1000 macro-particles were used for initial design studies, emitted over 1 ns, which is multiple RF periods (≈ 6). Two field maps shown in Fig. 4.6 were used for the simulations, taken from electromagnetic simulations of a $\beta = 1$ and a $\beta = 0.5$ cell. The $\beta = 0.5$ cell field map was used for the first 2 cells and the $\beta = 1$ cell field map for the remaining cells. These fields were then stitched together to form a 1D on-axis field. The first step was to scan the maximum field on axis for a 5 cell linac with the imported field maps stretched and set to $\beta = 0.35$ for the first cell (longer than $\beta_{inj} = 0.302$), to $\beta = 0.7$ for the second cell and to $\beta = 1$ for the remaining cells. The amplitudes of the first and second cells were set to $0.3A_{max}$ and $0.6A_{max}$, as estimated from Fig.

4.5. The maximum and average energies of the electrons were recorded at the end of the linac, shown in Fig. 4.7. This gave a rough idea of the field amplitude that would be required to reach the 2 MeV design energy.

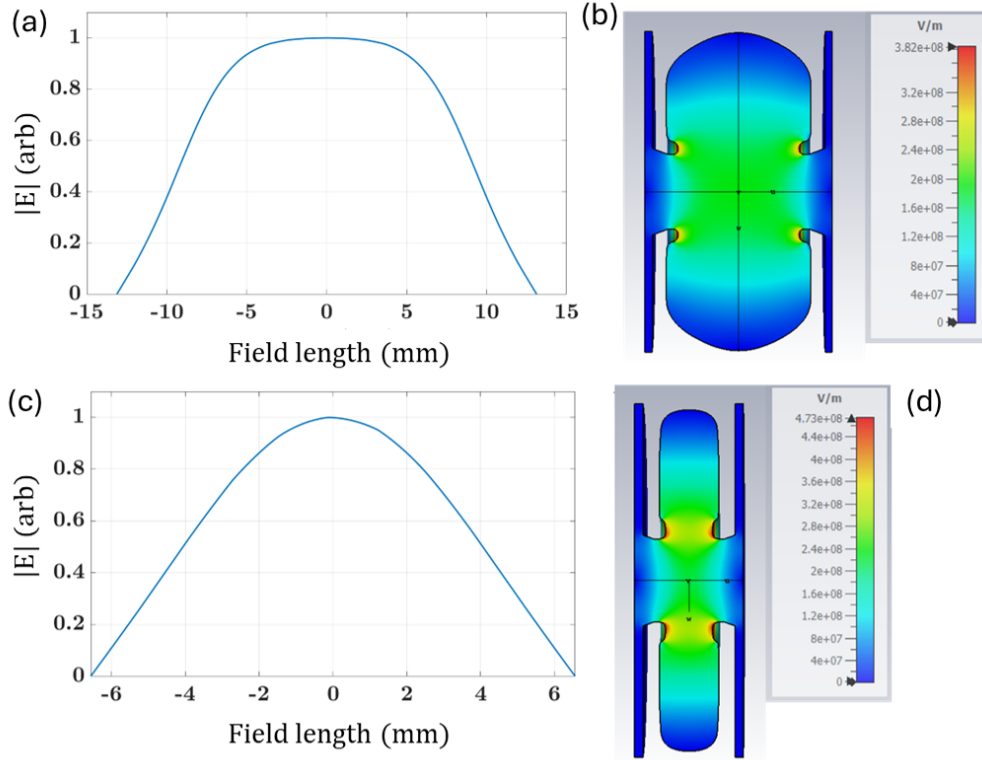


Figure 4.6: (a) $\beta = 1$ 1D field (b) $\beta = 1$ electric field contour (c) $\beta = 0.5$ 1D field (d) $\beta = 0.5$ electric field contour.




Figure 4.7: (a) Scanning of the maximum electric field amplitude. (b) Calculated required power to obtain this field.

4.4.1 MO Optimisation

As performing parameter scans of the field amplitudes and lengths is a time consuming process, the ideas from the previous cavity optimisations were extended to optimise the beam. An optimisation process was developed using Matlab and Astra. From the scan of the maximum field amplitude on axis, the required input power required to achieve the fields in the linac was calculated using the following:

$$P_{req} = \frac{P_{CST} E_{axis}^2}{E_{CST}^2} \quad (4.19)$$

Where E_{axis} is the value provided to Astra to scale the fields to, P_{CST} is the power in the CST frequency domain solver, which is 0.5 W and E_{CST} is the electric field value on axis in the $\beta = 1$ cell. The results are shown to the right of Fig. 4.7. To achieve 2 MeV, ≈ 0.8 MW would be required from the klystron, which was deemed acceptable for a multi-beam klystron after considering the losses in the waveguide network. Following this a multi-objective optimisation was set up using Matlab to drive ASTRA with the input parameters used for the optimisation and limits shown in Table. 4.2. The lengths and amplitudes of the first three cells were allowed to vary with the last two cell lengths set to $\beta = 1$ and $A = E_{max}$. The genetic algorithm in Matlab (a variation of NSGA-II) was used for the optimisation with an initial population size of 80, and a run of 200 generations. The other parameters set in the optimiser settings are shown in Table. 4.3, with a description of the function of each parameter. The optimisation objectives used for the optimisation were the average kinetic energy of the particles with energies > 1.5 MeV, the % of electrons captured with energies > 1.5 MeV, the total capture efficiency and the % of electrons that return to the cathode, i.e. that are not captured and end up back-streaming. The capture % is the number of particles that make it to the end of the linac with the specified energy. The fields were initially scaled to a maximum amplitude of 40 MV/m based on the previous maximum amplitude on axis scans.

The 2-D scatter plots showing the solution points are shown in Fig. 4.8. The

Parameter	limits
β cell 1	0.2-0.5
β cell 2	0.5-0.75
β cell 3	0.75-1
Amplitude cell 1/max amplitude	0.25-0.7
Amplitude cell 2/max amplitude	0.5-0.9
Amplitude cell 3/max amplitude	0.9-1

Table 4.2: Astra MO input parameters and initial limits.

Parameter	Value	Description
Population size	80	Size of population.
Pareto fraction	0.35	Sets the fraction of individuals to keep on the first Pareto front while the solver selects individuals from higher fronts.
Selection	Binary tournament	Selection method for crossover.
Creation function	ga creation uniform	Creates a random initial population with a uniform distribution.
Crossover function	@crossover intermediate	Creates children by taking a weighted average of the parents.
Crossover fraction	0.8	Fraction of the population at the next generation.
Max generations	200*Number of variables	Maximum number of generations the simulation runs for.
Function tolerance	1e-4	Algorithm stops if the average relative change in the best fitness function over max stall generations is less than function tolerance.
Max stall generations	100	See above.
Mutation function	1.43	Randomly generates directions that are adaptive with respect to the last successful or unsuccessful generation. The mutation chooses a direction and step length that satisfies constraints.

Table 4.3: Matlab gamultiobj parameters used for optimisation of cell lengths.

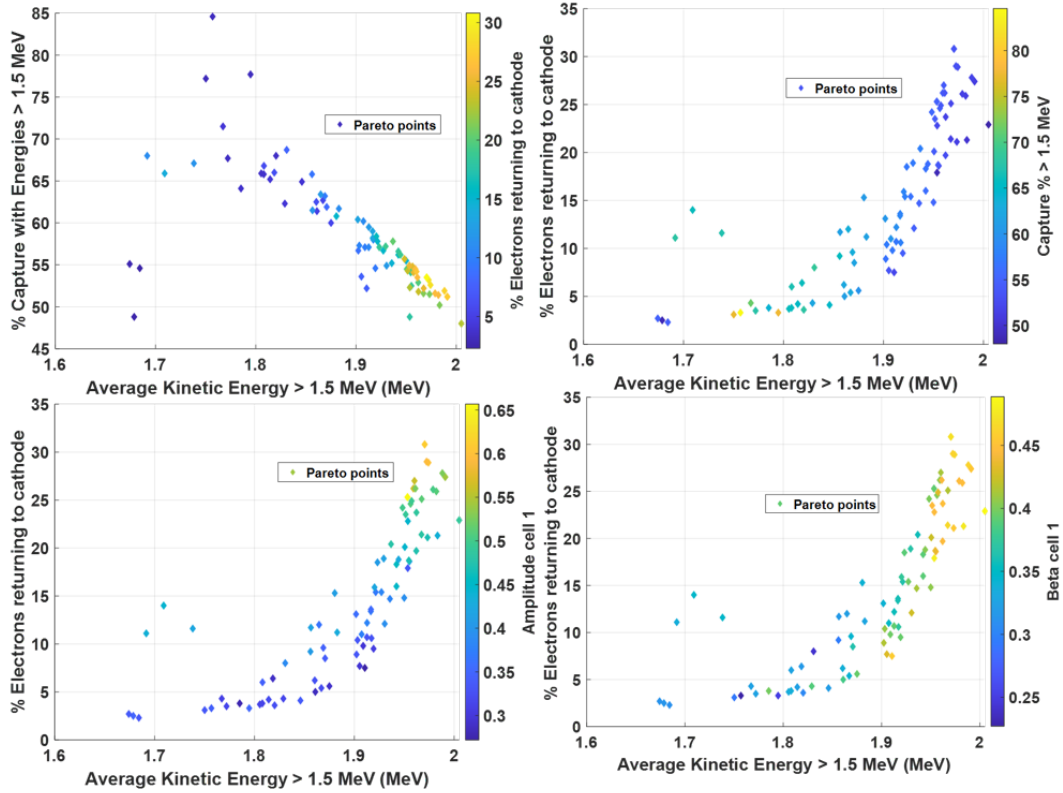


Figure 4.8: Scatter plots (last generation) for cell length optimization, showing how the average kinetic energy, % electrons returning to cathode and length/amplitude of the first cell interact.

points represent the Pareto points from the last generation of the optimisation routine. These figures show that in order to maximise the capture efficiency, the length and amplitude of the first cell must be reduced. However, as expected, this comes at the expense of the average kinetic energy, as the electrons are accelerated over a shorter distance [93]. As the average KE was slightly too low for the 40 MV/m case, the optimisation was repeated at 45 MV/m. Fig. 4.9 shows the full matrix of the optimization results for both cases. Again, this confirmed that the first cell length needed to be between $\beta = 0.25$ and $\beta = 0.35$. Also, when the first and second cell lengths are decreased, the third cell length must be slightly shorter than $\beta = 1$, somewhere between $\beta = 0.925$ and $\beta = 0.975$ to improve the transition into the last two $\beta = 1$ cells from the first and second cells, as the phase is still changing.

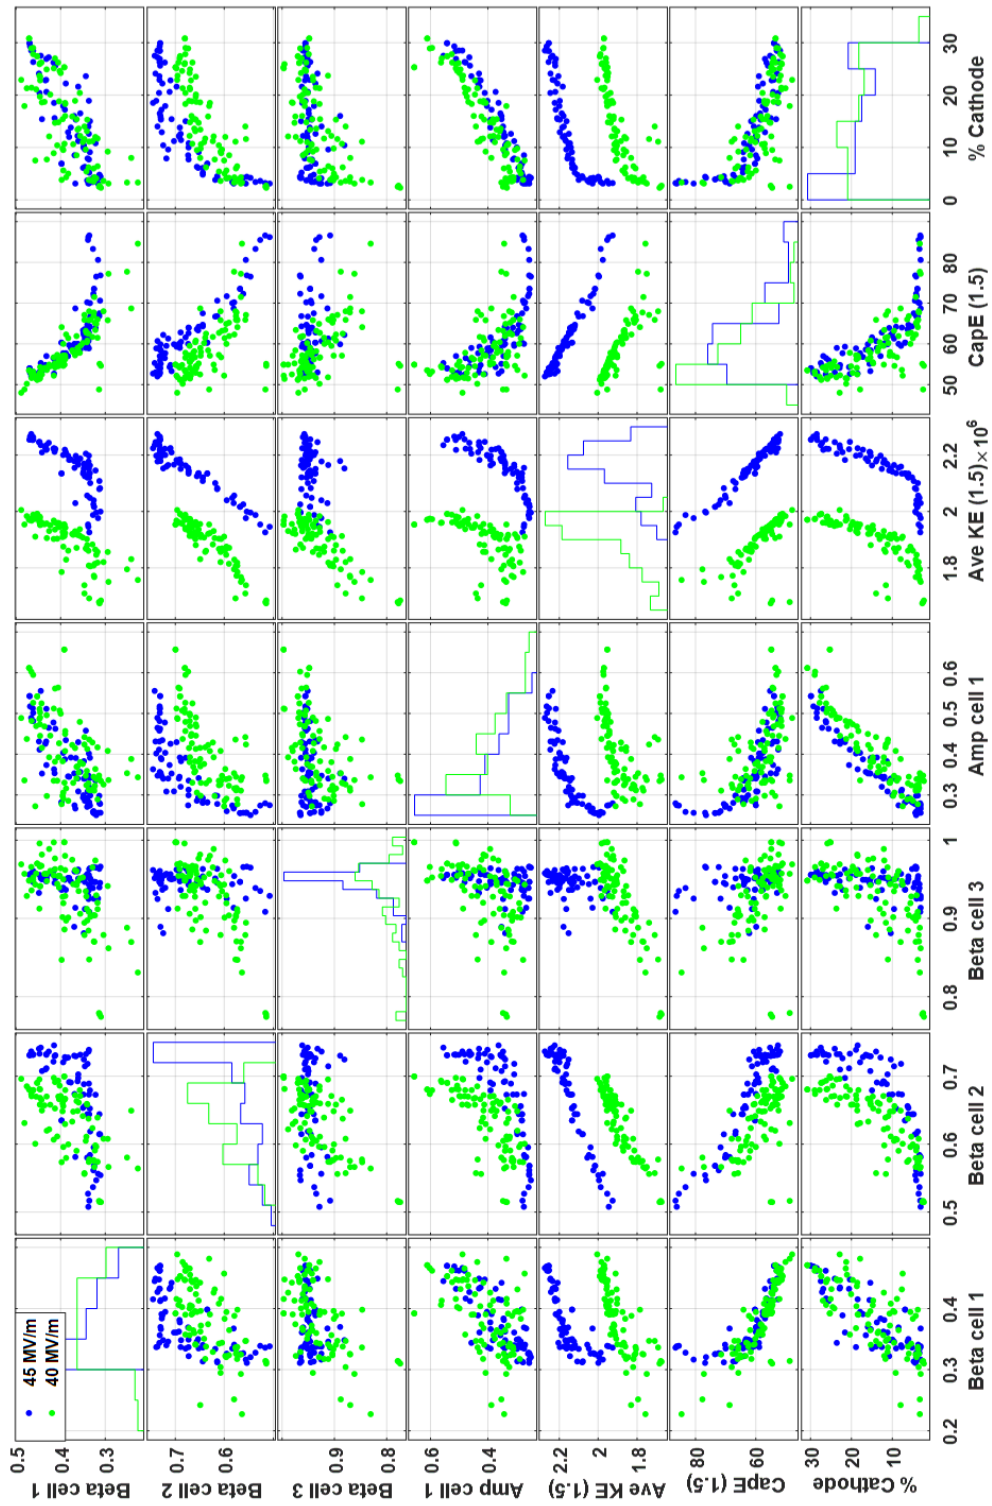


Figure 4.9: Matrix for cell length Optimization, showing Pareto points from last generation at 40 MV/m and 45 MV/m.

4.4.2 First Cell Design

Following the initial optimisation, it was envisioned that the most difficult part of achieving one of these designs would be obtaining a short first cell length and low amplitude, as there are limits to how short a cell can be made. An investigation into the feasibility of a very short first cell was performed to see what the limit on the length would be. An option to reduce the field length would be to just reduce the length of the cell without having nose cones until the desired β is achieved, as shown on the left of Fig. 4.11, for a $\beta = 0.44$ cell. Although this would reduce the field length, as the length is reduced, eventually, the aperture will be larger than the cell length, and the field will become dominated by the decay into the beam-pipe. The cell would also become stiffer, making it more difficult to tune and more susceptible to manufacturing tolerance errors due to the smaller dimensions. An example of a shorter $\beta = 0.35$ is shown in the centre of Fig. 4.11. When performing a bead-pull [98] it would also become difficult to measure the field if the length of the cell is small compared to the size of the bead; therefore, it should not be excessively short.

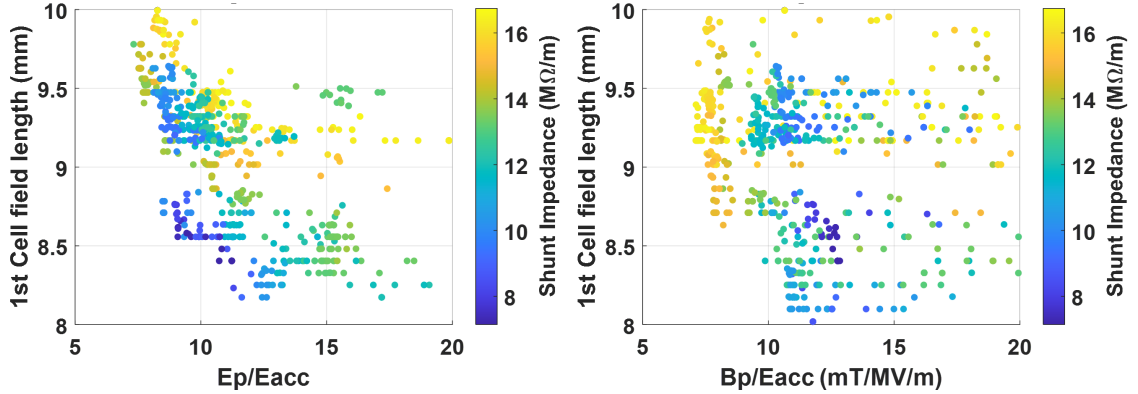


Figure 4.10: First cell optimization results. Length vs. E_{pk}/E_{acc} (left), Length vs. B_{pk}/E_{acc} (right).

In order to navigate these issues, another optimisation was performed for the first cell. First, the initial beam-pipe radius was further reduced to 2.5 mm from 5 mm in order to force the field to zero faster, and the cavity geometry was set up asymmetrically so that one side remained straight and the other was defined by a

NURBS spline. The cavity was defined in this way to give control over the distance between the first and second cell fields, as the aperture could not be changed by a large amount, and the distance between the cells was fixed by the coupling cell. This meant that an axisymmetric design was needed to push the fields away from the axis, giving fine control over the coupling between the first two cells. The optimisation was set up the same as the $\beta = 1$ cells but also to minimise the length of the field between the points where it decays to 10 % of the maximum value. The results of the optimisation are shown in Fig. 4.10. An example of the progression of the first cell design is shown in Fig. 4.11, with the on-axis electric field for the three cases shown in Fig. 4.12





Figure 4.11: Progression in the design of the low amplitude first cell.

For each of the cells shown in Fig. 4.11, the radius was scanned and the mode frequency recorded, so that values of $\Delta f/\Delta R$ could be found, these are shown in Table. 4.4. These values show that the spline model is most easily tuned as effects at the cavity radius change the frequency $\approx 20\%$ more than the other models. Fig. 4.12, shows the field on axis for the three models. It shows a reduction in field length for the optimised spline model of $\approx 35\%$ compared to a $\beta = 0.44$ cell and $\approx 20\%$ compared to the $\beta = 0.35$ cell. The final first cell was chosen as a trade-off between the peak fields, shunt impedance, and field length, although preference was given to a shorter design so that most of the design points that minimise the number of electrons returning to the cathode could be accessed.

Cavity	β	$\Delta f/\Delta R$ (GHz/mm)
1	0.35	0.28
2	0.44	0.288
Spline model	0.44	0.341

Table 4.4: Frequency change for outer radius change for the three first cell models.

Figure 4.12: First cell optimization results showing a reduction in field length whilst maintaining the same value of β on the outer cell.

4.4.3 Final Design Selection

After the first cell geometry had been finalised, the second and third cells were designed. The MO single cell optimisation was repeated for $\beta = 0.5$ so that the second cell could be selected and it was assumed that the third cell could be selected from the $\beta = 1$ optimisation pool, being only 5 % shorter. An initial second cell design selected is shown in Fig. 4.13. After putting the first three cells together to investigate the amplitudes, it became clear that having the second cell amplitude larger than $0.8A_{max}$ was infeasible; as the cell becomes shorter, the amplitude tends to decrease due to the increased coupling, and there was a limit to how much the coupling could be reduced again using chamfers.

To improve this the optimisation was repeated with an upper limit to the second cell amplitude of $0.8A_{max}$, the β of the third cell was fixed to 0.95 and the gradient allowed to increase to 45.5 MV/m to compensate for the decreased




Figure 4.13: Second cell design.


amplitudes. The Pareto fronts of the capture > 1.5 MeV vs. kinetic energy and number of electrons returning to the cathode vs. kinetic energy are shown in Fig. 4.14a and Fig. 4.14b respectively. Four optimisations are shown with various limits on the maximum field on axis, amplitude limits and β values for the cells. The final target design was selected from the green curve and the optimisation values for this result are shown in Table. 4.5.

Parameter	Value
β cell 1	0.3258
β cell 2	0.5486
β cell 3	0.95
Amplitude cell 1/max amplitude	0.303
Amplitude cell 2/max amplitude	0.7900
Amplitude cell 3/max amplitude	0.9438
Total Capture (%)	92.4
Capture > 1.5 MeV (%)	84.5
Average KE (MeV)	1.926


Table 4.5: Final design parameters selected from beam capture optimisation.

4.5 Coupled Periodic Standing Wave Cavities

In this section a theory of bi-periodic standing wave cavity coupling is presented to show how the coupling and frequencies of the cells are related to the amplitudes in the cells. A complete derivation is described in [99], and additional models describing similar analytical models can also be found in [100–103]. Fig. 4.15 shows a circuit



(a)



(b)

Figure 4.14: (a) Capture vs. kinetic energy Pareto fronts for cell length optimisation. (b) Electrons at cathode vs. kinetic energy Pareto fronts for cell length optimisation.

model for a bi-periodic cavity where the cells are identical and capacitively coupled through the iris. M1 represents the nearest neighbour coupling, M2 is the coupling between accelerating cells, and M3 is the coupling between the coupling cells.

Using Kirchoff's voltage law in the n^{th} and $n^{th} + 1$ cells gives the following

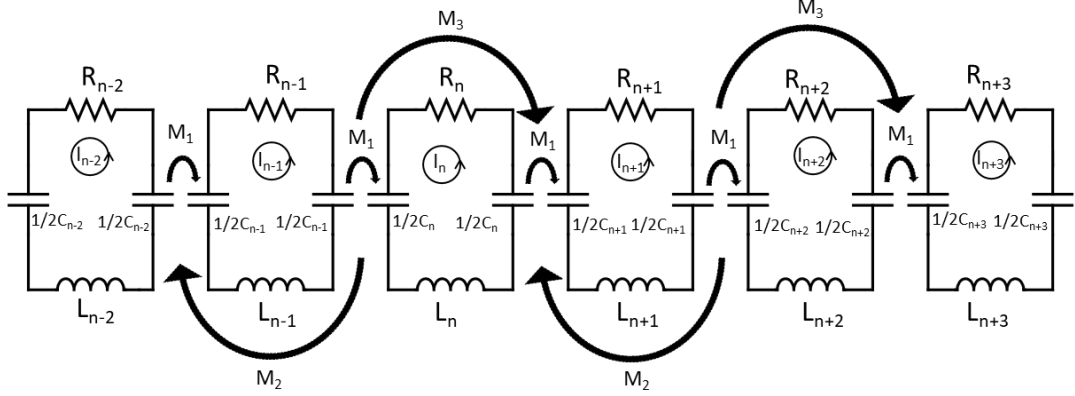


Figure 4.15: Chain of capacitively coupled resonators with nearest neighbour and next nearest neighbour coupling included.

$$0 = i_n \left(j\omega L_n + R_n + \frac{1}{j\omega C_n} \right) + \frac{M_1}{j\omega C_n} (i_{n-1} + i_{n+1}) + \frac{M_2}{j\omega C_n} (i_{n-2} + i_{n+2}) \quad (4.20)$$

$$0 = i_{n+1} \left(j\omega L_{n+1} + R_{n+1} + \frac{1}{j\omega C_{n+1}} \right) + \frac{M_1}{j\omega C_{n+1}} (i_n + i_{n+2}) + \frac{M_3}{j\omega C_{n+1}} (i_{n-1} + i_{n+3}) \quad (4.21)$$

Dividing through by $jL_n\omega$ and $jL_{n+1}\omega$, letting $M_1 = \frac{k_1}{2}$, $M_2 = \frac{k_2}{2}$, $M_3 = \frac{k_3}{2}$, assuming the cavities are lossless, and using the relationships

$$\omega_1 = \frac{1}{\sqrt{L_n C_n}} \quad \text{and} \quad \omega_2 = \frac{1}{\sqrt{L_{n+1} C_{n+1}}} \quad (4.22)$$

gives

$$0 = i_n \left(1 - \frac{\omega_1^2}{\omega^2} \right) - \frac{\omega_2^2 k_1}{2\omega^2} (i_n + i_{n+2}) - \frac{\omega_2^2 k_3}{2\omega^2} (i_{n-1} + i_{n+3}) \quad (4.23)$$

Then current solutions of the form

$$i_n(q) = A \cos(n\phi) \exp^{j\omega_q t} \quad (4.24)$$

$$i_{n+1}(q) = B \cos[(n + 1)\phi] \exp^{j\omega_q t} \quad (4.25)$$

are assumed, where $\phi = \frac{\pi q}{N-1}$, $0 \leq q \leq (N-1)$, and $1 \leq n \leq N$. Here, q is the mode number, N is the total number of cells, and n is the cell number. Making these substitutions and then simplifying, the dispersion relation including nearest and next nearest neighbour coupling can be obtained as

$$k_1^2 \cos^2 \phi = (1 - \omega^2/\omega_1^2 + k_2 \cos 2\phi)(1 - \omega^2/\omega_2^2 + k_3 \cos 2\phi) \quad (4.26)$$

The stopband is given by

$$sb = f_2 \sqrt{1 - k_3} - f_1 \sqrt{1 - k_2} \quad (4.27)$$

Where:

$$f_{1,2} = \omega_{1,2}/2\pi \quad (4.28)$$

4.5.1 Lossy Bi-Periodic Equivalent Circuit

Fig. 4.16 shows the model for a capacitively coupled cavity with a driving voltage and end cells. Following the same approach using Kirchhoff's voltage loops as before, the following is obtained

$$V_d = i_1 \left(R_1 + \frac{1}{sC_1} + sL_1 \right) + \frac{i_2 M_1}{sC_1} + \frac{i_{n-1} M_2}{sC_1} \quad (4.29)$$

$$0 = i_2 \left(R_2 + \frac{1}{sC_2} + sL_2 \right) + \frac{i_1 M_1}{sC_2} + \frac{i_{n-1} M_1}{sC_2} + \frac{i_n M_3}{sC_2} \quad (4.30)$$

$$0 = i_{n-1} \left(R_{n-1} + \frac{1}{sC_{n-1}} + sL_{n-1} \right) + \frac{i_2 M_1}{sC_{n-1}} + \frac{i_1 M_2}{sC_{n-1}} + \frac{i_n M_1}{sC_{n-1}} \quad (4.31)$$

$$0 = i_n \left(R_n + \frac{1}{sC_n} + sL_n \right) + \frac{i_2 M_3}{sC_n} + \frac{i_{n-1} M_1}{sC_n} \quad (4.32)$$

Again replacing variables with resonant frequencies




Figure 4.16: (Top), Bi-periodic cavity chain showing nearest neighbour and next nearest neighbour coupling. (Bottom), 4 cell model with end cells and driving term V_d .

$$\frac{V_d}{L_1} = i_1 \left(\frac{\omega_1}{Q_1} - j \left(\frac{\omega_1^2 - \omega^2}{\omega} \right) \right) - i_2 j \frac{\omega_1^2}{\omega} M_1 - i_{n-1} j \frac{\omega_1^2}{\omega} M_2 \quad (4.33)$$

$$0 = i_2 \left(\frac{\omega_2}{Q_2} - j \left(\frac{\omega_2^2 - \omega^2}{\omega} \right) \right) - i_1 j \frac{\omega_2^2}{\omega} M_1 - i_{n-1} j \frac{\omega_2^2}{\omega} M_1 - i_n j \frac{\omega_2^2}{\omega} M_3 \quad (4.34)$$

$$0 = i_{n-1} \left(\frac{\omega_{n-1}}{Q_{n-1}} - j \left(\frac{\omega_{n-1}^2 - \omega^2}{\omega} \right) \right) - i_2 j \frac{\omega_{n-1}^2}{\omega} M_1 - i_1 j \frac{\omega_{n-1}^2}{\omega} M_2 - i_n j \frac{\omega_{n-1}^2}{\omega} M_1 \quad (4.35)$$

$$0 = i_n \left(\frac{\omega_n}{Q_n} - j \left(\frac{\omega_n^2 - \omega^2}{\omega} \right) \right) - i_2 j \frac{\omega_n^2}{\omega} M_3 - i_{n-1} j \frac{\omega_n^2}{\omega} M_1 \quad (4.36)$$

This can then be written in matrix form, with the impedance matrix being inverted in order to solve for the currents in each cell

$$\begin{bmatrix} \frac{\omega_1}{Q_1} - j \left(\frac{\omega_1^2 - \omega^2}{\omega} \right) & -j \frac{\omega_1^2}{\omega} M_1 & -j \frac{\omega_1^2}{\omega} M_2 & 0 \\ -j \frac{\omega_2^2}{\omega} M_1 & \frac{\omega_2}{Q_2} - j \left(\frac{\omega_2^2 - \omega^2}{\omega} \right) & -j \frac{\omega_2^2}{\omega} M_1 & -j \frac{\omega_2^2}{\omega} M_3 \\ -j \frac{\omega_{n-1}^2}{\omega} M_2 & -j \frac{\omega_{n-1}^2}{\omega} M_1 & \frac{\omega_{n-1}}{Q_{n-1}} - j \left(\frac{\omega_{n-1}^2 - \omega^2}{\omega} \right) & -j \frac{\omega_{n-1}^2}{\omega} M_1 \\ 0 & -j \frac{\omega_n^2}{\omega} M_3 & -j \frac{\omega_n^2}{\omega} M_1 & \frac{\omega_n}{Q_n} - j \left(\frac{\omega_n^2 - \omega^2}{\omega} \right) \end{bmatrix}^{-1} \begin{bmatrix} \frac{V_d}{L_1} \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} i_1 \\ i_2 \\ i_{n-1} \\ i_n \end{bmatrix} \quad (4.37)$$

This shows how the amplitude in the cells is related to the coupling and frequencies of the cells. This can then be solved using the process described in [99], for an on-tune 9-cell bi-periodic cavity at 5.712 GHz, driven by V_d in the centre cell, the amplitudes in all cells are shown in Fig. 4.17. The amplitude in the coupling cells is higher in cavities 4 and 6 due to the losses along the linac. From equation 4.21 it can understand that as the coupling from the previous cell/next cell is changed, there must be a corresponding change in the current (amplitude) of the cell to ensure that the sum of the voltages is still zero. Using this method, the required coupling between the cells that would give the amplitude distribution shown in Table. 4.5 for maximum capture could be estimated.




Figure 4.17: Current amplitude in all 9 cells for a perfectly tuned bi-periodic capacitively coupled cavity.

The coupling constants between the first cells were scanned in order to look at the effects on the fields and find the required amplitudes. The final field displayed in Fig. 4.18 shows the field amplitude distribution with a 62 % increase in coupling

between cells 1-2, a 9 % increase between cells 3-4 and a 3.2 % increase between cells 5-6. This method could also be used to investigate detuning effects of the individual cells on the amplitudes, which would be very useful when tuning a manufactured structure as one could compare bead-pull measurements with these simulations.




Figure 4.18: Current amplitude in all 9 cells after tuning the coupling between the first few cells.

4.5.2 Amplitude Tuning

The first three cells were put together with the correct lengths as shown in Fig. 4.19, with the resulting $\pi/2$ mode field for different chamfer angles between the third and second cells shown below. The first step to obtaining the correct amplitudes was to adjust the parameters of the first cell so that they matched the length and amplitude to the design values. The parameters used are shown in Fig. 4.20, and the effects are shown in Fig. 4.21 and Fig. 4.22. The parameter cg_{min} was used to adjust the length and a combination of aperture radius and blend radius was used to ensure that the ratio of the first to second field was correct. Once this had been achieved, the chamfer on the third cell was used to change the amplitude ratio between the first and second cells and the third cell as shown in the bottom of Fig. 4.19.

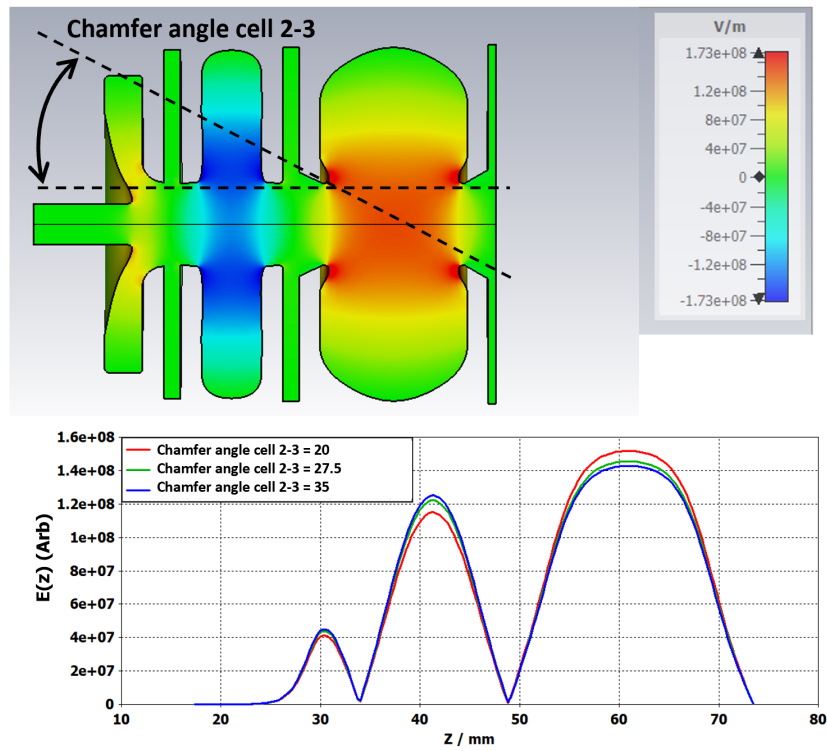


Figure 4.19: First three cells (top) and field amplitude of the $\pi/2$ mode for various values of the chamfer angle between the third and second cell (bottom).




Figure 4.20: First cell parameters used for tuning.




Figure 4.21: First cell length for different values of $c_{g_{min}}$ which increases the nose cone, pushing the mode field further up the cell. This shows how the length and amplitude can be tuned.

4.6 Coupler Design

The following sections describe the rest of the linac design process, including the coupler design and linac tuning. It was decided that power would be coupled into the linac from the third cell, therefore it would need to be modified to incorporate the power coupler.

The quality factor of a resonator is defined as the ratio of the initial stored energy to the energy lost in one RF cycle. It is essentially a measure of the power lost in the cavity walls due to the currents flowing there and the finite resistivity of copper. The intrinsic quality factor of the cavity is defined as

$$Q_0 = \frac{\omega U}{P_d} \quad (4.38)$$

where ω is the angular frequency, P_d is the power dissipated in the cavity walls and U is the stored energy given by

$$U = \frac{1}{2} \mu_0 \int |H|^2 dV \quad (4.39)$$

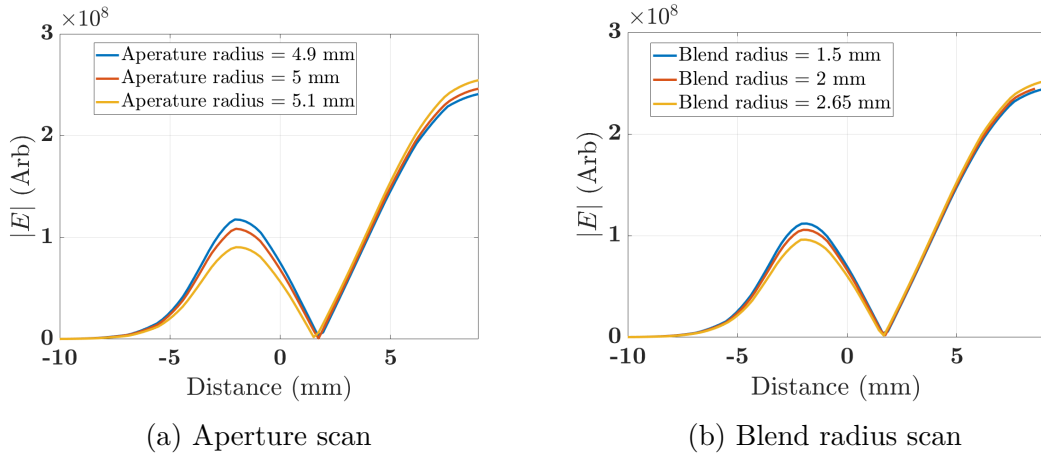


Figure 4.22: Effect of aperture radius and blend radius parameters on length and amplitude of the field in the first cell.

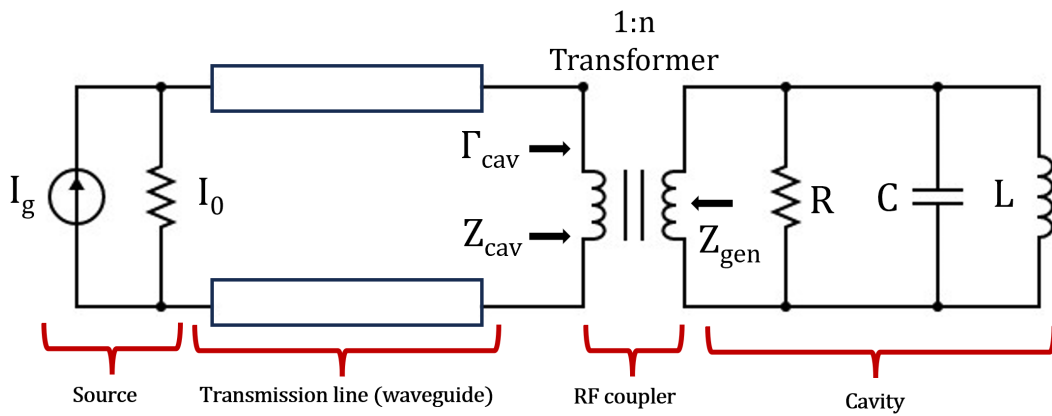


Figure 4.23: Circuit model showing waveguide and cavity connected by a coupler.

The cavity coupler and waveguide can be modelled with a circuit as shown in Fig. 4.23. The RLC network represents the cavity, the Klystron is the current generator, the transmission line is the waveguide, and the 1:n transformer is the RF coupler. [104] To the generator, the cavity looks like an impedance termination given by

$$Z_{cav} = \frac{R/n^2}{1 + jQ_0\delta} \quad (4.40)$$

where δ is the ratio of the frequency difference from the resonance ($\omega - \omega_0$) to the resonance frequency ω_0

$$\delta = \frac{\omega - \omega_0}{\omega_0} \quad (4.41)$$

As the cavity must be connected to a waveguide, more power will be lost before reaching the cavity. So that this is taken into account, the loaded quality factor is defined as

$$Q_l = \frac{\omega U}{P_d + P_{ext}} \quad (4.42)$$

Where

$$\frac{1}{Q_l} = \frac{1}{Q_0} + \frac{1}{Q_{ext}} \quad (4.43)$$

Where P_{ext} is the power flowing out of the cavity into the load and Q_e is the external quality factor defined as

$$Q_e = \frac{\omega U}{P_{ext}} \quad (4.44)$$

The coupling factor is defined as

$$\beta_c = \frac{Q_0}{Q_e} = \frac{P_{ext}}{P_d} = \frac{R}{n^2 Z_0} \quad (4.45)$$

which measures the ratio of the power supplied to the cavity to the power dissipated in the walls. The reflection coefficient seen by the source for a zero length power line is given by

$$\Gamma_{cav} = \frac{\beta_c - 1 - jQ_0\delta}{\beta_c + 1 + jQ_0\delta} \quad (4.46)$$

which reduces to

$$\Gamma_{cav} = \frac{\beta_c - 1}{\beta_c + 1} \quad (4.47)$$

when the cavity is on resonance ($\omega = \omega_0$). Ideally, the reflection of the cavity should be minimised, or $\beta_c = 1$, so that all the power supplied is delivered to the cavity. The coupling factor scales with the stored energy; therefore, for identical cells it would scale directly with the number of cells. As a first approximation using a one cell model with a coupler for a 5-cell cavity would require $\beta_c = \approx 5$, although this

would need to be tuned as the cells are not identical. The single-cell coupler model is shown in Fig. 4.26, along with a cut of the electric and magnetic fields through the centre. The waveguide selected for the design was WR159 (WG13), which has an operating band of 4.90 GHz to 7.05 GHz and dimensions of 40.386 mm by 20.193 mm. The length of the third cell not including the wall thickness was only 17.75 mm, therefore a taper was required to get from the waveguide into the third cell; this is shown in Fig. 4.24.




Figure 4.24: Taper to get from cell width to WR159 waveguide.

A simulation was set up to investigate the S-parameters of the taper to ensure adequate transmission into the linac. The results are shown in Fig. 4.25. Increased transmission is seen until the length of the taper approaches the guide wavelength (39 mm), meaning that the ideal length for the taper was around 35 mm.




Figure 4.25: S_{11} vs. taper length for the taper design.

As the introduction of the coupler onto one of the cells perturbs the fields and causes a dipole component of the field, it was also decided to include the vacuum pumping port in a similar configuration to the coupler but on the opposite side [104]. This should remove the dipole component and instead introduce a less harmful quadrupole component. To optimise the coupler, the width of the section that connects the coupler and the cavity was changed to achieve $Q_e = 2306$ as $Q_0 = 11546$, giving $\beta_c \approx 5$. The frequency of the cavity then had to be re-tuned as changes in the coupler perturb the frequency. This in turn changes the β_c , so an iterative approach was used until both were correct. One concern when designing

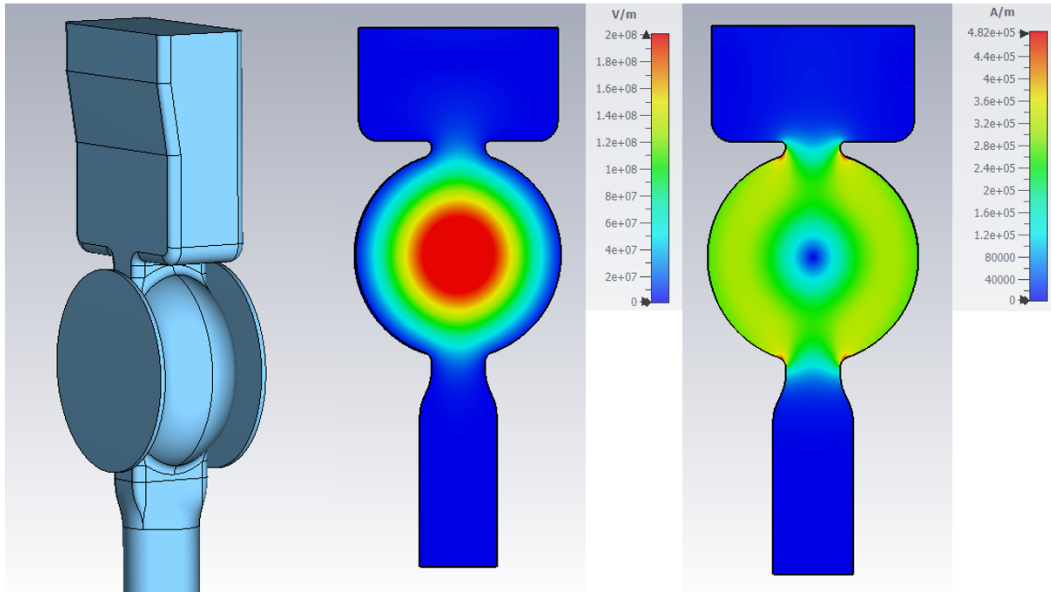


Figure 4.26: Coupler design with taper. Vacuum model (left), electric field through cell (centre), magnetic field strength through cell (right).

the coupler is the peak magnetic fields on the section where the coupler attaches to the cavity, as shown in Fig. 4.27. These high fields can cause surface fatigue from cyclic stresses, leading to arcing and breakdown in the coupler, so they must be avoided. The original value of B_{pk}/E_{acc} was 3.71, and with the addition of the coupler, this value increased to 6, thus by a factor of 1.61. The B_{pk}/E_{acc} limit for a 20° temperature rise at 26 MV/m is 8.8, which was considered a safe value.




Figure 4.27: Peak magnetic field on coupling slot.

4.7 Full Cavity Design

After the coupling cell and first cells had been designed, they were combined with two more $\beta = 1$ cells to create the full linac. The vacuum model of the linac and the contour of the z -component of the electric field are shown in Fig. 4.28.

A full frequency domain simulation was performed on the linac and it was found that the β_c of the linac needed to be adjusted as the system was over-coupled, meaning that the minimum S_{11} value was only -20 dB. The width of the coupler was scanned, with the middle cell frequency adjusted each time to compensate, and the full linac was then simulated to find the critical coupling point. The results are shown in Fig. 4.29, showing that the critical coupling occurs around $\beta_c = 4.55$ in the single-cell simulation. This was due to the fact that the coupling scales with the stored energy, and the first two linac cells are very short, which means that combining them equals something close to a stored energy of $0.55U_{\beta_1}$, where U_{β_1} is the stored energy of a $\beta = 1$ cell.

4.7.1 Tuned Model Comparison

Once the β_c of the linac had been tuned, the z -component of the electric field on axis was exported from CST and imported into Astra. The same simulation was performed, propagating the beam through the tuned CST fields instead of the single-cell combinations used previously. A comparison of the tuned fields from CST and

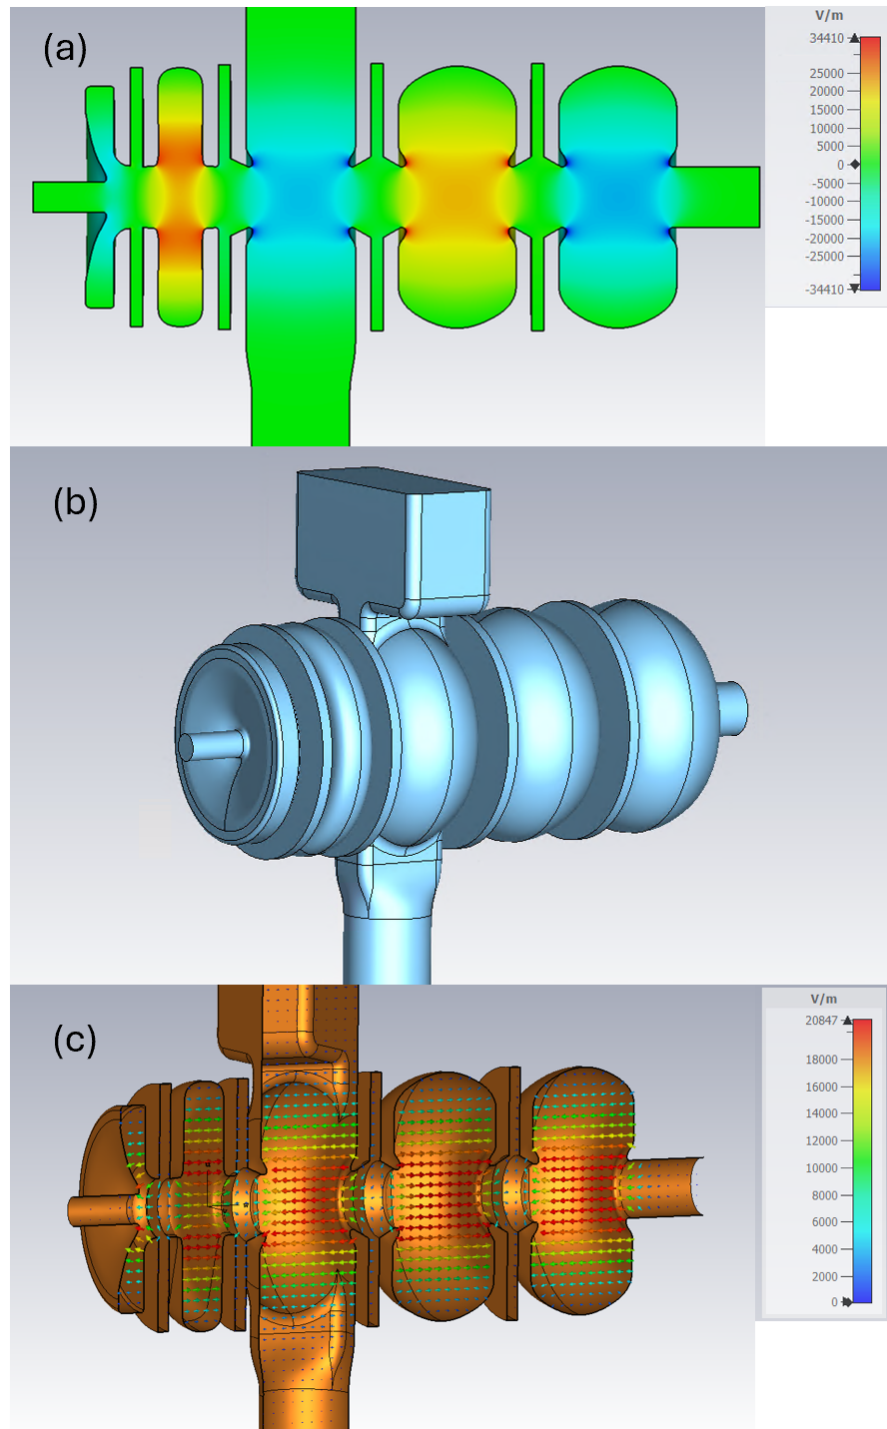


Figure 4.28: Final linac design. (a) Z-component of the electric field (b) Vacuum model (c) Absolute value of the electric field with copper walls.




Figure 4.29: Scan of β_c of a single cell model to find the critical coupling, leading $\beta_c = 1$ (critical coupling) for the full cavity.

the combined fields used for the Astra optimisation is shown in Fig. 4.30. The fields align very well, with the only differences being seen in the length and shape of the third cell and the decay of the last cell into the beampipe. A comparison of the optimisation outputs between the two simulations is shown in Table. 4.6, the values are very similar, with only a slight reduction of the total capture, and capture % > 1.5 MeV. There was also a 0.1 % increase in the number of electrons returning to the cathode and a 0.6 % in the number lost on the aperture. These increases were small, so the tuning was considered successful. These values could most likely be improved by changing the β of the third cell slightly to compensate for the decrease in length, due to the shifting of the zero crossing point.

Parameter	Individual fields	Tuned CST fields
Ave K.E > 1.5 MeV (MeV)	1.93	1.92
Total capture (%)	92.4	91.7
Capture % > 1.5 MeV (%)	84.5	83.7
Electrons lost $Z < Z_{min}$ (%)	2.7	2.8
Electrons lost Aperture (%)	4.8	5.4

Table 4.6: Single vs. tuned field comparison for $E = 45.13$ MV/m.

Fig. 4.31a shows the spectrum for the optimised Astra simulations, and Fig.





Figure 4.30: Comparison of stitched 1D fields used for optimisation and final tuned fields.

4.31b shows the spectrum of the electrons obtained at the end of the imported CST fields. Both show a similar structure, with the majority of the electrons falling between 1.9 and 2.1 MeV and a second small peak around 1.5 MeV.



(a) Histogram of final electron energies from optimised fields in Astra. (b) Histogram of final energies obtained from imported tuned CST fields.

4.8 Mode Tuning

Once the amplitudes and frequencies of the accelerating cells had been tuned, the coupling cells needed to be tuned to close the stop-band between the two $\pi/2$ modes. To tune the coupling cells, eigenmode simulations with magnetic boundary conditions were used on two halves of the accelerating cells, as shown in Fig. 4.32. For each pair of accelerating cells, the radius of the coupling cell was changed so that the mode inside the coupling cell had a frequency of 5.712 GHz.

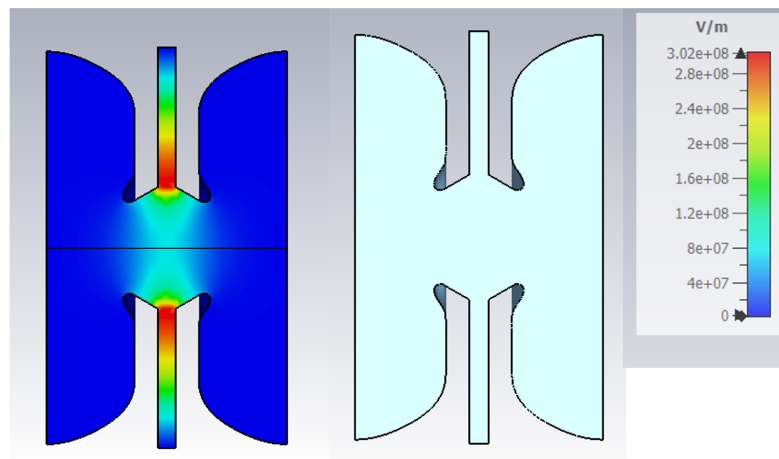


Figure 4.32: Model used to initially tune coupling cell $\pi/2$ mode.

Fig. 4.33 shows the S-parameter results from before and after the tuning on the coupling cells. Before tuning, the modes were not symmetric about the operating mode because the mode frequencies were shifted by the detuned coupling cells. The mode at 5.7044 GHz was also quite close to the operating mode of 5.712 GHz, which could cause issues as it may be excited when powering the structure. After tuning, the modes are equally spaced around the centre frequency, as desired. As power is coupled into these simulations from an accelerating cell, in order to check the modes that are excited inside the coupling cells, an eigenmode simulation had to be performed. The modes found from this simulation are shown in Fig. 4.34, where there are now 9 modes. All the modes, aside from the first and last, closely follow

the analytical dispersion curve, given by:

$$\omega_q = \sqrt{\omega_a^2 / (1 + k_c \cos(\pi q / 2N))} \quad (4.48)$$

where $N = 2$ for a 9 cell structure and q is the mode number. This indicated that the stop-band was closed. The first and last modes are far removed from the main curve because the coupling from the first to the second cell is very high to allow for a low amplitude first cell. Unfortunately, this separation is unavoidable, but it was anticipated that this would not cause any problems during linac operation.




Figure 4.33: S11 vs. frequency and location of modes from frequency domain solver, before and after tuning.

4.8.1 Beam Capture Tolerances

The resilience of the design to the tolerances on the cell lengths and field amplitudes was then explored to see how sensitive the design was. The three quantities explored were: the total capture efficiency, the capture efficiency of electrons with energies > 1.5 MeV, % of electrons that back-stream to the cathode and the average kinetic energy of the beam. The lengths and amplitudes of the first three cells were linearly scanned and the quantities of interest evaluated. The results of the amplitude scan of the first cell are shown in Fig. 4.35.




Figure 4.34: Plot showing a comparison of the 9 modes in the structure with the analytical dispersion for a 9-cell structure.

If the amplitude is too small, there is only a small effect on the capture and the number of back-streaming electrons. However, this is not true if the cell amplitude is too large. The % capture > 1.5 MeV and the total capture decrease rapidly with only a +10-15 % error on the amplitude. This needed to be kept in mind if the design were to be manufactured to ensure that the amplitude was as close to the working design point as possible to avoid degradation of linac performance. A similar story is seen in Fig. 4.36, where the results of the scan of β for the first cell are shown. There is little effect on the outputs when the cell is shorter than the design point, but the capture starts to drop off if the cell is ≈ 10 % too long, along with a doubling of the number of back-streaming electrons. It also became apparent that there was a reasonable variation in the results ± 0.4 % with only minor changes in length due to noise, which can be seen in the plot of the back-streaming electrons, as the design point falls below most of the other points close by. This meant that the final design would likely end up with a capture close to 3.2 % rather than 2.8 %.

The cell lengths and amplitudes of the second and third cells were also scanned, but they were relatively resilient for errors up to 15 %. The only parameter that

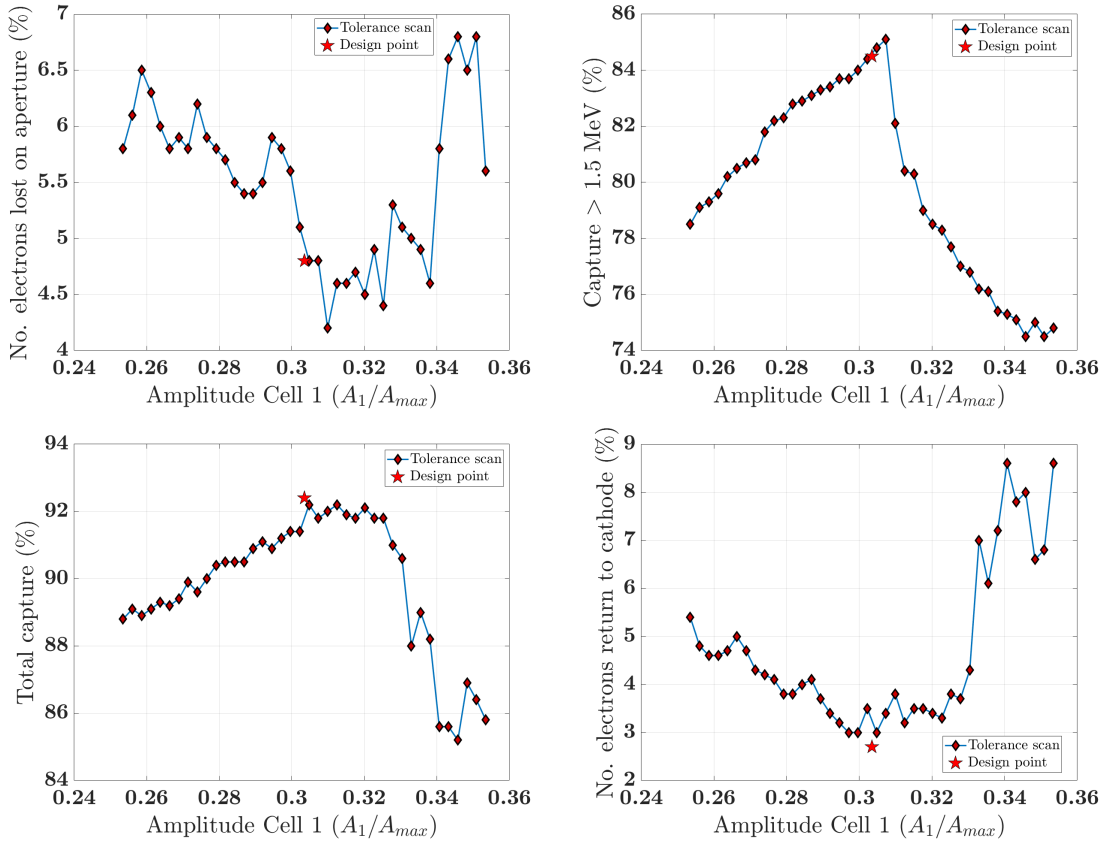


Figure 4.35: Effect of changing the amplitude of the first cell on they key outputs of interest.

was significantly affected was the capture of electrons with energies > 1.5 MeV as shown in Fig. 4.37. A positive error on the cell length causes electrons to slip behind the main bunch and end up being accelerated to a lower energy. A negative amplitude error on both cells, but particularly in the second cell, causes the peak of the recaptured electrons to fall below 1.5 MeV, explaining the sharp drop. As it was assumed that tuning the first cell would be the most challenging section of the linac, an investigation into mitigation strategies was made to see if errors in the first cell could be compensated for using the power supplied to the linac. The results of these scans are shown in Fig. 4.38. All three curves for both the total capture and back-streaming electrons have very similar shapes; this means that if there is, for example, a -10 % error in A_1 then this can be compensated for by increasing the power. This relaxes the tuning requirements and means that the design offers some

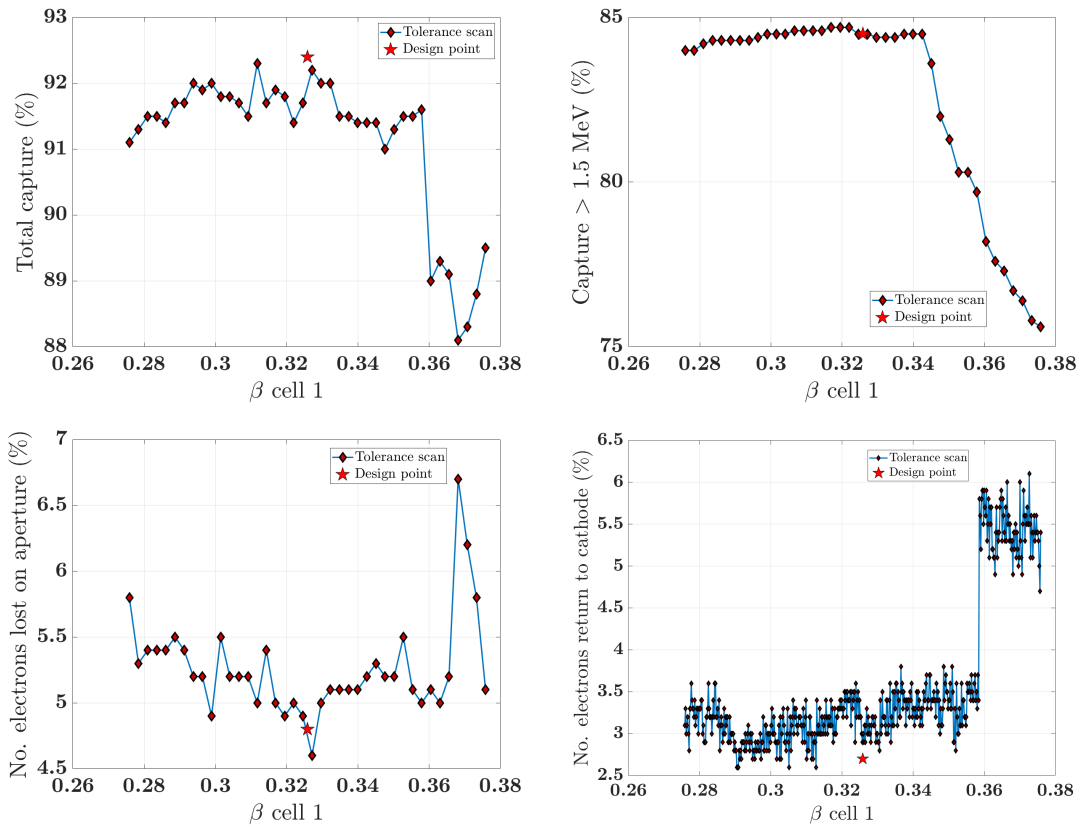


Figure 4.36: Effect of changing the length of the first cell on they key outputs of interest.




Figure 4.37: Effect of errors on cell 2 and cell 3. Cell length error (left) and cell amplitude (right).

flexibility, as the high capture can mostly be recovered using the gradient. This, of course, comes at the expense of higher peak fields for an increase in power and a

mean energy less than 2 MeV for a decrease in power, so this method would not work for very large errors. Similar scans of the power were performed for the tuned CST fields, with the results being very similar. These scans also show that the design is relatively resilient to small changes in the power, and what matters most is that the ratio of the amplitudes in the cells, especially in the first cell, is as close to the design point as possible.

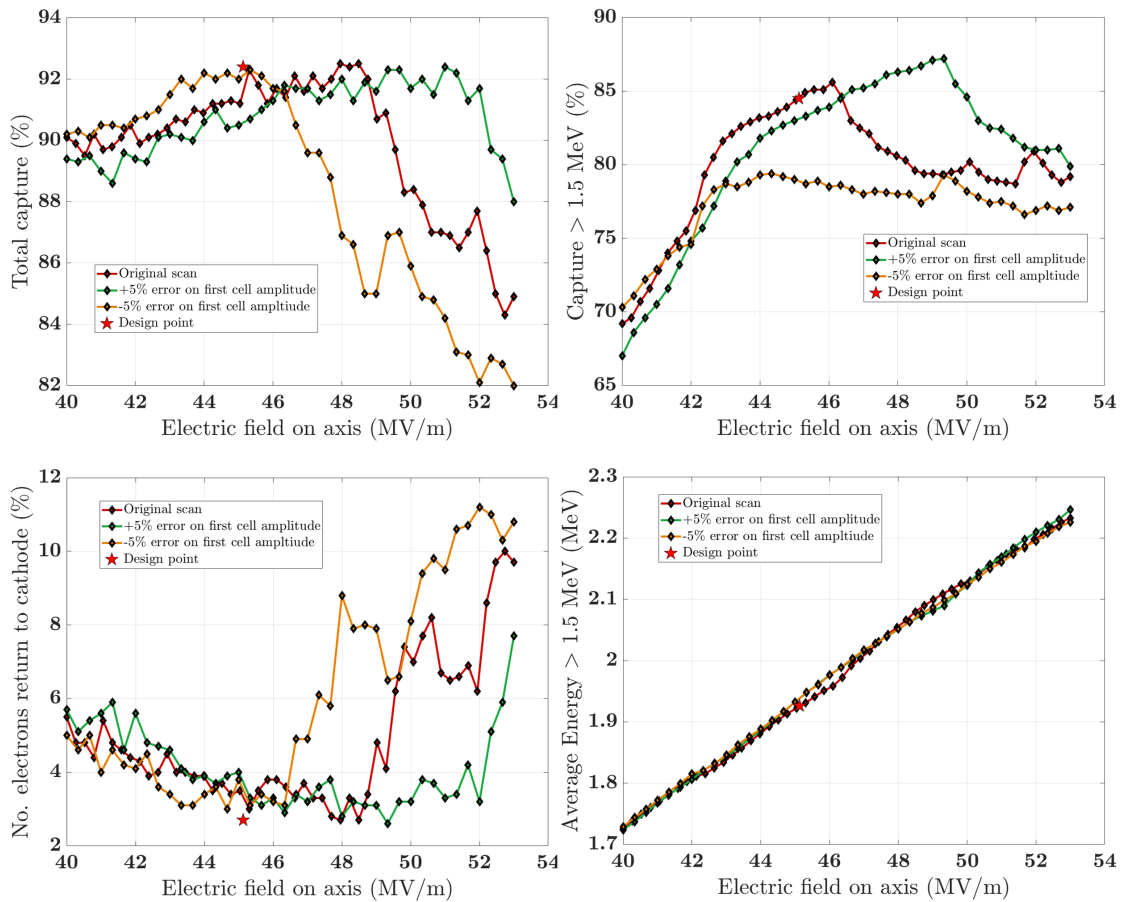


Figure 4.38: Scan of field on axis for a $\pm 5\%$ error on the first cell amplitude.

4.9 Particle-in-Cell Design Verification

As the Astra optimisation was performed in 1D, a further design verification was performed in the particle-in-cell (PIC) solver of CST studio suite. Both the electric

and magnetic strength field maps from the tuned CST model were exported on a grid in 3D, imported into the PIC solver and scaled to the required accelerating gradient to match the Astra simulations. These fields are then interpolated onto the hexahedral mesh used in the PIC solver. The imported fields are shown in Fig. 4.39.




Figure 4.39: Imported fields from tuned CST model in PIC solver. (Top) Electric field, (Bottom) magnetic field.

The input macro particle distribution used for the Astra simulations were modified in order to transform it into the (*.pit) file format that CST requires. Each row represents the particle number, and there are 10 columns: coordinates x , y , z , momenta $(\beta\gamma)_x$, $(\beta\gamma)_y$, $(\beta\gamma)_z$, particle mass, charge, macro-particle charge, and time. 100k macro-particles were used, with a macro-particle charge of 1×10^{-15} . The input distribution was the same as the one used for the ASTRA simulations, and was 1 ns long, with a current of 100 mA and a 25 keV energy. The mesh was set to 60 cpw, leading to 3.6M mesh cells. The simulation was run for 3 ns, enough time for there to be no particles remaining at the end of the simulation. Phase-space monitors and particle position monitors were set up to look at the bunch formation, along with a 2D monitor at the end of the linac to measure the energy and number of particles that reach the end. A shot from the particle position monitor is shown

in Fig. 4.40, showing the bunches forming and confirming that the final energy is around 2 MeV at the end of the linac.




Figure 4.40: Particle position monitor snapshot showing bunch formation and exiting bunch at 2 MeV.

Parameter	Tuned CST fields in Astra	3D PIC simulation
Ave K.E > 1.5 MeV (MeV)	1.92	1.95
Total capture (%)	91.7	91.8
Capture % > 1.5 MeV (%)	83.7	82.1
Electrons lost $Z < Z_{min}$ (%)	2.8	4.2
Electrons lost Aperture (%)	5.4	4

Table 4.7: Tuned field 1D Astra simulation vs. full 3D PIC results for $E = 45.13$ MV/m.

Fig. 4.41 shows the bunching structure recorded at the end of the linac on a 2D particle monitor with a 0.01 ns resolution time. The tails of the bunches are clearly seen with some of the electrons slipping into the preceding bunches. The final energy distribution of electron energies at the end of the linac is shown in Fig. 4.42. It is reasonably similar to those obtained using Astra shown in Fig. 4.31a and 4.31b, although the second peak at 1.5 MeV is missing and is now distributed evenly at lower energies. Table. 4.7 gives a comparison of the results obtained from the 1D Astra simulations and the 3D PIC simulations. The average kinetic energies and total capture align well, although the PIC simulation predicts 4.2 % electrons returning to the cathode, as opposed to 2.8 % from the 1D field maps. This is most likely due to the CST simulations including the entirety of the decaying field into

the beampipe, whereas the Astra field import was cut off before the full decay. This issues could be rectified by remodelling the fields with the correct decay into the beampipe and then repeating the optimisation. The key result was that the total capture efficiency is consistent, so the maximum percentage of electrons that can return to the cathode is 8 %, still a significant improvement over common industrial linac designs.




Figure 4.41: Particle energy vs. time recorded at the end of the linac showing the bunches leaving the structure at 2 MeV.




Figure 4.42: Histogram of the beam at the end of the linac in the PIC simulation.

Chapter 5

X-ray Target Design

The next design step after the final linac design had been completed was the design of the electron-photon conversion target. X-ray targets generate Bremsstrahlung radiation through the Coulomb interaction between the incident electron beam and the nuclei of the materials in the target. The electrons in the electron beam are decelerated and this lost kinetic energy is given off in the form of heat and photon radiation. As the spectrum produced by the target is a function of the electron beam energy, the thickness of the target, and the target material atomic number (Z_a number), these are typically the main design parameters of the target.

5.1 Target Theory

For a thick target where the range of electrons (the greatest distance of penetration in a given direction) is much smaller than the thickness of the target, the intensity of the spectrum produced from the target can be written as follows

$$I(h\nu) = CZ_a(E_b - h\nu) \quad (5.1)$$

Where E_b is the kinetic energy of the electron beam, Z_a is the atomic number of the target, $h\nu$ is the photon energy and C is a proportionally constant [105]. Different

materials have different electron-photon conversion efficiencies at different energies, with some example materials and conversion efficiencies shown in Fig. 5.1. Typically, targets are made from a high Z material to maximise the photon yield, with tungsten being a popular choice due to its high melting point (3422 °C) meaning that it can withstand a large temperature rise due to the energy lost as heat from the electron beam [106]. Because of this, tungsten was the main material investigated for the target design.

Electron radiation yield, $Y(T)^a$									
T (MeV)	C	Water	Air	Al	Fe	Mo	W	Pb	U
0.1	0.0005	0.0006	0.0007	0.0014	0.0031	0.0054	0.0103	0.0116	0.0131
0.2	0.0009	0.0010	0.0011	0.0022	0.0052	0.0093	0.0187	0.0212	0.0242
0.5	0.0018	0.0020	0.0022	0.0043	0.0100	0.0179	0.0371	0.0424	0.0491
1.0	0.0034	0.0036	0.0040	0.0764	0.0170	0.0297	0.0603	0.0684	0.0792
1.5	0.0050	0.0053	0.0058	0.0110	0.0239	0.0410	0.0802	0.0901	0.1035
2.0	0.0067	0.0071	0.0077	0.0145	0.0310	0.0519	0.0986	0.1096	0.1249
4.0	0.0142	0.0149	0.0158	0.0292	0.0595	0.1136	0.1625	0.1761	0.1955
6.0	0.0222	0.0233	0.0242	0.0444	0.0874	0.1325	0.2157	0.2304	0.2518
8.0	0.0304	0.0319	0.0327	0.0596	0.1139	0.1677	0.2612	0.2765	0.2990
10.0	0.0387	0.0406	0.0411	0.0745	0.1389	0.1999	0.3006	0.3162	0.3394
15.0	0.0595	0.0622	0.0618	0.1105	0.1951	0.2689	0.3800	0.3955	0.4193
20.0	0.0798	0.0833	0.0817	0.1438	0.2435	0.3252	0.4403	0.4555	0.4790
50.0	0.1856	0.1920	0.1825	0.2959	0.4328	0.5249	0.6316	0.6439	0.6635
100.0	0.3181	0.3190	0.3022	0.4448	0.5848	0.6673	0.7526	0.7617	0.7766

^a $Y(T)$ is the fraction of the initial kinetic energy T lost as bremsstrahlung. The yields do not account for the minor effects of radiation losses from energetic secondary electrons (delta rays) produced during the deceleration of the primary electron.

Source: [ICRU 1984].

Figure 5.1: X-ray conversion efficiency for different energies and materials [107].

5.1.1 X-ray Production Processes

When the electrons in the beam hit the target, some of them will come very close to the nucleus, causing a deceleration and change in direction due to the positive charge in the nucleus. This process converts the kinetic energy of the incident electrons into radiation, and is called breaking radiation or bremsstrahlung. The closer the incident electron is to the nucleus, the higher the X-ray photon energy, but the probability of an electron being closer to the nucleus also decreases with decreasing distance. This gives rise to a characteristic X-ray spectrum as shown in Fig. 5.3, with a high intensity at lower energies and a much reduced intensity as the curve approaches the values of the highest photon energy.




Figure 5.2: Description of the bremsstrahlung process. Image sourced from [108].

Fig. 5.3 also shows examples of characteristic X-rays that are generated when energetic incident electrons interact with the electrons in the outer electron shells of the target material and knock these electrons out of the shells. The target atoms then become unstable, and electrons from outer shells transition to fill the gaps and release radiation in the process, thereby giving peaks in the X-ray spectrum.




Figure 5.3: Typical X-ray spectrum showing characteristic X-rays. [109].

5.2 G4beamline Simulation Setup

In order to optimise the target with respect to the beam coming from the final linac design, the beam parameters were exported at the end of the linac in ASTRA. Using this beam, a Monte Carlo simulation was setup using G4beamline, which is a Geant4-based particle tracking software [110]. Geant4 simulates the passage of particles through matter by splitting the trajectories into discrete steps. At each step, it calculates the probability of various interaction processes occurring based on the properties of the particle and the material it is passing through. These probabilities are determined using Monte Carlo sampling techniques. This allows it to solve a multitude of problems involving particles and matter including detector simulations for high-energy physics, energy deposition problems or for solving sections of accelerator beamlines.

A histogram of the beam size exported at the end of the linac is shown in Fig. 5.4. The momentum of the beam from Astra was split into 500 bins, ranging from 0-2.8 MeV/c, and the number of electrons residing in each bin was found. Using this with the parameters shown in Table. 5.1, a beam was constructed in G4beamline in Gaussian format using the parameters exported from Astra.

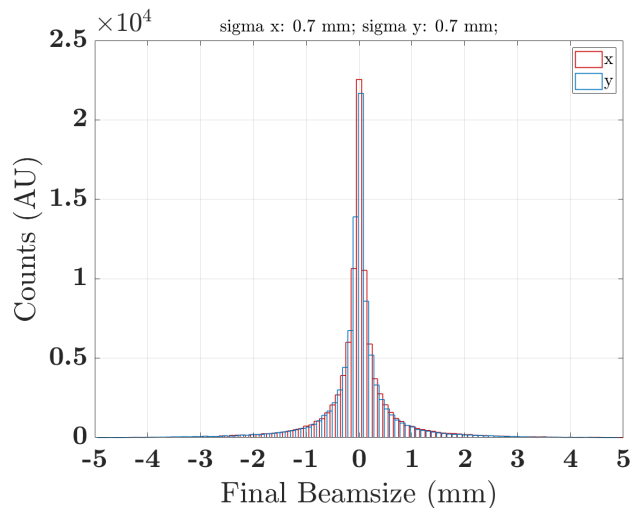


Figure 5.4: Histogram of the final beam size at the end of the linac.

The rest of the geometry was setup as shown in Fig. 5.5. The target is a

Quantity	Value	Units
σ_X	0.7558	mm
σ_Y	0.7312	mm
σ_{XP}	5.75	dxdz
σ_{YP}	5.5683	dydz

Table 5.1: G4beamline gaussian beam parameters.

composite style target consisting of two disks, one made from tungsten, and one from copper. Although tungsten has a high melting point, its thermal conductivity is 174 W/mK at 20 °C compared to coppers 401 W/mK, therefore to effectively remove the heat deposited by the beam, a copper backing is used. It is typical practice to measure X-rays at a distance of 1 m from the source, so the detecting region was placed 1 m from the target, as a 1 m diameter disk.

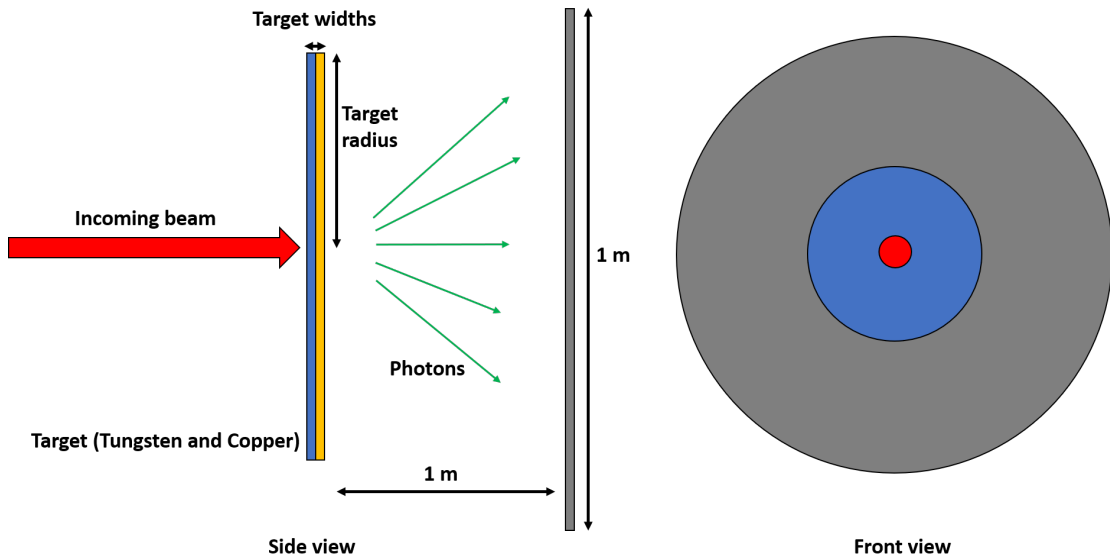


Figure 5.5: Simple diagram of target simulations showing composite two layer target, incoming beam and X-ray detection area.

A visualisation of the process of the beam hitting the target is shown in Fig. 5.6, where the red trajectories are the electron paths, which are mostly scattered from the target and end up going in the reverse direction. The green lines represent the photons, and the cross-sectional view in Fig. 5.6b shows where the beam is generated and the paths of some of the photons as they interact with the tungsten

and copper.

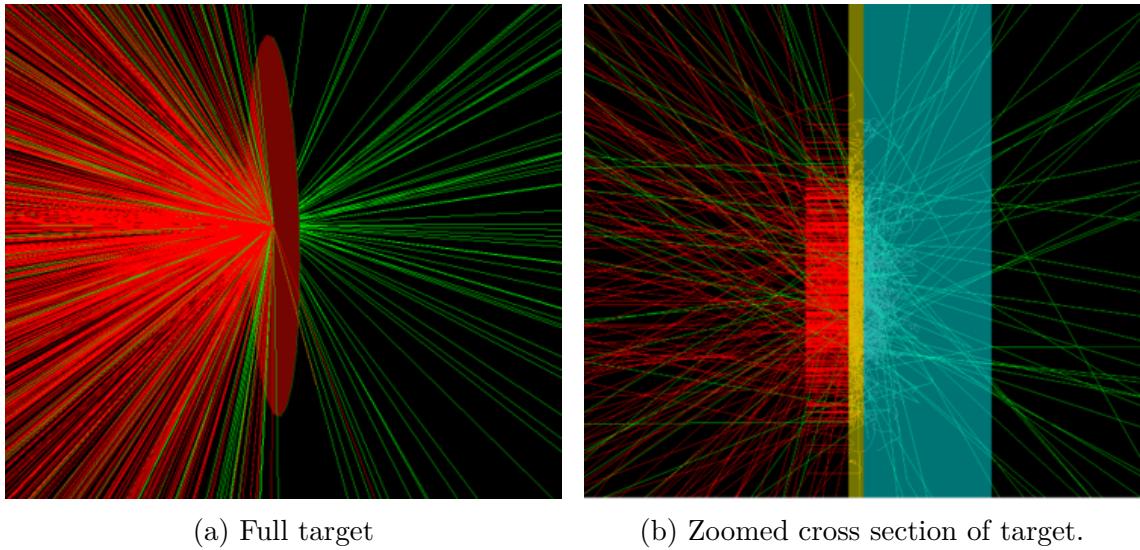


Figure 5.6: Images of simulation from G4beamline, red lines represent electron trajectories and green lines represent photon trajectories.

Fig. 5.7 shows the low energy section of the spectrum of a 2 MeV Gaussian beam hitting a tungsten target, the various peaks are given by the different characteristic energies obtained from the various shell vacancies being filled. A description of the process that gives rise to each peak is given in Table. 5.2




Figure 5.7: Low energy section of X-ray spectrum of a tungsten target showing characteristic X-rays.

Letter	Energy (keV)	Description
A	8.9	L-series peak
B	58	K-series peak ($K_{\alpha 1}$)
C	59.3	K-series peak ($K_{\alpha 2}$)
D	67.2	K-series peak ($K_{\beta 1}$)
E	69.1	K-series peak ($K_{\beta 1}$)
F	511	Electron-positron annihilation

Table 5.2: Various characteristic X-ray peaks for a tungsten target.

5.3 X-ray Target Optimisation

The only parameters available to optimise the target were the thickness of each of the layers. The tungsten needed to be tuned to optimise the X-ray output of the target, and the copper needed to be tuned as a trade-off between the target heating and X-ray output. If the target is too thin, then there is not enough opportunity for the incident electrons to interact with the target atoms and many electrons will exit the target. As the target becomes thicker, photons generated at the front of the target end up giving up their energy or scattering before reaching the detector; therefore, there should be an optimum thickness for the target. Ideally, the electron leakage from the target should also be minimal, so detectors were also set up behind the target to measure this.

Three thicknesses of the copper backing were investigated, and for each copper thickness a scan of the tungsten thickness carried out. An example of the spectra generated for one of the scans is shown in Fig. 5.8. The figure shows that increasing the thickness of the tungsten reduces the number of photons at lower energies, reducing the total X-ray yield. The results showing the X-ray yield from all the scans are shown in Fig. 5.9a, showing that the yield increases from 0.05 mm to 0.2 mm for all three copper thicknesses and then starts to drop off. The electron leakage is shown in Fig. 5.9b, for both the 1.5 mm and 2 mm copper backings there is a constant electron leakage, with an increase only seen when the copper backing is 1 mm and the tungsten thickness is < 0.2 mm.






Figure 5.8: Example of spectra for a tungsten (W) thickness scan with a 1 mm copper (Cu) backing. Np/Max represents the number of photons normalised to the maximum number from the scan.

5.4 Copper Target

Before the target design was finalised, an investigation was carried out to determine whether a pure copper target could be used. This would improve the cooling of the target, as it would be a solid piece, improving the heat flow from the target, rather than having two pieces with a thermal contact resistance between them. The results of the scan are shown in Fig. 5.10. The target needed to be at least 1.25 mm thick in order to minimise the electron leakage, but the maximum X-ray yield occurred around 0.75 mm. A comparison of the generated spectrum for the copper only target and the tungsten-copper target is shown in Fig. 5.11, the X-ray yield at low energies < 0.2 MeV is higher for the copper only target, but for the rest of the spectrum, the tungsten copper target provides a higher yield. As it was anticipated that the heating of the target would not be excessive due to the reasonably low



(a) X-ray yield vs. tungsten thickness for different copper backing thicknesses.



(b) Electron leakage measured behind the target vs. tungsten thickness for different copper backing thicknesses.

repetition rate, it was decided that the X-ray yield should not be compromised and that a tungsten-copper composite target should be used.




Figure 5.10: Results for a copper only target. (Top) X-ray yield vs. copper thickness (Bottom) Electron leakage vs. copper thickness

5.5 X-ray Target Heating

Although a 1.5 mm copper backing provided enough thickness to minimise the electron leakage, the target heating needed to be explored to find a thickness of copper suitable for removing the heat generated by the beam. The parameters used for the preliminary study are shown in Table 5.3. Ansys mechanical simulations [111] were set up in order to calculate the heating of the target. The initial geometry for the simulations is shown in Fig. 5.12, which is a 3.6° slice to take advantage of the symmetry of the problem. The section in the centre of the tungsten was cut out as it would be used to deposit the heat generated by the beam in the simulation, its




Figure 5.11: Comparison of spectra from a 1.75 mm copper target and a copper-tungsten composite target.

radius is representative of the FWHM size of the beam. The physics preference was set to non-linear mechanical, an element size of 1×10^{-4} was used, leading to 6200 elements for the slice.

Parameter	Value	Units
Copper thickness	1	mm
Tungsten thickness	0.2	mm
Pulse length	3	μs
Peak current	100	mA
Energy	2	MeV
Duty cycle	0.06	%
Beam FWHM	2	mm

Table 5.3: Parameters used for composite target heating study.

The peak power for the beam was calculated using

$$P_{pk} = V_b I_b \tag{5.2}$$

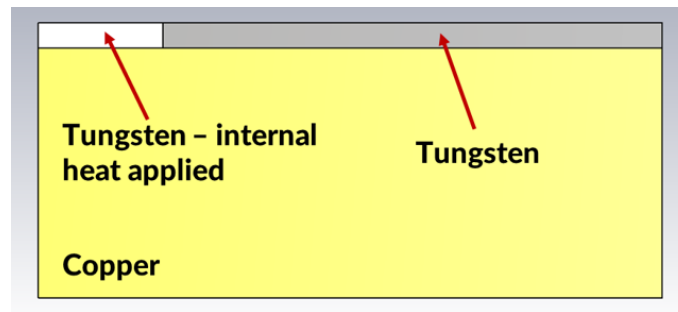


Figure 5.12: Geometry of initial target heating simulations showing tungsten layer, copper layer, and tungsten section where the heat is generated.

where V_b is the beam voltage and I_b is the peak beam current. This is then multiplied by the conversion efficiency of tungsten from Fig. 5.1, which is approximately 0.902, as only 9.8 % of the initial kinetic energy is lost as bremsstrahlung for a 2 MeV beam. This gave a power of approximately 180 kW, which was then converted into a power density using the volume of the section where the heat was to be applied.




Figure 5.13: Transient heating applied to element of tungsten to represent heating by the beam. (Top) A single pulse, (Bottom) multiple pulses at the 200 Hz repetition rate.

The temporal pulse distribution that was applied to the heating element is shown

in Fig. 5.13. The individual pulses are $3 \mu\text{s}$ long, and these come at the 200 Hz repetition rate. When the pulse is high, the peak power is applied to the central element. The thermal distributions at the start and end of the first and last pulses are shown in Fig. 5.14. At the start of the pulse, all of the heating is in the tungsten element, which then dissipates into the rest of the tungsten and copper, with the average temperature remaining low at the end of the first pulse. The maximum temperature is seen at the copper wall which is at the water temperature of $35 \text{ }^\circ\text{C}$. However, at the end of the last pulse, the average temperature of the target has risen at the beginning of the pulse, and at the end of the pulse, the $35 \text{ }^\circ\text{C}$ water is now providing cooling, as there is a temperature gradient across the entire target. The right boundary of the simulation is held at one temperature ($35 \text{ }^\circ\text{C}$) which is the assumed temperature of the cooling water that would be used to cool the target.

The transient thermal simulation was run until the average heating of the target reached a steady state value, with the maximum temperature in the target vs. time for the target shown in Fig. 5.15 for two different boundary temperatures. Every pulse there is a sharp rise in temperature and then a decay until the next pulse arrives; this decay reaches a higher minimal value with every pulse until the system reaches a point where the maximum and minimum target temperatures do not change at $\approx 0.1 \text{ s}$. Changing the boundary temperature reduces both the maximum temperature and the minimum temperature. The same water used to cool the linac was expected to be also used for the target, so it was likely that the water temperature would be between $25 \text{ }^\circ\text{C}$ and $35 \text{ }^\circ\text{C}$.

5.6 Energy Deposition Model

Although the cylindrical model for the deposition of energy into the target provided some insight into the target heating, a more accurate method was employed using the measured energy deposition in the target from G4beamline. Fig. 5.16 shows how the energy deposited in the target was measured. The target was split into

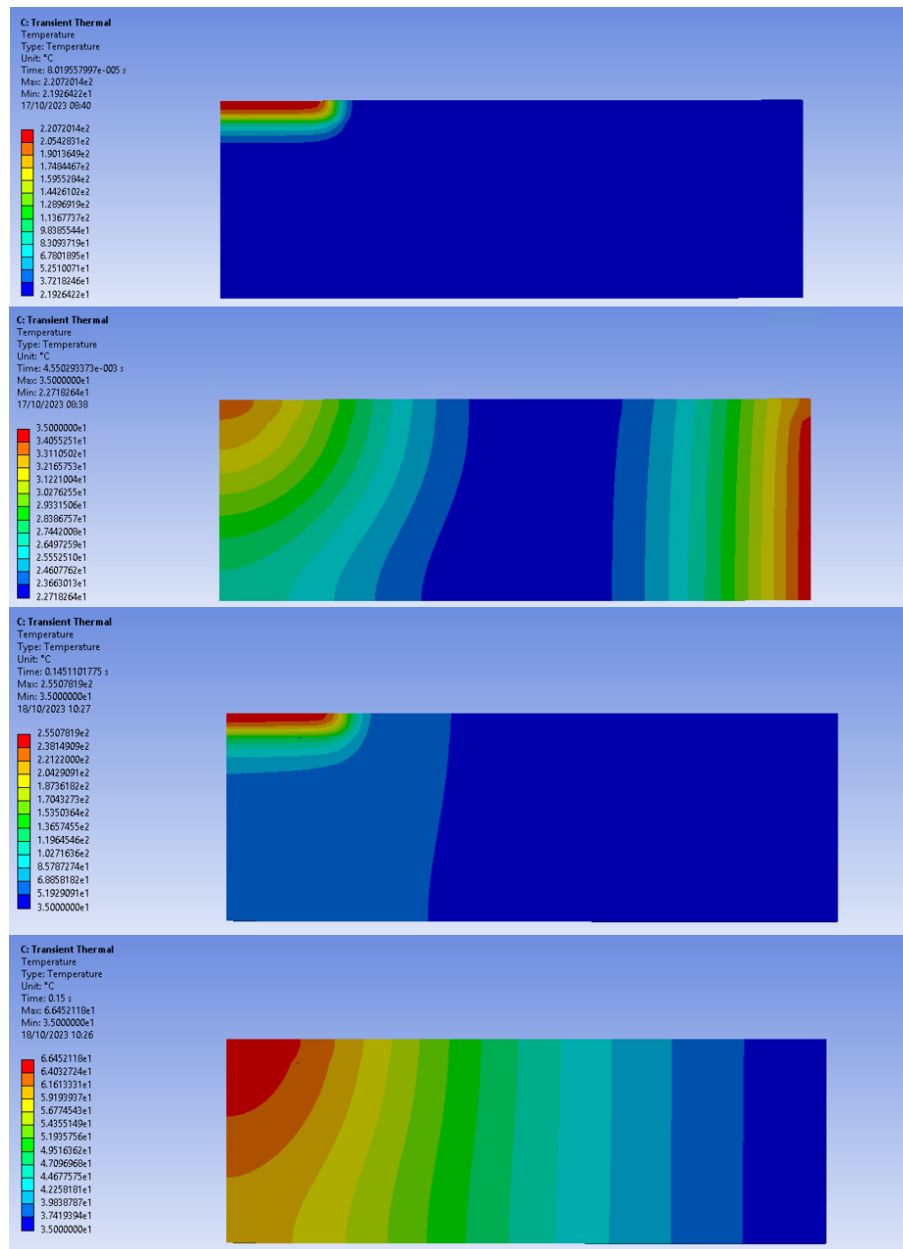


Figure 5.14: Heat distribution within the target. From top to bottom, start of the first pulse, after the decay of the first pulse, start of the last pulse, after the decay of the last pulse.

slices longitudinally and into concentric rings radially, with the energy deposited recorded for each section. The maximum energy is deposited around 0.05 mm into the target, with the radial deposition following a Gaussian shape. Deposition into copper is also shown, which decays to close to zero by 1 mm into the target.

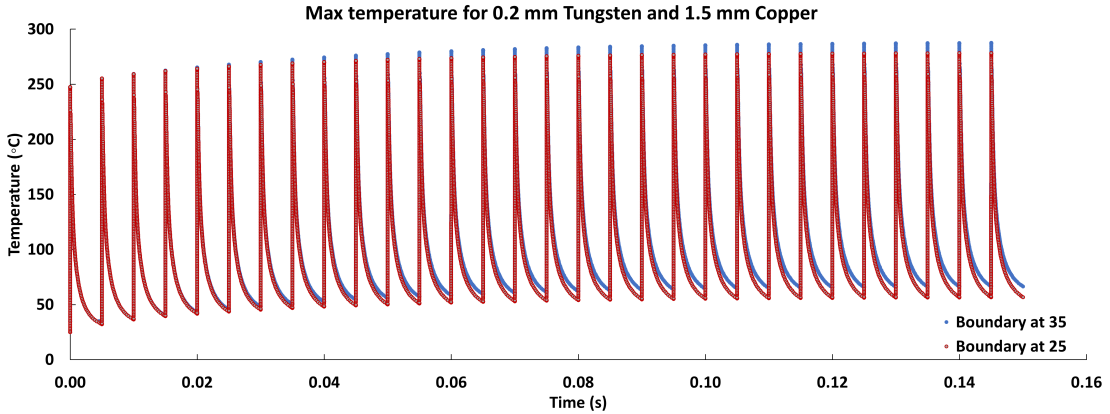


Figure 5.15: Maximum temperature in the target vs. time for two different boundary temperatures.

The radial deposition was fit to a Gaussian of the form

$$y(x) = Ae^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (5.3)$$

and then this curve was used to define the Gaussian geometry shown in Fig. 5.17. Then each section was given a power density corresponding to the measured energy deposition. A similar set of results to the initial cylindrical shape are shown in Fig. 5.18. The initial distribution decays to a very similar shape as the cylindrical model, despite the very different initial distribution. The key difference between the models becomes clear when the maximum temperature in the target vs. time are compared, as shown in Fig. 5.19. There is a nearly 30 °C difference in the maximum heating of the target, showing that the cylindrical model overestimates the maximum temperature by around 12 %.

5.6.1 Optimal Target Thickness

The temperature of the copper backing should be kept well below 600°C in order to minimize grain growth, and ideally, both the tungsten and copper should be kept below 1/2 of the melting point (so 540 °C for copper and 1700 °C for tungsten) [112]. With this in mind, a scan of the thickness of the copper backing was performed,

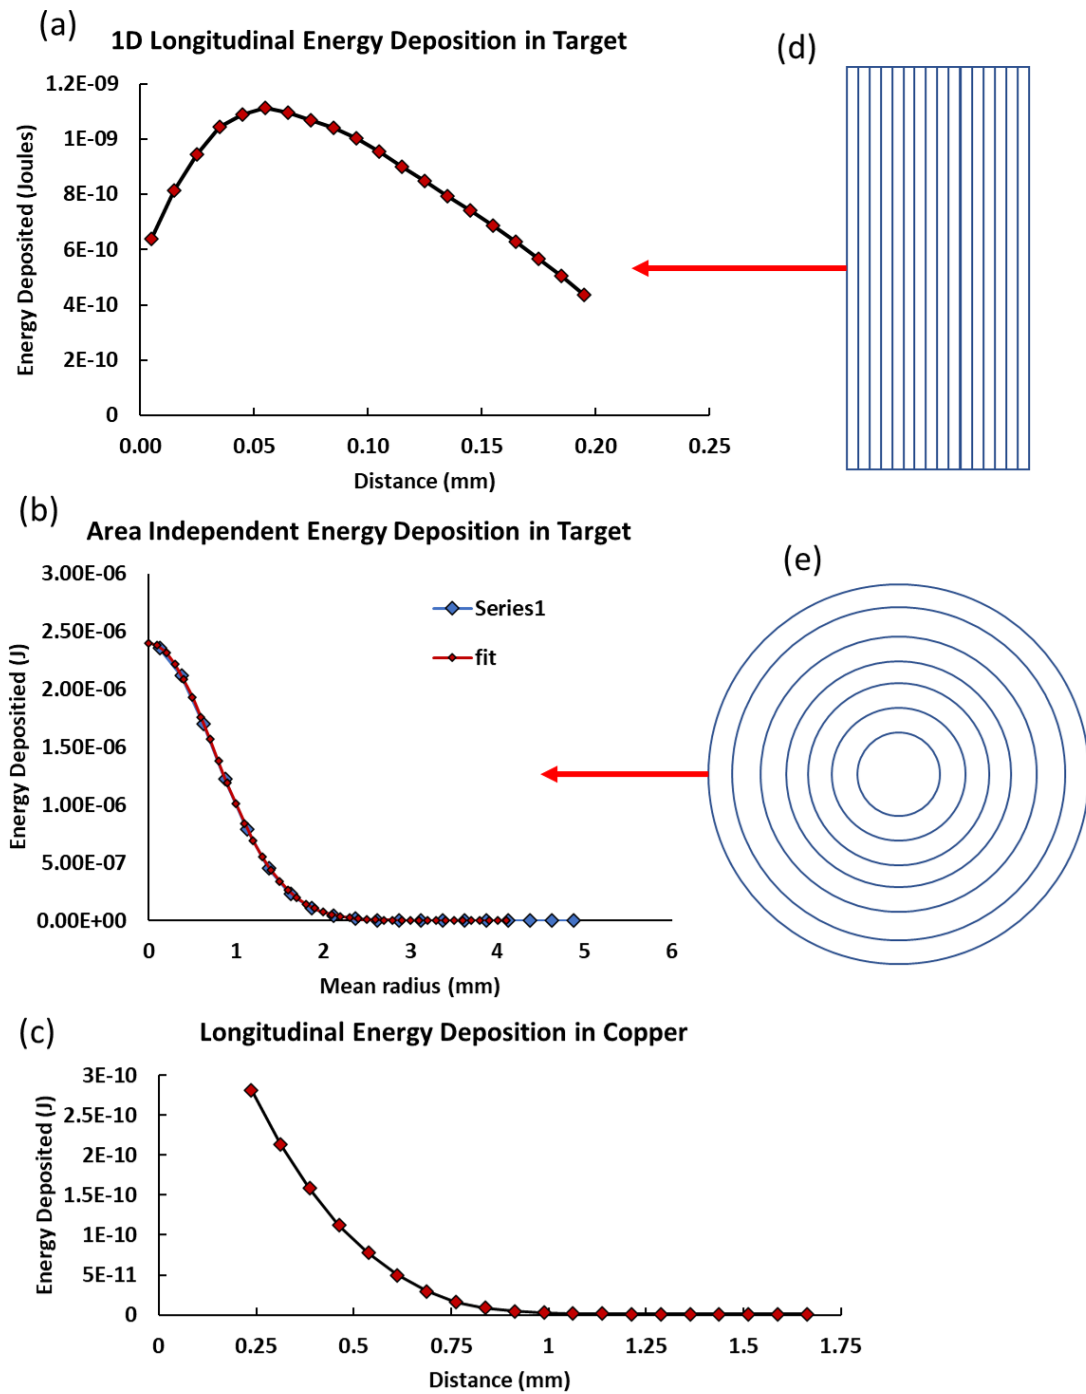


Figure 5.16: (a) Longitudinal energy deposition in the target. (b) Area independent energy deposition radially in target. (c) Longitudinal energy deposition in copper backing. (d) Longitudinal slices used to measure energy deposition. (e) Concentric segments used to measure radial energy deposition.




Figure 5.17: Gaussian heating model used for a more accurate assessment of target heating.




Figure 5.18: Gaussian heating thermal distributions. (Top) start of the last pulse, (Bottom) after the decay of the last pulse.

using the Gaussian deposition model, and the results are shown in Fig. 5.20. When the copper backing is only 1 mm thick, the maximum temperature reaches nearly 350 °C. Once the thickness is increased to 1.5 mm, this maximum temperature drops by ≈ 100 °C, but a further increase to 2 mm gives a much smaller drop in the maximum temperature. If the copper is too thin, then there is a bottleneck between where the bulk of the energy is deposited and the water cooling. Once the thickness reaches a certain threshold, this bottleneck disappears, and further thickness increases do not provide significantly more cooling.




Figure 5.19: Comparison of target heating for a single cylinder and the gaussian model.




Figure 5.20: Comparison of maximum target temperature vs time for three copper backing thicknesses, using a gaussian deposition model.

5.7 Final Design

The final thickness of the copper backing was chosen as 1.5 mm, sufficient to reduce electron leakage and provide enough thermal mass for cooling. As a check, to see the quality of the X-rays being generated, another G4beamline simulation was carried out, using a pure Gaussian electron beam centred at 2 MeV. This was then compared with the X-ray spectrum generated by the beam from the tuned CST field Astra model, as shown in Fig. 5.21. This comparison shows that the beam from the linac

provides an X-ray performance very similar to a Gaussian beam at 2 MeV, showing that the quality of the optimised beam would be sufficient for X-ray production.




Figure 5.21: Comparison of final X-ray spectrum 1 m away from the target for a Gaussian beam and the beam generated from the tuned CST Astra model.

Chapter 6

Linac Thermal and Mechanical Analysis

This chapter details all the thermal considerations that were made for the linac design. This includes an analysis of thermal effects during the manufacturing process, single-cell simulations to find the minimum acceptable cavity wall thickness, and a full thermal analysis, including computational fluid dynamics (CFD) studies to ensure the full linac heating during operation was at an acceptable level.

6.1 Basic Heat Transfer Theory

There are three different ways for heat to be transferred: conduction, convection, and radiation. Heat will always be transferred from a warmer area to a colder one, and with no sources, will stop when the temperatures of the two areas have equalised. The three processes are briefly described below:

6.1.1 Conduction

Conduction is the transfer of energy through particle collisions and diffusion of molecules during random motion. It can occur in solids, liquids, and gases. In solids, the energy is transported by free electrons. It is described by Fourier's law

of heat conduction

$$\dot{Q}_{cond} = -kA \frac{dT}{dx} \quad (W) \quad (6.1)$$

where \dot{Q}_{cond} is the rate of heat conduction, k is the thermal conductivity, A is the area normal to the heat flow and dT/dx is the temperature gradient. This reveals that the rate of heat flow is dependent on the thickness of the material, the surface area, the difference in temperature between the two regions and the type of material. The thermal conductivity of a material is defined as the rate of heat transfer through a unit thickness of material per unit area per unit temperature difference and can vary widely between materials, for example the difference between copper and air at room temperature is four orders of magnitude [113].

6.1.2 Convection

Convection is the mechanism of heat transfer of a solid surface and a fluid or gas moving over it, combining the effects of conduction and the movement of the fluid [113]. There are two types: forced convection, where a fluid is forced over a surface externally, such as using a pump to move water or a fan to blow air, and natural convection, where the flow is caused by natural density differences in the fluid due to temperature gradients. It is described by Newton's law of cooling

$$\dot{Q}_{cond} = hA_s(T_s - T_\infty) \quad (W) \quad (6.2)$$

where h is the convection heat transfer coefficient (HTC), A_s is the surface area, T_s is the temperature of the surface and T_∞ is the bulk fluid temperature far from the surface. The value of the HTC depends on several factors including surface geometry, fluid properties, fluid velocity, and the type of fluid motion [113].

6.1.3 Radiation

The final mechanism is radiation, which is the transfer of energy through EM waves. It does not require the surfaces to be touching and is an important consideration for the thermal loads of superconducting cavities [114]. The heat exchange between two surfaces is the balance between the radiation emitted from the first and absorbed by the second and the radiation emitted by the second and absorbed by the first, it is given by [115]

$$Q_{1-2} = \epsilon\sigma(T_1^4 - T_2^4)A_1F_{12} \quad (W) \quad (6.3)$$

Where Q_{1-2} is the heat transfer from surface 1 to surface 2, ϵ is the emissivity of the surface (between 0 and 1), σ is the Stefan-Boltzmann constant, A_1 is the surface area of the first surface, and F_{12} is the geometric view factor which is the fraction of the total radiation leaving the first body being absorbed by the second, it depends on the orientation of the surfaces and how well they can 'see' each other.

6.2 Thermal Effects During Manufacture

An important aspect considered when deciding on the final thickness of the wall was the potential thermal effects, both from the cavity manufacturing process and from heating during operation. Historically, many methods of forming normal conducting RF structures have been used, including: machining, rolling, forging, spinning, and hydroforming [116]. Machining is generally preferred as it is more accurate, with recent advancements in precision machining using diamond tools giving μm precision [117]. Once the cells have been made, they need to be joined, and here a number of techniques also exist, the key ones being brazing, diffusion bonding, electroplating, electron-beam welding, and inert-gas arc welding [116]. Both types of welding can cause defects in the joints that can limit the gradient and lead to increased breakdowns; therefore, brazing and diffusion bonding are generally preferred for normal conducting structures, and newer methods are also being investigated, such

as Tungsten Inert Gas (TIG) welding [118]. Currently, a typical copper cavity is manufactured by first machining a block of oxygen-free high thermal conductivity (OFHC) copper to create one cell of the structure. These cells are then stacked and combined using brazing or diffusion bonding [116]. These methods can also be used to attach two halves of a full structure instead of individual cells if required. Brazing requires the use of a filler material such as BVAg-8 (72 % silver, 28 % copper, $T_m = 779^\circ\text{C}$) that has a lower melting temperature than copper [119]. It is placed around the edges of the two halves, and the structure is heated, melting the filler material that then flows into the joints through capillary action, electrically connecting the structure's two halves. A basic schematic of this method is shown in Fig. 6.1. Diffusion bonding, on the other hand, requires the copper to be held at




Figure 6.1: Schematic of basic brazing technique.

a slightly lower temperature of approximately 1030°C , but for significantly longer (1 hr 30 vs. 20 min). Large pressures are also applied to bond the halves together which can increase the risk of deformation and material evaporation. In addition, a good surface finish, on the order of nm is required for diffusion bonding, increasing the cost. As diffusion bonding does not require a filler metal between the joining surfaces, potentially dangerous residues contaminating the cavity and bad joints can be avoided. It is normally more costly to manufacture the cavity using diffusion bonding, and it was known that high gradients would not be required, so for this

design, it was assumed that the process of brazing would be used. A typical brazing




Figure 6.2: Typical temperature vs. time brazing curve [120].

temperature curve is shown in Fig. 6.2. The structure is heated in several stages, eventually reaching a temperature very close to its melting point of 1084 °C. It is well known that when metals are heated to high temperatures and subjected to low stresses, they can undergo diffusional flow, which could cause a deformation of the nose cone in the design [121]. It was assumed that this effect would be more significant if the wall thickness was very thin, reducing the load handling capacity of the wall, and therefore simulations would need to be performed to validate this assumption and find how much deformation would occur. Fig. 6.3, shows the two types of diffusion that can occur at high temperatures and low stresses, grain boundary diffusion [122] (Coble creep) at lower temperatures and lattice diffusion (Nabarro-Herring creep) [123] at higher temperatures. The formulas for the strain rates for both types of diffusion are given in Fig. 6.4. It is interesting to note that Coble creep has a dependence of $1/d^3$ on the grain size within the copper and Nabarro-Herring has a dependence of $1/d^2$, meaning that the grain size and the way the material is initially prepared has a large effect on the strain rates experienced during brazing.

In order to perform simulations of this non-linear process, the temperature dependent properties of OFHC copper had to be obtained so that they could be used as material properties in the simulations. The first properties of the materials

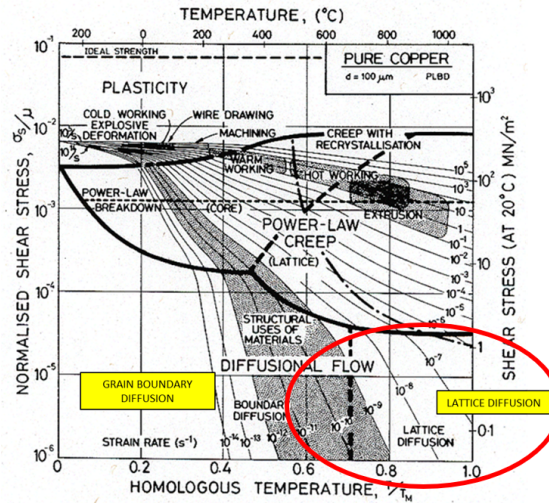


Figure 6.3: High temperature diffusion mechanisms [121].

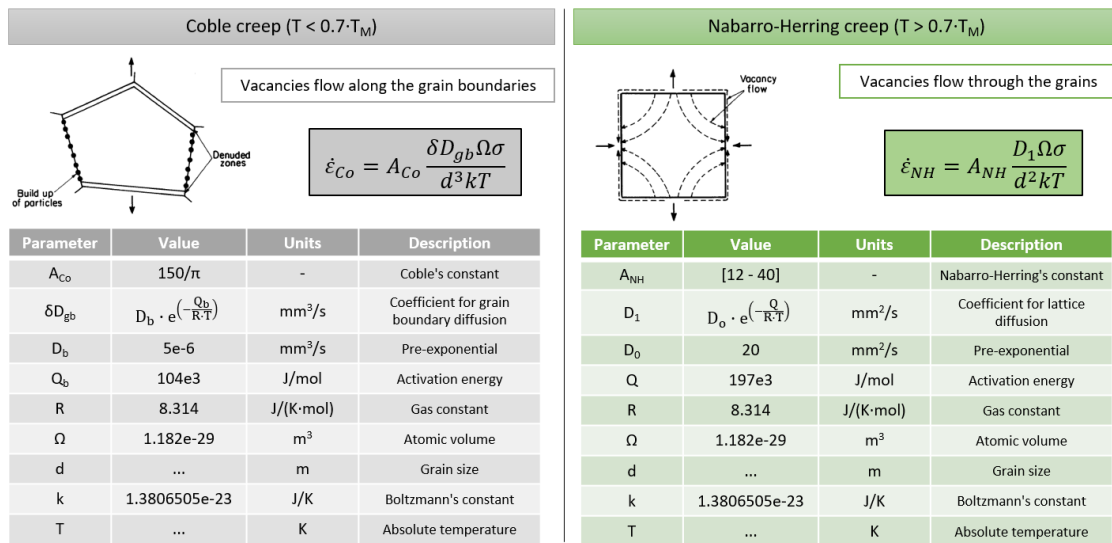


Figure 6.4: Low pressure, high temperature diffusion mechanisms. Image sourced from [124].

obtained where the temperature-dependent diffusion coefficients [125, 126] shown in Fig. 6.5. From this, the strain rates for both types were calculated using the formulas in Fig. 6.4, we can clearly see that Nabarro-Herring creep becomes the dominant mechanism when $T > 760 \text{ }^\circ\text{C}$ as expected. In order to use this information in Ansys




Figure 6.5: Diffusion coefficients (left) and strain rates (right).

Mechanical [127], the strain rate had to be in the following form

$$\dot{\epsilon}_{cr} = x_1 \sigma^{x_2} t^{x_3} e^{-x_4/T} \quad (6.4)$$

where $\dot{\epsilon}_{cr}$ is the strain rate, σ is the applied stress, t is the time dependence, T is the temperature and x_{1-4} are the constants that needed to be found by fitting to experimental data. In order to incorporate both types of stress, the strain rates had to be summed, which then meant a piecewise fit had to be used to find the fit constants. By splitting the strain rate into 3 sections, a good fit was obtained as shown in Fig. 6.6.

6.2.1 Ansys Results

The Ansys simulations were set up as shown in Fig. 6.7 with a symmetric quarter model of one cell. The element size was set to 1×10^{-4} and the physics preference set to non-linear mechanical. A frictionless support was used at the bottom with a pressure equal to 2 kg applied at the top, which would typically be used to keep the cavities still during the brazing process. The brazing curve in Fig. 6.2 was imported and applied as a thermal condition to the model. The temperature-dependent material properties were then imported with the values used shown in




Figure 6.6: Piece-wise fit to strain rate.

Fig. 6.9 and a nonlinear coupled thermal-structural simulation set up to solve the problem.




Figure 6.7: Single cell thermal deformation model (left), Parameters in model (centre) deformation after heating (right).

A typical curve showing the maximum (green), minimum (red) and average (blue) deformation in the vertical direction is shown in Fig. 6.8, we can see that the deformation follows the temperature curve nicely, but once the temperature is removed, some residual deformation is left which is the deformation due to diffusion. An example of where the maximum deformation occurs is shown in Fig. 6.7, where we can clearly see the nose cone sagging in the centre. Finally, various parameters were scanned to gauge the effect on deformation, including wall thickness, shell




Figure 6.8: Deformation vs. time Ansys result, maximum (green), minimum (red) and average (blue) deformation in the vertical direction.

thickness, and assumed grain size in copper; the results are presented in Table 6.1. Decreasing the wall thickness and the shell thickness both increased deformation, but the effect was relatively minor; the largest effect was the grain size, which, when reduced to $40 \mu m$ from $60 \mu m$, caused the deformation to double. However, it was judged that all the deformation values were most likely acceptable based on previous work measuring iris deformation during diffusion bonding [128]. Therefore, it was unlikely that deformation during manufacture would be a limiting factor for the wall thickness, as long as the shell is thick enough.

Grain size (μm)	Wall thickness (mm)	Shell (mm)	Max deformation (μm)
40	2	5	6.54
60	2	5	3.21
80	2	5	1.53
60	1	5	3.8
60	2	5	3.21
60	3	5	2.74
60	2	2	3.05
60	2	5	2.77
60	2	10	2.67

Table 6.1: Effect of parameters on deformation.

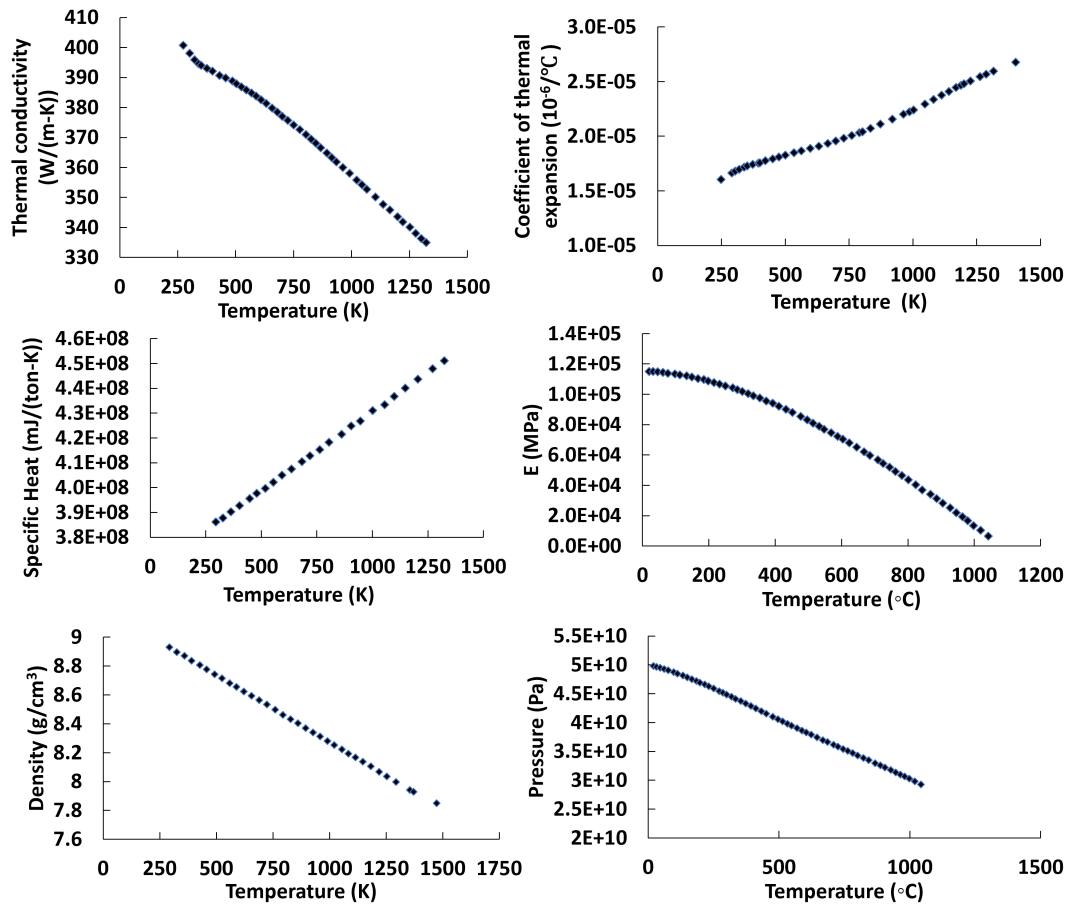


Figure 6.9: Temperature dependent material properties of copper obtained from [124, 129]. Top left to bottom right: thermal conductivity, coefficient of thermal expansion, specific heat capacity, Young’s modulus, density, pressure.

6.2.2 Heating During Operation

Another heating effect that was considered was RF heating during operation. When the cavity is in operation, RF power will dissipate in the walls, eventually causing thermal deformation and an expansion of the cavity. The key consideration is whether enough heat can be extracted from the nose cone region if the wall thickness is made excessively thin. Thermal simulations were set up using CST studio suite using both the electromagnetic and thermal solvers [130]. The cavity surface losses were calculated using the electromagnetic fields generated in the eigenmode solver,

and then imported as thermal losses in the thermal solver. Before the fields are imported, they were scaled using an assumed average input power of 1 kW (estimated using the duty cycle and peak power). The heat capacity of the copper used was 0.39 kJ/K/m and the electrical conductivity used was 5.8×10^7 S/m. Four cooling channels were inserted into the corners of the structure, each with a radius of 7 mm; these pipes were fixed at 35 °C to represent the distilled water while the rest of the copper had an initial temperature of 35 °C. The approximation of a fixed water temperature is used as there should be a small temperature rise in the water across a short section of the pipe. Holding the water temperature constant represents an infinite flow of water, and will only give an estimate of the temperature gradient due to the thermal conductivity of the copper, which was the value of interest when simulating the single cell. The automatic mesh refinement was used, with a 10 % refinement on each pass, and a maximum energy deviation of 0.001, a typical number of elements was 70000 tetrahedrons. The results for three different wall thicknesses along with the thermal model are shown in Fig. 6.10, where it can be seen that reducing the wall thickness significantly increases the temperature increase of the nose cone region in steady-state.

The exact results of the increase in temperature are plotted in Fig. 6.11. Originally a wall thickness of 2 mm was chosen for the design, but after consideration and discussion with the industrial partners, it was decided that the wall thickness should be increased as the design needed to be as robust as possible. Increasing the wall thickness from 2 mm to 2.5 mm reduced the temperature by ≈ 0.8 °C, while going from 2.5 mm to 3 mm only reduced the temperature by ≈ 0.2 °C. There is a trade-off as the shunt impedance is reduced as the wall thickness is increased, also shown in Fig. 6.11. It was eventually concluded that 2.5 mm would be more appropriate, as there was a diminishing benefit in increasing the wall thickness further to 3 mm, and the shunt impedance drops off significantly for a wall thickness of 3 mm. The diminishing thermal benefit is due to the fact that as the wall thickness reaches the thickness of the nose cone, most of the heat can be

dissipated away efficiently as the thermal bottleneck is removed.




Figure 6.10: Heating during operation (single cell). Left to right: 3 mm wall thickness, 2 mm wall thickness, 1 mm wall thickness, cooling pipe arrangement.




Figure 6.11: Nose cone temperature rise and shunt impedance vs. wall thickness for the single cell geometry.

6.3 Single Cell Frequency Analysis

A single cell thermal simulation was also set up to investigate single-cell frequency shifts. The $\beta = 1$ single cell was re-simulated in Ansys HFSS in order to have the surface losses to use in the thermal solver. Ansys was used as it has the capability of performing coupled thermal-mechanical-electromagnetic simulations. These losses were scaled to the input power for one cell ≈ 160 W and then imported into the

thermal solver. As before, the pipe sections were held at 35 °C. Once the thermal distribution was found, this was transferred to the mechanical solver to estimate the deformation in the cell. The ends of the cell were fixed to allow the cell to deform transversely. Once the deformation had been found, it was then re-imported back into the electromagnetic solver in order to look at the effects on the frequency. The steps of this process are shown in Fig. 6.12 with the electric field plotted on the deformed mesh shown in Fig. 6.13. The frequency shift obtained was 2.35 MHz which is relatively large, on the order of the cavity bandwidth. For security linacs auto-frequency tuning is typically used where the frequency of the magnetron is adjusted using a motor and tuning device, which allows for the compensations of such large frequency shifts as long as the linac field flatness remains acceptable [131]. This is in contrast to HEP machines where the water temperature is adjusted to bring the frequency of the linac to the nominal value, as it must remain constant for the beam synchronisation requirements. Although fixing the walls in the longitudinal plane is not a bad assumption, it causes excessive deformation in the transverse plane leading to a large shift; in reality this shift was expected to be smaller for each cell, but this needed to be confirmed with a full linac analysis.

6.4 Full Linac Thermal Analysis

During operation, when RF power is dissipated into the walls of the cavity, this heating causes a thermal deformation of the cavity, which in turn causes a frequency shift and can also affect the amplitudes in the cells, leading to asymmetric fields and degradation of the field flatness. Ideally, the thermally induced frequency shift should be within the bandwidth of the linac so that operation is not affected, although this shift can be compensated somewhat by adjusting the inlet water temperature. The deformation and field errors are dependent on the input power, the linac geometry, the placement of the pipes, and the properties of the cooling fluid. For this case, the heat load will be significantly smaller in the first few cells due to

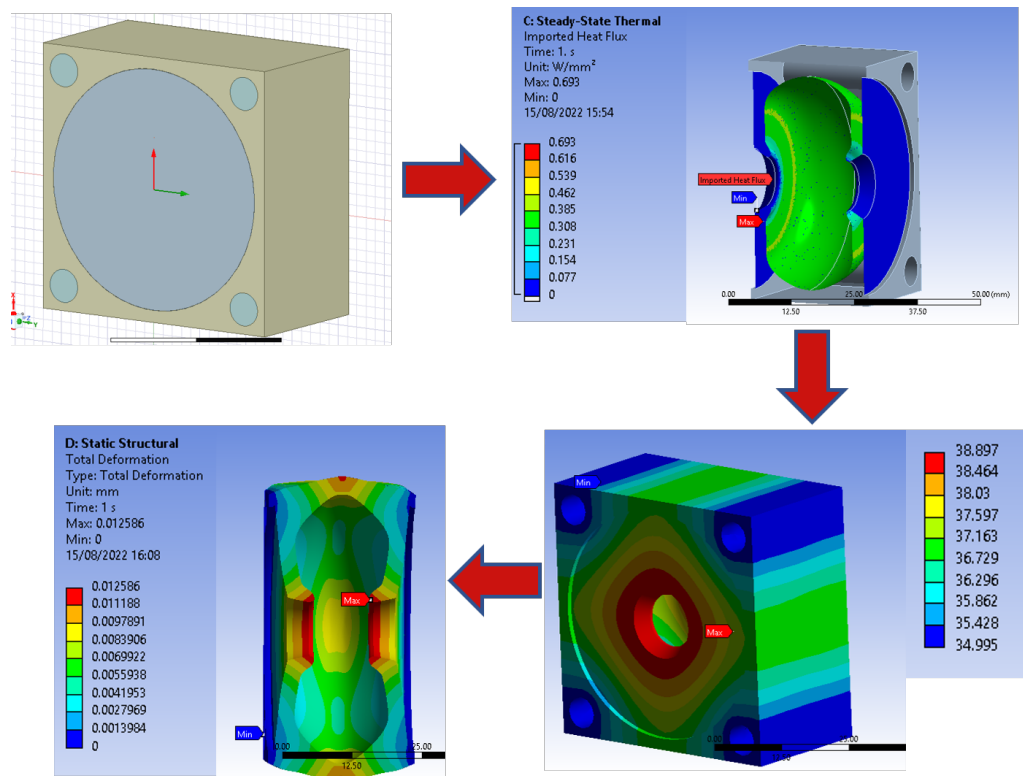


Figure 6.12: Steps required for single cell frequency shift analysis. (Top left) Single cell model in HFSS. (Top right) Imported heat flux from HFSS into the thermal solver. (Bottom left) Single cell thermal deformation. (Bottom right) Temperature distribution from thermal solver.




Figure 6.13: (Left) Displaced electric field after thermal deformation. (Right) Displacement due to deformation.

the reduced amplitude there, which could cause a larger effect on the field flatness compared with a constant gradient cavity. Therefore, a full thermal analysis was

required to quantify this effect. This section details the studies that were performed on a full linac under different operating conditions. It should be noted that this thermal analysis was performed on a previous design iteration, so there were slight differences in the first and second cells, although it is assumed that these differences will not affect the thermal performance significantly since the input power, field distribution, and cell shapes were very similar.

6.4.1 Forced Convection

The typical case for forced convection is the flow of water through tubes, to provide heating or cooling. This is also the case for normal conducting cavities, so the key theory for flow in tubes is presented here to give context for the next sections in the chapter.

6.4.1.1 Mean values

The analysis of fluid flow is normally approximated to be a one-dimensional problem, and the properties of the fluid are assumed to be uniform across the cross-section. The fluid velocity across a section of pipe with fluid flowing changes from zero on the surface of the pipe to a maximum in the centre; a mean velocity is defined from the conservation of mass

$$\dot{m} = \rho V_m A = \int_A \rho V(r, x) dA \quad (6.5)$$

where \dot{m} is the mass flow rate, ρ is the fluid density, A is the cross-sectional area and $V(r, x)$ is the velocity profile. As fluid flows through a pipe and is heated or cooled, a temperature gradient forms through the cross section from the surface to a maxima or minima in the centre of the fluid flow. A mean temperature is defined from the conservation of energy

$$\dot{E} = \dot{m} C_p T_m = \int_{\dot{m}} C_p T \delta \dot{m} = \int_A \rho C_p T V dA \quad (6.6)$$

where C_p is the specific heat of the fluid [113]. The fluid properties are evaluated at a bulk mean fluid temperature, an average of the inlet and outlet temperatures

$$T_b = \frac{T_{i,m} + T_{o,m}}{2} \quad (6.7)$$

where $T_{i,m}$ and $T_{o,m}$ are the mean inlet and outlet temperatures, respectively.

6.4.1.2 Boundary Layers

As fluid flows through pipes, the particles closest to the edge will eventually come to a complete stop due to frictional contact with the surface of the wall, which then causes friction for the next layer. For the mass flow rate to remain constant, the velocity of the fluid at the centre has to increase and a velocity boundary layer is formed between the wall and the fluid. The thickness of this layer increases until it finally reaches the centre of the pipe and the distance it takes to reach the centre is called the entrance length. It is typically around 10 diameters long, so for long straight pipes, this region is negligible compared to the length of the pipe [113]. Along with the velocity boundary layer there is also a thermal boundary layer where the temperature profile varies (along with the HTC) until it also becomes fully developed. These layers are shown in Fig. 6.14.

6.4.2 Heat and Cooling for Linacs

The following formulas apply to straight pipe geometries and describe the heat transfer between the pipe and a distributed heat source, which is the power lost on the cavity walls. The average total temperature rise in the water used for cooling the structure can be estimated from the conservation of energy for steady state flow using [81]:

$$\Delta T = \frac{\dot{Q}}{C_p \cdot \dot{m}}, \quad (6.8)$$

where \dot{Q} is the heat added to the system (i.e. the power dissipated in the cavity walls), C_p is the specific heat of the fluid, ΔT is $T_{o,m} - T_{i,m}$ (the average outlet




Figure 6.14: Boundary layer formation at the start of a pipe, also showing the variation of the HTC in this region. Image sourced from [113].

temperature minus the average inlet temperature) and \dot{m} is the mass flow rate of the fluid given by:

$$\dot{m} = \rho V_m A, \quad (6.9)$$

where V_m is the mean fluid velocity and A is the cross-sectional area of the pipe. There are two options for considering the thermal conditions at the surface

- Constant surface temperature, $T_s = \text{constant}$
- Constant surface heat flux, $\dot{q}_s = \text{constant}$

The surface heat flux is defined as

$$\dot{q}_s = h_{local}(T_s - T_m) \quad (6.10)$$

Where h_{local} is the local heat transfer coefficient at a certain location. For the linac case, we assumed that the surface temperature was constant for the initial analysis. For an average input power of 800 W and a mass flow rate of 0.08 kg/s, a water

temperature rise of 2.39 °C is obtained. The temperature difference between the fluid and the surface of the walls of the copper pipes, for an assumed constant surface temperature, is given by:

$$\Delta T = \frac{\dot{Q}}{hL\pi D}, \quad (6.11)$$

where L is the length of the pipe system, D is the hydraulic diameter. This temperature difference between the fluid and the surface decays exponentially along the length if the surface temperature is assumed to be constant, meaning that after a certain length there will be no improvement in the heat transfer for an increasing length of the tube. In practice, neither assumption is true for our case where the surface temperature and the surface heat flux will change along the length, though this analytical analysis is useful in the first instance to estimate the temperature rises and gradients for the system. In the formula above, h is the average heat transfer coefficient (HTC):

$$h = \frac{Nu \cdot \kappa}{D}, \quad (6.12)$$

where κ is the thermal conductivity of the fluid and Nu is the Nusselt number given by the Gnielinski equation for turbulent flow in tubes:

$$Nu = \frac{(\frac{f}{8})(Re - 1000)Pr}{1 + 12.7(\frac{f}{8})^{1/2}(Pr^{2/3})}. \quad (6.13)$$

Here, Re is the Reynold's number, which provides information about whether the flow is turbulent or not. It is given by:

$$Re = \frac{\rho v D}{\eta} = \frac{994.1 \cdot 2.85 \cdot 0.06}{7.26 \times 10^{-4}} = 23400 \quad (6.14)$$

where ρ is the density of water, v is the fluid velocity and η is the dynamic viscosity. For the case of 6 mm pipes, a 0.08 kg/s mass flow rate and water at 35 °C a Reynold's number of 23400 is obtained, which is well over the threshold for turbulent flow (\approx

4000). The Prandtl number describes the ratio of viscous to thermal diffusion and is given by:

$$Pr = \frac{\eta \cdot c_p}{\kappa}. \quad (6.15)$$

Finally, f is the Darcy friction factor, which, for a completely filled pipe with $Re > 4000$, is given by:

$$\frac{1}{\sqrt{f}} = -2 \log \left(\frac{\epsilon}{3.7D_h} + \frac{2.51}{Re\sqrt{f}} \right) \quad (6.16)$$

Where ϵ is the pipe's roughness height, and D_h is the hydraulic diameter of the pipe. As the formula is implicit, it is normally solved numerically, although recently accurate explicit approximations have been made [132], the one used here gives f as

$$f = \left(\frac{8.128943 + A_1}{8.128943A_0 - 0.85859209A_1 \ln\left(\frac{A_1}{3.7099535Re}\right)} \right)^2 \quad (6.17)$$

where

$$A_0 = -0.79638 \ln \left(\frac{\epsilon/D}{8.208} + \frac{7.3357}{Re} \right) \quad (6.18)$$

and

$$A_1 = Re(\epsilon/D) + 9.3120665A_0 \quad (6.19)$$

To calculate the required pumping power, first, the pressure drop along the pipe is calculated using

$$\Delta P = f \frac{L}{D} \frac{\rho V_m^2}{2} \quad (6.20)$$

then the required pumping power is found from

$$\dot{W}_{pump} = \dot{V} \Delta P \quad (6.21)$$

where W_{pump} is the required pumping power and \dot{V} is the volumetric flow rate. An estimate for the pumping power required for the linac was found to be less than 5 W for a pipe diameter of 6 mm and a length of 3 metres. The input parameters and initial calculated values are summarised in Table 6.2.

Parameter	Value	Units
Pipe diameter	6	mm
Water density	994.1	kg/m ³
Initial mass flow rate	0.08	kg/s
Fluid velocity	2.85	m/s
Average input power	800	W
Dynamic viscosity	7.26×10^{-4}	Ns/m ²
Prandtl number	4.96	-
Reynolds number	23399	-
Friction coefficient	0.02550	-
Nusselt number	149.55	-
Heat transfer coefficient	15478	W/(m ² K)
Outlet water temperature rise ($T_{o,m} - T_{i,m}$)	2.38	°C
Boundary temperature rise ($T_s - T_m$)	4.68	°C

Table 6.2: Initial analytical calculations of thermal properties.

6.4.3 HFSS Results

In order to compare the analytically calculated values with those obtained from Ansys, the first step was to re-simulate the RF structure in HFSS to make sure that the results were the same as those obtained from the original CST design. The reason for doing this is that HFSS is linked to the Ansys mechanical solver, allowing for the correct surface losses to be imported, and for the effects on the frequency to be quantified after thermal deformation. The vacuum model was set up in HFSS and ran with the driven modal simulation type, with the mesh set to a 5 μm surface deviation and a 10-degree normal deviation. A comparison to the results obtained in CST for S_{11} and the absolute value of the longitudinal electric field is shown in Fig. 6.15. The electric field in the transverse plane is also shown in Fig. 6.16. The S-parameters aligned well although a slight difference in amplitude was observed in the second and last cells between the models. This difference is small and could be due to slight mesh differences between the models, causing more reflections and, therefore, a difference in the field pattern.




Figure 6.15: (Left) Comparison of S_{11} from the CST and HFSS models. (Right) Comparison of the the absolute value of E_z on axis for the CST and HFSS models.




Figure 6.16: Plot of the absolute value of the HFSS electric field in the x-y plane.

6.5 Steady State Thermal Analysis

The full thermal analysis of the linac was performed using Ansys products. HFSS was used for the RF simulations, the mechanical solver for the thermal and structural simulations, and Ansys Fluent for the CFD analysis [127]. It was selected because it has the capability of performing full coupled RF-thermal-mechanical simulations that allows the investigation of the effects of thermal deformation on the frequency and fields, and it has previously been used successfully to analyse many RF cavity and accelerator system designs [133–139].

The general process for the simulation flow is shown in Fig. 6.17. Once the RF surface losses have been obtained, they are imported into the steady-state (SS) solver and mapped to the inner surface of the cavity. Typically, the heat transfer

coefficient is calculated using the formulas given in section 6.4.2, applied to the boundaries of the fluid and the solid as a constant value and then the thermal solver is used to find the temperature distribution map; this is then imported into the structural solver to find the stresses on the cavity surface and the displacement. The displaced mesh can then be re-imported back into HFSS to close the loop and find the effects on the frequency and fields. A more accurate approach would be to perform CFD simulations of the pipes, including the real water flow, in order to find the thermal distribution, but this process can be time and resource-consuming, so using analytical approximations would be preferred. Both methods were explored for the design, and a comparison was made between analytical calculations of the heat transfer coefficient and those obtained from a CFD analysis.

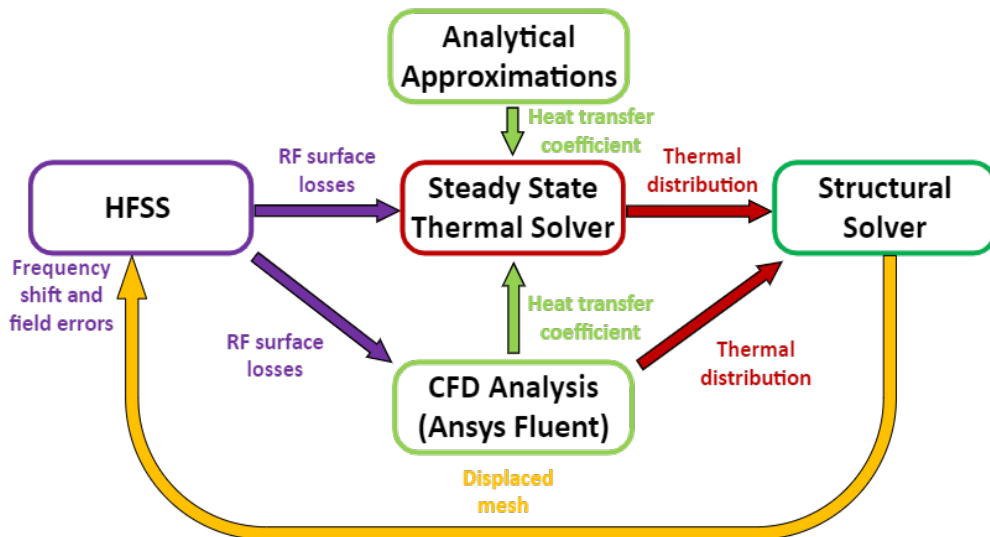


Figure 6.17: Simulation process flow for analysing the thermal deformation of RF cavities showing two options for obtaining the thermal distribution.

6.6 SS Simulation Setup

Once the HFSS simulation had been run, the surface losses were calculated on the internal faces of the linac. These losses are shown in Fig. 6.18 with the maximum losses being observed around the coupler and in the second cell as these areas were

where the magnetic fields were strongest.




Figure 6.18: RF surface losses calculated in HFSS.

The steady state thermal solver was then setup as shown in Fig. 6.19. To simulate the water, thermal elements (SURF152 [140, 141]) were used along the lengths of the pipes, with a mass flow rate and a water inlet temperature set at nominal values of 0.08 kg/s and 35 °C, respectively. The heat transfer coefficient for the pipes was set to 15478 W/m²K as calculated previously. The convective heat transfer coefficient to air was set to 10 W/m²K, which is reasonable for a well ventilated room as described in [142]. The RF losses were mapped to the surface labelled E in Fig. 6.19 and scaled to the average linac input power of 800 W. The mapped and scaled losses are shown in Fig. 6.20, the small specs shown are where there were inconsistencies in the mapped mesh elements meaning these elements received no flux, but they are very small and were considered unlikely to effect the results significantly. The resulting temperature distribution for 800 W of input power is shown in Fig. 6.21, with the temperature along the thermal elements in the water shown in Fig. 6.22. The largest temperature rises are observed on the second cell and fourth cell iris' with approximately a 3.7 °C gradient seen between the second cell iris and the outer wall. The temperature increase of the water was

2.37 °C, as expected from the calculated values in Table. 6.2.




Figure 6.19: Setup of steady state thermal simulation showing imported flux faces, convection to the water pipes and air, and the inlet of the pipes.




Figure 6.20: Mapped and scaled RF surface losses from HFSS into the SS solver, maximum values are seen on the slots connecting to the coupler.

6.7 Thermal Deformation

Once the temperature distribution had been obtained, a static structural simulation was setup to investigate the deformation resulting from the heating. The thermal distribution was mapped from the SS solver to the structural mesh. For the




Figure 6.21: Steady state thermal temperature distribution.




Figure 6.22: Temperature rise along thermal elements in SS solver.

structural solver, a fixed displacement boundary must be placed in at least one location in order to constrain the model adequately. Three cases were considered as shown in Fig. 6.23. For the first, the structure was fixed at the two ends, in this case the maximum displacement is seen on the vacuum port and coupler as the structure stretches in the transverse plane. For the second case, the boundaries were placed on the edges of the coupling port and vacuum port, the maximum deformation is now observed at the ends of the structure as the structure deforms longitudinally. In the final case the structure was fixed at all four points and the resulting deformation looks like a combination of the two previous cases.

The equivalent von Mises stress for the three cases is shown in Fig. 6.24. The yield stress for fully annealed OFE copper is around 62 MPa [73], and the maximum stress seen on any of the models is 56 MPa so well below the plastic deformation




Figure 6.23: Total deformation for three different boundary cases. (Left, 800x magnification), Fixed at ends, (Centre, 500x magnification), Fixed top and bottom, (Right, 800x magnification), Fixed at all four locations.

limit.




Figure 6.24: Von Mises stress for the three boundary cases. Total deformation for three different boundary cases. (Left) Fixed at ends (Centre) Fixed top and bottom (Right) Fixed at all four locations.

6.8 Frequency and Field Errors

The final step in the analysis process was to re-simulate the structure in HFSS using the deformed mesh to look at the frequency shift and field errors. The displaced mesh was taken from the structural solver and mapped to the HFSS mesh, and the resulting deformed mesh is shown in Fig. 6.25 for the case where the structure is fixed at the two ends. An example of the perturbed fields is shown in Fig. 6.26.

Once the electric fields from the deformed mesh had been obtained, they could be compared to the original to find the field error along the length of the linac. The field error was calculated using:




Figure 6.25: Deformed mesh in HFSS, used to re-simulate structure.




Figure 6.26: Absolute value of the electric field after simulating the deformed mesh in HFSS.

$$\% \text{ field error} = \frac{E_{deformed}(z) - E_{original}(z)}{E_{original}} \times 100 \quad (6.22)$$

This field error is shown in Fig. 6.27 for the three fixed displacement cases. The areas where the error seems to asymptotically increase is due to a slight shift in the zero crossing position of the fields. This is shown more clearly in Fig. 6.29 where the absolute error in the fields is shown for two of the cases, when the field is deformed slightly the position of the minima shifts leading to a divide by 0, causing a large percentage error. The maximum field error was always observed in the first cell, up to $\approx 1\%$ for the cases where the structure is fixed at the ends. The second cell and last cell saw approximately a 0.5% error for these cases and cells 3/4 were

relatively unaffected. Fig. 6.28 shows the effect of increasing the linac power from 800 W to 1.2 kW for the case where all four sides are fixed. The error is increased by approximately 60 % in the first cell, 50 % in the second cells with a smaller impact seen in the other cells.




Figure 6.27: Electric field error on axis after mechanical deformation.




Figure 6.28: Electric field error on axis after mechanical deformation for varying average power levels in the linac.

6.8.1 Frequency Errors

As well as comparing the errors in the fields, the frequency shift of the $\pi/2$ mode can also be found for all the cases, being listed in Table. 6.3. For all the cases,




Figure 6.29: Absolute error on the electric field on axis after mechanical deformation.

approximately a 1 MHz frequency shift is observed with a 1.5 MHz shift when the power is increased to 1.2 kW, leading to a shift of 1.23 MHz per kilowatt of power.

Model	Δf (MHz)
Fixed at ends (800 W)	0.97
Fixed top + bottom (800 W)	1
Fixed at 4 points (800 W)	1.01
Fixed at 4 points (1.2 kW)	1.5

Table 6.3: Comparison of frequency shifts for three different fixed displacement options.

6.9 CFD Analysis

The previous SS simulations were performed using a single value of the heat transfer coefficient, but, in reality, it varies along the length of the pipe and on different parts of the walls. The only accurate way to take this into account is to perform a CFD simulation of the system that includes the turbulent water flow in the pipes. These simulations can be time and resource-intensive, so if accurate results can be obtained using approximations, this would be preferred. This section details the CFD simulations that were carried out on the linac and compares them with approximations that are normally made.

6.9.1 Simulation Setup

The simulation setup is shown in Fig. 6.30. The same RF surface losses were imported, scaled to 800 W then applied to the red surface. At the inlet, the velocity of the water entering the pipes was set to 2.85 m/s and the temperature to 35 ° C. The turbulence model used for the simulation was the well-known SST $k-\omega$ model [143, 144] with the inlet turbulence defined by the turbulence intensity given by

$$I = 0.16(Re)^{-1/8} = 4.5\%. \quad (6.23)$$

[145], and the hydraulic diameter of 6 mm, recommended for fully developed internal flows.

6.9.2 Meshing

There are generally two options for simulating turbulent flow, which dictates how the mesh is set up.

1. Use standard wall functions that are designed for high Re flows, these can be used as the velocity distribution near the wall is similar for almost all turbulent flows. This approach uses approximations for the distance between the wall and the first mesh cell. They generally use a coarser mesh and require y^+ to be in the range of $30 < y^+ < 300$.
2. Resolve the viscous sublayer using a fine mesh around the interface between the fluid and the solid. Requires $y^+ \approx 1$.

The dimensionless wall distance y^+ is given by

$$y^+ = \frac{yu_\tau}{\nu} \quad (6.24)$$

where y is the absolute distance from the wall, u_τ is the friction velocity and ν is the kinematic viscosity. To calculate the placement of the first mesh cell, the skin

friction coefficient C_f is first estimated using the following

$$C_f \approx 0.079Re^{-1/4} \quad (6.25)$$

Then the wall shear stress is found using

$$\tau_w = \frac{1}{2}C_f\rho U_\infty^2 \quad (6.26)$$

where ρ is the density of the fluid and U_∞^2 is the free flowing fluid velocity. The friction velocity can then be found from

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}} \quad (6.27)$$

and finally the position of the first mesh cell as

$$y = \frac{y^+\nu}{u_\tau} \quad (6.28)$$

Therefore for a y^+ of 1 to resolve the sublayer a first mesh cell height of $4 \mu m$ is required. Although the SST $k-\omega$ model is y^+ insensitive, the mesh was generated to ensure that the value y^+ was adequate to resolve the boundary layer. A slice of the fluent mesh is shown in Fig. 6.31, with the very fine mesh around the boundaries of the water pipes shown in blue. A total of 9.6 million cells were included in the mesh. The material properties of the water were set as polynomials, fit over the temperature range of interest with the constants found using the data shown in Fig. 6.32.

6.9.3 Results

The simulation was run with the absolute convergence criteria of $1e-4$ for all quantities (continuity, velocities, k and ω) and $1e-8$ for the energy, it took 240 iterations to satisfy all convergence criteria. The temperature distribution observed




Figure 6.30: CFD simulation setup in Ansys fluent. (Left) Inner structure for surface loss mapping and cooling pipes showing inlet and outlet. (Right) Outer wall and cooling pipe walls.




Figure 6.31: Ansys fluent mesh. (Left) Slice through the x-y plane at the start of the structure showing the fine mesh around the boundary layer of the pipes. (Right) Mesh on outer wall.

in Fluent on the y-z plane is shown in Fig. 6.33. Comparing this to the one obtained in the SS solver in Fig. 6.21 the temperature distribution is very similar with the hottest areas occurring in the same places, as expected. There is a discrepancy in the temperature of ≈ 0.5 °C, the main reason for this being the difference between a constant HTC value applied over the pipes for the SS case instead of varying everywhere on the pipe walls and the assumption of a linear temperature increase in the water along the pipes. Fig. 6.34 shows a slice through through the x-y plane of the structure with two difference scales applied. The inlet is the pipe in the top




Figure 6.32: Water properties and fits used in Ansys fluent.




Figure 6.33: Slice of the y-z plane showing the resulting temperature distribution obtained using Ansys fluent.

right and the outlet on the top left, taking the average water temperature at the outlet a total temperature increase of 2.36 °C was found, very close to the predicted value of 2.38 °C. Fig. 6.35 shows a zoomed version of the outlet pipe, showing the boundary layer temperature rise compared with the bulk of the water. Taking the average temperature difference across this boundary yielded 4.32 °C, 8.3 % lower than the estimate calculated in Section. 6.4.2.

6.10 Validity of HTC Approximations

Again, the general relationship for convective heat transfer between a moving fluid and a surface can be defined by:

$$q_h = h_{local}(T_s - T_f) \quad (6.29)$$




Figure 6.34: Slice through the structure close to the inlet with two temperature scales. (Left) Slice through the structure showing the temperatures in the water pipes. (Right) Slice through the structure showing the temperature in the copper.




Figure 6.35: Slice through pipe outlet showing the boundary layer temperature rise in the water.

where q_h is the heat flux, h_{local} is the local convective heat transfer coefficient, T_s is the surface temperature and T_f is the fluid reference temperature. [146] This means that it depends on the type of surface, the type of fluid, the velocity of the fluid, and the direction of flow [147]. As mentioned previously, the formulas used to estimate the HTC in Section 6.4.2 are for straight pipe approximations, which lead to slight differences in the estimated boundary layer and temperature distributions. To see if an improved approximation could be found, investigations were carried out to estimate the HTC in curved-pipe geometries.

6.10.1 Heat Transfer in Curved Geometries

The Nusselt number for curved pipe geometries for liquids which have a high Prandtl number has been found from a series of papers as follows [148]:

$$Nu_c Pr^{-0.4} = \frac{1}{41} Re^{5/6} \left(\frac{a}{R} \right)^{1/12} \left\{ 1 + \frac{0.061}{[Re(a/R)^{2.5}]^{1/6}} \right\} \quad (6.30)$$

where Nu_c is the Nusselt number in a curved pipe geometry, Pr is the Prandtl number, Re is the Reynolds number, a is the radius of the pipe and R is the radius of curvature of the pipe axis [148]. Fig. 6.36 shows a comparison of the Nusselt number calculated using this formula compared to that of straight pipes. In the turbulent regime, the Nusselt number is $\approx 30\%$ larger than the straight-pipe approximation. The Nusselt number was calculated for the pipe geometry used in the linac, with 6 mm pipes and a 10 mm radius of curvature, which turned out to be 199, 25% larger than the straight pipe value. Most of the the pipe geometry is straight, however, to get a value for a final Nusselt number, the two values were combined in proportion to their fractional length for the pipes (17% for the curves and 83% for the straight sections). This gave a final Nusselt number of 159, which led to an HTC value of 16348 W/m²K, 6% larger than the straight-pipe estimate. The surface Nusselt number measured on one of the corners of the pipes is shown in Fig. 6.37. Similar to the estimate, the value of the straight sections of the pipes is between 130 - 160 and closer to 200 at the corners and on the short sections connecting two corners, due to the increased turbulence there.

The heat transfer coefficient is a derived quantity, which means that it depends on the choice of reference temperature for the calculation [150]. To remain consistent with the analytical calculations, the bulk mean fluid temperature was used for the fluent HTC calculation, given by

$$T_b = \frac{T_{m,i} + T_{m,o}}{2} \quad (6.31)$$

where $T_{m,i}$ and $T_{m,o}$ are the mean inlet and mean outlet water temperatures




Figure 6.36: Comparison of Nu_c for curved and straight pipes, image from [149].




Figure 6.37: Surface Nusselt number, showing variation in the corners of the pipes compared with the straight sections.

respectively. Using this definition, the heat transfer coefficient along the length of the pipes is shown in Fig. 6.38, where the lower values along the straight sections and higher values near the corners are clearly seen. Fig. 6.39 shows the value of the HTC for all of the nodes along the surface of the pipes with a 26 period moving average included. Averaging all the points gave a value of $16423 \text{ W/m}^2\text{K}$. Using this value in Equation 6.11 lead to an estimated boundary layer temperature rise of $4.43 \text{ }^\circ\text{C}$, much closer to the measured fluent value of $4.32 \text{ }^\circ\text{C}$.

The reason for these variations along the pipes can found by examining Fig. 6.40 where there are significant dips and increases in the velocity and temperature around the bends of the pipes, before normalising again around half way along the




Figure 6.38: HTC values along the pipes. (Left) Outer pipes visible (Right) Slice through all the pipes showing the values on the inner surface.




Figure 6.39: HTC values at all of the nodes along the pipes with a moving average of 26 periods shown in red.

straight section. One would therefore expect differences between the SS solver using a linear increase in temperature for the water and a constant HTC value over the surface.




Figure 6.40: (Left) Water velocity along the pipes showing variation around the corners. (Right) Water temperature showing the difference between the SS linear temperature increase approximation and the real temperature distribution in the water.

A further investigation was started to compare the different models and to look

at the approximation validity. Ideally, a comparison should be drawn between something that would be measurable in the real system, so the temperature was measured and averaged over a small area on the outer shell of the linac as shown in Fig. 6.41, where a temperature probe could be attached to the side of the linac. This was then used as a metric for comparisons between models.




Figure 6.41: Temperature on external wall of the linac with the area used for the average temperature measurement shown in red.

A number of simulations were performed with various values for the HTC, including importing the calculated HTC from fluent into the SS solver and applying it to the boundary between the water and the copper wall. For each simulation, the temperature was measured on the side of the linac. The results are presented in Fig. 6.42. The dark blue points show how the temperature on the surface changes with a single value of the HTC applied to the entire wall. The yellow point is the temperature obtained using the analytical value of HTC without bends applied to the wall in the SS solver (45.548 °C). Comparing this with the value measured directly in fluent (45.415 °C) gives an error of 0.3 %. Including the bends in the formula and applying the new analytical HTC value reduces the error by more than half to 0.14 %. The results of the SS solver with the HTC imported from fluent gives very similar results to the value measured directly in fluent (0.015 % difference), confirming that the difference is a direct results of the HTC variation along the pipes and not due to the simulation setup. A comparison was also made between the maximum temperature reached on the second iris which is where the maximum




Figure 6.42: Temperature on outer wall vs value of the HTC for various different simulation configurations.

temperatures were found. The comparison is shown in Fig. 6.43. By holding the pipe walls at 35 °C the maximum iris temperature is severely underestimated, with nearly a 10 °C difference seen between it and the other cases. Including convective heat transfer and using the initial estimated HTC of 15478 W/m²K reduces the difference, but still, the maximum temperature is underestimated by ≈ 1 °C due to the colder water temperature and lower value of the HTC. Increasing the water temperature to an average of the estimated inlet and outlet temperatures, i.e 36.2 °C causes both the cases with the original HTC and improved HTC values to overestimate the maximum temperature, although they are much closer so this approximation would most likely be adequate for most cases where there are some safety margins in the allowable maximum temperature. Using the water elements and the improved HTC gives a maximum temperature very close to the one obtained from fluent, and so should be used for designs where the maximum temperature is important, but there are not the resources or time available to do a full CFD analysis. For cases where one is operating close to the maximum temperature limits, such as high repetition rate systems, then a full CFD analysis would be recommended to provide the most accurate results.

The results presented here confirmed that it is appropriate to estimate the HTC




Figure 6.43: Comparison of the maximum temperature on the iris for different simulation setups.

and perform SS simulations in Ansys fluent without having to perform detailed CFD simulations. The error induced by the straight-pipe approximation was only equivalent to a 0.3 % error on a measurable temperature on the linac surface. This could then be reduced by half by incorporating the variation of the Nusselt number in the bends, as described.

6.11 Operational Considerations

Generally, there are two options to deal with the thermally induced frequency shift when the linac is operating, each useful for different applications:

1. Design the linac to be slightly smaller than required so that when it is running at the operating temperature, the thermal expansion is accounted for, and the linac ends up at the operating frequency. This is typically used for cases where linacs need to be synchronised with an existing beam, though the average power of the linac will now be fixed.
2. Assuming that the frequency shift is still within the bandwidth of the RF source, the frequency of the RF input can be adjusted to match the new resonant frequency of the linac, this is useful for systems where the operating

frequency can afford to be adjusted such as in short industrial linacs as there are no external synchronisation requirements.

3. Once the linac is operating, the temperature of the cooling water is adjusted at the inlet to bring the frequency back to the required value.

6.11.1 Frequency Adjustment

As there was no need to externally synchronise the linac with a beam, it was decided that the system could be tuned using the cooling water. The SS solver and the structural solver were used initially to find the frequency shift with an inlet water temperature of 35 °C; This was found to be around 1.5 MHz, for a water flow rate of 3 L/min. Fig. 6.44 shows the operating frequency of the linac vs. the inlet temperature of the water. To bring the frequency back to 5.712 GHz, the inlet water would need to be adjusted from 35 °C to 18 °C. Ideally the temperature of the inlet




Figure 6.44: Frequency vs. inlet water temperature showing the change required to bring the linac back on frequency.

water should remain either above or below room temperature so that the system is not constantly swapping between heating and cooling. With this in mind, a power scan was performed with the water flow rate set to 5 L/min, a reasonably high value that avoids severe cavitation [151]. For each power level, the water temperature was adjusted from 35 °C down to whatever temperature was required to bring the

resonance frequency back to 5.712 GHz, while not going below 20 °C. The results are shown in Fig. 6.45. It shows that the system can run with powers up to 1.2 kW before it is not possible to remove the frequency shift by adjusting the water inlet temperature.




Figure 6.45: Frequency shift vs. average power after reducing the inlet temperature to the value required to bring the linac back on resonance, down to 20 °C.

6.11.2 Power Limits

Next, the ultimate power limits for the linac were assessed. It was decided that the limit would be set by a maximum external temperature of the linac as a safety consideration. For various values of the water flow rate, the power was increased until the external temperature of the linac averaged over the hottest spot shown previously reached 60 °C. The results are shown in Fig. 6.46. With a flow rate of 1 L/min, the maximum power limit was 800 W, this increases to a maximum of 2.65 kW with a flow rate of 5 L/min. Finally, Fig. 6.47 shows curves of outer shell temperature vs. inlet temperatures for various power levels and a flow rate of 5 L/min. The value of 5 L/min is a conservative estimate of a safe fluid flow rate that will avoid cavitation in susceptible areas in the water network such as the pump. It is possible to run at higher flow rates (> 10 L/min), but this system needed to be as robust as possible so a lower value was chosen.




Figure 6.46: Maximum average input power to the structure before the outer linac wall reaches 60 °C vs. the water flow rate.




Figure 6.47: Temperature on outer shell vs. water inlet temperature for various average input powers and a flow rate of 5 L/min.

If the system is operating in a warmer or much colder environment where the chiller can be operated at 20 degrees to avoid swapping between cooling and heating then the linac would be able to operate up to powers of 2.65 kW with a 60 °C temperature limit on the cavity surface for a flow rate of 5 L/min. In many cases cargo scanning systems are operated outside with large temperature changes meaning that the optimal temperature and power handling will be location dependent. The maximum power limit for the linac of 1.2 kW is given for an inlet temperature of 35 °C.

Chapter 7

RF System Design

7.1 Multi-View Linac System

Once the linac design had been completed, the design of a system that could generate multi-view X-ray images was started. This design would include the orientations of the linacs, the specifications required for the RF source, and the RF network needed to connect the linacs.

7.1.1 Initial RF System Considerations

The initial design goal of the project was to design a linac that could be potentially used for a multi-view cargo system, where multiple linacs would be fired to generate X-ray images from different angles, potentially allowing quasi-3D images to be generated to provide more information about the cargo being scanned compared with traditional 2D images. The main obstacle to developing such a system was the question of how to design an RF system that would allow the timings between when the linacs were fired to be changed so that they could be synchronised to generate good-quality X-ray images. One easy solution to this problem would be to have three separate RF sources, each connected to a separate linac, allowing full control of the linac firing time. This presents a problem, though, as the RF source would be either the most expensive (in the cast of klystrons) or the component with the

shortest lifetime (in the case of magnetrons) of industrial cargo scanning systems, so having three would significantly increase the cost or reduce the reliability of the whole system.

In theory, multiple linacs could be driven from one RF source, assuming that it can run at the required duty cycle and that high-power RF splitters are available. The linacs could be connected through differing lengths of RF waveguide to change the delay between them, but there would only be tens of nanoseconds delay between the linacs for a few metres of standard rectangular waveguide, causing the system to be unable to distinguish between the linacs when they are firing. Ideally, the delay between the firings should be on the order of a few μs , which would require up to 500 m of waveguide, making it an infeasible solution. The initial idea to overcome this problem was to use some other type of delay line between the source and each linac, as shown in Fig. 7.1.

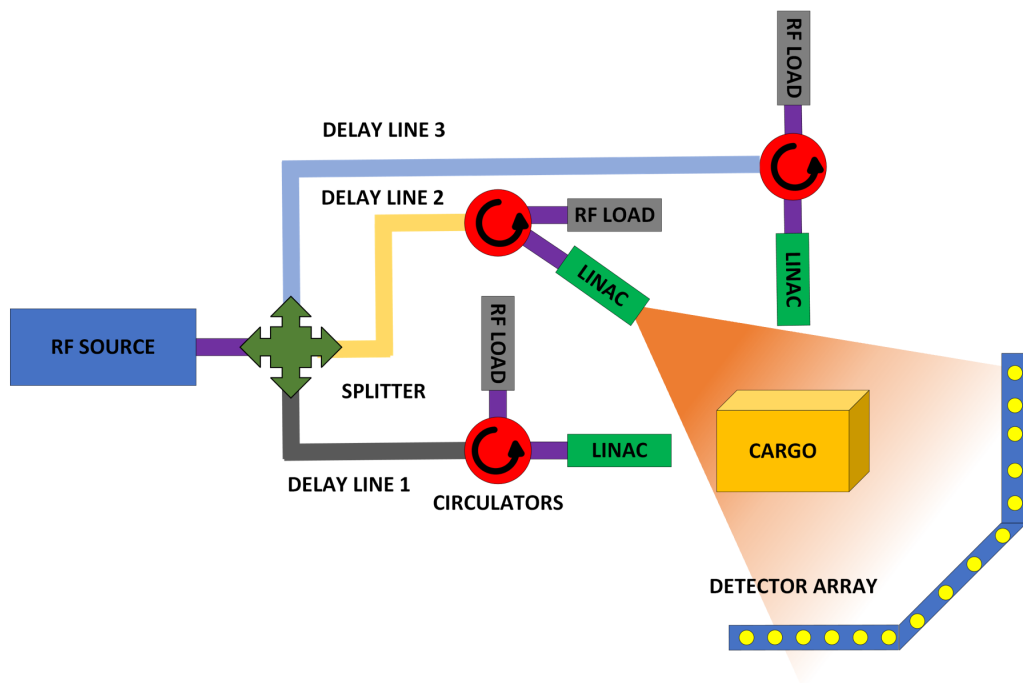


Figure 7.1: Initial design concept using delay lines and circulators to allow for the use of only one RF source.

7.1.2 Delay Lines

Various methods for delaying the RF pulses were considered with the potential options shown in Table. 7.1, along with a brief description of how they would be implemented to create a delay line. The solution needed to be cheap, which ruled out the use of pulse compressors or HOM resonant delay lines, which would also still require long waveguides [152–154]. The system needed to be able to cope with powers in the megawatt range, making the use of dielectric-loaded waveguides unlikely. A high-power RF switch would be an ideal solution as the power could be diverted to any of the lines for synchronisation, but a switch with the required power handling and switching speed could not be found. The fastest switch with the highest power found was a 1 MW peak switch with a switching time of 150 μs [155], or a 300 kW peak switch with a switching time of $< 50 \mu\text{s}$ [156], both falling short of the requirements.

Option	Delay method	Potential issues
Pulse compressors	Use three pulse compressors with different fill times, releasing the pulses with the required delay.	Expensive option, requires long design times.
HOM resonant delay lines	Cut delay line length by N times, where N is the number of HOM's.	Would still require 10's of metres of waveguide, complex design process.
Waveguide near cut-off	Group velocity of RF is reduced the closer to the cut-off frequency it is.	Potentially very lossy.
Dielectric loaded waveguides	Use dielectric to modify the group velocity.	Power handling concerns.
Travelling wave cavity	Use cavities to modify group velocity.	Large power dissipation in cavity walls.
High power RF switches	Use one input line then switch to other lines to create delay.	Mostly exist in the ms switching range, μs range is required for the linacs.

Table 7.1: Initial delay line considerations.

The most interesting options from a cost and ease of manufacture perspective

were to use travelling wave cavities or design a waveguide and operate it close to cut-off to reduce the group velocity. The main concern with these methods was whether they would be too lossy to allow enough RF power to reach the linacs from the RF source. Fig. 7.2 shows the power loss in circular waveguide vs. normalised




Figure 7.2: Power loss vs. normalised frequency for various higher-order waveguide modes, operating 0.3 MHz from cut-off.

frequency for 8 different TE/TM modes. For each mode, the radius corresponding to a cut-off frequency of 5.7117 GHz was found, and then the conductor loss for the appropriate mode found. The cut-off frequency of TE modes can be found using

$$(f_c)_{mn}^{TEz} = \frac{\chi'_{mn}}{2\pi a \sqrt{\mu\epsilon}} \quad (7.1)$$

Where a is the radius of the waveguide and χ'_{mn} are the zeros of the derivative $J'_m(\chi'_{mn}) = 0$ ($n=1,2,3,\dots$) of the Bessel function $J_m(x)$. The cutoff frequency for TM modes is given by

$$(f_c)_{mn}^{TMz} = \frac{\chi_{mn}}{2\pi a \sqrt{\mu\epsilon}} \quad (7.2)$$

Where χ_{mn} are the zeros of $J_m(\chi_{mn}) = 0$ ($n=1,2,3,\dots$) of the Bessel function $J_m(x)$. [157] Using these, the cutoff frequency can be set to the required value, and the

corresponding value of the waveguide radius found for each mode. The conductor attenuation coefficient for the TE modes is given by

$$(\alpha_c)_{mn}^{TEz} = \frac{R_s}{a\eta\sqrt{1 - \left(\frac{f_c}{f}\right)^2}} \left[\left(\frac{f_c}{f}\right)^2 + \frac{m^2}{(\chi')^2 - m^2} \right] Np/m \quad (7.3)$$

and for the TM modes it is given by

$$(\alpha_c)_{mn}^{TMz} = \frac{R_s}{a\eta} \frac{1}{\sqrt{1 - \left(\frac{f_c}{f}\right)^2}} \quad (7.4)$$

Where R_s is the surface resistance, f_c is the cut-off frequency, a is the waveguide radius and η is the impedance of free space. The group velocity in the waveguide is given by

$$v_g = c\sqrt{1 - \left(\frac{f_c}{f}\right)^2} \quad (7.5)$$

Then the group delay can be found using the following

$$t_{delay} = \frac{l}{v_g} \quad (7.6)$$

where l is the length of the guide.

As expected, the loss for the TE₀₁/TE₀₂ modes decreases as the frequency is increased, but close to the cut-off value, higher-order TM_{X1} modes were found to be the lowest loss. By operating the waveguide close to its cut-off frequency, a delay can be introduced as shown in Fig. 7.3. Operating with a v_g of 0.01 %c (0.3 MHz from cut-off) would require a length of ≈ 6 m, a much more realistic value. However, to achieve this delay, a complicated mode converter would be required to switch to the TM_{10,1} mode, and almost 50 % of the input power would be lost in the delay line. Cooling would be required for the waveguide, and the cost of the RF source would increase as extra power would be needed to compensate for these losses.




Figure 7.3: (Left) Percentage of power lost vs. delay in waveguide for 8 circular waveguide modes. (Right) Loss vs. v_g/c for 8 circular waveguide modes.

7.2 Frequency-Division Multiplexing

Frequency division multiplexing (FDM) has a long history of use within the communications industry as it allows for combining multiple signals with different frequencies into a single signal and transmitting them simultaneously [158, 159]. A good example of this is sending multiple TV channels to people's homes without interference between the channels [159]. A simple schematic showing the basic concept of FDM is shown in Fig. 7.4. The initial signals are mixed with a carrier frequency using RF mixers [160] and combined into one signal. This signal may then be modulated if necessary to shift the entire set of frequencies. The signal then passes through one transmission channel, is filtered, and then is demultiplexed at the other end in order to recover the original signals.

7.3 Frequency Switching

As the design needed to minimise the cost to make industrial use of the system viable, the next iterations of ideas explored tried to minimise the number of additional components and the cost of implementation. The next solution considered is shown in Fig. 7.5. This method required an RF source that would be capable of frequency

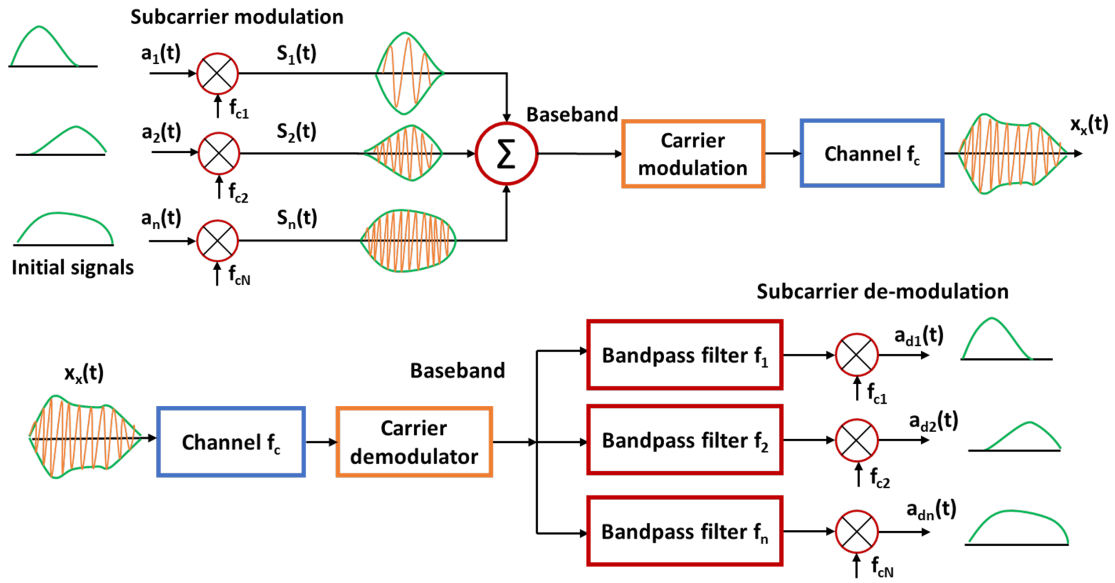


Figure 7.4: FDM schematic showing modulation, transmission, then demultiplexing.

adjustment within a ≈ 10 MHz range, such as multi-beam klystrons [161–163] or cross-field amplifiers [164]. It draws from the idea of FDM, in that all three signals will use the same transmission line, with the cavities acting as the multiplexers at the other end. The key difference is that the signals would be separated in time with a variable delay to allow for synchronisation of the linacs.

Each linac would be tuned to a slightly different frequency with circulators connecting each linac to the source. The frequency of the RF source would be switched twice within the pulse to create 3 sections with different frequencies within the main pulse. The pulse would travel through the first circulator and reach the first linac, where the first part of the pulse corresponding to the frequency of the first linac would be absorbed to power the linac. The remaining power would be reflected and go back to the circulator, then make its way to the second circulator to repeat this process. The second part of the pulse at the second frequency would be absorbed by linac 2, with the final part going to linac 3. The remaining signal that is not absorbed by any of the linacs would end up at the final RF load. As the frequency of the linacs is adjusted, it could be possible that one of the frequencies of the other modes in one of the linacs would correspond with the operating frequency

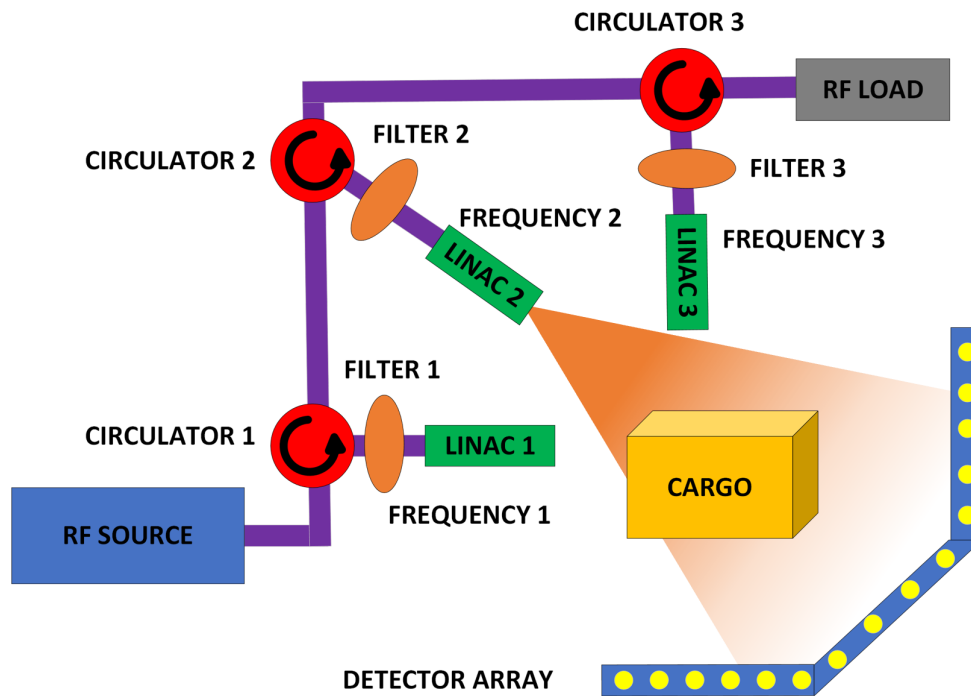


Figure 7.5: Frequency switching RF system idea, using three linacs tuned to slightly different frequencies and an RF source with a large enough bandwidth to accommodate them.

of the pulse, meaning it could be absorbed in the wrong linac and power the wrong mode. One way to avoid this would be to tune the frequencies in such a way so that this does not occur or to include filter cavities before each linac that only allow a restricted band of frequencies to pass. This would become more important for linacs with a large number of cells as the resonances are closer together. It was decided that the tuning method would be used initially as there were only five cells in the linac design, meaning that it should be possible to avoid overlapping modes.

7.4 RF Pulse

It was envisioned that the RF pulse would be split as shown in Fig. 7.6, with a rise time set by the RF source time, and then 3 sections of the pulse each $\approx 3 \mu\text{s}$

in length. The frequency would be switched using an LLRF feedback system to adjust the input frequency to the klystron, which should be capable of switching the frequency in under 100 ns [165, 166].




Figure 7.6: RF pulse split into three sections with the frequency being changed every 3 μs by switching the klystron frequency using a LLRF system.

7.5 Simulation Verification

Some preliminary simulations were set up in order to verify the RF system design and observe the reflected and transmitted signals from the linacs. The system was modelled using the circuit diagram shown in Fig. 7.7. The linacs were modelled as RLC circuits, and the input coupler and probes were modelled as transformers, with coupling beta's β_e and β_t .




Figure 7.7: Circuit diagram of the modelled RF system showing the linac modelled as an RLC section and transformers representing the input coupler and probe.

For the purposes of the simulations, only the S-parameters of the system were of interest. The S_{11} and S_{21} for the two port circuit in terms of the Q-factors and coupling coefficients are given by

$$S_{11} = \frac{\beta_e - \beta_t - 1 - j2Q_0\delta}{1 + \beta_e + \beta_t + j2Q_0\delta} \quad (7.7)$$

$$S_{21} = \frac{2\sqrt{\beta_e\beta_t}}{1 + \beta_e + \beta_t + j2Q_0\delta} \quad (7.8)$$

Where the coupling coefficients are given by

$$\beta_e = \frac{Q_0}{Q_e} \quad (7.9)$$

and

$$\beta_t = \frac{Q_0}{Q_t} \quad (7.10)$$

and

$$\delta = \frac{f - f_0}{f_0} \quad (7.11)$$

.

The relationship between Q_0 and Q_L is given by:

$$Q_0 = Q_L(1 + \beta_e + \beta_t) \quad (7.12)$$

7.5.1 Signal Generation

The pulse was shape was created for the signal envelope, as shown in Fig. 7.6 and the three carriers applied to the correct sections, with the resulting frequency spectrum of the input signal shown in Fig. 7.8.

The signal was generated by creating a square wave envelope and multiplying it with a time signal with the required frequency content. The process of generating the frequency spectrum of a simple square wave is derived below.

A centred square wave of amplitude A and duration $2T$ is defined as




Figure 7.8: Frequency domain spectrum of the input signal showing frequency content at 3 frequencies, separated by 3 MHz.

$$x(t) = \begin{cases} A & \text{for } -T \leq t \leq T \\ 0 & \text{otherwise} \end{cases} \quad (7.13)$$

Here, the square wave extends symmetrically from $-T$ to T . The Fourier Transform $X(f)$ of the time-domain signal $x(t)$ is given by

$$X(f) = \int_{-\infty}^{\infty} x(t)e^{-j2\pi ft} dt \quad (7.14)$$

Since the signal $x(t) = A$ only for $-T \leq t \leq T$, the limits of the integral are restricted

$$X(f) = \int_{-T}^T Ae^{-j2\pi ft} dt \quad (7.15)$$

Solving the integral

$$X(f) = A \int_{-T}^T e^{-j2\pi ft} dt \quad (7.16)$$

This is a standard integral for an exponential function. The solution is

$$X(f) = A \left[\frac{e^{-j2\pi ft}}{-j2\pi f} \right]_{-T}^T \quad (7.17)$$

Evaluating this at the limits $t = T$ and $t = -T$

$$X(f) = A \left(\frac{e^{-j2\pi fT} - e^{j2\pi fT}}{-j2\pi f} \right) \quad (7.18)$$

Simplifying using Euler's identity

$$e^{j\theta} = \cos(\theta) + j \sin(\theta) \quad \text{and} \quad e^{-j\theta} = \cos(\theta) - j \sin(\theta) \quad (7.19)$$

Then, $e^{-j2\pi fT} - e^{j2\pi fT}$ becomes

$$e^{-j2\pi fT} - e^{j2\pi fT} = -2j \sin(2\pi fT) \quad (7.20)$$

Substituting this into the expression for $X(f)$

$$X(f) = A \left(\frac{-2j \sin(2\pi fT)}{-j2\pi f} \right) \quad (7.21)$$

Simplifying further

$$X(f) = A \cdot \frac{\sin(2\pi fT)}{\pi f} \quad (7.22)$$

Which can be recognised as the sinc function

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \quad (7.23)$$

The FFT was taken of a simple square wave, with the result shown in in Fig. 7.9. This confirmed that the signal spectrum shown in Fig. 7.8 was reasonable, with the slight asymmetries being due to the rising and falling edges of the pulse.

The Wigner-Ville and smoothed pseudo-Wigner-Ville distributions of the input signal are shown in Fig. 7.10. The Wigner-Ville distribution is a time-frequency analysis tool that provides high-resolution representations of a signal's energy distribution across both time and frequency, allowing for the visualisation of complex signals. It was used to ensure that the signals had the correct frequency content




Figure 7.9: (Top) Simple 10 μ s square wave. (Bottom) FFT of the square wave showing sinc function.

and that the switched at the correct times.

The Wigner-Ville Distribution (WVD) of a signal $x(t)$ is given by the following equation:

$$W_x(t, f) = \int_{-\infty}^{\infty} x\left(t + \frac{\tau}{2}\right) x^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau \quad (7.24)$$

Where, $W_x(t, f)$ is the WVD as a function of time t and frequency f , $x(t)$ is the signal, $x^*(t)$ is the complex conjugate of the signal, τ is a time-lag variable, and $e^{-j2\pi f\tau}$ is the Fourier kernel, providing the frequency dependence. The pulse length was reduced to 1 μ s to reduce memory requirements, with switching occurring at 360 ns and 660 ns. The three frequencies are clearly seen at the correct location along the pulse.

The next question to answer was whether the three pulses would fit within the bandwidth of the RF source. A search was performed for a source with the correct power and bandwidth requirements, and one example found was the Nelson NC4 [167]. It has a peak RF power output of 3.5-4 MW and a bandwidth of over 10 MHz. It also allows for a 16 μ s pulse width for a pulse repetition rate of 180 pps, which would allow the three different frequency sections to fit comfortably within




Figure 7.10: Wigner-Ville distributions of the input signal, over a shorter length of $1 \mu\text{s}$, showing the three different frequencies in time. (Top) Wigner-Ville distribution. (Bottom) Smoothed pseudo Wigner-Ville distribution.

the pulse. Fig. 7.11 shows the accelerating mode S_{11} resonance curve, which has been shifted by ± 3 MHz, showing that all three curves can fit within a 10 MHz bandwidth with minimal overlap. The S-parameters generated by Equations 7.7 and 7.8 were concatenated using ABCD matrices with the ideal circulator S-parameters shown below:

$$[S] = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \quad (7.25)$$

The S-parameters were then multiplied by the Fourier transform of the input signal




Figure 7.11: S_{11} curves for three linacs with 3 MHz separation between them, all within a 10 MHz bandwidth.

(Fig. 7.6) and then inverse Fourier transformed to find the final time domain pulse shape. This process used the convolution theorem which states that convolution in the time domain corresponds to multiplication in the frequency domain, and vice versa. Given two time-domain signals $x(t)$ and $h(t)$, the convolution of these signals in the time domain is denoted as $y(t) = (x * h)(t)$, where $*$ represents the convolution operation. The convolution theorem states that

$$(x * h)(t) \leftrightarrow X(f) \cdot H(f) \quad (7.26)$$

Where, $X(f) = \text{FFT}(x(t))$ is the Fourier Transform of $x(t)$, $H(f) = \text{FFT}(h(t))$ is the Fourier Transform of $h(t)$ and $Y(f) = X(f) \cdot H(f)$ is the product of the two Fourier Transforms in the frequency domain. To get the convolution result back into the time domain, the inverse FFT (IFFT) is used on the product of the two Fourier Transforms:

$$y(t) = \text{IFFT}(X(f) \cdot H(f)) \quad (7.27)$$

This process is much more efficient than convolving the time domain signals.

The results of the initial simulations are shown in Fig. 7.12. The plots on the left show the envelopes of the reflected signals from each cavity, and the plots on the

right show the envelopes of the transmitted signal measured by the probe, assumed to have a quality factor of approximately 50000. It is clear from the reflected signal from linac 1 that the first part of the pulse from 0.5 to 3 μs has been accepted into the linac. Oscillations were observed in the reflected signals at the transition point from the first to second and second to third frequencies, where power leaking from the first and second cavities mixes with the reflected signals, causing frequency beating. These ripples do not appear in the cavity voltage signal as they are filtered out, so they should not affect the beam.

The transmitted signals are relatively unaffected and the pulses are stable once they reach their final amplitudes after $\approx 1 \mu\text{s}$. For the signal transmitted into linac 1, the second and third frequencies had relative amplitudes of 18 % and 9 %, respectively. Reducing Q_0 reduces the ripple, as shown in Fig. 7.13, which is expected as the cavity response is reduced at the frequency at which it is being excited. Even for higher Q_0 values of 16000, the transmitted pulses into linacs 2 and 3 do not ring for longer than 1.5 μs . A scan of Δf was also carried out, showing the frequency difference between the resonances, with the results shown in Fig. 7.14. The pulses are stable until the difference in frequencies is reduced to 1 or 2 MHz, and then power from the second and third portions of the pulse starts to leak into the first linac significantly. Going to a 5 MHz difference does improve the leakage, but not significantly and ideally they would be as far away as possible, provided that all three pulses can fit within the bandwidth of the chosen RF source.

7.5.2 Validation

In order to validate the design of the RF system, some more advanced simulations were performed. For this, a model was made with a small probe inserted in the end of the linac beampipe, as shown in Fig. 7.15. The probe length was then tuned to the correct coupling factor, with $Q_t \approx 50000$. The system was then simulated in the frequency domain, with the S-parameters generated shown in Fig. 7.16. These S-parameters were then used in place of the formulas given in Section 7.5. For

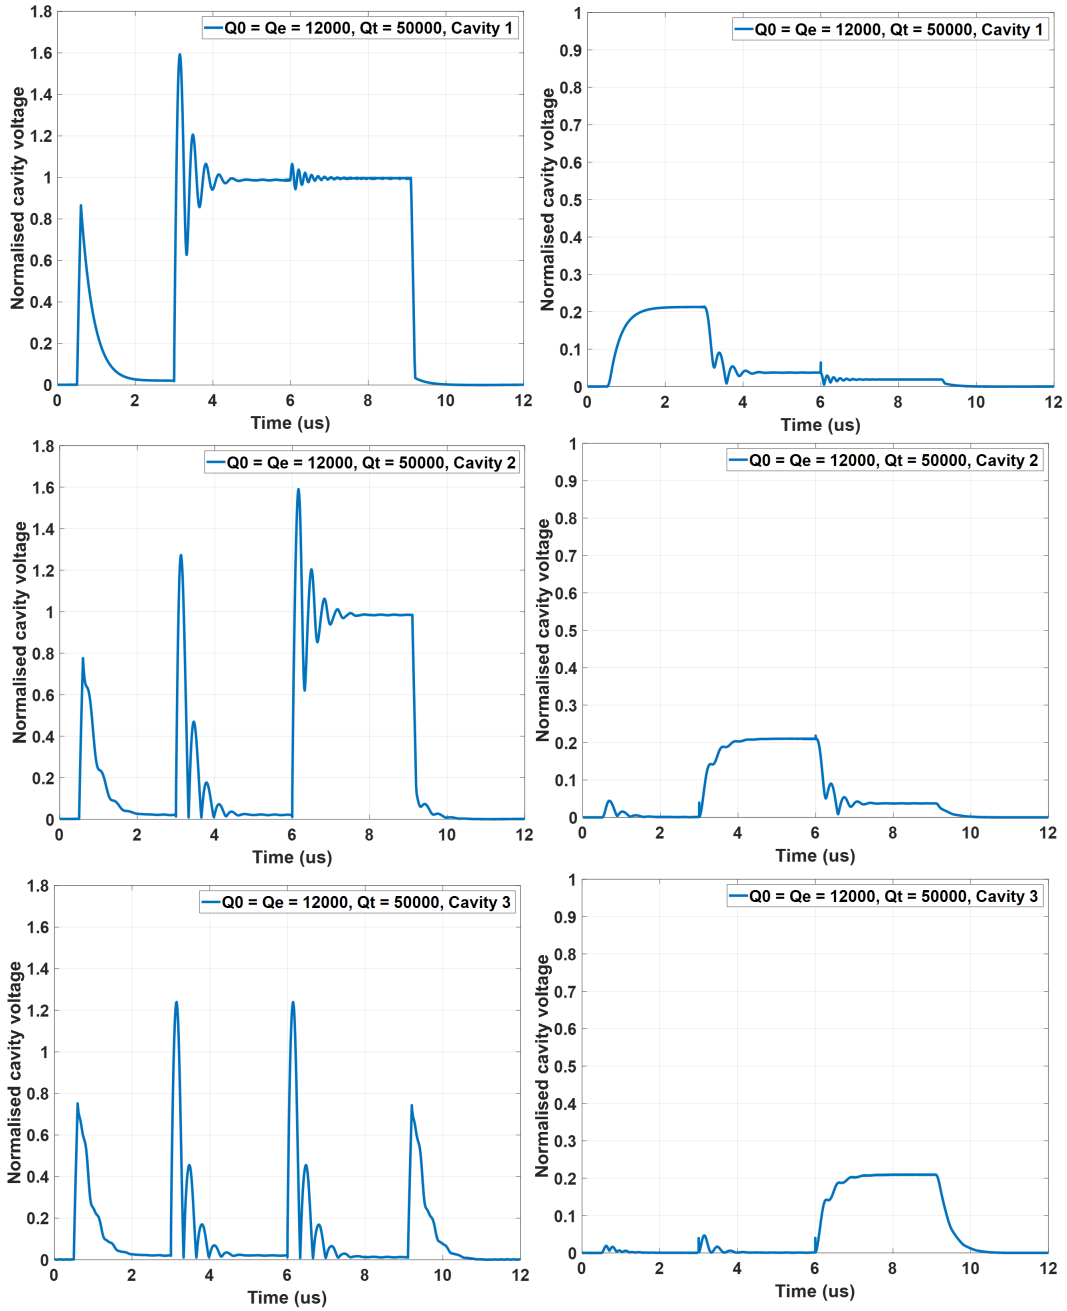


Figure 7.12: Simulation results of three linacs and three circulators, with $Q_0 = Q_e = 12000$ and $Q_t = 50000$. (Top) Reflected (left) and transmitted (right) voltage signals from the first linac. (Middle) Reflected (left) and transmitted (right) voltage signals from the second linac. (Bottom) Reflected (left) and transmitted (right) voltage signals from the third linac.

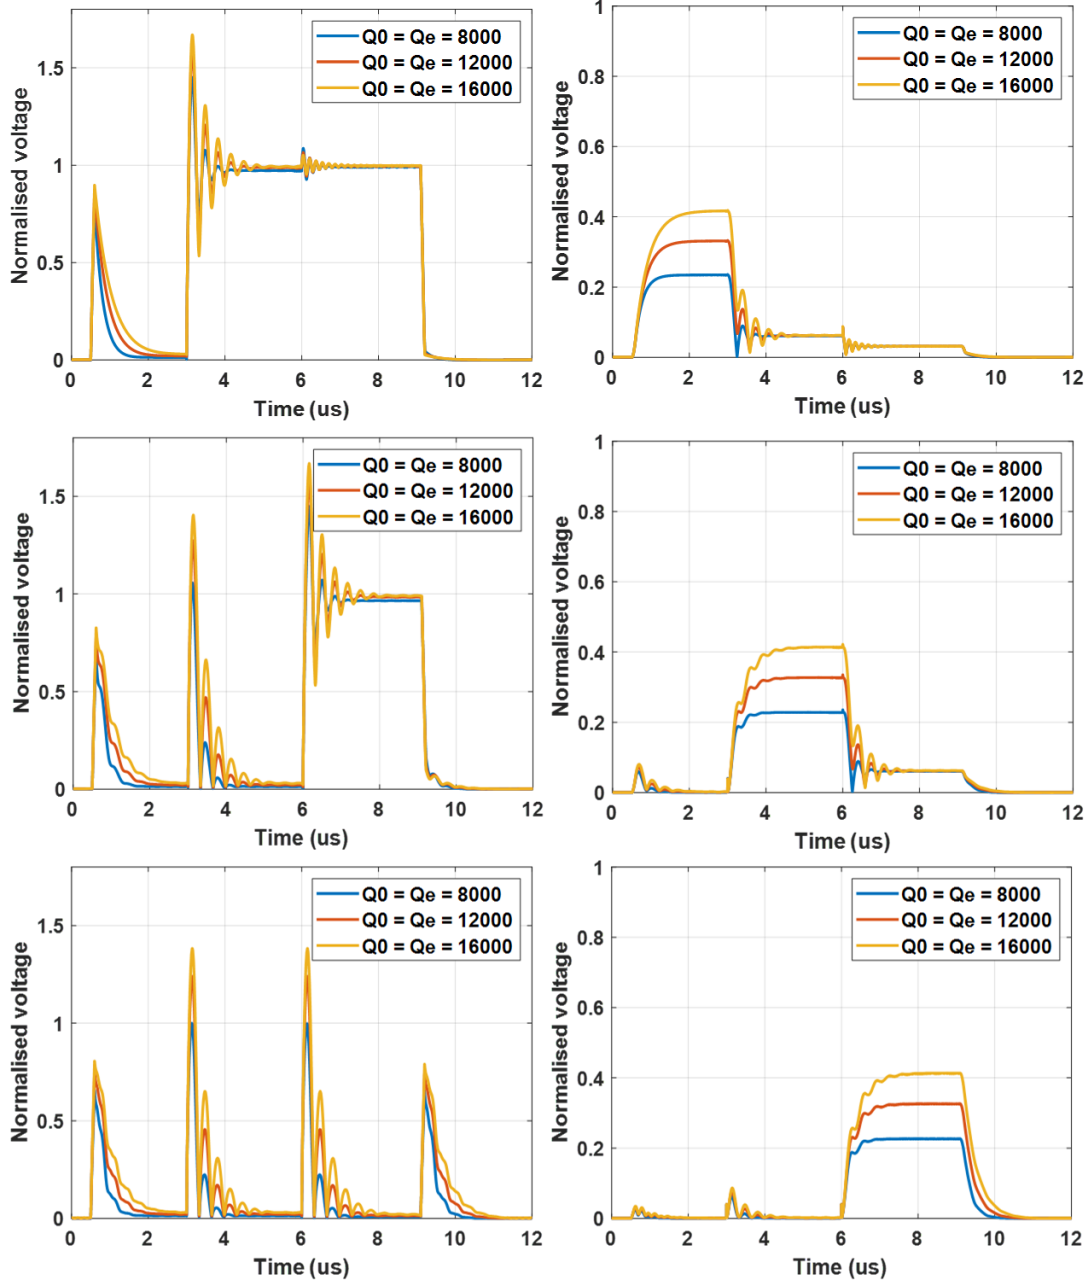


Figure 7.13: Simulation results of three linacs and three circulators, with varying Q_0 and Q_e and $Q_t = 50000$. (Top) Reflected (left) and transmitted (right) voltage signals from the first linac. (Middle) Reflected (left) and transmitted (right) voltage signals from the second linac. (Bottom) Reflected (left) and transmitted (right) voltage signals from the third linac.

the circulator, realistic S-parameters were found for a C-band circulator, shown in Fig. 7.17. They were found using a CST ferrite circulator example, scaled to C-

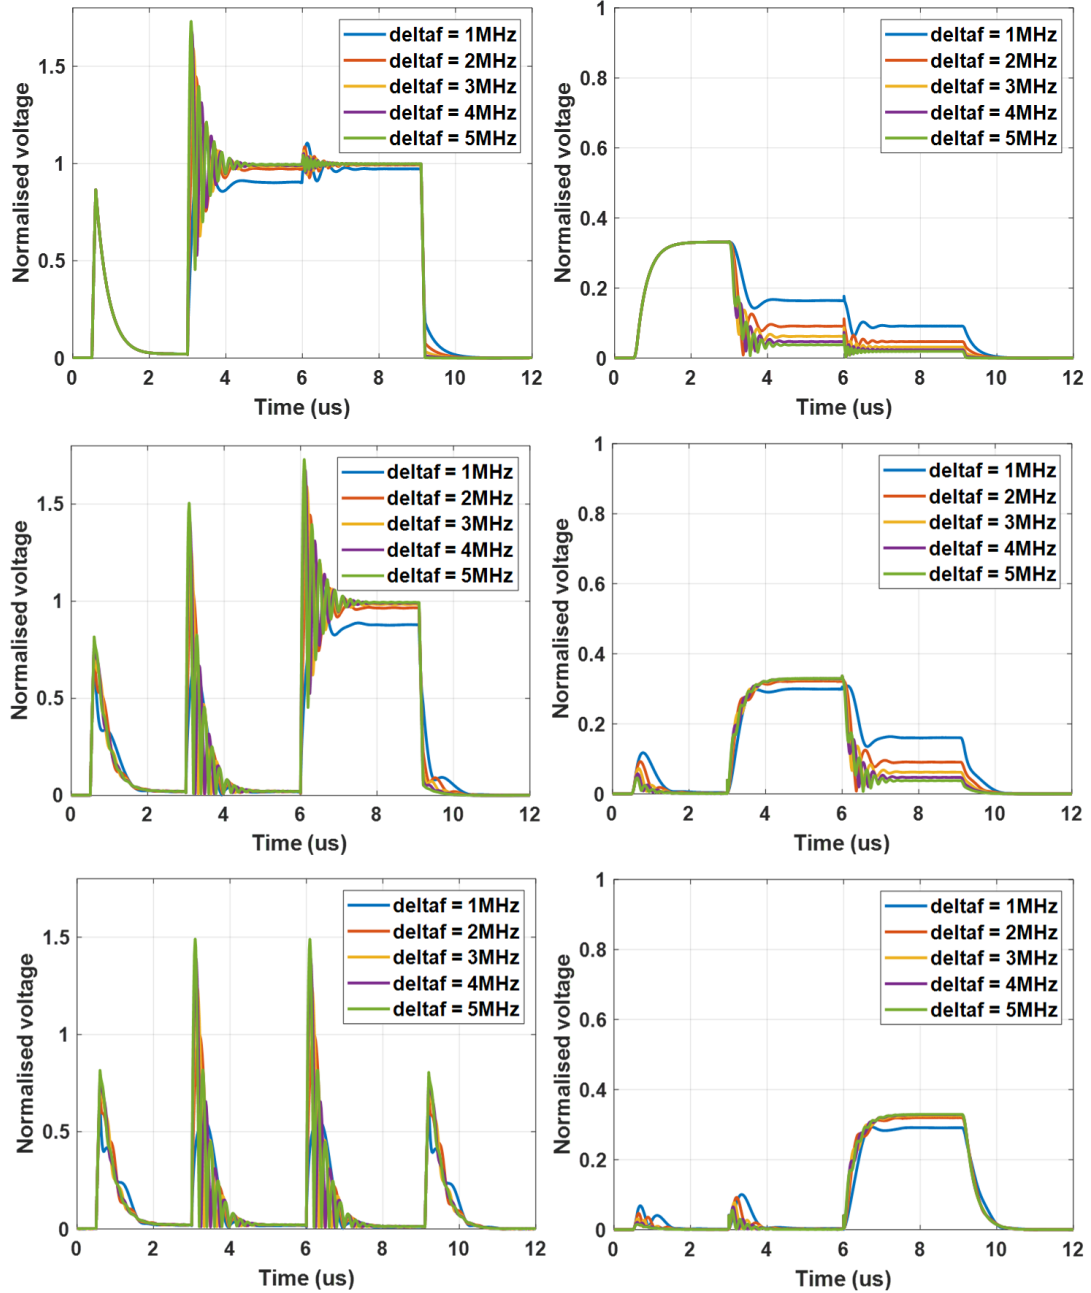


Figure 7.14: Simulation results of three linacs and three circulators, with $Q_0 = Q_e = 12000$ and $Q_t = 50000$ for various values of the frequency difference between the resonances. (Top) Reflected (left) and transmitted (right) voltage signals from the first linac. (Middle) Reflected (left) and transmitted (right) voltage signals from the second linac. (Bottom) Reflected (left) and transmitted (right) voltage signals from the third linac.

band, with similar S-parameters being measured for comparable circulators [168]. Simulating the system using the same input parameters as before led to the results shown in Fig. 7.18 for the reflected and transmitted signals from linac 1. When these are compared to the signals in Fig. 7.12, good agreement is seen in the shape of the pulses, with the only major difference being the amplitude of the signal, as expected, as more losses have been introduced in the simulation.

7.5.3 Power Requirements

To assess whether sufficient power would reach the final linac in the full system, Matlab's Simulink RF Blockset™ was used, with the circuit diagram for the system shown in Fig. 7.19. Three signals were generated, one at each frequency, and then they were combined. The signals then pass through a transmission line, tuned to the correct losses and delay for 9 m of WR-159 waveguide (which is assumed to be the distance from the source to the first linac. There are then two sections of 2 m waveguide connecting each linac. The simulations were performed using circuit envelope simulations, allowing one to simulate high-frequency signals quickly, as the carrier frequencies are handled analytically. The signal that reaches linac 3, after two circulators, 4 m of waveguide and the two other linacs, is shown in Fig. 7.20. As expected, the amplitude of the signal has dropped to around 80 % of its initial value, which would mean that ≈ 70 % of the power reaches the final linac. This suggests that to power the linac with 1 MW, a source capable of delivering 1.3 MW would be required.

7.6 Experimental Verification

A small-scale experimental verification of the frequency switching method was carried out independently by colleagues (Matthew Jones) at Daresbury Laboratory U.K. using a spare 1.3 GHz EMMA S-band cavity. [169] They devised a frequency shifting system using a 15 MHz IF, that could adjust the frequency in 0.291 Hz




Figure 7.15: (Left) Coaxial probe inserted into exit beampipe. (Right) Electric field showing coupling to probe.




Figure 7.16: S11 and S12 for coaxial probe inserted into cavity.

increments [170]. The signal generated by this system is shown in Fig. 7.21, showing similarities to the analytical spectrum shown in Fig. 7.8.

7.6.1 Results

The frequency shifting system was then tested with a spare 1.3 GHz cavity. The IF was mixed up to 1.2975 GHz, 1.3 GHz and 1.3025 GHz, chosen to be within the range of the test cavity. Measurements of the forward and reflected pulses were taken using a directional coupler, and the transmitted power was measured using a cavity field probe.




Figure 7.17: Realistic S-parameters for a C-band circulator.




Figure 7.18: Reflected (left) and Transmitted (right) signals from the first linac, including realistic circulator S-parameters.

The signals measured from the system are shown in Fig. 7.22a. As the cavity had a fill time longer than the pulse length of $2 \mu\text{s}$, the signal never reaches a flat top in the transmitted pulse. It is clear that the first part of the signal has been accepted into the cavity, with the second and third parts being reflected as desired. The frequency beating also appears in the reflected signal, validating the previous simulation results. Only one cavity was available at the time of the test, but this test proved that the method could be implemented with a LLRF system and that it works. High power tests are now planned to test the system with three cavities




Figure 7.20: Signal reaching linac 3 after two circulators and reflection from two linacs showing $\approx 80\%$ of the initial pulse reaching the third linac.





Figure 7.21: Measured signal showing the possibility of adjusting the frequency within the pulse. (Top) Time domain signal (Bottom) Frequency spectrum of the signal. Image reprinted with permission [170].


and MW power levels.

7.6.2 Validating Experimental Results

As a final check in order to verify that the simulations were able to capture similar behaviour to the experimental tests, simulations were repeated using the analytical code and with the EMMA cavity properties shown in Table. 7.2. The



(a) Forward, reflected and transmitted power measured from the first cavity, using the frequency shifting method. Image reprinted with permission. [170]



(b) Simulation results with cavity parameters adjusted to match EMMA cavity, scaled to similar voltages as those measured by colleagues in the experimental tests.

frequencies of the cavities were adjusted to 1.3 GHz, and the Q factors were changed to replicate those of the EMMA cavity used in the tests. A noise level similar to the noise measured from the experiment was also applied to the signal. The agreement between the code and measurements is excellent as shown in Fig. 7.22b, with a similar structure observed in both the reflected and transmitted signals. This gave confidence that a full system with three cavities and higher powers is feasible and realisable with existing LLRF techniques. The code developed can also be used to simulate a variety of scenarios for a full design, such as varying lengths of transmission lines, incorporating S-parameters from measurements of real components that would be used in a real system and modelling a variety of cavity types.

Quantity	Value
Frequency	1.3 GHz
Q_0	20500
Q_e	19500
Q_L	10000

Table 7.2: EMMA cavity parameters used to compare simulations and experimental results.

Chapter 8

X-Ray Imaging and 3D Reconstruction Techniques

8.1 Image Reconstruction

This chapter details the validation work completed on reconstructing quasi-3D X-ray images using only a few views. Probably the most well-known use of image reconstruction is in Computed Tomography (CT), which was developed in the early 1970s and uses acquired X-ray projections to generate medically useful images. It usually involves taking X-rays from multiple angles around the body of a patient and then using reconstruction methods to produce cross-sectional images of the internal parts of the body [171, 172].

CT image reconstruction has evolved over the years, with advancements in both hardware and software contributing to improved image quality and reduced dosage requirements. These same principles are also being leveraged for security applications, with some modifications to how images are generated and processed. Typically in medical CT, the source and detector are rotated around the object of interest (the patient) in about 30 seconds. The object can also be rotated, and this is typically done in laboratories for material defect analysis or dimensional metrology of smaller items and can take a few minutes or longer to complete the scan [173].

For security applications, there is an emphasis on throughput, so faster methods are required. One of these new methods is real-time tomography (RRT) [174]. RRT systems have been developed for baggage scanners and use hundreds of smaller X-ray sources arranged in a circle, which fire sequentially but quickly as the items transverse through the ring of sources. Traditional reconstruction methods can then be used to generate an image [174].

There is now an interest in using similar methods for scanning larger cargo, such as vehicles and unit load devices (ULDs). As the linac-based sources are much more expensive than smaller X-ray sources, it is not feasible to have hundreds of them; therefore, a feasibility study was required to see whether using three or four sources would be sufficient to generate useful reconstructions of cargo in 3D. Due to the large throughput requirement, the scanning speed is essential, so putting a large vehicle onto a rotating platform or having rotating sources were not considered options.

8.2 2D Theory

The basic theory behind image reconstruction is provided here. The DUT is scanned using X-rays, and the intensity is recorded in detectors, as shown in Fig. 8.1. Lambert's law gives the transmitted intensity:

$$I(x, y) = I_0 \exp\left(-\int_{path} \mu(x, y) ds\right) \quad (8.1)$$

where I_0 is the incident beam intensity and $\mu(x, y) = f(x, y)$ is the 2D function to be rebuilt [175]. When the material is homogenous, this reduces to the exponential attenuation law:

$$I = I_0 e^{-\mu d} \quad (8.2)$$




Figure 8.1: Geometry of object being scanned, it is rotated through θ , and the intensity of the transmission is stored for each rotation [175].

8.2.1 Radon Transforms

The Radon transform is a mathematical operation integral to the process of converting X-ray attenuation measurements into meaningful image representations. Mathematically, the Radon transform R of a function $\mu(x, y)$ is the integral of that function along a straight line described by distance t from the origin and angle θ relative to the x -axis:

$$R(\theta, s) = \int_{-\infty}^{\infty} \mu(x, y) \delta(x \cos \theta + y \sin \theta - t) dx dy \quad (8.3)$$

Where t is the distance from the origin along the line, and θ represents the angle of rotation, which could be the X-ray source rotation or the rotation of the object. δ is the Dirac delta function which enforces that the integral is evaluated along the line $t = x \cos \theta + y \sin \theta$. The Radon transform accumulates X-ray attenuation values along lines at various angles, resulting in a set of projections of the object [176–178].

8.2.2 Filtered Back Projection (FBP)

Once the set of projections has been obtained, they are back-projected to reconstruct the image; this is essentially undoing the forward projection operation by mapping

from the detector back to the source. The filtered back projection algorithm is the inverse of the Radon transform given as:

$$\mu(x, y) = \int_0^\pi \int_{-\infty}^{\infty} g(\theta, s) \delta(x \cos \theta + y \sin \theta - t) d\theta ds \quad (8.4)$$

where $g(\theta, s)$ is found by convolving the projections with a filter to remove blurring. This uses similar methods to those described in Chapter 7, where the filtering is performed in the frequency domain by taking the FFT and multiplying, before taking the IFFT to transform back to the spatial domain. The back projection integrates the filtered projections at different angles to reconstruct the original image [176–178].

8.2.3 Number of Projections Required

The Niquist-Shannon sampling theorem determines the number of projections required to reconstruct an image successfully. The theorem states that to avoid aliasing, the sampling rate must be at least twice the highest frequency present in the signal, where for CT, the frequency is the changes in the attenuation along the various directions of the object. For parallel rays, a sampling of S points per projection line requires P number of projections to satisfy the Niquist-Shannon constraint [179]. A good rule of thumb used is that if we assume $\mu(x, y)$ has a diameter D , and the required resolution is R , then the number of projections P required for an accurate reconstruction is

$$P > \frac{\pi D}{R} \quad (8.5)$$

[179]. This means that if the number of projections is reduced then the diameter of the object to be scanned must also be reduced to obtain the same quality of reconstruction. An example of a 2D reconstruction with varying numbers of views is shown in Fig. 8.2. Clearly, when there are only 3 views, there are many artefacts in the image, with the severity of the artefacts improving as the number of projections

increases. The question to be answered was whether reconstructions with few views could be post-processed to gain useful information about the object being scanned.




Figure 8.2: 2D reconstruction using FBP method. (A) Original image (B) 1 projection (C) 3 projections (D) 4 projections (E) 16 projections (F) 32 projections (G) 64 projections [175].

8.3 Algorithms

Three algorithms were used to investigate this problem, as they were part of a simulation package for 3D cone beam reconstruction [180]. The first was standard FBP, as discussed above, which is less computationally expensive than other methods but is more prone to artefacts. The others are discussed below.

8.3.1 Simultaneous Algebraic Reconstruction Technique

Both the Simultaneous Algebraic Reconstruction Technique (SART) and the Maximum Likelihood Expectation Maximisation Technique (MLEM) are iterative in nature, so they refine the image over a number of iterations to reduce the amount of noise and improve the resolution. They start with an initial guess of the image which could be uniform or a low-resolution approximation of the image. They then update the image by correcting it iteratively based on the difference between the image and that calculated from all the projections. A system of linear equations is solved for each pixel, which has contributions from all the projections. Because

of this, they are generally more computationally expensive than FBP, but for cases where there are only a few views to be analysed, they were considered as options for improving the reconstructions. The update equation for SART is given as:

$$x_{i,j}^{k+1} = x_{i,j}^k + \lambda \frac{p_{i,j} - \sum_{k=1}^N A_{i,j,k} x_{i,j}^k}{\sum_{k=1}^N A_{i,j,k}^2} \quad (8.6)$$

where $x_{i,j}^k$ is the voxel value at the i -th row and j -th column of the k -th iteration, $p_{i,j}$ is the measured projection, $A_{i,j,k}$ is the system matrix element for the i -th row, j -th column and k -th iteration and λ is a relaxation parameter [181].

8.3.2 Maximum Likelihood Expectation Maximisation

The MLEM algorithm is a statistical iterative reconstruction algorithm based on the estimation of the maximum likelihood. It assumes that the probability that a photon emitted at pixel j is detected at detection bin i follows a Poisson distribution [182]. The update equation is given by:

$$x_{i,j}^{k+1} = \frac{x_j^k}{\sum_i c_{i,j}} \sum_i c_{i,j} \frac{y_i}{\sum_k c_{ik} \lambda_k} \quad (8.7)$$

where $c_{i,j}$ is the probability that a photon emitted at pixel j is detected at detection bin i and y_i are the raw counts at the detector [183]. This method is generally useful when there is a lot of noise in the image. sparse data sets or when the dose is not high enough [184].

8.4 Methods

To test whether images could be reconstructed with few views, a databank of X-ray images was required. Instead of using real X-rays, the images were simulated using a 3D ray-casting technique [185]. First a 3D model of an object was created in CST studio suite, with an example shown in Fig. 8.3, a script then rotates the object by a set angle and exports the geometry in stereolithography (STL) format. The STL

file is then read into MATLAB and voxelised into a 3D image array [186, 187].




Figure 8.3: CAD model of arbitrary object to test reconstruction.

The geometry was then placed in a simulation setup as shown in Fig. 8.4. The number of pixels for the detector was chosen to be 512×512 pixels. L_2 was initially set to 1 m and the detector panel dimensions were 750×750 mm. The value of L_1 is calculated to ensure that the object fits within the beam as it rotates through 360° . The code used to simulate the X-rays is part of a simulation package written by Michael Behr [188]. Examples of the simulated X-rays for two angles are shown




Figure 8.4: Setup for simulating X-rays of an arbitrary object using ray-tracing.

in Fig. 8.5 where the density is shown in colour (yellow for more dense and blue for less dense).




Figure 8.5: Simulated X-rays of arbitrary object using ray-tracing (Left) Front view (Right) 45 degrees rotation.

8.5 Reconstruction Simulations

The object was rotated 360 degrees with X-rays simulated in 1 degree intervals. These X-ray images were then imported into a 3D cone beam reconstruction package [180]. For the first tests, the Feldkamp-Davis-Kress (FDK) algorithm was used for the reconstruction, which relies on the FBP method [189]. The initial reconstructions are shown in Fig. 8.6, with the vol3d v2 code used to visualise the 3D data [190].




Figure 8.6: (Left) 3D view of reconstruction, showing artefacts around the edges. (Right) Slice through the centre of 3D reconstruction showing reconstructed object in the centre.

As expected, artefacts were present in the image as shown on the left of Fig. 8.6; if a cross-section is taken, it becomes clear that the object was reconstructed, but that extra processing would be required to improve the image. The first step was to remove all the points with values < 0 , this gave the image shown in Fig. 8.7. The reconstruction was much clearer, but artefacts are still present close to the object.




Figure 8.7: 3D image of reconstruction with points < 0 removed. Artefacts are still present, but the form of the object is visible.

Points were then successively removed if they fell below a certain intensity threshold until only points that make up the object remained. Of course, this requires prior knowledge of what the object should look like, but this step could be automated, where the intensity threshold is gradually lowered at the post-processing stage. The final images of the reconstruction are shown in Fig. 8.8.




Figure 8.8: Reconstruction of an object using 360 views in 1-degree increments with the FDK algorithm, after post-processing.

8.6 Threat Reconstruction

A more complicated object was then chosen to test the reconstruction further. A simple weapon model was made, using only the outer shape, with no internal parts or mechanisms, to reduce the complexity of the reconstruction to start with. A simulated X-ray of the model is shown in Fig. 8.9.




Figure 8.9: Simulated X-ray of realistic threat object with no internal parts.

The reconstruction of the threat using the FDK algorithm and 180 views over 360 degrees is shown in Fig. 8.10. The reconstruction was successful, but there were clear artefacts at the edges where the rays crossed boundaries that weren't parallel with the rays.




Figure 8.10: Reconstruction of threat using 180 views split around 360 degrees.

The number of views was then reduced to 5, with X-rays simulated at 0° , 45° , 90° , 135° and 180° . The resulting reconstruction using the FDK algorithm is shown in Fig. 8.11. Even after processing, the images were noisy with many artefacts, and it is hard to see the structure of the object. It became clear that the FDK algorithm would not give the best results when there were only a few views, and so the iterative methods were explored next. The 5-view reconstruction was run




Figure 8.11: Reconstruction of the threat using 5 views and FDK algorithm.

again, this time using the MLEM algorithm for reconstruction. The algorithm was run for 50 iterations, a trade-off between convergence and processing time. The reconstructed image is shown to the left of Fig. 8.12. After processing, the

reconstruction can reproduce the threat, and key features such as the trigger and barrel end can be easily identified. Two views were removed, leaving 3 views, which should be the minimum number required to reconstruct a quasi-3D image. X-ray images simulated at 0° , 45° and 90° were used, with the reconstruction shown to the right of Fig. 8.12. Again, the threat can be identified in form, but some details, such as the end of the barrel, have been lost.




Figure 8.12: (Left) Reconstruction using 5 views and MLEM algorithm. (Right) Reconstruction using 3 views and MLEM algorithm.

8.7 Density Variations

After the successful reconstruction of the threat using five and three views, sections were removed from the inside of the model to create some internal structure and density variation. The new model with included density variations is shown in Fig. 8.13. The reconstruction with 5 views is shown in Fig. 8.14 and 3 views in Fig. 8.15, with both the MLEM and SART algorithms tested. With the 5-view reconstruction, the density variations are identifiable and well reconstructed. When the number of views is reduced to 3, some density information is lost due to the amount of post-processing required to remove enough artefacts to make the threat identifiable. The MLEM and SART algorithms performed equally well, and some density information is visible, but not enough to be confident in post-processing.




Figure 8.13: CAD model of threat with density variations included.

Although these reconstructions do not provide the level of detail that may prove useful in a CT linac system, they point to the possibility that successful few-view reconstructions are at least possible.




Figure 8.14: Reconstruction of threat with density variations using 5 views and MLEM algorithm.

8.8 Distance Variations

The final reconstructions performed were to see if multiple objects with some distance between them, such as boxes, could be reconstructed and whether objects inside the boxes could be identified. This would be closer to a real scenario where




Figure 8.15: Reconstruction of threat with density variations using 3 views. (Left) MLEM (Right) SART

cargo is being scanned at borders or ports. The first model with a set distance between multiple objects is shown in Fig. 8.16. The model consisted of three boxes, two of comparable size and one smaller. The reconstructions were performed with the MLEM algorithm, with the 3 view image shown in Fig. 8.17. The placement of the boxes can be visualised, but the smaller box is not reconstructed at all, and it is unclear what the shape is. The choice of X-ray angles is critical here; if the boxes are positioned so that only one of the X-rays contains information about the smaller box, then the shape will be distorted in the reconstruction, as seen in the image. The main objective was to assess whether the position of the boxes relative to each other would be visible, and in this regard, the reconstruction can be considered successful.

Fig. 8.18 shows the final model that was tested. For this setup, the boxes were made to have a wall thickness, and then an arbitrary object was placed inside each box to represent containers holding items. An example of a simulated X-ray for this setup is shown in Fig. 8.19.

The image reconstruction using 3 views and the MLEM algorithm is shown in Fig. 8.20. The reconstruction required a lot of points to be removed and adjustments of the scales in order for the internal structure to be made out. Despite this, it was




Figure 8.16: CAD model of boxes to test the ability of the algorithms to reconstruct positions accurately.




Figure 8.17: Reconstruction of boxes using 3 views and MLEM algorithm. Reconstruction of the smaller box is unsuccessful due to angles chosen for X-ray simulations.

clear that some information about the internal objects is available, the sphere and cylinder can be identified, but not the block in the third box. By slowly changing the threshold during the processing, the walls of the boxes can be removed, helping to visualise the internal structure. One could imagine an automated process where the threshold is changed and an algorithm for threat identification is run for each value of the threshold.




Figure 8.18: CAD model of boxes that includes a wall and an object inside.




Figure 8.19: Simulated X-ray of boxes with objects inside.

8.9 Further Work

The work presented in this chapter shows that some level of quasi-3D reconstruction is possible using a low number of views (3-5), but significant image processing is required to remove artefacts. There are several challenges that need to be overcome before such a system can be used to provide enough information to be useful in cargo scanning applications. These include:

- Identification of the best positioning of the 3-4 linac system for reconstructing




Figure 8.20: Reconstruction of boxes using 3 views and MLEM algorithm.

the types of object of interest.

- Investigation of different algorithms that are more applicable to few-view reconstruction.
- Finding the optimum post-processing procedure that enables the information of interest to be extracted from reconstructed images.

Due to the recent push in the medical field to reduce patient dosage when performing CT scans, several reconstruction methods for limited-view analyses have been developed. These methods techniques are employed to balance image quality with a reduction in the radiation that the patient would see for a full 360° scan. These techniques include optimisation-based approaches such as those described in [191] and many approaches employing Convolutional Neural Networks (CNNs), which look the most promising. There are several ways that CNNs can aid in the reconstruction process, as shown in Fig. 8.21.

The first option is that they can be trained on images that have been reconstructed using algorithms described in this section, such as SART [192] or FBP methods [193]. After the reconstruction, they are used to extract and suppress the artefacts that are generated to provide cleaner images. The second option is

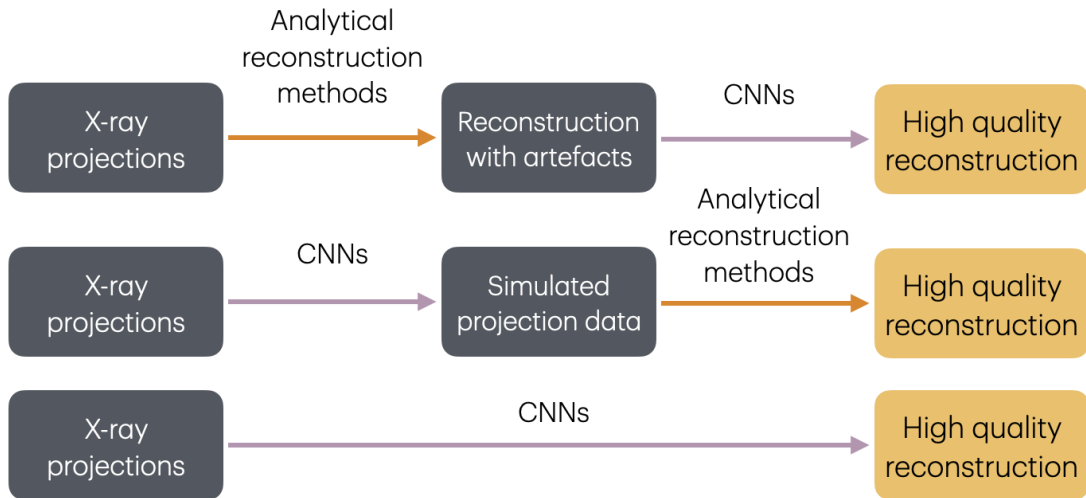


Figure 8.21: Methods of employing CNNs for limited-view CT reconstructions.

to train CNNs to synthesise images from angles that were not originally taken to build a full set; then, standard reconstruction methods are applied. Some examples of the application of this method are given in [194, 195]. These approaches can also be combined, but require much larger data sets as training needs to be performed twice [196].

The final option is to go straight from the limited-view X-ray images to reconstructions, which is a more recent development [197]. This currently requires a very high processing cost, but this represents an ideal solution. These methods could be combined with the recent work on Threat Image Projection in the cargo scanning industry, where simulated X-rays of various threats or objects of interest are projected into real X-ray images of cargo [198]. This is being explored as there is a limited supply of X-ray images that include objects of interest, whereas X-ray images of empty or standard containers are easy to obtain. Employing this method would allow for the generation of the large datasets that will be required to test the methods described here.

Chapter 9

Conclusions

The main research questions that this thesis aimed to answer were as follows:

1. Whether a 2 MeV, compact (< 15 cm), C-band linac could be designed to be as efficient as possible within a set of given design constraints (bi-periodic, $\pi/2$ mode, and low peak fields) using the latest modelling and optimisation techniques.
2. Whether this linac could achieve a high beam capture efficiency (> 70 %) from a thermionic cathode using re-capture methods.
3. If the linac could be incorporated into a low-cost RF system, where multiple linacs can be fired and synchronised to generate quasi-3D CT images of light cargo or small vehicles.

Throughout the thesis, the design has been developed and shown through analytical calculations and simulation-based tools that it would fulfil these requirements, with several novel design techniques and methodologies being applied. The key highlights and contributions to future designs that have resulted from the work presented in the thesis are listed below:

- The extensive use of MO optimisation and NURBS modelling techniques facilitated a Pareto analysis design approach where the key trade-offs between

objectives could be easily visualised. This enabled an optimum design within the constraints of the peak fields required by the industrial partners.

- The incorporation of chamfers on the cells allowed for an improvement of 28 % in the coupling constant, with little effect on the shunt impedance and peak fields. These chamfers also allowed for fine-tuning of the RF amplitudes in the cells, a tuning parameter typically unavailable in other linac designs.
- Using MO optimisation and re-capture methods, the design achieved a high capture efficiency of 90 %, with minimal cathode back-bombardment. This approach will improve cathode lifetimes and reduce the frequency of component replacement on many low-energy linac systems.
- A comprehensive thermal analysis and set of thermal simulations were performed, which gave detailed insights into the operating limits of the linac. Various methods for performing a thermal analysis were explored and compared. Incorporating bends into the analytical formulas used to estimate the heat transfer coefficient will allow other designers in the future to reduce the amount of time and effort required to accurately estimate the thermal performance of other linac designs.
- The linac design has been integrated into a novel full RF system design that would be capable of operating three linacs simultaneously using only one RF source. This method uses frequency division multiplexing techniques, delivering three frequencies to three linacs with minimal loss and field distortion. Initial experimental measurements have shown that this method is viable. This will help reduce the cost and complexity of future systems where multiple linacs need to be operated and synchronised.
- A study of basic 3D image reconstruction techniques on a series of simulated X-rays showed that it should be possible to generate quasi-3D images using only a few X-ray angles, potentially opening up the possibility of low-cost

linac-based CT scanning systems.

In the following sections, the conclusions of each chapter are presented, the key contributions and results are highlighted, and the limitations of this research are discussed. Finally, improvements to the design are suggested, along with future work that would be required to realise the full design.

9.1 RF Design and Optimisation

The design specifications for the linac were set out in Chapter 3, with the main requirements from the industrial partners being that it was a C-band system, would provide a 2 MeV beam, and would use similar values of the pulse length and electron gun voltages that are currently in use. They also required the design to be as robust as possible, minimise the cost of manufacturing and operation where possible and have a small footprint for the potential use as a portable system. These specifications influenced many of the major design decisions, with the key ones being the type of cavity and mode of operation chosen (bi-periodic $\pi/2$ mode, on-axis coupling), the conservative peak field limits (peak electric field < 100 MV/m, peak $S_c < 1$ W/ μm^2 , peak magnetic field < 229 mT) and the frequency (5.712 GHz). In Chapter 3, the RF design of a 2 MeV, 5-cell bi-periodic linac was presented, detailing the optimisation process used to give the best RF performance possible within these requirements. The final parameters of the $\beta = 1$ single-cell design are given in Table. 9.1, showing that they meet the requirements with all values of the peak fields falling below or are close to the limits imposed. It was found that the final design was limited by the peak electric field requirement, with overhead on the other two peak fields. Although no discussions were held about whether the peak electric field limit could be pushed higher, if it could have been increased to 125 MV/m, a shorter design would have been possible with a higher gradient, reducing the total footprint. If this is considered an option in the future, then the Pareto optimal sets generated during the design can easily be used to select a new geometry with minimal effort.

Quantity	Value
Frequency	5.712 GHz
Gradient, $\beta = 1$ cell	26 MV/m
Peak electric field, E_{pk}	103 MV/m
Peak magnetic field, B_{pk}	96 mT
Modified Poynting vector, S_c	0.73 W/ μm^2
Shunt impedance per unit length, Z	89.4 M Ω /m
Coupling constant, k_c	0.9 %
Quality factor, Q_0	12400
Aperture radius	5 mm
Wall thickness	2.5 mm
Coupling cavity length	2 mm

Table 9.1: Final $\beta = 1$ cell parameters.

A comparison of the design with some typical X-band structures is given in Table 9.2. The key differences include the ability to sustain higher maximum electric fields and smaller aperture radii leading to an improved shunt impedance. Although there are clear benefits in moving to X-band, the design presented in this thesis should be competitive, both in terms of cost and size. This is mostly due to the more relaxed tolerances on the design and the availability of good performance C-band RF sources.

Quantity	C-band Design	X-band 1	X-band 2	X-band 3
Frequency (GHz)	5.712	9.3	9.3	12
Type	SW	SW	SW	TW
Gradient (MV/m)	26	10	-	100
Energy (Mev)	2	1	6	-
E_{pk} (MV/m)	104	48.3	-	245
B_{pk} (mT)	156	91	-	623
Z (M Ω /m)	89.4	-	148	100.5
k_c (%)	0.93	5	2.5	-
Quality factor, Q_0	12400	7328	8800	6183
Aperture radius (mm)	5	5	2	2.75
Input power (MW)	1	0.5	1.8	63.8
Number cells	5	8	19	24

Table 9.2: Comparison with typical X-band linacs. X-band 1 [38], X-band 2 [199], X-band 3 [200]. Parameters are given as an average for the TW structure.

These values were achieved using the latest modelling and optimisation techniques, namely:

- The inclusion of NURBS modelling methods allowed for a 10 % improvement in the shunt impedance at low values of E_{pk}/E_{acc} compared with traditional modelling methods. These shapes were considered with the knowledge that modern manufacturing techniques allow for machining complex geometries.
- The use of MO optimisation algorithms to generate Pareto optimal sets showing the trade-offs in the peak fields and shunt impedance for any cavity geometry.
- The separate optimisation of each individual cell length instead of scaling one or two lengths. This allowed each cell to operate at the E_{pk}/E_a limit.

There was also a question of whether on-axis coupling was the best choice for the design. Due to the requirement that the weight of the shielding required should be minimised, on-axis coupling was selected. Using side-coupled cells would have potentially allowed for the performance to be improved further as the shunt impedance and k_c can be increased using this method. If the shielding required is less of a concern, such as for a fixed non-portable system, then a design using side-coupled cells should be used.

In Chapter 3, the use of chamfers to increase the coupling between the individual accelerating cells was also introduced. Although blends have been investigated before to increase the coupling between cells, as far as the author is aware, a variable length/angle chamfer between cells has not been previously exploited for an industrial linac design. The chamfers were primarily introduced to squeeze out extra cell-cell coupling without increasing the beampipe radius and sacrificing the shunt impedance. It was later realised that this has the added benefit of providing an extra tuning parameter for adjusting the field amplitudes. This method should be incorporated into designs that require a specific RF field amplitude configuration.

The tuning of the final linac was performed manually, with some insights from the circuit theory presented in Chapter 4. This process was quite time-consuming due to the large number of variables (cell radii, chamfers, coupling port parameters) and constraints (field flatness, frequency, minimum S_{11}). Modern methods such as machine learning/neural networks could be used to train a model of a cavity, which may then allow for automatic cavity tuning. This was attempted during the course of the design, but due to the large sensitivity of some of the outputs, the model was unable to predict all the cavity parameters accurately enough to be useful.

For cell optimisation, including considerations of the coupling constant, the coupling proxy developed in Chapter 3 only provided a rough estimate for the coupling constant within a small range of chamfer angles and lengths. Ideally, the real value of the coupling would be calculated within the optimisation to ensure that the design point selected was actually the best in terms of the various trade-offs. Including the coupling constant in the calculation would require extra computational steps as the cell would need to be re-simulated with different boundary conditions, but this removes the possibility of selecting a sub-optimal design. This would be a recommendation for designers in the future.

9.2 Note on Multi-objective Optimisation

MO methods were used extensively in the design of the linac, both for the electromagnetic design and capture optimisation, and they have the potential to be applied to many other areas of linac design. Although multi-objective optimisation methods have previously been used on some parts of linac design, no comprehensive studies have been performed that showcase the power of the approach for entire linac designs. Having the ability to perform a single set of optimisations and then being able to quickly select a new optimal design for a linac (where the design specifications are often changed) is extremely powerful and is an improvement over traditional methods. One of the limitations of the optimisation in the thesis is

that only well-known, established algorithms were used; there is definitely scope to investigate other algorithms potentially allowing for improved convergence speeds. The recommendation would be to use MO methods when there is a set of conflicting objectives and where it would be beneficial to visualise the trade-offs between them. When optimising cavities, performing the initial optimisations with a smaller mesh (30 cells per wavelength, for example) is also recommended. This enables a much faster first-pass optimisation and provides a set of input variables that roughly lead to optimum designs. These points can then be used as a seed for a second optimisation with a more accurate mesh to provide the final results. This also enables the optimiser to test different algorithms, input limits, and objective weightings at a much faster rate.

9.3 Beam Capture Optimisation

The optimisation detailed in Chapter 4 builds upon recent studies attempting to achieve the highest possible capture efficiency from thermionic cathodes in industrial linacs. These studies have shown that high capture efficiencies are possible by appropriately adjusting the first few cell lengths and amplitudes. Typically, values ranging from 40-60 % are normally achieved in industrial linacs, with the design presented here achieving over 90 % using re-capture methods. The incorporation of MO optimisation techniques here, too, allowed for rapid design adjustments and provided extra insight into the effects of the cell lengths and amplitudes, which may have been hard to realise using parameter scans alone. For example, the third cell length needed to be slightly shorter than $\beta=1$ to improve the transition into the rest of the cells. The beam was optimised such that 84 % of the electrons reaching the target were considered to have useful energies (> 1.5 MeV), and the average energy of these useful electrons was 1.93 MeV, close to the 2 MeV design point. One of the concerns that arose from the design was whether such a low amplitude (30 % of the maximum amplitude) in the first cell would be achievable and whether

it would be too sensitive to manufacturing errors, causing amplitude and length errors. A tolerance analysis was performed, showing that the performance was not degraded with a $\pm 10\%$ error on the length but that the design was very sensitive to the amplitude of the first cell, with the percentage of electrons returning to the cathode doubling for a 10% amplitude error. This loss was found to be recoverable by adjusting power to the linac so that the amplitude of the first cell is closer to the design point, at the expense of higher fields in the other cells or a slightly reduced beam energy. This recapture method should be employed in industrial linac designs where a thermionic cathode is used, as it can significantly improve cathode lifetimes and reduce the cost of replacing components. This is especially important for linacs operating in remote regions, as security linacs often are, or areas lacking linac experts/technicians available for repairs.

One of the limitations of the optimisation approach employed here was that beam optimisation was performed separately from the RF optimisation; in the future, there could be the potential for a combined approach where the RF fields are optimised in parallel with the beam for a set of desired final beam properties. Many commercial code suppliers are already working towards unified approaches where many solver types are integrated, but it is normally the case that there are varying levels of accuracy for accelerator-specific applications for the different code families. The author believes developing one optimisation framework using the most applicable codes for linac design problems could be a worthwhile pursuit.

9.4 Linac Thermal Analyses

Typically, a thermal analysis of the RF cavity is performed at the end of the design as a check, but they do not always incorporate CFD simulations to encapsulate the effects of water flowing in the cooling pipes. As a result, large safety factors are often used to avoid large temperature rises and ensure safe operation.

For the linac design presented here, it was found that to keep the frequency of the

linac constant using a minimum water temperature of 20°C, 1.2 kW average power would be permissible for a 5 L/min flow rate. Assuming that the source frequency had a bandwidth large enough to accommodate a 1MHz frequency shift, this limit is increased to 1.6 kW. A mechanical analysis found that the thermal deformation at 800 W input power was in the region of 5-10 μm , causing a maximum error of 1 %, seen in the first cell. When the power is increased to 1.2 kW, the field error was 1.5 %. The analysis in Chapter 4 suggested that a 1.5 % amplitude error should not affect the beam significantly, though this was not checked directly due to time constraints. Under normal operating conditions (800 W average input power), there should be no thermal limitations when running. If the powers are pushed higher, then the linac would likely be limited by the ability to shift the frequency within the range of the source rather than field errors reducing the beam performance.

As there are few comprehensive studies of the optimum ways to both estimate and simulate thermal effects in low-power industrial linacs, a variety of methods for modelling the thermal heating effects in the linac were considered in Chapter 6, these included:

- Representing the water with a fixed temperature applied to the walls.
- Using a SS thermal solver with a constant water temperature and a constant applied HTC calculated using analytical approximations.
- Using a SS thermal solver with fluid solver elements, giving a real temperature rise along the pipes and a constant applied HTC calculated using analytical approximations.
- Full CFD analysis that incorporates fluid flow in the pipes.

The results presented from these analyses show that a CFD analysis is not required for the majority of linac designs. When the design is not required to operate close to the limits for copper, such as at very high repetition rates (> 800 Hz) or high powers, some error in the estimated temperature distribution is acceptable. The

introduction of the effects of pipe bends on the calculations of the Nusselt number gave an error of 0.14 % in a temperature that could be measured on the cavity surface. Not including these effects only leads to an error of 0.3 %, so there is an insignificant difference for most cases. Still, incorporating this method into future designs will result in more accurate estimates of the HTC than the 1D formulas provide, allowing for safety factors to be reduced slightly and a performance boost.

The author recommends that for the large majority of low power linac designs, using an analytical approximation for the HTC, and then using the average temperature of the water for the fluid in the pipes would provide enough accuracy. If more accurate results are required, using fluid elements instead of a constant temperature is recommended, with only a moderate increase in complexity and resource requirements. Finally, a full CFD analysis is recommended for cases where designs operate close to thermal limits. The calculations presented in the thesis need to be verified with an experimental campaign, where the temperature of the linac is measured using probes and compared with the estimated values. Due to time constraints, there were several other limitations of the study that will need to be addressed:

- Only one type of pipe geometry and placement configuration was explored. The radii of the pipe bends were fixed, though it is likely that there would have been an optimal value for the flow rates considered. There may also be other configurations that reduce differential heating across the linac. It is also not clear whether the approximations would work as well for smaller radii pipes or shorter straight sections, where the approximations may deviate further from reality. Considerations of whether the pipes should be placed internally or externally and how they would be manufactured should also be made.
- It is unclear how well the approximations will work with higher flow rates or for different turbulence regimes.

- Only a few turbulence models were tested for the CFD simulations; this would need further investigation for higher flow rates and different pipe geometries.

9.5 RF System Design

Chapter 7 introduced an RF system design allowing three linacs to be fired with a chosen delay time between them using only one RF source. This would significantly reduce the costs of constructing a high-energy CT scanning system capable of scanning containers or light vehicles. Various options were considered, including resonant delay lines and waveguides operating in higher-order modes. The final design selected is based on a frequency division multiplexing approach whereby one RF pulse will contain three frequencies separated by approximately 3 MHz. The pulse would then travel through a series of circulators while providing power to three linacs that will be tuned to the three frequencies. This idea was validated through both code-based simulations and Simulink. By incorporating realistic S-parameters for the linacs, circulators and waveguides, it was shown that 70 % of the initial power in the last section of the pulse would reach the final linac and that the signal reaching each linac would be stable. One of the shortcomings of the RF system design is that there was limited opportunity to test this method at higher powers. The low-power tests performed by colleagues at Daresbury Laboratory, U.K., look promising, as they agree with simulations of a similar system, but it remains to be seen whether the system will work with three cavities and at higher powers. To validate this, three cavities must be manufactured and integrated with a LLRF control system and high-power RF source.

9.6 Quasi-3D Image Generation using Linacs

The topic of 3D image reconstruction using only a few X-ray angles was only briefly covered in Chapter 8, and there are a large number of unknowns for how a system

using only a few angles could perform in practice. The algorithms explored are likely not optimal for few-view X-ray analysis, and this is something that could be explored in much greater detail. Despite this, the results using the basic well-known algorithms and 3D reconstruction techniques were able to generate adequate 3D reconstructions from simulated X-ray images of various objects.

9.7 Future Work

The first step towards realising the design presented in this thesis would be to manufacture a prototype cavity. Once manufactured, RF tests should be performed, including S-parameter measurements to validate the resonant frequency and a bead-pull measurement to check the on-axis electric field. It is likely that the field amplitude is not initially correct, and the cavity will require tuning; this could be informed using the coupling model described in Chapter 4, allowing for the correct adjustments to be made to match the desired amplitudes in each of the cells. The cavity must then be integrated into the full system, including the X-ray target, thermionic cathode, RF source, HV supply, modulator and cooling network. The high capture efficiency would be validated through current measurements and temperature probes used to validate the thermal analysis from Chapter 6. Preliminary tests for the validation of the cavity multiplexing system could be performed in the new RF bunker at Daresbury Laboratory, U.K. This bunker will operate with an S-band Klystron but has a circulator that would be appropriate for testing at least the high power reflected signal from a pulse containing three frequencies ($2.998 \text{ GHz} \pm 3 \text{ MHz}$). The current LLRF system would need to be modified to include the three-frequency generation, which can then be sent to the solid-state power amplifier (40 MHz BW) before going to the Klystron.

The forward and reflected power could be measured using existing directional couplers and compared with simulations of a similar configuration. The experimental setup for such a test is shown in Fig. 9.1. Three frequencies would be generated at

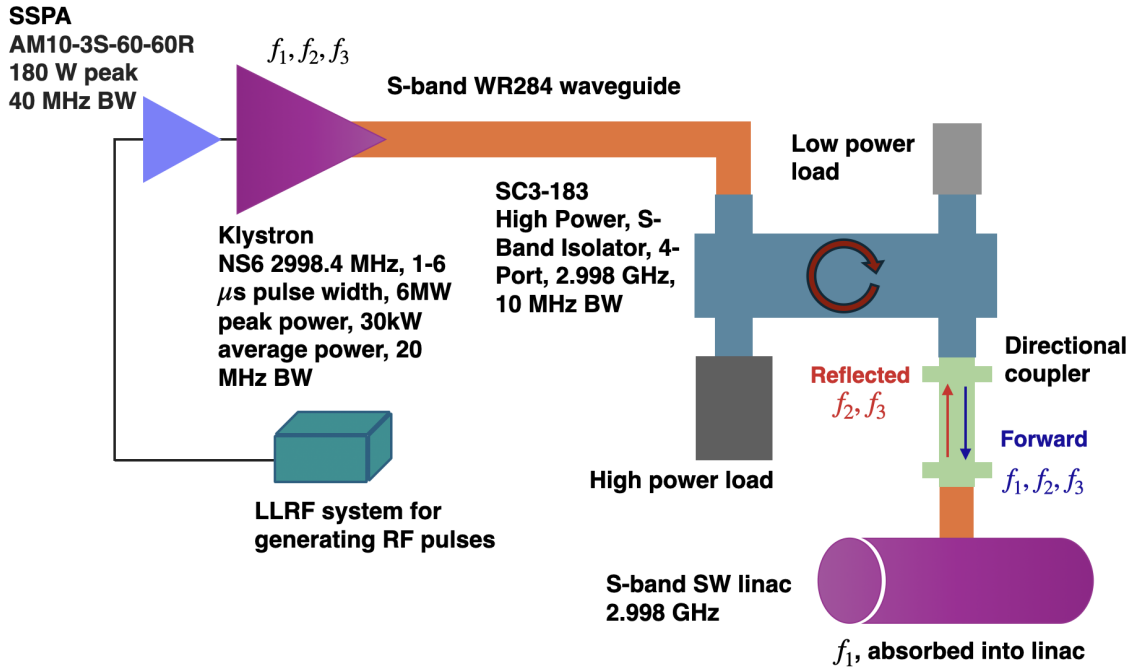


Figure 9.1: Experimental setup for testing FDM technique using RF bunker at Daresbury laboratory.

2.995 GHz, 2.998 GHz and 3.001 GHz and used to generate a 5 μs pulse with these three frequencies on separate parts of the pulse. These would travel to one S-band linac, and the forward and reflected signals from the linac would be measured using a directional coupler. The expected results from a simulation of a similar arrangement are shown in Fig. 9.2. The Q factors of the linac were set to the following values: $Q_0 = 13000$, $Q_e = 12500$, $Q_t = 50000$, representative of the first cavity to be tested in the RF bunker. Measurements of taken at the directional coupler connected to the linac should be compared with these simulations to verify the results.

9.8 Final Remarks

The research presented in this thesis aimed to investigate a high capture efficiency, low-cost, 2 MeV, C-band linac solution for 3D scanning applications. The design presented is capable of achieving a 90 % beam capture efficiency and has been integrated into a linac CT system design, which will be able to generate 3D images.




Figure 9.2: (Left) Simulated reflected pulse from S-band linac showing first frequency (2.995 GHz) being absorbed into linac. (Right) Transmitted pulse into S-band linac showing linac accepting first frequency.

Once built and shown to be successful then these types of systems can be used to improve security at borders, ports and airports worldwide. The techniques and methodologies developed throughout the thesis are also not security linac specific, and given the widespread usage of low-energy linacs, they are applicable to any low-energy linac design.

References

- [1] P. J. Bryant. “A Brief history and review of accelerators”. In: *CERN Accelerator School: Course on General Accelerator Physics*. 1992, pp. 1–16.
- [2] Rolf Wideröe. “Über ein neues Prinzip zur Herstellung hoher Spannungen”. In: *Archiv für Elektrotechnik* 21 (1928), pp. 387–406.
- [3] The Insight Partners. *Linear Accelerator Market Forecast to 2027 - COVID-19 Impact and Global Analysis by Energy Range, Application, Industry, and Geography*. The Insight Partners, 2020.
- [4] Serguei Chatrchyan, EA de Wolf, P Van Mechelen, et al. “The CMS experiment at the CERN LHC”. In: *Journal of instrumentation.-Bristol, 2006, currens* 3 (2008), S08004.
- [5] Settimio Mobilio, Federico Boscherini, and Carlo Meneghini. *Synchrotron Radiation*. Springer, 2016.
- [6] IA Bataev et al. “Surface hardening of steels with carbon by non-vacuum electron-beam processing”. In: *Surface and Coatings Technology* 242 (2014), pp. 164–169.
- [7] Andrzej G Chmielewski et al. “Recent developments in the application of electron accelerators for polymer processing”. In: *Radiation Physics and Chemistry* 94 (2014), pp. 147–150.
- [8] AG Chmielewski, T Sadat, and Z Zimek. “Electron accelerators for radiation sterilization”. In: *Trends in radiation sterilization of health care products* (2008).
- [9] Robert W Hamm and Marianne E Hamm. *Industrial Accelerators and Their Applications*. WORLD SCIENTIFIC, 2012. DOI: 10.1142/7745. eprint: <https://www.worldscientific.com/doi/pdf/10.1142/7745>. URL: <https://www.worldscientific.com/doi/abs/10.1142/7745>.
- [10] *Workshop on Energy and Environmental Applications of Accelerators*. Tech. rep. USDOE Office of Science, Washington, DC (United States), 2015.
- [11] Samy Hanna. *RF linear accelerators for medical and industrial applications*. Artech House, 2012.
- [12] Robert Apsimon et al. “JACOW: Electron Beam Based Leather Tanning”. In: *JACoW LINAC 2022* (2022), pp. 645–649.

-
- [13] Peter St John. “Counterterrorism Policy-Making: The Case of Aircraft Hijacking, 1968—1988”. In: *Democratic Responses to International Terrorism*. Brill Nijhoff, 1991, pp. 67–121.
- [14] Alona E Evans. “Aircraft Hijacking: What is Being Done”. In: *American Journal of International Law* 67.4 (1973), pp. 641–671.
- [15] David J Haas. “Electronic security screening: Its origin with aviation security 1968–1973”. In: *Journal of applied security research* 5.4 (2010), pp. 460–532.
- [16] Joann Peterson and Alan Treat. “The post-9/11 global framework for cargo security”. In: *J. Int’l Com. & Econ.* 2 (2009), p. 1.
- [17] William T Eckles. “Importing Nuclear Weapons Through the Selectively-Permeable Border of the United States”. In: *InterAgency Journal* 8.2 (2017).
- [18] Jessica Romero. “Prevention of maritime terrorism: The container security initiative”. In: *Chi. J. Int’l L.* 4 (2003), p. 597.
- [19] Daphne F Jackson and David J Hawkes. “X-ray attenuation coefficients of elements and mixtures”. In: *Physics Reports* 70.3 (1981), pp. 169–233.
- [20] DJ Mistry, TA Cross, CR Weatherup, et al. “RF sub-systems for cargo and vehicle inspection”. In: *WEIB02, in Proceedings, IPAC 14* (2014).
- [21] Conyers Herring and MH Nichols. “Thermionic emission”. In: *Reviews of modern physics* 21.2 (1949), p. 185.
- [22] Albert W Hull. “The magnetron”. In: *Journal of the American Institute of Electrical Engineers* 40.9 (1921), pp. 715–723.
- [23] Gongyin Chen et al. “Linatron mi6, the x-ray source for cargo inspection”. In: *Physics Procedia* 66 (2015), pp. 68–74.
- [24] Rapiscan Systems. *Rapiscan Products, Cargo and Vehicle Inspection, EAGLE P60*. Ed. by Rapiscan Systems. URL: <https://www.rapiscan-ase.com/products/portal/eagle-p60-cargo-inspection> (visited on 01/03/2024).
- [25] Rapiscan Systems. *Rapiscan Products, Cargo and Vehicle Inspection, EAGLE G60*. Ed. by Rapiscan Systems. URL: <https://www.rapiscansystems.com/en/products/rapiscan-eagle-g60> (visited on 01/03/2024).
- [26] Rapiscan Systems. *Rapiscan Products, Cargo and Vehicle Inspection, EAGLE R60*. Ed. by Rapiscan Systems. URL: <https://www.rapiscansystems.com/en/products/rapiscan-eagle-r60> (visited on 01/03/2024).
- [27] Thomas W Rogers et al. “Automated x-ray image analysis for cargo security: Critical review and future promise”. In: *Journal of X-ray science and technology* 25.1 (2017), pp. 33–56.
- [28] S Ogorodnikov and V Petrunin. “Processing of interlaced images in 4–10 MeV dual energy customs system for material recognition”. In: *Physical Review Special Topics-Accelerators and Beams* 5.10 (2002), p. 104701.

- [29] S Kutsaev et al. “Electron accelerators for novel cargo inspection methods”. In: *Physics Procedia* 90 (2017), pp. 115–125.
- [30] Y Liu, BD Sowerby, and JR Tickner. “Comparison of neutron and high-energy X-ray dual-beam radiography for air cargo inspection”. In: *Applied Radiation and Isotopes* 66.4 (2008), pp. 463–473.
- [31] Aleksandr Y Saverskiy, Dan-Cristian Dinca, and J Martin Rommel. “Cargo and container X-ray inspection with intra-pulse multi-energy method for material discrimination”. In: *Physics Procedia* 66 (2015), pp. 232–241.
- [32] AA Zavadtsev et al. “A dual-energy linac cargo inspection system”. In: *Instruments and Experimental Techniques* 54.2 (2011), pp. 241–248.
- [33] George Zentai. “X-ray imaging for homeland security”. In: *International Journal of Signal and Imaging Systems Engineering* 3.1 (2010), pp. 13–20.
- [34] Thorsten M Buzug. “Computed tomography”. In: *Springer handbook of medical technology*. Springer, 2011, pp. 311–342.
- [35] Godfrey N Hounsfield. “Computed medical imaging”. In: *Science* 210.4465 (1980), pp. 22–28.
- [36] Olive Emil Wetter. “Imaging in airport security: Past, present, future, and the link to forensic and clinical radiology”. In: *Journal of Forensic Radiology and Imaging* 1.4 (2013), pp. 152–160.
- [37] Polad M Shikhaliev. “Large-scale MV CT for cargo imaging: a feasibility study”. In: *Nuclear instruments and methods in physics research section a: accelerators, spectrometers, detectors and associated equipment* 904 (2018), pp. 35–43.
- [38] M Jenkins et al. “Prototype 1 MeV X-band linac for aviation cargo inspection”. In: *Physical Review Accelerators and Beams* 22.2 (2019), p. 020101.
- [39] Dr Graeme Burt. *Linacs for Cargo Screening, CERN High gradient Day 2015*. https://indico.cern.ch/event/336335/contributions/789053/attachments/657779/904349/Security_CLIC_Jan15.pdf. Accessed: 2021-07-20. 2015.
- [40] Youngwoo Joo et al. “Design of 9/6 MeV S-band Electron Linac Structure with 1.5 Bunching Cells”. In: *9th International Particle Accelerator Conference*. June 2018. DOI: 10.18429/JACoW-IPAC2018-THPMK139.
- [41] Ahmadiannamin Sasan et al. “Design and simulation of side coupled six MeV linac for X-ray cargo inspection”. In: *Proc. IPAC’14* (2014), pp. 3844–3846.
- [42] Ch Tang et al. “Low-energy linacs and their applications in Tsinghua University”. In: *Proc. LINAC 2006* (2006), pp. 256–258.
- [43] Wei Bai et al. “Physical design of 2 MeV standing wave accelerating tube”. In: *High Power Laser and Particle Beams* 17.6 (2005), pp. 917–920.

-
- [44] L. Yu. Ovchinnikova and V. I. Shvedunov. “Electron Accelerator for Radiation Therapy with Beam Energy 6-20 MeV”. In: (July 2020). arXiv: 2007.05987 [physics.acc-ph].
- [45] H Yang et al. “Design of Compact C-band Standing-wave Accelerating structure Enhancing RF Phase Focusing”. In: *Proc. of IPAC12* (2012), pp. 20–25.
- [46] Jiahang Shao et al. “Development of a C-band 6 MeV standing-wave linear accelerator”. In: *Physical Review Special Topics-Accelerators and Beams* 16.9 (2013), p. 090102.
- [47] Yu Kubyshin et al. “C-band linac optimization for a race-track microtron”. In: *EPAC*. Vol. 8. 2008, p. 778.
- [48] E Tanabe et al. “Medical applications of C-band accelerator technologies”. In: *Proc. Linear Accel. Conf.* 1998, pp. 627–629.
- [49] *NELSON products*. <https://www.nelsoncreated.com/services>. Accessed: 24/07/21.
- [50] *CPI magnetrons*. <https://www.cpii.com/product.cfm/8/2/372>. Accessed: 24/07/21.
- [51] Hideyuki Obata, Kuniyoshi Furumoto, Hiroyuki Miyamoto, et al. “State of the art advanced magnetrons for accelerator RF power source”. In: *Proc. LINAC*. 2016, pp. 405–407.
- [52] Yuichiro Kamino et al. “Development of an ultrasmall-band linear accelerator guide for a four-dimensional image-guided radiotherapy system with a gimbaled x-ray head”. In: *Medical physics* 34.5 (2007), pp. 1797–1808.
- [53] Zhenxing Tang, Yuanji Pei, and Jian Pang. “Optimal design of a standing-wave accelerating tube with a high shunt impedance based on a genetic algorithm”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 790 (2015), pp. 19–27.
- [54] M El-Ashmawy et al. “Overall quality comparison of c-band and x-band medical linacs”. In: *The 14th symposium on Accelerator Science and Technology, Tsukuba, Japan*. 2003.
- [55] Merrill I Skolnik. “Introduction to radar”. In: *Radar handbook 2* (1962), p. 21.
- [56] *Mevex Linear Accelerators (Linacs)*. <https://mevex.com/mevex-ebeam-xray-linacs/>. Accessed: 24/07/21.
- [57] Sergey V Kutsaev et al. “Electron linac with deep energy control for Adaptive Rail Cargo Inspection System”. In: *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. IEEE. 2015, pp. 1–7.
- [58] Anatoli Arodzero et al. “MIXI: mobile intelligent X-ray inspection system”. In: *IEEE Transactions on Nuclear Science* 64.7 (2016), pp. 1629–1634.

- [59] *Z-PORTAL® for trucks*. <https://www.rapiscansystems.com/en/products/ase-z-portal-for-trucks>. Accessed:24/07/21.
- [60] *NUCTECH MT series*. <http://www.nuctech.com/en/SitePages/ThDetailPage.aspx?nk=PAS&k=DGAGCF>. Accessed: 24/07/21.
- [61] Anatoli Arodzero et al. “High speed, low dose, intelligent X-ray cargo inspection”. In: *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. IEEE. 2015, pp. 1–7.
- [62] Huaibi Chen, Yaohong Liu, and Tang. xuh. “Electron Linacs for cargo inspection and other industrial applications”. In: vol. 10. 44. 2009, p. 11. URL: <https://api.semanticscholar.org/CorpusID:201818778>.
- [63] Thomas P Wangler. *RF Linear accelerators*. John Wiley & Sons, 2008.
- [64] Rob Appleby et al. *The science and technology of particle accelerators*. Taylor & Francis, 2020.
- [65] Sergey Belomestnykh and Valery Shemelin. “High-beta cavity design—a tutorial”. In: *SRF2005 Workshop. Ithaca, NY, USA: Cornell University*. 2005.
- [66] Alexej Grudiev, S Calatroni, and W Wuensch. “New local field quantity describing the high gradient limit of accelerating structures”. In: *Physical Review Special Topics-Accelerators and Beams* 12.10 (2009), p. 102001.
- [67] Kyrre Ness Sjobak, Alexej Grudiev, and Erik Adli. “New Criterion for Shape Optimization of Normal-Conducting Accelerator Cells for High-Gradient Applications”. In: *27th International Linear Accelerator Conference*. 2014, MOPP028.
- [68] Walter Wuensch. “High-gradient breakdown in normal-conducting RF cavities”. In: *EPAC02 (2002)*, pp. 134–138.
- [69] Eugenio Senes et al. “Results of the Beam-Loading Breakdown Rate Experiment at the CLIC Test Facility CTF3”. In: *8th Int. Particle Accelerator Conf.(IPAC’17), Copenhagen, Denmark, 14–19 May, 2017*. JACOW, Geneva, Switzerland. 2017, pp. 1348–1351.
- [70] A Degiovanni et al. “High gradient RF test results of S-band and C-band cavities for medical linear accelerators”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 890 (2018), pp. 1–7.
- [71] S. Verdú Andrés. “High-gradient accelerating structure studies and their application in hadrontherapy”. In: 2013. URL: <https://api.semanticscholar.org/CorpusID:54138092>.
- [72] David Peace Pritzkau. “RF pulsed heating”. PhD thesis. Stanford University, 2001.

-
- [73] David P Pritzkau and Robert H Siemann. “Experimental study of rf pulsed heating on oxygen free electronic copper”. In: *Physical Review Special Topics-Accelerators and Beams* 5.11 (2002), p. 112002.
- [74] Lisa Laurent et al. “Experimental study of rf pulsed heating”. In: *Physical review special topics-accelerators and beams* 14.4 (2011), p. 041001.
- [75] Eite Tiesinga et al. “CODATA recommended values of the fundamental physical constants: 2018”. In: *Rev. Mod. Phys.* 93 (2 2021), p. 025010. DOI: 10.1103/RevModPhys.93.025010. URL: <https://link.aps.org/doi/10.1103/RevModPhys.93.025010>.
- [76] VA Dolgashev and SG Tantawi. “Study of basic breakdown phenomena in high gradient vacuum structures”. In: *Invited Talk at 25th International Linear Accelerator Conference, LINAC10, FR105 Tsukuba, Japan*. Vol. 2. 2010.
- [77] Jim Norem, Zeke Insepov, and I Konkashbaev. “Triggers for RF breakdown”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 537.3 (2005), pp. 510–520.
- [78] Valery A Dolgashev. “Design Criteria for High-Gradient Radio-Frequency Linacs”. In: *Applied Sciences* 13.19 (2023), p. 10849.
- [79] David M Pozar. *Microwave engineering: theory and techniques*. John wiley & sons, 2021.
- [80] Gregory A Loew and Richard Talman. “Elementary principles of linear accelerators”. In: *AIP Conference Proceedings*. Vol. 105. 1. American Institute of Physics. 1983, pp. 1–91.
- [81] Sam Pitman. *Optimisation studies for a high gradient proton Linac for application in proton imaging: ProBE: Proton Boosting Linac for imaging and therapy*. Lancaster University (United Kingdom), 2019.
- [82] S.V. Kutsaev et al. “Design of hybrid electron linac with standing wave buncher and traveling wave structure”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 636.1 (2011), pp. 13–30. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2011.01.047>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900211001252>.
- [83] I. R. R. Shinton, R. M. Jones, and N. Juntong. “SRF Cavity Geometry Optimization for the ILC with Minimized Surface E.M. Fields and Superior Bandwidth”. In: *Proc. PAC’09* (Vancouver, Canada, May 2009). Particle Accelerator Conference 23. JACoW Publishing, Geneva, Switzerland, 2009, pp. 3300–3302. URL: <https://jacow.org/PAC2009/papers/TH5PFP045.pdf>.

- [84] Roger M Jones. “Linear Accelerators: Theory and Practical Applications: WEEK 3”. In: *Cockcroft Institute Lecture Programme 12* (2007).
- [85] Sami Tantawi et al. “Design and demonstration of a distributed-coupling linear accelerator structure”. In: *Physical Review Accelerators and Beams* 23.9 (2020), p. 092001.
- [86] MH Nasr and SG Tantawi. *New Geometrical-Optimization Approach using Splines for Enhanced Accelerator Cavities’ Performance*. Tech. rep. SLAC National Accelerator Lab., Menlo Park, CA (United States), 2018.
- [87] S. Smith et al. “Multiobjective optimization and Pareto front visualization techniques applied to normal conducting rf accelerating structures”. In: *Phys. Rev. Accel. Beams* 25 (6 2022), p. 062002. DOI: 10.1103/PhysRevAccelBeams.25.062002. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.25.062002>.
- [88] Dassault Systèmes. *CST STUDIO SUITE*. Version R2024x. Jan. 16, 2025. URL: <https://www.3ds.com/products-services/simulia/products/cst-studio-suite/>.
- [89] Dassault Systèmes. *ISIGHT & THE SIMULIA EXECUTION ENGINE*. Version 2021x. Apr. 25, 2021. URL: <https://www.3ds.com/products-services/simulia/products/isight-simulia-execution-engine/>.
- [90] K. Deb et al. “A fast and elitist multiobjective genetic algorithm: NSGA-II”. In: *IEEE Transactions on Evolutionary Computation* 6.2 (2002), pp. 182–197. DOI: 10.1109/4235.996017.
- [91] Santosh Tiwari et al. “AMGA: an archive-based micro genetic algorithm for multi-objective optimization”. In: *Proceedings of the 10th annual conference on Genetic and evolutionary computation*. 2008, pp. 729–736.
- [92] D. Alesini. “Linac”. In: *CAS - CERN Accelerator School 2019: Introduction to Accelerator Physics*. Mar. 2021. arXiv: 2103.16500 [physics.acc-ph].
- [93] Sadiq Setiniyaz et al. “Pushing the capture limit of thermionic gun linacs”. In: *Physical Review Accelerators and Beams* 24.8 (2021), p. 080401.
- [94] Sergey V Kutsaev. “Electron bunchers for industrial RF linear accelerators: theory and design guide”. In: *The European Physical Journal Plus* 136.4 (2021), pp. 1–73.
- [95] SG Anderson et al. “Velocity bunching of high-brightness electron beams”. In: *Physical review special topics-accelerators and beams* 8.1 (2005), p. 014401.
- [96] Toshiteru Kii et al. “Experiment and analysis on back-bombardment effect in thermionic RF gun”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 475.1-3 (2001), pp. 588–592.

-
- [97] Klaus Flottmann, Steven Lidia, and Philippe Piot. *Recent improvements to the ASTRA particle tracking code*. Tech. rep. Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), 2003.
- [98] T Khabiboulline et al. “A new tuning method for traveling wave structures”. In: *Proceedings Particle Accelerator Conference*. Vol. 3. IEEE. 1995, pp. 1666–1668.
- [99] Stephen James Vanderet. “Modeling a Standing Wave Linear Accelerator’s Dispersion Relationship and Field Profile via an Equivalent Circuit and Numerical Analysis”. MA thesis. Indiana University, 2011.
- [100] J Sekutowicz, S Kulinski, and M Pachan. “Frequency-Dependent Capacitive-Inductive Model for Axially Coupled”. In: (1985). URL: <https://cds.cern.ch/record/897558>.
- [101] Dario Laneve et al. “Electromagnetic design of microwave cavities for side-coupled linear accelerators: A hybrid numerical/analytical approach”. In: *IEEE Transactions on Nuclear Science* 65.8 (2018), pp. 2233–2239.
- [102] DE Nagle, EA Knapp, and BC Knapp. “Coupled resonator model for standing wave accelerator tanks”. In: *Review of Scientific Instruments* 38.11 (1967), pp. 1583–1587.
- [103] Eugeniusz Plawski, Stanislaw Kulinski, and Marcin Wojciechowski. “Five Cell Method of Tuning Biperiodic Linear Standing Wave $\pi/2$ Accelerating Structures”. In: *Particle Accelerator Conference (PAC 09)*. 2010, TH5PFP092.
- [104] David Alesini. “Power coupling”. In: *arXiv preprint arXiv:1112.3201* (2011).
- [105] ERVIN B Podgorsak et al. “Treatment machines for external beam radiotherapy”. In: *IAEA Radiation Oncology Physics: A Handbook for Teachers And Students International Atomic Energy Agency, Vienna* (2005).
- [106] Yu A Trunev et al. “Heating of tungsten target by intense pulse electron beam”. In: *AIP Conference Proceedings*. Vol. 1771. 1. AIP Publishing LLC. 2016, p. 060016.
- [107] Martin J Berger and Stephen M Seltzer. *Stopping powers and ranges of electrons and positrons*. Tech. rep. US. Nat. Bureau Stand. Inst. Basic Stand., 1982.
- [108] J Anthony Seibert. “X-ray imaging physics for nuclear medicine technologists. Part 1: Basic principles of x-ray production”. In: *Journal of nuclear medicine technology* 32.3 (2004), pp. 139–147.
- [109] Rita Joana da Cruz Roque. “X-Ray Imaging Using 100 μm Thick Gas Electron Multipliers Operating in Kr-CO₂ Mixtures”. PhD thesis. Universidade de Coimbra (Portugal), 2019.
- [110] Thomas J Roberts and Daniel M Kaplan. “G4beamline simulation program for matter-dominated beamlines”. In: *2007 IEEE Particle Accelerator Conference (PAC)*. IEEE. 2007, pp. 3468–3470.

- [111] Ansys®. *Academic Research Workbench, Release 21.2*. <https://www.ansys.com/en-gb/products>. Computer code. 2021.
- [112] Jinghui Wang et al. “Thermal limits on MV x-ray production by bremsstrahlung targets in the context of novel linear accelerators”. In: *Medical physics* 44.12 (2017), pp. 6610–6620.
- [113] Yunus A Cengel. *Heat and mass transfer*. 2000.
- [114] Federico Carra et al. “Assessment of thermal loads in the CERN SPS crab cavities cryomodule1”. In: *Journal of Physics: Conference Series*. Vol. 874. 1. IOP Publishing. 2017, p. 012005.
- [115] Vittorio Parma. “Cryostat design”. In: *arXiv preprint arXiv:1501.07154* (2015).
- [116] Ian H Wilson. “Cavity construction techniques”. In: (1992).
- [117] S Atieh et al. “Machining and Characterizing X-Band RF-Structures for CLIC”. In: *CERN-ATS-2011-238, CLIC-Note-916* (2011). URL: <https://cds.cern.ch/record/1404999>.
- [118] VA Dolgashev et al. “Materials and technological processes for High-Gradient accelerating structures: new results from mechanical tests of an innovative braze-free cavity.” In: *Journal of Instrumentation* 15.01 (2020), P01029.
- [119] SR Ghodke et al. “Machining and brazing of accelerating RF cavity”. In: *2014 International Symposium on Discharges and Electrical Insulation in Vacuum (ISDEIV)*. IEEE. 2014, pp. 101–104.
- [120] Sauza Bedolla Joel. personal communication. Mar. 31, 2021.
- [121] Harold J Frost and Michael F Ashby. *Deformation mechanism maps: the plasticity and creep of metals and ceramics*. Pergamon press, 1982.
- [122] RW Balluffi. “Grain boundary diffusion mechanisms in metals”. In: *Metalurgical transactions B* 13.4 (1982), pp. 527–553.
- [123] David M Owen and Terence G Langdon. “Low stress creep behavior: An examination of Nabarro—Herring and Harper—Dorn creep”. In: *Materials Science and Engineering: A* 216.1-2 (1996), pp. 20–29.
- [124] F. Rossi. *Investigation on the thermo-mechanical deformations related to the diffusion bonding process*. Accessed: 2021-03-20. 2012.
- [125] Daniel B Butrymowicz, John R Manning, and Michael E Read. “Diffusion in Copper and Copper Alloys, Part II. Copper-Silver and Copper-Gold Systems”. In: *Journal of Physical and Chemical Reference Data* 3.2 (1974), pp. 527–602.
- [126] Hans Magnusson and Karin Frisk. *Self-diffusion and impurity diffusion of hydrogen, oxygen, sulphur and phosphorus in copper*. Svensk Kärnbränslehantering AB, Swedish Nuclear Fuel and Waste Management, 2013.

-
- [127] ANSYS. *Ansys Mechanical Finite Element Analysis (FEA) Software for Structural Engineering*. Version R2020. Apr. 25, 2021. URL: <https://www.ansys.com/en-gb>.
- [128] Antti Moilanen. “Creep effects in diffusion bonding of oxygen-free copper”. PhD thesis. Aalto U., 2013.
- [129] Clair Upthegrove and Henry Lewis Burghoff. “Elevated-temperature properties of coppers and copper-base alloys”. In: *ASTM PAPERS* (1956).
- [130] Dassault Systèmes. *CST MPhysics® Studio*. Version R2021x. Apr. 25, 2021. URL: <https://www.3ds.com/products/simulia/cst-studio-suite/electromagnetic-simulation-solvers>.
- [131] Sungsu Cha et al. “Development of an automatic frequency control system for an X-band (= 9300 MHz) RF electron linear accelerator”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 855 (2017), pp. 102–108.
- [132] Viktor Mileikovskiy and Tetiana Tkachenko. “Precise explicit approximations of the Colebrook-White equation for engineering systems”. In: *International Scientific Conference EcoComfort and Current Issues of Civil Engineering*. Springer. 2020, pp. 303–310.
- [133] M Mohseni Kejani et al. “Multiphysics analysis of side-coupled RF cavity”. In: *Journal of Instrumentation* 14.07 (2019), P07001.
- [134] SV Kutsaev et al. “Design and multiphysics analysis of a 176 MHz continuous-wave radio-frequency quadrupole”. In: *Physical Review Special Topics-Accelerators and Beams* 17.7 (2014), p. 072001.
- [135] V Paramonov, B Militsyn, and A Skassyrskaya. “Procedure of simulations for Evaluations of Transient Parameters Stability for RF Cavities Due to Effects of Pulsed RF Heating”. In: *arXiv preprint arXiv:1812.10317* (2018).
- [136] Liu Yang, Xiaozhong He, and Hong Li. “Electromagnetic, Thermal, Structural Analysis for the RF-cavity of a Rhodotron Accelerator”. In: *13th Symposium on Accelerator Physics*. 2018, TUPH27. DOI: 10.18429/JACoW-SAP2017-TUPH27.
- [137] Jean Clifford Brutus et al. “Mechanical Design and 3-D Coupled RF, Thermal-Structural Analysis of Normal Conducting 704 MHz and 2.1 GHz Cavities for LEReC Linac”. In: *Proc. 7th Int. Particle Accelerator Conf.(IPAC’16)*. 2016, pp. 525–527.
- [138] V. Pettinacci et al. “Thermal-mechanical Analysis of the RF Structures for the ELI-NP Proposal”. In: *Proc. 5th International Particle Accelerator Conference (IPAC’14), Dresden, Germany, June 15-20, 2014*. 5. JACoW, 2014, pp. 3860–3862. URL: <http://jacow.org/ipac2014/papers/thpri043.pdf>.

- [139] N Hartman and RA Rimmer. “Electromagnetic, thermal, and structural analysis of RF cavities using ANSYS”. In: *PACS2001. Proceedings of the 2001 Particle Accelerator Conference (Cat. No. 01CH37268)*. Vol. 2. IEEE. 2001, pp. 912–914.
- [140] Achuth Rao. “Fluid—solid interaction analysis using ANSYS/multiphysics”. In: *Computational Fluid and Solid Mechanics 2003*. Elsevier, 2003, pp. 1492–1496.
- [141] Daniela Capatina et al. “Photon Beam Applied as Heat Flux on Irregular Surfaces in FEA”. In: *Mechanical Eng. Design of Synchrotron Radiation Equipment and Instrumentation (MEDSI'18), Paris, France, 25-29 June 2018*. JACOW Publishing, Geneva, Switzerland. 2018, pp. 214–216.
- [142] Cheuk Wing Edmond Lam et al. *Thermal control for the two-beam module of the compact linear collider with computational fluid dynamics simulations of conjugate heat transfer problems, with subsidiary finite-element simulations*. Tech. rep. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2642576>.
- [143] Florianr Menter. “Zonal two equation kw turbulence models for aerodynamic flows”. In: *23rd fluid dynamics, plasmadynamics, and lasers conference*. 1993, p. 2906.
- [144] Florian R Menter, Martin Kuntz, Robin Langtry, et al. “Ten years of industrial experience with the SST turbulence model”. In: *Turbulence, heat and mass transfer 4.1 (2003)*, pp. 625–632.
- [145] Heinrich Blasius. “Das aehnlichkeitsgesetz bei reibungsvorgängen in flüssigkeiten”. In: *Mitteilungen über Forschungsarbeiten auf dem Gebiete des Ingenieurwesens: insbesondere aus den Laboratorien der technischen Hochschulen*. Springer, 1913, pp. 1–41.
- [146] Adam Neale et al. “Determination of surface convective heat transfer coefficients by CFD”. In: *11th Canadian Conference on Building Science and Technology Banff, Alberta*. 2007.
- [147] Piotr Duda. “Heat Transfer Coefficient Distribution—A Review of Calculation Methods”. In: *Energies* 16.9 (2023), p. 3683.
- [148] Yasuo Mori and Wataru Nakayama. “Study of forced convective heat transfer in curved pipes (2nd report, turbulent region)”. In: *International journal of heat and mass transfer* 10.1 (1967), pp. 37–59.
- [149] Yasuo Mori and Wataru Nakayama. “Study on forced convective heat transfer in curved pipes:(3rd report, theoretical analysis under the condition of uniform wall temperature and practical formulae)”. In: *International journal of heat and mass transfer* 10.5 (1967), pp. 681–695.

-
- [150] XC Li, J Zhou, and K Aung. “On selection of reference temperature of heat transfer coefficient for complicated flows”. In: *Heat Transfer Summer Conference*. Vol. 42762. 2007, pp. 497–505.
- [151] Kai Papke, Carlo Rossi, and Graeme Campbell Burt. *Coupled RF-Thermo-Structural Analysis of CLIC Traveling Wave Accelerating Structures*. Tech. rep. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2730555>.
- [152] Sami G Tantawi. “Multimoded reflective delay lines and their application to resonant delay line rf pulse compression systems”. In: *Physical review special topics-accelerators and beams* 7.3 (2004), p. 032001.
- [153] SG Tantawi, NM Kroll, and K Fant. “RF components using over-moded rectangular waveguides for the next linear collider multi-moded delay line RF distribution system”. In: *Proceedings of the 1999 Particle Accelerator Conference (Cat. No. 99CH36366)*. Vol. 2. IEEE. 1999, pp. 1435–1437.
- [154] SG Tantawi et al. “Multimoded rf delay line distribution system for the Next Linear Collider”. In: *physical review special topics-accelerators and beams* 5.3 (2002), p. 032001.
- [155] RE Willoughby. “High-Power, SPDT, Fast Ferrite Switch”. In: *Journal of Applied Physics* 36.3 (1965), pp. 1247–1248.
- [156] Microwave Applications Group. *C-Band Rotary-Field Phase Shifters*. [Online; accessed 12-December-2023]. 2023. URL: <https://www.magsmx.com/index.html>.
- [157] Constantine A Balanis. *Advanced engineering electromagnetics*. John Wiley & Sons, 2012.
- [158] Stephen Weinstein and Paul Ebert. “Data transmission by frequency-division multiplexing using the discrete Fourier transform”. In: *IEEE transactions on Communication Technology* 19.5 (1971), pp. 628–634.
- [159] Owen E DeLange. “Wide-band optical communication systems: Part II—Frequency-division multiplexing”. In: *Proceedings of the IEEE* 58.10 (1970), pp. 1683–1690.
- [160] Gary Breed. “The mathematics of mixers: basic principles”. In: *High Frequency Electronics* (2011). URL: https://highfrequencyelectronics.com/Jan11/HFE0111_Tutorial.pdf.
- [161] Yaogen Ding et al. “S-band multibeam klystron with bandwidth of 10%”. In: *IEEE transactions on electron devices* 52.5 (2005), pp. 889–894.
- [162] Yaogen Ding et al. “An Overview of Multibeam Klystron Technology”. In: *IEEE Transactions on Electron Devices* (2023).
- [163] Yaogen Ding et al. “Research progress on C-band broadband multibeam klystron”. In: *IEEE transactions on electron devices* 54.4 (2007), pp. 624–631.

- [164] AS Gilmour. *Klystrons, traveling wave tubes, magnetrons, crossed-field amplifiers, and gyrotrons*. Artech House, 2011.
- [165] Sun-Hong Min et al. “Low-level RF control of a klystron for medical linear accelerator applications”. In: *AIP Advances* 9.2 (2019).
- [166] A H Regan, A S Rohlev, and C D Ziomek. “APT LLRF Control System Functionality and Architecture”. In: (1996). DOI: 10.5170/CERN-1996-007.225. URL: <https://cds.cern.ch/record/399795>.
- [167] Nelson. *Nelson Products*. Ed. by Nelson. URL: <https://www.nelsoncreated.com/services> (visited on 01/03/2024).
- [168] Hamid Reza Dalili Oskouei et al. “Design and construction of ferrite waveguide circulator with short wall in S band”. In: *The Journal of Engineering* 2023.8 (2023), e12291.
- [169] C D Beard et al. “EMMA RF Cavity Design and Prototype Testing at Daresbury”. In: (2008). URL: <https://cds.cern.ch/record/1182568>.
- [170] Matthew Jones. personal communication, Cockcroft Institute, Daresbury, Warrington WA4 4AD, U.K. Jan. 18, 2024.
- [171] Rodney A Brooks and Giovanni Di Chiro. “Theory of image reconstruction in computed tomography”. In: *Radiology* 117.3 (1975), pp. 561–572.
- [172] Godfrey N Hounsfield. “Computerized transverse axial scanning (tomography): Part 1. Description of system”. In: *The British journal of radiology* 46.552 (1973), pp. 1016–1022.
- [173] Robert N Yancey et al. “CT-assisted metrology for manufacturing applications”. In: *Nondestructive Evaluation for Process Control in Manufacturing*. Vol. 2948. SPIE. 1996, pp. 222–231.
- [174] Jason M Warnett et al. “Towards in-process x-ray CT for dimensional metrology”. In: *Measurement Science and Technology* 27.3 (2016), p. 035401.
- [175] Faycal Kharfi. “Mathematics and physics of computed tomography (CT): demonstrations and practical examples”. In: *Imaging and radioanalytical techniques in interdisciplinary research-fundamentals and cutting edge applications* 117 (2013), pp. 113–116.
- [176] Avinash C Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*. SIAM, 2001.
- [177] Frank Natterer. *The mathematics of computerized tomography*. SIAM, 2001.
- [178] Gabor T Herman. *Fundamentals of computerized tomography: image reconstruction from projections*. Springer Science & Business Media, 2009.
- [179] Dan E Dudgeon and Russell M Mersereau. *Multidimensional digital signal processing*. Prentice-hall, 1984.

-
- [180] Kyungsang Kim. *3D Cone beam CT (CBCT) projection backprojection FDK, iterative reconstruction*. <https://www.mathworks.com/matlabcentral/fileexchange/35548-3d-cone-beam-ct-cbct-projection-backprojection-fdk-iterative-reconstruction-matlab-examples>. [Online; MATLAB Central File Exchange, Retrieved January 24, 2024]. 2024.
- [181] Anders H Andersen and Avinash C Kak. “Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm”. In: *Ultrasonic imaging* 6.1 (1984), pp. 81–94.
- [182] David G Politte and Donald L Snyder. “Corrections for accidental coincidences and attenuation in maximum-likelihood image reconstruction for positron-emission tomography”. In: *IEEE Transactions on Medical Imaging* 10.1 (1991), pp. 82–89.
- [183] Johan Nuyts, Christian Michel, and Patrick Dupont. “Maximum-likelihood expectation-maximization reconstruction of sinograms with arbitrary noise distribution using NEC-transformations”. In: *IEEE transactions on medical imaging* 20.5 (2001), pp. 365–375.
- [184] Abdelwahhab Boudjelal et al. “A novel iterative mlem image reconstruction algorithm based on beltrami filter: Application to ect images”. In: *Tomography* 7.3 (2021), pp. 286–300.
- [185] Scott D Roth. “Ray casting for modeling solids”. In: *Computer graphics and image processing* 18.2 (1982), pp. 109–144.
- [186] Adam A. *Mesh voxelisation*. <https://www.mathworks.com/matlabcentral/fileexchange/27390-mesh-voxelisation>. [Online; MATLAB Central File Exchange, Retrieved January 24, 2024]. 2024.
- [187] Sandeep Patil and Bhallamudi Ravi. “Voxel-based representation, display and thickness analysis of intricate shapes”. In: *Ninth International Conference on Computer Aided Design and Computer Graphics (CAD-CG’05)*. IEEE. 2005, 6–pp.
- [188] Michael Behr. *XraySim, Toronto Rehabilitation Institute, University Health Network*. https://github.com/MichaelBehr/Xray_Sim.git. [Online; Retrieved January 24, 2024]. 2024.
- [189] Lee A Feldkamp, Lloyd C Davis, and James W Kress. “Practical cone-beam algorithm”. In: *Josa a* 1.6 (1984), pp. 612–619.
- [190] Oliver Woodford. *vol3d v2*. <https://www.mathworks.com/matlabcentral/fileexchange/22940-vol3d-v2>. [Online; MATLAB Central File Exchange, Retrieved January 24, 2024]. 2024.
- [191] Xiao Han et al. “Optimization-based reconstruction of sparse images from few-view projections”. In: *Physics in Medicine & Biology* 57.16 (2012), p. 5245.

- [192] Jiayi Wang et al. “Deep learning based image reconstruction algorithm for limited-angle translational computed tomography”. In: *PLOS ONE* 15 (Jan. 2020), pp. 1–20. URL: <https://doi.org/10.1371/journal.pone.0226963>.
- [193] Tobias Würfl et al. “Deep learning computed tomography: Learning projection-domain weights from image domain in limited angle problems”. In: *IEEE Transactions on medical imaging* 37.6 (2018), pp. 1454–1463.
- [194] Dimitrios Bellos et al. “A convolutional neural network for fast upsampling of undersampled tomograms in X-ray CT time-series using a representative highly sampled tomogram”. In: *Journal of Synchrotron Radiation* 26.3 (2019), pp. 839–853.
- [195] Chul Kyun Ahn et al. “A deep learning-enabled iterative reconstruction of ultra-low-dose CT: use of synthetic sinogram-based noise simulation technique”. In: *Medical Imaging 2018: Physics of Medical Imaging*. Vol. 10573. SPIE. 2018, pp. 809–814.
- [196] Ao Zheng et al. “A dual-domain deep learning-based reconstruction method for fully 3D sparse data helical CT”. In: *Physics in Medicine & Biology* 65.24 (2020), p. 245030.
- [197] Bo Zhu et al. “Image reconstruction by domain-transform manifold learning”. In: *Nature* 555.7697 (2018), pp. 487–492.
- [198] Thomas W. Rogers et al. “Threat Image Projection (TIP) into X-ray images of cargo containers for training humans and machines”. In: *2016 IEEE International Carnahan Conference on Security Technology (ICCST)*. 2016, pp. 1–7. DOI: 10.1109/CCST.2016.7815717.
- [199] SV Kutsaev et al. “Compact X-Band electron linac for radiotherapy and security applications”. In: *Radiation Physics and Chemistry* 185 (2021), p. 109494.
- [200] A Grudiev and W Wuensch. *Design of an X-band accelerating structure for the CLIC main linac*. Tech. rep. 2008.