

Pulsed Electromagnetic Field Generation using Timed Array Structures

A novel approach to radiated NEMP
hardness validation testing for
large naval platforms

N.H.W. Elshove
24-06-2026

Colophon

Pulsed Electromagnetic Field Generation using Timed Array Structures

Author:	N.H.W. Elshove
Student number:	6332099
Institution:	TU Delft
Programme:	MSc Electrical Engineering
Faculty / Department:	Faculteit Elektrotechniek, Wiskunde en Informatica Microwave Sensing, Signals and Systems (MS3)
Academic year:	2024 - 2026
Thesis conducted at:	Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek (TNO)
Unit:	Defence, Safety and Security
Department:	Electronic Warfare
Company location:	The Hague, Netherlands
Daily supervisor 1:	dr. ir. A.P.M. Zwamborn
Daily supervisor 2:	dr. ir. M.A. Echeverri Bautista
University supervisor:	Prof. DSc. A. Yarovoy
Version:	FINAL
Confidentiality:	Public
Thesis defended on:	24 June, 2026, at 15:00
Committee Chair:	Prof. DSc. A. Yarovoy
Committee Member	dr. ir. R.F. Remis
External Advisor	dr. ir. A.P.M. Zwamborn

In high power time domain electromagnetics, there is no such thing as a free lunch
Peter Zwamborn, 2025

Abstract

Bounded-wave NEMP simulators cannot accommodate a full-scale warship and the conventional radiating alternative, a large monocone, cannot deliver sufficient field strength at acceptable homogeneity over the testing surface. This thesis investigates an unconventional architecture for radiated NEMP hardness validation of large naval platforms: a timed array of directive, ultra-wideband radiators that deliberately trades full-platform illumination for local higher field strength and better homogeneity.

Starting from the NEMP standards, a set of field-strength, homogeneity and timing requirements is distilled, which may be interpreted as design specifications. The monocone baseline (EMIS-III-VPD) is characterised and shown to meet none of the three criteria, achieving only 1.28 kV/m peak field, roughly 10 dB amplitude inhomogeneity and 138 ns of peak-time spread across a testing surface of 200 m x 50 m at 100 m distance.

The design process is started by evaluating candidate antenna concepts that are all at least ultra-wideband and non-dispersive. It is revealed that directivity is the enabling factor for antennas in radiated NEMP testing, resulting in the reflector IRA to be selected as element in the array. Subsequently, this antenna is modelled at three scales in a full wave-solver, exposing the trade-off between aperture size and low-frequency operating limit. Using results from the full-wave solver, the array is synthesized and optimized in MATLAB using a multi-objective genetic algorithm. This was done over a far-field model anchored to the outcome from the full-wave solver, reducing runtime significantly. With peak field strength imposed as a minimum, the amplitude and timing homogeneity could be minimised simultaneously.

A single 12.2 m IRA fed at 500 kV radiates 4.5 kV/m. An optimised four-element array reaches 9.6 kV/m peak at 0.56 dB inhomogeneity and 1.06 ns peak-time spread over a 10×10 m surface at 100 m distance. The array is therefore a peak-field multiplier rather than a means of improving homogeneity: the single antenna is the most uniform configuration of any studied, and optimisation only recovers near-single-antenna uniformity up to a ceiling of roughly 9 kV/m, beyond which it degrades steeply. The 50 kV/m of the standards is unreachable by this approach, since the feasible element counts and feed voltages required are impractical and the binding limits, aperture size and voltage breakdown, lie outside the array. This contribution is not a finished simulator but a quantified baseline, and bounds on what a radiated NEMP test of a large platform can realistically achieve, as well as a framework within which future work may proceed from.

Preface

The project you hold in your hands (or more likely: on your screen) is my thesis to obtain the master's diploma in electrical engineering from the TU Delft. This project spans about 8 months, having started at the beginning of November 2025 until the defence at the end of June 2026, and is worth 45 ECTS. In making this project I prove I am worthy of obtaining my masters degree, as with the successful completion of the project I may call myself an electrical engineer.

The project has been conducted at TNO, and was therefore combined with an internship at TNO. TNO is a non-profit research organisation that exists by Dutch law (although not technically a public company) and it bridges the gap between academic research and practically implementable technology. This has the goal to improve the market position of the Netherlands, to improve safety and security and to sustainably enhance the well-being of society.

This project is part of a larger project at TNO to improve the knowledge on NEMP in general, and NEMP hardness validation in particular. The project is commissioned by the Dutch Ministry of Defence, who have renewed interest in the topic. This is what made it appealing to me: there is direct applicability and need for the existence of this research work.

Here, I want to take the time to thank Peter, who was my daily supervisor at TNO. It was also based upon his idea that this thesis was written. His enthusiasm for electrical engineering in general, and NEMP in particular, has also kept me motivated throughout this project. His wisdom and clarity has been invaluable to keep this work grounded and continuously aimed in the right direction. Thank you Peter!

Next, I would also like to take the time to thank Mario, who after being asked just one question on NEMP over a cup of coffee decided that despite his busy schedule, he also wanted to guide a master's student. His knowledge on computational EM-solvers and the meticulousness with which he could read my work has been of great value to the project.

Finally, I want to thank Professor Yarovoy. He was the university supervisor with whom I met monthly. Despite (or perhaps thanks to) meeting just monthly, his comments have often provided an interesting, new perspective. On numerous occasions this has led to significant improvement on the quality of the work, specifically in the representation of results obtained throughout.

I hope this work may prove useful to you.

Niels Elshove
June 25, 2026

Contents

Abstract	i
Preface	ii
1 Introduction	1
1.1 The Nuclear Electromagnetic Pulse	1
1.2 NEMP Hardness Validation Testing	2
1.3 Research Questions	4
1.4 Structure of the thesis	4
2 Desired Waveform	6
2.1 Waveform as defined by NEMP standards	6
2.1.1 Time domain	6
2.1.2 Frequency domain	8
2.1.3 Frequency limitation and time domain reconstruction	10
2.2 Other Waveform Parameters	12
2.3 Testing Surface Requirements	12
2.4 Conclusion	14
3 Baseline	15
3.1 The EMIS-III-VPD	15
3.2 First-order estimate	16
3.3 Simulation	19
3.3.1 Far-field considerations	19
3.3.2 Cone Impedance	20
3.3.3 End-termination model	20
3.3.4 Resistively loaded wire model	24
3.3.5 Solid cone model	28
3.3.6 Comparison and Observations	31
3.4 Performance Estimation	32
3.5 Simple two-element moncone array	38
3.6 Challenges of Previous NEMP Hardness Validation Tests	42
3.7 Discussion	43
3.8 Conclusion	45
4 Antenna Selection	46
4.1 Transmission-line model for travelling-wave radiators	46
4.2 Biconical Antenna	48
4.2.1 Antenna principle	48
4.2.2 Design considerations	48
4.2.3 First order estimations	49
4.2.4 Application of the transmission-line model	49
4.3 TEM Horn Antenna	50
4.3.1 Antenna principle	51
4.3.2 Design considerations	51
4.3.3 First order estimations	52

4.3.4	Application of the transmission-line model	53
4.4	Vivaldi / Tapered Slot Antenna	54
4.4.1	Working principle	55
4.4.2	Design considerations	55
4.5	Reflector IRA	56
4.5.1	Working principle	56
4.5.2	Design considerations	57
4.5.3	First order estimations	58
4.5.4	Application of the transmission-line model	59
4.6	Lens IRA	60
4.6.1	Working principle	60
4.6.2	Design considerations	61
4.7	Comparison	62
4.8	Discussion	64
4.9	Conclusion	64
5	Antenna Design	65
5.1	Design Parameter Selection	66
5.1.1	Number of Feed Arms	66
5.1.2	Feed Arm Angular Placement	67
5.1.3	F/D Ratio	67
5.1.4	Summary of High-Level Parameters	69
5.1.5	Reference Design	69
5.1.6	Antenna Dish Size	70
5.2	A note on the antenna far field	71
5.2.1	A time-domain criterion	71
5.2.2	Implications for the test distance	72
5.2.3	Implications for this thesis	74
5.3	Simulations	75
5.3.1	Antenna Geometry	75
5.3.2	Feeding Pulse	77
5.3.3	30 m model	78
5.3.4	12.2 m model	81
5.3.5	5 m model	83
5.3.6	Comparison and Observations	86
5.4	Performance Estimation	87
5.4.1	Frequency domain performance estimation	87
5.4.2	Peak field amplitude estimation	88
5.4.3	Time-domain performance estimation	89
5.4.4	Testing surface performance estimate	91
5.4.5	Near-field analysis of the 30 m model	94
5.5	Discussion	97
5.5.1	Antenna performance	97
5.5.2	Near-field considerations	98
5.5.3	Reflection on methodology	98
5.5.4	Antenna practicality	99
5.6	Conclusion	100
6	Array Design	101
6.1	Single Antenna Model	102
6.1.1	Single antenna impulse response	102
6.1.2	Idealized feeding pulse	104
6.1.3	Single antenna weighted impulse response	108
6.2	Manual array synthesis	109
6.2.1	4 antenna linear array	110
6.2.2	Observations	114
6.3	Array Optimization	116

6.3.1	Optimizing for two competing objectives	116
6.3.2	MATLAB Genetic-algorithm solver	117
6.3.3	Optimised array	118
6.3.4	The price of peak field strength	122
6.3.5	Feed-network and mechanical tolerances	123
6.3.6	Limitations and validity of the optimisation	126
6.4	Discussion	126
6.5	Conclusion	128
7	Discussion	129
7.1	Context of this thesis	129
7.2	Synthesis of central findings	130
7.2.1	What has been demonstrated	130
7.2.2	Assumptions, heuristics and their consequences	131
7.2.3	Practical radiated NEMP testing	132
7.2.4	Central findings	133
7.2.5	Novelty and contributions of this work	134
7.3	Future Work	135
8	Conclusion	140
	Bibliography	142

Chapter 1

Introduction

1.1 The Nuclear Electromagnetic Pulse

An exoatmospheric nuclear detonation causes an extreme electromagnetic (EM) environment referred to as a Nuclear Electromagnetic Pulse (NEMP). A component of this environment is extremely wideband, in some literature jestingly referred to as "DC-to-daylight," and of very high intensity. Furthermore, the NEMP illuminates anything in line-of-sight, which could be in the scale of entire continents. The wideband nature in combination with the high intensity makes NEMP a danger to all objects containing delicate electrical equipment, as the NEMP may induce currents and voltages on any electrically conducting surface much higher than safe operating limits [1]. This may cause large scale destruction or upset of electrical equipment if measures are not taken to protect against it.

A NEMP event contains three components: the early-time (E1), intermediate-time (E2) and late-time (E3) components. All three components of NEMP are caused by different mechanisms and are different in nature. The late-time component is similar in nature to a geomagnetic storm, lasting up to minutes. The long duration of this component means the pulse only contains the very lowest frequencies, meaning this is mainly a threat to extremely long objects, such as national power grids. To smaller objects this component is relatively harmless as the object size is significantly shorter than the shortest wavelength. The E2 component lasts microseconds and is similar in nature to lightning electromagnetic pulse (LEMP). This component is therefore harmless when proper measures to protect against LEMP have been taken. It is the E1 component, which lasts nanoseconds and hence contains an ultra wide bandwidth (UWB), that is most threatening to smaller objects [1]. This is the component of NEMP that is the focus of this thesis. When referring to NEMP in the remainder of this thesis, the E1 component is meant.

A NEMP is caused by the gamma rays that are emitted by a nuclear detonation. When this detonation occurs outside of the atmosphere, the gamma rays can travel indefinitely. A portion of this expanding shell of photons will strike the atmosphere, which is of increasing density closer to the Earth's surface. As the density increases, the chance of a collision between a photon and an electron orbiting an air particle increases. Such a collision knocks the electron out of its orbit and releases it to travel in roughly the direction the photon was travelling in, a process referred to as Compton Scattering [1]. A region of air at approximately 30 ± 6 kilometres altitude, within line-of-sight of the burst, is the region of air in which most of the photons are absorbed by the scattered electrons. This region is referred to as the source region of NEMP [2]. Under influence of the Earth's magnetic field these electrons will coherently be bent off their original path. This constitutes a varying electric current, which causes electromagnetic radiation. A NEMP is significantly more potent when caused by an exoatmospheric detonation because the electrons are generated in a less dense layer in the atmosphere, which allows them to travel further before slowing down and being reabsorbed by air particles, and hence generate an more intense pulse.

1.2 NEMP Hardness Validation Testing

Military equipment is designed with the threat of NEMP in mind. The ability to withstand a NEMP event, in other words to continue functioning during such an event, is referred to as hardness, which can be seen as the opposite of susceptibility. Note that hardness is not equivalent to immunity, in the sense that a military platform may still be considered NEMP hard despite experiencing significant disturbance and upset during and after a NEMP event. Designers will have specified that mission-critical equipment must survive with minimum upset, despite other, less critical equipment being heavily disturbed or even destroyed. When new military equipment is procured, the NEMP hardness (specifically the damping effect of the outer shells) is set in a requirement by a ministry of defence. A supplier of military equipment will deliver its product claiming it is NEMP hard, but it is generally up to the ministry of defence to validate this hardness.

NEMP hardness validation can occur through several methods which instil varying levels of confidence in users. It is eventually up to the users of the equipment to decide when they are sufficiently confident that their equipment will function during and after an NEMP event. A simple method of NEMP hardness validation is an injected cable test, in which the waveform that would be coupled into a certain cable is calculated and injected into that cable. This requires the calculation of the coupling mechanics into a certain cable with sufficient accuracy, which is not trivial. Furthermore, for complex and large objects such as aircraft or warships this is laborious. An alternative is a radiated test, in which the NEMP pulse is recreated and the device under test (DUT) is illuminated by it. When done properly this test is more accurate than a cable injector test, as it encapsulates all effects of a NEMP event, which cannot all be accounted for or predicted in cable injection tests.

However, the design of a good NEMP validation test is not trivial. Since the discovery of the potential of NEMP in 1962 [3], there have been many NEMP simulator designs. The main requirements for proper NEMP hardness validation is that the field-strength, homogeneity and the temporal fidelity are sufficient. It is currently impossible to achieve all three requirements at the same time, as certain metrics may be locked in a trade-off with other metrics. As such, a choice must be made to prioritise one at the cost of another metric. Balancing these trade-offs is the reason why NEMP validation testing is still a topic of scientific relevance to this day, and a perfect NEMP simulator is yet to be invented. The balancing of such trade-offs will be a recurring theme in this thesis.

There are three types of NEMP tests in which electromagnetic energy is radiated: a bounded wave simulator, a hybrid simulator and an unbounded (in literature referred to as radiating) test [4]. In an unbounded simulator the DUT is placed within line-of-sight of the simulator, characterised by its antenna, and illuminated by it. In a bounded simulator, the defining difference is that the DUT is placed inside the simulator, and the simulator acts as a waveguide that contains a wave propagating in TEM mode. A bounded wave simulator can only generate a vertically polarized wave. To generate a horizontally polarized wave in a bounded wave fashion, a hybrid simulator was designed. This simulator also has the DUT contained inside the simulator, but operates on more complex mechanics that are outside of the scope of this thesis. A bounded-wave simulator and a hybrid simulator are shown below in figure 1.1



(a) A Leopard II placed in a bounded wave NEMP simulator owned by the Bundeswehr [5]



(b) A hybrid NEMP simulator sold by Montena [6]

Figure 1.1: Examples of a bounded wave and hybrid simulators

No perfect simulator exists, but the NEMP simulator concept that best balances the trade-offs is the bounded wave simulator, in which a simulated NEMP pulse is generated at one end, travels in TEM mode along the length of the simulator and is terminated at the other end. The bounded nature of the simulator means that high field strength can be achieved at acceptable levels of field homogeneity and time domain fidelity. Furthermore, almost no electromagnetic energy is radiated on devices not under test, meaning the test can practically be achieved without requiring remote locations.

A problem with this type of NEMP simulator is it quickly becomes impractical for larger military platforms, as this type of simulator must be approximately three times larger than the DUT. This can be seen in figure 1.2, which shows a bounded wave simulator suspended under an existing bridge. This was a proof-of-concept developed with TNO's involvement (along with counterparts in the UK and Norway) in the 90s and it illustrates well that despite the enormous size of the simulator, the maximum size of the DUT is relatively small. A bounded wave simulator might be ideal, but for larger platforms it is unfeasible for practical implementation. For the same reason a hybrid simulator is not feasible.



Figure 1.2: A bounded wave simulator suspended under a bridge [7]

For larger military platforms, such as military shelters, large aircraft or warships, only an unbounded (radiating) NEMP simulator remains feasible. The nature of radiating such a wave in an unbounded manner into free space makes achieving the required field strength, homogeneity and temporal fidelity challenging, more so than with bounded wave simulators. The most common recent method used to validate NEMP hardness was to use a large monocone, effectively half of a biconical antenna placed above a conducting surface. The main problem with this design is the lack of directionality: this antenna is omnidirectional in the horizontal plane. This means homogeneity will only be achieved at very large distances from the antenna, which in turn means the field strength cannot be achieved. In fact, the lack of directionality means the required field strength will not be achieved by such an antenna at any significant distance without exceeding the air breakdown limit of 3 MV/m at the radiator feed point. Hence, the tests of the past did not achieve sufficient field strength, homogeneity or temporal fidelity to instil sufficient confidence in users in the modern day. There is a desire to develop a new NEMP hardness validation method that improves the level of confidence the users will have in their systems. The goal of this research is to improve upon existing methods to validate the NEMP hardness of warships in the Royal Netherlands Navy (RNLN).

1.3 Research Questions

With the two main conventional methods for NEMP hardness validation testing (which are the bounded-wave test and the radiated test using a monocone architecture) exhausted, this thesis explores whether an unconventional method may provide better results. This may be seen as both a feasibility study, after which a substantiated claim can be made as to whether or not this unconventional method is worth further research, and as a framework within which future work may be done.

This thesis explores whether the use of a time-domain antenna array focused on a limited surface area can provide better performance than the monocone antenna used until now. An array will introduce at least some directionality, which should make reaching the field strength and homogeneity requirements more straightforward. Furthermore, splitting the generation of such a pulse across several radiators means the air breakdown limit will not be reached as quickly. It is envisioned that the array will use its improved directionality to illuminate a surface area of limited size, as opposed to illuminating the entire ship. This introduces a fault in the test (there are no perfect simulators) as effects that occur when the entire DUT is illuminated will no longer be addressed. Already, this presents a trade-off: decreasing the illuminated surface area increases the field strength across that surface area, while the homogeneity across the entire DUT is reduced. The research questions that this thesis will address are:

To what extent can a timed array of directive, ultra-wideband radiators serve as a feasible radiated NEMP hardness validation method for large naval platforms?

To answer this question, and to structure this thesis, the following supporting questions are asked:

- What performance does the existing mono-cone baseline achieve, and can full-surface illumination of a large platform meet these criteria?
- Which ultra-wideband, non-dispersive radiator is best suited as an array element, and what trade-off governs its performance?
- How closely can a timed array of these elements satisfy the field-strength, homogeneity and timing criteria simultaneously, and what bounds that performance?

The intended novelty of this project is the use of a timed array of directive elements for the purpose of NEMP hardness validation. Furthermore, it is the first publicly available source detailing the design process from standards to selection of antenna concept to optimization and quantification of the array performance. Besides the findings of this work, it also provides a framework within which future work may be structured.

1.4 Structure of the thesis

The next chapter in this thesis describes the desired waveform. In this chapter the NEMP waveform as described in unclassified standards is highlighted, and from it all radiated waveform requirements for any NEMP simulator are distilled. The conclusion of this chapter is a list of characteristics the field radiated by the array will have at the testing surface.

In chapter 3 the monocone antenna used in previous NEMP validation campaigns for RNLN warships are described. An estimation of their performance is made, as there are limited public sources describing their performance. The estimated performance of the monocone will be used as the baseline of this thesis; any results made in this thesis will be compared to the baseline.

Chapter 4 uses the NEMP waveform as described in chapter 2 to generate a list of potential antenna concepts. The main quality of these antenna concepts is that they are non-dispersive and ultra-wideband (UWB). These antenna concepts are briefly introduced after which they are compared and contrasted based on the requirements set in chapter 2. One antenna concept is selected in this chapter and will eventually serve as the antenna element in the timed array.

In Chapter 5 the antenna concept selected in chapter 4 is modelled in a full-wave electromagnetic solver. Specifically, three models are made that are scaled versions of each other such that the fundamental trade-off with radiated NEMP simulators is highlighted: the choice between antenna practicality (defined by its size) and the low frequency operating limit. The results of this model are brought into Chapter 6, where the array will be designed.

The array is designed in chapter 6 using results obtained from chapter 5. First, a manually synthesized array is presented, as that gives insight into the performance drivers and trade-offs at play in the design of an antenna array for NEMP testing purposes. Then the antenna element layout, aiming points, timing delay and feed amplitude are varied such that the homogeneity is optimized over the testing surface.

From all the results gathered, Chapter 7 synthesizes the key results of this thesis. Here, the main qualitative takeaways are distilled and presented. Furthermore, the main caveats and assumptions that come with these results are mentioned again. Finally, potential future work is presented in that chapter. Chapter 8 contains the final conclusion of the thesis.

Chapter 2

Desired Waveform

Designing a NEMP simulator from scratch starts with defining the desired EM environment that the simulator will reproduce. The primary criterion for assessing the suitability of an antenna concept is how accurately it can reproduce a NEMP-like pulse in free space. Such pulses are defined in standards such as the IEC 61000-2-9 [18] and related specifications that describe the temporal shape and amplitude of the E1 environment. In addition to waveform fidelity, other important criteria include the practical feasibility of the antenna concept, particularly with regard to repeatability of tests and the availability of an existing design that can be further developed or adapted within the scope of this research. In this section the primary criterion, the desired waveform, is described. The description provided by the standards requires further distillation because the standards leave room for interpretation by the designer of a simulator. At the end of this chapter a concise list of requirements for the desired EM environment to be generated by the time domain array is given.

2.1 Waveform as defined by NEMP standards

IEC 61000-2-9 is a civilian standard that defines a standardized electromagnetic environment caused by a nuclear detonation at high altitude and it includes all three components of NEMP. Its goal is to ensure that all NEMP hardening and verification activities are based upon a consistent and traceable threat definition. A military standard also exists, MIL-STD-2169 [32], but it is classified and hence not suitable for use in this project. This project is only interested in verification of hardness against the E1 component of NEMP, which is described by a double-exponential (DEXP) waveform.

2.1.1 Time domain

The DEXP listed by the IEC has the following form [44]:

$$E(t) = E_0 k (e^{-\alpha t} - e^{-\beta t}) u(t - t_0) \quad (2.1)$$

In this equation, $E(t)$ gives the instantaneous electric field strength in V/m at time t in s , E_0 gives the maximum electric field strength in V/m and k is a dimensionless scaling factor to ensure the difference of exponentials reaches the desired maximum. Furthermore, α is the decay time constant and β is the rise time constant both in s^{-1} . Finally, $u(t)$ is a step function with time shift t_0 . The IEC 61000-2-9 standard defines the following values for these parameters:

Table 2.1: IEC 61000-2-9 defined values for parameters of the DEXP [18]

Parameter	Meaning	Value
E_0	Peak electric-field amplitude	50 kV/m
α	Decay time constant	$4 \times 10^7 s^{-1}$
β	Rise time constant	$6 \times 10^8 s^{-1}$
k	Normalisation factor	1.3

The peak electric field amplitude of 50 kV/m is unattainable for a single MSc project. However, a peak field strength that is too low will not invoke some protection measures on the DUT, which makes testing much less insightful considering the non-linear nature of such protection measures. In this thesis the goal is to achieve 5 kV/m in such a way that increasing the goal field strength is still possible in later research.

Filling in the values from table 2.1 into equation (2.1) gives the complete NEMP standard represented by a DEXP. It is visualised in figure 2.1. This figure shows the temporal distribution of E-field strength of the pulse. Markings are added at 10%, 50% and 90%, to aid in the visualisation of two metrics commonly used to describe such a transient pulse.




Figure 2.1: Visualisation of the standard NEMP DEXP waveform

The first of these metrics is the rise time, commonly described by the rise time from 10% electric field strength to 90% electric field strength, denoted by $t_{10\%-90\%}$. Evaluating the rise time gives:

$$t_{10\%-90\%} = 2.5 \text{ ns} \pm 0.5 \text{ ns} \quad (2.2)$$

The tolerance bound follows from MIL-STD-461 RS105 [34], a publicly available source on NEMP testing procedures that references MIL-STD-2169 [32] for the environment. The description of the environment in MIL-STD-461 RS105 [34] is equal in metrics to the pulse defined by the IEC [18], which implies the IEC defines the same pulse as the MIL-STD-2169 [32] and hence allows taking the tolerance bound from MIL-STD-461 RS105 [34] for this thesis.

The second metric is the pulse duration, which is defined as the time taken between the pulse passing the 50% of peak-field point during its rise and its decay. This is called the Full Width Half Maximum (FWHM). Measuring this from the figure gives:

$$t_{FWHM} = 23 \text{ ns} \pm 5 \text{ ns} \quad (2.3)$$

The tolerance bound is once again taken from MIL-STD-461 RS105 [34].

2.1.2 Frequency domain

The frequency content of the DEXP is also of interest. In order to quantify the spectral density of the DEXP, the Fourier Transform of the DEXP must be taken. The Fourier transform of the DEXP is given by [44] as:

$$E(j\omega) = E_0 \frac{\alpha - \beta}{(j\omega + \alpha)(j\omega + \beta)} \quad (2.4)$$

At very low frequencies, the magnitude spectrum is roughly flat; there is an equal amount of energy contained at all low frequencies. At low frequencies, the parameters α and β are significantly larger than $j\omega$, meaning there is negligible frequency dependence in the magnitude spectrum at low frequencies. As there are two time parameters defining the DEXP, α and β , there are two inflection points in the magnitude spectrum. These define three regions: the low-, intermediate-, and high-frequency region. The inflection points defining these areas are defined in [44] as :

$$f_1 = \frac{\alpha}{2\pi} = \frac{4 \times 10^7}{2\pi} \approx 6 \text{ MHz} \quad (2.5)$$

$$f_2 = \frac{\beta}{2\pi} = \frac{6 \times 10^8}{2\pi} \approx 100 \text{ MHz} \quad (2.6)$$

The low-frequency region is defined as $f < f_1$, the intermediate frequency region as $f_1 < f < f_2$ and the high frequency region as $f_2 < f$. At these inflection points the magnitude starts to roll-off because $j\omega$ becomes of significant size compared to α at f_1 and β at f_2 . In the intermediate frequency region the magnitude rolls off with a factor $1/\omega$, in the high frequency region the magnitude rolls off with $1/\omega^2$ [44]. This is visualized in figure 2.2, which shows the magnitude roll-off from 1 kHz to 100 GHz:




Figure 2.2: Visualisation of the standard NEMP DEXP magnitude spectrum

It is impractical to create an antenna that efficiently radiates the entire spectrum of the DEXP, especially considering the bandwidth it covers. Hence, for a practical NEMP hardness validation, the spectrum of interest must be defined. Defining the upper bound is trivial because of the $1/\omega^2$ roll-off of spectral content at higher frequencies. At some point, the spectral magnitude will be of sufficiently low power that it is considered harmless. For the purposes of this thesis, the 1 GHz mark is chosen, at which point the spectral density is a factor 10^{-3} W/Hz lower than the baseline level before the roll-off.

The lower bound of the spectrum of interest requires more attention and thought, as it is determined by the (DUT). The relevant question is which longest wavelength could efficiently couple into the DUT. Radiative coupling mechanisms are complicated and difficult to accurately predict for a complex object such as a ship, which is part of the reason why practical radiated testing is chosen over simulation or cable injection. Hence, the answer to this question is non-trivial and worthy of a thesis on its own.

To answer this question, numerous expert opinions have been gathered. The general consensus is that realistically speaking the longest external wiring will be no longer than approximately 30 meters. Furthermore, the largest resonant structures on a ship, such as the bridge and the hangar, which act as cavities, are also no larger than approximately 30 meters. From this number the lower bound of the relevant frequency content must be distilled.

To determine the lower bound of the spectrum of interest from the largest resonant structures and longest external cables it is useful to look at the radiation efficiency of dipole antenna with respect to wavelength. The dipole may serve as a useful analogy. A 30-meter-long cable being illuminated by a NEMP pulse will behave like a 30 meter long dipole antenna, which will give insight into which frequency components efficiently couple into the antenna [63]. It is known that a half-wavelength dipole is most efficient, but that around that frequency the antenna will still receive EM energy. Most notably at a quarter wavelength the antenna is equally efficient to a half-wavelength dipole if it is placed on an electrically conducting plane. The lower bound for antenna efficiency is once the antenna becomes electrically small, generally speaking when the wavelength is 10 times longer than the dipole. Hence, if the longest component of a ship at risk for coupling is 30 meters long, the longest wavelength that could couple into the object is 10 times longer, hence 300 meters or 1 MHz. The frequency spectrum of interest in this thesis is 1 MHz to 1 GHz.

To solidify the choice for the chosen frequency spectrum it is useful to look at previous NEMP testing facilities and the frequency content they define. IEC61000-4-32 [10] is a NEMP simulator compendium containing specifications of many existing and previous NEMP simulators, which is a useful source in this part. It does not provide an exact number for frequency range, but from the general dimensions of the antenna an estimate of the lower operating frequency bound can be made. Many NEMP simulators seem to efficiently operate starting from 10 MHz, which is a 30 meter wavelength. If the intended DUT of these simulators is 10 times smaller than what is intended here, it is a sound assumption that the desired frequency content should start 10 times lower.

Support is found in [62], which lists 100 kHz to 1 GHz as the frequency spectrum of interest for NEMP hardness validation testing on power facilities with a maximum DUT external cable length of 100 meters. At a different scale this report used an identical calculation. The Metatech report [1], which focuses on the NEMP impact on the power grid (without listing a size of the DUT), uses 1 MHz to 100 MHz as the relevant frequency spectrum. No other sources were found explicitly stating the operating range of a NEMP simulator, likely because they are classified.

For further validation of this choice the standard NEMP environment may be consulted. NEMP standards are not to be treated as equivalent to NEMP simulator specifications [21], but they may provide insight into whether the chosen frequency range is appropriate to the actual NEMP environment. The IEC61000-2-9 standard [18] states that 96% of NEMP energy is contained in the 100 kHz to 1 GHz range. The NATO AECPT 250 [64] states that 70% of NEMP energy lies in the 1 MHz to 1 GHz range. Hence, according to the standards the chosen frequency spectrum contains a significant amount of the total energy contained in a NEMP. Evidently, the 100 kHz to 1 MHz band contains significant energy, but at such long wavelengths it is likely harmless to the DUT. The omission of this band will introduce some form of error. The chosen frequency band of 1 MHz to 1 GHz however still contains 70% of the total energy in a NEMP.

MIL-STD-461 [34] lists that the generated frequency magnitude must deviate by no more than 10 dB from the theoretical NEMP spectral magnitude as shown in figure 2.2. This number will also be used in this thesis.

2.1.3 Frequency limitation and time domain reconstruction

Setting the desired frequency spectrum to just a portion of the pulse means the radiator that is to be designed will only have to radiate a limited spectrum; from 1 MHz to 1 GHz. In effect the radiator will behave as a band-pass filter (BPF). Such an alteration to the pulse in frequency domain will in the time domain. To visualize the impact of limiting the pulse spectrum, a BPF has been applied to the original pulse in frequency domain. This was done with a perfect BPF (with perfectly sharp cut-off frequencies).

This has been expanded upon with several additional visualisations run at higher lower-frequency limits of the BPF. This was done with the UWB nature of this design in mind, in which the lower limit of the desired waveform is the most critical. As a rule of thumb, antennas must be roughly as large as half their longest radiated wavelength. With a 1 MHz lower limit, this means an antenna height of at least 150 meters for ideal radiation, which may prove challenging in obtaining practical implementation. Hence, to visualise the trade-off in increasing the lower limit of the bandwidth, lower limits of 5 MHz, 10 MHz, 20 MHz, 50 MHz and 100 MHz have been added to this script.

The time domain reconstructions of the pulse with a BPF in the frequency domain is shown below in figure 2.3.




Figure 2.3: Reconstruction of time domain pulse after BPF in frequency domain

Figure 2.3 shows the reconstructed waveform within the finite observation window. The applied band-pass filter is implemented as an ideal, real-valued, symmetric frequency mask. Such a filter is not physically causal because its impulse response is necessarily two-sided in time and can produce pre-ringing. To prevent these non-causal artefacts from contaminating the reported time-domain metrics, the filtering is performed on a strongly zero-padded FFT grid and the inverse-FFT result is truncated back to the original record length, ensuring that any wrap-around or pre-echo components are shifted outside the plotted window and excluded from all subsequent numerical evaluations.

The most striking feature of the reconstructed signals that have had a BPF applied to them is the ringing, which is an artefact of the perfect BPF in the frequency domain. Realistically, an antenna would not behave as a perfect, 'brick wall,' BPF, but it would have a smoother roll-off. This would reduce, but not eliminate ringing as the BPF filters out the DC-component of the pulse, which also causes ringing. This causes the reconstructed pulses to be equal in magnitude on the positive and negative side, which is part of the cause of the ringing. Furthermore, this feature of the reconstructed pulses means that they can realistically be radiated, as there is no DC-component.

What also stands out in the reconstructed pulses is that their peak field strength amplitude is lower, reducing the peak field and shortening the rise time and decay time. This is a natural consequence of filtering out a portion of the pulse that contained energy. The parameters of the reconstructed pulses are laid out in table 2.2 on the next page.

Table 2.2: Effect of ideal band-pass filtering on DEXP pulse energy, fluence, rise-rate, and time-domain metrics

BPF range (MHz)	Energy retained (%)	Fluence (J/m ²)	max(dE/dt) (kV/m/ns)	Peak field (kV/m)	Rise time (ns)	FWHM (ns)
0.1–1000	96.85	1.109×10^{-1}	30.85	49.59	2.453	22.729
1–1000	81.70	9.354×10^{-2}	30.84	46.84	2.409	21.295
3.01–1000	70.00	8.015×10^{-2}	30.57	41.09	2.343	18.600
5–1000	41.39	4.740×10^{-2}	30.75	36.87	2.285	16.176
10–1000	20.95	2.399×10^{-2}	30.40	29.54	2.118	12.138
20–1000	8.42	9.640×10^{-3}	29.35	22.04	1.783	8.050
50–1000	1.92	2.201×10^{-3}	26.04	13.09	1.290	4.226
100–1000	0.45	5.106×10^{-4}	21.63	7.36	0.809	2.419

Table 2.2 shows that omitting the low frequency portion of the pulse has a large effect on the fidelity of the pulse. At a 10 MHz lower limit (which would imply an antenna size of 15 meters), the total energy retained is only 21%. At 100 MHz lower limit, less than one percent of the total original energy is retained. A NEMP pulse contains most of its energy below 100 MHz, as is also visualized in figure 2.2.

In table 2.2 the 100 kHz to 1 GHz range is included, which was omitted from figure 2.3 as it was too similar to the original pulse to give any meaningful insight. This was done in order to compare these metrics to those as described in IEC-61000-2-9 [18], which stated that 96% of energy is contained in the 100 kHz to 1 GHz range. This means that IEC-61000-2-9 comes to roughly the same value as this thesis. It also implies that below 100 kHz the NEMP pulse does not contain significant energy of interest.

NATO AECPT 250 [64] does not support these findings however, as it states that 70% of energy is contained in the 1 MHz to 1 GHz range, while in table 2.2 it is calculated that it is 81.70%. It is important to note here that the AECPT 250 bases its findings on measured NEMP waveforms, which may be susceptible to variations that the simulations are immune to. For instance, there may have been variations in the measured NEMP waveforms, uncertainties in instrumentation, non-ideal radiators or environmental effects. However, as seen in table 2.2, the 3.01 MHz lower limit corresponds to this 70% from the NATO AECPT 250, and the remainder of the parameters (apart from peak E field strength) are still within the constraints as defined in equations 2.2 and 2.3.

A further effect of the reduction of peak and total energy contained in the pulse is the distortion in time domain. Besides being observed in figure 2.3, it is now quantified in table 2.2. The lower peak results in a shorter rise time and a shorter overall pulse duration. It can be observed that dE/dt is conserved despite the filtering up to a lower limit of 20 MHz, after which this parameter is distorted. The magnitude of the dE/dt is determined by the high frequency limit, set at 1 GHz in this thesis. The full NEMP wave has a dE/dt equal to 35.83 kV/m/ns, so the upper limit of 1 GHz also brings some distortion.

The maximum dE/dt is the main driver behind inductive coupling. Inductive coupling is the type that can be explained with Faraday’s law, which states that electromotive force $\mathcal{E} = \oint_{\partial A} \mathbf{E} \cdot d\mathbf{l} = -\frac{d}{dt} \int_A \mathbf{B} \cdot d\mathbf{A}$ [14], and as the magnetic and electric field of an electromagnetic wave are related by the free space impedance, it is dependent on the dE/dt . Inductive coupling is however not what protection is sought against, as it is capacitive coupling that poses the larger threat [72]. Hence, it is maximum E field strength that is the threat defining parameter [73].

To conclude, the radiation of just a portion of the full spectrum of the NEMP causes distortion in time domain which can be seen in the form of reduction in peak and total energy and in shortening and ringing of the pulse. The dE/dt is still conserved up to a lower limit of 20 MHz, despite the distortion in pulse shape and only a tenth of the power remaining. This distortion with respect to the pulse defined in the standard is the result of it being infeasible to radiate the full spectrum. It therefore becomes part of the error made in simulation, keeping in mind the idea that no perfect simulator exists.

2.2 Other Waveform Parameters

The IEC-61000-2-9 [18] defines the temporal distribution of the waveform, but there are other parameters to consider. One of these is the homogeneity of the field as it strikes the DUT. As the NEMP environment is created in a layer in the atmosphere, by the time it reaches the Earth it may be assumed to behave as a plane wave. Furthermore, the field may be assumed to be homogeneous in electric field strength. The plane wave approximation only holds for a radiator at sufficient distance from the observation point, which is impractical considering the desired electric field strength. The MIL-STD-461G RS105 [34] defines that the electric field strength within the testing volume must not be lower than 50 kV/m and no more than twice as large ($-0 + 6$ dB). To keep the goal of this thesis viable and achievable, the goal is 5 kV/m ($-0 + 6$ dB).

Another parameter of the radiated field to consider is the polarization. The IEC61000-2-9 [18] states that the polarization of the NEMP environment is predominantly horizontal. This is dependent on burst location and burst point to observer orientation, as modelling of the generation of a NEMP event can show [66]. From this modelling it can be shown that horizontal polarization holds most of the time. However, it is possible to generate burst locations with burst to observer orientations in which the vertical polarization is dominant. A NASA test report [67], where modern aircraft were tested against high intensity radio frequencies to validate their continued functioning despite illumination by such an environment, showed that certain components responded to vertical polarization but not to horizontal and vice versa. Practically, there is a need to radiate both horizontal as well as vertical polarizations. The confidence a NEMP hardness validation test can instil on a user increases if both polarizations are tested.

The NEMP standard gives a description of NEMP-like environment, but it is not physically realizable as a radiated far-field. The time integral of the DEXP is non-zero, which corresponds to a large DC-component in frequency domain. The result is that practically speaking, the antenna will radiate a much longer pulse that does have a time integral equal to zero. The total waveform decomposes into three elements: the prepulse, the main impulse and the postpulse. The full radiated waveform must satisfy Maxwell's constraints on radiated fields, hence the prepulse and postpulse must be at least partly in opposite field strength to the main impulse. Furthermore, the prepulse and postpulse must be of significantly lower intensity and of significantly longer duration in order to not disturb the test [37]. The NEMP standard hence gives this project the bounds of the waveform shape and characteristics for the main impulse while allowing the prepulse and postpulse to be open to practical implementation.

2.3 Testing Surface Requirements

So far, attention has focused on the characteristics of the waveform that was to be generated. This section turns to the testing surface. This is the area which the NEMP-like impulse that will be generated will illuminate during the test. To still be practically achievable within the scope of this project, an area of 10 m by 10 m is envisioned to be the testing surface. The choice to define a testing surface and not a volume is deliberate and for the same reason. It is acknowledged that this is a comparatively small area, especially considering the broadside of a warship is in the order of hundreds of meters by tens of meters. However, the homogenous illumination of such a surface has not been achieved despite research since 1962, and it will not be achieved in the months allocated to this master's thesis. It is envisioned that if the illumination of an area of 10 m by 10 m is successful, the scaling to larger subsections of the warship is relatively straightforward.

Not illuminating the entire ship in one run of the test introduces errors. In IEC61000-4-23 [68], which concerns NEMP hardness testing variations, lists partial illumination testing (PARTES) as a potential method. It also notes that such a test is fundamentally incomplete, as potential synergistic effects are omitted. Furthermore, any effects that couple into components of the DUT that are of larger size than the testing surface are omitted or not properly addressed. For this last reason it is useful to have the testing surface at least comparable in size to the largest external cable or largest resonant structure. This condition is not met in this thesis, but in the current approach acceptable nonetheless because future expansion of the testing surface is envisioned.

Field homogeneity is further specified by the tolerance bounds given for rise-time and FWHM in equations 2.2 and 2.3. These do not just hold for any observation point, they must hold for the total DUT; or in this case for the testing surface. This means that the spherically expanding waveform must have an as flat as possible wavefront, which is specified by the most critical time requirement: the rise time tolerance bound of ± 0.5 ns. The time difference of arrival across the testing surface of 10 m x 10 m must be no greater than 0.5 ns.

With the size of the testing surface it is now possible to make a first order approximation of the required power. Knowing the fluence of the pulse is approximately 0.1145 J/m^2 , the total required energy over the surface of 100m^2 is 11.45 J. Dividing this by the pulse width of 23 ns gives the required peak power of the generator, giving 0.5 GW. This is the minimum required energy for a generator, as this estimation does not include any losses occurring in the path from generator to testing surface.

The testing surface must exist in the radiated far field of the source; that is, the test requires the wavefront to be flat enough to be considered a plane wave, electric- and magnetic fields to be perpendicular, and electric- and magnetic fields to decay with $1/r$. The exact distance from the radiator this happens is dependent on the antenna used, but a distance of 100 m is defined in the project proposal. This relatively large distance is used due to the coupling between wavelength and far-field distance; the NEMP-like pulse contains very low frequencies that have a long wavelength, hence increasing far-field distance. This may be adjusted to suit the needs of an eventual design: a shorter distance is better for achieving higher field strengths if the far field requirement allows it, but a longer distance may be needed if the far field requirement is not (yet) met at 100 m.

It is envisioned that this surface will be illuminated above the surface of sea water, as that is where the envisioned device under test (DUT) will spend most of its time. This will bring frequency dependent propagation effects. For now, these are assumed to be negligible. The frequency content of interest in this project is of sufficiently short wavelength that the flat seawater surface may be initially assumed to behave like a PEC.

2.4 Conclusion

The goal of this chapter was to define the waveform that is to be generated by the NEMP simulator that is to be designed in this thesis. The starting point for this was the NEMP standards, specifically the IEC61000-2-9 [18]. Further distillation of the requirements was done after this to fully define the specifications of the waveform that is to be generated by the time domain array NEMP simulator in this thesis. After all, NEMP standards are not equivalent to simulator requirements [21]. The requirements are as follows:

Table 2.3: NEMP Test Specification Requirements

Specification	Requirement	Note
Peak electric field strength	5 kV/m $-0/+6$ dB	The standards require 50 kV/m compared to theoretical DEXP
Rise time (10%–90%)	2.5 ± 0.5 ns	
Full width at half maximum	23 ± 5 ns	
Lower frequency bound	1 MHz	
Upper frequency bound	1 GHz	
Spectral magnitude tolerance	± 10 dB	
Polarization	Horizontal and Vertical	
Testing surface area	10 m $\times$ 10 m	To be tuned to antenna MIL-STD-461G RS105
Test distance from radiator	100 m	
Field homogeneity	$-0/+6$ dB	

This list should be seen as the goal of this thesis. However, the field-strength target has been reduced the desired field strength from 50 kV/m as required by the standards [18] to 5 kV/m in order to make this thesis project achievable. Further deviation from this list of requirements may be necessary to maintain practical viability. Any deviation will be considered with respect to the other requirements: some metrics are more important than others. The peak electric field strength for instance is needed to evoke the non-linear over-voltage effects, and hence is more important than for instance the field homogeneity or the exact time fidelity. The ability to radiate both polarizations is desirable, but if it requires a design that is practically unfeasible it must be given up. The lower relevant frequency bound is set by the object size in combination with a $\lambda/10$ coupling-efficiency criterion. However, a lower bound of 1 MHz efficient radiation would mean an antenna length of at least 150 meters, which is practically unfeasible. Instead, a choice may have to be made to deviate from this list of requirements by increasing the lower frequency bound and finding another solution for the coupling between 1 MHz and the new bound. These are all design choices that will be considered in the remainder of the design process in this thesis.

Chapter 3

Baseline

To be able to meaningfully quantify any results achieved in this thesis, they must be compared to a baseline. For this research, the baseline is the performance of the EMIS-III-VPD, the NEMP simulator used by the Royal Netherlands Navy in the 90s. In this chapter, this simulator will be described, modelled and its performance quantified.

No publicly available sources describe the performance of this antenna exactly. Therefore, using information that is available, a numerical model is made of this antenna. Numerical simulations run with this model will provide some numbers estimating its performance, as well as give insight into design choices and trade-offs made for this specific simulator. These numbers then serve as the baseline and form the conclusion of this chapter.

3.1 The EMIS-III-VPD

The "Electromagnetic Impuls Simulator" (EMIS) is the name given to the original three NEMP simulators developed by TNO between 1970 and 1992. The successor of this series of simulators was the EMPSIS (EMP Simulator for Ships) proof-of-concept, shown in figure 1.2. This was a bounded wave, TEM-cell like test that was proven to effectively generate a NEMP-like waveform while being impractical to implement and having a definitive upper limit of DUT size [8].

EMIS consists of three test set-up designs. The first was a proof of concept and the second a bounded wave, TEM-cell like test for army and Air Force vehicles. The third, EMIS-III, is of interest to this thesis as it was designed specifically for military platforms that could fit under EMIS-II. The total test consists of two antenna: a vertically polarized dipole (VPD) and a horizontally polarized dipole (HPD). The HPD could also be adjusted to attach into a hybrid simulator into which a DUT (of limited size) could be placed [9]. It was originally envisioned that all large platforms would be submitted to both the HPD and the VPD, however in practice only the VPD was used to test warships of the Dutch Navy. The EMIS-III-VPD is the simulator that provides the baseline for this thesis.

The EMIS-III-VPD is a monocone antenna with 20 meters height and 35 meter cone diameter [10]. A monocone is half of a biconical antenna, which exists virtually when the monocone is placed above an electrically conducting surface. As the antenna was placed on land, there must have been some form of electrically conducting mesh placed under the antenna that guides the wave to at least the sea water surface. A 500 kV Marx generator was placed in the tip of the cone. In figure 3.1 on the next page the set-up is shown with the M-frigate Hr. Ms. Van Nes in the test position. The antenna is placed on the grass to the left of the ship, where mainly the fibreglass pylons supporting the monocone and the pulse generator are visible.



Figure 3.1: The EMIS-III-VPD test set-up in Den Helder

3.2 First-order estimate

In this section the performance of the EMIS-III-VPD will be estimated to first order. There are no publicly available sources quantifying the performance of the antenna, so an estimation must be made in order to attain the baseline. There are sources that provide clues, such as the IEC-61000-4-32 [10] and the report of a test by the EMPRESS-II [11], an American monocone VPD placed on a barge. a first-order estimate is made using some "back-of-the-envelope" calculations, which will be followed by a simulation of the antenna in the next section.

The height of the antenna can be used to estimate the lower bound of the effective radiating bandwidth. For wavelengths that are more than twice as long as the length of the dipole the antenna is no longer efficient. Hence, a first-order estimate of the lower bound of the antenna is approximately:

$$f_{min} \approx \frac{c}{4h} = 3.75 \text{ MHz} \quad (3.1)$$

Note that this upper bound is determined by $\lambda/4$ with respect to the height of the antenna h . This is as the antenna uses an electrically conducting plane which virtually doubles the length of the antenna. It must also be noted that this is definitely not a hard cut-off, and lower frequencies still do radiate, albeit with lower efficiency than frequencies above 3.75 MHz

a first-order estimate of the upper frequency bound is challenging, as it is dependent on the tolerances of construction accuracy in the feeding region. This is difficult to estimate using images of the antenna. Furthermore, this number is relatively un insightful, as can be seen in table 2.2, where the shifting of the lower limit of the frequency band by MHz's caused significant alterations to the waveform.

If the EMIS-III-VPD had a requirement of upper operating frequency equal to that of this thesis, 1 GHz, then the tolerance can still be estimated. The upper frequency bound f_H is typically constrained by the electrical size of the feed region. To maintain quasi-TEM behaviour and avoid strong localized resonances at f_H , the smallest feed-region features should remain electrically small:

$$\ell_{feed} \lesssim \frac{\lambda_H}{20}, \quad \lambda_H = \frac{c}{f_H}. \quad (3.2)$$

With $f_H = 1$ GHz, $\lambda_H \approx 0.30$ m and thus $\ell_{\text{feed}} \lesssim 15$ mm. The UWB nature of the antenna yields a tens-of-meters tall structure to have mm constraints.

From figure 3.1 the distance between the generator and the DUT can be estimated. The M-frigate of the Karel Doorman class is known to be 14.88 meters wide, which when used as a scale reference allows estimating the distance from DUT centre point to pulse generator to be approximately 100 meters. In IEC-61000-4-32 [10] the electric far field strength is listed to be 2 kV/m at 100 meters distance, which can be taken straight from the report to this thesis as first-order estimate for the baseline value. It must be noted that the IEC report is a publicly available source, which is likely to have received the non-classified number for peak electric field strength. Furthermore, considering the desired bandwidth of a NEMP simulator in generator, it may be the case that at 100 meters the conditions are not yet purely far field.

The field homogeneity can be estimated with these numbers. The EMIS-III-VPD is designed to illuminate the entire DUT. Hence, the surface it radiates must be as large as the largest warship currently in service in the RNLN. This gives a testing surface that is rectangular, 200 meters long and 50 meters tall. The VPD is assumed to be placed so it is 100 meters away from the middle of the testing surface sitting on the ground. This test set up is schematically illustrated in figure 3.2 on the next page.




Figure 3.2: Schematic of the EMIS-III-VPD set up to radiate on the testing surface

Assuming that field strength is dependent only on spherical spreading by $1/r$, then the field strength at the edges is dependent only on the longer propagation distance, which is calculated on the next page:

$$\Delta r = \sqrt{100^2 + 100^2} - 100 = 41.42 \text{ m} \quad (3.3)$$

$$\frac{E_{\text{edge}}}{E_{\text{max}}} = \frac{100}{141.42} = 0.707 \quad (3.4)$$

$$E_{\text{edge}} = 0.707 \times E_{\text{max}} = 1.41 \text{ kV/m} \quad (3.5)$$

1.41 kV/m is 3 dB lower than the original 2 kV/m that was assumed as the maximum electric field strength. The difference in time of arrival can also be estimated from this geometry:

$$\Delta t = \frac{\Delta r}{c} = \frac{41.42}{3 \times 10^8} \approx 138 \text{ ns} \quad (3.6)$$

a first-order estimate of the performance of the antenna gives a maximum electric field strength of 1.41 kV/m – 0 + 3dB with a maximum difference in time of arrival of 138 ns and a frequency spectrum starting at ~ 3.75 MHz

Next, the power required at the feed-point can be estimated. Assuming far-field conditions and a directive gain G in the direction of interest, the RMS electric field magnitude at range r can be related to radiated power as

$$E(r) = \frac{\sqrt{30 P_{\text{rad}} G}}{r} \Rightarrow P_{\text{rad}} = \frac{(E(r) r)^2}{30 G}. \quad (3.7)$$

A biconical has a gain on the order of a half-wave dipole, i.e. $G \approx 1.6$ (≈ 2.1 dBi). The EMIS-III-VPD geometry is a monocone over an electrically conducting surface with height $h = 20$ m and cone diameter $D_c = 35$ m. The cone angle follows from the geometry. With radius $R = D_c/2 = 17.5$ m,

$$\theta = \arctan\left(\frac{R}{h}\right) \approx \arctan\left(\frac{17.5}{20}\right) \approx 41.2^\circ, \quad \alpha = 2\theta \approx 82.4^\circ. \quad (3.8)$$

Using the quasi-TEM approximation

$$Z_{\text{in}} \approx \frac{\eta_0}{\pi} \ln\left(\cot(\alpha/4)\right), \quad (3.9)$$

with $\eta_0 \approx 377 \Omega$ yields $R_{\text{in}} \approx 59 \Omega$.

Since the monocone-over-PEC radiates into half-space, a first-order gain adjustment relative to a dipole-like radiator is to take $G \approx 2 \times 1.6 \approx 3.2$. Using $E_{100} = 2$ kV/m, $r = 100$ m and $G = 3.2$ gives

$$P_{\text{rad}} = \frac{(E_{100} r)^2}{30 G} = \frac{(2 \times 10^3 \cdot 100)^2}{30 \cdot 3.2} \approx 4.17 \times 10^8 \text{ W}, \quad (3.10)$$

and with $R_{\text{in}} \approx 59 \Omega$ the peak feed voltage estimate becomes

$$V_{\text{pk}} \approx \sqrt{2 P_{\text{rad}} R_{\text{in}}} \approx \sqrt{2 \times 4.17 \times 10^8 \times 59} \approx 3.12 \times 10^5 \text{ V} \sim 0.31 \text{ MV}. \quad (3.11)$$

If this voltage is supported across a 15 mm feed region, the implied electric field strength at the feed point is

$$E_{\text{feed}} \approx \frac{V}{d} \approx \frac{3.12 \times 10^5}{0.015} \approx 2.08 \times 10^7 \text{ V/m} = 20.8 \text{ MV/m}, \quad (3.12)$$

This electric field strength is above the ~ 3 MV/m air breakdown limit, implying that when operated in the atmosphere the feed point would arc and short-circuit. However, it is known that the EMIS-III-VPD used SF_6 gas when operating, making the calculated electric field strength of 20.8 MV/m plausible.

These values give insight into the order of magnitude of the performance of the EMIS-III-VPD, and can be used to validate the simulations performed in section 3.3.

3.3 Simulation

Simulations of the EMIS-III-VPD are performed for two reasons. Firstly, it is to make a detailed estimation of the performance of the EMIS-III-VPD in order to accurately quantify any improvements proposed by this research. Furthermore, these simulations provide insight into the physical effect of end-termination and resistive loading on the broadband behaviour of an antenna. Simulations are done using the CST Microwave Studio program, which amongst others has a time-domain solver. The CST time-domain solver is based on the Finite Integration Technique (FIT), which discretizes Maxwell's equations in their integral form on a hexahedral grid. The electromagnetic fields are updated explicitly in time by computing the curl relations between electric and magnetic field components, allowing broadband responses such as S-parameters and radiation characteristics to be obtained from a single transient simulation [12].

The main challenge in simulating a structure such as the EMIS-III-VPD is building an efficient model. The large size of the antenna in combination with the ultra wideband desired frequency range gives rise to a huge number of discretization cells. This number can somewhat be reduced by making use of symmetry in the antenna's shape, as well as by using a coarse discretization step-size and locally refining around small objects, such as the feeding region, the termination region and the thin wires. A discretization step-size of 10 per smallest wavelength was made over the whole model, with local refinement to 20 around the feed region. Further improvement of model size was made by applying the 'thin wire model,' a standard CST function in which the radius of thin wires is assumed to be zero. This alleviates the model to discretize all around and in the wires, instead simplifying it to an infinitely thin current [12]. Finally, the PEC surface is modelled using a symmetry plane with zero electric field strength as definition.

To gain more understanding in what drives antenna operating characteristics, more than one model was made. The first model contains the simplest architecture with 8 thin wires each fitted with a single terminating lumped element. The second model expanded upon this by introducing resistive loading. Finally, a model of a solid cone was made, an idealized model to which the 8-wire models can be compared.

3.3.1 Far-field considerations

The antenna size (20 m height and 35 m diameter) raises a concern about the validity of evaluating the antenna (far-field) radiation patterns across the testing surface at 100 m distance. A more complete analysis of this issue is done in Chapter 5, section 5.2, but for now it may be assumed that the radiated electromagnetic energy has not yet settled into its asymptotic far field. This also concerns the post-processing methodology adopted both in this chapter for the EMIS-III-VPD and in Chapter 5 for the reflector IRA models. All time-domain pulse reconstructions and field-strength predictions presented in this thesis are derived from CST broadband far-field monitors. These monitors perform a near-field-to-far-field transformation on the recorded surface currents and return the asymptotic far-field pattern of the antenna [12]. The output of such a monitor at a stated reference distance r is, by construction, the asymptotic pattern $\vec{E}_\infty(\theta, \phi, f)$ scaled by $1/r$:

$$\vec{E}_{\text{CST}}(r, \theta, \phi, f) = \frac{\vec{E}_\infty(\theta, \phi, f)}{r}, \quad (3.13)$$

which is conceptually equivalent to placing the observer at infinity and rescaling the amplitude, rather than to evaluating the field at a finite distance r . Two approximations underlie this construction [12]: a planar wavefront at the observer, such that path-length differences across the aperture are evaluated only to first order in angle, and pure $1/r$ amplitude decay without near-field reactive contributions.

As will be discussed in more detail in section 5.2, near-field measurements are subject to an effect called near-field dispersion, in which contributions of the pulse originating from different locations on the antenna arrive at different times. This causes pulse spreading, artificially elongating the rise-time and the FWHM close to the antenna. Therefore, any results presented in this thesis on the EMIS-III-VPD must be taken lightly. The choice to retain the far-field-monitor methodology despite the above is deliberate; it is by far the most efficient metric to extract from CST, as the simulation intensity is significantly lower for the far-field assumption. Furthermore, synthesizing an array using an antenna is also significantly simplified using the far-field addition of contributions.

3.3.2 Cone Impedance

One problem that was encountered straight away was that early versions of the models presented in this section used an incorrect value for reference impedance, based on the cone impedance of the structure. It is stated in the IEC61000-4-32 [10] that the EMIS-III-VPD is a bicone with cone impedance 150Ω . As it is actually a monocone placed above an (assumed to be) perfectly electrically conducting (PEC) surface, the impedance of the monocone is half, yielding 75Ω . This value was taken to mean that all components on the antenna were tuned to 75Ω .


Initial results were odd, showing the signs of bad impedance match across the full frequency spectrum regardless of tuning the terminating resistances. Further investigation led to manually calculating the antenna impedance, using the following equation for cone impedance defined by Tesche [13]:

$$Z = 120 \ln \left(\cot \left(\frac{\theta}{2} \right) \right) \quad (3.14)$$


In this equation, cone impedance depends only on θ , which is the cone angle: the angle between the two flaring parts of the antenna. Half this angle is the angle between the ground plane and the monocone, which for the geometry of the EMIS-III-VPD is equal to 42.41° . Filling in gives a cone impedance of 60Ω . Tuning all components to 60Ω reference impedance eventually did give plausible results.

3.3.3 End-termination model

The first model made has the simplest architecture. It exhibits the same 8 wires that the EMIS-III-VPD had, which meet at the base of the antenna. The wires do not meet directly, instead a small solid cone of the same proportions as the antenna but only 0.3 m tall connects them, as shown in figure 3.3a. This was done in order to mitigate any numerical errors caused by the clipping of the wires at the base, which led to distortion of the radiation pattern and the S-parameters. This solid cone is fed by a discrete port at its tip, which covers the 0.1 meter long gap between the tip of the cone and the PEC surface. At the top, each wire has a lumped element connecting the wire to the ring, which is intended to terminate the current travelling through the wires, as shown in figure 3.3b.



(a) The solid cone base connecting the 8 wires



(b) The wire termination connecting to the ring

Figure 3.3: The base and termination of the first 8-wire EMIS-III-VPD model

To tune the end-termination of the wires efficiently, the model was slightly adapted to determine the impulse response of a single wire. One termination of the structure is replaced by a discrete port with a reference impedance of 60Ω , while the discrete port originally located at the base feeding gap is replaced by a lumped element resistor of 60Ω . The solver is then executed to observe the voltage and current response along the wire using standard voltage and current monitors. By analysing the ratio of voltage to current at an observation point, an estimate of the broadband matching behaviour of a single wire is obtained. Because the correct termination impedance is generally not known a priori, the procedure is performed iteratively by adjusting the termination resistance until reflections along the wire are

minimized, indicating that the termination matches the characteristic impedance of the structure. The iterative tuning procedure yielded an optimum termination of $60\ \Omega$, which is consistent with the analytical cone impedance estimate.

The structure was then replaced back into its correct form. The solver was executed again, this time with far field monitors set up at several observation frequencies. The results are presented below, starting with the S11 parameter over the full desired frequency range:




Figure 3.4: S11 parameter of the first model from 1 to 1000 MHz

While the S11 parameter shown in figure 3.4 does fluctuate over the desired frequency band, most notably showing deep troughs at 350 MHz and 615 MHz, it does remain below -15 dB for almost the full spectrum. These deep notches are attributed to structural resonances and frequency-dependent modal behaviour, for which the antenna input impedance approaches the $60\ \Omega$ reference particularly closely. Only at the very lowest frequencies, below approximately 50 MHz, the amount of energy reflected to the antenna feed is higher than -15 dB, approaching -10 dB at 1 MHz. At the low end of the band the antenna dimensions are small compared to the wavelength, so the structure behaves as an electrically short radiator and the input match degrades.

Next, the antenna radiation efficiency is inspected. It is presented below in figure 3.5, which shows the calculated radiation efficiency from 1 to 1000 MHz measured at 50 MHz intervals.






Figure 3.5: Radiation efficiency of the first model from 1 to 1000 MHz

The radiation efficiency is moderate at lower frequencies and reaches a minimum of approximately 0.6 near 200 MHz, after which it increases steadily toward unity at higher frequencies. The dip at 200 MHz corresponds to the peak in S11 in figure 3.4. Since the model uses only a $60\ \Omega$ end termination, this behaviour suggests that a significant fraction of the accepted power is dissipated in the terminal load at lower and intermediate frequencies. Inspection of the power loss in lumped elements confirms this. At higher frequencies the structure radiates more effectively, such that a larger fraction of the accepted power is converted into radiation before reaching the termination.


Next, the far field radiation patterns are inspected. At four frequencies these are presented in this paper; 1 MHz, 100 MHz, 500 MHz and 1000 MHz. This is shown in figure 3.6 below.




(a) 1 MHz



(b) 100 MHz



(c) 500 MHz



(d) 1000 MHz

Figure 3.6: Radiation patterns of the first EMIS-III-VPD model at four frequencies.

At 1 MHz, the antenna is electrically small; it is 20 meters tall while radiating a 300 meter wave. The radiation pattern therefore approaches that of an electrically short vertical monopole-like radiator above PEC, with a distinct lack of directivity. There is only a null along the z-axis, both in positive and negative directions.

From 15 MHz upwards the structure size becomes significant for the wavelength. Hence, from this point onwards the radiation pattern will reflect the intended functioning of the monocone antenna. At 100 MHz this can be seen: the radiation is most intense in the XY-plane and decreasing as the angle θ approaches zero. What is striking about this plot is the erratic nature of the radiation pattern, showing many peaks and troughs across short distances. This is a result of the wires being sparsely distributed, resulting in a significant distance between the wires at the top of the antenna. This causes significant phase difference between radiation departing from different wires, causing the erratic radiation pattern.

At 500 MHz the erratic nature of the radiation pattern has decreased. The higher frequency means that distances between peak and trough are shortened to the point the radiation pattern smoothens.

Finally, at 1000 MHz, grating lobes appear. Firstly, the main bundle of radiation is no longer directed in the XY-plane, but is instead elevated to some degrees above this plane. Secondly, a grating lobe is pointing upwards has appeared with a trough between these two bands of higher radiation intensity. This is attributed to higher-order current distributions on the electrically long structure.

Another interesting performance metric to inspect is the power absorbed in the terminating resistances. These resistances will convert absorbed power into heat, which, considering the high power the EMIS-III-VPD worked with, is interesting as it is a limiting factor of operation. The power absorbed by all resistances is plotted below on the next page in figure 3.7.

Note that this figure is scaled to an input power of 0.5 W, which means that while the trend-line is correct, the total absorption calculated must be scaled. Most power is absorbed around 200 MHz, where approximately 26 % of input power is absorbed by the lumped elements. On average, 6.78 % of accepted power is absorbed in the resistances. This can be used to then scale to total absorption: in equation 3.10 in the previous section it was estimated that 0.417 GW needed to be radiated to achieve the 2 kV/m at 100 m distance. For this specific model, it would require 0.72 GW needs to be supplied to it, as its efficiency is approximately 0.58 at its worst frequency. Assuming linear scaling of power absorption, this would mean that 0.2 GW is absorbed as a maximum. Note that this is peak absorbed power; not average power. The pulse duration the standard NEMP pulse 23 ns (FWHM), so the resistances absorb approximately 4.6 J of energy.




Figure 3.7: Power absorbed by the end termination from 1 to 1000 MHz

3.3.4 Resistively loaded wire model

A wire of finite length is inherently narrowband. The endpoint of the wire will always reflect a portion of the wave back towards the feed, hence forming a standing wave along the length of the wire. Close to the feed point this causes the local wave to be a superposition of two waves: the wave that is fed to the wire and the wave that is reflected from the termination, each travelling in opposite directions. The difference in distance travelled between the waves causes a phase difference, which hence causes the superposition to result in a local maximum or minimum. In the case of a wideband feed, the wavelength will also determine the phase difference, resulting in a portion of the feed frequency spectrum to have a maximum at a certain location while another portion of the frequency spectrum has a minimum at that same location. As the local voltage and current at the feed point together determine the input impedance, the input impedance therefore is frequency dependent. Note that [14]:

$$Z_{in} = Z_0 \frac{1 + \Gamma e^{-j2\beta L}}{1 - \Gamma e^{-j2\beta L}} \quad (3.15)$$

End termination, as applied in the first model, will dampen this effect, but never completely eradicate it. Hence, some form of ripple is always seen in the S11 parameter as a function of frequency.

Wu and King [15] proposed a method of resistive loading of a wire in 1965 to counter the standing wave effect for the improvement of the broadband characteristics of a dipole antenna. They proposed to segment the wire into a number of segments connected by resistances, with exponentially increasing resistance as the distance to the feed-point increased. This would allow for a forwards travelling wave from the feed to the end point to exist for long enough to radiate, while at the same time dampening it sufficiently to eradicate it before it is able to form a standing wave at the feed point.

It is never explicitly mentioned in open sources that the EMIS-III-VPD makes use of resistive loading on its wires to enhance its broadband performance. Therefore, the resistively loaded wire model is introduced here as a physically motivated comparison case. The first model is adapted to include resistive loading to show the effect of applying the Wu-King [15] method. In order to apply this method, each wire is segmented such that twenty uniformly distributed narrow gaps are formed, and these are consequently filled with lumped elements. The resistance of a given lumped element, R_k is given by:

$$R_k = k \ln \left(\frac{h - z_{k-1}}{h - z_k} \right) \quad (3.16)$$

In which k is the index of the lumped element, numbering from 0 to 20, h is the length of the wire and z_k is the position of the antenna. The lumped elements are uniformly distributed over the wire. To find the total resistance of a wire, a parameter sweep was conducted. For this, a distribution of resistances which together summed to 60 Ω , distributed according to equation 3.16, was applied to each of the wires. Each resistance was multiplied by a parameter, named 'multiplier,' and the simulation was run across several values for multiplier. One value for the multiplier provided the best match, and the process was repeated with new values for multiplier around the value that provided the best match previously. This way the best value was approached iteratively, albeit in a time-consuming manner. A slightly lower total value for termination was found to be optimal: $0.85 * 60 = 51\Omega$. The segmentation of the wires is shown in more detail below, in figure 3.8. The feeding and termination is the same as in the previous model.




Figure 3.8: A wire segmented by lumped elements of increasing resistance

As with the previous model, the S11 parameter is inspected first.




Figure 3.9: S11 parameter of the second model from 1 to 1000 MHz

The S11 parameter of the model with resistive loading, in figure 3.9, shows some differences with respect to with end-termination as in figure 3.4. Firstly, the very low end of the frequency spectrum shows significantly less reflection, only reaching to -10 dB at 1 MHz. It does however show a sharp peak in reflection around 200 MHz, which with end-termination was smeared out over a wider bandwidth at the same frequency. Averaged over the 1-1000 MHz band, the end-terminated model yielded a lower mean reflection coefficient than the resistively loaded model (about -20.2 dB versus -18.6 dB), indicating that the addition of distributed resistive loading does not improve the average broadband match in this implementation.

Next, the radiation efficiency is inspected. Once again, it is measured at 50 MHz intervals.




Figure 3.10: Radiation efficiency of the second model from 1 to 1000 MHz

The radiation efficiency with resistive loading, shown in figure 3.10, is lower over the whole band that it is for end-termination. Averaged over the 1–1000 MHz band, the end-terminated model exhibits a mean value of 0.87, while resistive loading has a mean of 0.81. Lower radiation efficiency means more electromagnetic energy is absorbed by resistances and convert to heat than is radiated by the antenna. This result contradicts the fact that the resistive loading model had lower total resistance over each wire, 51 Ω instead of 60 Ω , but the effect of having resistances earlier in the wire has more impact than lower total resistance. The general shape of the radiation efficiency across the band of interest is similar to that of the model with end-termination, as shown in figure 3.10.

Finally, the radiation patterns are inspected. For relevant comparison with the previous model, this is done at 1 MHz, 100 MHz, 500 MHz and 1000 MHz.




Figure 3.11: Radiation patterns of the second EMIS-III-VPD model at four frequencies.

The 1 MHz plot will be recognisable to the reader as the radiation pattern similar to a unit dipole, a consequence of an extremely short antenna with respect to the wavelength. The differences between resistive loading and end-termination start once the wavelength is proportional to the antenna size; around 15 MHz. At 100 MHz the radiation pattern still shows significant peaks and troughs, a consequence of the sparse wire distribution. However, these are cleaner and neatly distributed along the XY-plane, an improvement with respect to the previous model. The effect is more pronounced at 500 MHz, which shows a significantly cleaner radiation pattern than the previous model did. Terminating the waves earlier in the wire causes there to be less energy to be radiated at the top of the wire, hence causing the large distance difference between wires at this height to be a less significant effect. Finally, at 1000 MHz, the sidelobing that was seen with end-termination is gone. This too can be attributed to resistive loading: damping the wave in the wire before it can form a sidelobe.

Finally, the power absorbed by the resistances for the resistive loading is inspected:




Figure 3.12: Power absorbed by the resistive loading from 1 to 1000 MHz

Figure 3.12 shows the power absorbed by all resistances across the desired frequency spectrum by all resistances in the resistive loading system. It must be again noted that the total resistance per wire is lower than it was for end-termination: 51Ω for resistive loading and 60Ω for end-termination.

This figure is once again scaled to 0.5 W input power, so while the trend is correct the scale is not tuned to a NEMP-like pulse fed to the EMIS-III-VPD. As with the end-termination model, most power is absorbed in the 100 to 300 MHz band, but interestingly the resistive loading shows higher frequency dependency; nulls and peaks dominate the pattern. The peak absorption is 28.4 % at just under 200 MHz, the average absorption over the entire band is 6.6 %.

Using the same estimates as with the end-termination model, the peak absorbed power is estimated to be equal to 0.23 GW, results in 5.3 J absorbed by a single pulse.

3.3.5 Solid cone model

The final model that was made is a solid cone of the same proportions as the EMIS-III-VPD. This is the shape the EMIS-III-VPD approximates with wires, but it does not do so perfectly. Distortions are especially visible in the middle half of the frequency spectrum, where the distance between the wires is in the order of magnitude of half the wavelength of the emitted pulse, hence effectively causing super-positioning. To see what the far field would look like without distortion from the wires, the solid cone model is made. This model will show the idealised performance of this antenna




Figure 3.13: The solid cone model of the EMIS-III-VPD

The feeding of the cone is done in the same manner as in previous models. The termination is done by a smooth edge leading into a solid lumped element with the same radius as the top of the cone. The edge was deliberately smoothed in order to prevent edge scattering, which became very dominant in this architecture at higher frequencies. This was not a feature the actual EMIS-III-VPD had, but this is acceptable as this model is intended to show an idealized conical radiator.

First, the S11 parameter of the solid cone model is shown:




Figure 3.14: S11 parameter of the third model from 1 to 1000 MHz

The S11 parameter reveals a highly broadband response to an impulse. Especially at low frequency the S11 parameter is as low as -40 dB, and omitting one peak at 690 MHz the S11 parameter does not exceed -15 dB.

Next, the radiation efficiency is inspected.





Figure 3.15: Radiation efficiency of the third model from 1 to 1000 MHz

The radiation efficiency is once again only measured at 50 MHz intervals. Inspection reveals that the radiation efficiency of the solid cone model is much higher than it was for the 8-wire models, which suffered a dip in efficiency as low as 0.5 around 200 MHz. The solid cone model remains above 0.9 across the whole band, approaching 1 at the lower half of the desired spectrum.


Finally, the radiation patterns are inspected. This is done at the familiar 4 frequencies: 1 MHz, 100 MHz, 500 MHz and 1000 MHz. This is shown in figure 3.16, on the next page.

The 1 MHz plot once again shows a radiation pattern similar to a small dipole, with a distinct lack of directivity. At higher frequencies, specifically the 100 and 500 MHz plots, the radiation pattern is as it is intended: omnidirectional and smooth in the XY plane. At 1 GHz the pattern starts to break down: edge scattering is more prominent and in the XY plane there are more peaks and nulls than at lower frequencies. This is not perfect, but further optimization of this model is deemed irrelevant to this research.


The solid cone model has almost no need for its end-termination, as it almost perfectly radiates all the supplied power. Therefore, an inspection of the power absorbed in the end-termination of the antenna is deemed superfluous.




(a) 1 MHz



(b) 100 MHz



(c) 500 MHz



(d) 1000 MHz

Figure 3.16: Radiation patterns of the solid cone model at four frequencies.

3.3.6 Comparison and Observations

The three models presented in the previous subsections each represent a different level of idealisation of the EMIS-III-VPD. The solid cone model represents the theoretical ideal of a conical radiator, the end-terminated 8-wire model represents the closest approximation to the actual EMIS-III-VPD construction, and the resistively loaded 8-wire model introduces a physically motivated modification to investigate the effect of distributed damping on broadband performance. Comparing the three models across the presented metrics provides insight into which physical effects dominate the antenna's behaviour across the 1–1000 MHz band.

The S11 parameters of the three models reveal that the solid cone is by far the best matched structure, remaining below -15 dB across nearly the entire band and reaching as low as -40 dB at lower frequencies. Both 8-wire models show more reflection, particularly in the 100–300 MHz range where structural resonances are most pronounced. The end-terminated model exhibits broad, moderate reflection peaks, while the resistively loaded model concentrates the mismatch into a sharper peak near 200 MHz. Averaged over the band, the end-terminated model achieved a slightly lower mean reflection coefficient (-20.2 dB versus -18.6 dB), suggesting that the Wu-King resistive loading distribution, while effective at suppressing standing waves along the wires, does not improve the average broadband impedance match in this particular implementation. This may be because the resistive loading was optimised for travelling-wave behaviour on the individual wires rather than for the overall input impedance of the parallel combination of eight wires.

The radiation efficiency comparison is perhaps the most revealing. The solid cone model maintains an efficiency above 0.9 across the entire band, approaching unity at lower frequencies. Both 8-wire models suffer a significant dip near 200 MHz, with the end-terminated model reaching a minimum of approximately 0.6 and the resistively loaded model dropping even lower to approximately 0.5. The resistively loaded model exhibits lower mean radiation efficiency (0.81 versus 0.87 for end-termination) despite having lower total resistance per wire ($51\ \Omega$ versus $60\ \Omega$). This confirms that the placement of resistances along the wire is more significant than the total resistance: absorbing energy earlier in the propagation path removes it before it has the opportunity to radiate. This is the fundamental trade-off of the Wu-King loading scheme—it improves travelling-wave behaviour and reduces standing-wave effects at the cost of converting a larger fraction of input power into heat rather than radiation.

The radiation patterns expose the most significant difference between the 8-wire models and the solid cone. At low frequencies (1 MHz), all three models produce the expected electrically-small monopole pattern with little directivity, which is expected as the antenna is orders of magnitude smaller than the wavelength. In the intermediate frequency range (100–500 MHz), the solid cone produces the smooth, omnidirectional pattern expected from a conical radiator. Both 8-wire models show erratic radiation patterns at 100 MHz due to the large electrical distance between adjacent wires at the top of the cone, which causes significant phase differences between radiation contributions from different wires. However, the resistively loaded model produces noticeably cleaner patterns than the end-terminated model in this range, as the distributed loading reduces the current amplitude at the wire endpoints where the inter-wire spacing is largest. At 500 MHz this improvement becomes more pronounced: the resistively loaded model's radiation pattern is significantly smoother than that of the end-terminated model. At 1000 MHz, an interesting reversal occurs: the end-terminated model develops grating lobes attributed to higher-order current distributions on the electrically long structure, while the resistively loaded model suppresses these lobes by attenuating the current before it can contribute to such patterns.

The power absorbed in the lumped elements follows a consistent trend for both 8-wire models: absorption peaks in the 100–300 MHz range where a significant fraction of the wave energy reaches the terminations without being radiated. The peak absorption is comparable between the two models (approximately 26 % for end-termination and 28.4 % for resistive loading), while the average absorption across the band is similar as well (6.78 % and 6.6 % respectively). Scaling these values to the estimated 0.72 GW input power needed to generate 2 kV/m at 100 m gives peak energy absorption on the order of 4.6–5.3 J per pulse, which is modest enough that thermal management of the resistances is not expected to be a limiting factor for single-pulse operation.

To arrive at an estimation of peak electric field strength at the testing surface, the realised gain of each model was computed as the product of directivity, radiation efficiency and mismatch factor. The directivity was obtained by averaging the far-field directivity at $\theta = 90^\circ$ over all azimuth angles ϕ , which gives a representative value for the nominally omnidirectional monocone. The input voltage was taken as a DEXP pulse with 500 kV peak amplitude (consistent with the Marx generator), and the resulting time-domain electric field at 100 m was synthesised via inverse Fourier transform. The results are summarised in table 3.1.

Table 3.1: Estimated peak electric field strength at 100 m for the three EMIS-III-VPD models.

Model	Peak E -field at 100 m [kV/m]
End-terminated 8-wire	1.28
Resistively loaded 8-wire	1.31
Solid cone (ideal)	1.98

The solid cone model produces a peak field of approximately 2 kV/m, which is consistent with the value reported in IEC 61000-4-32 [10] for the EMIS-III-VPD at 100 m distance. This agreement serves as a validation of the CST model and the calculation methodology. Both wire models produce approximately 1.3 kV/m, which is 3.6 dB below the solid cone. This difference is attributable entirely to the lower radiation efficiency of the wire models: energy dissipated in the terminating resistances cannot contribute to the far field. Notably, the end-terminated and resistively loaded models produce nearly identical peak field strengths despite their different loading strategies, because the efficiency penalty of resistive loading is offset by its marginally improved pattern uniformity when averaged over azimuth.

In summary, the comparison reveals that the 8-wire construction of the EMIS-III-VPD introduces two main degradation mechanisms relative to the ideal conical radiator: reduced radiation efficiency due to ohmic losses in the terminations, and distorted radiation patterns due to the sparse wire distribution. Resistive loading mitigates the second effect (pattern distortion) at the expense of worsening the first (efficiency), presenting a clear trade-off. For NEMP simulation purposes, where temporal fidelity and pattern quality are important, the resistive loading approach yields a more predictable and cleaner radiation pattern, which may be preferable despite the lower efficiency. However, neither 8-wire model approaches the performance of the solid cone, indicating that the wire approximation is a significant source of degradation for a conical radiator operating over the 1–1000 MHz band. This observation is consistent with the general principle that wire approximations degrade as the inter-wire spacing becomes comparable to the wavelength.

3.4 Performance Estimation

With three models of the EMIS-III-VPD and a better understanding of their differences and ultra-wideband performance, it is now time to use the models to estimate the performance of the actual EMIS-III-VPD. The model of choice for this is the 8-wire model with end-termination, as this is as close to the real EMIS-III-VPD as possible. Despite this, other models will also be shown to further illustrate the differences between the models.

Firstly, the radiated frequency spectrum will be calculated and presented. As a starting point, the 500 kV Marx generator mentioned in IEC-61000-4-32 [10] is assumed to perfectly feed its potential difference across the antenna feeding gap. Any losses in the generator, connecting cables or feeding network are omitted, meaning that any results shown are optimistic. The antennas radiated power per frequency is multiplied by an averaged gain per frequency. The gain is averaged over the angle ϕ , meaning that the peaks and nulls caused by sparsity of the wires are omitted for this calculation, as that specifically will be inspected in more detail next. Only spreading losses are assumed between the antenna and a point 100 m away from the antenna. The result is Fourier transformed to give the radiated electric field strength per frequency over the desired band. The results are shown below in figure 3.17




Figure 3.17: Spectral content comparison of the three models across the desired band

Figure 3.17 on the previous page shows the spectral density of the three EMIS models across the desired frequency band. The theoretical DEXP at 5 kV/m is added as a dashed black line to show the goal of this thesis, as well as the accompanying 10 dB tolerance bound. What is clear is the lacking very low frequency performance of the EMIS-III-VPD: it is simply too small to efficiently radiate 1 MHz and therefore does not achieve the desired band. The low frequency performance is better for the solid cone than it is for the 8-wire models: the solid cone enters the tolerance band at 4 MHz while the 8-wire models do so at 6 MHz. From approximately 300 MHz onwards the 8-wire models have higher spectral density than the solid cone, as that is where they overcome the 100 - 300 MHz performance gap described earlier as well as edge scattering becoming more prominent for the solid cone model.

Over the whole band the EMIS-III-VPD is only barely within the ± 10 dB tolerance bound from the 5 kV/m theoretical limit. This is to be expected as the EMIS only actually achieved 2 kV/m (solid cone model) electrical field strength. It stands to reason that the low frequency performance is acceptable so long as the radiated electrical field strength is at least 5 kV/m: higher field strength across the whole band mitigates low frequency performance.

To visualise the spatial distribution of electric field strength across the testing surface, a heatmap was generated for the 8-wire with end termination model at several frequencies. The testing surface is defined as a vertical plane at $x = 100$ m from the antenna, spanning 200 m horizontally ($y = -100$ to $+100$ m) and 50 m vertically ($z = 0$ to 50 m), discretised into a grid of $201 \times 51 = 10,251$ points at 1 m spacing. This is as illustrated in figure 3.2. For each grid point, the distance r , polar angle θ and azimuth angle ϕ with respect to the antenna at the origin were computed from the Cartesian coordinates. The directivity at each (θ, ϕ) was obtained by two-dimensional linear interpolation over CST far-field data exported at seven elevation cuts ($\theta = 60^\circ$ to 90° in 5° steps), each containing 360 azimuth samples ($\phi = 0^\circ$ to 359°). The electric field magnitude at each grid point was then calculated as $|E(f)| = \sqrt{30 P_{\text{avail}}(f) G_{\text{real}}(f, \theta, \phi)} / r$, where the realised gain G_{real} incorporates the interpolated directivity, the radiation efficiency and the mismatch factor $(1 - |S_{11}|^2)$ at the frequency of interest. The $1/r$ dependence is evaluated per grid point, so the geometric spreading from the varying propagation distance across the testing surface is included. The result of this calculation is shown in figure 3.18 on the next page




Figure 3.18: Heatmap of EMIS-III-VPD radiation across 200x50 m vertical testing plane at 100 m distance

What figure 3.18 on the previous page clearly shows is the distinct lack of homogeneity across the 200x50 m testing area as defined in figure 3.2. This was already predicted in section 3.2, where a 3 dB path loss was predicted between the centre of the testing area and the edge due to the path length difference. This is reflected well in the 1 MHz plot, which shows the antenna behaving as an electrically small antenna.

However, what was unaccounted for in the first-order estimate was the effect of the sparsity of the 8-wire antenna. At frequencies where the antenna is not electrically small this effect causes large ripples of peaks and nulls across the testing area, with up to 12 dB differences over meters of distance separation. Inspection of all available frequency cuts shows that this effect is largest from 50 to 300 MHz, where the antenna is electrically large but the distance between the wires is roughly a wavelength long. This observation is consistent with what was observed earlier in this section.

The two other models were also observed in the same way. To compare their performance, the ratio between maximum and minimum electric field strength across the testing area is calculated and presented in dB in the table below:

Table 3.2: E-field homogeneity range across the 200 m $\times$ 50 m testing surface at 100 m distance

f [MHz]	End-terminated	Resistively loaded	Solid cone
1	4.0	4.0	4.0
5	3.5	3.0	2.9
10	10.6	10.1	6.9
20	9.7	9.2	6.2
50	17.9	9.6	5.6
75	13.9	13.3	5.4
100	20.2	11.9	3.5
150	16.8	13.2	3.6
200	17.1	7.3	3.7
250	10.7	14.8	4.0
300	7.8	5.9	3.8
350	7.7	7.0	4.1
400	6.6	5.9	3.8
450	8.1	6.1	3.9
500	7.3	6.0	4.2
550	7.1	5.3	3.9
600	10.2	5.4	4.6
650	11.6	7.4	7.1
700	8.9	9.5	7.6
750	7.5	14.4	7.8
800	6.1	10.2	9.4
850	8.3	10.4	8.8
900	9.0	9.3	9.7
950	9.3	8.4	9.4
1000	9.5	7.8	9.5
Average	9.9	8.6	5.8

What the table shows is the clear homogeneity improvement when a solid cone is used as opposed to the 8-wire models at lower frequencies. Only at very high frequencies does this difference become less pronounced, as from 600 MHz onwards the difference is no longer as large. It also shows that resistive loading generally speaking aids in improving homogeneity, with some exceptions. All in all, the average ratio between maximum and minimum electric field strength is 9.9 dB for the 8-wire model with end-termination, improved slightly to 8.6 dB with resistive loading and finally 5.8 dB for the solid cone model.

An additional point of view of the homogeneity can be visualised when considering the time difference of arrival across the testing area. In the first-order estimate done in section 3.2, it was estimated that the total time difference of arrival across the area would be 138 ns, which is significant for a pulse with a FWHM of just 23 ns. The pulse would have completely passed through the centre of the testing area while not yet having reached the edges. This is visualized in the plot below:




Figure 3.19: Time domain fidelity comparison

Figure 3.19 shows the E field intensity as a function of time at 4 different points on the testing surface as described in figure 3.2: all are at a height of 25 m above the horizontal plane, hence in the middle of the testing surface. They differ in being at different positions with respect to the boresight, taking steps of 25 m off boresight out towards the edge of the testing surface.

What can be seen clearly is the increasing time difference of arrival with equal step size off-boresight, a consequence of the circularly expanding wave from the omnidirectional radiator. Once the pulse at 50 m off boresight starts, the pulses that passed through boresight and 25 m off boresight have already decayed into late-time ringing. The 100 m off boresight starts almost exactly at 138 ns, as predicted in the first-order estimate.

The implication of this effect is that the different parts of the DUT are illuminated at different moments. Therefore, one cannot speak of 'total illumination' when discussing the EMIS-III-VPD, as the time difference of arrival of the pulse across the DUT is significantly longer than the pulse duration. Any effects that may be incurred when the full DUT is illuminated homogeneously are unlikely to occur when the time difference of arrival is as significant compared to the pulse-width.

What is also interesting to see is the difference in overall pulse shape between these pulses and the theoretical DEXP shape as described in chapter 2, specifically in figure 2.1 or in 2.3. This is an expected result that will be discussed in more detail on the next page, where time-domain fidelity in general is discussed.

Finally, the time-domain fidelity of the EMIS-III-VPD was inspected. A reconstruction of the time domain pulse was made from the spectral density of the radiated pulse. This is not ideal, as the antenna certainly radiates some energy outside of the 0.1 to 1000 MHz bound within which the radiated electromagnetic energy was evaluated. However, this loss is considered negligible. For this specific graphic, the DEXP pulse as described in chapter 2 was fed to the antenna. The results are displayed below:




Figure 3.20: Time domain fidelity comparison for EMIS-III-VPD

Figure 3.20 presents the reconstructed time-domain electric field at boresight on the test surface (100 m range, 25 m above the horizontal plane) for the three EMIS-III-VPD antenna models driven by a double-exponential (DEXP) voltage pulse. All three waveforms exhibit the expected bipolar signature of an impulse radiating antenna: a sharp positive peak corresponding to the fast rising edge of the feed pulse (where dV/dt is large and positive), followed by a zero crossing and a negative lobe produced by the slower exponential decay. The solid cone model, representing the theoretical ideal of a continuous conical conductor, achieves the highest peak field of approximately 1.7 kV/m. Its broad positive lobe and sustained negative tail reflect the wideband, low-loss nature of the solid structure, which radiates efficiently across the full spectrum of the DEXP pulse. However, the prolonged negative excursion—still at -0.4 kV/m at $t = 100$ ns—indicates that the unattenuated return current along the cone surface produces a significant late-time low-frequency contribution that would be undesirable in a practical EMP testing scenario. Further tuning of the solid cone model is necessary to fully eliminate this effect.

The end-terminated 8-wire model delivers a peak field of roughly 1.2 kV/m, approximately 70% of the ideal cone value. Its most notable feature is the rapid decay of the radiated pulse: the waveform settles to near zero within 25–30 ns, with negligible late-time ringing. This clean pulse shape is a direct consequence of the matched resistive termination at the wire ends, which absorbs the outward-travelling current wave as it reaches the rim of the antenna aperture, preventing reflections that would otherwise re-radiate and distort the pulse. The trade-off is clear: the termination dissipates a portion of the input energy, reducing the peak radiated field, but produces a compact, well-defined impulse that closely approximates the ideal monocycle waveform required for standardised EMP hardness testing.

The resistively loaded 8-wire model occupies a middle ground. Its initial peak is comparable to the end-terminated variant at approximately 1.2 kV/m, but the distributed resistive loading along the wire length is less effective at suppressing late-time oscillations. Residual ringing persists with an amplitude of 0.1–0.4 kV/m out to roughly 50–60 ns before decaying. This behaviour is consistent with the nature of distributed loading: while it progressively attenuates the current as it propagates along the wires, it does not present a single well-matched termination point, allowing partial reflections at each loading discontinuity. The result is a waveform that offers a compromise between the high peak field of the unterminated solid cone and the temporal cleanliness of the end-terminated design.

The main parameters of the time domain pulses are summed up in the table below:

Table 3.3: Timedomain fidelity metrics.

Model	Peak E [kV/m]	Rise time [ns]	FWHM [ns]	max dE/dt [kV/m/ns]
Solid cone (ideal)	1.96	5.3	19.4	7.25
End-terminated (8-wire)	1.26	2.45	10.2	32.40
Resistively loaded (8-wire)	1.29	2.55	9.8	37.65
Thesis requirement	5.0 (−0/ + 6 dB)	2.5 ± 0.5	23 ± 5	—

Note that in this table, the FWHM is measured for the envelope around the pulse, not the actual pulse itself. That way a unipolar pulse can be compared to a bipolar pulse.

3.5 Simple two-element monocone array

With the three models of the EMIS-III-VPD, it is tempting to do a first order validation of the idea that an array performs better than a single antenna. To analyse this, the radiated far fields of the EMIS-III-VPD from the 8-wire with end-termination model were coherently summed as if they originated from two separate positions. The model therefore includes no effects such as inter-element coupling, and it assumes perfect and identical feeding and timing. Two array configurations were considered: a broadside configuration with the two antennas placed at $y = \pm 50$ m, and a timed endfire configuration with both antennas on the x -axis separated by 40 m (respecting the 17.5 m cone base radius), where the rear antenna fires 133 ns earlier to compensate for its longer path to the centre of the testing surface. At each point on the testing surface, the complex electric field contributions from both antennas were summed as phasors, accounting for the propagation phase e^{-jkr} and the $1/r$ spreading from each antenna individually.

The broadside configuration is shown in figure 3.21 and the timed endfire configuration in figure 3.22 on the next two pages. Note that these images show E-field strength in kV/m, not in dB as in previously in figure 3.18.




Figure 3.21: Two-element broadside array (100 m separation) of the end-terminated 8-wire model




Figure 3.22: Two-element timed endfire array (40 m separation) of the end-terminated 8-wire model

At 1 MHz both configurations perform similarly. The wavelength of 300 m is much larger than the element separation in either case, so the two sources add nearly coherently across the entire testing surface. The peak field approximately doubles compared to the single antenna, rising from 1.28 kV/m to roughly 2.3 kV/m, consistent with the expected 6 dB gain from coherent addition of two identical sources minus a small correction for the increased propagation distance. The homogeneity at 1 MHz is 5.1 dB for both configurations, only marginally worse than the 4.0 dB of the single antenna.

At higher frequencies the two configurations diverge dramatically. The broadside configuration in figure 3.21 develops dense vertical interference fringes from 100 MHz onwards. This is a direct consequence of the two-source array factor: at 100 MHz the wavelength is 3 m and the 100 m separation corresponds to approximately 33 wavelengths. The path length difference from the two antennas varies rapidly across the horizontal extent of the testing surface, producing constructive and destructive interference with a fringe spacing of approximately 1.5 m. The resulting homogeneity degrades to 46 dB ratio at 100 MHz, far worse than the 20 dB of the single antenna. This pattern persists and intensifies at 200 and 500 MHz.

The timed endfire configuration in figure 3.22 shows a qualitatively different pattern. Because both antennas radiate along the same axis towards the testing surface, there are no lateral interference fringes. Instead, the interference pattern takes the form of concentric arcs centred on the boresight. This is because the time delay compensation is exact only at the centre of the testing surface; at other points the path length difference from the two antennas is not perfectly compensated, producing circular fringes whose spacing depends on frequency. At 100 MHz these arcs are broad and the central region of the testing surface shows improved field strength with reasonable uniformity. At 200 and 500 MHz the arcs become denser, although the pattern remains significantly cleaner than the broadside case.

Neither configuration solves the fundamental problem. The monocone is omnidirectional in the horizontal plane, meaning it radiates equal power in all directions regardless of array configuration. Any power not directed at the testing surface is wasted, and the array factor introduces interference patterns that worsen homogeneity at frequencies where the element spacing is larger than a wavelength. This is an inherent limitation of arraying omnidirectional elements for a wideband application: no single element spacing can avoid interference nulls across the full 1 MHz to 1 GHz band. At any frequency where the spacing exceeds approximately 0.7λ , nulls will appear within the angular extent of the testing surface. For 100 MHz this requires a spacing below 2.1 m, which is physically impossible when dealing with 35 m wide monocone antenna.

This is a fundamental problem with forming an array: the ultra-wideband nature of this application requires the antenna to be very large to accommodate the lowest frequencies, which in turn means the array spacing has a definite lower limit, meaning that the array will be sparse at the higher end of the frequency spectrum. A sparse array will by definition have an array factor with grating lobes [16], which will cause an omnidirectional antenna to 'leak' significant radiation into unwanted directions. This can be solved by using a directive antenna, as the total radiation pattern of an array is the sum of the radiation pattern of an element by its array factor (the radiation pattern of the array if all its elements were omnidirectional). Hence, a directive antenna may have nulls (or low radiation intensity) in the grating lobes of the array factor, hence mitigating this effect. This will improve homogeneity of an array.

The conclusion from this exercise is that simply duplicating the EMIS-III-VPD does not provide a viable path to improved NEMP hardness validation. While the peak field strength at the centre of the testing surface does increase, the spatial homogeneity degrades severely at the frequencies that matter most. What is needed is an antenna element with inherent directionality, so that the radiated energy is concentrated towards the testing surface rather than spread omnidirectionally. This is the motivation for the search for a directional antenna in the remainder of this thesis: its directional radiation pattern avoids the interference problem by ensuring that each element contributes primarily to the desired testing area, while the array timing ensures coherent addition across the testing surface. Furthermore, it is likely impossible to design a practically viable antenna that can radiate the desired electric field over an area of 200 m by 50 m. Shrinking test area size is necessary so long as practical viability remains desirable.

3.6 Challenges of Previous NEMP Hardness Validation Tests

In this section the problems encountered during NEMP hardness validation testing by the RNLN will be highlighted to give insight into where improvements may be made. This is not just with respect to the parameters of the electric far field, but also in for instance the practicality of the test. This section is based on expert opinions and recounts of the original tests using the EMIS-III-VPD in the 90s and 00s.

Most obviously, the EMIS-III-VPD radiates omnidirectionally. At the time, tests were conducted in the middle of the naval base, although seemingly no tertiary parties in the naval base were disturbed by the tests of the EMIS-III-VPD. This speaks to the low threat level of the antenna, but it also highlights the fact that if higher threat levels were to be achieved more measures have to be taken to not interfere others. The directivity of a new system aids in this, as the new system may be placed in close proximity to others without disturbing them so long as it is aimed away carefully. Despite this, a new location is likely needed, and therefore the system must be movable and deployable at another location without becoming too impractical.

Secondly, the EMIS-III-VPD used SF-6 gas at the feeding point and in the Marx generator in order to increase the voltage breakdown limit of the gas surrounding it. Besides being a highly effective dielectric, SF-6 is also 25000 times more potent as a greenhouse gas than an equivalent volume of CO₂. Hence, the European Union has banned its use in high voltage electronics if a proper recapturing mechanism is not installed [17], and hence a new solution must be found. While recapturing SF-6 is possible, it is unlikely to be the path of least resistance, and hence it is more fruitful to look for a solution that does not require a dielectric of this nature to be employed at the feed of an antenna. The use of an array may aid by spreading the peak electric field strength across several radiators.

The EMIS-III-VPD was not constructionally perfect. The supporting construction has to be electromagnetically invisible in the relevant spectrum as to not disturb the radiation pattern, and was hence likely made out of some form of plastic or fibreglass. This structure may have been impractical to assemble and disassemble, potentially being dangerous. Specifically, the supporting pillars had to be transportable and were hence made up of segments, which allegedly were difficult to assemble and disassemble. It is also said that after an assembly / disassembly cycle the pillars were deformed and therefore had a limited lifespan.

As the name alludes, the EMIS-III-VPD only radiates a vertically polarized waveform. As already stated in the previous chapter, an ideal NEMP test radiates both horizontally and vertically polarized pulses in order to fully chart any weaknesses in a DUT. This shortcoming is noted.

3.7 Discussion

There is a fundamental problem with the EMIS-III-VPD: its physical size prevents it from efficiently radiating electromagnetic energy at frequencies as low as 1 MHz. This was clear from figure 3.17, which showed the wire model to not come into the 10 dB bound until 6 MHz. Note that this is for the 5 kV/m limit, a tenth of the actual NEMP standard as prescribed in IEC-61000-2-9 [18]. Simply increasing the power is no solution, as that will push the potential difference across the feeding gap beyond the voltage breakdown limit. Using more than 1 EMIS-III-VPD antenna, practical considerations aside, only raises the maximum output power by approximately 3 dB, at the cost of considerable worsening of the homogeneity. A trade-off exists between the low limit of radiated frequency and the physical size of the antenna.

To gain some feel for this shortcoming, a modified version of the solid cone model was made: it was shrunk by a factor 10, making the model 2 m tall with a diameter of 3.5 m. Apart from resizing, no other changes are made to the model; feeding voltage remains identical. Establishing at what frequency this model would enter the region within the 10 dB tolerance of the DEXP allows more accurate determining a ratio between antenna diameter and lowest efficient radiation. The results for this are shown below, which can be compared to the full-size solid cone model.




Figure 3.23: E-field spectral density at 100 m for the full-size (20 m) and shrunk (2 m) solid cone models.

Figure 3.23 shows that the shrunk model enters the -10 dB tolerance band at a frequency approximately ten times higher than the full-size model, consistent with the factor-ten reduction in size. The frequency at which each model crosses the -10 dB boundary was determined by interpolation, and the corresponding electrical size D/λ at that crossing was computed. Averaging the two values yields $D/\lambda = 0.2873$ as the electrical size at which a conical radiator begins to produce spectral content within the -10 dB tolerance of the DEXP. This relationship below holds for specifically the solid cone model. However, other antenna will not deviate from this significantly, therefore this relationship can be used to estimate the antenna size needed.

$$f_{\min} = \frac{D/\lambda \cdot c}{D} = \frac{0.2873 \cdot c}{D} \quad (3.17)$$

where c is the speed of light. This relationship is plotted in figure 3.24 for antenna heights from 2.5 to 30 m.




Figure 3.24: Required antenna height as a function of the target lower frequency bound.

Figure 3.24 illustrates the trade-off between antenna size and low-frequency coverage. A 5 m height would only reach down to approximately 17 MHz, while a 10 m antenna extends this to roughly 9 MHz. To reach the 1 MHz lower bound specified in Table 3.4, an antenna height of approximately 83 m would be required, which is clearly not practically achievable. Even the 20 m height of the solid cone model only supports frequencies down to approximately 4 MHz within the -10 dB tolerance, which was reduced to 6 MHz in the case of the more realistic wire models of the EMIS-III-VPD, and even this was an optimistic estimate with a perfect feeding pulse generator.

The implication is that a design choice must be made to either accept a higher lower frequency bound than the 1 MHz specified in Chapter 2 for the sake of practical implementability, or to accept the impracticality of the design. This choice cannot be made by the author of this thesis: it is up to a Ministry of Defence to what extent a design can be considered feasible, specifically regarding its cost and size, and to what extent missing low frequency components can be considered acceptable. This work may however serve to illuminate the trade-off and present the different options for mitigation of a higher low frequency bound.

Besides the trade off between low frequency limit and antenna size, it must be accepted that a testing area of 200 x 50 m is unfeasible for the desired waveform as it is defined in Chapter 2. This was already stated in Chapter 2, but now this chapter motivates that choice. The trade-off is that the entire DUT is no longer illuminated at once, instead only a portion is. This means that any effects that occur when the entire DUT is illuminated by an NEMP are no longer simulated. The exact effect of this simplification is unknown, and outside of the scope of this thesis. However, the entire DUT was also not illuminated at once with the baseline method, and the illumination of the DUT at threat level over a limited but known area may be considered an improvement over illuminating below threat level inhomogeneously.

3.8 Conclusion

This chapter aimed to qualify and quantify the performance of the EMIS-III-VPD in order to establish a baseline. This baseline is the reference to which any work done in this thesis will be compared. To establish such a baseline, a CST model was made of the EMIS-III-VPD, which was verified with crude estimations of the antenna performance. Furthermore, some experiences of operating this simulator were gathered in order to gain some understanding in what challenges were encountered when operating this NEMP simulator.

The 8-wire model with end-termination is deemed as the most realistic antenna model for the EMIS-III-VPD, as there is no public source confirming the use of resistive loading. This model achieved 1.28 kV/m at 100 meters distance as peak E-field strength, which is an optimistic estimate as it was assumed that the 500 kV Marx generator was able to achieve that 500 kV across the feed gap of the antenna perfectly. This meant that the spectral density was below requirement up to 6 MHz, as well as falling out of the tolerance band in the 100 to 300 MHz gap in performance found due to ohmic losses in that band. The model showed an average ratio of maximum to minimum electric field strength of 9.9 dB across the 200 x 50 meter testing area. The homogeneity was at its worst between 50 and 300 MHz, as that is where the distance between wires is approximately one wavelength. The antenna performance summarized in a new column added to table 2.3, which gives the table below:

Table 3.4: NEMP Test Specification Requirements and estimated EMIS-III-VPD performance

Specification	Requirement	EMIS-III-VPD
Peak electric field strength	5 kV/m $-0 + 6$ dB	1.28 kV/m
Rise time (10%–90%)	2.5 ± 0.5 ns	2.45 ns
Full width at half maximum	23 ± 5 ns	24.1 ns
Lower frequency bound	1 MHz	~6 MHz
Upper frequency bound	1 GHz	~1 GHz
Spectral magnitude tolerance	± 10 dB	Exceeded 50-300 MHz
Polarisation	Horizontal and Vertical	Vertical only
Testing surface area	10 m $\times$ 10 m	200 m $\times$ 50 m
Test distance from radiator	100 m	100 m
Field homogeneity	$-0 + 6$ dB	9.9 dB avg.

This chapter has also proven that resistive loading of the wires of the EMIS-III-VPD may have improved its performance, albeit at the cost of increased design complexity and increased thermal loading of the wires, potentially reducing the antenna practicality. It was also shown that an 8-wire model to approximate a cone over the chosen bandwidth was sparse and introduced significant artefacts. Increasing wire density is likely to have alleviated this.

The EMIS-III-VPD design was impractical, using SF6 gas and being difficult to transport and set-up at a new location. Furthermore, its omnidirectionality makes selection of a testing location challenging. A more practical design is desired.

Now, the work will aim to improve this performance. This chapter identified two main trade-offs in the design of a radiated NEMP simulator for large DUT's. One is the size of the testing area, which this thesis keeps fixed at 10 by 10 meters. The trade-off here is between homogeneity and peak-electric field strength and the portion of the DUT tested at any given moment. This chapter has shown that the EMIS-III-VPD covered neither of these options; it did not illuminate at the desired peak-electric field strength nor did it illuminate the entire DUT at once. The other trade-off is the practicality (specifically size and cost) of the antenna and the lower frequency limit, for which this thesis will provide insight and potential mitigation of shortcomings. Using an array of antennas is useful for improving field strength (and hence low frequency limit) and homogeneity over a limited area. Further improvement may be found by using a directed antenna.

Chapter 4

Antenna Selection

The first step in designing a the time domain array is to select an antenna concept. Only a limited number of concepts are able to generate a NEMP-like pulse. The main challenge associated with replicating the NEMP standard impulse is the highly broadband nature of the pulse. Normally antennas are designed with a narrowband operating frequency in mind. Canonical examples of wideband antenna concepts include for instance log periodic antennas, commonly used in television antennas. These are however dispersive; high frequencies arrive at an observation point at a different time than lower frequencies, hence distorting the pulse in time domain. This is caused by reverberation, and an antenna that makes use of reverberation is not well suited when optimizing for pulse fidelity. Furthermore, the extreme field strength requirement of the testing area makes the design of such an antenna non-trivial. This combination of requirements gives rise to the consideration of the following antenna concepts: the Biconical Antenna, the vivaldi Antenna, the TEM horn, the Reflector IRA and the Lens IRA [60]. Literally, any of these antenna may be qualified as an impulse radiating antenna (IRA) in the sense that they all radiate an impulse, but in literature this is not how these antenna are described.

In this section, the possible choices will be highlighted, then compared and contrasted and at the end a choice will be made on which concept will serve as point of departure for the rest of this thesis. This will be done by shortly describing the working principle and the main antenna design trade-offs. A first order estimation of the performance of the antenna will be made for each antenna, which will give some insight into the viability of the concept. Furthermore, as all antennas are of the travelling-wave type, they can be approximated with a transmission-line model. This model is first set up, and then applied to each antenna. This gives further insight into the upper and lower frequency bounds, as well as shedding light on the nature of frequency dependence of the input impedance. Only if the antenna input impedance is (relatively) frequency independent over the desired frequency bound can it be considered suitable for this project. The primary criterion on which antennas will finally be assessed is the ability to reconstruct the standard NEMP waveform with the highest level of fidelity, but also the practical feasibility and the availability of an existing design will be taken into consideration. For reading ease the antenna will be presented in order of increasing directivity, as it will turn out that this is the defining feature of these antenna.

4.1 Transmission-line model for travelling-wave radiators

All antenna concepts considered in this chapter can be interpreted, at least to first order, as travelling-wave radiators. In that sense they resemble transmission lines that guide a pulse from a feed region towards an aperture, while gradually converting the guided pulse into a radiated pulse. The main difference between concepts is the manner in which the effective transmission-line impedance evolves from the feed to the aperture, and how effectively the structure behaves as a matched radiative termination. This subsection introduces a common framework to assess whether a given concept can provide a relatively frequency-independent input impedance over the required bandwidth.

In the idealized picture, the antenna supports a dominant quasi-TEM mode over most of its length, such that a local characteristic impedance $Z_c(z)$ may be defined along a coordinate z from feed to aperture.

For self-similar structures this impedance is primarily geometry-controlled and therefore only weakly dependent on frequency within the quasi-TEM regime. Therefore, when the upper and lower bounds of the frequency operating regime are set, the frequency dependence of local impedance may be set by the geometry of the antenna [71].

Even when $Z_c(z)$ is approximately frequency independent, the measurable input impedance $Z_{\text{in}}(f)$ generally is not. Any finite travelling-wave radiator is truncated: the wave launched at the feed encounters an end region where part of the energy is radiated and part is reflected back towards the source (or absorbed in a terminating resistance). A useful first-order representation is to model the antenna as a transmission line with an effective characteristic impedance Z_0 and an effective termination reflection coefficient $\Gamma_L(f)$ located a distance L from the feed. This gives the standard input-impedance relation

$$Z_{\text{in}}(f) = Z_0 \frac{1 + \Gamma_L(f) e^{-2j\beta L}}{1 - \Gamma_L(f) e^{-2j\beta L}}, \quad \beta = \frac{2\pi f}{c}. \quad (4.1)$$

When $|\Gamma_L(f)|$ is not small, the input impedance exhibits standing-wave ripple. The characteristic frequency spacing of this ripple scales as

$$\Delta f \sim \frac{c}{2L}. \quad (4.2)$$

Hence, obtaining a relatively flat input impedance requires not only an approximately constant Z_0 , but also a termination behaviour that keeps $|\Gamma_L(f)|$ small across the desired band.

Two practical bounds limit the frequency range over which a travelling-wave radiator can exhibit quasi frequency-independent impedance. At the low-frequency end, the structure becomes electrically small, radiation resistance decreases, and $Z_{\text{in}}(f)$ becomes strongly reactive. An order-of-magnitude criterion for leaving this regime is that the relevant electrical length is at least a quarter wavelength:

$$L \gtrsim \frac{\lambda}{4} \Rightarrow f \gtrsim f_{\text{low}} \sim \frac{c}{4L}. \quad (4.3)$$

At the high-frequency end, the quasi-TEM assumption fails when the feed region and smallest geometric features are no longer electrically small, causing higher-order mode content, localized resonances and reflections that perturb $Z_{\text{in}}(f)$. A practical criterion is

$$\ell_{\text{min}} \lesssim \frac{\lambda}{20} \Rightarrow f \lesssim f_{\text{high}} \sim \frac{c}{20 \ell_{\text{min}}}, \quad (4.4)$$

where ℓ_{min} denotes the smallest relevant dimension in the feed/transition region.

Within this framework, a travelling-wave antenna can be said to have a relatively frequency-independent input impedance over a band if, simultaneously:

1. The dominant field behaviour remains quasi-TEM, such that $Z_c(z)$ (and thus a representative Z_0) is largely geometry-controlled over the band;
2. The structure is electrically long enough at the low-frequency end, $f \gtrsim c/(4L)$;
3. The feed/transition region remains electrically small at the high-frequency end, $f \lesssim c/(20\ell_{\text{min}})$;
4. The effective termination is sufficiently absorptive/radiative such that $|\Gamma_L(f)| \ll 1$ over the band, limiting standing-wave ripple in $Z_{\text{in}}(f)$.

This immediately shows why extremely wide bandwidth requirements are challenging: a finite structure must simultaneously be physically large enough to avoid the electrically-small regime at low frequencies while maintaining mm - cm scale tolerances in the feed region for GHz radiation. In the following antenna specific subsections, this common framework is used to identify the dominant mechanism that limits the flat-impedance region for each concept.

4.2 Biconical Antenna

The biconical antenna is the oldest antenna concept that is up for consideration. It was invented by Sir Arthur Lodge in 1898 [56], and has since seen application in UWB communication technology and immunity testing. Furthermore, the concept has found application in numerous NEMP simulators. One of these is the hybrid NEMP simulator, as shown in figure 1.1b, in which the bicone in the centre radiates the higher frequency components and the resistively loaded wires attached to the cone radiate the lower frequency components. Another is the EMIS-III-VPD, the baseline for this thesis, and EMPRESS II, which both were specifically designed for NEMP hardness validation testing of warships [11]. A difference to be noted is that these concepts used a monocone on top of an electrically conducting plane that mirrored the other half of the cone.

4.2.1 Antenna principle

The biconical antenna is named after its basic structure: two conducting cones that are placed such that their tips point towards each other. Between these tips a stepped voltage source feeds the desired pulse onto the conducting cones. The continuously expanding cross-section of the cones act like a tapered transmission line, in which the impedance increases with cross-section. This results in a varying electric current, which causes broadband electromagnetic radiation [58].




Figure 4.1: Biconical antenna schematic

The general schematic of the biconical antenna is shown in figure 4.1. The antenna uses a travelling wave structure that resembles that of a transmission line. It also is practically speaking an impedance converter, that due to the angle of the cone converts from the feed impedance to the free space impedance. This specific schematic shows what is simply put a broadband dipole concept. A significant property of this antenna is that it is omnidirectional in one plane. This was also discussed in the previous section as being a design obstacle for the purpose of NEMP hardness validation tests.

4.2.2 Design considerations

It is tempting to consider taking half of the biconical antenna when designing a NEMP hardness test for ships. The sea surface can be assumed to behave like a perfect electrical conductor for the frequency band discussed, hence requiring only half of the antenna to be built as the other half will exist virtually in the mirror caused by the electrically conducting surface.

The lower limit of operating frequency for a biconical antenna is the size of the dipole-like waveguide, which is what determines the maximum wavelength the dipole can support. This comes at the cost of physical size, which may be tens of meters tall if low frequencies are desired. This may be mitigated by the use of wire or mesh guides, as shown in figure 4.1, at the cost of high frequency fidelity [59]. The upper operating frequency limit is determined by the feed region. At wavelengths where the feed region is electrically large the structure will no longer support the quasi TEM waveform between the two cones due to localized resonances and reflections. However, having a low upper bound on frequency gives rise to more tolerant design and fabrication criteria for the feeding region.

Biconical antenna do not necessarily have a terminating structure as depicted in figure 4.1, where the metallic lines starting at the stepped voltage source bend back into a terminating structure. Instead,

the cones may be terminated at that bending point and connected by a large ring. This makes the construction of such an antenna more practical, at the cost of more internal reflections and post pulse ringing [58]. The diameter of this ring is considered the diameter of the cone, which is decided by the required clearing time of the wave. In turn, the clearing time is dependent on the rise time of the pulse, as it must be longer than the rise time [59]. Hence, for a desired pulse, the desired clearing time can be determined, and hence the diameter of the cone. A larger diameter will necessitate a larger framework around the antenna to support the structure, giving rise to diffraction and reflection on and around the supporting structure.

4.2.3 First order estimations

It is possible to make some first order estimation of the minimum required size of the antenna based on the design considerations. Furthermore, to observe the effect of increasing directivity across this chapter, the feed-voltage required can be estimated by estimating the gain of the antenna. The order-of-magnitude calculation will speak to the feasibility of the design. As most of these calculations have already been done in detail for the EMIS-III-VPD in Chapter 3, this subsection will quickly present the results.

The lowest efficiently radiated frequency f_L is set by the largest dimension of the structure: $L \approx (0.25-0.5) \lambda_L$, where λ_L is the wavelength of the lowest operating frequency. For a biconical this yields $L \approx 75-150$ m, which is halved when the configuration of a moncone above an electrically conducting surface is used.

The upper frequency bound f_H is typically constrained by the electrical size of the feed region, which was calculated in chapter 3 to be equal to $\ell_{\text{feed}} \lesssim 15$ mm. The UWB nature of the antenna yields a tens-of-meters tall structure to have mm constraints.

Next, the power required at the feed-point can be estimated, as was also done for the EMIS-III-VPD. In the calculation done in chapter 3, the gain was estimated as $G \approx 1.6$ (≈ 2.1 dBi), which is doubled for the moncone over electrically conducting surface configuration. It is also independent of size of the antenna, so the feed power calculated in chapter 3 can be linearly scaled to generate the required 50 kV/m at the testing surface. 50 kV/m is used in this part of first order approximations despite the fact that the goal of the thesis is 5 kV/m as eventually 50 kV/m is what is required for a reliable NEMP hardness validation. Scaling the result from chapter 3 to generate 50 kV/m at 100 m from the antenna gives a peak feeding voltage equal to ~ 7.8 MV, which when supported across the 15 mm feed gives an electric field strength at the feed point equal to ~ 570 MV/m. This is so significantly larger than the air breakdown limit of 3 MV/m that it is unlikely a medium can be found within which this voltage will not arc.

4.2.4 Application of the transmission-line model

The biconical antenna is the most direct example of a travelling-wave radiator that supports the transmission-line interpretation. The two cones form a flared two-conductor structure that can support a quasi-TEM wave launched from the feed. The four criteria of the framework above can therefore be evaluated explicitly for this concept.

Within the quasi-TEM regime, the biconical is approximately self-similar and its effective line impedance is largely set by cone angle rather than frequency. D. Giri [58] used the following first-order expression for the characteristic impedance of an infinite biconical:

$$Z_0 \approx \frac{\eta_0}{\pi} \ln\left(\cot(\alpha/4)\right), \quad (4.5)$$

with $\eta_0 \approx 377 \Omega$ and α the full cone angle. This highlights that, in principle, a biconical can be designed to present a near-constant impedance over a broad frequency region by selecting the cone flare.

The lower bound of the "flat" impedance region is set by electrical length. When the structure is electrically small, radiation resistance decreases rapidly and the input impedance becomes strongly reactive.

Using the quarter-wavelength criterion of the general framework yields

$$f_{\text{low}} \sim \frac{c}{4L}, \quad (4.6)$$

where L denotes the relevant one-way path length from feed to the termination region. For a full biconical, L is comparable to half the tip-to-tip length; for a monocone-over-PEC, the physical height h corresponds approximately to $L \approx 2h$ in terms of effective dipole length. This criterion therefore predicts that achieving efficient radiation and a relatively stable input impedance down to 1 MHz requires L on the order of 75 m. Shorter structures shift the onset of efficient radiation to several MHz, consistent with the first-order geometric sizing arguments used earlier.

At the high-frequency end, deviations from quasi-TEM behaviour occur once the feed gap and smallest geometric features become electrically large. Applying the framework criterion gives

$$f_{\text{high}} \sim \frac{c}{20 \ell_{\text{min}}}, \quad (4.7)$$

where ℓ_{min} represents the smallest relevant feature in the feed/launch region. Supporting spectral content up to 1 GHz therefore implies mm - cm scale feed tolerances. In addition to strict dimensional constraints, localized resonances in the feed region and excitation of higher-order modes can introduce frequency-dependent perturbations to $Z_{\text{in}}(f)$ and generate post-pulse ringing in the time domain.

Even when the local impedance is geometry-controlled, the measured input impedance of a finite biconical depends on the effective termination at the end of the finite cone. Using the transmission-line representation,

$$Z_{\text{in}}(f) = Z_0 \frac{1 + \Gamma_L(f) e^{-2j\beta L}}{1 - \Gamma_L(f) e^{-2j\beta L}}, \quad \beta = \frac{2\pi f}{c}, \quad (4.8)$$

it follows that a small termination reflection coefficient $|\Gamma_L(f)| \ll 1$ is required for a relatively flat input impedance. In many practical biconicals the end region is not a matched termination but a geometric truncation, such as with the EMIS-III-VPD which had a ring connecting the ends of the cone. This must lead to imperfect matching, hence reflections and thus impedance ripple. The ripple spacing is on the order of

$$\Delta f \sim \frac{c}{2L}, \quad (4.9)$$

so longer structures exhibit denser ripple in frequency unless reflections are strongly suppressed. A resistive termination or resistive loading can reduce $|\Gamma_L(f)|$ and improve both impedance flatness and time-domain ringing behaviour, at the expense of radiated efficiency and/or high-frequency fidelity depending on implementation [58].

The transmission-line framework therefore leads to an assessment for the biconical. The concept can satisfy the first criterion (geometry-controlled impedance) over a broad region, and it can be designed to have a near-constant line impedance by carefully selecting the cone angle. However, covering the full 1 MHz - 1 GHz requirement within a frequency interval of relatively frequency-independent input impedance would require, simultaneously: a structure of $\mathcal{O}(10^2)$ m scale to avoid the electrically-small regime at the low-frequency end, and mm - cm feed-region tolerances to preserve quasi-TEM behaviour at the high-frequency end, and an effective termination with sufficiently small $|\Gamma_L(f)|$ over ten decades of frequency.

4.3 TEM Horn Antenna

The largest drawback of the biconical antenna was the total lack of directionality in the horizontal plane. Introducing some form of directionality will improve the performance of the antenna and reduce the feed power needed, as well as reducing the electromagnetic interference (EMI) in the region. Therefore, the next antenna under consideration is the TEM horn antenna. The TEM horn is one of the earliest UWB antenna, with significant literature coverage on its design and application, although a significant portion of that literature is for generic UWB purposes and not specifically on NEMP simulation [49]. The TEM horn has found extensive use in communication and radar engineering, as well as general EMC hardness testing [50].

4.3.1 Antenna principle

Electrically, the TEM horn antenna consists of two metallic plates that flare outwards to form a gradually expanding transmission line. Structurally there usually are two more metallic plates that connect the two previous metallic plates, hence forming a horn. This is omitted from figure 4.2 for clarity. The pulse, which is fed in TEM mode into this transmission line, will start to bulge outwards as the aperture widens. At the end of the antenna, the TEM wave is launched as a spherically expanding waveform. Fundamentally, this concept can be seen as an impedance transformer, where the metallic plates gradually allow the wave to change from the feed impedance to the free-space impedance. The antenna is visualized in the schematic shown in figure 4.2 below:




Figure 4.2: TEM horn antenna schematic

A stepped voltage source will feed a desired waveform in TEM onto the metallic plates. The plates guide the waveform into a spherically expanding radiating wave. The TEM horn can be modelled by treating each incremental length of the feed structure as generating electric and magnetic dipoles whose combined radiation forms a far field [51]. The summation of these dipoles eventually gives an analytic description of the electric far field.

The radiated field on boresight is proportional to the input voltage convolved with the horns impulse response. The horns impulse response contains two main parts: a narrow early time component and a wider late time component. The narrow early time component arises from the initial feeding of the waveform onto the metallic plates, which partly is immediately radiated. The broader, later time component arises from the actual radiation from the spherically expanding TEM waveform between the metallic plates [48]. Together the two components make up the main impulse.

The TEM horn preserves the temporal shape within the operating band of the horn, hence making it well suited for the purposes of this thesis.

4.3.2 Design considerations

The TEM horn is constrained in bandwidth. The lowest frequency the TEM horn can still radiate efficiently is determined by the ratio of the wavelength to the aperture size. Once the wavelength is significantly longer than the aperture height, the antenna will behave like a poorly matched, electrically short dipole [49]. The upper bound however is primarily constrained by the excitation of higher order modes. The TEM mode knows no cut-off frequency, but higher order modes do and will start to radiate once this frequency is exceeded. Reducing the smallest distance between the two metallic plates at the source region will increase the upper operating limits [50].

The angle between the two metallic plates is another design parameter to consider. The flaring angle determines the speed at which the antenna transforms the feed impedance into the free-space impedance, where a smaller angle causes a slower impedance change. Slower impedance change does mean better matching, lower dispersion [51] and better temporal fidelity [50], but at the cost of a larger and bulkier antenna. If the length of the antenna is fixed, then a lower flare angle will cause the operating frequency to shift upwards. Furthermore, in that case the flare angle would determine the size of the aperture, which determines the gain.

A different shape of the metallic plates may be chosen than the one visualized in figure 4.2, which is linearly expanding. A common variation is the use of an exponentially expanding horn. This results in smoother impedance transformation, which leads to lower internal reflections, improved pulse fidelity and a broader bandwidth, at the cost of more complex design and fabrication and a larger physical size [52]. Generally, the shape of the metallic plates may be adjusted to optimize for the desired waveform.

4.3.3 First order estimations

As with the biconical antenna, some first order estimations are made in order to quantify some of the design considerations at play. This will again give insight into the feasibility of the antenna design. The starting point is the well-known size - frequency relation: the lowest efficiently radiated frequency f_L is largely set by the largest dimension of the radiating structure. For a travelling-wave TEM horn, this is governed by the aperture dimensions. For a first estimate, the aperture should be electrically large enough to radiate efficiently, which may once again be taken as a fraction of a wavelength:

$$D_{\text{ap}} \sim (0.25\text{--}0.5) \lambda_L, \quad \lambda_L = \frac{c}{f_L}, \quad (4.10)$$

where D_{ap} denotes the largest linear dimension of the horn aperture (e.g. height or width). With $f_L = 1$ MHz, $\lambda_L \approx 300$ m, which yields

$$D_{\text{ap}} \sim 75\text{--}150 \text{ m}. \quad (4.11)$$

This immediately shows that a TEM horn that is intended to efficiently radiate spectral content down to 1 MHz would require aperture dimensions on the order of $\mathcal{O}(10^2)$ m. Smaller antenna no longer effectively radiate the very low frequencies.

At the opposite end of the bandwidth, the upper frequency bound f_H is primarily constrained by the electrical size of the feed region and by deviations from the quasi-TEM field distribution. To maintain quasi-TEM behaviour and to avoid strong localized resonances at f_H , the smallest feed-region features should remain electrically small:

$$\ell_{\text{feed}} \lesssim \frac{\lambda_H}{20}, \quad \lambda_H = \frac{c}{f_H}. \quad (4.12)$$

With $f_H = 1$ GHz, $\lambda_H \approx 0.30$ m and thus $\ell_{\text{feed}} \lesssim 15$ mm. Similar to the biconical concept, a structure that becomes tens of meters in size to support the low-frequency content simultaneously demands mm-scale constraints in the feed region to preserve GHz fidelity.

With the geometric bounds established, the next step is to estimate the power required at the feed-point for a specified far-field requirement. Assuming far-field conditions and a directive gain G in the direction of interest, the RMS electric field magnitude at range r relates to radiated power as

$$E(r) = \frac{\sqrt{30 P_{\text{rad}} G}}{r} \Rightarrow P_{\text{rad}} = \frac{(E(r) r)^2}{30 G}. \quad (4.13)$$

Unlike the biconical, a TEM horn exhibits increasing gain with frequency as the aperture becomes electrically larger. A convenient first-order estimate for the gain of an aperture-type radiator is

$$G(f) \approx \eta_{\text{ap}} \frac{4\pi A_{\text{ap}}}{\lambda^2} = \eta_{\text{ap}} \frac{4\pi A_{\text{ap}} f^2}{c^2}, \quad (4.14)$$

where A_{ap} is the horn aperture area and η_{ap} is an aperture efficiency factor (typically $\eta_{\text{ap}} < 1$). This expression makes explicit what was argued qualitatively earlier: higher directivity is obtained either by increasing aperture size or by operating at higher frequency content.

To convert radiated power into a feed-voltage estimate, the feed is approximated as being dominated by a real input resistance R_{in} :

$$P_{\text{rad}} \approx \frac{V_{\text{rms}}^2}{R_{\text{in}}} \Rightarrow V_{\text{rms}} \approx \sqrt{P_{\text{rad}} R_{\text{in}}}, \quad V_{\text{pk}} \approx \sqrt{2} V_{\text{rms}}. \quad (4.15)$$

For a practical TEM horn feed, R_{in} is commonly designed around standard generator impedances (e.g. 50Ω), hence $R_{\text{in}} = 50 \Omega$ is used as a representative value in the example below.

Combining the above expressions yields a compact first-order relation between the required field at 100 m and the required peak feed voltage:

$$V_{\text{pk}} \approx \sqrt{2 R_{\text{in}}} \frac{E_{100} r}{\sqrt{30 G}}, \quad r = 100 \text{ m}, \quad (4.16)$$

where $E_{100} = E(r = 100 \text{ m})$ and G may be taken as a representative value for the relevant frequency content. With the scenario of operating above an electrically conducting plane and the large wavelengths involved, it is tempting to select half a TEM-horn for this first order approximation. In that case, the radiation is confined to half-space, which to first order doubles the gain relative to the same aperture in free-space. To explore the limiting case where the horn is sized to be marginally efficient at 1 MHz, the virtual aperture is taken as $D_{\text{ap}} = 75 \text{ m}$ (i.e. $\sim 0.25\lambda$ at 1 MHz). Approximating a square virtual aperture gives $A_{\text{ap}} \approx 75 \times 75 = 5625 \text{ m}^2$. With $\eta_{\text{ap}} = 0.6$ and a representative frequency $f = 100 \text{ MHz}$ ($\lambda \approx 3 \text{ m}$), the free-space aperture estimate yields

$$G \approx \eta_{\text{ap}} \frac{4\pi A_{\text{ap}}}{\lambda^2} \approx 0.6 \frac{4\pi \cdot 5625}{3^2} \approx 4.7 \times 10^3, \quad (4.17)$$

This becomes $G_{\text{half}} \approx 2G \approx 9.4 \times 10^3$ for the half horn over an electrically conducting surface. Using $E_{100} = 50 \text{ kV/m}$, $r = 100 \text{ m}$ and G_{half} gives

$$P_{\text{rad}} = \frac{(E_{100} r)^2}{30 G_{\text{half}}} \approx \frac{(5 \times 10^6)^2}{30 \cdot 9.4 \times 10^3} \approx 8.9 \times 10^7 \text{ W}, \quad (4.18)$$

and with $R_{\text{in}} = 50 \Omega$,

$$V_{\text{pk}} \approx \sqrt{2 P_{\text{rad}} R_{\text{in}}} \approx \sqrt{2 \cdot 8.9 \times 10^7 \cdot 50} \approx 9.4 \times 10^4 \text{ V} \sim 94 \text{ kV}. \quad (4.19)$$

Across a 15 mm feed gap this implies

$$E_{\text{feed}} \approx \frac{V}{d} \approx \frac{9.4 \times 10^4}{0.015} \approx 6.3 \text{ MV/m}, \quad (4.20)$$

which remains above the $\sim 3 \text{ MV/m}$ air breakdown limit. Conversely, if the feed is limited to $E_{\text{feed}} = 3 \text{ MV/m}$ ($V_{\text{pk}} = 45 \text{ kV}$), then the achievable field at 100 m becomes $E_{100} \approx 24 \text{ kV/m}$. With some tweaking of the antenna or the feeding point, for instance by the submersion of an electronegative substance, this design is physically possible.

4.3.4 Application of the transmission-line model

A TEM horn can also be interpreted as a travelling-wave radiator: the feed launches a quasi-TEM wave between conducting plates which propagates towards an aperture while the field expands and progressively couples into free space. Compared to the biconical, the horn offers a more controlled aperture and higher directivity, but the same transmission-line considerations still determine whether the input impedance can remain relatively frequency independent over the required bandwidth.

In the idealized model, the horn supports a dominant quasi-TEM mode from the feed to the aperture, such that an effective characteristic impedance $Z_c(z)$ can be associated with the local plate spacing and conductor geometry. If the flare is smooth and self-similar, $Z_c(z)$ is geometry-controlled and only weakly dependent on frequency within the quasi-TEM regime. In practical designs the feed transition is commonly engineered to present a standard generator impedance, and the flare is chosen to avoid abrupt impedance steps that would generate reflections [14]. Hence, the TEM horn can often satisfy the first framework criterion over a substantial band: a quasi frequency-independent local impedance can be achieved through geometry and a well-designed feed transition.

The dominant low-frequency limitation for a horn is the electrical size of its aperture. When the aperture is electrically small, radiation is inefficient and the input impedance becomes reactive. Using the quarter-wavelength criterion of the general framework gives a first-order bound

$$f_{\text{low}} \sim \frac{c}{4L}, \quad (4.21)$$

where L represents the relevant length scale for travelling-wave radiation. For horns this scale is strongly linked to the largest aperture dimension D_{ap} : if $D_{\text{ap}} \ll \lambda$, the structure is in the electrically-small regime and does not radiate efficiently. Consequently, supporting spectral content down to 1 MHz implies aperture dimensions on the order of $\mathcal{O}(10^2)$ m, consistent with the sizing arguments used earlier. Any practically sized horn therefore shifts the onset of efficient radiation to higher frequencies (typically several MHz or higher depending on size).

At the high-frequency end, the horn departs from the quasi-TEM model when the feed region and smallest geometric features become electrically large and/or when higher-order modes are excited. Applying the framework criterion yields

$$f_{\text{high}} \sim \frac{c}{20 \ell_{\text{min}}}, \quad (4.22)$$

where ℓ_{min} denotes the smallest relevant dimension in the feed and transition region. In addition, in horn antenna higher-order mode behaviour becomes more likely as transverse dimensions become electrically large, which can introduce dispersive propagation and frequency-dependent impedance perturbations. In time domain, these effects can manifest as pulse distortion and post-pulse ringing.

In the transmission-line framework, the horn terminates at the aperture, where the guided wave transitions to free-space radiation. Unlike a matched resistive termination, this transition is by definition imperfect: diffraction at the aperture edges and any abrupt change in geometry introduces an effective reflection coefficient $\Gamma_L(f)$. This makes the horn input impedance sensitive to the flare profile, aperture shape, and any absorber or edge treatment. A first-order representation remains

$$Z_{\text{in}}(f) = Z_0 \frac{1 + \Gamma_L(f) e^{-2j\beta L}}{1 - \Gamma_L(f) e^{-2j\beta L}}, \quad \beta = \frac{2\pi f}{c}, \quad (4.23)$$

with an associated ripple spacing

$$\Delta f \sim \frac{c}{2L}. \quad (4.24)$$

For well-designed horns, the flare is chosen to make the transition to free space gradual, thereby reducing $|\Gamma_L(f)|$ over a broad band. Nevertheless, because the termination is fundamentally radiative and frequency dependent, a perfectly flat $Z_{\text{in}}(f)$ over extremely wide bandwidth is not expected.

The TEM horn can satisfy the “geometry-controlled impedance” criterion over a broad range and, relative to the biconical, offers more directivity and a feed impedance that is often easier to design around standard values. However, the same multi-scale constraints remain: achieving 1 MHz - 1 GHz within a region of relatively frequency-independent input impedance would require an extremely large aperture to avoid the electrically-small regime at the low-frequency end while simultaneously maintaining mm - cm level feed tolerances and controlling aperture reflections at the high-frequency end.

4.4 Vivaldi / Tapered Slot Antenna

The next antenna under consideration is the vivaldi antenna, which in literature is often referred to as a tapered slot antenna (TSA). It was first introduced by Gibson in 1979 [53] and has since seen application in ground penetrating radar, biomedical imaging, indoor positioning systems and UWB communication systems. It sees most application when small antenna size is a requirement, although scaling up to larger antenna is possible [54]. It is noted that this antenna is very similar to the TEM-horn antenna, but the differences are such that they are worth mentioning.

4.4.1 Working principle

The vivaldi consists of two metallic flares separated by a gradually widening slot, often seen in an exponential tapering profile. The vivaldi supports a quasi-TEM mode waveform that is fed in the narrowest point and starts to bulge as it travels along the expanding slot. At the end, it launches the quasi-TEM mode wave and it expands spherically in the endfire direction. The general layout of the antenna is shown in figure 4.3, on the next page.




Figure 4.3: vivaldi antenna schematic

Practically, the vivaldi antenna and the TEM horn antenna of section 4.3 are very similar. They use the same principle of feeding a quasi-TEM waveform between two metallic plates that are tapered such that the wave will radiate spherically outwards at the end. One key difference between the two antennas is that the vivaldi antenna uses only the two thin plates, allowing the antenna to be very flat. This characteristic means that it exhibits similar working principles to the TEM horn, while at the same time being easier and cheaper to fabricate and significantly easier to place in an array [54]. Very small models can also be integrated onto PCB's [55]. In figure 4.3, the area between the metallic plates are air. However, a choice may be made to use a dielectric substrate to print the antenna on, specifically when integrating with a PCB. For the purposes of this project, this is impractical and therefore not considered.

4.4.2 Design considerations

The vivaldi antenna itself has many of the same design considerations as the TEM horn antenna, which are not worth discussing in detail again. Generally speaking, the excitation of the vivaldi is more precarious, but not very different from TEM horn excitation [54]. The same holds for the first order estimation and the application of the transmission line framework, which therefore will not be repeated for the vivaldi antenna.

4.5 Reflector IRA

A further step up in directionality is made with the reflector Impulse Radiating Antenna (IRA), which the classical and most mature antenna concept for generating transient impulses. It is the antenna that is most studied for the purpose of generating NEMP-like pulses of all the antenna concepts that will be discussed, and hence is also the antenna that is generally meant when referring to an IRA [38]. It originates from the work of C.E. Baum and D.V. Giri in the Sensor and Simulation Notes series. The reflector IRA was first proposed in 1989 [35] and since has found application in UWB radar systems, biomedicine and in defence systems. Baum, Giri and Farr have summarized all work done on IRA's until 1998 in [29], an especially useful document for this section containing an ordered list of sources.

4.5.1 Working principle

The architecture of the reflector IRA consists of a wideband TEM feeding structure located at the focal point of a parabolic dish reflector. The stepped voltage source launches the wave onto the, in this case four, TEM feedlines. The stepped voltage source is located at or close to the focal point of the parabolic dish reflector. The flared nature of the TEM feedlines cause the TEM wave propagating between them to expand to a quasi-spherical shape. When this quasi-spherically expanding wave strikes the dish, the TEM wave is focused into a directed pulse. This means the wave exhibits far field characteristics at a relatively close distance to the antenna, considering the bandwidth at which the IRA operates [39]. The matching circuit serves to prevent internal reflections in the TEM feeding structure. The feedlines widen towards the dish for the same reason. This general architecture is shown in figure 4.4 below.




Figure 4.4: Reflector IRA Schematic [28]

The radiated electrical field from any aperture antenna, like the reflector IRA, can be found by the spatial integration of the fields over the aperture. This is as follows:

$$E_{far}(t) \propto \iint_{aperture} E_{ap}(r, t_r) dS \quad (4.25)$$

This is a logical consequence of the working principle of the parabolic reflector dish; it collects all the electromagnetic energy fed to it through the TEM feed and bundles it into approximately a plane wave. The time variation of the radiated far field is approximately the time derivative of the aperture field, a consequence of Maxwell's equations. Furthermore, it is a property of the IRA that the electric far field is the time derivative of the feeding voltage source [35]:

$$E_{far} \propto \frac{d}{dt} V_{feed} \quad (4.26)$$

Therefore, if a step voltage is fed to the IRA, it will produce a waveform that is the differentiated step; in the case of a perfect step this is an impulse [40]. Applying a DEXP will produce another DEXP, but with higher initial amplitude and a sharper early peak. This will shift the frequency content of the DEXP upwards, meaning higher frequencies are stronger while lower frequencies weaken. In other words, in time domain the antenna sharpens the DEXP form.

Apart from the desired main impulse, the reflector IRA will also radiate a prepulse and postpulse. Typically, the prepulse is caused by the weak direct radiation from the feeding region before the main spherical TEM wave illuminates the dish. This leakage arrives at an observation point earlier than the main pulse and appears as a low amplitude precursor [37]. The postpulse originates from reflections within the antennas structure, specifically scatters from the edge of the dish. This shows up as a low amplitude, longer duration peak after the main impulse at the observation point. The main impulse is bracketed by two faint and broad peaks in electromagnetic radiation.

A variation of the reflector IRA is the Half IRA (HIRA). It is literally that; half a reflector IRA on top of a reflecting ground plane. A HIRA works by replacing the lower half of the IRA's feed and reflector with a conductive ground plane that creates an electromagnetic mirror. This allows a single feed arm pair to launch a clean spherical TEM pulse, simplifying mechanics, reducing cost, and producing nearly the same radiated impulse as a full IRA as long as the ground plane is electrically large [39].

4.5.2 Design considerations

The main choice to be made in the design of the reflector IRA is the physical size of the dish. This determines the minimum wavelength that the structure supports without distortion, and hence determines the lower bound of the operating frequency range. Furthermore, a larger dish will produce a higher electric field strength in the far field. The trade-off lies in the practicality of a large antenna.

The upper bound of the operating frequency range is fundamentally constrained by the rise time of the pulse, which is what gives rise to the highest meaningful spectral component. However, the antenna can only radiate these if the dish has sufficiently small tolerances on its accuracy [36]. Furthermore, the feeding structure must be such that no distorting effects appear, such as higher order modes [30]. Between the lower and upper bounds of the operating frequency, the magnitude spectrum is approximately flat [29].

Another design choice to consider is the number and orientation of the TEM feeds. Even numbers of feeds are required, but the choices leading to the schematic in figure 4.4 are not necessarily ideal [38]. Two arms for instance is mechanically simpler and has less aperture blockage, while on the other hand being only symmetric in one plane. Furthermore, two TEM feeds may be more prone to common mode currents and the polarization purity is more sensitive. Four arms, like in figure 4.4, gives the TEM wave spherical symmetry, as well as allowing the antenna to radiate both vertically and horizontally polarized waves without the need for physically rotating the antenna. Adding more arms will increase the uniformity of the TEM mode, while at the same time increasing aperture blockage and construction complexity [38]. Additionally, the orientation of the TEM feeds may be altered such that the equal angles between the arms are changed. This may increase effective gain while reducing cross-polarization, aperture blocking and the appearance of common mode paths [41].

Finally, the motivation for choosing a full reflector IRA or a HIRA will be highlighted. Making use of the image theorem above an electrically conducting plane is tempting in this project as the DUT will be in seawater. Intrinsically, half a parabolic dish is less complex, heavy and expensive than a full parabolic dish. The feeding point of the antenna can now be placed on the ground plane as opposed to having to be suspended in front of the dish. However, the HIRA only radiates in a half-space, meaning the DUT must be placed at the surface of the ground plane for electric field strength and fidelity. Furthermore, the surface of the ground plane must be approximately flat, as any distortion of the surface will disturb the propagation of the wave, which is not a given for the sea surface, even in a harbour or sheltered area. Finally, the HIRA provides less freedom for use in an array, as the HIRA must be placed on the ground plane, hence allowing placement only on the surface without any vertical difference between antenna elements.

4.5.3 First order estimations

Having established the reflector IRA working principle and the dominant design drivers a brief first-order estimate is useful to quantify feasibility and to compare the effect of increased directivity against the earlier concepts. As before, the intent is not accuracy, but an order-of-magnitude check to gain understanding in the validity of the concept as well as gaining some feeling for the design considerations at hand.

A first sizing driver is again the lowest efficiently radiated frequency f_L . For an aperture-type radiator, the aperture must be electrically large, hence a common rule-of-thumb is

$$D \sim (0.25-0.5) \lambda_L, \quad \lambda_L = \frac{c}{f_L}, \quad (4.27)$$

where D is the reflector diameter. With $f_L = 1$ MHz, $\lambda_L \approx 300$ m, which yields

$$D \sim 75-150 \text{ m}. \quad (4.28)$$

Once again the half-IRA will be applied to this first order estimation, as that is a logical choice considering the environment and the wavelengths. This results in $D_{half} \sim 75-150$ m. This also reinforces the conclusion from the previous subsections: efficient radiation down to 1 MHz demands structures of $\mathcal{O}(10^2)$ m scale.

The upper end of the usable bandwidth is constrained by the rise time content and by how accurately the reflector and feed preserve the desired wavefront. In first order, the feed region must remain electrically small at f_H , and the reflector surface accuracy must be sufficiently tight to avoid phase errors across the aperture. For GHz fidelity this typically implies mm - cm level mechanical tolerances, which becomes a significant practical driver for large structures.

To estimate the feed voltage required for a specified field level at 100 m, the same far-field power relation is used:

$$E(r) = \frac{\sqrt{30 P_{\text{rad}} G}}{r} \Rightarrow P_{\text{rad}} = \frac{(E(r) r)^2}{30 G}. \quad (4.29)$$

For the reflector IRA, the gain is set primarily by the reflector aperture. A first-order aperture estimate is [14]

$$G(f) \approx \eta_{\text{ap}} \frac{4\pi A}{\lambda^2} = \eta_{\text{ap}} \frac{4\pi (\pi D^2/4)}{\lambda^2} = \eta_{\text{ap}} \frac{\pi^2 D^2}{\lambda^2}, \quad (4.30)$$

with A the physical aperture area and $\eta_{\text{ap}} < 1$ an efficiency factor accounting for illumination and blockage. Once again assuming the real part of the input impedance at the feed is approximately R_{in} , the feed-voltage estimate follows from

$$P_{\text{rad}} \approx \frac{V_{\text{rms}}^2}{R_{\text{in}}} \Rightarrow V_{\text{rms}} \approx \sqrt{P_{\text{rad}} R_{\text{in}}}, \quad V_{\text{pk}} \approx \sqrt{2} V_{\text{rms}}. \quad (4.31)$$

Combining yields the same convenient scaling relation used previously:

$$V_{\text{pk}} \approx \sqrt{2 R_{\text{in}}} \frac{E_{100} r}{\sqrt{30 G}}, \quad r = 100 \text{ m}. \quad (4.32)$$

It should be noted that, specific to the IRA, the radiated far-field waveform scales with dV_{feed}/dt (as stated in (4.26)), meaning that the time-domain fidelity is linked to the feeding structure of the antenna. Despite this, the estimate remains useful for comparing concepts at first order.

As an illustrative example, take $E_{100} = 50$ kV/m, $r = 100$ m and $R_{\text{in}} = 50 \Omega$. To remain consistent with the minimum size estimate driven by $f_L = 1$ MHz, the reflector diameter is taken as $D = 75$ m (i.e. $D \sim 0.25\lambda_L$), with an aperture efficiency $\eta_{\text{ap}} = 0.6$. For a HIRA, the reflector is placed above an electrically conducting ground plane and therefore radiates into half-space. To first order, this doubles

the gain in the main direction relative to the full-space case (as only half the energy is radiated), hence $G_{\text{HIRA}} \approx 2G$. At $f = 100$ MHz ($\lambda \approx 3$ m), the full-aperture estimate gives

$$G \approx \eta_{\text{ap}} \frac{\pi^2 D^2}{\lambda^2} \approx 0.6 \frac{\pi^2 (75)^2}{3^2} \approx 3.7 \times 10^3, \quad (4.33)$$

and therefore $G_{\text{HIRA}} \approx 7.4 \times 10^3$. The required radiated power becomes

$$P_{\text{rad}} = \frac{(E_{100r})^2}{30G_{\text{HIRA}}} \approx \frac{(5 \times 10^6)^2}{30 \times 7.4 \times 10^3} \approx 1.1 \times 10^8 \text{ W}, \quad (4.34)$$

and the corresponding peak feed voltage estimate is

$$V_{\text{pk}} \approx \sqrt{2P_{\text{rad}}R_{\text{in}}} \approx \sqrt{2 \times 1.1 \times 10^8 \times 50} \approx 1.1 \times 10^5 \text{ V} \sim 0.11 \text{ MV}. \quad (4.35)$$

If this voltage is assumed to be supported across a 15 mm feed region, the implied electric field strength is

$$E_{\text{feed}} \approx \frac{V}{d} \approx \frac{0.11 \times 10^6}{0.015} \approx 7.1 \times 10^6 \text{ V/m} = 7.1 \text{ MV/m}, \quad (4.36)$$

which remains above the ~ 3 MV/m air breakdown limit. Once again, to achieve 50 kV/m at 100 meters some adaptation has to be made to the feeding area to tolerate the voltage at the feeding region.

4.5.4 Application of the transmission-line model

The reflector IRA can also be interpreted within the travelling-wave transmission-line framework, but with an important distinction: the “transmission line” behaviour is associated with the TEM feed structure, while the parabolic reflector acts as a radiated pulse bundling aperture that redirects the expanding wave into a focused pulse. In other words, the dominant impedance behaviour is set by the feed network and its termination behaviour, whereas the reflector mainly influences gain, waveform shaping, and edge-scattering rather than acting as an impedance taper. Therefore, only the feeding structure can be considered in this model, which makes it applicable to both the reflector IRA and the lens IRA.

The feed of a reflector IRA consists of a multi-conductor TEM structure designed to support a broadband wave launched by a fast step source. Over the portion of the feed where the fields remain quasi-TEM, a geometry-controlled characteristic impedance can be defined, analogous to Z_0 in a transmission line. Practical IRA feeds are typically designed to interface to standard pulsed-power impedances, and the feed arms are flared to reduce abrupt discontinuities and suppress reflections [70]. As long as the dominant mode remains TEM and the geometry is electrically smooth, this part of the structure can satisfy the “geometry-controlled impedance” criterion of the general framework.

Even if the local feed impedance is approximately constant, the measurable input impedance is frequency dependent whenever the wave is not perfectly absorbed by radiation and/or by a deliberate termination. In the TL approximation, the feed can be modelled as an effective line of impedance Z_0 and length L_{feed} that is terminated by an effective load representing the combined effect of radiation and structural scattering. The corresponding input impedance follows the standard form

$$Z_{\text{in}}(f) = Z_0 \frac{1 + \Gamma_L(f)e^{-2j\beta L_{\text{feed}}}}{1 - \Gamma_L(f)e^{-2j\beta L_{\text{feed}}}}, \quad \beta = \frac{2\pi f}{c}, \quad (4.37)$$

where $\Gamma_L(f)$ is an effective reflection coefficient at the “end” of the feed network. For an actual IRA this $\Gamma_L(f)$ does not correspond to a simple lumped termination; instead it is caused by a combination of effects. For specifically a reflector IRA this includes imperfect radiative loading, internal reflections in the feeding network and scattering from the parabolic dish reflector [70]. For the lens IRA, this may be caused by internal reflections in the lens, frequency dependence of the propagation properties of the dielectric in the lens and scattering on the edges of the lens. Together these effects cause there to be not perfectly flat $Z_{\text{in}}(f)$.

The low-frequency bound is not set by the feed length but primarily by the electrical size of the focusing aperture, in this case the parabolic dish or lens. When the dish diameter is electrically small, the radiative loading becomes weak and reactive, which increases $|\Gamma_L(f)|$ and makes $Z_{\text{in}}(f)$ strongly frequency

dependent. This statement also holds for the lens. This is consistent with the earlier sizing arguments that relate the lowest efficiently radiated frequency to the reflector diameter.

At the high-frequency end, the limiting factor is the ability of the feed and reflector to preserve the intended quasi-TEM excitation. As the feed features become electrically large and tolerances become a significant fraction of a wavelength, higher-order mode content and localized resonances can disturb the effective radiation impedance and increase $|\Gamma_L(f)|$. In addition, the reflector surface accuracy and edge details become increasingly important for waveform fidelity and for suppressing spurious scattering contributions.

Within the transmission line framework, the reflector IRA can be viewed as a quasi-TEM feed line driving a frequency-dependent radiative termination associated with the reflector/lens. The concept can achieve a comparatively broad region of near-constant input impedance if the feed network and matching elements are designed to minimize reflections and if the reflector is electrically large over the band of interest. However, covering 1 MHz - 1 GHz with a relatively frequency-independent $Z_{in}(f)$ remains fundamentally challenging: the low-frequency end demands a very large aperture to provide sufficient radiative loading, while the high-frequency end demands mm - cm tolerances and careful suppression of feed-region resonances and scattering. In practice, the reflector IRA is therefore well suited to a wide but finite impedance window.

4.6 Lens IRA

The final antenna that is considered in this chapter is the lens IRA. It has a feeding structure similar to that of the reflector IRA, but instead of relying on a metallic reflecting dish to bundle and direct the spherical TEM wave into a plane wave, it relies on a dielectric lens. The use of such a lens in a NEMP simulator is very old [42], but despite this the lens IRA remains less prevalent in literature than the reflector IRA. While working NEMP simulators have been made with reflector IRA's, there has yet to be a design published using a lens.

4.6.1 Working principle

Similar to the reflector IRA, the lens IRA is an aperture antenna that is fed by a TEM feed. The aperture serves to focus the spherically expanding TEM wave into a radiating plane wave. The lens IRA differs in that it does this with a dielectric lens as opposed to a parabolic dish. A simple schematic is shown on the next page.




Figure 4.5: Schematic of Lens IRA [43]

Figure 4.5 shows schematically the main components of the Lens IRA. The reader will recognise the TEM feed, in this case two conical arms, that support a spherically expanding wave originating at the stepped voltage source. The spherically expanding TEM wave propagates into the dielectric lens, where it is focused into a plane wave. Instead of using reflection to focus the beam, the lens IRA uses refraction

at the lens-air interface to convert the spherically expanding TEM wave into a plane wave approximation.

The lens IRA is an aperture antenna, just like the reflector IRA. This means that the radiated electric field is proportional to the spatial integration over the aperture, so equation 4.25 holds. Furthermore, as an IRA, equation 4.26 also holds.

The lens IRA has a less significant prepulse in the far field than the reflector IRA. The feeding source is behind the aperture from the point of view of an observer in the boresight, so no electromagnetic energy can be radiated before the main impulse passes through the lens. This also means there is no aperture blocking. Consequently, the postpulse is more significant, despite the lack of metallic structures in the aperture causing edge diffraction, reflections and other disruptive effects. Instead it is the dielectric lens exhibiting internal reflections because of impedance mismatches, which will appear as a ringing or reverb after the impulse [29].

Absorption in the lens is by definition frequency dependent; certain frequencies will disperse more than others [46]. This necessitates predistortion of the feeding voltage in order to compensate internal frequency dependent losses in the lens. Furthermore, this may cause pulse stretching, reducing main impulse risetime and hence pulse fidelity [45].

A variation of the lens IRA as it is shown in figure 4.5 is a solid dielectric lens IRA. This is an IRA variation in which the entire TEM feeding structure is contained within the same dielectric as the lens [45]. A solid dielectric lens will have higher absorption losses, but it also only has one medium interface, reducing reflection losses. This means pulse fidelity may be lower while maximum electric field strength is higher compared to the lens IRA of figure 4.5.

4.6.2 Design considerations

In similar fashion to the reflector IRA, the minimum operating frequency is fundamentally determined by the size of the lens structure. The larger the lens diameter, the larger the minimum wavelength the lens can efficiently support. The weight of a dielectric lens IRA is non-trivial [47], so the wish to support very low frequencies may mean the lens IRA is an impractical choice.

The choice of lens material is a significant design parameter. The relative permittivity provides a trade-off: increasing it will reduce focal length at the cost of stronger internal reflections and higher sensitivity to inhomogeneity in the dielectric. The frequency dependency of relative permittivity will cause pulse stretching. Furthermore, the loss tangent of the dielectric governs attenuation due to heating [45].

The lens geometry, independent of size, is another trade-off a designer will have to balance. A wider lens will provide a higher gain at the cost of higher weight, more reflection and absorption losses and a longer and stronger postpulse [47]. The exact shape will also determine the production complexity and hence cost.

As the lens IRA is similar in feeding and aperture mechanics to the reflector IRA, the first-order estimation and the application of the transmission line model calculations are omitted due to practical considerations.

4.7 Comparison

In this section, the comparison between all the antenna concepts discussed will be made in order to eventually choose the concept best suited for this thesis project. This will be a qualitative comparison based on the design trade-offs as they were shortly illustrated in the previous sections. The metrics used, as well as their weight, to compare these antenna are a result of the requirements for the radiated field presented in chapter 2. At the end, the points will be summed up and the highest scoring antenna will be selected.

The most important metric chosen is named "radiation pattern," for which 20 points may be gained. This is an umbrella term that will encapsulate the field homogeneity (both field strength and time difference of arrival across the testing area), peak field strength and practical antenna testing ease with respect to directivity. The higher the antenna directivity, the better its radiated field homogeneity, the higher its peak field strength (given the same feed) and the easier it is to design a test in a crowded environment where others may not be disturbed.

10 points may be scored for the time domain fidelity (the opposite of time domain disturbance of the pulse), the design simplicity (how straightforward is it to design a working antenna) and construction practicality (how well can a design be practically implemented). 5 points may be scored on the ability to radiate both polarizations and the availability of literature on the antenna concept.

Table 4.1: Comparison of antenna concepts for NEMP generation

Metric	Biconical	TEM Horn	Vivaldi	Reflector IRA	Lens IRA
Radiation Pattern	2	12	10	20	16
Time Domain Fidelity	6	8	8	10	8
Design Simplicity	6	10	8	4	1
Construction Practicality	4	10	8	4	2
Polarization	1	2	2	5	5
Literature Availability	2	3	2	5	3
Total	23	45	38	48	35

A quick look at table 4.1 will show that the reflector IRA is best suited as antenna in radiated tests. The motivation for the given scores will be given now per metric.

The row titled radiation pattern scores the antenna concepts on their radiation pattern and its consequences for this project. The radiation-pattern score can be supported by the first-order gain and feed-voltage estimates made in the preceding sections. For a required field level at range r , the radiated power scales as

$$P_{\text{rad}} = \frac{(Er)^2}{30G},$$

so the required feed voltage scales approximately as $V_{\text{pk}} \propto 1/\sqrt{G}$. This immediately shows why directivity is important for the present application. The biconical/monocone concept has only low directive gain and therefore requires a peak feed voltage of order several MV to reach 50 kV/m at 100 m. In contrast, the aperture-type concepts considered here reduce the required voltage to the order of 10^5 V for comparable assumptions. The radiation-pattern score therefore mainly reflects the ability of the antenna to concentrate the available pulse energy onto the test surface.

This calculation strongly supports the low score assigned to the biconical antenna: most of the radiated energy is not directed towards the test surface, resulting in poor peak-field efficiency and practical test difficulties. The TEM horn and Vivaldi antenna improve this substantially because they are directional travelling-wave apertures. The reflector and lens IRA score higher still because they are aperture-focusing concepts. The reflector IRA receives the maximum score not only because of its aperture gain, but because the parabolic reflector provides strong wavefront control and a well-defined boresight pulse, making it the most suitable candidate for producing a high-field, localized, NEMP-like illumination region.

Next the time domain fidelity of the antennas is scored. This is the ability of the antenna to radiate a pulse as undisturbed as possible. Naturally, this is for a large part achieved through careful design of the antenna, but certain designs are more prone to disturbances than others. Specifically the lens IRA, the TEM horn, the vivaldi antenna and the biconical are more prone to disturbances in time domain than the reflector IRA. For the lens IRA this is because the lens will interact differently with different frequency components, hence being prone to dispersive radiation of the pulse. For the TEM horn and the vivaldi antenna careless design may cause higher order modes may be supported in the feeding structure of the antenna, hence causing notches in the frequency content. Finally, the biconical antenna itself may produce a pulse with high fidelity, but the feeding region and support structure cause disturbances. Careful design of the antenna can prevent most of these effects, so the score for antennas that are more prone to this is only slightly lower than the 10 points given to the reflector IRA.

The design simplicity metric encapsulates the ease with which a working design can be created. In this category the TEM horn gains 10 points for being the simplest design concept. The vivaldi, being very similar to the TEM horn, gains 8. The biconical antenna requires more tuning than the TEM horn and the vivaldi antenna, specifically for the feed region and the terminating structure, so it is granted 6 points. The IRA's perform worst in this category, specifically the lens IRA, as the dielectric lens requires significant tuning to efficiently implement.

The next metric is the construction practicality. It scores antenna concepts on the ease with which a working design can be built and implemented. The TEM horn once again scores the full 10 points as the simple design can easily be implemented in reality. The vivaldi is slightly lower as its slightly more complex shape requires extra attention. The biconical and reflector IRA score low on this metric due to the supporting structure required for both antenna, specifically the biconical antenna. The reflector dish of the reflector IRA must also consider wind loads generated by the dish. Finally, the lens IRA with its heavy and large dielectric lens scores lowest, as it requires significant supporting structure. Furthermore, a lens cannot simply be constructed from different components put together on the site due to the internal reflections that would cause, which amongst others makes transport impractical.

It must be noted that for all of these antennas, it holds that practicality is a function of frequency bound: the 1 MHz frequency bound requires aperture sizes of up to 150 meters - which is not feasible. The scores are given considering that the maximum size of the antenna is constrained to a feasible number, so comparison between the designs is possible.

The ability to emit both vertical and horizontal polarization is the next metric on which antenna may be scored. Only 5 points may be scored here as it is considered a less important metric. The reflector IRA and lens IRA score highest here, assuming that they have four feed-arms. If they have four feed-arms, both horizontal and vertical polarizations can be emitted at the same time. The TEM horn, vivaldi and biconical antenna can not do this, so they score lower. The TEM horn and vivaldi score higher than the biconical antenna because of the possibility of physically rotating the entire antenna in order to also radiate the other polarization, which is impractical for the TEM horn and vivaldi antenna but impossible for the biconical.

Finally, the literature availability of the antenna concepts is scored. 5 points can be scored as a maximum, which the reflector IRA achieves because the development of the antenna has been fully recorded in the Sensor and Simulation Notes by Baum, along with numerous publications. The Lens IRA and TEM horn score a 3, as they exist on Baums Notes as well as on publications, although not necessarily for the purposes of NEMP hardness validation testing. Finally the vivaldi and biconical antenna are niche and under-reported antenna concepts, hence scoring the lowest.

4.8 Discussion

The reflector IRA scored best on the metrics considered to qualitatively compare the antenna concepts on their suitability as an element in a time domain array for NEMP simulation. It scored particularly well due to its comparability high directivity, time domain fidelity, the ability to radiate multiple polarizations and the availability of literature which can support the design process. It lost out on points for its relatively complex design and the difficulties associated with practically implementing such a design. The antenna requires extensive tuning to minimize a post-pulse and the feeding region design is challenging. Furthermore, the practical implementation may be challenging considering the potential size of such an antenna when the minimum required frequency is 1 MHz. Going forwards, this requirement may have to be shifted upwards in order to practically implement such a design. Furthermore, a solid sheet dish will experience extensive wind loads and require a large supporting structure, as well as being difficult to transport. The upper frequency bound may need to be lowered as well as the time domain fidelity in order to accommodate a dish structure made up out of wires or a mesh. Overall, within the weighting of metrics adopted in this thesis, the reflector IRA emerges as the most suitable candidate, although its performance at the lowest frequencies remains constrained by practical dish size.

The TEM horn is closely second. It does not perform as well as the reflector IRA does in radiation pattern or time domain fidelity, but it makes up for it by being a simple and practical design. This concept may be preferable in projects where mechanical feasibility or modular deployment outweigh the need for maximum directivity and early-time pulse shaping. The vivaldi antenna may be the antenna of choice for very small scales, and hence very high frequencies. This antenna is less suitable for the frequency range and test geometry considered here.

The lens IRA lost out on simplicity and practicality, without gaining in fidelity and radiation pattern. The lens introduces significant mechanical and dispersive challenges and does not appear to offer a clear performance benefit for the frequency range and test geometry considered here, resulting in a net lower score. Perhaps that is the reason there has yet to be a NEMP simulator to be designed using a lens as a focusing structure.

Interestingly, the biconical antenna scored worst of all antennas considered. This is interesting because this is the design used by the EMIS-III-VPD, the baseline design for this project. This contrast illustrates how strongly directivity impacts homogeneity and peak field requirements for the chosen test configuration. The biconical lost out most points in its radiation pattern, as it is omnidirectional. This is logical as this was also the biggest problem with this set up.

4.9 Conclusion

The goal of this chapter was to choose an antenna concept best suited for the generation of a NEMP-like pulse. The reflector IRA, lens IRA, TEM horn, vivaldi and biconical antennas were considered. Their working principles are all highlighted in order to get some feeling for the characteristics of these antennas. Furthermore, their design considerations have been discussed in order to be more acquainted with the trade-offs that are at play in the design of such an antenna. These give insight into the antennas' general ability to produce such a NEMP like waveform and what metrics are neglected when others are prioritised.

A choice for antenna concept was made by qualitatively comparing the antenna concepts. For this, a scoring table was used, in which concepts could score points based on the importance of certain metrics for this project. The reflector IRA is the antenna concept that is chosen as it has scored the most points in table 4.1, and it is the antenna on which this thesis will continue.

The TEM horn antenna was closely second, scoring particularly well as a simple and practical design. This makes it an interesting antenna to consider for future research when these are the metrics that have the biggest priority, above the directionality of the antenna. Interestingly, the biconical antenna scored lowest of all antennas on this list. This is mainly because it is not a directional antenna. It is tempting to suggest that the use of any directional antenna will provide an improvement in the quality of a NEMP hardness validation test.

Chapter 5

Antenna Design

A motivated choice has been made to use the reflector IRA as the antenna element in the time domain array that will serve as NEMP simulator. The choice was mainly motivated by the reflector IRAs well-suited radiation pattern, which allows the concept to generate a field with high field strength and homogeneity. Furthermore, the concept has a high level of time domain fidelity, is able to generate pulses in both polarizations and is well described in literature. The challenge of such an antenna lies in the design complexity and practical implementation, which will require careful thought and design. This is what will be covered in this chapter.

Design of the reflector IRA will be done to specification as distilled from the NEMP standards in Chapter 2. For convenience, the main specifications are summed up again below:

Table 5.1: NEMP Test Specification Requirements

Specification	Requirement	Note
Peak electric field strength	5 kV/m $-0/+6$ dB	The standards require 50 kV/m compared to theoretical DEXP
Rise time (10%–90%)	2.5 ± 0.5 ns	
Full width at half maximum	23 ± 5 ns	
Lower frequency bound	1 MHz	
Upper frequency bound	1 GHz	
Spectral magnitude tolerance	± 10 dB	
Polarization	Horizontal and Vertical	To be tuned to antenna MIL-STD-461G RS105
Testing surface area	10 m $\times$ 10 m	
Test distance from radiator	100 m	
Field homogeneity	$-0/+6$ dB	

It is explicitly not a goal of this thesis to optimize a reflector IRA design for the purpose of NEMP testing. The goal is only to obtain a working model of such a radiator in order to produce numerically calculated far fields that can serve as input for the design of the time domain array. A working model is sufficient and therefore desired above a perfect model. For this reason, the approach taken in this chapter is not to optimize each geometric parameter from first principles, but to:

1. review the literature on reflector IRAs to identify the high-level parameters that any well-performing design for NEMP simulation should have;
2. identify an existing, well-characterised reflector IRA from the literature whose parameters match those identified, and adopt its detailed geometry as a template;
3. scale this template to the dish diameters necessary to facilitate the desired wavelengths.

This approach avoids the considerable design and optimisation effort associated with custom-tailoring a reflector IRA for NEMP simulation, while still producing a physically meaningful radiator suitable for use in the array design that follows in later chapters.

5.1 Design Parameter Selection

Before addressing the dish size, which is the dominant trade-off between low-frequency coverage and practical feasibility, there are several high-level design parameters of the reflector IRA that must be selected. These are the focal-length-to-diameter ratio (F/D), the number of feed arms and the feed impedance. While these parameters interact with one another and with the dish size, their selection can be motivated independently from the literature on reflector IRAs. This section provides that motivation and fixes these parameters at the level needed to identify a suitable reference design in section 5.1.5. The detailed geometry is not derived here, but is instead inherited from a to be selected reference design.

5.1.1 Number of Feed Arms

The number of feed arms N determines the symmetry of the TEM field launched towards the reflector, the feed impedance, the aperture blockage and the polarisation capability of the antenna. In the IRA literature, designs with two, four, six and eight arms have been studied [28, 38, 41].

The feed arms are arranged in pairs that together support a TEM mode between them. Each pair has a characteristic impedance Z_{pair} that is determined by the arm geometry. For a reflector IRA this is typically around 400Ω [70]. Because the pairs are connected in parallel at the feed point, the total feed impedance is:

$$Z_{\text{feed}} = \frac{Z_{\text{pair}}}{N/2} \quad (5.1)$$

This gives $Z_{\text{feed}} = 400 \Omega$ for two arms, 200Ω for four arms, 133Ω for six arms and 100Ω for eight arms.

The quality of the TEM field distribution over the aperture is characterised by the geometry factor f_g , which appears in the expression for the on-boresight far-field impulse response [28]:

$$E_{\text{far}}(r, t) = \frac{D \cdot f_g}{4cr} \frac{dV_{\text{feed}}}{dt} \quad (5.2)$$

where D is the dish diameter, c the speed of light, r the observation distance and V_{feed} the feed voltage. The geometry factor f_g quantifies how effectively the feed arms distribute the TEM field across the aperture, and it increases with the number of arms: $f_g \approx 0.41$ for two arms, $f_g \approx 0.58$ for four arms, $f_g \approx 0.65$ for six arms and $f_g \approx 0.68$ for eight arms [28, 38]. There are diminishing returns beyond four arms.

Adding arms also increases aperture blockage, as each metallic arm casts a shadow on the reflector. For four arms on a large dish, the total blocked area is a small fraction of the aperture and the effect on gain and waveform fidelity is negligible. Farr et al. [41] treat blockage as a secondary effect in their design optimisation of four-arm IRAs, and the JOLT system [39] operated with (the electrical equivalent of) four arms without reporting blockage as a limiting factor. For six or more arms the blocked fraction grows, but even then the effect remains modest provided the dish diameter is much larger than the arm width.

From the perspective of this thesis, the most significant consideration is the ability to radiate both horizontally and vertically polarised pulses, as this was identified in section 2.2 as being important for NEMP hardness validation. A two-arm feed can only radiate a single polarisation. A four-arm feed with two orthogonal pairs can radiate both horizontal and vertical polarisation without mechanically rotating the antenna, by selectively exciting one pair at a time. This capability was explicitly demonstrated by Baum et al. in the JOLT system [39].

Taking all of the above into account, four arms are selected for this design. The four-arm configuration provides a good geometry factor ($f_g = 0.58$), acceptable feed impedance ($Z_{\text{feed}} = 200 \Omega$), negligible blockage, and dual-polarisation capability. While six or eight arms offer slightly higher f_g values, the improvements are marginal and come at the cost of increased construction complexity and lower feed impedance, which requires higher currents for the same radiated power. The four-arm configuration is also the most thoroughly documented in the IRA literature [28, 29, 39, 41], which supports the design process within the scope of this thesis.

5.1.2 Feed Arm Angular Placement

With four arms selected, the angular placement of these arms must be considered. Tyo [84] showed that the feed-arm geometry of a four-arm crossed-coplanar reflector IRA can be optimized by allowing the angular position of the arms to vary. Baum [85] also showed that increasing the feed-arm angle from the conventional $\phi_0 = 45^\circ$ case can provide a potential performance improvement, with $\phi_0 = 60^\circ$ being a notable case. Bowen et al. [41] experimentally confirmed this trend for solid and collapsible IRAs by comparing feed arms placed at $\pm 45^\circ$ and $\pm 30^\circ$ with respect to the dominant polarisation axis. For the collapsible IRA, this resulted in an increase in peak gain from approximately 23 dBi to 24 dBi.

For this thesis, dual-polarisation capability is required to test both horizontally and vertically polarised NEMP environments. This means the two arm pairs must be orthogonal. If the pairs are not at 90° , a pulse fired on one pair will partially excite the other, contaminating the polarisation purity of the test. The standard 90° configuration is therefore adopted. The modest gain improvement from non-orthogonal placement does not justify the loss of clean polarisation separation in a NEMP hardness validation context.

5.1.3 F/D Ratio

The ratio of focal length F to dish diameter D determines the shape of the parabolic reflector and has several consequences for the antenna performance. A low F/D ratio corresponds to a deep dish in which the focal point is close to the vertex, while a high F/D ratio gives a shallow dish with the focal point further away. Three effects are of interest: the aperture efficiency, the waveform fidelity and the practical feasibility of the structure.

Aperture Efficiency

The TEM feed produces a spherically expanding wave from the focal point. As this wave propagates towards the dish, the electric field amplitude decays as $1/r_s$, where r_s is the distance from the focus to a point on the reflector surface. For a paraboloid with focal length F , this gives an aperture field distribution [28]:

$$E_{\text{ap}}(\rho) = \frac{E_0}{1 + \left(\frac{\rho}{2F}\right)^2} \quad (5.3)$$

where ρ is the radial coordinate in the aperture plane. This distribution is nearly uniform for high F/D ratios and tapers towards the rim for low F/D ratios.

The aperture efficiency is the product of the illumination efficiency η_{ill} , the spillover efficiency η_{spill} and the blockage efficiency η_{block} . The illumination efficiency quantifies the penalty from the non-uniform aperture distribution and is close to unity for all practical F/D values due to the favourable nature of the TEM feed illumination. The spillover efficiency is the fraction of the power radiated by the feed that is intercepted by the dish. Defining $q = D/(4F) = 1/(4 \cdot \text{F/D})$ and $\theta_{\text{edge}} = 2 \arctan(q)$ as the half-angle subtended by the dish rim as seen from the focus, the spillover efficiency, assuming isotropic radiation from the feed, is:

$$\eta_{\text{spill}} = 1 - \cos(\theta_{\text{edge}}) \quad (5.4)$$

This efficiency drops steeply with increasing F/D. At $\text{F/D} = 0.25$ the dish subtends the full forward hemisphere ($\theta_{\text{edge}} = 90^\circ$) and $\eta_{\text{spill}} = 1$. At $\text{F/D} = 0.40$ this drops to $\eta_{\text{spill}} \approx 0.56$, meaning that 44% of the feed power misses the dish entirely. The combined aperture efficiency therefore favours a low F/D ratio from a pure power-capture standpoint.

Waveform Fidelity

The choice of F/D ratio directly affects the pre- and postpulse of the radiated waveform. As discussed in section 4.5.1, the radiated waveform consists of three components: a low-amplitude prepulse, the main impulse, and a postpulse caused by edge diffraction from the dish rim.

The prepulse duration is approximately $\tau_{\text{pre}} \approx 2F/c$, which is the round-trip delay between the direct feed leakage and the reflected wave arriving at the focal point [37]. A higher F/D ratio results in a longer

prepulse. Since the total area of the prepulse is constrained by Maxwell's equations (the time integral of the far field must be zero), a longer prepulse has a lower amplitude, which is desirable. The postpulse amplitude is proportional to the edge illumination level, which is the aperture field strength at the dish rim relative to the centre. From equation (5.3), the edge taper is:

$$\frac{E_{\text{edge}}}{E_{\text{centre}}} = \frac{1}{1 + q^2} \quad (5.5)$$

With q short-hand notation for $q = D/4F$ [86]. At $F/D = 0.25$ the edge taper is -6 dB, at $F/D = 0.40$ it is approximately -2.9 dB and at $F/D = 0.50$ it is -1.9 dB. A stronger edge taper (more negative value) means a weaker postpulse, which is desirable for waveform fidelity. The postpulse arrives one clearing time $\tau_c = D/(2c)$ after the main impulse, which for a 14 m dish is 23 ns; the FWHM of the desired NEMP-like pulse. The postpulse arriving during the tail of the main pulse and will distort the waveform if its amplitude is too large. From this perspective, a low F/D ratio is preferred as it provides stronger edge taper and hence a weaker postpulse.

There is hence a conflict: lower F/D gives stronger edge taper (smaller postpulse) but shorter prepulse duration (larger prepulse amplitude). Both the pre- and postpulse are undesired perturbations to the main impulse. However, the postpulse is generally the more problematic of the two, as it arrives during the time window where the NEMP pulse shape is evaluated against the standard [29]. The prepulse, while of larger amplitude at low F/D , precedes the main pulse and does not overlap with the evaluation window.

Practical Feasibility

Higher F/D ratios place the feed point further from the dish, resulting in a longer feed arm assembly that must be mechanically supported. Conversely, very low F/D ratios produce a deep dish that is difficult to manufacture with the surface accuracy required for 1 GHz operation, and that experiences higher wind loads due to its concave shape. Both extremes present practical difficulties. One may look at existing designs to see how this trade-off was dealt with in a practical setting.

Selection

The literature on reflector IRAs shows a convergence towards F/D values in the range of 0.35 to 0.45. Farr [41] recommends this range when optimising for high gain and low impedance. The JOLT system, which is the most powerful reflector HIRA documented in the open literature, used $F/D = 0.4$ [39]. Bowen et al. [41] used similar values in their experimental optimisation of feed arm placement. The 46 cm reference IRA analysed in detail by Taylor and Sarkar [40], which will be adopted as the geometric template for this design in section 5.1.5, also uses $F/D = 0.40$.

5.1.4 Summary of High-Level Parameters

The high-level design parameters selected for the reflector IRA element are summarised in table 5.2. These parameters define the family of reflector IRAs to which the antenna of this thesis belongs, but they do not yet specify a complete design. Several geometric details remain open: the exact shape of the feed-arm cross-section along its length, the geometry of the feed-point transition where the coaxial generator output splits into the four arms, the manner in which the arms terminate at the dish rim, and any matching elements between the arm tips and the rim. Each of these affects the broadband impedance match and hence the time-domain fidelity of the radiated pulse. Rather than re-deriving these details from first principles, an existing reflector IRA whose parameters match those of table 5.2 is adopted as the geometric template in the next subsection.

Table 5.2: Selected high-level reflector IRA design parameters

Parameter	Value	Justification
Number of feed arms N	4	Dual polarisation
Arm angular placement	90° separation	Clean H/V polarisation
Feed impedance Z_{feed}	200Ω	$Z_{\text{pair}}/(N/2)$, eq. (5.1)
Geometry factor f_g	0.58	Farr SSN-329 [28]
F/D ratio	0.40	Literature range, JOLT [39]
Dish diameter D	To be determined	Section 5.1.6

5.1.5 Reference Design

The parameters fixed in the previous subsections describe a family of reflector IRAs rather than a specific antenna. The detailed geometry that turns this family into a buildable design; the precise taper profile of the feed arms, the geometry of the coaxial-to-four-arm transition at the focal point, and the configuration of the matching elements at the dish rim, has a strong effect on the broadband impedance match and the time-domain fidelity of the radiated pulse. Optimising these details is a non-trivial design problem in its own right, and as stated in the introduction of this chapter it is explicitly not a goal of this thesis to perform such an optimisation. Instead, the strategy adopted here is to identify an existing, well-characterised reflector IRA from the literature that already matches the parameters in table 5.2, and to adopt its detailed geometry as a template that is then scaled to the dish diameters relevant for this work.

The reference design selected for this purpose is the 46 cm diameter four-arm reflector IRA analysed by Taylor and Sarkar in Sensor and Simulation Note 528 [40]. In their work they show that it is possible to replicate the measured time and frequency domain EM results of an existing IRA, described in [74], [41], [75], [28], [76], [77], [78], in electromagnetic simulations. This makes Sensor and Simulation Note 528 ideal, as it both describes the antenna developed by Farr in great detail as well as describing the steps taken to simulate this antenna correctly.

This antenna shares all of the high-level parameters of table 5.2: it has four feed arms in two orthogonal pairs at 90° separation, an F/D ratio of 0.40, and a feed impedance of 200Ω realised by two pairs of 400Ω coplanar arms in parallel. The full electromagnetic analysis of the antenna is provided in [40], including the impedance match across a broad frequency range and the on-boresight time-domain response, which makes it a well-validated starting point. Furthermore, the design itself is closely related to the line of work by Farr and co-workers from which the high-level parameters in table 5.2 were drawn, so adopting this geometry is consistent with the rest of the literature trail used in this chapter.

The aspects of [40] that are adopted directly are:

- the linear taper profile of the four feed arms from the focal point to the dish rim;
- the geometry of the feed-point transition;
- the configuration of the matching elements at the dish rim that terminate the feed arms.

The features that are not adopted from [40] are the absolute dimensions: the 46 cm diameter of the reference antenna places its lower frequency bound far above the 1 MHz target of this thesis, as the antenna scaling argument of chapter 3 makes clear. The reference design is therefore used as a geometric template, scaled up to the dish diameter required for NEMP simulation.

When the reference design is scaled by a factor $s = D/D_{\text{ref}}$, where $D_{\text{ref}} = 0.46$ m, all geometric features scale linearly with s . The F/D ratio, the feed impedance, the geometry factor f_g and the edge taper are all dimensionless and therefore invariant under this scaling, which means the high-level parameters of table 5.2 are preserved. The radiated electric field on boresight, given by equation (5.2), scales linearly with D for a fixed feed voltage, while the lower frequency bound at which the antenna radiates efficiently scales as $1/D$, consistent with the relationship derived in section 5.1.6 of chapter 3.

What does not scale linearly is the upper frequency bound and the associated mechanical tolerances. The upper frequency limit of a reflector IRA is set by the absolute size of the smallest features in the feed-point transition rather than by the dish diameter. The 46 cm reference antenna is reported to operate up to several GHz in [40], so the smallest feed-point features are on the millimetre scale. When the antenna is scaled up to metre-scale diameters, these features become several centimetres, which would normally lower the upper frequency bound below the 1 GHz target of this thesis. For the present work this discrepancy is accepted: as discussed in chapter 2, the spectral content of the NEMP environment is heavily concentrated below 100 MHz, and the loss of fidelity above several hundred MHz is a minor effect compared to the gain in low-frequency performance from scaling up the dish. Re-optimisation of the feed-point transition for higher upper frequency performance is left as future work, consistent with the scope statement at the start of this chapter.

5.1.6 Antenna Dish Size

As was shown in chapter 3, there exists a fundamental problem with radiated NEMP testing: the antenna must be impractically large in order to efficiently radiate the low frequency component of the NEMP-like pulse. This was an inherent problem with the EMIS-III-VPD, the baseline for this thesis, which could simply not radiate the very lowest desired frequencies efficiently. Radiated NEMP testing contains a fundamental trade-off between antenna feasibility and ability to efficiently radiate the low frequency component of the pulse.

It is not up to the author of this thesis to determine where the line between feasible or infeasible lies. The motivation behind this choice depends on a multitude of factors, such as cost, construction capability, available test range size, power availability, or whether it is desired to have a transportable facility. However, this thesis may instead illuminate this trade-off somewhat. If the trade-off is more clearly parametrized, it may aid those who have to make this decision.

Therefore, this thesis will present three scaled reflector IRA sizes. Each will represent a possible choice for where the feasible/infeasible line lies. The antennas simulated have a dish diameter of 5, 12.2 and 30 m. These numbers are based on figure 3.24, in which the EMIS-III-VPD model was scaled up and down to estimate its low frequency operating limit, as well as from the far field consideration discussion done in the next section. From figure 3.24, it was determined that the EMIS-III-VPD had to be 83 m tall to radiate 1 MHz effectively. While it was initially the ambition to show simulations of this antenna scale to show what would be possible in an ideal, limitless world. Unfortunately, even in simulations, this size is impossible to implement. The largest feasible numerical model has a 30 m dish diameter, which will come to serve as an optimistic, upper bound for antenna size. Such an antenna is estimated to radiate upwards of 3 MHz, according to the calculation done for figure 3.24

The 5 m dish size, the smallest, will represent the practical option. Such an antenna is still comparatively huge, but relative to the wavelengths desired it is the least costly, transportable and practical option. It is estimated to radiate efficiently from 17 MHz onwards, according to the calculations done for 3.24. Finally, the 12.2 m dish size will provide a compromise between these two antennas. This specific size is selected as we will come to see that it is exactly the size allowed to be in the antenna far-field at 100 m distance. It is expected to radiate from 7 MHz onwards, which may still be more considering the improved directivity of the IRA.

5.2 A note on the antenna far field

A question must be addressed that was raised but not fully resolved in Chapter 3 for the EMIS-III-VPD: does evaluating the radiated field at 100 m satisfy far-field assumptions for the antennas under consideration? It was noted in section 3.2 that the 100 m test distance may not place the EMIS-III-VPD in its asymptotic far field, given the ultra-wideband nature of the radiated pulse. For the reflector IRA models considered in chapter 5, this concern carries directly across, and is amplified: the largest reflector diameter to be considered is a significant fraction of the testing distance. The applicability of far-field assumptions therefore needs explicit treatment.

The conventional far-field distance for an aperture antenna is the Fraunhofer distance [14],

$$r_{\text{ff}} = \frac{2D^2}{\lambda}, \quad (5.6)$$

where D is the largest aperture dimension and λ the wavelength. This criterion is derived from the requirement that the path-length difference between the centre and the edge of the aperture, evaluated at the observation point, remains below $\lambda/16$, which corresponds to a phase error of $\pi/8$ across the aperture [14]. It is a frequency-domain criterion and assumes a continuous-wave (CW) excitation. Equation (5.6) is dominated by the highest frequency of interest, since shorter wavelengths give larger far-field distances. Evaluated at the upper bound of the desired spectrum, $f_H = 1$ GHz ($\lambda \approx 0.30$ m), equation (5.6) yields 167 m, 1.0 km and 6.0 km for the 5 m, 12.2 m and 30 m models respectively. By this criterion, none of the antennas are in the far field at 100 m distance.

This conclusion is, however, overly conservative for an impulse-radiating antenna. The reflector IRA does not radiate a CW signal at 1 GHz; it radiates a transient whose spectral content is heavily concentrated below 100 MHz, as established in Chapter 2. Applying the Fraunhofer criterion at 1 GHz therefore demands far-field conditions for spectral components that contain only a small fraction of the radiated energy, and ignores the fundamentally transient nature of the radiated waveform. A different approach is more suitable. Note that the term 'near-field' will often be used in this section. With that, the radiative near field is meant, also known as the Fresnel zone, and not the reactive near field.

5.2.1 A time-domain criterion

A more appropriate criterion follows directly from the time-domain behaviour of the radiated pulse. The relevant question is not whether the phase error across the aperture is small at every frequency, but whether the time-of-arrival difference between contributions from different parts of the aperture is small compared to the rise time of the pulse. If this condition is met, the on-axis pulse shape at the observation point is preserved; if it is not, the finite size of the aperture stretches and distorts the pulse, an effect known as near-field dispersion [79].

For an aperture of largest dimension D and an on-axis observation distance r_p , the maximum path-length difference between a contribution from the centre of the aperture and one from the rim is approximately

$$\Delta R = \sqrt{r_p^2 + (D/2)^2} - r_p \approx \frac{D^2}{8r_p}, \quad r_p \gg D/2, \quad (5.7)$$

with a corresponding time-of-arrival difference

$$\Delta t = \frac{\Delta R}{c} \approx \frac{D^2}{8cr_p}. \quad (5.8)$$

Sabath [79] shows that, for the radiated pulse shape to remain distance-independent within an acceptable tolerance, Δt must be small compared to the rise time τ_r of the radiated signal. Adopting his practical criterion that the rise time should be at least six times the longest time-of-arrival difference, $\tau_r \geq 6\Delta t$, and substituting equation (5.8) yields the transient far-field distance for a generic aperture antenna;

$$r_p \geq r_{\text{ff,td}} = \frac{3(D/2)^2}{c\tau_r} = \frac{3D^2}{4c\tau_r}. \quad (5.9)$$

This expression has been validated experimentally by Sabath on a 1.8 m-diameter half-IRA driven by a 100 ps rise time double-exponential pulse [79], and reformulated and extended in a more recent treatment by Wang and Xie [80], who note that the classical frequency-domain criterion is generally insufficient for high-power UWB pulsed antennas because it fails to capture the rise-time-driven dispersion that governs the boresight pulse shape.

For reflector IRAs specifically, a slightly more permissive criterion can be derived from the focusing property of the dish. Baum et al. [39] establish that the boresight pulse from a reflector IRA reaches its asymptotic shape once:

$$r_p \geq r_{\text{ff,IRA}} = \frac{D^2}{2 c \tau_r}. \quad (5.10)$$

Requiring that this aperture-to-observer time spread remain no larger than one quarter of the rise time, $\Delta t \leq \tau_r/4$, and solving for r_p recovers equation (5.10) directly. Applying the same procedure to the Sabath criterion of $\tau_r \geq 6 \Delta t$ reproduces his expression $r_p \geq 3D^2/(4c\tau_r)$, so the factor 1.5 between the two thresholds reflects a difference in stringency on rise-time preservation rather than a difference in physical mechanism.

What distinguishes the two criteria in practice is therefore not the algebra but the implicit assumption about path equalisation inside the antenna. The aperture-to-observer leg is captured by both. The focus-to-aperture leg, however, is path-equalised by the parabolic dish and contributes negligibly for a well-formed IRA, which is what justifies adopting the more permissive prefactor in equation (5.10). For aperture antennas without such focusing, such as TEM horns and biconicals, additional time-of-arrival spread is accumulated within the antenna structure before the wave reaches the aperture plane, and the Sabath criterion (or a similarly conservative bound) is the appropriate choice. The EMIS-III-VPD baseline of Chapter 3 is evaluated under the Sabath criterion only for this reason.

Two further features of these expressions are worth emphasising. First, unlike the Fraunhofer criterion, the transient far-field distance does not contain a wavelength, but a rise time: the relevant time scale of the pulse, not its highest spectral component (although the two parameters are of course linked). Second, both $r_{\text{ff,td}}$ and $r_{\text{ff,IRA}}$ scale as D^2/τ_r , so longer rise times relax the criterion proportionally. For the 2.5 ns rise time of the NEMP standard DEXP set in chapter 2, equations (5.9) and (5.10) yield the values listed in Table 5.3, alongside the Fraunhofer distances at 1 GHz for comparison. The IRA-specific criterion is applied only to the three reflector IRA models; the EMIS-III-VPD baseline is not a focused aperture and is evaluated under the Sabath criterion only.

Table 5.3: Far-field limit for the EMIS-III-VPD and the three reflector IRA models.

Model	D [m]	r_{ff} (Fraunhofer)	$r_{\text{ff,td}}$ (Sabath)	$r_{\text{ff,IRA}}$ (Baum)
EMIS-III-VPD baseline	20	2.7 km	400 m	–
IRA 5 m	5	167 m	25 m	16.7 m
IRA 12.2 m	12.2	1.0 km	149 m	99 m
IRA 30 m	30	6.0 km	0.9 km	600 m

Note that the Fraunhofer column is evaluated at 1 GHz, the highest frequency in what is defined as the relevant frequency spectrum in Chapter 2 of this thesis. The Sabath column uses $\tau_r = 2.5$ ns in equation (5.9). Equally, the Baum column uses the same rise time in equation (5.10) and applies only to reflector IRAs. The largest dimension D is the cone height for the EMIS-III-VPD and the dish diameter for the reflector IRAs.

5.2.2 Implications for the test distance

Table 5.3 makes clear that the 100 m test distance does not place any of the antennas considered fully in their far field by both time-domain criteria, with the exception of the 5 m and 12.2 m IRA. The 5 m IRA comfortably satisfies both criteria by a factor of four or six, depending on the calculation used. The EMIS-III-VPD baseline misses the Sabath criterion by a factor of four. The 12.2 m IRA is a deliberately

chosen borderline case: its Baum far field begins at 99 m, so it satisfies the IRA-specific criterion essentially exactly at the 100 m test distance by construction, while still missing the more conservative Sabath criterion by a factor of about 1.5. The IRA-specific Baum criterion is more permissive but does not change the qualitative picture for the largest reflector: the 30 m IRA still misses it by a factor of six. Under the Fraunhofer criterion the situation is yet more pessimistic for all four antennas, although as established this is not a suitable criterion for UWB transient antennas.

This result is not a fault of the antennas, but an unavoidable consequence of the size required to efficiently radiate the lowest frequencies of the desired NEMP spectrum, in combination with the requirement to deliver useful field strength at a practically reachable distance. The same dish size that brings the lower operating frequency down to a few MHz also pushes the transient far field out to kilometre scales for the chosen rise time, while moving the test surface out to kilometres would render the field-strength target unattainable. This presents yet another trade-off in radiated NEMP tests: either ensure far-field criteria are met or ensure peak field strength is met. As shown, in the case of the EMIS-III-VPD neither are met, but an attempt was made to sacrifice far-field guarantee for peak field strength.

The peak electric field intensity is not affected by this, it scales with $1/r$ within this limit and thereby not at risk of falling below the tolerances stated for this thesis. The time-domain fidelity however is at risk due to the near-field dispersion. This causes pulse stretching, causing the rise-time and FWHM to be longer, so much so that it may fall outside of the bounds of the NEMP hardness tolerances as stated in the standards, specifically MIL-STD-461G RS105 [34]. Furthermore, near field dispersion may cause inhomogeneity across the testing surface, as contributions from different sections of antenna arrive at different times.

Placing the testing surface within the far-field limit is still a viable strategy. Clearly so, as the EMIS-III-VPD also had its DUT placed in its near-field, however no description of any characterisation of this problem has been found, let alone a mitigation strategy. The same holds for the EMPRESS-II [11]. Therefore, one must search elsewhere for the characterisation or mitigation of this error.

A mitigation strategy may be implemented as to overcome the dispersion caused by radiating within the near-field. If the dispersion caused by measuring in the near-field is quantified, it may also be corrected. One possible strategy is to adjust the pulse fed to it, a form of pre-distortion, in order to ensure that the pulse arriving at the measuring location is true to the desired DEXP. Regarding the distortion experienced from the feeding pulse to the measuring point as a (linear) transfer function, it may be corrected at the feed. This is unlikely to fully alleviate near-field dispersion, but it may help.

Another possibility is to adjust the lensing structure, in this case the dish, in order to 'aim' it at the testing surface. This would mean elongating the F/D ratio by not using a parabolic dish but, for instance, a prolate-spheroidal dish in order to not aim towards infinity, but to a point closer by, proposed by Baum in [33]. The prolate-spheroid, a three-dimensional ellipsoid, has its feeding point on one of its two focal points and its measurement location at the other focal point. Such a construction has been built before for the purpose of medical radiation therapy in [83] supported by [82]. Further development of this idea is outside of the scope of this thesis, but it is good to keep in mind that such mitigation strategies exist. Mitigation makes near-field testing possible and a viable method to achieve high field strengths and a low low-frequency limit.

The 100 m test distance specified in Chapter 2 is therefore retained as a deliberate compromise, identical to that implicitly accepted in the EMIS-III-VPD campaigns: spatial spreading at this distance is accompanied by some amount of near-field dispersion, as the radiated pulse will not yet have settled to its far-field shape. As the goal of the thesis is to illuminate a significantly smaller area (5 m x 5 m as opposed to 200 x 50 m), this potential dispersion incurred by the set distance is accepted. Taking the Baum criterion for the far-field limit, the largest dish size whose far field begins at 100 m for a 2.5 ns pulse has a diameter of 12.2 m; this is precisely the motivation for including the 12.2 m reflector IRA among the models studied, as the largest dish that remains marginally far-field-compliant at the chosen test distance under the IRA criterion.

5.2.3 Implications for this thesis

For reflector IRAs the consequences of these approximations are partially mitigated by the focusing property of the dish. The on-boresight prompt amplitude of an IRA scales cleanly as $1/r$ at any range, in the near and far field alike, and the field-range product rE_{far} is the standard figure of merit reported for such antennas precisely because it is conserved across the near field [39]. The CST output therefore represents the prompt peak amplitude faithfully even when the test distance lies inside the transient near field. What it does not represent faithfully is the pulse shape: the time-of-arrival difference $\Delta t = D^2/(8cr_p)$ between contributions from the dish centre and the dish rim arriving at an on-axis observer is implicitly set to zero in the far-field calculation, and this difference is precisely what stretches the trailing edge of the pulse and degrades rise time at the observer.

A natural threshold against which to evaluate Δt is provided by the rise-time tolerance adopted in Chapter 2. The MIL-STD-461G RS105 specification [34] admits a rise-time deviation of at most ± 0.5 ns from the nominal 2.5 ns. For the simulated waveform to remain within this tolerance, Δt at the test point must be small compared to 0.5 ns. The Sabath criterion $\tau_r \geq 6 \Delta t$ that underlies equation (5.9) corresponds to $\Delta t \leq 0.42$ ns for the present rise time [79], in close agreement with the MIL-STD tolerance and thereby providing independent support for the threshold. Table 5.4 lists Δt at 100 m for each antenna considered in this thesis, together with its ratios to both the rise time and the tolerance bound.

Table 5.4: Time-of-arrival difference between dish-centre and dish-rim contributions at an on-axis observer 100 m from the antenna, expressed as a fraction of the 2.5 ns rise time and of the ± 0.5 ns tolerance.

Model	D [m]	Δt at 100 m [ns]	$\Delta t/\tau_r$	$\Delta t/0.5$ ns
EMIS-III-VPD baseline	20	1.67	0.67	3.3
IRA 5 m	5	0.10	0.04	0.20
IRA 12.2 m	12.2	0.62	0.25	1.2
IRA 30 m	30	3.75	1.50	7.5

For the 5 m IRA the path-difference contribution is 0.10 ns, comfortably within the rise-time tolerance, and the CST output may be taken at face value. For the 12.2 m IRA the contribution is 0.62 ns, essentially equal to the Baum threshold of $\tau_r/4 = 0.625$ ns — consistent with this antenna sitting exactly at its IRA far-field limit at 100 m — and it exceeds the strict ± 0.5 ns tolerance only marginally, by a factor of 1.2. For the EMIS-III-VPD baseline Δt is already 3.3 times the tolerance bound; the simulated rise time at 100 m is therefore an underestimate of what would be measured, by an amount of the same order as the rise time itself. For the 30 m IRA the path-difference contribution exceeds the tolerance by a factor of 7.5 and the rise time itself by 50%; in this case the simulated pulse shape at 100 m bears little quantitative relation to what would be measured at that range, and only the prompt amplitude and the asymptotic pulse shape it represents are physically meaningful.

The implication for the analyses that follow is that, for the three reflector IRA models, the simulated peak E-field at 100 m is a faithful representation of the prompt boresight amplitude, by virtue of the IRA's $1/r$ amplitude scaling property. The simulated rise time and FWHM, however, must be read as scaled far-field values: a lower bound on the achievable rise time and a lower bound on the FWHM, with these bounds becoming progressively less tight as the antenna size increases. For the 30 m IRA near-field dispersion alone is sufficient to push the rise time well outside the MIL-STD tolerance, even if the simulation itself reports values inside it; for the 12.2 m IRA the dispersion is marginal, exceeding the strict ± 0.5 ns bound only slightly while still satisfying the Baum far-field criterion by construction. For the EMIS-III-VPD baseline, which is not a focused aperture and does not benefit from the IRA's more tolerant near-field limit, both the simulated peak field and the pulse-shape parameters carry near-field corrections, although these are modest given the relatively small ratio $\Delta t/\tau_r = 0.67$. Quantifying the actual local pulse shape at the test surface in the near-field regime would require an analytical Kirchhoff-integral propagation from the simulated aperture field distribution to each test point, evaluated without the far-field approximation, or equivalently the application of the closed-form on-boresight near-field expression for a reflector IRA given in [39]; this is left as future work.

As already stated in chapter 3, the choice to retain far-field monitors is deliberate, despite the analysis done above. The far-field pattern is physically the correct quantity for later synthesization of the contributions made by different elements in the array, therefore of value to the remainder of this thesis. Furthermore, it is the most efficient approach in handling the results produced by CST.

5.3 Simulations

Three models of three different sizes (5, 12.2 and 30 m dish diameter) were made using scaled versions of the design geometry as listed in Sensor and Simulation Notes [40]. Each size represents a choice in the trade-off between antenna practicality and lowest operating frequency limit. In this section the model is discussed and results are shown.

5.3.1 Antenna Geometry

As discussed in the previous section, the geometry of the antenna is not designed in this work. Instead, a working antenna geometry is taken [40] and scaled to suit the needs of this thesis. The geometry remains constant across all three models, only the size of the antenna is scaled. The general layout of this antenna is shown below:





Figure 5.1: Geometry of the reflector IRA


Figure 5.1 shows the reflector IRA implemented in CST [12]. Two main components can clearly be seen: the parabolic dish reflector and the feed arms. Although in the image the illusion is that the top of the dish is a horizontal plane, but this is in fact hollow. All components are PEC. The feed arms can be seen to taper outwards with the same specification as in [40]. Equally, the feed arms taper back into a single point where a lumped element connects the arms with the dish. At the top of the arms, the feed points do not touch. Instead, a gap is left between them with a discrete element connecting the arms. Opposite pairs of arms are connected. All gaps, for the lumped elements and for the feed, are equal in size: all are 0.05 m. This does mean that the very ends of the arms are of slightly varying geometry in order to accommodate the equal size of gap across the different sizes of antenna. This is done in order to ensure that the elements added for feeding and termination are in fact lumped; they are so short they are insignificant to wavelength. At 0.05 m they are 1/12th of the highest simulated frequency. A closer look is taken at the feed-point and the termination in figure 5.2 on the next page.

In figure 5.2a the feed point of the IRA is shown. The ends of the arms are adjusted such that the gap is always 0.05 m in length. Opposite arms are connected, as can be seen in the image. In this image, only one pair of arms is connected, the other two remain unconnected. This means that in that specific case only one polarization is radiated.

Figure 5.2b shows the end termination connecting the arm to the dish. This component is crucial for the effective operation of the reflector IRA. Its goal is to terminate the wave travelling along the arm as completely as possible, in order to prevent reflections. Connecting it with the dish allows any electromagnetic energy not dissipated in the lumped element to travel onwards in the dish instead of reflecting back into the arm.



(a) The feed point of the IRA



(b) Connection between feed arm and dish

Figure 5.2: A more detailed look at the geometry of the reflector IRA

The matching of these elements is a crucial part of antenna design. It is known from [40] that this specific geometry has an arm impedance of $200\ \Omega$, so the initial value for discrete port and lumped element impedance was set to this value. Further refinement was done by scanning over different values for a limited frequency range and coarse cell size. It was found iteratively that a value of $225\ \Omega$ yielded lower S11 and radiation efficiency across the frequency range for the numerical model, hence a better broadband match. This difference may be due to a multitude of reasons; this model may have slightly different feeding and terminating geometry, it may have different feedgap and terminating lumped element size, its numerical solver may introduce (different) artefacts, etc. Note it is not the goal to accurately reconstruct the IRA as it is presented in [40], instead [40] is the starting point from where a working model can be made.

A final important note to make about the antenna design is the way in which CST handles an analytical shape, such as the dish, with respect to its solver meshing. The parabolic reflector is defined as an ideal analytical surface, but the transient solver employed throughout this work discretises the geometry onto a structured hexahedral grid using the Finite Integration Technique, so the continuous paraboloid is approximated by the mesh before the fields are computed [12]. On a purely staircased grid this approximation would replace the smooth dish with a series of cell-aligned steps, and for a reflector IRA this is not just a cosmetic concern: the entire focusing principle relies on the paraboloid equalising the optical path lengths from feed to aperture plane to within a small fraction of the rise time, so any mesh-induced departure of the reflecting surface from the ideal profile causes a spread of arrival times at the aperture and hence into a numerically induced broadening of the impulse and rise time, mechanically identical to the near-field dispersion discussed above. CST mitigates this through its Perfect Boundary Approximation and Thin Sheet Technique, which resolve curved and thin perfectly conducting boundaries at sub-cell resolution rather than snapping them to the grid, recovering the curvature of the dish without the cell counts that a staircase representation would demand [12]. Even so, the surface must be resolved finely enough that the residual geometric error is small. Experimentally it was found that using less than 10 cells per (maximum) wavelength yields results with some numerical distortion at high frequency.

5.3.2 Feeding Pulse

In the previous chapter it was discussed that it is a characteristic of the reflector IRA that the time domain E field is the derivative of the voltage fed to it. This implies that if it is desired to radiate a DEXP, as described in chapter 2, then the voltage fed at the feeding gap must be the integral of the desired E field. Hence, the integral of the desired DEXP must be calculated.

As a reminder, the standard NEMP electric field is approximated by a double-exponential waveform,

$$E(t) = E_0 k (e^{-\alpha t} - e^{-\beta t}) u(t), \quad (5.11)$$

where E_0 is the peak field amplitude, k is a normalisation factor, α is the decay constant, β is the rise constant, and $u(t)$ is the unit step function. The time integral of this waveform from 0 to t is

$$\int_0^t E(\tau) d\tau = E_0 k \int_0^t (e^{-\alpha\tau} - e^{-\beta\tau}) d\tau. \quad (5.12)$$

Carrying out the integration gives

$$\int_0^t E(\tau) d\tau = E_0 k \left[\frac{1 - e^{-\alpha t}}{\alpha} - \frac{1 - e^{-\beta t}}{\beta} \right]. \quad (5.13)$$

This can also be written as

$$\int_0^t E(\tau) d\tau = E_0 k \left[\left(\frac{1}{\alpha} - \frac{1}{\beta} \right) - \frac{e^{-\alpha t}}{\alpha} + \frac{e^{-\beta t}}{\beta} \right]. \quad (5.14)$$

For use as a normalised feed waveform for an IRA, the integrated waveform can be divided by its final value. The final value is obtained by taking the limit as $t \rightarrow \infty$:

$$\lim_{t \rightarrow \infty} \int_0^t (e^{-\alpha\tau} - e^{-\beta\tau}) d\tau = \frac{1}{\alpha} - \frac{1}{\beta} = \frac{\beta - \alpha}{\alpha\beta}. \quad (5.15)$$

The normalised integrated DEXP waveform is therefore

$$V_{\text{feed, norm}}(t) = 1 - \frac{\beta}{\beta - \alpha} e^{-\alpha t} + \frac{\alpha}{\beta - \alpha} e^{-\beta t}, \quad t \geq 0. \quad (5.16)$$

For $t < 0$, the feed waveform is taken as zero:

$$V_{\text{feed, norm}}(t) = 0, \quad t < 0. \quad (5.17)$$

Plugging in the values of α and β from Chapter 2, the following is obtained:

$$V_{\text{feed, norm}}(t) = 1 - 1.0714e^{-4 \times 10^7 t} + 0.0714e^{-6 \times 10^8 t}, \quad t \geq 0. \quad (5.18)$$

The normalized result is plotted, yielding:

The signal is made to artificially decay back to zero from 400 ns to 1000 ns, as simply feeding the voltage as described in 5.18 is unacceptable for the solver to initiate a run in CST.

A comment must be made on the artificial decay between 400 and 1000 ns. As mentioned, it has been introduced in order for the CST solver to accept it, but any alteration made to the excitation signal will carry through to the eventually radiated signal. The roll-off radiates a slow, low-amplitude, negative-going lobe spanning roughly 400 ns to 1000 ns, separated from the wanted impulse by several hundred nanoseconds; because the two are disjoint in time the artefact does not distort the shape of the radiated DEXP itself and can, if required, be removed by time gating before further processing. The decay time of 600 ns places its spectral content at $1/600\text{ns} = 1.7$ MHz, which is within the frequency band of interest, but as will be shown, below the effective radiation limit for the antennas presented here. The amplitude and spectral position of this artefact are governed entirely by the steepness and smoothness of the decay: a faster decay produces a larger, higher-frequency lobe that risks encroaching on the lower edge of the band, whereas a slower, smoother roll-off pushes it down in both amplitude and frequency. The duration of this plateau will only determine the temporal separation between this tail and the main impulse, as well as lengthening the time taken for numerical simulations.




Figure 5.3: Excitation signal fed to the IRA models

As a first trial, this waveform will be fed to the reflector IRA in order to achieve the desired DEXP in the E far field. The deriving property of the IRA only holds for a perfect IRA, which is unlikely to be obtained in this chapter. Hence, to simulate the DEXP to within the tolerances as described in chapter 2 further tuning is needed. With the geometry of the antenna now implemented and discussed, the results from simulating the three scaled models can now be presented. This is done in the following three subsections, after which conclusions can be drawn.

5.3.3 30 m model

Analysis of the three models will be done in similar fashion to the presentation of the three models made to estimate the performance of the EMIS-III-VPD. First, the basic initial results gathered will be presented; the S11 parameter at the feed, the radiation efficiency (both as a function of frequency) and the radiation patterns inspected at different frequencies. This is done for each of the three models, after which in the next section a qualitative comparison and analysis of these models is done. This section will also feature all the post-processing results, such as the time domain fidelity and the frequency magnitude spectrum.

First, the S11 parameter of the 30 m model as a function of frequency is inspected:




Figure 5.4: S11 parameter of the 30 m IRA model

Generally, the S11 parameter of the 30 m model as presented in figure 5.4 reveals a good broadband match. There are two exceptions, where the S11 parameter exceeds -10 dB, at which point 90% of energy fed to the antenna is accepted by it. This is very briefly at 10 and 20 MHz, where the S11 parameter reaches -9 dB, and at a broader band between 870 and 930 MHz where the S11 reaches as high as -3 dB. Tuning of the impedances at the end termination and of the reference impedance did not alleviate this. It is likely that a geometric change to the end termination region and the feeding region is needed

to fully remove this peak of antenna mismatch. However, it is deemed that this peak is acceptable for the purposes of this thesis; the bulk of energy of the NEMP is contained in the very lowest portion of the relevant frequency spectrum and therefore this region of mismatch between 870 and 930 MHz does not seriously influence the pulse fidelity.

Apart from this region of antenna mismatch, the antenna shows excellent broadband match, especially in the region from 50 to 750 MHz, where it is settled below -20 dB, corresponding to 99% of feeding power being accepted by the antenna. It is concluded so far that the model is operating as intended.

Next, the radiation efficiency as a function of frequency is inspected:





Figure 5.5: Radiation efficiency of the 30 m IRA model

Figure 5.5 shows that the radiation efficiency of the 30 m model reveals a less impressive broadband match than the S11 alluded to. The radiation efficiency is excellent between approximately 30 MHz and 350 MHz, comfortably above 0.9 between this range. Below 30 MHz it is seen to be steeply rising as the antenna starts to 'work' as the extremely long wavelength at the lower end of the relevant frequency spectrum starts to fit into the arms and dish of the antenna. At higher frequencies however, a decay is seen. This does not rhyme with the excellent broadband match seen in the S-parameter, implying that it is caused by another effect.


It is suspected that the numerical meshing of the antenna model is causing the drop in radiation efficiency. This may be in the feed-arms, which would impact the S11 and thereby influence the radiation efficiency, although nothing alluding to this has been seen. It may also be the meshing of the analytical shape of the dish is causing this decay in radiation efficiency. A very fine mesh is needed to preserve the shape and effectiveness of the parabolic dish, but the size of the model means this leads to huge simulations. To make simulating this antenna viable, it was chosen to reduce the mesh density on the dish to just 5 per wavelength at 1 GHz, away from the recommended 10 [12]. The coarser mesh means that numerical artefacts that arise when meshing are of significant size to disturb the analytical shape in this higher frequency band, leading to inefficient radiation. This is reflected by the radiation patterns that are inspected next, shown in figure 5.6.

Finally, the radiation pattern of this antenna is inspected at four different frequencies: 1, 100, 500 and 1000 MHz. These are the same 4 frequencies used in the EMIS-III-VPD model analysis. These are shown in figure 5.6. At 1 MHz the IRA can be seen to be operating a frequency still too low to make the IRA 'work' as intended. While it shows a peak in the +z direction, with a main lobe magnitude of 5 dB, its 3 dB beamwidth as wide as 128°. In the next figure, 5.6b, the antenna has started to 'work.' Here it is shown to have a strong main lobe in the +z direction of 23.2 dB and a beamwidth of 7°. At 500 MHz this pattern of improving antenna performance with increasing frequency is continued, with a main lobe intensity of 33.7 dB and a 3 dB beamwidth of 1.4°. The improvement at 1000 MHz is not significant, here the main lobe intensity is 29.7 dB and the 3 dB beamwidth is 1.9°. This is slightly less than it was at 500 MHz, likely because of the coarse meshing issue described above.


Note that in order to make this numerical model solvable with reasonable memory use, the meshing density had to be reduced from the ideal 10 per wavelength (at highest frequency) to 5 per wavelength. Practically, this means that above 500 MHz there are less mesh-cells per wavelength that recommended. Any values presented in this thesis above 500 MHz for the 30 m model cannot be compared directly with any certainty.




(a) 1 MHz



(b) 100 MHz



(c) 500 MHz



(d) 1000 MHz

Figure 5.6: Radiation patterns of the 30 m reflector IRA model at selected frequencies.

5.3.4 12.2 m model

The 12.2 m model is intended to show a balanced choice between practicality and low frequency operating limit, while being the largest possible antenna that still operates in the far field at 100 m. It is estimated to radiate from 7 MHz onwards, based on the EMIS-III-VPD calculations, while at the same time being almost a third as large as the 30 m model. As with the previous model, first the S11 parameter is inspected:




Figure 5.7: S11 parameter of the 12.2 m IRA model

At first glance, the 12.2 m model has a good broadband S parameter. It is shown to be below -10 dB for almost the entire range, only appearing above it for a very short band around 20 MHz, where it does not exceed -8 dB. Promptly, the 12.2 m model has an excellent match at about 90 MHz, after which it can be seen to slowly worsen with increasing frequency, however never exceeding the -10 dB limit associated with a good match. From this plot it can be concluded that the 12.2 m model has a good broadband match, accepting most of the power fed to it across the band of interest.

Next, the radiation efficiency is inspected:






Figure 5.8: Radiation efficiency of the 12.2 m IRA model

The radiation efficiency, shown in figure 5.8, shows an almost perfect match between 100 and 800 MHz. Below this value it is sharply increasing from zero. This very low radiation efficiency for very low frequencies is to be expected, as here the antenna is not yet 'working' as its size is significantly smaller than the wavelength. What is striking is that the radiation efficiency drops from 800 MHz onwards, dropping below 0.9 beyond 900 MHz. It seems then that this inefficiency is frequency dependent.


The radiation pattern of the 12.2 m model inspected at 4 frequencies, the same as for the 30 m model (1, 100, 500 and 1000 MHz). This is shown in the familiar gridlike pattern on the next page in figure 5.9, which show the antenna directivity. Figure 5.9a shows the radiation pattern for 1 MHz. What can clearly be seen is that the antenna is showing infinitely small dipole behaviour at this frequency, a consequence of the large wavelength (150 m) compared to the antenna diameter. While the entire upper half of the plot is a deep red, indicating high intensity, the maximum directivity is only 4.72 dB, with 3 dB beamwidth of 135° .




(a) 1 MHz



(b) 100 MHz



(c) 500 MHz



(d) 1000 MHz

Figure 5.9: Radiation patterns of the 30 m reflector IRA model at selected frequencies.

Once the wavelength shortens to the point it is roughly the same size as the antenna the antenna will start to 'work,' as the wave will fit inside its feeding structure. At 100 MHz, in figure 5.9b, this is seen clearly. A strong spike in the $+z$ direction, with an intensity of 15.8 dB, with a 3dB beamwidth of 14.1° . Strong, wide sidelobes can be seen in the $\pm x$ direction. This is caused by imperfect termination of the wave at the end of the feed-arm, causing edge scattering off the arm or the dish. Their location is a consequence of the fact that for this specific screenshot these arms were the ones that were excited. If the other feedarm pair was excited, this pattern would be rotated 90° .

The 500 MHz plot, depicted on figure 5.9c, shows an even more idealized radiation pattern; in the sense that the desired direction shows a much stronger peak, narrower peak. The maximum directivity, in the $+z$ direction, is 29.3 dB. The 3 dB beamwidth is 3° . Sidelobing in $\pm x$ direction is still present but much narrower and lower intensity. Finally, the 1000 MHz plot does show a continuation in this pattern of improvement with increasing frequency. This plot has a maximum directivity of 33.1 dB and a 3 dB beamwidth is 1.8° .

5.3.5 5 m model

The 5 m diameter model will represent the case in which the constraints that determine what is feasible to be very tight. It may be that the testing facility is cramped, allowing space for only a small antenna. Or the costs may be limited, allowing only for the construction of a small antenna. If there is a desire for a transportable antenna or an otherwise temporary facility, a smaller antenna would be ideal. If any or all of these are the case, the 5 meter dish diameter will illustrate what the effect is of these choices on NEMP simulator performances.

As is familiar by now, the S11 parameter is inspected first. This is presented below:




Figure 5.10: S11 parameter of the 5 m IRA model

The S11 parameter, shown in figure 5.10, reveals a good broadband match for the antenna across the desired frequency range; it is shown to be below 10 dB across almost the full range. Only at the very highest frequencies does it pass through this 10 dB limit. There are a few dips in S11 indicating near perfect match, where the S11 parameter drops significantly, specifically at 190 and 320 MHz. While this is not unwelcome, the goal is a broadband match. Generally speaking, this antenna shows a much smoother S11 than the previous antenna. This is a consequence of the significantly shorter antenna size.

The radiation efficiency, shown in figure 5.11, shows excellent radiation efficiency above 200 MHz. Below this frequency the radiation efficiency is sharply rising from 0 to 50 MHz, after which it slowly climbs towards 200 MHz. This is to be expected as the antenna only starts radiating above about 17 MHz, according to the prediction made in Chapter 3. As this model uses the same number of discretized elements in the dish and arms, the effect of a dropping radiation efficiency at increasing frequency is not seen in this limited frequency range.





Figure 5.11: S11 parameter of the 30 m IRA model


The radiation pattern of the 5 m model inspected at the familiar four frequencies, shown in figure 5.12. First, in figure 5.12a, the 1 MHz radiation pattern is shown. Once again, this antenna is too small for the radiating wave. In fact, the directivity is 2.71 dB, with a 100° beamwidth. This is implying very little is being radiated at this very low frequency as the IRA is not yet working as intended, which was also seen in the radiation efficiency.

At 100 MHz the wavelength is 3 m, approximately the same size as this antenna diameter. This is shown in figure 5.12b, where the antenna can be seen to start to work. In its maximum it achieves 9.6 dB directivity and a 3dB beamwidth of 29.4° . It does however show a strong sidelobe all the way around this main lobe, approximately 45° off the $+z$ direction. This is thought to be the result of imperfect termination leading to edge scattering, a known phenomenon of the reflector IRA. The strength of this at 100 MHz with respect to this effect at other frequencies is thought to be due to scaling of the main lobe intensity, where the main lobe is significantly stronger at higher frequencies while this sidelobe is of equal strength all over the frequency spectrum. The sidelobe level is -7.4 dB at this frequency.


At 500 MHz the 5 m dish starts to yield more recognisable results; a strong and narrow main lobe with familiar sidelobing in the $\pm x$ direction, as was the case in figure 5.9b. The directivity in the main lobe is at 25.2 dB. Here, the 3 dB beamwidth is 8.5° . Finally, at 1000 MHz, this is at 25.9 dB. The 3dB beamwidth is 4.6° , almost half of what it was at 500 MHz.




(a) 1 MHz



(b) 100 MHz



(c) 500 MHz



(d) 1000 MHz

Figure 5.12: Radiation patterns of the 5 m reflector IRA model at selected frequencies.

5.3.6 Comparison and Observations

With the initial results of the antenna models obtained and discussed, this subsection will aim to summarize the observations made on these models. Furthermore, some deeper analysis using post-processing of these results will be done in the next section, in similar fashion to section 3.3.6. This will make use of conclusions drawn from this part.

The biggest difference between the EMIS-III-VPD model results and the reflector IRA results is the directionality. The directivity, and thereby the gain and realized gain, are significantly higher than they were for the omnidirectional EMIS model. This was after all the driving force behind the selection of this antenna in Chapter 4. With the reflector IRA there is an actual beam to speak of, including figures for beamwidth. This beamwidth decreases with increasing frequency, and is largest for the biggest antenna. In the next section the effect this has on the far field strength will be discussed in more detail.

Overall, it was observed that the reflector IRA possess a good broadband match. The geometry of the antenna seems to be working as intended, providing a good S11 and radiation efficiency across the desired band, which then in turn is supported by the radiation patterns. Naturally, the size of the antenna comes with constraints; the 5 m model simply does not radiate 1 MHz. Note that the antenna is 60 times shorter than that wavelength. Low frequencies remain in a trade-off with the size (and thereby practicality) of the antenna. This will be explored in more detail in the next section.

A numerical artefact that was discovered during analysis of the antenna models, although the exact source of the artefact cannot be determined with certainty. The solver takes the analytical shape of the arms and dish and discretizes it into discrete elements, named meshcells. To capture the analytical shape of the arms and dish the meshcells must be electrically small. However, for the larger models, there was an upper limit on the available memory in the computer. Therefore, a limited number was used, which led to lower radiation efficiency and scattering of the high frequency component. Whether the S11 parameter has also been affected by this cannot be said with certainty. It may impact time domain analysis, by diffusing the high frequency component it may cause the pulse to have a more shallow leading edge and a lower peak value. Future, improved models may consider using separate models for each scale, or finding a method whereby the dish and arms are made of single, defined elements.

In the next section, the performance of these antenna elements will be inspected. The results presented in this section will be subjected to further processing to yield deeper analysis in their performance and operation.

5.4 Performance Estimation

With the initial antenna models discussed and presented, post-processing of these results can be applied to gain deeper insight in their broadband and time-domain behaviour, which in turn is used to estimate their performance. This section follows the layout used for the EMIS-III-VPD in Chapter 3: the far field magnitude spectrum and time-domain fidelity are inspected, and the field intensity across the full testing area is evaluated. Note that all of the figures presented in this section are based on numerical simulations with necessary heuristics and assumptions made, resulting in potentially optimistic results. They are all therefore deemed estimates of the performance of an IRA.

5.4.1 Frequency domain performance estimation

To inspect if these antennas achieve the frequency domain performance as required, their magnitude spectrum is inspected. To do this, broadband, transient far field monitor results are exported from CST. These contain, amongst others, the realized gain of the antennas for each inspected frequency. For a fair comparison with the EMIS-III-VPD and with the goals for this thesis, the same feeding generator set up is assumed. The results of this calculation yield the figure below:




Figure 5.13: Spectral Density of the three IRA models at 100 m

Note that unlike figure 3.17, which began at 1 MHz, this figure starts at 0.1 MHz. Figure 5.13 is one of the most insightful images regarding the main trade-off between the three antenna images: the antenna practicality versus the lower frequency operating limit. Straight away it may be observed that increasing the antenna dish size also shifts the lower frequency operating limit downwards. The 30 m model enters the green tolerance bound set for this thesis at 4 MHz, the 12.2 m dish at approximately 12 MHz and the 5 m model at 33 MHz.

Note that the numerical antenna simulations done to generate this image are not perfect. The goal is to give insight in the trade-off between antenna practicality (determined by the dish size in this case) and the lower frequency operating limit. The performance of the IRA as presented in figure 5.13 may be improved by enhancement and tuning of the models as well as of the simulation.

What is striking about figure 5.13 is the fact that all three antennas don't just achieve the theoretical limit as set by the standards, they exceed it in certain frequency bands. For larger dishes it is at lower frequencies; for the 30 m model from 70 MHz onwards, for the 12.2 m model from 160 MHz onwards and for the 5 m model from 600 MHz onwards. Note that the actual radiated pulse does not necessarily follow this specific line; it has been generated using the realized gain across the frequency band of interest. When aiming to radiate an actual DEXP, the pulse shape of the feeding voltage will ensure that less energy is radiated in the high frequency portion, ensuring the radiated E-field spectral density remains within tolerances.

This effect does however speak to an interesting problem with the reflector IRA specifically, but any aperture antenna in general: the spectral density of the antenna follows exactly the opposite shape of that what is desired by the DEXP. The DEXP has most of its energy in the low frequency portion and its spectral density decays at higher frequencies, while an aperture antenna will have increasing gain with shorting wavelength, as was also seen in the previous sections describing the radiation patterns of the antenna models. So long as the antenna is built with tolerances matching the operating frequency, its gain will increase with frequency.

5.4.2 Peak field amplitude estimation

Before moving on to the time domain performance estimate, it is good to inspect the peak E-field strengths generated by each of the antenna at 100 m. With most of the energy of the DEXP contained by the first 100 MHz of the relevant frequency spectrum, the peak E-field strengths vary for each of the antenna sizes. This means that when inspecting the time domain E field plots at 100 m in the next section, they will be of significantly varying peak amplitudes, as was the case in figure 2.3. Therefore, in the next section, the E field plots will be normalized. To compensate, this section is dedicated to the peak field amplitudes.

For fair comparison with the EMIS-III-VPD it is assumed that the IRA's are fed by the same 500 kV Marx generator that was placed at the base of the EMIS-III-VPD. The peak amplitude is simply read from far field time domain plot generated by CST when the antenna is fed by a DEXP, the results of which are presented in the table below:

Table 5.5: Estimated peak electric field strength at 100 m for a 500 kV feed voltage.

Model	Estimated peak field at 100 m
5 m diameter	2.65 kV/m
12.2 m diameter	8.10 kV/m
30 m diameter	12.93 kV/m

Increasing the dish diameter directly increases the peak E field strength, as is to be expected. The 12.2 m model already achieves the 5 kV/m field strength goal set for this thesis, while the 30 m model actually exceeds the $-0 + 6$ dB tolerance (which translates to 5-10 kV/m) set for the 5 kV/m goal. This is not considered a problem; the 5 kV/m goal was reduced from the 50 kV/m taken from the standards. To achieve 50 kV/m, either a bigger dish or more antennas are needed.

Note that this is based on the far field results of the antennas, which have been measured at a point in infinity and scaled back to 100 m by a factor $1/r$. This holds for the 5 m and 12.2 m models, but does not hold for the 30 m model. The 30 m model is still in the near field at 100 m distance, and this negatively affects the peak field strength by the effect of near-field dispersion. This is shown in a later section.

Also note that these are peak E-field results. Altering the pulse shape, such as by using the integral of the DEXP at the feed and thereby broadening the energy generated by the pulse generator, will smear out energy over a longer pulse. Inefficiencies such as a large pre- or postpulse may also lower this value, as well as any propagation effects that are not accounted for here.

5.4.3 Time-domain performance estimation

Next, it is interesting to see what the effect of the antenna dish size is on the time-domain fidelity of the radiated pulse, and how that fidelity depends on the choice of feeding waveform. As shown in the previous subsection, the three dish sizes yield very different lower-frequency operating limits, and as discussed previously (figure 2.3, table 2.2) the loss of low-frequency content distorts the time-domain shape and reduces the energy of the pulse. Two analyses are therefore presented in this subsection. The first uses the standard NEMP DEXP as the feed, allowing direct comparison with the analysis of the EMIS-III-VPD done in chapter 3. The second uses the integrated-DEXP feed waveform of equation (5.18), which would be the starting point for tuning the feeding pulse for these antennas because the reflector IRA differentiates the pulse fed to it at the boresight observer, so feeding the integral of the desired field returns the desired pulse at the receiver.

A challenge in direct comparison between these antennas is that their different sizes yield significantly different peak field strengths at the boresight. This is because, as discussed in detail in Chapter 2, the DEXP contains most of its energy in the lower portion of its frequency range, thereby making the antenna size (and therefore lower frequency operating limit) a significant factor in peak radiated field strength. This was not a problem for the EMIS-III-VPD comparisons as the models compared all had the same size. Therefore, all comparisons done will be done in normalized plots as to maintain fairness and ease of making observation. The peak electric field strength comparison will be done in the next subsection.

DEXP feed

The reconstructed boresight waveforms for the three reflector IRA models, fed by the standard DEXP, are shown in figure 5.14. The traces are obtained by inverse Fourier transform of the radiated spectrum at boresight at the testing distance (100 m). To isolate differences in pulse shape from differences in absolute peak field, each trace has been amplitude-normalised, with the sign of the remainder of the waveform preserved, and time-shifted so that the three peaks coincide.




Figure 5.14: Time domain comparison of the IRA models at boresight fed by the standard DEXP, equation 5.11

All three traces share the bipolar signature expected of the differentiating action of an IRA driven by a DEXP: a sharp positive peak associated with the rising edge of the feed pulse, a zero crossing, and a negative lobe produced by the slow exponential decay of the feed. This is precisely the response that the integrated-DEXP feed of the next analysis is intended to undo. A closer look at the time scale involved with these pulses reveal that they are of significantly shorter duration than the pulse intended to be radiated; a FWHM of just 2.4 ns, 1.1 ns and 2.6 ns for the 5 m, 12.2 m and 30 m model respectively. The 12.2 m model has such a relatively short pulsewidth because it shows a notch in its decay, dropping to

below 0.5 and then rising slightly again, thereby disturbing the FWHM calculation slightly. The rise-times of the three models are 0.7 ns, 0.8 ns and 2.4 ns, respectively. Only the 30 m model therefore achieves the required rise-time as constrained by the standards.

All three models show some form of prepulse. This is initially negative for the 30 m and 12.2 m models, but this has deliberately been left out of figure 5.14 in order to more closely examine the main impulse. The 5 m model, with a significantly shorter prepulse, is still negative in the figure. In the case of the 12.2 m and 30 m models, this seems to be a prepulse in the classical sense: EM energy originating directly from the feed arms of the IRA, arriving at the measuring point before the main impulse that has been bundled by the dish arrives. Therefore it is not illogical that the 30 m dish produces a significantly longer prepulse than the 12.2 m dish, which in turn is longer than the 5 m dish. The 5 m dish prepulse is also much more intense as the same feeding pulse is radiated from the arms over a shorter distance, meaning it is not smeared out over the travel-time over the longer arms of the longer dishes. In shape however, it no longer looks like a prepulse in the classical sense of the word.

The decay of each pulse is similar: sharp (with respect to the DEXP fed to it, a consequence of the differentiating effect of the IRA) and with a negative tail. Some notches are visible in the decay, especially in the larger models. This is attributed to ringing caused by numerical imperfections in the meshing of the solver, an effect also discussed with respect to the radiation pattern and efficiency of the antenna models previously. Furthermore, the tail of each postpulse is negative, a consequence of not being able to radiate the DC-component of the DEXP fed to it and therefore mathematically forced to have a time-integral equal to zero.

Comparing this pulse to the time domain reconstructions done for the EMIS-III-VPD is possible as both antenna received the same feeding pulse. The main impulses presented here are significantly shorter than the ones generated by the EMIS-III-VPD baseline, a consequence of the differentiating characteristic of the reflector IRA. The EMIS-III-VPD also showed a bipolar shape within the envelope of the main impulse for the wire models, which is not reflected by any of the IRA models. In order for the IRA to radiate a pulse more like the pulse demanded by the standards, a different pulse must be fed to it.

Integrated-DEXP feed

Figure 5.15 shows the same three antennas fed instead by the integrated-DEXP waveform of equation 5.15. The processing of the trace is identical to figure 5.16, with one noticeable difference: they are no longer time aligned to the E-field peak but to the first zero-crossing. The contrast with figure 5.14 is striking, revealing that the change in feeding waveform is giving the desired results.




Figure 5.15: Time domain comparison of the IRA models at boresight fed by the integrated DEXP

All three models now produce roughly the unipolar shape as desired by the NEMP standards and defined in Chapter 2. The FWHM of each model is 15.4 ns, 33.4 ns and 26.6 ns for the 5 m, 12.2 m and 30 m respectively. The shorter FWHM of the 5 m model is to be expected, it is consistent with table 2.2, which showed the effect of bandpass filtering the DEXP waveform, which is essentially what happens when an antenna has a higher low operating frequency limit than the desired pulse requires. However, the 12.2 m model having a longer pulse duration than the 30 m model is inconsistent with this. Looking at the figure show that the 12.2 m model has a local maximum right as it is approaching the half maximum value, which stretches the FWHM.

The rise-times across the three models are 5.5 ns, 2.7 ns and 1.9 ns for the 5 m, 12.2 m and 30 m models respectively. This is also inconsistent with the band-pass filtering calculations done in Chapter 2, table 2.2. There, the risetimes were shorter with higher operating frequencies as a result of the peak value being lower, resulting in the peak being reached earlier. Here, however, the pulses are roughly equally as long, at least to a first order compared to the pulses of figure 2.3. It speaks to the fact that while the analogy of band-pass filtering is similar to the actual effect of having a limited band in radiation to a first order, reality is more complicated.

The pre-pulses, in the literal sense of the word, of the antenna are significantly sharper and more intense than they were for the previous pulse shape. None of the three antennas show the low intensity, simmering EM energy arriving before the main impulse arrives. In fact, in the case of the 5 m and 12.2 m model, the prepulse is more intense than the main impulse. Only the 30 m model has a stronger main impulse than the prepulse, but its prepulse is still up to 50% intensity of the main impulse. While some bipolarity is expected from this feeding pulse, the intensity of these pulses are of such a nature they reduce the energy in the main impulse. It would seem then that the small IRA's are radiating predominantly from the feed-arms (and perhaps edge diffraction) and less from the dish. This is attributed to the fact that for this specific pulse it is the very low frequency components of the pulse that are dominant, and at these very low frequencies the antenna dish is simply too small to 'work' as intended. In the standards there is no specific limit to the intensity of the prepulse or postpulse, but this is of significant intensity and proximity to the main impulse that it will disturb the test. Further tuning of the feeding voltage is required.

5.4.4 Testing surface performance estimate

As with the EMIS-III-VPD, it is insightful to gain insight into the homogeneity (or lack thereof) with which the antenna radiates the testing surface. The EMIS-III-VPD was deployed to illuminate an area of 200 m across and 50 m upwards from the horizontal plane at a distance of 100 m. In chapter 3 it was concluded that this antenna did not do so within the tolerance bounds for homogeneity as required by the standards. It was concluded in that chapter that an area of that size is infeasible to illuminate with the required homogeneity while still complying with all other tolerances. Therefore, it is the goal of this thesis to illuminate a significantly smaller area of just 5 m x 5 m at the same 100 m distance. This suits the reflector IRA to some extent, as was seen in the previous sections, the 3 dB beamwidth was as narrow as 2° at high frequencies. 4° is the angular offset between boresight and the corner of a 10 m x 10 m surface (with other words, a point 7.07 m off boresight at 100 m). This means that at high frequencies, corners of the testing surface will fall outside of the 3 dB beamwidth. At those points, the electric field will contain relatively less high frequency content, as that is being bundled better by the antenna.

As the reflector IRA is explicitly not intended for use as an antenna to illuminate the full testing surface the EMIS-III-VPD was intended for, it is not fair to discredit its performance on the large 200 m x 50 m testing surface. The reason the word discredit is used in this context is because of the highly directional nature of the reflector IRA. A heatmap, similar to for instance figure 3.18, would simply be a very intense speck at the boresight of the antenna surrounded by significantly lower intensity radiation in the form of scattering from the dish-edge, feed-arms or imperfections in the meshing on the dish. Therefore, only the 5 m x 5 m area will be presented in this thesis. The heatmaps generated will be calculated for different frequencies to give an idea of the way this depends on frequency. The heatmaps are calculated for the 12.2 m model and presented at 50, 300, 500 and 1000 MHz, and shown on the next page.




Figure 5.16: Heatmap of the E-field intensity from the 12.2 m IRA model illuminating a 10 x 10 m testing area at 100 m distance. Each gridsquare is 1 x 1 m. Orange indicates -3dB

The heatmaps of the 12.2 dish diameter model are presented at four frequencies. These were selected in order to show how the broad band of interest has an effect on the homogeneity with which the 10 m x 10 m testing area is illuminated by this dish. It was already observed that with increasing frequency the directionality is increasing, thereby decreasing the 3 dB beamwidth. This is reflected in figure 5.16 by the red circle, indicating the highest intensity E field, being displayed with decreasing size in the testing surface. At the highest frequency, 1000 MHz in figure 5.16d, the corners of the area are 20 dB lower intensity than the centre, or a tenth the E field strength in kV/m with respect to the centre. The orange area indicates 3 dB difference, approximately a factor 0.8 in kV/m. The colourscheme holds for all figures, so at 500 MHz the corner has an intensity of -10 dB, in other words a factor 0.3 with respect to the centre. The extreme case presented here, 50 MHz in figure 5.16a, has a corner intensity of -0.13 dB, a factor 0.985.

The other models have very similar patterns. They differ in the exact frequencies at which the plots above are generated, as the larger model has a more directional main lobe at a lower frequency, and vice versa. It is good to note that the method chosen to generate these heatmaps, using far-field monitors evaluated at infinity and scaled by $1/r$ back to the observation point, will produce results that differ from what one may image, or in fact reality. A larger antenna produces a narrower main lobe at infinity at the same frequency with respect to a smaller antenna. When observing at 100 m distance, CST will show that a larger antenna will have a more pencil-like beam, while the smaller antenna has a more diffused beam. In reality however, at such short range to the antenna, the beam of the large antenna will not yet have narrowed into its asymptotical form. In the case of the 12.2 m antenna, it may be assumed that at 100 m distance the antenna has settled into its far field. However, this is the limit, and therefore must be taken lightly. Only with adjustment of the focal length, such as proposed with the prolate-spheroid IRA in [33], will a larger antenna produce a narrower beam at 100 m.

Inspection of spot frequencies is useful for understanding how the antenna directivity, which is frequency dependent, affects the homogeneity over the testing surface. However, to fully evaluate the homogeneity of the antenna over a surface, the full bandwidth must be summed up and shown in a single figure. This is done in chapter 6, section 6.1.

Another metric of homogeneity is the time domain analysis of the pulse at different points on the testing surface. As was presented in figure 3.19, the time difference of arrival across a testing surface can be significant, so significant in the case of the EMIS-III-VPD that the pulse had cleanly passed through the centre of the testing surface before it had even reached the edges. With a smaller testing surface comes smaller time difference of arrival, where the estimated time difference of arrival between the centre and the corner of the testing surface is 0.83 ns, approximately equally as long as the tolerance on the DEXP risetime. A time domain reconstruction of the pulse at the centre, the corner of a 5 m x 5 m and corner of a 10 m x 10 m area are calculated for the 12.2 m model fed by the DEXP and presented on the next page.




Figure 5.17: The time difference of arrival across the 10 m x 10 m testing surface visualized using the DEXP fed 12.2 m model

Figure 5.17 shows the time domain reconstructions made for the 12.2 m models at three points: in the very centre of the testing surface, at the corner of the testing surface, and halfway between these two points. The scale of the E-field strength is normalized, as it was in the time domain performance estimation section above, but the pulse recorded at the corner is scaled to the pulse recorded at the centre. It can be seen to achieve a peak value 0.625 of the peak value of the centre, which when consulting table 5.5 would be 5.06 kV/m, still above the goal of the thesis. It reaches this peak 0.83 ns after the peak of the centre point. Note that the distance between these two pulses may seem slightly skewed by figure 5.17, the pulse fed to the IRA to generate this figure is the DEXP, which was already discussed to be significantly shorter than its integrated-DEXP fed counterpart. For that counterpart, the FWHM of the 12.2 m model was 33.4 ns, which was artificially stretched by a late time 'bump' in the pulse decay. It may be observed that despite the peaks occurring later and at lower field values, both off centre pulses fit roughly within the 'envelope' of the main impulse.

The pulse shape is slightly altered at the off-boresight corner point. This is a consequence of the high-frequency component of the pulse being less present in that area, as the high frequency content is more effectively bundled by the dish. The consequence is a wider pulse, reflected by both a larger FWHM (2.95 ns instead of 1.07 ns, a factor 2.76 longer) and a longer rise-time (1.8 ns instead of 0.84 ns, a factor 2.14 longer). The sluggish nature of the pulse may be directly attributed to the high frequency content being absent.

5.4.5 Near-field analysis of the 30 m model

As stated in section 5.2, the 30 m model is the only model of the three presented in this chapter that does not satisfy the far field conditions at 100 m, with the 5 m model comfortably in the far field at that distance and the 12.2 m model a limiting case. The near-field dispersion, which will manifest its self in the form of pulse-widening, caused by employing an antenna too large at a distance too close by for far field conditions, does not necessarily disqualify it for use as an NEMP simulator. In fact, the larger antenna size means lower a low frequency operating limit, and the dispersion also means a more diffuse high frequency radiation pattern. Therefore, this specific model's performance will also be analysed, using a different method to the analysis done above to show its true performance.




Figure 5.18: The time difference of arrival across the 10 m x 10 m testing surface visualized using the DEXP fed 30 m model using nearfield probes

In discussing the initial results of the 30 m model, specifically the radiation efficiency, that to make a model this size viable for simulation the level of discretisation had to be reduced. This resulted in numerical artefacts appearing above approximately 500 MHz. The model cannot be trusted to provide completely accurate results for this reason, specifically in the high frequency region. However, the DEXP has such a significant portion of its energy located in the very lowest portion of the relevant frequency spectrum that qualitative observations on near-field behaviour can be made. Exact numerical results must be taken lightly and with this caution in mind.

Figure 5.18 has been constructed by use of E-field probes placed at two points on the testing surface, the boresight and a point 7.07 m off boresight, which corresponds to the corner of the 10 m x 10 m testing area. As the far field numerical calculation simplifications are not used, this manner of near field analysis is robust. A consequence is that these probes cause an exponential increase in calculation size, as the entire area between antenna and probe has to be traversed by meshcells.

The time difference between the peak values is 2 ns, which is longer than it is in the far field case of figure 5.17. This is not quite a fair comparison, as in the near field the pulse is stretched due to the near-field dispersion.

A full analysis of the near field electric field distribution over the testing surface is unfortunately not possible within this study. The two probes inserted in the numerical solver to generate the findings of figure 5.18 already pushed the solver to its limits. It is left as future work.

In figure 5.18, the boresight impulse has a rise time of 4.95 ns and a FWHM of 4.88 ns. The 7.07 m off-boresight pulse has a risetime of 5.14 ns and a FWHM of 5.17 ns. As a reminder; in figure 5.14, which was the far field counterpart, the 30 m pulse had a 2.4 ns rise time and a 2.6 ns FWHM. The peak field strength recorded by the E-field probe in the boresight is just 7.97 kV/m, compared to 12.93 kV/m recorded by the far field method. The lower peak E field strength may be attributed to the wider pulse, both in time domain as well as in the directivity sense.

This means that for this specific case, the 30 m model is no better than the 12.2 m model regarding peak electric field strength. Its only improvement is its better low frequency operating limit, but it pays for this with a loss in time domain fidelity caused by near-field dispersion.

The DEXP-integral fed result was also inspected. The main impulse at boresight has a rise-time of 2.95 ns and a FWHM of 18.32 ns, as opposed to 1.9 ns and 26.6 ns respectively for the far field estimate. The factor difference between the near-field and far-field estimates are not as significant as they were for the DEXP fed wave, but a degree of pulse dispersion has still occurred. The off-boresight point had a rise-time of 2.88 ns and a FWHM of 27.66 ns, occurring just 0.44 ns after the boresight impulse. The off-boresight had a peak value of 0.71 with respect to the main impulse.

The effect of near-field dispersion is significant. The pulse is widened by almost a factor 2, while the peak field strength is reduced by a factor 0.6, with the measuring point a factor 6 away from the far field condition (far field is achieved after 600 m for a 2.5 ns rise time). It may be concluded that simply assuming far field conditions of this antenna at 100 m will cause a significant discrepancy.

The pulse widening is the cause of the time difference of arrival between contributions originating from different parts of the large dish. This widening is also the cause for the lower field strength, as the total pulse energy is smeared over a longer time. The gain of the main lobe is not yet finalized and increasing with distance from the antenna.

Finally, for validation of the discussion done in section 5.2, in which 12.2 m was discussed to be at the limit of the antenna far field, the same near-field analysis has been done for the 12.2 m antenna. This yielded a FWHM of 2.2 ns and a rise time of 0.9 ns, compared to 1.1 ns and 0.8 ns with the far field assumption. There then seems to be a discrepancy between the FWHM's, but upon inspection it was revealed that the 'notch' seen in figure 5.14 in the decay caused an artificially short FWHM. This same notch is present in the near field probe analysis, but it is of slightly different intensity whereby it does not cut-off the FWHM measurement. It is deemed that while cutting it close, the 12.2 m model suffices to the far field assumption.

5.5 Discussion

The goal of this chapter was to simply produce a working model of the reflector IRA in order to gain insight into the effect of directivity in radiated NEMP tests, as well as to illustrate the trade-off between antenna practicality and lower frequency operating limit. It was explicitly not the goal to design the definitive, optimized reflector IRA that suits the needs of a radiated NEMP tests perfectly. This section aims to reflect on the work done in this chapter to achieve that goal, and to address the research gaps between this working model and the definitive IRA. Furthermore, the discussion regarding the trade-off between the practicality, driven mainly by the dish diameter, and the lower frequency operating limit is continued in this section.

The main takeaway from the quantification of the performance of the EMIS-III-VPD is that the antenna is simply too small to meaningfully radiate any electromagnetic energy at frequencies as low as 1 MHz. This chapter has not changed that; despite the improved directivity of the reflector IRA with respect to the omnidirectional monocone it still would need to be enormous to effectively radiate the very lowest frequencies. Therefore, a choice remains to be made about the point at which increased impracticality no longer justifies a lower low-frequency operating limit. This chapter has illuminated what the consequences of different choices are by presenting three different sizes of reflector IRA models. Future work may examine to what extent the 1 MHz lower limit selected for this thesis is necessary, or whether a higher low frequency bound will suffice.

5.5.1 Antenna performance

Not radiating the very lowest frequencies has an effect on peak electric field strength and on the time domain fidelity of the pulse. The peak electric field strength drops when the lower frequency operating limit is increased as that is precisely where most of the energy of the radiated pulse is contained. The pulse shape is altered in two distinct ways by the radiating properties of the reflector IRA. The first, as already noted, is the loss of low-frequency content that arises from the antenna being too small with respect to the very lowest wavelengths of the desired pulse, altering the pulseshape. The second is the differentiating action of the IRA itself: as discussed in chapter 4 and confirmed numerically in section 5.4.3, the time-domain field is approximately the derivative of the voltage fed at the feed gap. Feeding the standard DEXP directly therefore produces a sharp, narrow pulse with a FWHM in the order of 1 to 2.5 ns across the three models, far short of the 23 ns required by the standards. Feeding the integral of the DEXP, derived in section 5.3.2, recovered approximately the desired unipolar shape with FWHM values between 15 and 33 ns. This does however come at the cost of a strongly amplified prepulse, in some cases exceeding the main impulse in intensity. The prepulse is attributed to low-frequency content radiated directly from the feed arms before the dish has had time to focus the bulk of the pulse onto the boresight observer, and is most severe for the smaller models that radiate this low-frequency content inefficiently anyway. Further tuning of the feed waveform is therefore required before the radiated pulse can be considered fully compliant with the standards; this may take the form of pre-distortion of the feed waveform or modification of the antenna geometry to suppress direct radiation from the arms. This is left as future work.

The reflector IRA's achieve better low frequency performance relative to their size than the EMIS-III-VPD monocone did. This is a consequence of the antenna directivity, which shifts the spectral density lines presented in figure 5.13 upwards as more electromagnetic energy is radiated in the desired direction. As a comparison, the EMIS-III-VPD is a half of a bicone 40 m tall and radiated from 6 MHz upwards within the required spectral density (and dropped out of the green range between 50 MHz and 300 MHz). The 30 m reflector IRA, electrically three-quarters the size of the EMIS-III-VPD, radiated within the limit from 4 MHz onwards. Higher directivity improves spectral density without altering antenna size.

All three reflector IRA's showed improving directivity with increasing frequency. A consequence of this is that from a certain frequency onwards the spectral content increases up to a point where it exceeds the 10 dB tolerance of the NEMP. This is not an issue, as this plot was generated using the realized gain of the antenna, meaning the traces shown in figure 5.13 give an upper bound to antenna performance. Feeding a pulse with a spectral content distribution will alter the actual radiation of the antenna. In other words, the pulse has to be tuned to the antenna radiating it to ensure the correct spectral distribution is achieved.

The numerical implementation of the models itself imposed further limitations. The choice to use a single reference geometry scaled to three different sizes meant that the meshing density, expressed in cells per wavelength at 1 GHz, was not uniform across the three models. For the larger dishes the cell size grows linearly with the scaling factor while the highest frequency of interest remains fixed, leading to fewer cells per wavelength on the dish surface. This manifested most clearly in the 30 m model, where the radiation efficiency dropped from approximately 0.9 in the 30 to 350 MHz band to as low as 0.2 above 850 MHz despite a good broadband S11 match. The numerical staircasing of the parabolic dish surface, even with the Perfect Boundary Approximation employed by CST, is thought to be the dominant cause. For a definitive design, separate models with mesh densities scaled to the dish size should be used. This too is left as future work.

5.5.2 Near-field considerations

A separate concern that emerged from this chapter is the validity of the far-field assumption at the chosen 100 m test distance. The time-domain criteria of section 5.2 place the 30 m model firmly outside its far field at this range, and the 12.2 m model was chosen exactly at the Baum far-field limit. The implication is that the simulated rise time and FWHM of the largest model carry a near-field dispersion contribution that is not represented in the far-field monitor output. This was made tangible in the near field analysis section, in which the time domain fidelity of the near field probe was compared to the far field estimates. A widening with a factor 2 and an almost halving of the peak electric field strength was seen. This leads to the conclusion that an increase in antenna size brings no improvement of antenna performance except for a lower frequency operating limit.

A full near field analysis was unfortunately not possible within this study, so it was not possible to examine the near field energy electric field distribution of the 30 m model. It is suggested that the near field dispersion that causes pulse widening in time domain may also cause main lobe widening, also causes better homogeneity over the narrow 10 m x 10 m cross-section of the main lobe. A full 2D near field analysis of the testing surface must be conducted in order to explore this idea further.

The near field dispersion should not be shunned or avoided in as an effect. With the antenna size required for a lower frequency bound of 1 MHz in combination with the peak field strength required means that an eventual NEMP radiated test may choose to remain in the near-field by design. It is certainly possible to estimate the effect of it and pre-distort the feeding signal in order to counteract its effects, while at the same time taking advantage of the suggested improved homogeneity. This may make an increasing size of antenna beneficial to the performance of the IRA as a radiating NEMP simulator, which is not a conclusion that is supported by this chapter. The potential design and evaluation of effectiveness of this is left as future work.

5.5.3 Reflection on methodology

As the goal was only to obtain a working model, and explicitly not to optimise the antenna for radiated NEMP testing, it was chosen to take an existing, working design and implement it into the numerical EM solving environment. This has saved a significant amount of time and kept the research work within its scope. However, it means the results presented in this chapter and the next are based on unoptimized antenna design. The use of a different dish shape, different feed arm geometry, a different feeding pulse may give results significantly better than achieved in this thesis. This is left as future work.

The performance estimates of this chapter must also be interpreted as optimistic. They were generated under the assumption of the perfect 500 kV pulse also fed to the EMIS-III-VPD was delivered to the feed gap. A real pulser will fall short of this in both peak field strength as well as time domain fidelity. The peak field strengths in table 5.5 should therefore be read as upper bounds, and the time-domain fidelity reported is best-achievable rather than typical. The implications of realistic pulser performance, in particular the timing jitter between pulsers in an array, will be returned to in chapter 6.

Another caveat to be mentioned here is the fact that the 30 m model had a lower mesh density, meaning that any qualitative results above 500 MHz cannot be taken with a high degree of certainty. Future work may seek for solutions that make the numerical solver more efficient, or somehow mitigate the lower mesh density.

Perhaps the most useful finding of this chapter for the work that follows is the homogeneity analysis of section 5.4. While a single reflector IRA does not illuminate the $10\text{ m} \times 10\text{ m}$ testing surface within the $-0/+6\text{ dB}$ tolerance at the higher frequencies, it does meet this tolerance over a $5\text{ m} \times 5\text{ m}$ sub-area centred on its boresight. This is a direct consequence of the geometric match between the narrowest 3 dB beamwidth of the antenna models (2°) and the angular extent of a $5\text{ m} \times 5\text{ m}$ surface at 100 m, which corresponds to 2° off boresight at the corner. The reflector IRA element is therefore well-suited to illuminating a $5\text{ m} \times 5\text{ m}$ area within tolerance, and the question of how to extend coverage to the full $10\text{ m} \times 10\text{ m}$ specification becomes one of tiling four such areas with four antennas. This forms the natural starting point for the array design of chapter 6.

5.5.4 Antenna practicality

It is a goal of this thesis to present a final result that is not completely infeasible. It is envisioned that this is in some part a feasibility study, which means that feasibility is a goal set for this thesis. However, it is not up to the author of this thesis to qualify what makes a feasible or infeasible design. Therefore, as was done in more detail in this chapter for the trade-off between antenna practicality and lower frequency operating limit, this subsection will examine what further steps may be taken to improve the feasibility of any of the three designs presented in this chapter.

The central issue with radiated NEMP testing is the size of the antenna required. The weight, high frequency tolerances and supporting construction considerations are significant. One effective and quick solution may be to halve the antenna and place it above an electrically conducting plane, effectively creating a HIRA. This has precedent, as this was the design choice made for the EMIS-III-VPD as well. There, the electrically conducting plane was implemented using a metallic mesh that was placed between the feed point and the DUT. Besides halving the antenna size, the HIRA would also no longer require a complex and aperture blocking construction to place the feeding network high up in the dish. Instead, it could be fed from under the electrically conducting plane, further reducing design complexity and unpracticality.

The central issue with the HIRA is that the desire to radiate both horizontally and vertically polarized waveforms has to be given up to some degree. Having both a vertically and a horizontally polarized waveform requires two pairs of feedarms to be orthogonal. However, with the HIRA, the horizontally oriented set of arms will be removed, as that is where the electrically conducting plane will dissect the dish. Near-horizontally polarized may still be possible, or a full 45° rotation, but the clean H- and V-polarizations are lost.

The practicality of the antenna may be further improved by the use of a non-solid dish. In the model, the dish is an infinitely thin, PEC analytical sheet. Electrically, there is no requirement for the sheet to be solid, so long as the density of the meshing is sufficient with respect to the upper frequency operating limit. At the 1 GHz (0.3 m wavelength) the density must be around one tenth of the wavelength, resulting in approximately 0.03 m wire spacing. For low frequency applications, wire or meshing of sufficient density is used to reduce weight and wind-load of the dish. Another advantage is that the potential risk of heat damage from bundling sunlight at the feed point is reduced to near zero by implementing meshing. This too could be applied to the reflector IRA, reducing weight and manufacturing cost significantly.

Implementing a wire or meshing solution on the dish may be convenient if transportability is a requirement important to the design. With a wire or mesh dish, the dish can potentially be made to be collapsible, a design feature showed in one of the reflector IRA's from the Sensor and Simulation Notes [41]. This would allow the operator of the simulator to make use of temporary facilities, which may aid in the practicality of the simulator.

5.6 Conclusion

The goal of this chapter was to obtain working reflector IRA model which could then be used to both estimate the potential performance of such an antenna and to shed light on the major trade-off between antenna practicality and lower frequency operating limit, which was a key problem for the EMIS-III-VPD baseline of this thesis. A summary of the results achieved by this chapter are presented in the table below:

Table 5.6: The goals, baseline and results of this chapter summarized

Specification	Req.	EMIS-III-VPD	5 m IRA	12.2 m IRA	30 m IRA
Peak E-field [kV/m]	5 -0 + 6 dB	1.28	2.65	8.10	7.97*
Rise time [ns]	2.5 ± 0.5	2.45	5.5	2.7	2.95*
FWHM [ns]	23 ± 5	24.5	15.4	33.4	18.32*
Lower freq. bound	1 MHz	~6 MHz	~33 MHz	~12 MHz	~4 MHz
Upper freq. bound	1 GHz	~1 GHz	~1 GHz	~1 GHz	~1 GHz
Spectral tol.	NEMP ± 10 dB	exceeded 50–300 MHz	within or above	within or above	within or above
Polarization	H and V	V only	H and V	H and V	H and V
Testing surface area	10 m $\times$ 10 m	200 m $\times$ 50 m	10 m $\times$ 10 m	10 m $\times$ 10 m	10 m $\times$ 10 m
Test distance from radiator	100 m	100 m	100 m	100 m	100 m
Field homogeneity	-0/ + 6 dB	9.9 dB avg.	meets (5 $\times$ 5)	meets (5 $\times$ 5)	–

* Probed near-field values.

Table 5.6 presents the quantification of results gathered in this thesis. The leftmost two columns shows the specifications set to be achieved by this thesis, which result from chapter 2. The EMIS-III-VPD results, obtained in chapter 3, are shown in the next column. This allows direct comparison with the three reflector IRA models presented in this chapter.

The most important improvement achieved with respect to the EMIS-III-VPD is the addition of directionality of the IRA. This manifests itself in improved peak field strength, improved lower frequency operating limit (at an equally sized antenna), and improved homogeneity, or rather more predictable field intensity distribution.

Another conclusion from this thesis is the lack of increased antenna performance with increased antenna size when the testing surface is in the near-field. Near-field dispersion causes a drop in peak electric field strength as well as pulse stretching in time domain. Therefore, increasing antenna size will only improve performance if the testing surface is in the far field. If the testing surface distance is fixed, the best performing antenna is the largest that still satisfies far-field conditions. Further research may look to mitigation or adaptation mechanisms that overcome the loss caused by near-field dispersion.

While achieving the desired field strength at a single point in the boresight of the dish, the IRAs do not yet illuminate the desired testing surface with sufficient homogeneity. The array, in which one of these IRA's will be the element, will be specifically designed to improve homogeneity over the testing surface.

Chapter 6

Array Design

The goal of this chapter is to answer the central question this thesis is aimed at answering: What is needed for a time-domain array to generate an NEMP-like pulse over an area of 10 m x 10 m at 100 m distance such that the pulse in that area complies with the tolerances set in Chapter 2 of this thesis? The previous chapters have given the building blocks to answer this question properly: a baseline performance has been set, an antenna has been selected to serve as an element and this antenna element has been implemented in CST and its performance charted numerically. Now this chapter will explore what an array of these antenna will look like and what is necessary to implement such a structure in real life. The tolerances of for instance the feeding network required to achieve homogeneity will speak to the feasibility of this approach to NEMP testing. It is expected that besides a list of requirements this chapter, as previous chapters did as well, will raise more questions that must be addressed in future work.

It is the goal of this chapter to implement such an array. The goal is explicitly not to finalize the design, there will not be an optimal and final antenna array design for NEMP simulation at the end of this chapter. Instead, insight will be developed on what the performance of such an array might be, as well as what constrains its performance and what is needed for practical implementation. Therefore, simplifications will be made in order to obtain a working model. This working model can always be expanded upon at a later moment in time, potentially in future work. At the end of this chapter, one may say make a substantiated judgement to pursue this approach to NEMP testing, or to disqualify it as infeasible or underperforming.

First, the results of the working antenna are implemented in MATLAB. This environment allows for quicker and easier calculations on the radiated far field of the antenna, negating the need for a full rerun of the EM time-domain numerical solver. Furthermore, the previous chapter had no definitive idealized feeding voltage to generate the desired NEMP-like pulse in the far field. By implementing this model in MATLAB, this pulse may be obtained.

This chapter starts with a manually synthesized array. Obtaining a working model where the full path from antenna feed to testing surface is defined will give insights into what the drivers and considerations of successful antenna synthesis is, as well as providing the foundation upon which further improvement may be found.

Next, the foundational manually synthesized model is improved upon by reformulating the antenna array synthesis as an optimization problem. The non-optimal, manually designed array would give a pessimistic outlook on the performance of the array. Optimizing the array performance will give insight into what may be achieved by pursuing arraying for radiated NEMP testing. Furthermore, this chapter will explore what is necessary from a feeding network such that the array performs as predicted here.

6.1 Single Antenna Model

In the previous chapter, a working CST model of three sizes of the same IRA were developed. While these provided insight into their performance and in the choice between antenna practicality and low frequency operating limit, the full time-domain solver takes hours to run and therefore is clunky and slow for further development in this section. For more efficient further work, the far fields can be exported from CST and imported into MATLAB, which negates the need for the full time domain EM solver to run again. Potential near field arraying, (which holds when larger antennas or a shorter testing distance is used) is left as future work.

As it is the far fields that are needed, the 12.2 m model is selected as antenna element in this section. It is not that this antenna is ideal, but it is the limiting case in which far field constraints still hold, which means that it is the biggest size of antenna given this rise time in which the far fields are valid at 100 m distance.

Another problem with the CST models is that no finalized feed pulse has been discerned. The differentiating property of the IRA led to a run of the models using the integral of the DEXP, but this did not yield ideal results. It was shown that the three models exhibited significant bipolar behaviour, with a large negative pulse before the main impulse, and in the case of the two smaller models it was larger than the main impulse. Using one of the two feeding pulses presented in Chapter 5 would distort the discussion on the results of the array, as any effects seen may be attributed to either the imperfect feeding pulse or to the array. An optimal feeding pulse is therefore desirable to separate arraying effects from feeding pulse effects.

The chosen method for this chapter is to assume that the antenna behaves as a linear, time invariant (LTI) system, in which the complex spectrum of the feeding voltage is related to the complex spectrum of the electric field strength at a certain point in the far field by a transfer function. Transforming to time domain, the IFFT of the transfer function becomes the impulse response of the antenna plus the path. This allows for quick and light calculation loads while simultaneously allows determining the ideal feeding pulse to achieve the best possible result for matching the desired NEMP like pulse given the impulse response of the band limited antenna.

6.1.1 Single antenna impulse response

In implementing the exported far field results from CST into MATLAB, it is assumed that the transfer function from feeding voltage to the E field at a certain point at a certain frequency behaves as an LTI system. This way, for any point p with coordinates (x,y,z) , the electric field strength at that point, E_p is given by the multiplication of the feeding voltage $V_{standard}$ and the point dependant transfer function H_p encapsulating the effects of the antenna and the propagation path between feeding voltage and the point p :

$$E_p = H_p V_{standard} \quad (6.1)$$

This transfer function will be used extensively throughout this chapter. To obtain that transfer function, a deconvolution would have to take place between the feeding voltage and the electric field strength at that point. However, with the standard feeding voltage of CST, this deconvolution already happens internally, meaning the exported result already is the transfer function at that frequency:

$$E_{pCST} = E_p / V_{standard} = H_p \quad (6.2)$$

To visualize the full spectrum homogeneity, the same 2D cuts used to generate the figure in section 5.4.4 are taken at 5 MHz intervals over the relevant frequency spectrum. These 2D cuts are per-frequency far-field exports from the CST model of the 12.2 m IRA. The far field results are generated by the standard CST feeding voltage, a broadband Gaussian pulse. A useful feature of this pulse is that when it is used, the solver will automatically normalize the frequency domain result to the complex spectrum of the standard feeding pulse [12]. This means that any result obtained has units 1/m, as opposed to V/m, as the unit of field strength. Practically, this means that the result is the transfer function of the antenna for an arbitrary feed.

These 2D cuts are imported into MATLAB, and linear interpolation is done for each point on the surface with 0.5 MHz interval. This makes 2000 frequencies for which there are 2D cuts. For each point on

the 2D surface, the complex spectrum is split into a two-sided spectrum and the Inverse Fast Fourier Transform (IFFT) is taken. This gives a time domain signal for each point on the surface, which can be seen as a band-limited impulse response of the antenna ($h_p(t)$). For added clarity, the process is illustrated graphically below:




Figure 6.1: The process of obtaining the antenna impulse response

Figure 6.1 shows graphically how the antenna impulse response is obtained. The leftmost figure shows a single 10 m x 10 m cut at 100 m at a single frequency, the antenna transfer function. A certain point on that 2D cut is highlighted. For that certain point, all frequency cuts may be placed in a single array, shown in the next figure, where all 2D cuts can be seen to be 'stacked' on top of each other in order of frequency. The IFFT of this array is the (band limited) time domain result, in other words the (band-limited) impulse response. To visualize this, which will be done in figure 6.2, the maximum value of that impulse response for that point is taken and consequently plotted

The intensity of each point on this time domain signal is corrected for the difference in propagation distance between boresight and the actual point. A CST far-field cut is angular: every exported point lies on a sphere of radius 100 m around the antenna. The testing surface however is a flat plane at $z=100$ m, its points lie at slightly larger distances (up to 0.25 m further away in the corners of the 10 m x 10 m area). The far-field $1/r$ converts each constant-radius pattern value to the point's true distance. The peak of each time domain signal is recorded, transformed into dB loss with respect to the highest value of the surface. This highest value happens to be on the boresight point. The result of this is shown on the next page in dB with respect to the highest value.

There are two things to note. Firstly, this heatmap shows the peak E field strength of the impulse response of the 12.2 m antenna in dB. It is not an instantaneous 2D cut of the testing surface as the pulse is propagating through that area. Next, the heatmap in figure 6.2 is irrespective of the spectral density of the pulse fed to it. It shows the contribution of each radiation pattern of each of the relevant frequencies equally. That is why it is shown to have such an intense and narrow peak, with the -3 dB area having a radius of approximately 2.5 m. This is because from approximately 25 MHz onwards, the wavelength is in the same order of magnitude as the antenna dish size, meaning the directivity is large and the main bundle is narrow, meaning all the contributions from that frequency onwards show a narrow main bundle. Only two effects therefore are visible in the impulse response; the bundling of energy into the boresight direction and band-pass filtering effect of the antenna, as discussed in Chapter 2. This impulse response describes how the antenna and propagation path transform an arbitrary input signal. In frequency-domain terms, the field at each point is obtained by multiplying the local antenna response by the chosen feed spectrum, after which an inverse Fourier transform gives the corresponding time-domain pulse.




Figure 6.2: The impulse response of the 12.2 m IRA over a 10 x 10 m surface in boresight at 100 m

6.1.2 Idealized feeding pulse

Plotting the antenna peak impulse response irrespective of spectral magnitude, as in figure 6.2, is not representative for the use-case of this thesis. That figure weights each frequency equally, whereas in reality almost all of the energy of the NEMP sits in the lowest part of the relevant spectrum, the 1 to 100 MHz band, which is precisely the part the dish either does not bundle, or bundles only inefficiently. To judge the field the antenna would actually radiate, it must therefore be excited with a realistic feed waveform. In this section, an idealized feed pulse is distilled from the antenna impulse response. It will be the feed pulse that is optimal only for the boresight point. Optimizing for the whole area simultaneously will be left as future work. The approach taken here proceeds in two stages. First, a feed waveform is derived that drives the single antenna to radiate as close to the desired NEMP pulse as the band-limited dish allows. This waveform is fed to the full-wave CST model. Second, the radiated boresight waveform that CST produces is taken as the ground truth, and the field across the full testing surface is reconstructed by anchoring to it. This section describes both stages in turn.

The first stage begins by characterising the antenna itself with a 500 kV feed. Treating the antenna and its propagation path as an LTI system, the radiated boresight field is related to the feed by a transfer function $H_{bs}(f)$. The patterns were computed in CST under a specific port excitation $V_{exc}(t)$ (the CST standard pulse but with 500 kV feed value), so the raw pattern data still carries the spectrum of that excitation. To isolate the response of the antenna alone, the excitation spectrum is divided out, $H(f) = E_{pat}(f)/V_{exc}(f)$, after which the boresight value $H_{bs}(f)$ is obtained. This transfer function $H_{bs}(f)$ is the same quantity that is visualized in figure 6.2, with the notable difference that in that figure the maximum of the time domain result was shown, while here the frequency domain values are used. What remains for this stage is the choice of the desired radiated waveform. The natural target is the analytical double-exponential of Chapter 2,

$$E_{DEXP}(t) = k (e^{-\alpha t} - e^{-\beta t}) u(t),$$

whose rise is set by $1/\beta$ and whose slow tail is set by $1/\alpha$.

The choice of β deserves comment. A faster target rise demands more high-frequency content in the feed, and that high-frequency content is exactly what drives the direct radiation from the IRA feed lines, which appears as a sharp prepulse ahead of the main lobe in the radiated boresight field. To shift energy out of this prepulse and into the smoother main lobe, β is lowered from its original value to 3×10^8 1/s, corresponding to a target rise time of roughly 7 ns. This value for β was found by iterative experimentations in running through this process, and may therefore not be the ideal value yet. This is a deliberate trade between prepulse suppression and main-lobe sharpness, and it ties directly to the differentiating behaviour of the reflector IRA discussed in section 5.4.3. As we will see, the induced sluggishness will be better than expected.

With both the transfer function and the target in hand, the required feed spectrum follows from inverse filtering, which in its ideal form would be $V_{\text{feed}}(f) = E_{\text{DEXP}}(f)/H_{\text{bs}}(f)$. A direct division is, however, not usable: wherever $|H_{\text{bs}}(f)|$ is small the antenna radiates inefficiently, and a direct inversion would call for an unrealistically large feed amplitude purely to overcome that inefficiency. The inversion is therefore regularised with a Tikhonov-type filter, inspired by [19] and [20],

$$V_{\text{feed}}(f) = E_{\text{DEXP}}(f) \frac{H_{\text{bs}}^*(f)}{|H_{\text{bs}}(f)|^2 + \epsilon}, \quad \epsilon = (0.05 \max |H_{\text{bs}}(f)|)^2.$$

Frequencies at which the boresight response falls below 5% of its maximum are thereby suppressed rather than inverted, so the method reproduces the DEXP target only over the band the antenna can actually radiate and yields a band-limited approximation elsewhere. This prevents the calculated $V_{\text{feed}}(f)$ to contain massive spectral content in the bands at which the antenna is inefficient in an attempt to overcome it. While that would mathematically be correct, from an engineering viewpoint that is meaningless. The feed waveform obtained from the inverse transform is, finally, scaled so that its peak does not exceed 500 kV, the limit of the available pulse generator that has been assumed to work perfectly throughout this thesis. The resulting waveform is shown in figure 6.3 below:




Figure 6.3: Idealized feed waveform required to radiate the (band-limited) analytical DEXP in boresight

It should be stressed that this is a numerical construct only; the feasibility of generating it with a real pulser is outside the scope of this thesis and is left as future work. For reference, feeding the raw DEXP shape directly to the antenna, without inverse filtering and at the same 500 kV cap, produces a noticeably lower and more distorted boresight pulse, which confirms that the inverse filter is doing useful work.

This feed waveform is then imported into the CST model of the 12.2 m IRA, and the time-domain solver is run. The result is the radiated boresight waveform $E_{\text{CST,bs}}(t)$, which is taken as the ground truth for the remainder of the analysis. Being the output of a causal full-wave solver, it carries the correct phase and includes effects the analytical treatment cannot capture, such as the residual feed-line prepulse, edge diffraction and any reverberation. In effect, $E_{\text{CST,bs}}(t)$ is the band-limited DEXP that the 12.2 m dish genuinely radiates in boresight, imperfections and all.

The second stage of this process uses this measured waveform to reconstruct the field across the full testing surface whereby the reconstruction is anchored directly to $E_{\text{CST},\text{bs}}(t)$, as the CST result may be considered the ground truth for the purposes of this chapter. The first stage of this process produced two results that the second stage builds on. The first is the ideal feeding pulse itself. The second is the field that this pulse actually radiates at boresight, $E_{\text{CST},\text{bs}}(t)$, obtained by loading the ideal pulse back into CST and reading off the boresight probe. Because this waveform comes straight out of the full time-domain solver, without any of the assumptions made when the work is moved into MATLAB, it is the most reliable field available anywhere in this chapter, and for that reason it is treated as the ground truth from here on.

The difficulty is that this ground truth exists at only a single point. The aim of the chapter is to describe the field over the entire 10×10 m testing surface, and obtaining $E_{\text{CST},\text{bs}}$ at boresight required a full solve. Repeating that solve for every point on the surface would bring back exactly the hours-long runtimes that the move to MATLAB was meant to avoid, so this route is not practical. What is needed instead is a way of taking the one reliable waveform at boresight and carrying it out to every other point on the surface without solving the full problem again. The information needed to do this is already available in the 2D far field cuts $E_{\text{pat}}(P, f)$ that were imported earlier. These cuts describe how the radiated field changes from one direction to the next across the surface. On their own, however, they are not the field that is sought: they were produced by CST's standard pulse rather than by the ideal feed. The cuts therefore contain the correct shape of the pattern across the surface, but neither the correct drive nor the correct scaling. The reconstruction takes the shape from the cuts and takes the drive and the scaling from the boresight ground truth, and combines the two.

The combination starts by asking a simple question: what feed would have been required to turn the boresight value of the pattern into the measured boresight waveform? Writing the boresight value of the pattern as $E_{\text{pat}}(0, 0, f)$, where $(0, 0)$ are the coordinates of the centre of the 2D cut, this effective feed is just the ratio of the two,

$$\text{feed}_{\text{eff}}(f) = \frac{E_{\text{CST},\text{bs}}(f)}{E_{\text{pat}}(0, 0, f)}. \quad (6.3)$$

At a few frequencies the denominator passes through a true null or drops away at the edge of the band, where this division would otherwise blow up. To keep it finite, $|E_{\text{pat}}(0, 0, f)|$ is floored at 1% of its maximum at those few bins, with its phase left untouched. This floor is purely a numerical guard and has nothing to do with the 5% Tikhonov regularisation used when deriving the feed in the first stage; the two act on different quantities for different reasons and should not be confused.

With this effective feed in hand, the field at any other surface point P follows from the same LTI product used throughout the chapter, now with $E_{\text{pat}}(P, f)$ in place of the boresight value,

$$E(P, f) = E_{\text{pat}}(P, f) \text{feed}_{\text{eff}}(f) = E_{\text{CST},\text{bs}}(f) \frac{E_{\text{pat}}(P, f)}{E_{\text{pat}}(0, 0, f)}. \quad (6.4)$$

The second form shows the reconstruction of the field. The reconstructed field at P is simply the measured boresight pulse multiplied, frequency by frequency, by the ratio of the pattern at P to the pattern at the centre. At boresight this ratio is equal to one, the reconstruction returns $E_{\text{CST},\text{bs}}(f)$ unchanged, and the boresight pulse is therefore the measured CST trace by construction. Away from boresight, that same pulse is reshaped according to how much stronger or weaker the pattern is at P , and how much earlier or later it arrives, compared with the centre.

Anchoring the reconstruction to the measured boresight waveform, rather than rebuilding it from the analytical target, is a deliberate choice that follows from the behaviour of an earlier version of the model. That version took the more direct route of synthesising the boresight field by passing the analytical DEXP target through the regularised inverse filter, which works out to $E_{\text{DEXP}}(f) |H_{\text{bs}}(f)|^2 / (|H_{\text{bs}}(f)|^2 + \epsilon)$. This expression is the target multiplied by a real, zero-phase filter, and a zero-phase filter is symmetric in time. It therefore places energy in front of the main pulse, at times before a real, causal antenna would have radiated anything, which appeared as a pronounced negative prepulse ahead of the main impulse. The same $|H_{\text{bs}}|^2 / (|H_{\text{bs}}|^2 + \epsilon)$ factor also pulls the field down wherever the antenna responds weakly, and in earlier tests this was found to suppress the main peak by roughly 38% and to narrow the pulse. Since the

only thing that differs between that earlier version and the present one is what sits at boresight, replacing the synthesised waveform with the real CST solve removes both problems at once: the causality comes straight from the solver, and the floored ratio leaves the boresight pulse at its true height and shape. Every other point on the surface inherits that same corrected behaviour, because every point is built from it.

With $E(P, f)$ known at all frequencies, the time-domain field at each surface point is recovered by the same inverse transform used for the boresight pulse, so a single trustworthy CST solve is enough to produce a reconstructed waveform everywhere on the testing surface. The IFFT is applied and now the time-domain result is obtained at each point. The resulting time domain plot is shown on the next page in figure 6.4, as well as a comparison with the CST results.




Figure 6.4: $E(t)$ on two points of the 10 x 10 m testing surface at 100 m, calculated using the antenna impulse response method and overlaid with the corresponding CST results.

Visually, the MATLAB model neatly overlaps with the CST results in the main impulse, around which the time scale has been centred. The two do deviate in amplitude outside the main impulse, specifically in the post-pulse. The deviation is attributed to the fact that the MATLAB reconstruction is only done for the 1 to 1000 MHz range for which there exist CST results, while the time domain solver encapsulates the full spectral content. This is not considered critical for a qualitative comparison of the time domain behaviour nor for a quantitative comparison of the main impulse results. The CST results do show a notch in the $E(t)$ at -4 ns, which skews the rise time comparison as this notch is not recreated by the MATLAB model.

The reconstruction is validated against the independent CST run, shown in figure 6.4. At boresight the comparison should be, by construction, an identity check, since the model should simply return $E_{\text{CST}, \text{bs}}$. This is nearly the case: the boresight pulse calculated by CST has a peak value of 4.4 kV/m, a FWHM of 5.27 ns and a risetime of 2.01 ns. The MATLAB impulse response method has a boresight pulse with a 4.58 kV/m peak, a 5.33 ns FWHM and a 2.38 ns risetime. The discrepancy is not considered critical, and attributed to the limited spectral content of the MATLAB model.

The boresight pulse metrics may further be compared against the band-pass-filtered DEXP analysis of Chapter 2, table 2.2. Here it was calculated that for a 10 - 1000 MHz BPF, the NEMP has a peak field strength of 29.54 kV/m, a rise time of 2.118 ns and a FWHM of 12.138 ns. Ignoring the peak field strength for now, the rise time seems acceptable, but the FWHM is significantly shorter. The 10 - 1000 MHz BPF is the selected comparison as figure 5.13 suggests the 12.2 m antenna is radiating efficiently from 12 MHz onwards; that is to say, within the ± 10 dB tolerance. However, in reality, the FWHM is much more similar to a value between the BPF with a lower limit of 20 MHz (8 ns) or even 50 MHz (4 ns). This implies that while the antenna may be within the ± 10 dB tolerance, that is not to say the antenna is radiating efficiently. As the Tikhonov-type filter is explicitly implemented to prevent overcompensation of weak frequencies feed voltage, this is not corrected by the feed.

The meaningful test of the angular-ratio propagation is the off-axis comparison, for which the model prediction at the corner of the surface is compared against a separate CST trace recorded at 4.04° off boresight, the angle the corner point of the testing surface makes with the centre point at 100 m. Over the main impulse, the model reproduces the CST corner trace with a peak of 4.34 kV/m against 4.27 kV/m, a FWHM of 5.65 ns against 5.65 ns, and a rise time of 2.72 ns against 3.20 ns. The discrepancy may be attributed again attributed to the missing frequency content, but another effect may also play a role; different frequencies are affected differently by moving off-boresight. This was shown in figure 5.16, where the high frequency content responded more severely to angular shifts than the low frequency content. This is somewhat captured by the interpolated 2D cuts from the previous section, but the interpolation is not linear in reality. To truly resolve this there must be real data between measuring points, not interpolated data. For the purposes of this thesis however this is sufficient.

6.1.3 Single antenna weighted impulse response

The method described in the previous subsection for calculating the time domain response of the antenna on two points of the testing surface using the spectral content of an idealized feed and the impulse response of the antenna can be expanded to create a full, 2D cut of the homogeneity of a single antenna radiating the idealized pulse. The result is shown below:




Figure 6.5: The 2D E field cut of the 12.2 m IRA over a 10 x 10 m surface in boresight at 100 m, scaled in dB

The effect of using a weighted spectral density is striking. The idealized feeding pulse, which occupies the very lowest portion of the relevant frequency spectrum, concentrates its E field contribution to that frequency band. At those frequencies, the 12.2 m antenna dish is not yet forming the narrow pencil-like beam seen in figure 6.2, and this makes the E field distribution significantly more homogeneous over the 10 m x 10 m area. The maximum difference, found between boresight and any corner point, is just -0.37 dB, whereas the impulse response of the antenna over the entire relevant frequency band resulted in a -6 dB loss with respect to boresight at a radius of 3 m off-boresight.

The peak field strength, found at boresight, is 'just' 4.53 kV/m, which is below the 5 kV/m goal for this thesis. Furthermore, the time difference of the peak value across the testing surface is 0.80 ns (between boresight and corner point). That is more than the rise-time tolerance of 0.5 ns. With an array, it may be attempted to improve on the peak field strength and the time domain homogeneity. Figure 6.5 and its accompanying metrics are then the baseline that any array may try to improve upon. This will be done in the next section.

6.2 Manual array synthesis

Before a more optimized array is designed, an array is first manually assembled. This is done in order to obtain a working script for element contribution and inspection of parameters such as homogeneity and time domain fidelity. This also gives more understanding in what the drivers and limitations of this specific array are. Once this has been finalized, work can be done to optimize the now manually selected parameters.

An array may be designed in an infinite number of combinations. To avoid laboriously going through each possibility, some constraints may be set beforehand. For instance, the physical antenna size prevents antennas from overlapping. Another constraint may be that, considering practicality is desired for this system, the antennas may only be placed on the horizontal plane. This prevents the need for any constructionally complex solutions being needed to suspend an antenna high above the horizontal plane, or worse, below it. Furthermore, it is assumed the antenna and testing surface are floating in an infinite vacuum, eliminating the need to simulate a surface for now.

The antenna selected for this chapter is the 12.2 m dish diameter IRA presented in the previous chapter. This antenna was explicitly chosen as the upper limit of antenna diameter for complying to the far field criterion at 100 m, the distance at which the testing surface will be from the array. It was verified to satisfy far field conditions at this point by comparing electric field probes with far field calculations, which were found to have minimal differences. Array synthesis in the far field is significantly simpler than in the near-field, saving on problem formulation complexity and on simulation runtime. A larger antenna does not necessarily bring an improvement on radiated field performance, as shown in the previous chapter.

A logical starting point for the antenna array design following chapter 5 is to use the array to improve the homogeneity across the 10 m x 10 m testing surface. The antenna does achieve the 5 kV/m electric field strength goal comfortably in boresight, but in the heatmap analysis done in section 5.4.4 it was shown that a single 12.2 m dish does not achieve homogeneity across the full area as its beam is too narrow at high frequencies. However, its narrowest possible beamwidth of 2° (occurring at the highest relevant frequency of 1 GHz) happens to correspond to homogeneously radiating a 5 m x 5 m area, with at most a -3 dB difference between the boresight and corner point of that surface at that frequency.

6.2.1 4 antenna linear array

To improve the figure 6.5, one may initially synthesize a linear array of 4 antennas, each placed with 12.2 m distance between their centres so they are placed dish-edge to dish-edge, 6.1 and 18.2 m off the boresight direction axis that is the centre of the array. Each of the antennas is aimed at a different quadrant on the 10 m x 10 m testing surface 100 m away, and timing differences are imposed such that each pulse arrives at the respective quadrant at the same time. No amplitude tapering is done over the arrays for this initial synthesis. The antennas in the array are all fed with the idealized pulse as obtained in the last section. This yields the following heat map of electric field strength at 100 m:




Figure 6.6: Heatmap of normalized peak electric field strength at 100 m generated by a 4 element linear array of 12.2 m dish diameter antennas, each antenna aimed at a quadrant of the 10 x 10 m testing surface

Figure 6.6 shows the heatmap of the normalized peak electric field strength over the 10 x 10 m testing area at 100 m distance. The point at which each antenna is 'aimed' is shown by the white crosshairs. The intended 5 m x 5 m quadrants at which each antenna in the array is aimed at is shown by the white dashed line.

The figure has been synthesized with the same method applied to figure 6.5; the idealized feeding pulse spectral content multiplied with the antenna impulse response, and the IFFT'd back into time domain response. The difference here is that the contributions from four antennas are added together at each point after the conversion to time domain.

The intensity pattern shows a vertical band of strong intensity through the middle of the plot, decaying as you move horizontally outwards. The pattern is symmetric in the $x=0$ line. This corresponds to the electric field originating from the IRA aimed at both the left and right quadrant constructively interfering in the middle. The result is a peak field intensity of 17.39 kV/m, occurring just above the very centre point of the surface. This off-centre maximum value is caused by the fact that the outer IRA's are pointing at the bottom two quadrants, while the top IRA's are pointing to the top two quadrants. The added path-length causes a slightly lower total contribution by these two IRA's, resulting in a slightly skewed peak value point.

The peak value is almost 4 times the peak value of a single antenna, and in fact exceeds the goal of this thesis (5 kV/m -0 +6 dB) as defined in chapter 2. This non-critical for now as this goal had been dropped from the 50 kV/m defined in the standards.

The corner values in the bottom both have the lowest peak electric field amplitude; 11.534 kV/m. This corresponds to a total range over the area of 5.858 kV/m, or a 3.57 dB decay. This is a significantly larger range between peak and through than a single antenna radiated.

The heatmap of figure 6.6 reports only the peak amplitude of the time-domain waveform synthesised at each surface point, leaving the full waveform itself unused. Since NEMP testing is concerned with the shape of the radiated pulse as well as its peak value, the same per-point IFFT can be exploited further: for every point on the $10\text{ m} \times 10\text{ m}$ surface the complete $E(\vec{r}_s, t)$ is retained, and a small set of waveform metrics is extracted from it.

The simplest of these is the time at which $|E(\vec{r}_s, t)|$ reaches its maximum. Plotted across the testing surface this produces figure 6.7, with the absolute time of the earliest peak anywhere on the surface chosen as the zero reference. The earliest peak occurs slightly off-centre, and the timing spread across the full surface is just 1.56 ns, with a 5th to 95th percentile width of 1.19 ns. This is comfortably shorter than the FWHM of the radiated pulse discussed, but about three times the value of the tolerance within which the rise-time falls. This means that while the pulses over this area are not fully temporally separated, and thereby do contribute to a larger total electric field value (as supported by the almost 4 times larger peak in figure 6.6, they are too far apart to not disturb the rise time. It can be said that the four contributions overlap but are not perfectly aligned: enough overlap for partial coherent summation, but not so well aligned that they sum as a single rigid pulse




Figure 6.7: Relative time of peak $|E(t)|$ across the testing surface for the four-element array with quadrant aiming. The earliest peak anywhere on the surface is taken as the zero reference.

The pattern of peak time over the testing surface is interesting. The two outer sides show the earliest peak points overall, as this section of the testing surface is closer to the outer antennas than the actual centre of the testing surface is. Besides these two areas, the centre of the testing surface has the earliest peaks, which is geometrically logical. The aim points themselves are comparatively late.

The fact that the inner two antennas are aiming at the top two quadrants is even more apparent in this figure. The peak time difference is at its worst in the quadrants that the furthers away antennas are pointing, showing up to 1.5 ns difference here.

Beyond the arrival time, the pulse shape itself is characterised at each surface point by two further metrics: the 10% to 90% rise time and the FWHM of the main impulse. Plotting these over the testing surface gives insight into the time domain fidelity over the testing surface:




Figure 6.8: Time-domain pulse characteristics over the 10×10 m testing surface for the four-element linear array.

The rise-time distribution is similar to the peak value timing distribution presented on the previous page. The implication is that there where the peak arrives early, it is relatively uncompromised by contributions from other antennas, while there where the peak is late, it is in part because the rise time has been smeared out by contributing antennas with different time-difference of arrivals. The rise time is shortest at the very centre of the figure, where it is 3.06 ns, and worst in the two quadrants that the outer antennas are aiming at, where it is 3.69 ns. This gives a spread of 0.62 ns.

The FWHM has a similar distribution to the peak electric field strength homogeneity plot shown in figure 6.6, with the centre band having the lowest FWHM and the outer bands the longest. The lowest recorded FWHM is 5.60 ns, the longest is 8.90 ns. This is a range of just 3.3 ns, significantly more than the range of the rise-time. Note that the wider FWHM on the edges comes from timing smear, not from regaining low-frequency content: the corner pulse is therefore more corrupted, not more NEMP-like.

In frequency domain, the results are also to be inspected, as was also done for the antenna models in Chapter 3 and 5. The layout of the figure will be familiar to readers, and will show the spectral magnitude of the 4-element array on two points of the testing surface: the boresight point and the corner. The NEMP standard at 5 kV/m peak is also added, as well as the 10 dB tolerance around the spectral magnitude allowed by the standards. Figure 6.9 is shown on the next page.




Figure 6.9: Spectral magnitude of the manually synthesized 4 element array

Figure 6.9 shows a deep notch for the boresight point at 500 MHz, dropping the spectral content to well below the 10 dB tolerance. This is attributed to destructive interference at and around that frequency. This notch was not initially detected by inspection of the time domain result at that point, and such a notch in frequency domain is explicitly warned for in Sensor and Simulation Note 538 [21]. The corner point does not have that notch.

The boresight and corner points can be seen to deviate from the 20 MHz point onwards, a consequence of the increased directivity of the antenna from around that frequency onwards. Both the corner and boresight traces can be seen to exceed the 10 dB limit around the NEMP standard, which corresponds to the fact that this set-up achieves 17.39 kV/m, while the NEMP standard remains at 5 kV/m.

Both traces enter the green 10 dB tolerance at 4.5 MHz, significantly better than 12 MHz estimated in chapter 5, figure 5.13. A direct comparison is not quite fair, figure 5.13 was calculated using the realized gain multiplied with an arbitrary 500 kV pulser, while figure 6.9 is the result of a full-wave simulation with an idealized feeding voltage fed to a 500 kV pulser. However, this does speak to the fact that increasing the number of antennas, which increases the total energy radiated, also shifts the spectral magnitude curve upwards and therefore the rising edge intersection with the tolerance to a lower frequency.

Another figure may be made to show the frequency domain performance of the antenna: a cross-section may be cut across the testing surface, from boresight to corner, to show the spectral magnitude as a function of offset of the boresight point. This is plotted on the next page on figure 6.10.




Figure 6.10: Cross-section of the spectral magnitude of the manually synthesized 4 element array

To still show the variations in spectral magnitude between 4 and about 500 MHz, the limit of the colour-scheme was set to -40 dB of the peak. The consequence is that everything below that value appears as zero, even though figure 6.9 showed that there is in fact still energy in those frequencies.

The highest intensity can be seen between 10 and 100 MHz, which is a band that is relatively uniform across the 5 m cut from zero (boresight) to 5 m (edge). The aim point, at $x=2.5$, is also shown as the white dashed line, although in fact the antennas are aimed to $y = \pm 2.5$, and not $y=0$ where this cut is. The uniformity that seems to hold for until approximately 100 MHz no longer holds as well at the higher frequencies: from 100 MHz onwards there is more variation in spectral content across the x axis. Specifically the high frequency content seems to be more present close to the boresight than away from it, with the exception of very high frequency content around the aim point. This would result in more sluggish behaviour of the time domain pulse away from the aim points or the boresight.

6.2.2 Observations

With a first linear array model constructed, some initial observations may be made. Firstly, there is a trade-off between the electric field peak value and homogeneity over the testing surface. A single radiator provided 0.37 dB decay over the entire surface, while 4 antennas aimed at the 4 quadrants resulted in 3.57 dB decay. Aiming all antennas at the centre point again, like the single antenna was, would recover the 0.37 dB decay in peak value over the surface.

The motivation for aiming at the four off-centre quadrant points was to achieve increased time domain homogeneity over the testing surface (the word homogeneity is specifically used here instead of fidelity, as this has already been discussed to be the closest this antenna will get to the true NEMP as possible. The time domain fidelity is coupled to the antenna and its low frequency performance and cannot be recovered with an array). With the four antennas aimed neatly at the four quadrants the expectation was that total time difference between peaks would improve, as well as the spread in time domain metrics such as rise time and FWHM. However, this was not the case. The arrival of components from different antennas at different times resulted in an effect similar to near-field dispersion: the pulses are smeared out in time.

A summary of the performance of the 4 antenna linear array aimed at quadrants is presented below:

Table 6.1: Comparison between the single antenna and the four-element array.

Metric	Goal	Single antenna	Array	Verdict
Peak field, surface max	5 kV/m $-0/+6$ dB	4.53 kV/m	17.39 kV/m	array $\sim 3.85\times$ better
Amplitude homogeneity	–	0.37 dB	3.57 dB	array $\sim 10\times$ worse
Peak-time spread, surface	–	0.80 ns	1.56 ns	array $\sim 2\times$ worse
Rise time, best point		2.38 ns	3.06 ns	array ~ 0.7 ns slower
Rise time, worst point	2.5 ± 0.5 ns	3.20 ns	3.69 ns	array ~ 0.5 ns slower
FWHM, best point		5.33 ns	5.60 ns	array essentially same
FWHM, worst point	23 ± 5 ns	5.65 ns	8.90 ns	array 3.25 ns wider
FWHM range across surface		0.32 ns	3.30 ns	array $\sim 10\times$ more variable

Looking at table 6.1, it would seem that the manually placed 4 element linear array trades time domain integrity, peak-time spread and amplitude homogeneity just for improved peak field amplitude.

The amplitude homogeneity may be solved by making clever use of amplitude tapering over the elements. By limiting the output from the inner antenna elements the top/bottom offset can be reduced. This will have limited effect on the time domain integrity over the testing surface.

More speculatively, the homogeneity over a 10 m x 10 m area may also be improved by tuning the aiming points to outside of the testing surface. If each antenna were aimed at a point in line with the corners of the testing surface, the 10 m x 10 m testing surface contained by the aim points may be more homogeneous than when the antennas aim to within the testing surface.

Finally, it may be hypothesized that the 10 m x 10 m area may be constraining the performance of the 4 element linear array. A single antenna already performed quite well, considering its limited low frequency capability forcing pulse metrics to miss the goals as set by this thesis. By adding to the amount of antennas radiating to a limited surface area, the disparity in peaks and troughs caused by constructive and destructive interference will only increase. Each antenna could instead be used to aim at a 10 m x 10 m area (or larger) itself, the total testing surface might be increased to N times 10 x 10 m (or larger).

This first model has shown that the shortcomings of the 12.2 m dish diameter IRA cannot be resolved with more antennas. Fundamentally, the antenna cannot radiate effectively below approximately 25 MHz, and that is reflected in its time domain fidelity. For the remainder of this work, a selection of parameters must be made that is susceptible to improvement by arraying. The chosen parameters are: peak electric field strength and peak time spread, as well as rise time and FWHM with respect to the single antenna performance.

With the 4 element linear array MATLAB model now shown to be working and producing results we can further explore the potential of the 4-element array by testing more combinations of time-delay, aim points, amplitude weights and antenna locations. By opening up the amount of variables that may be tuned, the optimization of the array becomes laborious. This can be overcome by rephrasing the MATLAB model as an optimization problem, in which the variables of the 4 element array are to be tuned such that time domain and homogeneity metrics are optimized, constrained by the performance of a single antenna. This will be done in the next section.

6.3 Array Optimization

The previous section has given the tools and insight that allow reformulating the array synthesis as an optimization problem. An observation from the previous section was that there exists a three-way tension between peak electric field strength, electric field homogeneity and the spread in peak-field arrival time across the testing surface. A single antenna outperformed the manually implemented array on both field homogeneity and peak-time spread, but with the selected waveform it does not achieve the 5 kV/m goal of this thesis. Adding antennas raises the achievable field strength, but at the cost of worse peak homogeneity and a larger spread in peak time.

An array is therefore necessary to exceed 5 kV/m, given the selected waveform and the assumed feed-generator strength. Treating the 5 kV/m level as a constraint that any candidate design must satisfy, instead of as a quantity to be maximised, reduces the optimization to two competing objectives: the trade-off between peak-field homogeneity and peak-time spread. This reflects the engineering reality that the field strength does not need to be as high as possible; it only needs to clear a threshold, beyond which homogeneity and timing are what matter.

As was shown in figure 6.8, the rise time and FWHM of the radiated pulse are tied to the peak-time spread and peak homogeneity respectively: where contributions from the four antennas arrive misaligned in time, the local pulse is smeared, which simultaneously widens the FWHM and degrades the homogeneity. The rise time and FWHM are consequences of the same arrival-time mismatch that the two chosen objectives already capture, and are strongly correlated with them. They are for that reason not carried as separate objectives. Doing so would add dimensions to the trade-off without adding independent information, and would make the result far harder to interpret, while the underlying alignment is already addressed by minimising the peak-time spread. It must be stressed that this correlation is one-directional and bounded: better alignment reduces time-domain smear broadening, but neither the rise time nor the FWHM can be improved beyond the hard limit set by the antenna's own band-limited response, which was established in section 6.1.1 to be the binding constraint on time-domain fidelity. The rise time and FWHM are instead retained as constraints, bounding the feasible region so that no candidate design is accepted that smears the pulse beyond a set tolerance, even if it scores well on the two objectives.

6.3.1 Optimizing for two competing objectives

The reformulated problem asks for the antenna positions, aiming points, feed amplitudes and feed delays that simultaneously minimise two quantities: the amplitude inhomogeneity (the variation in peak field strength across the testing surface) and the peak-time spread (the variation in pulse-arrival time across that surface). The difficulty is that these two objectives are in conflict. There is in general no single design that minimises both at once (except a single antenna, which does not achieve the minimum field strength); a design that improves homogeneity tends to worsen the timing spread, and vice versa. This is precisely the conflict already seen between the single antenna and the manual array.

When a problem has a single objective, the notion of an optimum is unambiguous: one design is better than another if it scores lower, and the best design is the one with the lowest score. With two conflicting objectives this ordering breaks down. If design A has better homogeneity but worse timing than design B, neither is unambiguously superior; which one is preferable depends on how the two objectives are weighed against each other, a judgement that belongs to the designer rather than to the optimizer. As with the ambiguous 'practicality' of the design, this final choice is outside of the scope of this thesis.

To handle this, the concept of dominance is used. A design is said to dominate another if it is at least as good in both objectives and strictly better in at least one. In other words, a dominated design can be improved in one respect without being made worse in any other, so there is never a reason to prefer it. A design that is not dominated by any other is called Pareto-optimal: it cannot be improved in either homogeneity or timing without giving up something in the other. There is not one such design but a whole set of them, each representing a different compromise between the two objectives. Plotted with one objective on each axis, this set forms a curve known as the Pareto front.

The Pareto front is the central result of a two-objective optimization. Every point on it is a design that is optimal in the sense that it cannot be improved in one objective without cost to the other; moving along the front trades one objective for the other. Crucially, the front does not select a single design. Instead it makes the trade-off explicit and quantitative: it shows exactly how many decibels of homogeneity must be surrendered to gain a given reduction in timing spread, allowing a final design to be chosen with full knowledge of what is being given up. For this thesis the front is the more valuable output, as the goal is to characterise what such an array can and cannot achieve, rather than to commit to one specific design.

The mathematical framework and terminology of multi-objective optimization and Pareto optimality used here follow the standard treatment [22] [23].

6.3.2 MATLAB Genetic-algorithm solver

The Pareto front is computed using a multi-objective genetic algorithm, as implemented in MATLAB's `gamultiobj` routine. A genetic algorithm is a population-based search method inspired by natural selection. Instead of refining a single candidate, it maintains a whole population of candidate designs and improves them over successive iterations, called generations. In each generation the designs are evaluated against the two objectives; the better-performing ones are preferentially retained and combined to produce new candidates (by analogy with biological crossover and mutation), while dominated candidates are discarded. Over many generations the population migrates towards, and spreads out along, the Pareto front [24] [25].

A genetic algorithm is well suited to this problem for two reasons. Firstly, it returns an entire population spread along the Pareto front in a single run, rather than one point at a time, which is exactly the output sought here [25]. Secondly, it makes no assumption that the objectives are smooth or differentiable functions of the design variables. This matters because the objectives here are evaluated through the full broadband time-domain reconstruction of section 6.1.1, including the per-point IFFT and peak detection, which is not available in closed form and need not vary smoothly as antennas are moved or re-aimed. Gradient-based methods, which rely on such smoothness, are poorly suited to this kind of objective, whereas a population-based search is not troubled by it. The cost is that a genetic algorithm is stochastic and does not guarantee that the true Pareto front has been found; it can only be run until the population ceases to improve, which is the convergence criterion adopted here. Concretely this means that any results presented here are not guaranteed to be the true optimum. This is non-critical for this the phase of NEMP research this thesis is in, as it is sufficient for showing general trends. Furthermore, the runtime of the optimizer is constrained by the evaluation of the 2D cut of the field at 100 m on the testing surface given an array configuration, not by the calculation time of the optimizer. Despite already being exported to MATLAB and using the heuristic angular reconstruction, the runtime of a full optimization is still hours long.

The design vector exposed to the optimizer comprises, for each of the four antennas, its position along the array axis, its aiming point on the testing-surface plane, its feed amplitude (as a fraction of the maximum available feed) and its feed delay. For calculation speed, the position of the antenna elements has been forced to be symmetric in the boresight axis. The two objectives, both minimised, are the amplitude inhomogeneity across the testing surface and the peak-time spread across it. The constraints are: a lower bound on the minimum peak field of 5 kV/m, ensuring the thesis goal is met everywhere on the surface; upper bounds on the rise time and FWHM, rejecting designs that smear the pulse beyond tolerance; and a geometric non-overlap constraint requiring the centres of adjacent dishes to be separated by at least one dish diameter (12.2 m), since two antennas cannot occupy the same space. A number of simplifications carried over from the manual synthesis remain in force: the antennas are confined to the horizontal axis, the array and testing surface float in free space, and the feed pulse is taken to be perfectly reproducible across the four channels. No feeding network (and all accompanying uncertainties) are incorporated as of yet.

6.3.3 Optimised array

Running the optimization at the goal 5 kV/m floor returns the Pareto front of designs that are not simultaneously improved in both homogeneity and timing. A single balanced design is selected from that front for inspection; the design that lies closest to the ideal (lowest-inhomogeneity, lowest-spread) corner of the front is taken as representative. In literature, this point is referred to as the knee. Its parameters are given in table 6.2.

The optimizer converges on a configuration that is mirror-symmetric about the array axis, as imposed, but with two further features that were not imposed and are therefore informative. First, the vertical component of every aiming point is driven to within a few centimetres of zero: the optimizer discards the quadrant-style vertical aiming of the manual array of section 6.2. This implies that the quadrant aiming, while perhaps intuitive, is clearly suboptimal. Second, the two outer antennas are strongly de-weighted relative to the two inner ones, to roughly a quarter of their amplitude, while the inner pair runs at almost full amplitude. The array effectively suppresses its outer elements: fewer strong coherent contributors produce less spatial interference ripple and less spread in arrival time, which is exactly what the two objectives reward. Each pair is aimed slightly inward (left pointing right, right pointing left), creating a crossed-beam pattern with peaks offset from the array centre. The resulting spatial distribution is distinctly different from a single antenna (twin lobes rather than a single central peak) but achieves only marginal gains in homogeneity at the balanced operating point.

Table 6.2: Converged symmetric four-element design at the 5 kV/m field floor

Antenna	Antenna offset x [m]	Aim x [m]	Aim y [m]	Delay [ns]	Amplitude Taper
1 (outer)	-19.37	2.01	0.49	3.23	0.262
2 (inner)	-6.19	2.24	0.11	0.00	0.983
3 (inner)	6.19	-2.24	0.11	0.00	0.983
4 (outer)	19.37	-2.01	0.49	3.23	0.262

The peak-field homogeneity and arrival-time distribution of this design are shown in figure 6.11.




Figure 6.11: Peak electric-field homogeneity and pulse-arrival-time distribution over the 10×10 m testing surface for the converged symmetric array at the 5 kV/m floor. White crosses mark the four aiming points.

The peak-field map of figure 6.11 shows the twin-lobed structure produced by the crossed inner beams, with a shallow central seam; the worst-to-best variation over the surface is 0.56 dB. The arrival-time map shows a corresponding worst-case spread of 1.06 ns across the surface. To visualize the performance of the optimized array, time-domain E-field traces all over the testing surface are taken and added together in the figure below:




Figure 6.12: Time domain E field ensemble over the testing surface

Figure 6.12 visually shows the effect of the time domain spread of waveforms arriving over the testing surface. The grey traces are 1024 time domain traces taken uniformly over the 10 m x 10 m testing surface. They are bound by two black lines which show the pointwise maximum- and minimum (so not necessarily from the same trace per point). Finally, in blue, the boresight point is shown. The conclusion that can be drawn from this image is that despite the heat maps exaggerating small differences in time domain results, this image shows a far more optimistic result. The largest difference between the traces is found at the peak point, with the lowest peak value being 8.8 kV/m compared to the 9.5 kV/m of the peak value. What is interesting about the peak is the boresight trace: it shows a dip with two shoulders. That staggered peak is attributed to the design of the optimized array, which has additional time domain delay for the two outer antennas (3.23 ns). This was likely done to improve the time domain homogeneity, but the result in time domain E field is this staggered peak. Apart from the peak area, the traces all follow each other closely, implying good homogeneity over the testing surface.

It is good to now compare these results with the manually synthesized array and the single antenna. Relative to the manual quadrant array of section 6.2, the optimised design achieves far better homogeneity (-0.56 dB to -3.57 dB) and a smaller timing spread (1.06 ns to 1.56 ns), though at a markedly lower peak field (~ 9.5 kV/m against the manual array's ~ 17.4 kV/m). The optimiser has, in effect, spent the peak field it holds in surplus relative to the 5 kV/m goal on uniformity and timing instead: the manual array runs all four elements at full amplitude and so reaches a high but very uneven field, whereas the optimised design deliberately de-weights its outer elements to trade that excess peak for a flatter, better-timed distribution. Relative to the single antenna the result is a trade rather than a total improvement. The balanced design delivers roughly twice the peak field of one dish (~ 9.5 kV/m against 4.53 kV/m), but its homogeneity (0.56 dB) and timing spread (1.06 ns) are slightly worse than the single antenna's 0.37 dB and 0.80 ns. The array's unambiguous advantage over a single element is therefore peak field; its homogeneity and timing are, at the balanced operating point, marginally poorer.

The frequency domain performance is also inspected, in similar fashion to the manually synthesized array:




Figure 6.13: Spectral magnitude of the optimized array

Figure 6.13 shows the spectral magnitude, at two points, of the optimized array in comparison to the NEMP standard. The spectral magnitude can be seen to be significantly closer to the 5 kV/m peak value standard trace, a consequence of using the additional peak field strength to obtain better peak value and timing homogeneity. The large notch at 500 MHz is gone, although the corner point does show some interference at the very highest frequencies which drop it down outside of the 10 dB tolerance. The value at which both traces enter the 10 dB limit is still 4.5 MHz, despite the peak value being significantly lower (9.5 kV/m as opposed to 17.4 kV/m).

The figure showing the cross-section of the spectral content is also constructed, in the figure on the next page.




Figure 6.14: Cross-section of the spectral magnitude of the optimized array

The cross-section is very similar to the one of the manually synthesized array, specifically in the low frequency portion until approximately 100 MHz. Here the optimized array shows similar uniformity across the cut as the manual array did. The difference is in the high frequency content, where once again it is centred around the bore-sight point and the aim points. With changing aim points, the high frequency content shifts with it accordingly.

6.3.4 The price of peak field strength

In the previous section it was shown that the optimizer trades peak field strength for homogeneity and peak time delay. The 5 kV/m floor is only the threshold this thesis happens to target; it is interesting to ask how the achievable homogeneity and timing change as progressively more peak field is demanded. The optimization was therefore repeated for a series of minimum-field floors from 5 to 12 kV/m. For efficiency, each floor was initialised from the converged population of the next-lower floor: because a lower floor is a strictly looser constraint, its solutions remain available to the tighter floor, which both accelerates convergence and guarantees that the achievable performance cannot improve as the floor is raised. The resulting family of Pareto fronts is shown in figure 6.15.




Figure 6.15: Pareto fronts of amplitude inhomogeneity against peak-time spread, one per minimum peak-field floor. Both axes are the 5th–95th percentile spread across the testing surface.

Up to about 9 kV/m the fronts overlap closely: demanding more peak field in this range costs essentially nothing in either objective, because the homogeneity-optimal configuration already produces this much field as a by-product. Above roughly 10 kV/m the fronts move sharply up and to the right, as satisfying the higher field floor forces the outer antennas back towards full amplitude, reintroducing the very interference ripple and timing spread that de-weighting them had suppressed. Reading the best achievable value of each objective off these fronts and plotting it against the field floor gives the summary of figure 6.16, which is the central quantitative result of this optimization exercise.

Note that the values plotted for 'best achievable' in figure 6.16 are read from opposite corners of the pareto front. That is to say, the antenna array optimized for either homogeneity or for timing, never for both simultaneously. The shape of figure 6.16 is the answer to whether arraying is worthwhile. Up to a field floor of approximately 9 kV/m, the achievable homogeneity is essentially flat and sits close to the single-antenna value: in this regime the array delivers more peak field than one dish at little or no cost in homogeneity. Beyond roughly 9–10 kV/m the homogeneity cost rises sharply, rising past 1 dB beyond 10 kV/m to roughly 1.7 dB at 12 kV/m, several times the single-antenna figure. The timing cost behaves similarly: the best achievable timing actually sits slightly below the single antenna up to about 9 kV/m, before degrading steeply at higher field demand. Beyond roughly 10 kV/m the field floor can be met




Figure 6.16: Best achievable amplitude inhomogeneity (left) and peak-time spread (right) as a function of the minimum peak-field floor demanded. Both are the 5th–95th percentile spread across the surface, not the worst-case (max–min) values quoted in the text and in Table 6.3

only by allowing the surface field to become increasingly non-uniform: the best balanced design spans 11.0–13.4 kV/m across the surface at the 11 kV/m floor, and 12.0–15.5 kV/m at 12 kV/m (the highest floor tested). The configuration thus remains geometrically feasible up to at least 12 kV/m, but only by sacrificing the homogeneity improvement that motivated the array.

Two conclusions follow. First, there is a usable operating regime, up to about 9 kV/m, in which the four-element array is genuinely worthwhile: it multiplies the peak field while preserving close to the homogeneity of a single element. Second, while no hard feasibility ceiling was reached within the tested range (designs satisfying all imposed constraints exist up to at least 12 kV/m) there is a clear practical ceiling: above roughly 9–10 kV/m the homogeneity degrades so steeply that higher field levels are of little practical value. Even the highest feasible field tested, 12 kV/m, remains far below the 50 kV/m of the original standard discussed in chapter 2, which quantitatively confirms that arraying a band-limited reflector IRA is a peak-field multiplier with a firm practical bound, not a route to standards-level field strength.

6.3.5 Feed-network and mechanical tolerances

Design practicality has been a desired throughout this thesis, and has been kept in mind in every chapter so far. The optimised design of section 6.3.3 prescribes specific feed amplitudes, feed delays and antenna positions; the question for practical implementation is how tightly each of these must be held before the homogeneity and timing degrade unacceptably. This was assessed by perturbing each parameter set about its nominal value with independent Gaussian jitter and recording the resulting distribution of the two objectives. 200 trials were done with 4 different values of σ , and the resulting spread in amplitude homogeneity and in peak time was recorded. Of the 200 trials, the 95th percentile was taken and plotted, which prevents skewing the results with a highly unlikely extremity. This is presented in figure 6.17 on the next page.


The amplitude tolerance is the most important, because the design’s homogeneity is produced almost entirely by the deliberate amplitude imbalance between inner and outer antennas. Figure 6.17a shows

that the homogeneity degrades only gently with amplitude jitter: a one-sigma amplitude error of 2% raises the worst-case inhomogeneity from 0.55 to only 0.60 dB, and even 5% keeps it below 0.7 dB. The timing spread is more sensitive, its upper tail beginning to open beyond about 5% jitter. An amplitude reproducibility of order a few percent across the four channels is therefore sufficient.


The delay tolerance, figure 6.17b, requires care in interpretation. Under the conservative worst-case metric the timing spread degrades quickly, suggesting delays must be held to a few tens of picoseconds. However, this is in large part an artefact of the worst-case (max-min) metric, which is set by a single extreme surface point. The median spread degrades far more gently, and under that a one-sigma delay jitter of 0.1–0.2 ns is tolerable. If sub nanosecond delays over the entire surface is required, the worst case metric can be retained and picosecond jitter must be enforced. However, if a certain percentile to remain within the nanosecond is sufficient, a much more relaxed tolerance may be applied to the delay jitter.

The positional tolerance, figure 6.17c, is the least demanding electrically: even a one-sigma position error of 0.5 m changes the homogeneity by under 0.08 dB. The binding limit on antenna placement is not electrical at all but mechanical: the inner antennas sit close to the dish-diameter non-overlap limit, leaving only about 0.1 m of clearance on each side, so positioning is constrained by physical collision rather than by field uniformity.


Taken together, these three sweeps constitute the feed-network specification that this chapter set out to find. To realise the optimised array, the feed network must deliver, across its four channels, an amplitude reproducibility of order a few percent, an inter-channel delay accuracy of order 0.1 ns (tighter if worst-case uniformity is demanded), with antenna positioning limited by mechanical clearance rather than by the field. None of these is exotic by the standards of pulsed-power practice [26] [27], which suggests that the limiting factor for this approach is not the feed network but the antenna element itself.



(a) Amplitude jitter



(b) Delay jitter



(c) Position jitter

Figure 6.17: Degradation of the two objectives against jitter.

6.3.6 Limitations and validity of the optimisation

Several caveats bound the results above. The forward model propagates each antenna's radiation using the per-frequency far-field patterns of section 6.1.1, linearly interpolated in frequency. Furthermore, the model assumes no inter-element effects, such as mutual coupling, are occurring; the model treats the antennas elements as completely independent. The optimised designs should be read as model optima rather than as exact predictions.

Finally, the entire study is constrained to four antennas placed on a single horizontal axis, radiating onto a fixed 10×10 m surface, all simplifications inherited from the manual synthesis. The practical field ceiling found in section 6.3.4 is therefore a property of this specific configuration, not of arraying in general. The option of more elements, and the possibility of a larger or subdivided testing surface are all directions in which the ceiling might be raised, and are taken up in the discussion of potential future work.

6.4 Discussion

An interesting observation from this chapter is the effect of the NEMP's spectral content on the radiation pattern of the 12.2 m antenna over the frequency of interest. The antenna starts to effectively radiate from approximately 25 MHz upwards, meaning the impulse response was a sharp and focused beam at 100 m as a result of the 975 MHz that were effectively bundled by the antenna dish. However, the NEMP has energy located mainly in the portion of the antenna pulse that is not effectively bundled, and therefore the radiation pattern of the antenna radiating that specific pulse is significantly more homogeneous than the impulse response. A large antenna would bundle the low frequency content more effectively, provided that the far-field condition is met.

Time domain fidelity is a challenge for any radiated NEMP test, also for the 12.2 m antenna dish. The very low frequency content that contains the bulk of the NEMP's energy is difficult to radiate as it requires an enormous antenna. By not radiating it, the time domain fidelity is altered by definition, as the low frequency content adds peak field value as well as pulse-width. By placing a band-limited antenna in an array, this problem is not solved. The hardest physical limit on time domain fidelity is not the array; it's the dish size. An N-element IRA array can give roughly N times the boresight peak, but it does not extend the radiated spectrum downward. Each element still radiates inefficiently once the wavelength no longer fits in the dish, approximately 25 MHz for 12.2 m.

The 50 kV/m peak field strength limit, as required by the NEMP standards, demands either more antennas fed with the same pulse generator, or a higher peak voltage from the pulse generator. The latter is a linear extrapolation: the 4-element array, fed by a 500 kV pulser, of this thesis achieved 9 kV/m before homogeneity dropped significantly. Extrapolating gives a requirement of 2.78 MV per antenna. If peak value is the only metric that mattered, then increasing the amount of elements is a linear extrapolation as they all add coherently in the centre. A single antenna achieved 4.58 kV/m, so 12 elements would exceed 50 kV/m. Taking the best achievable homogeneity cap at 9 kV/m for 4 antennas would require 22 elements.

However, peak field value is not the only metric that matters, and therefore increasing the amount of elements in the array to increase the peak field value is not a linear extrapolation. Each new element brings additional interference ripple over the testing surface, thereby increasing the complexity of the homogeneity demand. Besides, a 22 element long array is already 256 m wide, and it is likely that more antennas are needed. It is infeasible to achieve 50 kV/m with the 12.2 m antenna array.

Finally, a last observation made in this thesis is the existence of 'free' peak field strength design space with respect to the two homogeneity parameters (peak field strength and peak timing spread). Shifting the threshold from 5 kV/m to 9 kV/m came at almost no cost to the homogeneity parameters.

This chapter has treated the optimization of the array for the purpose of NEMP testing as a three-way tension between peak field strength, electric field strength homogeneity and peak timing spread over the surface. By fixing peak field strength as a constraint, in similar fashion to the standards, the array may be optimized for the best combination of electric field strength and timing homogeneity. The modelling done in this chapter was limited, just enough to shed light on some design trade-offs. By expanding on the design variables that the optimizer may tune, the optimization may improve further on performance. Potential expansion on the work done in this chapter will be discussed in the coming paragraphs:

For this thesis it was chosen to contain the element placing on a horizontal line because practical implementability is desired. However, the optimization may be expanded upon by allowing varying vertical placements of the antenna elements. Potentially, a grid or ring like structure of elements results in improved homogeneity, at the cost of practicality.

The feed waveform used has been derived and then locked in since. Varying parameters of the feed waveform may allow electronically improving array performance without the need of physically changing the layout. The use of pre distortion on certain elements may counteract time domain smearing of the pulse or interference ripple over the testing surface.

The number of elements has been fixed, which was done in context of practical viability. However, it is possible to leave it open as a design parameter to be tuned. It is likely that the peak field strength floor chosen will dictate the number of antennas required.

The antenna array operating in a vacuum assumption has allowed the model to be implemented with relative simplicity, albeit unrealistically. For more realism, a perfectly conducting ground floor may be added to the calculations. This has precedent: the EMIS-III-VPD had a metallic mesh suspended at its base, connecting it with the DUT, such that the monocone structure electrically appears as a full bicone. Even if in the final design of the antenna array the full IRA is used and not a HIRA, an electrically conducting mesh could still provide certainty on the reflectivity of the surface. With the antenna and testing surface so close to the reflecting ground floor, the antenna will see a significant image. What's more, the relative extra distance travelled results in time difference of arrival in the order of magnitude of nanoseconds, well within the FWHM of the NEMP. Adding the reflecting ground surface will add realism to the model and give further insight into antenna-array performance.

One parameter that has not been discussed, and therefore not optimized for in this chapter, is the spectral content of the antenna. The NEMP standards demand a certain spectral content within tolerance, as has been shown in chapter 2, and tested against in chapter 3 and 5, but it has been left out of this optimization exercise. An array (and antenna element combination) may have excellent homogeneity (both in time and peak distribution) as well as time-domain fidelity, but it may contain spectral 'notches' that are not accounted for. Therefore, in future work, the spectral behaviour of an antenna array may be incorporated in its optimization.

Finally, expansion of the work done here may be done by including antenna tolerances in the optimization. In this chapter, the antenna was first optimized, after which the tolerances for the feed and placement of the elements were explored. This means that robustness of the design depends on the achievable tolerances of its components, which may mean the goals are not reached. Instead, these tolerances may be imposed on the structure beforehand, after which the optimizer comes to a solution that is robust by design.

6.5 Conclusion

This chapter reformulated the synthesis of a four-element 12.2 m IRA array as a multi-objective optimization problem. Recognising that the radiated peak field, its spatial homogeneity, and the spread in pulse-arrival time across the testing surface are in mutual tension, and that the NEMP standards specify a peak-field threshold rather than a maximum, the peak field was imposed as a constraint and the two remaining quantities (amplitude inhomogeneity and peak-time spread) were minimised simultaneously. A broadband time-domain model, which was anchored to the CST far-field response of the single antenna, evaluated candidate designs, and a multi-objective genetic algorithm mapped the Pareto front of achievable trade-offs. The results are summarized in Table 6.3.

Table 6.3: Summary of results gathered in this chapter

Configuration	Peak field [kV/m]	Min. field [kV/m]	Amplitude inhomog. [dB]	Peak-time spread [ns]
Single antenna	4.5	4.3	0.37	0.80
Manual quadrant array	17.4	11.5	3.57	1.56
Optimised array	9.6	9.0	0.56	1.06

The central finding is that the three objectives cannot be jointly optimised by an antenna array. A single antenna achieves the best homogeneity and timing of any configuration considered, but at the selected 500 kV feed it radiates only 4.5 kV/m and so fails the 5 kV/m goal everywhere on the surface. An array is therefore necessary to clear that threshold, yet every element added to do so introduces interference ripple and additional arrival-time spread. The optimised array resolves this far better than the manually synthesised array of the previous section, improving homogeneity from 3.6 to 0.6 dB and timing spread from 1.6 to 1.1 ns while still delivering close to twice the single-antenna peak field. However, it does not beat a single antenna on homogeneity or timing; its sole advantage over one dish is peak field. It achieves this by strongly de-weighting the outer elements and aiming the inner pair across the array centre, spending the peak field it holds in surplus on uniformity.

Repeating the optimization across a range of peak-field floors quantifies the price of peak field. Up to a floor of roughly 9 kV/m the achievable homogeneity is essentially unchanged and remains close to the single-antenna value, so additional field is obtained at negligible cost. Above approximately 9 kV/m the homogeneity degrades steeply; feasible designs exist up to at least 12 kV/m, but only by allowing the surface field to become increasingly non-uniform. There is thus no usable path to the 50 kV/m of the standards through element count alone; maintaining homogeneity at that level would require of order twenty elements spanning some 250 m, which is impractical. The alternative, a linear scaling of the feed voltage, would demand approximately 2.8 MV per element: a severe but bounded requirement, ultimately limited by voltage breakdown rather than by the radiated physics. The hardest limit on time-domain fidelity is not the array but the dish; an N -element array multiplies the boresight peak by approximately N , but the array scales the magnitude, not the shape; the crossing point improves only because the absolute level rises. The shape remains what it is.

A tolerance analysis of the optimised design showed that the required feed-network precision is demanding: amplitude reproducibility of a few percent and sub-nanosecond inter-channel timing are necessary. Element positioning is limited by mechanical clearance rather than by field uniformity, allowing for less stringent tolerances. All values reported here are within-model optima derived from the interpolated far-field reconstruction; the selected design remains to be validated against an independent full-wave simulation before its performance can be quoted as realisable.

Chapter 7

Discussion

This thesis aimed to explore the feasibility of a timed array of ultra-wideband, directive radiators as a testing set-up for an NEMP hardness validation test for large (naval) platforms. With all the results gathered to conclude this thesis it is tempting to simply wrap up with a simple answer to the research question. However, it is good to qualify these results and dwell on the context of them to this thesis and to the broader setting of NEMP testing. This chapter's aim is to reflect on the work done and presented in this thesis in more detail, qualifying it and suggesting areas for improvement. Furthermore, the potential areas for future work are to be discussed as well.

7.1 Context of this thesis

To the author's knowledge, this is the first publicly documented research project specifically addressing radiated NEMP hardness validation of large naval platforms since the early 1990s. In Western Europe, the fall of the Berlin Wall in 1989 and the Soviet Union in 1991 signified the end of almost 50 years of tension between East and West. In the following years, the cost of keeping a defence force was deemed unnecessarily high, and budget cuts were implemented in most Western European countries reducing the size of the standing armed forces. The focus for the task of armed forces shifted from defence of own and allied territory by deterrence to international peacekeeping operations, which simply did not require the size or technological sophistication of the armed forces kept in the early '90s. As the general thought was that the armed forces of NATO member states would never face a nuclear power again like they did in the Cold War, costs could be saved by reducing or even fully discarding NEMP hardness requirements. It is simply not defensible to spend a large amount of money on protection against a threat so unlikely to be encountered as the NEMP, despite all the consequences if it did ever occur.

The idea that NEMP hardness is redundant and unnecessarily expensive is gone. With increasing world tension in general, sparked by the annexation of Crimea in 2014 and the full-scale invasion of 2022 in particular, the idea that Western Europe will one day face a nuclear capable enemy lives once again. Defence spending has increased significantly over the last 4 years, which trickles down to budget for research and development in the military sector.

That is where this thesis comes in. It is part of the revival of NEMP research in general, specifically radiated NEMP hardness validation testing. As it is a topic that has practically stood still for over 30 years, it has had to reinvent the wheel to some degree. Work has been done since the Cold War on IRAs, which is why they were a good choice for use of antenna for this thesis, but they were for very different purposes than NEMP hardness validation. Methods and computational power for numerical electromagnetic computations have improved vastly since the early '90s, which makes research into this topic now much more potent than it used to be. By keeping this thesis broad, and only shallowly diving into each of the topics covered by this thesis, doors are opened to new avenues along which research into NEMP hardness validation using radiated NEMP testing can continue. By its broadness, this thesis provides a framework on which future work can be built, and it provides context and bounds for any spin-off works. The role in the revival of this topic that this thesis plays is the high-level overview it provides on what is necessary before practical implementation of this idea may happen.

7.2 Synthesis of central findings

In this section, the central findings of this thesis will be distilled from the results gathered in this thesis. These will eventually form the backbone of the conclusion, the next and final chapter of this work. This will be split into several subsections; showing what has been demonstrated in this thesis, the assumptions that those results are based on, considerations for practical implementation and finally the central results

7.2.1 What has been demonstrated

In chapter 2 of this thesis, the relevant standards on the NEMP and NEMP hardness validation have been consulted. They describe how a perfect NEMP test can be conducted, which is by definition impossible for any realistic set-up. From them, design specifications have been distilled for the work done in this thesis. They have been catered specifically to the DUT in mind, a large naval platform, such that the work done remains anchored to the goal of the thesis. As one of the outcomes of this work is to provide a framework within which further research may be done, this is key in bounding this and future work. It was shown that amongst other things, a successful NEMP test achieves three key parameters: illumination at sufficient field strength with sufficient peak field strength homogeneity and time-of-peak homogeneity across the testing surface. These three main parameters are supplemented by parameters such as field polarization, frequency domain fidelity, frequency bounds and time domain fidelity.

The baseline study of the performance of the EMIS-III-VPD, which attempted to illuminate the entire DUT in a single pulse, achieved none of the three main parameters that constitute a successful test. The EMIS-III-VPD did not illuminate the DUT with sufficient peak electric field strength, achieving just 1.28 kV/m, nor did it achieve homogeneity in peak field value across the testing surface, with an average of 10 dB difference in peak electric field strength, nor did it achieve time domain homogeneity, with up to 138 ns difference in time of peak between centre and edge of the testing surface. Furthermore, the spectral magnitude dropped outside of the tolerance between 50 MHz and 300 MHz, and the EMIS-III-VPD is only capable of vertically polarized waveforms. It is concluded in this chapter that full illumination of a large naval platform with sufficient peak field strength and homogeneity is practically impossible. Therefore, this thesis does not aim to do so, instead accepting that the full DUT is no longer illuminated as a trade for improved homogeneity and peak field strength. With this in mind, a concept for an NEMP test is proposed: a time-domain array of directive elements.

The EMIS-III-VPD was a monocone antenna over a reflecting electric plane, and it was an ineffective antenna. Its omnidirectionality meant that the peak electric field is by definition low as it is smeared out over the entire horizontal plane. Furthermore, specifically the EMIS-III-VPD was ill-suited to the selected test configuration in operation. Chapter 4 of this thesis aimed to find an antenna better suited as antenna to serve as element in the time-domain array. Based on the design specifications that have been distilled in chapter 2 and the experiences from chapter 3, parameters were set up to quantify the estimated performance of all antennas that qualified for this work (UWB, non-dispersive antennas). The winner was the reflector IRA, selected for its high directivity, good time-domain fidelity, ability to support dual polarization and the maturity of the concept in the academic world. The TEM horn antenna was a close second, and may be useful in future work for another application.

The reflector IRA was taken and modelled in a full-wave microwave simulator. Three models were made, which were scaled versions of each other, such that a key trade-off could be shown: the antenna practicality (to say: its size) vs the low frequency operating limit. This antenna provided significantly better results than the EMIS-III-VPD did, specifically a higher peak field strength and better time-domain fidelity. The reflector IRA is shown to be a promising building block for use as an antenna element in a time-domain array. The 12.2 m model was used in the subsequent chapter 6, which covers the array, as this was shown to be the limiting case if it is desired that the testing surface is in the far field.

The arraying of the 12.2 m model showed that arraying can be used as a peak field strength multiplier. The boresight peak scales roughly linearly with the number of elements aimed at the same point. The improvement of the peak-field strength comes at the cost of the reduction in peak field homogeneity and in peak time homogeneity. Without optimization of the array, the added peak field strength comes at the cost of significant reduction in homogeneity by interference and by timing errors.

The synthesis of an array for the purpose of NEMP testing can be formulated as a multi-objective optimization problem. Instead of maximizing peak field strength, it was taken to be a constraint, and given that the peak field strength must be at least 5 kV/m, the amplitude and peak-time homogeneity can be optimized for. This design step is conceptually useful because it reveals an important aspect of the nature of an NEMP test: the peak field strength does not need to be indefinitely increased, it must be high enough while simultaneously balancing amplitude and peak time homogeneity.

A hard physical constraint remains for any radiated NEMP test: it cannot efficiently radiate the low frequency component of NEMP while simultaneously being a feasible practically implementable antenna. While arraying from 1 to 4 12.2 m reflector IRAs improved the point at which the spectral magnitude entered the bound as set by the NEMP standards, the time domain waveform still showed signs of missing low frequency content: a relatively short FWHM and sharp decay. No amount of arraying will alleviate this missing low frequency content. The question that remains is whether this is critical for the reliability of an NEMP test.

7.2.2 Assumptions, heuristics and their consequences

In this thesis, models were used to answer the central questions of this work. All models are imperfect: to some extent depending on heuristics and/or assumptions to feasibly work and provide insights. Some of the more egregious shortcuts taken in this thesis, as well as their consequences, will be touched on again.

The EMIS-III-VPD model is the first presented here, and will therefore be covered first. Firstly, it must be noted the model is built based on very limited publicly available information. The fact it is an 8-wire structure for instance is based on the only a picture of the set-up available to the author. Therefore, most modelling choices are based on estimations made by the author and not on real design choices made by the actual designers of the antenna. While an idea may be obtained of its performance and limitations based on this thesis, quantified results of its performance must always be treated as an estimate. This is part of the reason three models were made for the EMIS-III-VPD; to gain some further understanding in the design parameters and trade-offs at play for such a structure. Eventually it was deemed the end-termination model is most likely the closest to the real structure, and therefore its numerical results are taken as the baseline. While these figures may not be exact, they do give an approximation.

In [10] it was stated that the EMIS-III-VPD was fed by a 500 kV pulser. In chapter 3 this was assumed to be a Marx generator, as that is the pulser used in the EMPRESS-II [11], but this has never been verified in a publicly available source. What the '500 kV' number means exactly or how that speaks to pulser performance is unclear, but for this thesis it was taken to mean the pulser achieves a peak value of 500 kV. It was further assumed that the 500 kV Marx generator operates perfectly; that is to say, its peak value is 500 kV every single run. This means that the performance estimates achieved by the EMIS-III-VPD are upper estimates given a perfectly functioning generator. Furthermore, throughout this thesis the same generator has been used so fair comparison can be made between all antennas. How to practically fit a 500 kV Marx generator to the feed-point of a reflector IRA is left as future work, as is the choice for or design of a true pulse generator given the requirements needed for it.

Every model has taken heuristics in the incorporation of the surroundings in full-wave and MATLAB simulations. The EMIS-III-VPD models were placed directly on top of a perfectly electrically conducting surface, as was the testing surface. The reflector IRA models were suspended in a vacuum. The effect of the sea-surface, or an imperfect electrically conducting plane, will have significant effect on the antenna performance on the testing surface, as well as on the optimizer for the array. These effects are left for future work, but their absence from the models must be kept in mind.

It was known that the far field assumption did not hold for the EMIS-III-VPD, nor for the 30 m reflector IRA model. Despite this, it has still been presented throughout this thesis. It is by far the most efficient result to retrieve from CST, and it is the most straightforward value to use for array synthesization. It must therefore be kept in mind that numerical far field results of these antenna are not fully watertight, instead they provide an indication of the order of magnitude of the true near field quantity.

The size of the antenna models, specifically the EMIS-III-VPD models and the 30 m reflector IRA model caused enormous, infeasible calculation sizes. Therefore, the choice was made to reduce numerical solver accuracy, specifically by reducing the density of discrete mesh cells within which Maxwell's equations are solved in the time domain solver. This is likely to have resulted in numerical artefacts appearing in either the feed or the radiation of the antennas, which skew the results.

The array model introduces another important simplification. The optimized array was not simulated as a complete four-antenna structure in CST. Instead, the radiation pattern of a single antenna was exported and reused several times in MATLAB. This makes the optimization feasible, but it assumes that the antennas do not affect each other. In reality, mutual coupling, scattering between antenna structures, shadowing, feed-line interaction and changes in input impedance may occur. These effects are not included in the array synthesis. The optimized array should therefore be interpreted as an optimum of the simplified model, not as a guaranteed optimum of a real physical array. A full-wave simulation of the complete optimized array would be needed before the numerical results can be treated as final.

Finally, the work has not been experimentally validated. All quantitative results in this thesis come from simulations, analytical calculations or MATLAB post-processing. These methods are useful for exploring concepts, but they do not replace measurement. Real antennas include construction tolerances, imperfect materials, connectors, supports, feed cables, grounding structures and environmental effects. These may all influence the pulse. A scaled experiment or low-voltage demonstrator would be needed to check whether the simulated impulse responses, spectral behaviour and array interference effects occur in the same way in practice. Until such validation has been done, the numerical results should be interpreted as indicative rather than definitive.

7.2.3 Practical radiated NEMP testing

Practicality has remained a goal kept throughout this thesis. It is envisioned that eventually this thesis will stand at the beginning of a (very) long line of future work to revive this topic, after which the practical implementation of a method for validating the NEMP hardness of a large naval platform. It is for that reason that in this work the peak electric field strength and the testing surface size were reduced, in order to keep practicality as a design feature.

Furthermore, an antenna of limited size was selected for use as antenna element in the array: an antenna diameter of 'just' 12.2 m can be considered small considering the goal is to radiate from 1 MHz onwards. While this goal was not achieved, as will be discussed, the antenna still does radiate enough low frequency content for meaningful results while simultaneously not having an antenna height in the hundreds of meters. Furthermore, strategies to practically implement the antenna while preserving the electrical shape have been discussed in that section, such as the use of a HIRA and implementation of a mesh as the reflector dish.

This thesis achieved the 5 kV/m peak field strength with some margin. However, this was reduced from the 50 kV/m demanded by the standards for practicality and feasibility. To achieve 50 kV/m with the set-up presented in this thesis, each antenna would either have to be fed by a feed-gap voltage of ≈ 2.8 MV/m or by ≈ 20 antennas spanning more than 250 m. Neither is practical or feasible, therefore it can be deemed unlikely that 50 kV/m is ever practically achieved by any combination of antennas. Actual implementation of NEMP testing would require mitigation of this loss. This is left as future work.

A further limitation is that the work in this thesis only evaluates the incident electromagnetic field on the testing surface. The actual purpose of NEMP hardness validation is to determine whether representative currents and voltages are induced in the DUT, and whether critical systems remain functional. No detailed ship model, cable coupling model, aperture coupling model or internal equipment response was included. Therefore, it cannot yet be concluded that improved incident-field homogeneity directly corresponds to improved validation confidence. This connection between incident-field quality and DUT response remains one of the most important steps before practical implementation.

The tolerance study of the optimised design showed that the feed network is not the limiting factor. Amplitude reproducibility of order a few percent across the four channels, inter-channel delay accuracy of order 0.1 ns under a median timing metric (tighter only if worst-case uniformity over every surface point is demanded), and element positioning bounded by mechanical clearance rather than by field uniformity are all within established pulsed-power practice. The binding limits of the approach are physical: the dish aperture sets the spectral floor and pulse width, and voltage breakdown sets the peak field.

7.2.4 Central findings

The central problem is that there is not such a thing as a perfect NEMP hardness validation test. It is simply not possible to achieve every available metric required for exact replication of the NEMP simultaneously. By definition, some choices must be made to prioritise certain parameters at the cost of others. Most importantly, a NEMP-like pulse must be radiated with sufficient peak electric field strength homogeneously over the testing area. The standard for NEMP testing is the bounded wave simulator, a waveguide type architecture where the DUT is placed within the structure. This type of simulator does achieve sufficient electric field strength with good homogeneity. Furthermore, as no electromagnetic energy is radiated, the very low frequency component is effortlessly present within the bounded wave simulator. The problem is that large DUT's, such as warships, do not fit under or in bounded wave structure, and therefore radiated tests are the method of choice for such DUT's.

Radiated NEMP simulators have a defining problem: it is impossible to radiate DC and it is impractical and difficult to radiate the very low frequency component. This is not an issue that can be resolved within this thesis, but the trade-off between the low-frequency operating limit and the antenna practicality (defined by its size) is one that occurs throughout this thesis. The radiated NEMP test that is the standard and baseline for this thesis is the EMIS-III-VPD, a moncone over a PEC surface. This antenna aimed to illuminate the entire DUT simultaneously, trying to achieve the required peak electric field strength homogeneously over the entire surface.

The first main finding of this thesis is that it has been shown that within the assumptions and constraints considered in this thesis, full-illumination radiated testing of a large naval DUT does not appear to be a promising path toward simultaneously satisfying field-strength, amplitude-homogeneity and time-domain requirements. The required peak field strength and homogeneity cannot be achieved simultaneously over a surface as large as required for the DUT in mind. This thesis deliberately makes the choice to trade full scale illumination for improved peak field strength and homogeneity over a limited surface. This means effects that full illumination testing would capture are no longer generated. However, as was shown in chapter 3 of this thesis, the EMIS-III-VPD also did not illuminate the entire DUT simultaneously due to the nature of the circularly expanding pulse.

The second main finding of this thesis is that the directivity of the antenna is the enabling property, while the aperture is the limiting one. Of the UWB, non-dispersive antennas, the reflector IRA was selected. This was mainly for its directivity, but also for its time domain fidelity, dual polarisation capability and the extensive literature availability. Three models of the reflector IRA were made of varying sizes, and this further illustrated the central problem with radiated NEMP testing: the low frequency limit locked in a trade-off with antenna performance. The time domain fidelity is dependent on the low frequency component: it provides the longer tail and slower decay that pulls the FWHM into the requirement. Arraying does not improve on this: while the spectral content may be shifted upwards, which lower (and thereby improves) the point at which the spectral magnitude is within limits, it does not improve the time domain fidelity associated with it.

Finally, in designing the array, it was observed that optimization of the array might bring homogeneity close to the value of a single antenna, but at the cost of peak electric field strength. In this thesis, the three main parameters for a successful NEMP test: peak field strength, amplitude homogeneity and peak time homogeneity, were reduced to a trade-off between two by treating the peak electric field strength as a constraint to be satisfied. Thereby, it was shown that trading peak electric field strength for homogeneity, results of similar homogeneity can be achieved by an array at higher peak field strengths. This holds up until a ceiling value, found to be 9 kV/m for this specific thesis, after which the homogeneity decays quickly.

7.2.5 Novelty and contributions of this work

The novelty of this thesis lies less in any single result or component than in the way established components are combined and placed in the setting of a problem that has received little (publicly available) attention. The reflector IRA is a mature radiator [28–30], timed pulsed arrays whose element delays and amplitudes are tuned by optimization are established [31], and radiated NEMP testing of ships has historical precedent in the mono-cone NAVES/EMPRESS simulator [21] [11], as well as the EMIS-III-VPD, the baseline for this work. This thesis does not claim any of these individually as new. Its contribution is the following set of combinations and framings, each stated to the best of the author’s knowledge of the publicly available literature.

First, this work applies a timed array of directive reflector IRAs to the radiated NEMP hardness validation of large naval platforms. The prior radiated simulator for ships, EMIS-III-VPD, was an omnidirectional mono-cone, and the preceding NEMP simulator study for ships, EMPSIS, was a bounded-wave architecture [8]; neither used a timed directive array. The novel element is the use of directivity and timing to illuminate a constrained surface, together with the translation of the bounded-wave-oriented NEMP standards [18, 32] into a set of testable criteria for a deliberately partial-illumination radiated test.

Second, this thesis frames the synthesis of that array as a constrained multi-objective optimization. The optimization of timed arrays by adjusting element delays and amplitudes is itself established and has been applied, for example, to flat-top energy patterns [31]. The contribution here is the optimization was reversed by the NEMP standards: because the standards specify a peak-field threshold rather than a maximum [21] [18], the peak field is included as a constraint to be satisfied while amplitude homogeneity and peak-time spread are minimised. Mapping the resulting Pareto front and sweeping the peak-field floor then yields a quantified ceiling on field strength beyond which homogeneity collapses. To the author’s knowledge, this constraint-based framing has not been applied to radiated NEMP simulator synthesis in the open literature.

Third, this thesis provides what is, to the author’s knowledge, the first public quantitative performance estimate of the Cold-War-era naval radiated NEMP simulator EMIS-III-VPD [10]. The performance of these mono-cone simulators is undocumented in open sources, reconstructing it, even as an estimate bounded by the assumptions of Chapter 3, supplies the baseline against which the proposed array is judged.

Fourth, and more narrowly, the CST-anchored two-stage field reconstruction developed in Chapter 6 is a methodological contribution on its own. Anchoring the broadband array model to a single causal full-wave boresight solve and propagating across the surface by an angular ratio made the optimization significantly faster without a full-wave solve of the complete four-element structure, at the cost of neglecting inter-element coupling.

These contributions are exploratory rather than definitive. After all, this study was a feasibility study into a novel, non-conventional method. The value then lies in establishing a framework and a baseline on which the future work, outlined below, can build. This is consistent with the deliberately broad scope of this thesis. The framework set out by this work then is of highest value to future research.

7.3 Future Work

If anything, this thesis has caused more questions to be asked than it has been answered. This is partly by design, the thesis skimmed over the broad topic of radiated NEMP hardness validation testing without diving in depth into a specific subtopic, which supports the role this thesis plays in the revival of this topic. To aid this further, future work based on this thesis is listed in a categorical way below. This list is not definitive, as more research questions may be thought of based on this work. These are the first that came to mind and the ones with the biggest impact on furthering research in this topic.

Radiated partial illumination testing versus full illumination testing

This thesis has aimed at illuminating a 10 m x 10 m testing surface at 100 m distance, with the idea to eventually sweep that surface over the larger DUT. The key problem with this approach is that any effects caused by full homogeneous illumination of the DUT are no longer instigated by the partial test. This loss is accepted in this thesis, as it is traded for improved field strength, time domain fidelity and homogeneity over a limited surface. As was shown, attempts to illuminate the full 200 m x 50 m testing surface by the EMIS-III-VPD failed; they did not achieve the required field strength, nor the required homogeneity, nor the required time domain fidelity.

However, this thesis has taken that loss without any quantification, let alone mitigation, of that loss. Without this, one cannot with full certainty establish whether a DUT is NEMP hard. In public literature, this effect has not been discussed yet. Therefore, the following research question is suggested:

What is the effect of partial illumination with respect to full illumination in the context of NEMP hardness validation testing?

For this, it is envisioned that practical test is implemented. The computational power required to calculate the full EM effects of a complex object is significant. It is suggested that a system is constructed that is able to be placed both fully and partially under a bounded-wave NEMP simulator. Such a system may be two transportable containers linked with insulated electrical cables. Measurements may be taken inside both containers, while both are inside and while one is inside and one is outside the simulator. The difference between the two constitutes a transfer-function, which in turn may then be extrapolated to quantify the loss on a larger DUT, such as a warship. An outcome may be that specific additional testing is required, which may be done with for instance cable-injection.

Deliberate near-field radiated NEMP testing

After modelling the three reflector IRAs, which each represented a different choice in the antenna practicality versus lower frequency operating limit, this thesis made the deliberate choice to continue work using the 12.2 m model. That antenna was the limiting case for where far field conditions are met, which simplified post-processing and antenna array design significantly. However, a near-field analysis was conducted of the 30 m model. This showed that the pulse was temporally wider and of lower peak amplitude due to an effect called near-field dispersion. It is thought that while this does cause pulse widening in time domain as well as lower directivity (which is not yet the correct term in the near-field, but the bundling of electromagnetic radiation is meant), this may improve homogeneity. Therefore, the following question is suggested:

Can controlled near-field operation improve the homogeneity of a radiated NEMP test while maintaining acceptable time-domain fidelity?

A full, at least 2D, near-field analysis of a large IRA may quantify the effect of near-field dispersion. This quantification can be used to reverse-engineer the feeding pulse to the antenna in order to pre-distort the signal such that the expected increased homogeneity is experienced while counteracting the decreased time domain fidelity.

As a high peak electric field strength as well as a low low-frequency operating limit is desired, the deliberate choice for a near-field radiated NEMP test may be taken, such a study may provide the mitigation necessary to conduct near-field radiated NEMP testing successfully.

NEMP test uncertainty budget

In previous NEMP tests, specifically radiated tests such as the EMIS-III-VPD and the EMPRESS-II, it was difficult to quantify the exact uncertainty of the test. There was insufficient probing to fully quantify the field radiated by the antenna, and therefore it was impossible to make a quantified statement of the exact uncertainty of the test. As this thesis has revealed, in the case of a monocone antenna this is amplified by the highly irregular radiation pattern it produces. With modern computational EM solutions it may be possible to chart the uncertainty with much higher accuracy and detail. Therefore, the following research question is proposed:

How can the uncertainty of a radiated NEMP hardness validation test be quantified in terms of field strength, waveform fidelity, homogeneity, timing and polarization?

It is envisioned that besides a strong foundation in numerical answers, that the numerical solution is also validated against practical measurements. This would further solidify the confidence a user may place in the operation of the simulator.

Optimization of the reflector IRA for radiated NEMP testing

In chapter 4 of this thesis the reflector IRA was deemed the most suited antenna concept from a list of potential non-dispersive, UWB antennas. In chapter 5 a working model of this antenna was implemented in CST, which took design parameters from other research works. This saved significant optimization time for this thesis, at the cost of it being unlikely that this antenna design is optimal. Therefore, the following research question is suggested:

What reflector IRA geometry gives the best compromise between peak field strength, bandwidth, time-domain fidelity, homogeneity and practical size for radiated NEMP testing?

Tuning of parameters such as the feed arm geometry, the feeding point, the termination point, the F/D ratio, reflector edge shape, and dish diameter may produce better results than this thesis has presented. An interesting concept that may be explored under this research question is the implosion IRA, proposed by Baum in [33], in which the dish of a reflector-like antenna is not aimed at infinity but at a point at a finite distance to the dish. Furthermore, this work may be complemented by the work that is suggested in section 7.3, in which the design is deliberately designed for a testing surface in the near field.

Expansion of testing surface to a testing volume

This thesis has made use of a 2D testing surface over which it was the goal to have the waveform criteria set for a reliable NEMP test have to hold. This would hold for the analogy with radiating at the vertical face of the hull of a warship. However, for true NEMP testing this has to hold over a 3D volume. This is also how modern NEMP simulators define their testing area. To overcome this expansion upon the work done, the following question is suggested:

How can the 2D testing surface be expanded into a 3D testing volume, over which the same criteria must hold?

It must be kept in mind that a real NEMP event will propagate as a plane wave. Ideally, the testing volume criteria will reflect this, where it is stated in what manner the pulse propagates through the volume. Furthermore, the array optimization question must then be expanded to overcome this, as now homogeneity and time domain fidelity must be achieved through a volume as opposed to on just a surface.

Predistortion on the feedpulse

Originally, it was envisioned that this topic would be covered by this thesis. Unfortunately, constrained by time and talent, the topic is left here. The following research question is proposed:

How can predistortion be employed in the feed of an antenna for radiated NEMP hardness validation testing to improve the time domain fidelity while constrained by pulser capability?

Once the full system from feeding pulse to time domain electric field in the testing surface has been defined, it may be considered a linear, time-invariant system. If the analytically calculated ideal waveform is placed at the feed, the system convolutes it into a radiated waveform on the testing surface, which may or may not be the desired waveform. By inverting the system with the desired waveform, the waveform necessary for that generation is yielded. This is the method applied in section 6.1.1. The more accurately defined the system is, such as by incorporating a surface, propagation effects, inefficiencies in the antenna, etc, the more accurate the yielded feed waveform is. This then immediately yields another question: How can such a pulse be generated? Or: what concessions must be made such that the pulse can be generated?

time-domain array feeding

The central goal of this thesis was to explore the possibility of the use of a time-domain array for radiated NEMP testing. The focus of this thesis has been on the electromagnetics, treating the antenna back-end and feeding network as an optimally performing black box. However, as practical implementation is envisioned, that black box must be opened. The following question is posed:

How can we feed a time-domain array of antenna elements for the purpose of radiated NEMP testing?

To answer this question, the exact requirements for such a feeding network must be set. Once that is done, several approaches may be taken to design that network. Considering the timing requirements for accurate pulse reconstruction on the testing surface, one topic to explore is the question whether or not the pulse is generated centrally or separately for each antenna.

Missing low frequency component in radiated NEMP testing

A fundamental problem of a radiated NEMP simulator with respect to its bounded wave counterpart is the lack of the low frequency and DC-component. An antenna physically cannot radiate the very lowest frequencies, which contain a significant portion of the NEMP's power, while in a bounded wave test there is no radiation and therefore there is a low frequency and DC-component. This loss has been accepted throughout this thesis without quantification, let alone mitigation.

There is an interesting opposition with regard to this topic. On the one hand, the DC and low frequency component contains much of the NEMP's energy, making this the most potent portion of the spectral content of NEMP. On the other hand, the very long wavelengths that contain all that energy cannot possibly couple efficiently into the comparatively small DUT. In Chapter 2 it was discussed that the largest possible structure on a warship, with which conducting metallic structures or external cables are meant, is approximately 30 meters in length. Using the small-dipole analogy for such a structure, it was estimated that an incident NEMP would couple into this structure with some efficiency up to wavelengths approximately 10 times the structure length (assuming polarization is matched), resulting in a lower frequency bound of interest of 1 MHz. The opposing effect is that in this region, increasing wavelength results in increased power in the spectral content, but also decreased coupling efficiency.

This thesis has taken 1 MHz as the lower limit of the frequency bound of interest, but this may be more or less optimistic (with which a higher lower bound is meant, considering the antenna sizes required and the desire for practicality). This thesis suggests the following research question:

What is the loss made by a radiated NEMP test regarding the low-frequency content with respect to a bounded wave NEMP test?

This question could be answered with practical comparison. The construction of an antenna capable of a radiated NEMP test with a well-defined low frequency limit (and therefore time domain fidelity) will allow testing a DUT with both this antenna as well as by an existing bounded wave simulator. It must be noted that the DUT must contain components that are of sufficient size to potentially allow coupling of low-frequency spectral content, such as the 30 meters length used in this thesis. A sub-question that the suggested work could explore is to what extent the 30 metre estimate used here is realistic for a modern warship.

Extrapolation technique for low peak-field values in NEMP testing

Any NEMP test, but specifically radiated NEMP tests, have difficulty overcoming the 50 kV/m peak electric field strength limit. The extreme electric field strength required locally around the feeding point quickly exceeds the voltage air breakdown limit, and mitigating measures, such as the use of a material with a higher voltage breakdown limit, do not completely resolve the issue. In the baseline of this thesis, the EMIS-III-VPD, 50 kV/m was not achieved at the testing surface. To overcome this, field-strength measurements were done both inside and outside of the DUT such that a transfer-function could be set up describing the dampening effect of the system. With that transfer function in hand, extrapolation may be applied to estimate the dampening of the DUT at 50 kV/m external peak field strength.

The problem with this method is that the problem is non-linear. Protection measures taken against NEMP, such as over-voltage protectors placed on points-of-entry (POEs) will only activate if enough voltage is induced on that cable. Radiating so far below threat-level will mean that such measures are not triggered. Furthermore, the dampening effect of the outer shell of the DUT may be so strong that radiating low field strengths results will mean that the dampened field on the inside is below recording device sensitivity.

By using a directional antenna and arraying it, the peak field strength can be increased straightforwardly, meaning values higher than the estimated 1.28 kV/m of the EMIS-III-VPD can be achieved. However, 50 kV/m remains a challenge for radiated NEMP tests and comes at the cost of decreased homogeneity. Therefore, future work could investigate improving the extrapolation techniques instead of increasing peak electric field strength. The following question is suggested:

How can we extrapolate results from a sub-threat NEMP test to still validate hardness with confidence?

To capture non-linear effects, it is envisioned that practical testing is done to compare results between radiated and bounded-wave NEMP tests. By the nature of the bounded wave architecture, these simulators do comfortably exceed 50 kV/m. By comparing results between the bounded wave test and the radiated test, non-linear effects may be encapsulated by a transfer function.

Inclusion of the sea surface and conducting ground plane

In this thesis, all models were assumed to float in a vacuum. Naturally, in a real NEMP test, this assumption no longer holds and must be accounted for. Therefore:

How does the presence of sea water or an artificial conducting ground plane affect peak field strength, homogeneity, spectral content and time-domain fidelity in a radiated NEMP array test?

This question is asked specifically because the answer will also be the solution to the follow-up question: "how can we compensate for it?" In the EMIS-III-VPD, an electrically reflecting mesh was placed under the antenna and out towards the DUT, which was done to make use of the image theorem to halve the antenna size, but also has effect on the radiated far field. Say the centre of the radiating antenna is at 5 m above the horizontal plane, and the centre of the testing surface is also at 5 m above the horizontal plane at 100 m distance, then the indirect path via a ground reflection is 3.2 ns longer than the direct path: well within the FWHM and approximately as large as the rise-time. Knowing the effect of the sea-surface or the reflecting ground plane will also allow compensation for it or design with it in mind.

Idealized Pulse

To accurately assess the performance of the 4 element IRA, an idealized pulse was constructed that, when fed in CST back into the antenna, provided the best possible performance of the IRA despite its limited band. The two pulses fed to the IRA and the EMIS-III-VPD before this, the DEXP and the integral of the DEXP, provided results that were clearly suboptimal. The result was a numerical construct, holding no real physical meaning and not bound by real pulser limitations or artefacts. Before actual implementation of this idea, the truly optimal pulse is needed:

What is the ideal pulseform to be fed to an optimized 4 element array of antennas for the purpose of NEMP testing?

It is envisioned that the pulse will be constrained by true pulser capabilities. Furthermore, it may also be optimized constrained by antenna, mutual coupling and propagation effects between pulser and field. The project may be expanded to explore the possibility of feeding different pulses to different antennas.

Spectral-domain array optimization

This work aimed at optimizing the array for two competing parameters: peak E field homogeneity and peak time homogeneity. Spectral magnitude distribution across the testing surface was inspected but not optimized for. Therefore, this thesis proposes:

How can spectral magnitude requirements be incorporated directly into the optimization of a time-domain NEMP antenna array?

This way, notches in the spectral magnitude, as seen in the manual array in figure 6.9 can be prevented by incorporating spectral magnitude demands in the cost-function of the optimizer. As another design parameter is added to the optimizer, the optimizer is likely to yield lower performance on the parameters already included, as design space must also be allocated to the spectral magnitude. However, the final result will by its definition perform better across all metrics.

Incident field quality and the NEMP test

This thesis as set out to improve the quality of the radiated field, as defined by metrics such as peak field strength, amplitude homogeneity and peak time homogeneity, as close to the NEMP standards as possible. What remains to be explored is to what extent that is useful. At what point will improvement of the field metrics no longer significantly result in the increase of confidence with which the qualification for NEMP hardness can be made? Therefore, the follow research question is asked:

To what extent is it useful to improve field metrics to approach the true NEMP waveform for the improvement of the quality of an NEMP hardness validation test?

It is envisioned that the results of this work give further bounds to the goals as they are set in Chapter 2 of this thesis. That way over-engineering is prevented and eventually a working, practically implementable design can be delivered.

Chapter 8

Conclusion

This thesis set out to explore the feasibility of a timed array of ultra-wideband, directive radiators as an architecture for radiated NEMP hardness validation testing of large naval platforms. This was motivated by a lack of feasibility or quality from existing, conventional architectures: existing bounded-wave simulators cannot accommodate a full-scale warship, and increasing their size is impractical and therefore infeasible. The conventional method for testing warships was therefore to use large monocone antennas, but these did not achieve sufficient field strength with sufficient homogeneity over the testing surface to instil confidence in their operators. Therefore, this thesis has explored the feasibility of an unconventional simulator architecture: a time domain array of directive antennas aimed at a limited surface area. This thesis makes the deliberate choice to no longer intend to illuminate the entire DUT in a single test: instead choosing to trade testing surface size for increased homogeneity and peak electric field strength.

To approach this problem, the following research question has been asked:

To what extent can a timed array of directive, ultra-wideband radiators serve as a feasible radiated NEMP hardness validation method for large naval platforms?

To answer this question, several subquestions were posed in the first chapter of this thesis. Each of these has been answered in the preceding chapters. This work started by consulting the NEMP standards, where from these design parameters were distilled that aimed to simultaneously give the user of the simulator sufficient confidence in the hardness of the DUT while keeping the design feasible and eventually practically implementable. Next, the baseline, the EMIS-III-VPD monocone, performance was quantified in order to provide a comparison to any improvement achieved in this thesis. An antenna concept was selected from a shortlist of non-dispersive, UWB antennas, and this antenna was modelled. Finally, an array of these antennas was optimized for amplitude and peak-time homogeneity while constrained by a peak electric field strength value. In answering this question, several models were built and their results quantified. A summary of the results obtained in this thesis is presented below:

Table 8.1: Summary of key numerical results across the configurations studied in this thesis

Configuration	Surface [m]	Peak field [kV/m]	Min. field [kV/m]	Ampl. inhomog. [dB]	Peak-time spread [ns]
EMIS-III-VPD	200×50	1.28	—	9.9*	138
Single IRA (12.2 m)	10×10	4.5	4.3	0.37	0.80
Manual quadrant array	10×10	17.4	11.5	3.57	1.56
Optimised array	10×10	9.6	9.0	0.56	1.06

*Band-average max/min

These results quantify the improvement every model aimed at a limited testing surface size achieves over the EMIS-III-VPD baseline model. From these numerical findings the following main findings may be distilled.

- Full-illumination radiated NEMP hardness validation testing within the set constraints for homogeneity and peak electric field strength is unfeasible for large naval platforms. The baseline did not achieve any of the requirements set by this thesis.
- Antenna directivity is the enabling parameter and aperture size the limiting one. A practical aperture sacrifices the low-frequency content that governs time-domain fidelity, and arraying scales the radiated spectrum in magnitude but not in shape.
- An array is a peak-field multiplier, not a means of improving homogeneity. A single antenna is the most homogeneous and best-timed configuration of any studied; every element added to raise the field costs homogeneity and timing. Optimization minimises that penalty, recovering near-single-antenna homogeneity, but only up to a ceiling of roughly 9 kV/m, beyond which uniformity degrades steeply.
- The 50 kV/m of the standards is unreachable by this approach: neither feasible element counts nor feasible feed voltages close the gap. The binding limits are physical and external to the array, namely the dish aperture and voltage breakdown, whereas the required feed-network precision is demanding but within established pulsed-power practice.

The answer to the central research question is therefore not a definitive one, which is in line with the exploratory nature of this work. Within the constraints and assumptions this thesis has worked with, this work has shown that a timed array of directive antennas is a promising method for NEMP hardness validation testing. The array achieved good homogeneity over the field with simultaneously sufficient electric field strength for this thesis. It achieved this by trading the goal to fully illuminate the DUT, and instead by focusing on a limited testing surface.

The contribution of this thesis is not a finished simulator but a framework and a baseline for a problem that had received little open attention since the Cold War. It establishes which avenues are worth pursuing and which are not, and it bounds what may reasonably be expected of a radiated NEMP test of a large platform. The timed directive array is shown to be a credible building block, but one whose ceiling is now known. Significant research and design effort is still needed before such an NEMP simulator can be implemented. In aiding this effort, this work leaves a number of suggested research questions that may be explored, although this thesis underlines the fact that this list is not definitive.

Within that scope, the study has done what it set out to do. It has shown that the idea is worth pursuing, it has mapped the terrain on which that pursuit must proceed, and it has handed the questions it could not answer to the future work that follows.

Bibliography

- [1] E. Savage, J. Gilbert, and W. Radasky, “The early-time (E1) high-altitude electromagnetic pulse (HEMP) and its impact on the U.S. power grid,” Tech. Rep. Meta-R-320, Metatech Corporation, Goleta, CA, USA, 2010.
- [2] C. L. Longmire, “Justification and verification of high-altitude EMP theory, part I,” tech. rep., Mission Research Corporation, Santa Barbara, California, June 1986. Theoretical Note 368.
- [3] C. N. Vittitoe, “Did high-altitude EMP cause the Hawaiian streetlight incident?,” Technical Report SAND88-0043C, Sandia National Laboratories, Albuquerque, NM, USA, June 1989. System Design and Assessment Note 31.
- [4] B. Daout, “Selection of pulse simulators for HEMP / NEMP tests,” technical note TN18, Montena Technology SA, Rossens, Switzerland, n.d. Accessed: 2025-12-17.
- [5] Bundeswehr, “WIS nimmt modernen NEMP-Generator in Betrieb,” 2023.
- [6] Absolute EMC LLC, “HPD1000 outdoor horizontally polarized NEMP simulator,” 2025. Product page for a 10 m height, MIL-STD compliant NEMP simulator system.
- [7] Museum Waalsdorp, “Voormalige onderzoeksfaciliteiten: Elektro-Magnetische Puls (EMP) faciliteiten (1970–1993),” 2025.
- [8] C. de Haan, W. Pont, and M. Ouwens, “Feasibility study of a NEMP simulator for ships (EMPSIS),” Tech. Rep. FEL-91-B249, TNO Physics and Electronics Laboratory, September 1991.
- [9] P. Savat, “De Elektro-Magnetische Puls,” Natuur en Techniek, August 1982.
- [10] International Electrotechnical Commission, “Electromagnetic compatibility (EMC) – part 4-32: Testing and measurement techniques – high-altitude electromagnetic pulse (HEMP) simulator compendium,” Tech. Rep. IEC/TR 61000-4-32:2002(E), IEC, Geneva, Switzerland, 2002.
- [11] J. M. Nelson, C. S. Jones, and L. A. Poston, “USS Anzio EMPRESS II trial,” Naval Engineers Journal, vol. 108, no. 3, pp. 315–328, 1996.
- [12] Dassault Systèmes, SIMULIA, CST Studio Suite: High Frequency Simulation, 2025. CST Studio Suite documentation manual.
- [13] F. Tesche, “Numerical determination of a step response of a thin-wire scattering element arbitrarily located above a perfectly conducting ground plane,” tech. rep., Northrop Corporate Laboratories, Pasadena, California, USA, 1972. Sensor and Simulation Note 141.
- [14] D. M. Pozar, Microwave Engineering. Hoboken, New Jersey: John Wiley & Sons, 4th ed., 2012.
- [15] T. T. Wu and R. W. P. King, “The cylindrical antenna with nonreflecting resistive loading,” IEEE Transactions on Antennas and Propagation, vol. 13, pp. 369–373, May 1965.
- [16] C. A. Balanis, Antenna Theory: Analysis and Design. Hoboken, NJ: John Wiley & Sons, 4 ed., 2016.
- [17] European Commission, “F-gases in equipment and products,” 2024. European Commission, Directorate-General for Climate Action. Accessed: 2026-03-04.

- [18] “Electromagnetic compatibility (EMC) – part 2-9: Environment – description of HEMP environment – radiated disturbance,” *iec 61000-2-9:2025*, International Electrotechnical Commission (IEC), Geneva, Switzerland, 2025.
- [19] D. M. Pozar, “Waveform optimizations for ultra-wideband radio systems,” *IEEE Transactions on Antennas and Propagation*, vol. 51, no. 9, pp. 2335–2345, 2003.
- [20] C. Ge, R. Zhang, Y. Jiang, L. Cai, and B. Li, “Antenna de-embedding using deconvolution with Tikhonov regularization for mmWave channel measurement,” *IEEE Transactions on Antennas and Propagation*, vol. 70, no. 8, pp. 7024–7035, 2022.
- [21] D. V. Giri, W. D. Prather, and C. E. Baum, “The relationship between NEMP standards and simulator performance specifications,” tech. rep., Air Force Research Laboratory and University of New Mexico, Albuquerque, NM, 2009. Sensor and Simulation Note 538.
- [22] K. Deb, *Multi-Objective Optimization using Evolutionary Algorithms*. Wiley-Interscience Series in Systems and Optimization, Chichester, UK: John Wiley & Sons, 2001.
- [23] K. Miettinen, *Nonlinear Multiobjective Optimization*, vol. 12 of *International Series in Operations Research & Management Science*. Boston, MA: Kluwer Academic Publishers, 1999.
- [24] The MathWorks, Inc., *gamultiobj: Find Pareto Front of Multiple Fitness Functions Using Genetic Algorithm*, 2026. MATLAB Global Optimization Toolbox documentation. Accessed: 2026-05-30.
- [25] The MathWorks, Inc., *gamultiobj Algorithm*, 2026. MATLAB Global Optimization Toolbox documentation. Accessed: 2026-05-30.
- [26] B. Liang, G. Zhang, W. Zhang, W. Zhang, F. Wu, Y. Li, C. Zhu, and N. Yuan, “Synchronization of full-solid high-stability nanosecond pulse sources,” *High Power Laser and Particle Beams*, vol. 18, no. 6, pp. 1041–1045, 2006.
- [27] J. R. Mayes, W. J. Carey, W. C. Nunnally, and L. Altgilbers, “Sub-nanosecond jitter operation of marx generators,” in *Proceedings of the 13th IEEE International Pulsed Power Conference*, 2001.
- [28] E. G. Farr, “Analysis of the impulse radiating antenna,” Sensor and Simulation Note 329, Farr Research, Inc., Albuquerque, New Mexico, July 1991.
- [29] E. G. Farr, C. E. Baum, and D. V. Giri, “Review of impulse-radiating antennas,” tech. rep., Air Force Research Laboratory DEHP, Mar. 1998. Published on Farr-Research Inc.
- [30] M. Manteghi and Y. Rahmat-Samii, “On the characterization of a reflector impulse radiating antenna (IRA): Full-wave analysis and measured results,” *IEEE Transactions on Antennas and Propagation*, vol. 54, pp. 812–822, mar 2006.
- [31] A. Reyna, M. A. Panduro, et al., “Synthesis of timed antenna arrays for flat-top energy patterns,” *IETE Technical Review*, vol. 34, no. 1, 2017.
- [32] “High-altitude electromagnetic pulse (HEMP) environment,” *mil-std-2169c*, U.S. Department of Defense, Washington, D.C., December 2005. Defines the E1, E2, and E3 environments of a nuclear electromagnetic pulse (NEMP).
- [33] C. E. Baum, “Producing large transient electromagnetic fields in a small region: An electromagnetic implosion,” tech. rep., Air Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, NM, Aug. 2005. Sensor and Simulation Note 501.
- [34] “Department of defense interface standard: Requirements for the control of electromagnetic interference characteristics of subsystems and equipment,” *mil-std-461g*, requirement rs105, U.S. Department of Defense, Washington, D.C., December 2015.
- [35] C. E. Baum, “Radiation of impulse-like transient fields,” Sensor and Simulation Note 321, Air Force Weapons Laboratory, Kirtland AFB, New Mexico, Nov. 1989.

- [36] J. Rhebergen, A. Zwamborn, and D. Giri, "Design of an ultra-wideband ground-penetrating radar system using an impulse radiating antenna," Tijdschrift van het Nederlands Elektronika- en Radiogenootschap, Jan. 1999. Sensor and Simulation Note 575.
- [37] E. G. Farr and C. E. Baum, "Prepulse associated with the TEM feed of an impulse-radiating antenna," tech. rep., Farr Research & Philips Laboratory, Mar. 1992. Sensor & Simulation Note 337.
- [38] M. F. Miletić and D. I. Olćan, "Impulse radiating antenna with six feeding arms and a tapered balun," IEEE Transactions on Antennas and Propagation, vol. 70, pp. 6414–6422, Aug. 2022.
- [39] C. E. Baum, W. L. Baker, W. D. Prather, J. M. Lehr, J. P. O'Loughlin, D. V. Giri, I. D. Smith, R. Altes, J. Fockler, D. McLemore, M. D. Abdalla, and M. C. Skipper, "JOLT: A highly directive, very intensive, impulse-like radiator," Proceedings of the IEEE, vol. 92, no. 7, pp. 1096–1109, 2004.
- [40] M. C. Taylor and T. K. Sarkar, "An electromagnetic analysis of the 46 cm diameter impulse radiating antenna (IRA)," Tech. Rep. Sensor and Simulation Note 528, Syracuse University, Electrical Engineering and Computer Science Department, May 2008.
- [41] L. H. Bowen, E. G. Farr, C. E. Baum, T. C. Tran, and W. D. Prather, "Experimental results of optimizing the location of feed arms in a collapsible IRA and a solid IRA," tech. rep., Air Force Research Laboratory, Directed Energy Directorate, Nov. 2000. Sensor & Simulation Note No. 450.
- [42] C. E. Baum, "A lens technique for transitioning waves between conical and cylindrical transmission lines," tech. rep., Air Force Weapons Laboratory, Kirtland Air Force Base, New Mexico, Jan. 1967. Sensor & Simulation Note 32.
- [43] E. G. Farr, "Off-boresight field of a lens IRA," tech. rep., Farr Research, Albuquerque, New Mexico, Oct. 1994. Sensor & Simulation Note 370.
- [44] C. E. Baum, "Some considerations concerning analytic EMP criteria waveforms," Tech. Rep. TN-285, Air Force Weapons Laboratory, Kirtland Air Force Base, New Mexico, oct 1976. Theoretical Note 285.
- [45] L. H. Bowen and E. G. Farr, "A solid dielectric lens impulse radiating antenna surrounded by a cylindrical shroud," tech. rep., Air Force Research Laboratory, Kirtland AFB, NM, 2001. Sensor and Simulation Note 487.
- [46] E. G. Farr, "Development of a reflector IRA and a solid dielectric lens IRA, part II: Antenna measurements and signal processing," tech. rep., Phillips Laboratory, Kirtland AFB, NM, 1993. Sensor and Simulation Note 401.
- [47] W. S. Bigelow and E. G. Farr, "Lens TEM horns for sidelobe reduction," tech. rep., Air Force Research Laboratory, Kirtland AFB, NM, 2004. Sensor and Simulation Note 490.
- [48] B. Scheers, M. Acheroy, and A. V. Vorst, "Time-domain simulation and characterisation of TEM horns using a normalised impulse response," IEE Proceedings – Microwaves, Antennas and Propagation, vol. 147, pp. 463–468, Dec. 2000.
- [49] Z. N. Chen, "Development of ultra-wideband antennas," Journal of Electromagnetic Engineering and Science, vol. 13, no. 2, pp. 63–73, 2013. Open access. Available online: <https://pdfs.semanticscholar.org/ee23/9614b90a5d68195ffcc6a779d4a49df7f64c.pdf>.
- [50] C. A. Grosvenor, R. T. Johnk, D. R. Novotny, S. Canales, B. Davis, and J. Veneman, "TEM horn antenna design principles," Tech. Rep. NIST Technical Note 1544, National Institute of Standards and Technology, Boulder, CO, USA, January 2007.
- [51] E. G. Farr and C. E. Baum, "A simple model for small angle TEM horns," tech. rep., Air Force Weapons Laboratory, Kirtland AFB, NM, May 1992. Sensor and Simulation Note 340.
- [52] A. Mallahzadeh and F. Karshenas, "Modified TEM horn antenna for broadband applications," Progress In Electromagnetics Research, vol. 90, pp. 105–119, 01 2009.
- [53] P. J. Gibson, "The Vivaldi aerial," in Proc. 9th European Microwave Conference, (Brighton, UK), pp. 101–105, 1979.

- [54] A. Bhattacharjee, A. Bhawal, A. Karmakar, A. Saha, and D. Bhattacharya, "Vivaldi antennas: a historical review and current state of art," International Journal of Microwave and Wireless Technologies, vol. 13, no. 8, pp. 833–850, 2021.
- [55] A. K. Yerrola et al., "Design of wideband conformal Vivaldi antenna array," in URSI Radio Science Meeting, 2022.
- [56] O. J. Lodge, "Electric telegraphy," 1898. Filed February 1, 1898.
- [57] R. Kichouliya and M. J. Thomas, "Experimental simulation of low level hybrid electromagnetic pulse (EMP) for vulnerable studies on electronic systems and cables," in 2016 International Conference on Microelectronics, Computing and Communications (MicroCom), (Durgapur, India), IEEE, 2016.
- [58] D. V. Giri, "Time-domain radiated fields of a resistively loaded bi-conical antenna based on a transmission-line model," tech. rep., Pro-Tech, April 1994. Sensor and Simulation Note 366.
- [59] S. Li, J. Li, F. Wu, Q. Huang, Y. Gong, T. Yang, J. Liu, X. Zhou, and A. Qiu, "A hybrid simulator with novel bicone-wire grid antenna," Review of Scientific Instruments, vol. 95, March 2024. Open Access.
- [60] B. Cadilhon, B. Cassany, P. Modin, J.-C. Diot, V. Bertrand, and L. Pécastaing, "Ultra wideband antennas for high pulsed power applications," in Ultra Wideband Communications - Novel Trends - Antennas and Propagation (M. Matin, ed.), ch. 15, London: IntechOpen, 2011.
- [61] V. Gurevich, A Complementary View of HEMP for Electrical Engineers. Research Publish Journals, 2023. PDF available at researchpublish.com.
- [62] A. Qiu and J. Li, "The threats and research prospects of high-altitude electromagnetic pulse on power facilities," Smart Power & Energy Security, vol. 1, no. 1, pp. 1–11, 2025.
- [63] E. F. Vance, "Electromagnetic-pulse handbook for electric power systems," technical report no. dna 3466f, Stanford Research Institute / U.S. Defense Nuclear Agency, 1975.
- [64] "AECTP 250: Electrical and electromagnetic environmental conditions," Tech. Rep. AECTP-250, Edition 3, North Atlantic Treaty Organization, Brussels, Belgium, 2009. Allied Environmental Conditions and Test Publication; defines electrical and electromagnetic environments including HEMP/H-PEM.
- [65] "AECTP 500: Electromagnetic environmental effects tests and verification," Tech. Rep. AECTP-500, Edition E, Version 1, North Atlantic Treaty Organization, Brussels, Belgium, 2016. Defines electromagnetic environmental effects (E3) test methods, including for HEMP and HPEM.
- [66] M. Ianoz, B. I. C. Nicoara, and W. A. Radasky, "Modeling of an EMP conducted environment," IEEE Transactions on Electromagnetic Compatibility, vol. 38, pp. 400–413, Aug. 1996.
- [67] K. J. Moeller, K. L. Dudley, C. C. Quach, and S. E. Koppen, "In-flight characterization of the electromagnetic environment inside an airliner," Tech. Rep. NASA/TP-2001-210831, National Aeronautics and Space Administration, Langley Research Center, Hampton, Virginia, Mar. 2001.
- [68] International Electrotechnical Commission, "Electromagnetic compatibility (EMC) – part 4-23: Testing and measurement techniques – test methods for protective devices for HEMP and other radiated disturbances," Tech. Rep. IEC 61000-4-23:2016, Geneva, Switzerland, 2016. Consolidated version (with Amendment 1, 2025).
- [69] ETS-Lindgren, "Model 3110C biconical antenna (datasheet)." Datasheet, n.d. Accessed: 2026-01-19.
- [70] E. G. Farr, "Optimizing the feed impedance of impulse radiating antennas, part I: Reflector IRAs," tech. rep., Summa Foundation, Jan. 1993. Sensor and Simulation Note 354.
- [71] V. H. Rumsey, "Frequency independent antennas," Technical Report 20, Antenna Laboratory, Electrical Engineering Research Laboratory, University of Illinois, Oct. 1957.

- [72] Cybersecurity and Infrastructure Security Agency (CISA), “Electromagnetic pulse (EMP) protection and resilience guidelines,” tech. rep., U.S. Department of Homeland Security, Feb. 2019.
- [73] Department of Defense, “MIL-STD-188-125-2: High-altitude electromagnetic pulse (HEMP) protection for ground-based C4I facilities performing critical, time-urgent missions, Part 2: Transportable systems,” Tech. Rep. MIL-STD-188-125-2, U.S. Department of Defense, Mar. 1999.
- [74] E. G. Farr, L. H. Bowen, G. R. Salo, J. S. Gwynne, C. E. Baum, W. D. Prather, and T. C. Tran, “Studies of an impulse radiating antenna and a pulse radiating antenna element for SAR and target identification applications,” tech. rep., Air Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, New Mexico, Mar. 2000. Sensor and Simulation Note 442.
- [75] E. G. Farr, “IRA-2 dimensions.” Farr Research, Inc. internal document, Mar. 2000. Provided with minor edits in the appendix of Sensor and Simulation Note 528.
- [76] K. Kim and W. R. S. Jr., “Numerical analysis of the impulse-radiating antenna,” tech. rep., Georgia Institute of Technology, June 2003. Sensor and Simulation Note 474.
- [77] D. V. Giri, H. Lackner, I. D. Smith, D. W. Morton, C. E. Baum, J. R. Marek, W. D. Prather, and D. W. Scholfield, “Design, fabrication, and testing of a paraboloidal reflector antenna and pulser system for impulse-like waveforms,” IEEE Transactions on Plasma Science, vol. 25, pp. 318–326, Apr. 1997.
- [78] L. H. Bowen, E. G. Farr, and D. I. Lawry, “A dual-polarity impulse radiating antenna,” tech. rep., Air Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, New Mexico, Oct. 2003. Sensor and Simulation Note 479.
- [79] F. Sabath, “Near field dispersion of impulse radiation,” in Proceedings of the XXVIIth URSI General Assembly, (Maastricht, The Netherlands), p. paper p0580, International Union of Radio Science (URSI), Aug. 2002.
- [80] S.-f. Wang and Y.-z. Xie, “Far-field boundary estimation for the high-power UWB pulsed antennas,” IET Microwaves, Antennas & Propagation, vol. 12, no. 15, pp. 2199–2205, 2018.
- [81] O. V. Mikheev, S. A. Podosenov, K. Y. Sakharov, Y. G. Svekis, and V. A. Turkin, “New method for calculating pulse radiation from an antenna with a reflector,” IEEE Transactions on Electromagnetic Compatibility, vol. 39, pp. 48–54, Feb. 1997.
- [82] S. Altunc, C. E. Baum, C. G. Christodoulou, E. Schamiloglu, and C. J. Buchenauer, “Focal waveforms for various source waveforms driving a prolate-spheroidal impulse radiating antenna (IRA),” Radio Science, vol. 43, no. 4, p. RS4S13, 2008.
- [83] C. E. Baum, “Focal waveform of a prolate-spheroidal impulse-radiating antenna,” Radio Science, vol. 42, no. 6, 2007.
- [84] J. S. Tyo, “Optimization of the TEM feed structure for four-arm reflector impulse radiating antennas,” IEEE Transactions on Antennas and Propagation, vol. 49, pp. 607–614, Apr. 2001.
- [85] C. E. Baum, “Selection of angles between planes of TEM feed arms of an IRA,” tech. rep., Air Force Research Laboratory, Directed Energy Directorate, Aug. 1998. Sensor and Simulation Note 425.
- [86] C. E. Baum, “Aperture efficiencies for iras,” Tech. Rep. 328, Sensor and Simulation Notes, June 1991.