

Studies of Breakdown and Pre-Breakdown Phenomena in High Gradient Accelerating Structures

Jan Paszkiewicz
St. John's College, Oxford



*Thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy in Particle Physics at the University of Oxford*

Trinity Term 2020



Abstract

Vacuum breakdown is a complex process and an important limiting factor of the performance of normal-conducting high-gradient particle accelerators, and can result in loss of luminosity in particle collider applications, as well as damage to accelerating structures. The work presented here was done in the context of the Compact Linear Collider (CLIC) study, but is also relevant to a variety of applications such as medical linear accelerators or high-electric-field vacuum electronics. The aspects of vacuum breakdown discussed in this thesis together provide the theoretical basis for important technological parameters for the design of normal-conducting high-gradient devices: the influence of radio-frequency design and material properties on the achievable field, and the mechanism of conditioning.

The first part of this thesis discusses the development of an improved quantitative limit which determines the maximum accelerating gradient at which a given structure geometry could operate, with the intention of guiding the design of improved accelerating structures. It models the coupling of radio-frequency power to a breakdown, giving a value for surface electric field when loaded by a nascent breakdown. Calculations were performed on various cases that were tested experimentally, showing excellent consistency and the potential to become a very general model of vacuum breakdowns.

The second part presents an experimental study of dislocation dynamics in copper surfaces subject to high electric fields, to better understand the mechanism of the nucleation of breakdowns. This is believed to involve stochastic deformation of microscopic features under high-electric-field stress. Field-emitted current from radio-frequency structures, as well as electrodes subject to a static electric field were precisely measured, revealing small fluctuations in the latter case. The dependence of the rate of events on the surface field and the distribution of time intervals between events was found to match prior theoretical predictions. The rate of fluctuations was also found to decrease with the cumulative number of voltage pulses applied, supporting the idea of work hardening due to stress from pulsed electric fields as a mechanism for conditioning.

Acknowledgements

I wish to express my gratitude to Dr. Walter Wuensch, my CERN supervisor, and Professor Philip Burrows, my Oxford supervisor, for providing the opportunity to work with the CLIC RF group, many interesting and stimulating discussions, advice both technical and otherwise, great patience, and opportunities to attend numerous conferences around the world. I am grateful to CERN, the John Adams Institute, and the Department of Physics at the University of Oxford for supporting my work throughout the course of my DPhil.

I am grateful to my colleagues at the XBoxes and the CLIC group for introducing me to the facility, helping me with my work and most importantly for making my time with this group a memory I will cherish forever, in particular but not only Dr. Benjamin Woolley, Dr. Matteo Volpi, Dr. Samantha Pitman, Kamil Szypuła, Amelia Edwards, David Bañón Caballero, Iaroslava Profatilova, and Ruth Peacock.

Finally, I would like to thank my parents Adam and Barbara for their unwavering support and faith in me.

Contents

1	Introduction	9
1.1	Particle Colliders	9
1.2	The Compact Linear Collider	12
1.3	RF Accelerating Structures for CLIC	14
1.4	The High-Gradient X-Band Test Stands at CERN	17
1.5	DC Spark Systems	20
1.6	Vacuum Breakdowns	21
1.7	Electron Field Emission	25
1.8	Thesis Outline	27
2	Breakdown Criterion	29
2.1	Motivation	29
2.1.1	Kilpatrick's Criterion / Peak Electric Field	30
2.1.2	P/C	31
2.1.3	Modified Poynting Vector	34
2.1.4	An Improved Quantity	37
2.2	Proposed Model	37
2.2.1	Analytical Calculation of R_{bd}	44
2.3	Numerical Calculations and Experimental Results	54
2.3.1	Simulated CLIC-Like Accelerating Structures	54
2.3.2	Antenna Length Dependence in Numerical Simulations	67
2.3.3	Comparison of the Analytical Model with Numerical Simulations	70

2.3.4	CLIC Crab Cavity	73
2.3.5	T24 Structures	79
2.3.6	S-Band Backward Travelling Wave Structure	86
2.3.7	Choice of k Value	93
2.4	Improvements of the Method	96
2.4.1	Use of Z-Parameters	97
2.4.2	Finite-Time Broadband Quantity	99
2.5	Extension to DC	111
2.5.1	Breakdown Times	113
2.5.2	Breakdown Locations	114
2.5.3	Gap Size Dependence of Maximum Field	116
2.5.4	E^* for a Theoretical DC Geometry	118
2.5.5	Conclusions	128
2.6	Conclusions and Future Work	128
2.6.1	ITER Neutral Beam Injector	129
2.6.2	Choke-Mode Accelerating Structures	131
3	Dark Current Fluctuations	135
3.1	Motivation	135
3.1.1	Breakdown Statistics and Conditioning	135
3.1.2	Dislocation-Driven Breakdown Nucleation	136
3.1.3	Dark Current Fluctuations	139
3.1.4	Past Measurements	141
3.1.5	Approach	143
3.2	Searches in RF Structures	144
3.2.1	Experimental Setup	144
3.2.2	Signal Acquisition	147
3.2.3	Data Analysis	158
3.2.4	Conclusion	186
3.3	Measurement in the DC Spark System	187

<i>CONTENTS</i>	7
3.4 Discussion of Results	212
4 Conclusions	217

Chapter 1

Introduction

1.1 Particle Colliders

Particle colliders are machines which collide two high-energy beams of particles, in order to study the properties of subatomic particles and their interactions based on the behaviour of the collision products. The Large Hadron Collider is the largest currently operational particle collider, with a circumference of 27 km and a maximum dipole magnet strength of 8.3 T, allowing it to produce proton-proton collisions at a world-record Centre-of-Mass (CoM) energy of 13 TeV [1]. In 2012, the Higgs boson was experimentally observed for the first time at the LHC, completing the Standard Model of particle physics [2].

Following the discovery of the Higgs boson, there is a strong interest in the particle physics community in studying its properties, such as its mass, spin, couplings, and charge conjugation parity (CP) symmetry with much greater precision than is currently possible with the LHC. This is being partially addressed with the High Luminosity upgrade of the LHC, which will increase the integrated luminosity of the LHC to $250 \text{ fb}^{-1}/\text{year}$ [3], compared with a total integrated luminosity of 189 fb^{-1} obtained to date from the start of the LHC's operation [4]. The luminosity L may be defined for a circular collider as [5]:

$$L = \frac{f_{cross} N_1 N_2}{4\pi\sigma_x\sigma_y}, \quad (1.1)$$

where f_{cross} is the bunch crossing frequency, N_1 and N_2 are the numbers of particles in each of the colliding bunches, and σ_x and σ_y are the standard deviations of the transverse position profile of the beam, in the horizontal and vertical planes respectively. The number of Higgs bosons observed is directly proportional to the luminosity of the machine. This upgrade will thus allow the confidence margins on the measured parameters of the Higgs boson to be reduced due to the improved statistics.

The use of proton-proton collisions at the LHC is not ideal for precision measurements of the properties of the Higgs boson, due to the quantum chromodynamic (QCD) background produced. This arises from the large number of particles produced as a result of the strong interaction in the proton-proton collisions that are not scientifically interesting, yet still saturate the particle detectors and require a great deal of computing power to analyse and discard. Also, as protons are composed of quarks and gluons, the initial state of the colliding particles is not precisely known. Only the total energy of the fundamental particles comprising the proton is known, but not the way it is distributed between them at the time of collision, limiting the precision of the overall measurement. Leptons, on the other hand, do not interact with the strong force, meaning that lepton collisions would not directly produce QCD background, providing a much ‘cleaner’ signal. Since leptons are fundamental particles, the problem of the initial state of the colliding particles is also solved, allowing greatly improved kinematic control of the collisions.

These properties make a lepton collider a very attractive choice for a future high-energy-physics machine. Proposals for future lepton colliders include the Compact Linear Collider (CLIC), the International Linear Collider (ILC), the Future Circular Collider (FCC-ee), as well as the Circular Electron-Positron Collider (CEPC). The maximum CoM energy of each of these proposed machines is given in Table 1.1. For each option, the energy has been chosen to allow the production of Higgs bosons, and

Machine	Max CoM Energy (GeV)
CLIC (1st stage)	380
ILC (1st stage)	250
FCC-ee	360
CEPC	240

Table 1.1: Maximum centre of mass energies of proposed future e+e- colliders [6].

in the case of CLIC and FCC-ee, top-antitop quark pairs. All of these machines would collide electrons with positrons, since they are the only stable leptons.

Producing high-energy electron-positron collisions in a circular collider poses a technical challenge due to synchrotron radiation, which causes a beam of charged particles to lose energy when bent into a circular trajectory. The power lost by such a beam, P_{loss} , scales as [5]:

$$P_{loss} \propto \frac{E^4}{m_0^4 \rho^2}, \quad (1.2)$$

where E is the particle energy, m_0 is the rest mass, and ρ is the bending radius. Since the rest mass of an electron is about 1800 times smaller than that of a proton, the power lost by an electron due to synchrotron radiation would be thirteen orders of magnitude greater than that of a proton at the same energy. Thus, a circular lepton collider would need to be very large in diameter to reach energies relevant for the production of Higgs bosons without impractically large synchrotron losses. Because of this, proposals for future circular electron-positron colliders feature very large machines such as FCC with a circumference of 100 km.

The alternative to a large circular collider would be a linear collider, which would not suffer from any limitations in energy due to synchrotron losses. However, a linear configuration means that the particles cannot be recirculated and can only pass once through the machine. In order to reach the energies proposed with a linear collider, it would either need to be very long and/or have a very high accelerating gradient, which quantifies the energy gain per unit length.

1.2 The Compact Linear Collider

The Compact Linear Collider is one of the proposed electron-positron colliders that could be built as a successor to the LHC. It is, as its name suggests, a linear machine. A key feature of the CLIC design is its very high accelerating gradient of 100 MV/m which allows it to have a relatively short length. By comparison, the ILC has a design gradient of 31.5 MV/m [8]. Examples of high-gradient machines currently in operation include the European X-Ray Free Electron Laser (E-XFEL) with a design gradient of 24 MV/m [9], and the SwissFEL with a gradient of 26 MV/m [10]. The 100 MV/m figure for CLIC thus represents a significant step up over the current state of the art.

Because CLIC is linear, it has the advantage of being able to be constructed in stages, with an increasing CoM energy at each stage. The proposed first stage of CLIC would have an energy of 380 GeV and an overall length of 11 km, while the final stage would reach an energy of 3 TeV with a length of 50 km [11].

The CLIC study includes proposals for a two-beam acceleration scheme, or a single-beam, klystron-based machine. The two-beam, 3 TeV variant is depicted in Fig. 1.1. In this scheme, two high-current drive beams (one each for the electron and positron linacs) are accelerated to a relatively low energy of 2.4 GeV. Each drive beam would then be directed to Power Extraction and Transfer Structures (PETS), which would extract power from the drive beam, decelerating it, and transfer this power to the accelerating structures of the main linacs. The electron and positron beams accelerated by the main linacs would collide at the interaction point. The two-beam scheme would be more efficient and more cost-effective on large scales than generating radio-frequency (RF) power directly with 12 GHz klystrons, which would be the case in the klystron-based option.

The high-gradient acceleration technology developed as part of the CLIC study has been finding more and more applications beyond CLIC itself. These include medical linacs, X-ray free-electron lasers (XFELs), and inverse Compton scattering sources. In

each of these, there is a demand for building linear accelerators that are as compact as possible, though not on as large a scale as CLIC.

1.3 RF Accelerating Structures for CLIC

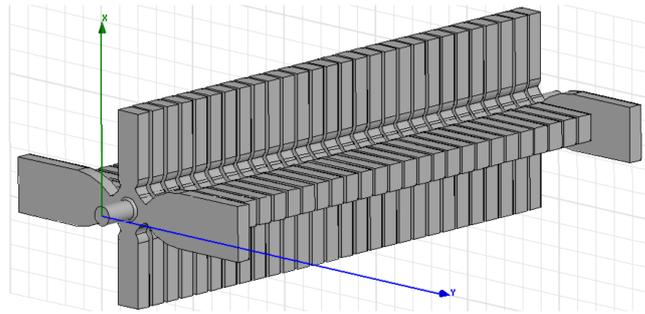


Figure 1.2: The internal volume of the nominal accelerating structure design for the 3 TeV version of CLIC [12].



Figure 1.3: A copper disc containing a single cell with wakefield damping features.

The CLIC main linacs will consist mainly of accelerating structures, in which oscillating electric fields are set up to accelerate the electrons and positrons to high energies. Each structure consists of a number of cells which act as resonators at microwave frequencies. When a bunch of electrons or positrons passes through the cell, it gets accelerated by the longitudinal component of the electric field in the cell. Here, the cells are electrically coupled to one another to form a periodic structure. The accelerating structures in CLIC are travelling-wave structures, meaning that they have both an input and output port, and when a radiofrequency (RF) wave is injected at the input port it will eventually propagate to the output port. This is in contrast to the standing-wave structures used in many particle accelerators, which have only one port and build

up a standing wave when fed with RF power. Standing-wave structures are typically weakly coupled to transmission lines or waveguides to minimise energy losses and thus attain a high quality factor and thus high power efficiency. This, however, makes the filling time of such structures very long, making them unsuitable for very short RF pulses. On the other hand, CLIC's use of travelling-wave structures allows the use of very short RF pulses (below 200 ns) so that the accelerating gradient can be made as high as possible while keeping the breakdown rate acceptably low [13].

The internal volume of a CLIC structure is shown in Fig. 1.2. The structures will be manufactured out of copper discs, an example of which is depicted in Fig. 1.3, which are to be diffusion bonded together to form a full structure. The nominal structure design has twenty-six regular cells and two coupling cells (which couple the structure to the input and output waveguides).

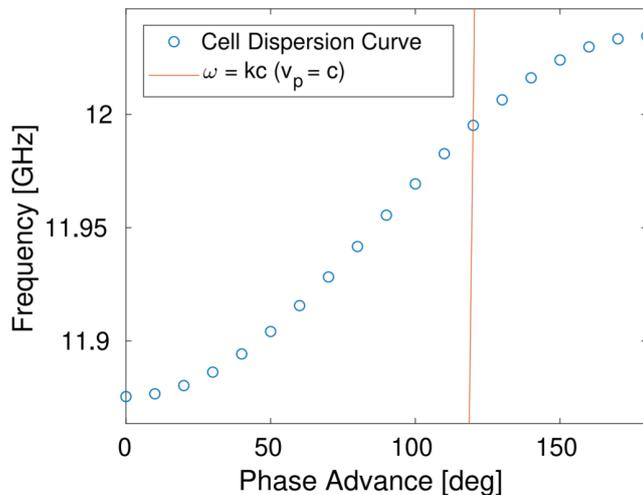


Figure 1.4: Frequency in GHz vs. phase advance in degrees. Blue circles: dispersion relation of the middle cell in a TD26CC prototype structure. Red line: a constant phase velocity of c , representing relativistic electrons [14].

In order to maximise the energy gained by a bunch crossing a cell, it has to be synchronised with the oscillating electric field in the cell such that the field is close to its peak when the bunch crosses the cell. For this synchronicity to apply for all of the cells in a structure, it has to be designed such that the phase velocity, which is the velocity at which the point of peak amplitude propagates through the structure, matches that

of the particle bunch. Since the beams within the main linacs are highly relativistic, the phase velocity has to be c .

Other key parameters of CLIC accelerating structures include the operating frequency of 11.994 GHz, and 120° phase advance from cell to cell [12]. A dispersion curve for one of the cells in a prototype is shown in Fig. 1.4, showing how the synchronicity condition is met at the design frequency of 11.994 GHz. The structures for the 380 GeV initial stage of CLIC are designed to operate at an average beam-loaded accelerating gradient of 72 MV/m, whereas the structures used in the upgrade to 3 TeV are designed for an average beam-loaded gradient of 100 MV/m [11].

The 100 MV/m gradient specification precludes the use of superconducting materials, as the superconductive state cannot exist when a magnetic field exceeding the critical magnetic field is applied. The critical magnetic flux for pure niobium is 180 mT. Since the peak magnetic flux within a resonant cavity is directly proportional to the accelerating gradient, the critical magnetic flux places a limit on the maximum gradient. The proportionality constant between the two, B_{peak}/E_{acc} , is a function of the cavity geometry. The geometry of the Low-Loss Cavity design for the ILC puts the ultimate accelerating gradient at 50 MV/m [15]. To overcome the gradient limitation posed by quench, CLIC structures are instead made of copper and operated at room temperature. In contrast with superconducting structures, one of the main limitations on the achievable gradient in normal-conducting structures is the phenomenon of vacuum breakdown. A significant amount of research and development effort has been spent on overcoming this limitation to maximise the achievable gradient, both in the context of CLIC and other particle accelerators around the world.

Breakdowns are undesirable as they can not only damage the accelerating structures, but also apply an undesired transverse kick to the particle beam, which results in the beams not colliding at the interaction point, and thus a loss of luminosity. The specifications for CLIC limit the allowable breakdown rate (BDR) to 3×10^{-7} breakdowns per pulse per meter of length, in order to keep loss of luminosity below 1% [16].

1.4 The High-Gradient X-Band Test Stands at CERN



Figure 1.5: The bunker housing XBox 2 and XBox 3 [17].

Prototype CLIC accelerating structures are tested under high power without beam at the High Gradient X-Band Test Facility at CERN, which consists of the test stands XBox 1, 2 and 3. XBox 2 and 3 are depicted in Fig. 1.5. One of the goals of these tests is to characterise the breakdown performance of the prototypes in order to guide the design process for future structures.

In these test stands, RF pulses with a power of up to 55 MW and a duration of up to 200 ns are generated using a network of high-power klystrons and pulse compressors [18]. XBox 1 and 2 can operate at a repetition rate of up to 50 Hz, while XBox 3 can deliver up to 200 Hz per structure under test. The structures under test, along with the waveguide networks, are operated under ultra-high vacuum (UHV) conditions, with the pressure being kept below 10^{-8} mbar during RF pulsing. The test stands are equipped with diagnostics including a number of directional couplers at various points of the high-power waveguide network, Faraday cups along the beam axis at each end of every structure under test, a radiation monitor, and optical fibres which serve to detect charged particles via Cherenkov radiation.

An integral aspect of the operation of the XBoxes and the prototype structures is a process known as ‘conditioning’. It has been found that a freshly manufactured accelerating structure undergoes breakdown at relatively low gradients far below the design value. In order to reach nominal operating conditions, the structures tested at the XBoxes undergo a several-month-long conditioning process in which they gradually become more resilient to breakdowns after repeated exposures to high fields [19]. Over this period, the electric field and pulse length are gradually increased up to the nominal parameters. The process of conditioning is not yet fully understood, though it is believed that important roles are played by the removal of surface contaminants, and especially by work hardening of the copper surface driven by the mechanical stress associated with the large pulsed electric field, as will be described in Sec. 3.

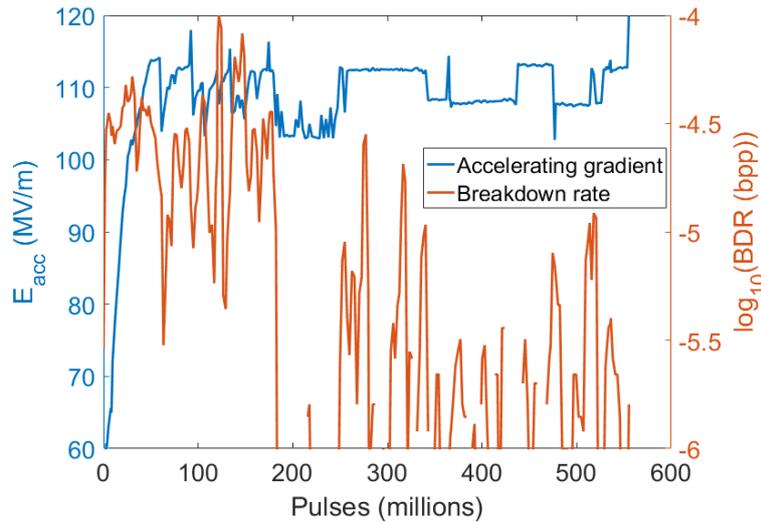


Figure 1.6: The conditioning history of the T24PSI2 structure when it was tested in XBox 2. Blue plot: unloaded accelerating gradient in MV/m vs. cumulative number of RF pulses applied. Red plot: logarithm of the breakdown rate vs. cumulative number of RF pulses applied

The amount of RF power sent to a structure under test on each pulse is controlled by a conditioning algorithm which monitors the current BDR [20], and attempts to keep it at roughly 1×10^{-5} breakdowns per pulse. Fig. 1.6 shows the conditioning history of a structure tested in XBox 2; between the 0 and 180 million pulse marks, the RF power, and thus the accelerating gradient, were gradually increased by the conditioning algorithm. The RF pulse length was also manually increased from 50 ns

to 200 ns in steps of 50 ns. The breakdown rate in this period stayed below 1×10^{-4} breakdowns per pulse, showing an improvement in breakdown performance. After the 180 million pulse mark, the structure was deemed fully conditioned, as the maximum gradient for a given breakdown rate stopped increasing with additional conditioning. After this, a series of breakdown rate measurements were performed under different conditions. This corresponds to a drop in the BDR below 1×10^{-5} breakdowns per pulse, demonstrating stable performance at high accelerating gradients.

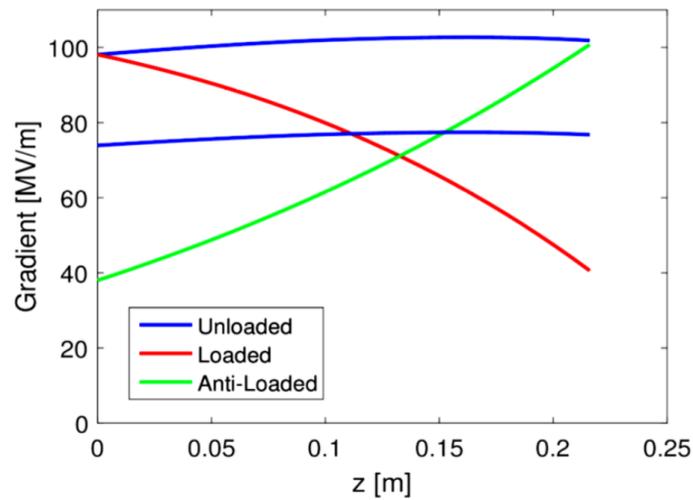


Figure 1.7: Local accelerating gradient for a prototype CLIC structure in MV/m vs. longitudinal position in m. Top blue line: with no beam loading, at an input power of 43.3 MW. Bottom blue line: with no beam loading, at an input power of 24.6 MW. Red: with loading by 1.6 A of beam current and 43.3 MW input power. Green: anti-loaded (i.e. with the RF phase chosen such that the beam is decelerated by the structure) with 1.6 A of beam current and 6.5 MW input power [21].

The nominal design for CLIC accelerating structures specifies very strong beam loading to maximise the RF-to-beam power efficiency. This means that a significant proportion of the RF power incident on each accelerating structure is absorbed by the beam, rather than being absorbed by the RF loads connected to the structure outputs. As CLIC structures are travelling-wave structures, the RF power from the input ports of the structure must flow through each cell in turn to reach the output. As it does so, some of the power is lost in each cell due to the resistance in the copper walls of the structure. If a beam is present, additional power is transferred from the RF wave to the beam in each cell. This means that there is a significant difference in the power flowing through

each cell between the loaded and unloaded cases, and this has an influence on the gradient profile. Fig. 1.7 compares gradient profiles of a prototype CLIC accelerating structure under various beam loading conditions. Because the gradient profile of a travelling-wave structure also changes under beam loading, it is not sufficient to simply increase the input power to compensate for the power transferred to the beam. This means that the breakdown performance of a structure under testing without beam will not be exactly the same as that with beam loading, both in terms of overall breakdown rate and the spatial distribution of breakdowns [21].

1.5 DC Spark Systems

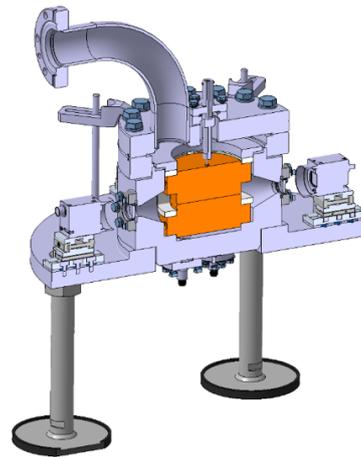


Figure 1.8: The Large Electrode System (LES) at CERN. Left: Photograph of the vacuum chamber. Right: Computer model showing the copper electrodes inside the chamber.

The DC spark systems at CERN are a complementary experiment designed to study vacuum breakdowns in a less complex environment than the XBoxes. The systems currently in operation at CERN are known as the Large Electrode Systems (LES) 2 and 3 [22], shown in Fig. 1.8. In these systems, vacuum breakdowns on the surface of copper electrodes are studied. The electrodes are manufactured using the same material with the same mechanical tolerances and subject to the same thermal cycles as CLIC accelerating structures to ensure the applicability of the results to the design of CLIC structures. Instead of RF pulses, the electrodes are subject to a pulsed DC voltage

which creates electric fields of up to 100 MV/m in the gap between the electrodes. The electrodes can have a circular area of up to 60 mm in diameter exposed to high electric fields, and the gap can be varied between 20 and 100 μm by using different ceramic spacers between the electrodes.

The electrodes in the LES exhibit similar conditioning behaviour to CLIC accelerating structures, needing a number of high-voltage pulses to be applied before a high electric field can be reached at a low breakdown rate. One advantage of the LES over the XBoxes is the much higher maximum pulse rate. Thanks to the use of a semiconductor-based pulse generator which can operate at up to 6 kHz [23], the conditioning time can be reduced from months, typical in the XBoxes, to days.

1.6 Vacuum Breakdowns

It has been established that vacuum breakdowns are a significant limiting factor of the maximum accelerating gradient of normal-conducting structures. Their study is of relevance to the design of high-gradient accelerators, as well as other applications including fusion power, ion thrusters, and electrical vacuum interruptors.

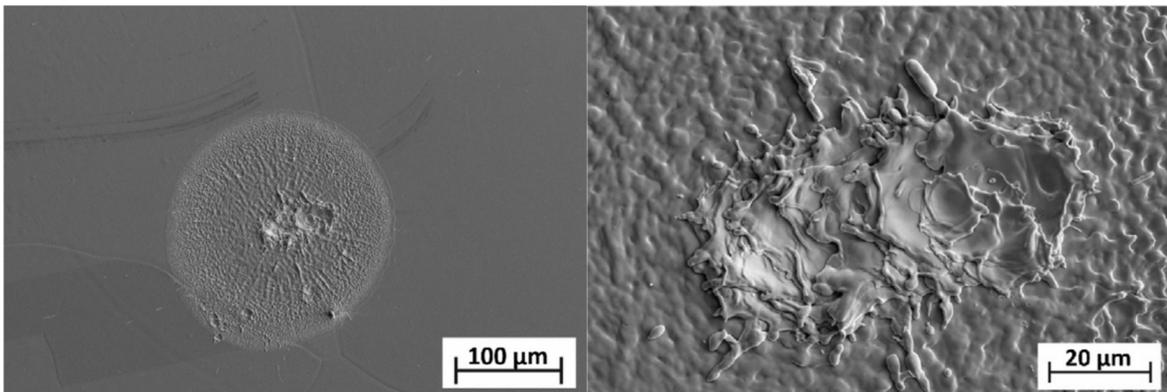


Figure 1.9: A breakdown crater on an electrode tested in the LES [24].

Apart from the loss in luminosity that occurs when the beam is present, another undesirable result of breakdowns in an accelerating structure is damage to the structure's surface. Fig. 1.9 shows an example of a crater left behind by a breakdown. This is caused by local melting of the copper surface by the very high temperature plasma [25].

The plasma, being conductive, also reflects any RF power incident on the structure back to the source. Excessive reflected power may cause damage to certain high-power RF components and power sources as well. In the XBox test stands, the timing of the reflected signal relative to the drop in power transmitted to the outputs of the structure is used to determine the approximate location of the breakdown along the beam axis.

Fig. 1.10 shows the stages of a vacuum breakdown. Field emission is believed to play an important role in initiating a breakdown (see Fig. 1.10, box a). This is the emission of electrons into the vacuum via tunnelling when the copper surface is subject to very large electric fields. Though the emission sites are on the order of nanometers in size, the current density can be very high, and cause significant Ohmic heating of the emission site. This heating can cause the emission of neutral copper atoms into the vacuum as well. The neutral copper gas just above the emitter can become ionised by collisions with the field-emitted electrons, which rapidly gain energy due to the high electric field (see Fig. 1.10, box b). The positive copper ions then become accelerated back towards the surface by the electric field, and bombard the emission site causing it to heat up and begin melting (see Fig. 1.10, box c). This causes the rapid emission of more neutral atoms which also become ionised, as well as a significant increase in the rate of emission of electrons due to thermionic emission. This creates a rapidly growing plasma sheath above the emission site, while further bombardment of the surface with ions and clusters of neutral atoms expands the area of the surface affected. The breakdown then stabilises into a quasi-neutral column of plasma with a thin plasma sheath across which the full potential of the plasma is held, bordering the copper surface. Recombination of ions and electrons is responsible for the visible glow of the breakdown plasma (see Fig. 1.10, box d). The plasma will continue burning until the power source is switched off, after which it will cool down and the electrons and ions recombine into neutral atoms, which then settle on the surroundings (see Fig. 1.10, box e). The crater is left behind, often with sharp features on its edges which enhance the surface field on their tips, leading a large field emitted current, and a relatively

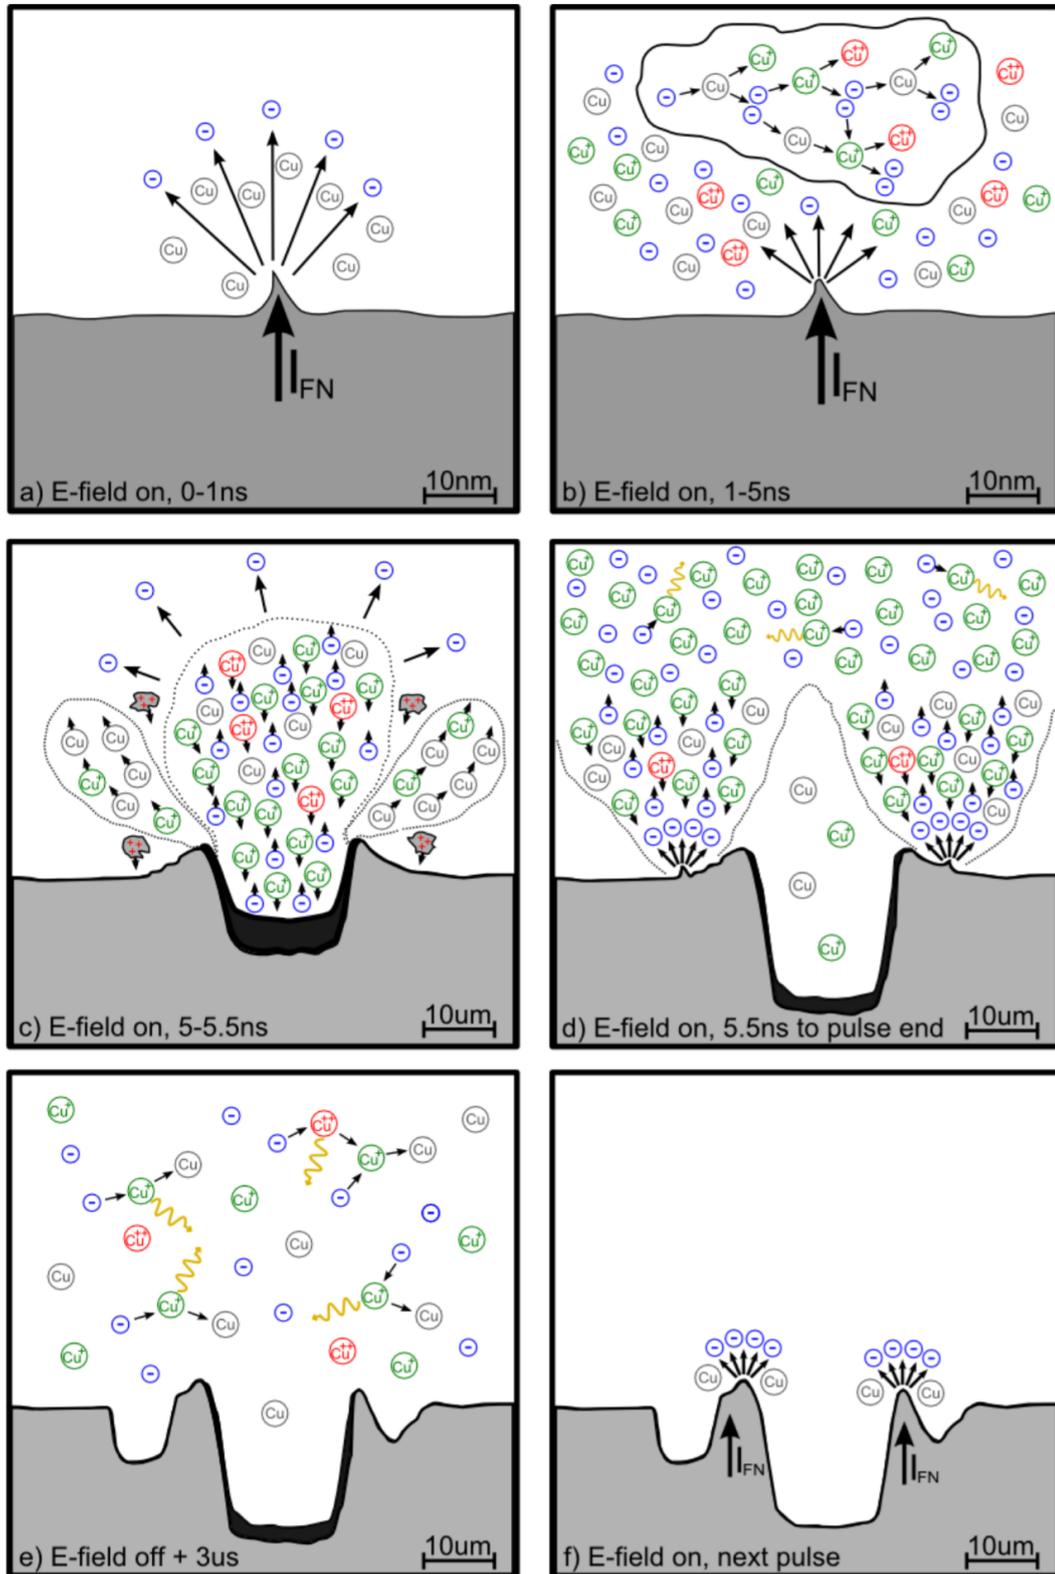


Figure 1.10: The stages of a vacuum breakdown [26]. For a description of panels a) - f), see text.

high probability of a breakdown occurring at that location again (see Fig. 1.10, box f) [27].

Breakdowns in CLIC accelerating structures and the LES occur randomly, with no known indicator of an imminent breakdown. Experiments have shown [28] that, for a certain state of conditioning, the mean breakdown rate BDR of a sample varies approximately as:

$$\text{BDR} \propto E^{30} \tau^5, \quad (1.3)$$

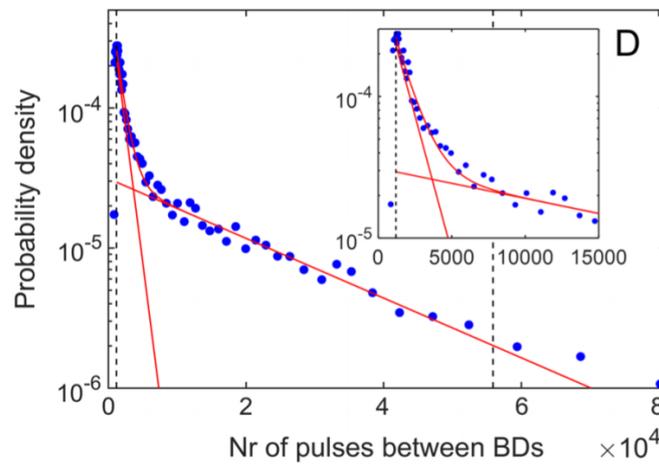


Figure 1.11: Probability density of the interval between two consecutive breakdowns being a certain number of pulses vs. number of pulses between consecutive breakdowns. Blue points: experimental data. Red lines: double-exponential fit [29].

where E is the surface electric field and τ is the pulse length. The proportionality constant depends, among other factors, on the conditioning of the structure and its geometry. Fig. 1.11 shows the probability density function (PDFs) of the number of pulses between successive breakdowns in an experiment conducted with the LES. Two exponential PDFs, which appear as straight lines on semi-logarithmic axes, have been fitted to the data. An exponential distribution of time intervals implies a Poisson point process, i.e. one in which events occur continuously and independently at a constant average rate. The fact that there are two distributions implies that the breakdown rate is enhanced for a certain period after a breakdown before returning to the ‘normal’

breakdown rate. This results in the tendency for breakdowns to be grouped temporally in clusters. Such behaviour is consistent with the description presented earlier and illustrated in Fig. 1.10, where it was stated that the sharp edges of breakdown craters increase the probability of a breakdown. The breakdown rate does eventually drop back to the value it had just before the beginning of the cluster, showing the effect of conditioning.

This behaviour is also shown in Fig. 1.12, which shows the change in longitudinal position between consecutive breakdowns in a CLIC structure. The double-exponential PDF can be seen as before, and the top plot shows that a breakdown that follows another closely in time tends to also occur nearby spatially.

1.7 Electron Field Emission

Field emission occurs when the surface of a conductor is subjected to high electric fields, which allow electrons within the metal to escape into the vacuum via quantum tunnelling. This current, known as dark current in the context of accelerating structures, becomes significant at macroscopic surface fields of several hundred MV/m. Dark current can be measured in both CLIC accelerating structures and the LES. Although field emission is thought to be the mechanism responsible for the onset of a breakdown, dark current can be measured on every pulse at high enough field levels. This suggests that there is a large population of stable field emitters at all times that do not undergo thermal runaway. The time-averaged emitted current \bar{I}_F for a sinusoidally varying electric field is given by [30]:

$$\bar{I}_F = \frac{5.7 \times 10^{-12} \times 10^{4.52\phi - 0.5} A_c (\beta E_0)^{2.5}}{\phi^{1.75}} \exp\left(-\frac{6.53 \times 10^9 \times \phi^{1.5}}{\beta E_0}\right), \quad (1.4)$$

where E_0 is the macroscopic surface electric field, ϕ is the material workfunction, and A_c is the emitting area. Typically, the current behaves as if the local surface field is a factor β higher than the macroscopic field, which is included in (1.4). A typical

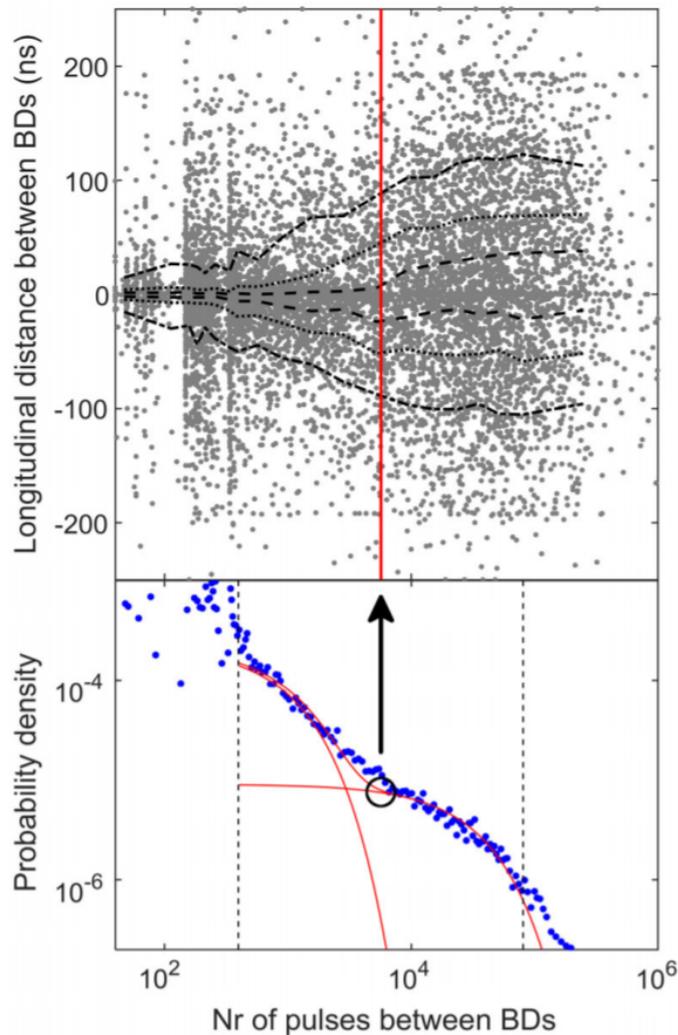


Figure 1.12: Top: the change in longitudinal position between consecutive breakdowns, in units of signal propagation time, vs. number of pulses between the breakdowns. Grey points represent individual breakdowns. Dashed, dotted, and dash-dot lines show, respectively, the intervals containing the middle 25%, 50%, and 75% of the data. Bottom: probability density vs. number of pulses between consecutive breakdowns. The red lines represent a double-exponential fit to the data, showing the two individual exponential distributions and their sum. The black circle marks the value of number of pulses between breakdowns for which there is, according to the double-exponential fit, a 95% probability that a breakdown is a primary breakdown rather than a follow-up of the preceding one (i.e. the point at which the exponential distribution representing secondary breakdowns has 5% of the value of the sum). This value of number of pulses between breakdowns is marked as a vertical red line in the top plot [29].

value for β measured in CLIC structures is around 30. This effect can be caused by local geometric field enhancement, or by a local lowering of the workfunction due to impurities or atomic-scale features [31]. The β factor can be determined experimentally using:

$$\frac{d(\log_{10} I_F/E^{2.5})}{d(1/E)} = -\frac{2.84 \times 10^9 \phi^{1.5}}{\beta} \quad (1.5)$$

It should be noted that a linear scaling of the measured current does not affect the measured β . In RF accelerating structures, it is expected that the vast majority of the field-emitted current is emitted from the tips of the irises, where the local electric fields are the highest.

1.8 Thesis Outline

As detailed earlier in this chapter, vacuum breakdown is a complex process, involving many scales and phenomena, which is difficult to study. It is nevertheless worthwhile to do so due to its important role in the overall performance of devices incorporating a high electric field, in this case accelerating structures. This thesis focuses on two different stages of the evolution of a breakdown, both of which play roles in the overall breakdown rate of the device in question.

The stage considered in Ch. 2 is the onset stage in which thermal runaway of an emission site has begun, and rapid expansion of the breakdown plasma is under way. The breakdown consumes the most power in this stage. The coupling of RF power to the breakdown to feed this expansion depends on the exact geometry of the structure, and it is proposed that this is the mechanism that determines the ultimate field that a device can be operated at. A physical model for this coupling is presented, and both analytical and numerical calculation results are shown and compared with experimental data from high-power tests of different structures, showing good consistency in a diverse range of conditions.

Ch. 3, on the other hand, focuses on the nucleation stage. This is the stage before thermal runaway, and involves the formation of field emitters which could in turn evolve into a full breakdown. The statistics of the formation of these field emitters drives the statistics of breakdowns. Thus, any dependence of nucleation events on factors such as material properties or temperature will have a corresponding impact on the breakdown rate, making it a process worth understanding well. The work detailed in this chapter constitutes the first direct measurement of the dislocation dynamics behind this nucleation process. This was done by precisely measuring small fluctuations in the field-emitted current from RF structures and the DC spark system. The measurement results were compared with previous theoretical work, showing good consistency with predictions.

Chapter 2

Breakdown Criterion

2.1 Motivation

The accelerating gradient is a key performance parameter of RF accelerating structures [32]. High gradient allows compact high-energy linear accelerators to be built, which are of interest to the particle collider community, research accelerators such as X-ray Free-Electron Lasers (XFELs) [33] and inverse-Compton-scattering sources [34], as well as for medical and industrial applications. Vacuum breakdown is usually the phenomenon that limits the achievable gradient in state-of-the-art machines. In operational machines, a maximum practical gradient is usually specified below which breakdowns, which behave stochastically, are not expected to occur at a significant rate. The gradient limit is thus defined by the gradient at which the breakdown rate is acceptably low from the perspective of machine operation, rather than a well-defined boundary.

The maximum achievable gradient can depend on the operating frequency, material, design geometry, manufacturing process, and conditioning state. The choice of frequency and geometry is part of the design process, and a way of quantifying the maximum gradient for a given geometry is important for optimising the design. There are usually design goals other than gradient: power efficiency is a very important one, and

wakefield damping is relevant for CLIC [12] but this is not the case for every machine. Several quantities have been proposed and used as measures for breakdown limits over the years, a few of which are detailed in the following sections.

2.1.1 Kilpatrick's Criterion / Peak Electric Field

Kilpatrick defined [35] a breakdown criterion for RF cavities which is a function of the surface electric field and the maximum energy with which a free charged particle could strike the cavity wall. A simplified form of this criterion [36] is frequently used, in which the particle energy term was eliminated by making it a function of the electric field and frequency. Specifically, the particle energy value used is the energy gained by a hydrogen ion with zero initial velocity over a half RF cycle whose amplitude is equal to the peak surface electric field of the cavity. This simplified form is:

$$f = 1.64 \text{ MHz} \cdot (E \text{ [MV/m]})^2 \exp\left(\frac{-8.5}{E \text{ [MV/m]}}\right), \quad (2.1)$$

where f is the RF frequency and E is the peak surface electric field. This relation is plotted in Fig. 2.1. At higher frequencies, there is less time available for a particle to be accelerated before the direction of the field changes, limiting the maximum particle energy for a given field. Thus, it predicts that higher-frequency structures should be able to attain higher peak surface fields, varying roughly as the square root of frequency. Kilpatrick's criterion is based on data obtained in an era before clean vacuum systems were common, resulting in very conservative values by modern standards [32]. It can be seen in Fig. 2.1 that both the nominal and experimentally measured CLIC prototypes have greatly exceeded the predicted maximum surface-field value.

Studies of structures that have undergone high-power testing often show clustering of breakdown craters around local maxima of surface electric field [38]–[40], suggesting that it does indeed play a significant role, though the scaling may not necessarily be correct. Furthermore, field emission, which mainly occurs in locations of high electric field, has a well-understood role in the breakdown process. However, as can be de-

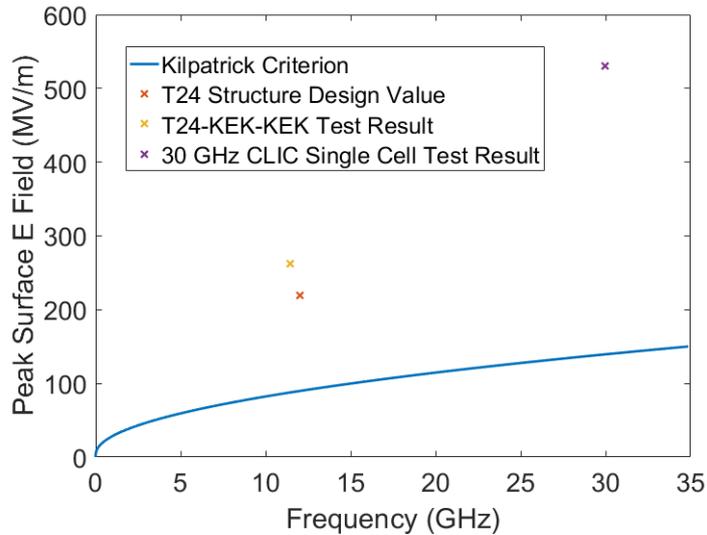


Figure 2.1: Peak surface electric field in MV/m vs. RF frequency in GHz. Blue curve: maximum surface electric field allowed by the formulation of Kilpatrick’s criterion in (2.1). Red cross: peak electric field in a CLIC T24 structure as designed. Yellow cross: peak electric field in the best-performing T24 structure tested. Purple cross: peak electric field in a 30 GHz single-cell test [37].

duced from the large disparity in maximum surface field reached in different designs of accelerating structures, the surface electric field and frequency cannot be the only factors determining the breakdown limit.

2.1.2 P/C

One alternative to Kilpatrick’s criterion is the quantity P/C , which was developed in the framework of the CLIC study based on test results from several structures developed and tested for linear colliders [41]. At the maximum achievable gradient, the peak value of the ratio of the incident RF power P to the circumference of the cell iris C was found to have a greater consistency between different designs than the peak surface electric field, see Fig. 2.2. In this figure, the square root of P/C was plotted for a more valid comparison with the surface electric field, since the surface field in a given structure is directly proportional to the square root of the incident power. For comparisons of structures of different operating frequencies, a frequency term f was also included in the quantity to take into account the tendency of higher frequency structures to reach higher gradients.

N	Name	f (GHz)	$\Delta\phi$ ($^\circ$)	v_g/c (%)	L (m)
1	DDS1 [42]	11.424	120	11.7-3.0	1.8
2	T53vg5R [42]	11.424	120	5.0-3.3	0.53
3	T53vg5MC [42]	11.424	120	3.3-1.6	0.53
4	H90vg3 [42]	11.424	150	3.1-1.9	0.9
5	H60vg3 [42]	11.424	150	3.0-1.2	0.6
6	H60vg3S18 [42], [43]	11.424	150	3.3-1.2	0.6
7	H60vg3S17 [42], [43]	11.424	150	3.6-1.0	0.6
8	H75vg4S18	11.424	150	4.0-1.0	0.75
9	H60vg4S17 [42], [43]	11.424	150	4.5-1.0	0.6
10	HDX11 [44]	11.424	60	5.1	0.05
11	CLIC X-band [45]	11.424	120	1.1	0.23
12	T18vg2.6 [46]	11.424	120	2.6-1.0	0.18
13	SW20a3.75 [43]	11.424	180	0	0.2
14	SW1a5.65t4.6 [47]	11.424	180	0	0.013
15	SW1a3.75t2.6 [47]	11.424	180	0	0.013
16	SW1a3.75t1.66 [47]	11.424	180	0	0.013
17	$2\pi/3$ [48]	29.985	120	4.7	0.1
18	$\pi/2$ [49]	29.985	90	7.4	0.1
19	HDS60 [50]	29.985	60	8.0-5.1	0.1
20	HDS60-Back [50]	29.985	60	5.1-8.0	0.1
21	PETS9mm [51]	29.985	120	39.8	0.4

Table 2.1: Structure parameters used for comparisons of breakdown quantities. From left to right: structure number used for identification in Fig. 2.2, name and reference, frequency f , RF phase advance per cell $\Delta\phi$, group velocity normalised to the speed of light v_g/c , and structure length L [42].

P/C is very easy to calculate but not very useful for designing the geometry as it is not a local quantity. This means that it has one value for a given cell of a structure, rather than a value for each point on the surface. Thus, it can at best specify which iris the most breakdowns are expected to occur at. It also does not provide guidance as to how to optimise the geometry beyond varying the iris circumference. Another disadvantage of P/C is that it does not apply to standing-wave structures since the power flow is very small in such structures, serving only to replenish energy lost due to wall resistance or to the beam.

It is interesting to note that in travelling-wave structures, the power flow and group velocity are closely related [32]:

$$P = \frac{v_g a Q}{\omega R} E_{acc}^2, \quad (2.2)$$

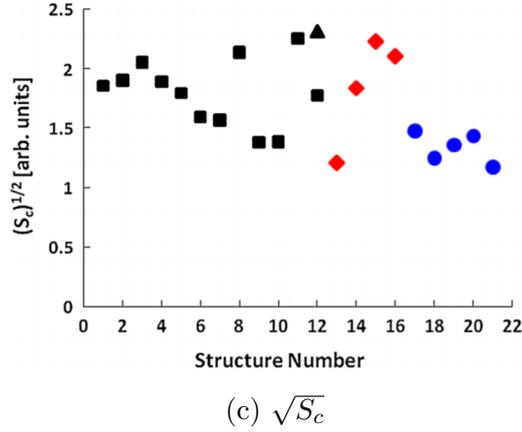
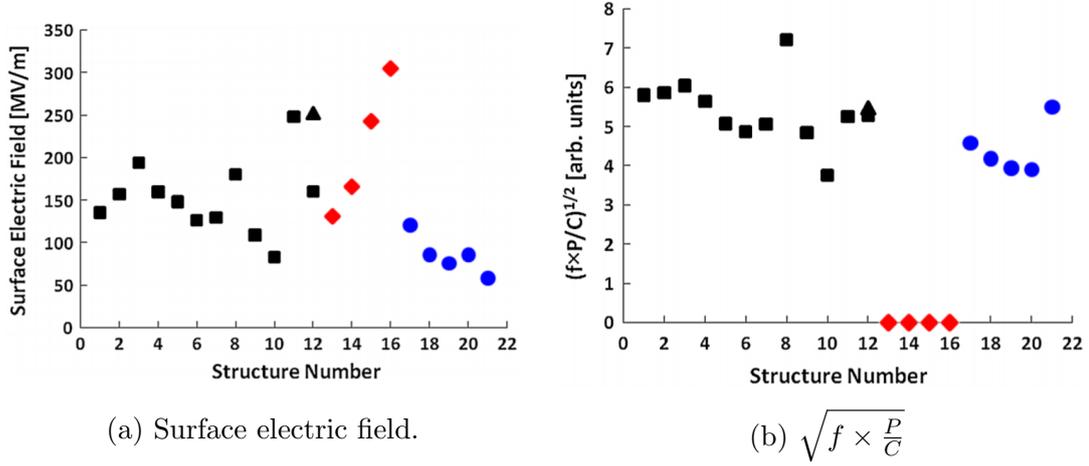


Figure 2.2: Maximum values of various quantities reached experimentally vs. structure geometry. Details of each structure are given in Table 2.1. Top left: surface electric field. Top right: $\sqrt{f \times \frac{P}{C}}$. Bottom: $\sqrt{S_c}$. Black markers represent X-band travelling-wave structures, red markers represent X-band standing-wave structures, and blue markers represent 30 GHz travelling-wave structures. Since $\sqrt{f \times \frac{P}{C}}$ does not apply to standing-wave structures, they are assigned zero values in the top-right plot. The results have been normalised to a pulse length of 200 ns and a breakdown rate of 10^{-6} breakdowns per pulse [42].

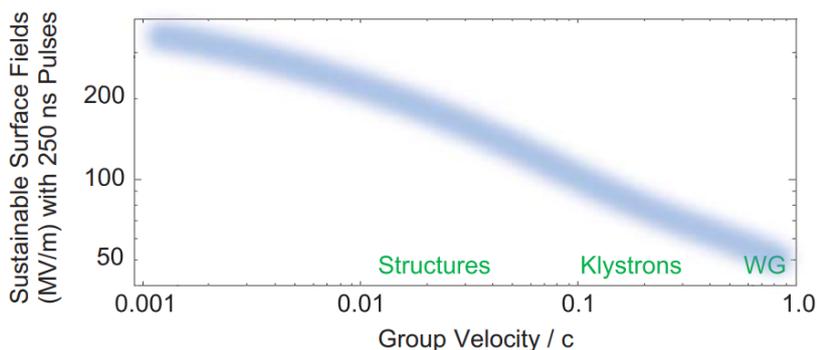


Figure 2.3: Maximum achievable surface electric field with 250 ns long RF pulses and an acceptable breakdown rate of about 5×10^{-7} breakdowns per pulse vs. group velocity [52].

where P is the incident RF power, v_g is the group velocity, E_{acc} is the accelerating gradient, Q is the quality factor, R is the shunt impedance in Ω , and a is the cell length. This, together with the idea of the breakdown limit being defined by power flow, implies a dependence of gradient on group velocity, with higher group velocities giving lower surface and accelerating fields. This has been noted by others in the context of linear accelerators [53] [52]. An example is shown in Fig. 2.3.

2.1.3 Modified Poynting Vector

The Modified Poynting Vector S_c [42] is an extension of the P/C quantity that preserves the emphasis on power flow, rather than surface electric field, as the quantity limiting the maximum achievable gradient. In contrast to P/C , it is a local quantity with a unique value at every location within the structure, which makes it more useful in the design process. It is defined as:

$$S_c = \Re(\mathbf{S}) + \frac{1}{6}\Im(\mathbf{S}), \quad (2.3)$$

where $\mathbf{S} = \mathbf{E} \times \mathbf{H}$ is the Poynting vector. Both \mathbf{S} and S_c have units of power density. The real part corresponds to sustained power flowing in and out of the cavity, while the imaginary part corresponds to reactive power flow, i.e. energy stored in the cavity circulating between the high electric field region near the iris and the high magnetic

field region near the cavity wall every RF cycle. As the power flow in standing-wave cavities has a very small real part and large imaginary part, the $\frac{1}{6}\Im(\mathbf{S})$ term in the definition of S_c prevents it from predicting unrealistically large gradients for standing-wave structures and travelling-wave structures with very low group velocities, as was the case with P/C . Physically, S_c represents the RF power density available for heating field-emission sites to allow them to evolve into full breakdowns. A peak value of 5 MW/mm^2 has been recommended as a maximum allowable value in accelerating structure design [42]. Very well performing structures have managed to slightly exceed this value [18] [54] [55].

S_c works well as a breakdown criterion for travelling-wave structures operating in the TM_{010} mode [42], encompassing all the CLIC accelerating structure prototypes that have been tested, and gives consistent values for the maximum gradients of 12 GHz and 30 GHz CLIC prototypes without an explicit frequency-dependence term, further strengthening the idea that power flow is a fundamental limit for breakdown.

An inconsistency with S_c was found when it was applied to the prototype CLIC Crab Cavity [56] tested in X-Box 2, as the breakdown crater locations identified in *post-mortem* analysis were much more consistent with the spatial distribution of electric field in each cell, rather than the distribution of S_c , as shown in Fig. 2.4 [38]. The fact that this prototype was a deflecting structure operating in the TM_{110} mode made the E and S_c distributions easy to distinguish. In contrast, in accelerating structures operating in the TM_{010} mode, the locations of peak E and S_c coincide, and tests of such structures cannot provide this information.

Despite the incorrect breakdown locations, the breakdown-limited maximum field predicted by S_c appears to be consistent with other accelerating structures [57]. The Crab Cavity was tested up to an input power of 52 MW with a 120 ns pulse length for a short period with a breakdown rate of 10^{-4} breakdowns per pulse [18], corresponding to a peak S_c value of 6.66 MW/mm^2 . For comparison, the peak S_c value reached in the T24PSIN1 structure was 5.13 MW/mm^2 at an input power of 55 MW with a 200 ns

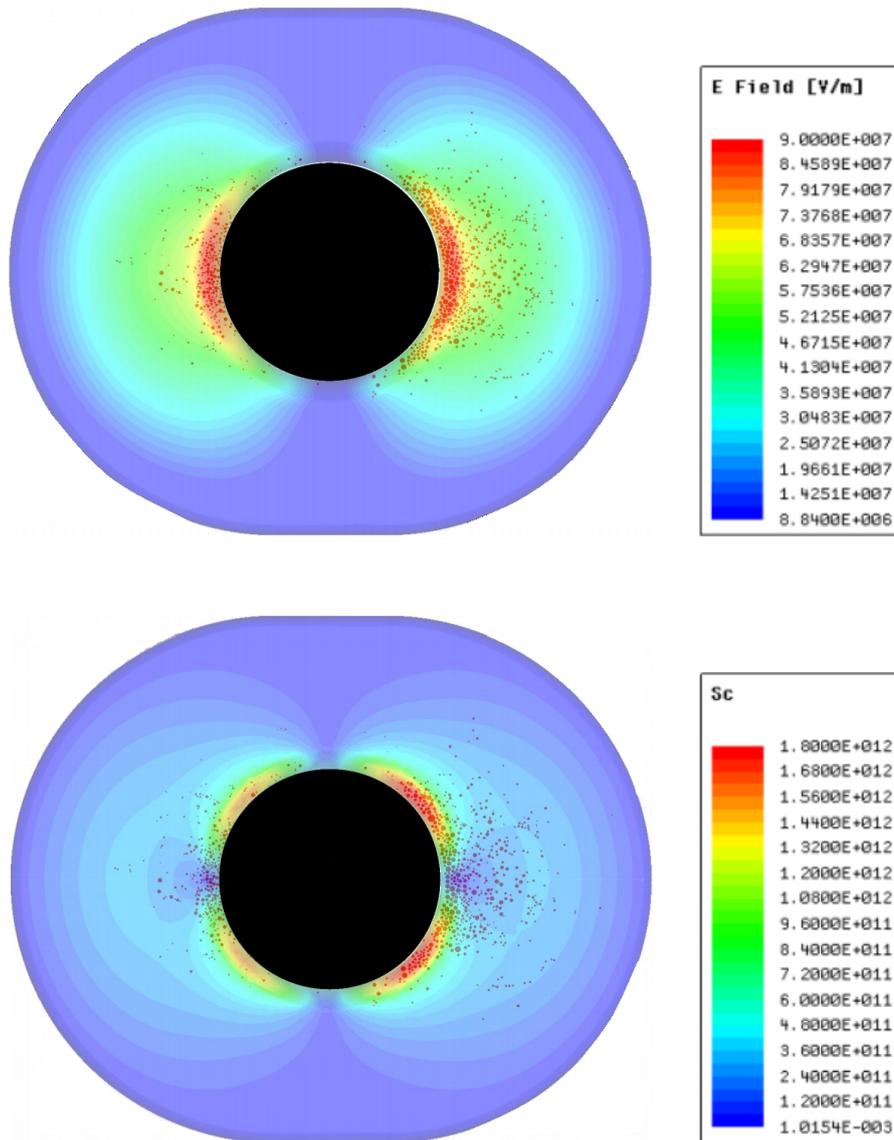


Figure 2.4: Breakdown locations, represented by black dots, vs. position transverse to the beam in the second cell of the CLIC Crab Cavity, obtained from *post-mortem* analysis [38]. Due to the close proximity of the input waveguide coupled to the first cell, there is a left-right asymmetry present in the field pattern and breakdown locations. Top: overlaid with surface electric field vs. position transverse to the beam. Bottom: Overlaid with S_c vs. transverse position.

pulse length and a breakdown rate of 7.5×10^{-5} breakdowns per pulse [58]. Thus, the peak S_c achieved in the two structures is reasonably consistent, taking into account the short, run limiting the accuracy of measuring the breakdown rate of the Crab Cavity. However, the disagreement between the distribution of the S_c and the breakdown locations is an indication that the phenomenon is not yet fully understood.

2.1.4 An Improved Quantity

Sections 2.1.2 and 2.1.3 show that there is much evidence pointing to power flow playing an important role in defining the breakdown threshold. However, the inconsistency in the Crab Cavity result shows that in reality, the contributions of the electric and magnetic fields are not interchangeable in the way they are in (2.3). An improved quantity should reflect this. It should also work for both standing-wave and travelling-wave structures at all frequencies.

Furthermore, all the quantities related to power flow that have been discussed so far are not consistent with DC experiments as they predict that breakdowns should never occur at zero frequency since there is no flow of power in this case. On the other hand, Kilpatrick's criterion predicts breakdowns should occur at zero electric field in DC experiments. The Large Electrode System (LES) routinely achieves peak DC surface electric fields in excess of 80 MV/m [59]. It would be desirable for an improved quantity to be able to resolve this discrepancy.

2.2 Proposed Model

As an attempt to meet the requirements set out in Sec. 2.1.4, a new quantity is proposed that considers the surface electric field when loaded by the initial stages of the breakdown. The breakdown-loaded electric field is a function of the unloaded electric field, which should help make it more consistent with the *post-mortem* results of the CLIC Crab Cavity, yet still depends on the supply of power to the breakdown site, which has been established to be an important factor in defining the breakdown

limit. It should be noted that the breakdown-loaded field distribution is not in general the same as the beam-loaded field distribution, as a breakdown typically only involves one cell.

It is thought that for a breakdown to fully develop after nucleation, emission of charged particles from the breakdown site needs to be sustained for a critical period of time. This is supported by atomic-level simulations [60], [61] of the initial stages of vacuum breakdowns, which show that the large density of emitted current causes a runaway process as the emitter temperature increases due to Joule heating and the Nottingham effect [62], causing a further increase in emitted current. The simulations show that about 200 ps of heating is needed for the runaway process to begin. The emitted particles are expected to be mostly electrons due to the short timescales of the initial stages of breakdown [63]. Copper ions and neutral atoms are also emitted and play a crucial role in the later stages of breakdown (see Sec. 1.6), but they are not accelerated over large distances by RF fields, meaning that they do not play a significant role in power absorption. The emitted electrons are accelerated by the electric field, absorbing energy from the field. Before the emitter heats up enough for thermionic emission to take over from field emission as the main electron emission mechanism, a large influx of power (up to tens of megawatts) is required to keep the electric field strong enough to sustain the emission of electrons [64]. Only the electric field is directly relevant here as a magnetic field cannot change the energy of a particle, and thus does not cause any power to flow from the field to the electrons, as the magnetic force is always perpendicular to the velocity. This can be deduced from the equation for the Lorentz force [32]:

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \quad (2.4)$$

where q is the particle's charge, and \mathbf{v} is the particle's velocity. The proposed model includes the assumption that a larger surface electric field results in more emitted current and that a larger current produces a larger opposing electric field, causing the net surface field to drop. These opposing effects eventually result in a quasi-equilibrium

state. It is proposed that the surface electric field at this equilibrium point is the quantity that determines the breakdown rate in a high-electric-field RF device. This equilibrium value of electric field is intended to be used in a manner analogous to S_c for predicting breakdown rates and the design of high-field devices such as accelerating structures.

The calculation of the loading of the electric field during the onset of a breakdown is performed by simulating the complete structure geometry with a small monopole antenna placed at the location of the hypothetical breakdown site. The antenna represents the stream of charged particles emitted from the developing breakdown site. This is a very simplified model of the situation in a real breakdown, in which the particles continue absorbing energy from the field as they travel through the resonant cell along complicated trajectories that require a Particle-in-Cell (PIC) simulation to accurately calculate [63]. However, this simplified model still provides valuable insight which is discussed below.

The coupling between the antenna and the RF input source can be obtained from this simulation, giving an impedance value which relates the voltage across the antenna to the current flowing through it. The calculation is performed at the design frequency of the structure, i.e. a sinusoidal current oscillating at this frequency is assumed to flow through the antenna. The voltage across the antenna is equal to the integral of the electric field along the antenna:

$$Z_{bd} = \frac{V}{I} = \frac{\int_{antenna} E_n dl}{I}, \quad (2.5)$$

where Z_{bd} is the aforementioned impedance value, V is the voltage across the antenna, I is the current flowing in the antenna, and E_n is the component of the electric field normal to the metal surface and along the antenna. For an antenna length l_{ant} much smaller than the RF wavelength, the electric field can be assumed to be almost constant along the length of the antenna and the integral can be reduced to a product of the electric field and the antenna length:

$$Z_{bd} \approx \frac{E_n l_{ant}}{I} \quad (2.6)$$

Only the real part of the resulting impedance Z_{bd} will be used in the following steps since only this corresponds to power delivered to the particles, whereas the imaginary part corresponds to power delivered but then removed in the same cycle resulting in no net power flow. The real part will be referred to from now on as breakdown resistance R_{bd} . By linear superposition of the fields set up by the incident power at the input port and those set up by the current in the antenna, the local electric field at the breakdown site as a function of emitted current can be determined. Since Maxwell's equations are linear, this will always result in a linear function requiring two parameters to fully describe the local surface electric field as a function of breakdown loading:

$$E_{surf} = E_0 - R_{bd}I, \quad (2.7)$$

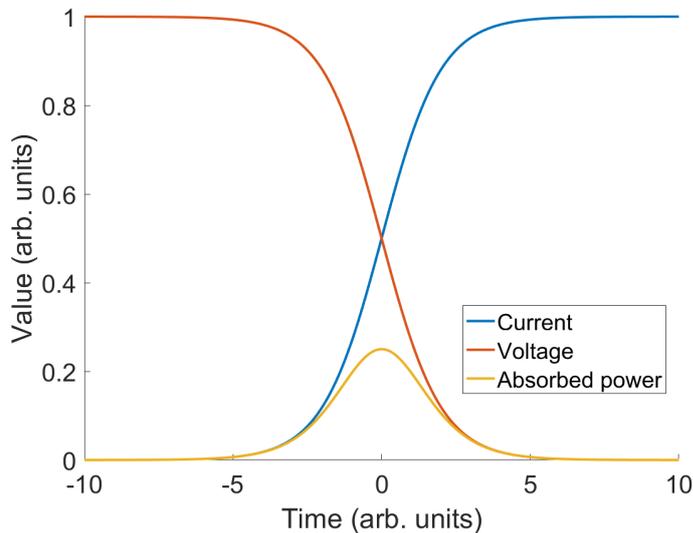


Figure 2.5: Current emitted (red) by, voltage across (blue), and power absorbed (yellow) by an idealised breakdown vs. time.

where E_{surf} is the beam-loaded surface electric field and E_0 is the unloaded electric field set up by the driver. The other component of the model is an emission function that gives the current emitted by the developing breakdown for a given local electric

field. This function should model the emitter during the breakdown onset stage, during which the power flow is the greatest.

An illustration of the time evolution of the current and voltages in an idealised breakdown is shown in Fig. 2.5, which graphically explains why this stage is the most relevant for a limiting quantity concerned with power flow. Three stages can be identified: an initial stage of low current emission at high electric fields, by a mechanism such as field emission; the onset of the breakdown in which the thermal runaway process begins in earnest, leading to a large increase in emitted current and collapse of the electric field; and a stage in which the plasma has fully developed, resulting in very large currents at low voltage [65]. The power absorbed by the breakdown, being the product of the current and voltage, is greatest during the onset phase.

Based on this argument, the emission function cannot be the Fowler-Nordheim equation [30] (which gives the density of the field-emitted current from a surface as a function of the applied electric field) since it is expected that processes other than field emission begin to dominate the current flow at the time of breakdown onset. In this work, it was assumed that the emission function should be the same for all structures made of the same material, since it is probably governed by the type of emitter that is likely to form on the surface. The current must be a monotonically increasing function of the surface electric field to produce the expected behaviour of higher E_0 always resulting in a higher breakdown rate. The functional form and parameters have been manually adjusted to give the best consistency with experimental results, but remain the same for all calculations. Since there is no direct way of measuring this function with the experimental setups available, other than matching predicted breakdown locations and field levels with experimental results, it may not be unique. It is likely that the emission function changes with conditioning, since both properties of the material and the achievable gradient at a given breakdown rate are known to change with conditioning. It may even vary slightly with location within a given structure, since breakdowns do not generally all occur at a single point, but have some spread in location. The func-

tion discussed here represents an ultimate limit for well-conditioned structures. Using a power law has given good results:

$$I = kE_{surf}^n, \quad (2.8)$$

where I is the current emitted by the breakdown site, and k and n are fitting parameters. The use of such a function, as opposed to say a field emission-like characteristic, is justified if the emitted current is limited by space charge. This is the case for very large current densities typical of breakdowns [66]. The Child-Langmuir law [67], given in (2.9), which describes the current density J in a uniform vacuum diode of length d , is equivalent to (2.8) if $n = 3/2$ and d is fixed. An example emission function with $n = 3/2$ is plotted in Fig. 2.6.

$$J \propto \frac{V^{3/2}}{d^2} \quad (2.9)$$

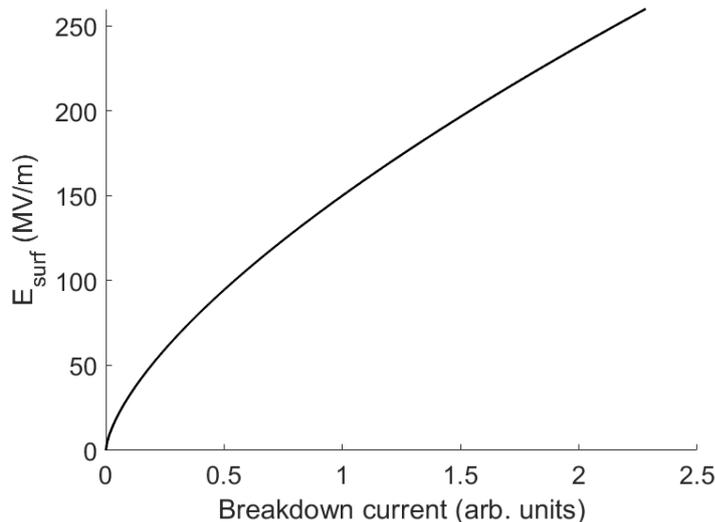


Figure 2.6: Surface electric field in MV/m vs. current emitted by a breakdown as given by (2.8).

The intersection of the loading and emission functions defines the equilibrium values of surface electric field and emitted current that this system will settle at, denoted E^* and I^* respectively. This is effectively a solution of the simultaneous equations (2.7) and

(2.8), i.e. $I^* = k(E^* - R_{bd}I^*)^n$. The unloaded electric field is proportional to the square root of the incident power, which means that the loaded electric field also depends on the incident power. The maximum power at which the structure can be safely operated is thus the value at which the peak E^* in the structure just reaches some threshold value. Some examples are shown in Fig. 2.7, referred to as a breakdown loading plot. The blue line represents a breakdown site with an E_0 of 200 MV/m and an arbitrary value of $R_{bd} = R_0$. The red line represents a breakdown site with the same E_0 but a smaller R_{bd} , resulting in a larger E^* . The yellow line represents a breakdown site with the same R_{bd} as the first case, but with a reduced E_0 of 150 MV/m. This results in a lower E^* . This shows how in general, with all other variables kept constant, either a larger E_0 or a lower R_{bd} will result in a higher breakdown rate. The threshold value of E^* should be independent of the device being analysed, and a significant portion of this chapter will be devoted to verifying this by various methods.

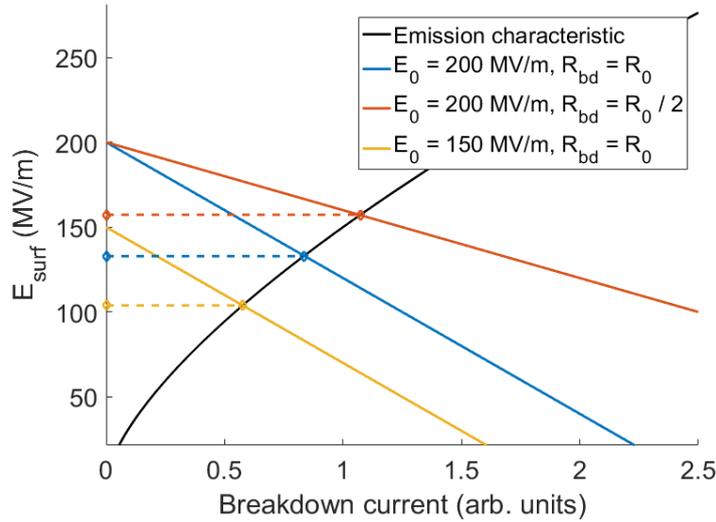


Figure 2.7: Calculation of E^* on a plot of surface electric field in MV/m vs. current emitted by a breakdown. Black curve: emission function as given by (2.8). Solid blue line: a breakdown site with an unloaded field of 200 MV/m and a breakdown resistance R_0 . Solid red line: a breakdown site with an unloaded field of 200 MV/m and a breakdown resistance $R_0/2$. Solid yellow line: a breakdown site with an unloaded field of 150 MV/m and a breakdown resistance R_0 . Dashed lines: the resulting E^* for the three cases.

The calculation must be repeated for every potential breakdown site under consideration, and the location which yields the greatest value of E^* is likely to be the site of the greatest concentration of breakdowns.

2.2.1 Analytical Calculation of R_{bd}

A key part of calculating E^* for a given geometry is calculating the value of R_{bd} for each potential breakdown location. The unloaded surface field E_0 with a given RF power input to a given geometry can be easily calculated using a number of readily available finite-element simulation tools, and is routinely done during the design process of many high-field devices. R_{bd} , on the other hand, is not a universally known quantity in the field, and calculating it is not as straightforward. One way to do this is via an analytical circuit model.

Circuit models are frequently used to model accelerating cavities as they are much simpler than a finite-element simulation of the physical geometry [32]. Using a circuit model to describe breakdown impedance allows the dependence of R_{bd} on the properties of the cells to be shown explicitly without having to resort to a pseudo-experimental approach with numerical calculations. The results are also general and apply to all possible geometries within the limits of the model's validity. A similar style of model is used here, though with the intention of calculating the coupling of power to a breakdown in a structure. Here, the breakdown is modelled as an extra port in addition to the usual input and output ports, with a nonlinear impedance connected to it that obeys the emission function discussed in Sec. 2.2. The circuit models discussed below only consider one of the resonant modes of each cell, but real cavities have many modes which may get excited by a physically small breakdown event. However, since it is common practice to only excite one mode of an RF structure and ensure that the remaining modes are adequately damped, it is believed that considering only the driven mode should be sufficient to model the broad characteristics of the structure in question.

Single-Frequency Model

An analytical model valid only at the resonant frequency of an accelerating structure will be considered first for simplicity. R_{bd} calculated in this way is sufficient for calculating a steady-state value of E^* . The start of a breakdown is a transient process involving the depletion of stored energy from the cell, which is not modelled here. The use of this model leads to some discrepancies from experimental observations, detailed in Sec. 2.3.1. A more complex model that is valid over a range of frequencies, important for modelling transient behaviour relevant to structures with very low group velocities, will be discussed afterwards.

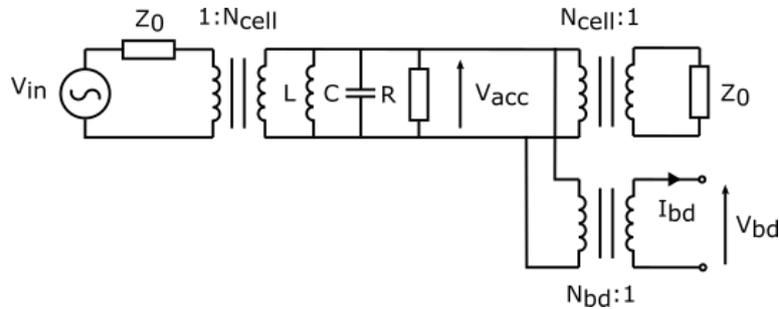


Figure 2.8: Equivalent circuit of a single cell in a travelling-wave configuration.

The simplest possible travelling-wave structure will be considered here, consisting of one cell coupled to a power source and an impedance-matched load. A schematic of the circuit model of such a structure is shown in Fig. 2.8. The cell is represented by a parallel inductor-capacitor-resistor (LCR) circuit, and the couplings to the power source and load are represented by ideal transformers of ratio N_{cell} . Both the power source (represented by the voltage source) and the load are matched to a characteristic impedance of Z_0 . A third coupling to the cell is included in this circuit, represented by transformer of voltage ratio N_{bd} , representing the breakdown. The breakdown current and voltage across the antenna are represented by I_{bd} and V_{bd} respectively. If I_{bd} is forced to be zero for all values of V_{bd} , no power flows to the breakdown and the cell behaves as it would in a test without loading from a particle beam or breakdown. R_{bd} in this model is the equivalent source resistance as measured from the breakdown port. The component values are scaled such that the voltage across the LCR circuit is the

accelerating voltage V_{acc} of the cell. The power P flowing through the cell, assuming zero breakdown current, can be expressed in terms of circuit parameters as:

$$P = \frac{V_{in}^2}{4Z_0} \quad (2.10)$$

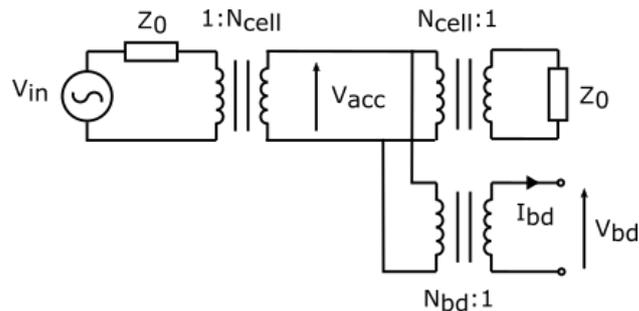


Figure 2.9: Simplified equivalent circuit of a single cell in a travelling-wave configuration at the resonant frequency.

The circuit model can be simplified further by assuming that the cell is strongly over-coupled, meaning that the power lost through the output coupling to the load is much larger than the resistive power losses in the cell. This is often the case with travelling-wave structures [32]. This means that the power loss in the resistance is negligible and one can assume that $R = \infty$ and thus remove the resistor from the circuit entirely. The drive frequency will also be assumed to be equal to the resonant frequency of the cell, $f_r = 1/2\pi\sqrt{LC}$ [68]. At this frequency, the currents flowing through the inductor and capacitor are equal and opposite. Thus, the impedance of the two components in parallel becomes infinite, and they too can be removed from the schematic. This simplified schematic is shown in Fig. 2.9.

It can be easily seen that in this circuit, $V_{acc} = N_{bd}V_{bd}$. The value of N_{bd} depends on the location of the breakdown within the cell, which determines the electric field along an antenna for a given accelerating voltage. The exact relation is:

$$N_{bd} = \frac{V_{acc}}{V_{bd}} = \frac{a}{l_{ant}} \frac{E_{acc}}{E_0}, \quad (2.11)$$

where a is the length of the cell, l_{ant} is the length of the antenna, E_0 is the surface electric field at the location of the antenna and E_{acc} is the accelerating gradient. With some circuit manipulations, the output impedance of the circuit as seen by the breakdown port can also be expressed in terms of circuit parameters:

$$R_{bd} = \frac{Z_0 N_{cell}^2}{2N_{bd}^2} \quad (2.12)$$

With no breakdown current, the accelerating voltage can be expressed in terms of the input voltage as $V_{acc} = N_{cell}V_{in}/2$. With this, (2.10), and (2.2), the relationship between input power and accelerating voltage can be derived:

$$\frac{V_{acc}^2}{P} = \frac{a\omega}{v_g} \frac{R}{Q} = N_{cell}^2 Z_0, \quad (2.13)$$

where P is the incident power from the source, ω is the operating frequency, v_g is the group velocity, R is the shunt impedance, and Q is the quality factor of the structure. R_{bd} can thus be expressed entirely in terms of quantities obtainable from an eigenmode solution of the electromagnetic wave equation in the cell geometry, i.e.:

$$\left(c^2 \nabla^2 - \frac{\partial^2}{\partial t^2} \right) \begin{pmatrix} \mathbf{E} \\ \mathbf{B} \end{pmatrix} = 0. \quad (2.14)$$

The relationship between R_{bd} and the eigenmode quantities is:

$$R_{bd} = \frac{\omega l_{ant}^2}{2v_g a} \frac{R}{Q} \left(\frac{E_0}{E_{acc}} \right)^2. \quad (2.15)$$

The dependence of R_{bd} on R/Q , v_g and E_0/E_{acc} in finite-element simulations of travelling-wave structures of similar geometries to CLIC structures will be reviewed in Sec. 2.3.3. This circuit model does not take into account stored energy in any way, as evidenced by the lack of energy-storing circuit components such as inductors or capacitors.

Broadband Model

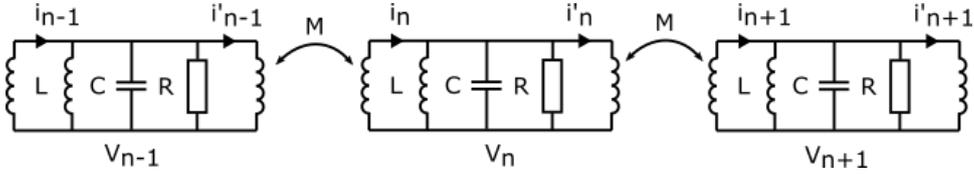


Figure 2.10: Equivalent circuit of an infinite array of identical coupled resonant cells.

The circuit model discussed above is only valid at a single frequency, for travelling-wave structures with a high group velocity and thus a negligible amount of stored energy. As will be discussed in Sec. 2.4.2, a model that does take stored energy into account would need to consider the full bandwidth of the structure. The fact that the structure has a finite bandwidth does not emerge from the model shown in Fig. 2.8 due to the use of ideal transformers for coupling, which have the same behaviour at any frequency. An appropriate circuit model should thus replace the ideal transformers with a frequency-dependent coupling, and may not omit reactive elements as was done in Fig. 2.9.

The array is assumed to have a travelling wave flowing through it, with a phase advance ka from cell to cell. Thus, the n^{th} cell is assumed to have a voltage v_n across it and current i_n defined as:

$$\begin{aligned} v_n &= v_0 e^{jkan} \\ i_n &= i_0 e^{jkan} \end{aligned} \quad (2.16)$$

In the above and subsequent equations, $j = \sqrt{-1}$ denotes the imaginary unit to avoid confusion with the current i .

Expressing the voltage across the $(n+1)^{\text{th}}$ cell as a function of the current i'_n in the n^{th} cell, and applying Kirchoff's Current Law [69] to the n^{th} cell, one can obtain an expression for the voltage across the $n+1^{\text{th}}$ cell:

$$v_{n+1} = j\omega M i'_n = j\omega M \left(i_n - \frac{v_n}{j\omega L} - j\omega C v_n - \frac{v_n}{R} \right) \quad (2.17)$$

One can also define the characteristic impedance Z_0 of this transmission line as:

$$Z_0 = \frac{v_n}{i_n} = -j\omega M e^{jka} \quad (2.18)$$

The coupling of multiple resonant cells results in a structure with a passband rather than a single resonant frequency as with an uncoupled cell. Equations (2.17) and (2.18) can be combined to obtain a dispersion relation for the wave flowing through the transmission line, in terms of circuit parameters:

$$v_n e^{jka} = j\omega M \left(\frac{v_n}{-j\omega M e^{jka}} - \frac{v_n}{j\omega L} - j\omega C v_n - \frac{v_n}{R} \right) \quad (2.19)$$

$$e^{jka} + e^{-jka} = 2 \cos(ka) = -j\omega M \left(\frac{1}{j\omega L} + j\omega C + \frac{1}{R} \right) \quad (2.20)$$

The following standard definitions of resonance frequency ω_r and quality factor Q for a parallel LCR circuit are used:

$$\omega_r = \frac{1}{\sqrt{LC}} \quad (2.21)$$

$$Q = R \sqrt{\frac{L}{C}} \quad (2.22)$$

By substituting 0 and π for the phase advance ka into the dispersion relation (2.20) and defining the resulting values of ω as ω_0 and ω_π respectively, an expression for M may also be obtained:

$$M = \frac{4\omega_0}{\omega_\pi^2 - \omega_0^2} \cdot \frac{R}{Q} \quad (2.23)$$

(2.20) through (2.23) can be used to obtain a generalised dispersion relation for an infinite array of coupled cells, expressed without any circuit values:

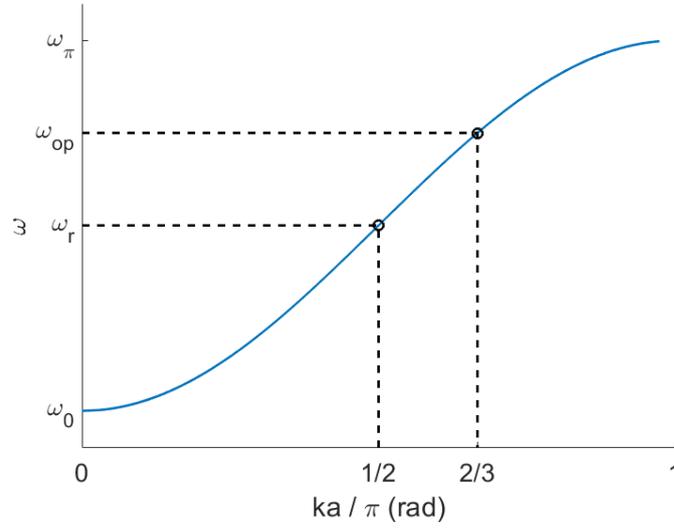


Figure 2.11: Angular RF frequency vs. phase advance per cell in radians for the circuit model shown in Fig. 2.10. ω_{op} is the operating frequency for a structure with a $2\pi/3$ phase advance.

$$\omega^2 = \omega_r^2 - \frac{\omega_\pi^2 - \omega_0^2}{2} \cdot \cos(ka) \quad (2.24)$$

This curve is plotted in Fig. 2.11. This resembles the dispersion curve of a typical travelling-wave structure [32], confirming that its behaviour has been modelled correctly. Performing a differentiation with respect to k , one can also obtain an expression for the group velocity:

$$v_g = \frac{d\omega}{dk} = \frac{\omega_\pi^2 - \omega_0^2}{4\omega} \cdot a \sin(ka) \quad (2.25)$$

Using (2.23) and (2.24), one can eliminate the parameters M and ka from the expression for characteristic impedance (2.18), yielding:

$$\Re(Z_0(\omega)) = \frac{R}{Q} \frac{4\omega\omega_r}{\omega_\pi^2 - \omega_0^2} \left(1 - 2 \frac{\omega_r^2 - \omega^2}{\omega_\pi^2 - \omega_0^2} \right)^{\frac{1}{2}} \quad (2.26)$$

Substituting $\omega = \omega_r$ into the above results in an expression that confirms the consistency of the broadband model with (2.13) which was derived from the single-frequency model:

$$\Re(Z_0(\omega_r)) = \frac{a\omega_r R}{v_g Q} \quad (2.27)$$

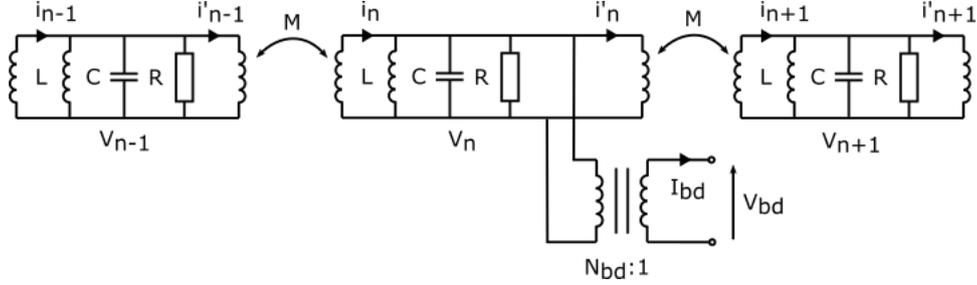


Figure 2.12: Equivalent circuit of an infinite array of identical coupled resonant cells, with a port in one cell representing the coupling to a breakdown.

To obtain a value for the equivalent source resistance for a breakdown in this model, one of the cells can be assumed to have an additional coupling via an ideal transformer, similar to that in the single-frequency model. This is presented in Fig. 2.12. The expression for characteristic impedance derived in (2.26) allows the equivalent source resistance as measured from the breakdown port to be easily calculated as follows, in a manner analogous to that of the single-frequency model in (2.12):

$$R_{eq,an}(\omega) = \frac{\Re(Z_0(\omega))}{2N_{bd}^2} = \frac{R}{Q} \frac{2\omega\omega_r}{\omega_\pi^2 - \omega_0^2} \left(1 - 2 \frac{\omega_r^2 - \omega^2}{\omega_\pi^2 - \omega_0^2}\right)^{\frac{1}{2}} \left(\frac{l_{ant}}{a}\right)^2 \left(\frac{E_0}{E_{acc}}\right)^2 \quad (2.28)$$

Discussion of Antenna Lengths

The analytical model exhibits a quadratic dependence of R_{bd} on the antenna length used, i.e.: $R_{bd} \propto l_{ant}^2$. It resembles the behaviour of the radiation resistance of a small Hertzian dipole of length $\delta l \ll \lambda$ [70]:

$$R_{rad} = \frac{2\pi}{3} \zeta_0 \left(\frac{\delta l}{\lambda}\right)^2, \quad (2.29)$$

where R_{rad} is the radiation resistance, and ζ_0 is the impedance of free space. The antenna length is the only parameter in the expressions derived for R_{bd} not already

defined by the geometry of the structure being studied. Thus, it is not immediately clear what value of antenna length should be chosen. It may not be practical in general to use a particular antenna length in all cases, due to limitations imposed by a particular geometry. The issue of choosing a value for antenna length can be circumvented by introducing a quantity R_0 that is invariant with l_{ant} , such that:

$$R_{bd} = R_0 l_{ant}^2 \quad (2.30)$$

As per the definition of R_{rad} in (2.6),

$$\begin{aligned} \Re(V_{ant}) &= I_{ant} R_{bd} \\ \Re(V_{ant}) &= I_{ant} R_0 l_{ant}^2 \\ \Re(E_n) &= \frac{\Re(V_{ant})}{l_{ant}} = I_{ant} l_{ant} R_0 \end{aligned} \quad (2.31)$$

Thus, with a breakdown loading plot where the axes are electric field and the product $I_{ant} l_{ant}$, the plot and results become independent of antenna length. In order to ensure that the emission function is also independent of antenna length, it can be defined as follows:

$$I_{ant} l_{ant} = k_0 E_{surf}^n, \quad (2.32)$$

where $k = k_0 l_{ant}$, with k being the proportionality constant relating emitted current and surface electric field in (2.8). This property will be used later in this chapter to compare simulation results using different antenna lengths.

The particles emitted during a breakdown will continue to absorb energy from the local electric field as long until they collide with the wall of the structure or escape. Thus, it should follow that the amount of energy absorbed depends on the length of the particle trajectory. The fact that the quantity $I_{ant} l_{ant}$ depends on antenna length reflects this behaviour, and it seems reasonable to choose l_{ant} to be a value representative of the

mean distance travelled by the emitted particles. It is difficult to determine the exact physical length of a breakdown during the onset phase in every possible geometry and breakdown location without a detailed Particle-in-Cell (PIC) simulation but it is most reasonable to assume it scales approximately linearly with the size of the cell, and therefore the free-space wavelength, i.e. $l_{ant} \propto \lambda_{fs}$. Thus, it may be reasonable to redefine R_0 with a wavelength term:

$$R_{bd} = R_0 \left(\frac{l_{ant}}{\lambda_{fs}} \right)^2, \quad (2.33)$$

where R_0 is purely a function of the surrounding geometry. This equation resembles (2.29), and it is expected that $R_0 = \frac{8\pi}{3}\zeta_0$ for an antenna placed on an infinite flat conductive plane, with ζ_0 being the impedance of free space. Since such a conductive plane is effectively a plane of symmetry, this should be equivalent to a Hertzian dipole of length $2l_{ant}$ in free space. If such a dependence on wavelength is adopted, (2.31) should thus be modified to obtain:

$$\Re(E_n) = I_{ant} l_{ant} \frac{R_0}{\lambda_{fs}^2}, \quad (2.34)$$

which gives a frequency dependence to E^* . In other words, if a given geometry is scaled in size by a factor a and the RF frequency scaled by $1/a$ to match, whilst keeping E_0 constant, the value of E^* may change. This property will be discussed further in Sec. 2.3.2 and used in Sec. 2.3.6. This behaviour is accessible to validation by experiment, the results of which would provide important insight into the process of breakdown onset.

2.3 Numerical Calculations and Experimental Results

A series of numerical calculations of E^* in specific structure geometries was performed to further study its behaviour and compare its behaviour to experimental results. The cases that will be discussed in this section are:

- A set of simplified, theoretical CLIC-like X-band accelerating structures designed to investigate the dependence of E^* on aperture size (Sec. 2.3.1).
- The CLIC Crab Cavity, an X-band deflecting structure tested at the XBoxes (Sec. 2.3.4).
- A T24 structure, one of a series of prototype CLIC X-band accelerating structures without wakefield damping tested at the XBoxes (Sec. 2.3.5).
- The BTW structure, a high-gradient S-band structure tested at CERN intended for use in a medical proton linac (Sec. 2.3.6).

2.3.1 Simulated CLIC-Like Accelerating Structures

As a first step towards verifying the applicability of the E^* quantity, one needs to check the results it produces for consistency with S_c for travelling-wave structures operating in the TM_{010} mode at X-band, a regime in which S_c shows very good agreement with experimental results [42]. This avoids the issue of going back to the original data and is based on the previously determined validity of predictions of S_c .

A representative CLIC accelerating structure-like geometry was simulated to compare E^* and S_c values at various locations inside it. A later step will be to apply the model to cases where S_c did not fit experimental results.

Cell Design

In order to represent an accelerating structure for the CLIC main linac, an operating frequency of 12 GHz was chosen for simplicity, but still very close to CLIC's 11.994 GHz.

The phase advance per cell and phase velocity were also chosen to match those of CLIC structures. A list of parameters is given in Table 2.2, and the shape of one of the cells is given in Fig. 2.15. Since CLIC structures are typically tapered in aperture size [42], a range of aperture radii from 2.5 mm to 6 mm was simulated. The cell radius was adjusted with the help of CST Microwave Studio's [71] eigenmode solver to ensure that a phase advance of 120° per cell at a frequency of 12 GHz was maintained in each case. The dependence of a variety of properties obtained from the eigenmode solutions as a function of aperture size is shown in Fig. 2.13, which, together with the relationship in (2.2), shows that as the aperture size increases, the amount of power needed for a given accelerating field increases. Examples of field distributions are shown in Fig. 2.14. The cells were radially symmetric and had no wakefield damping features.

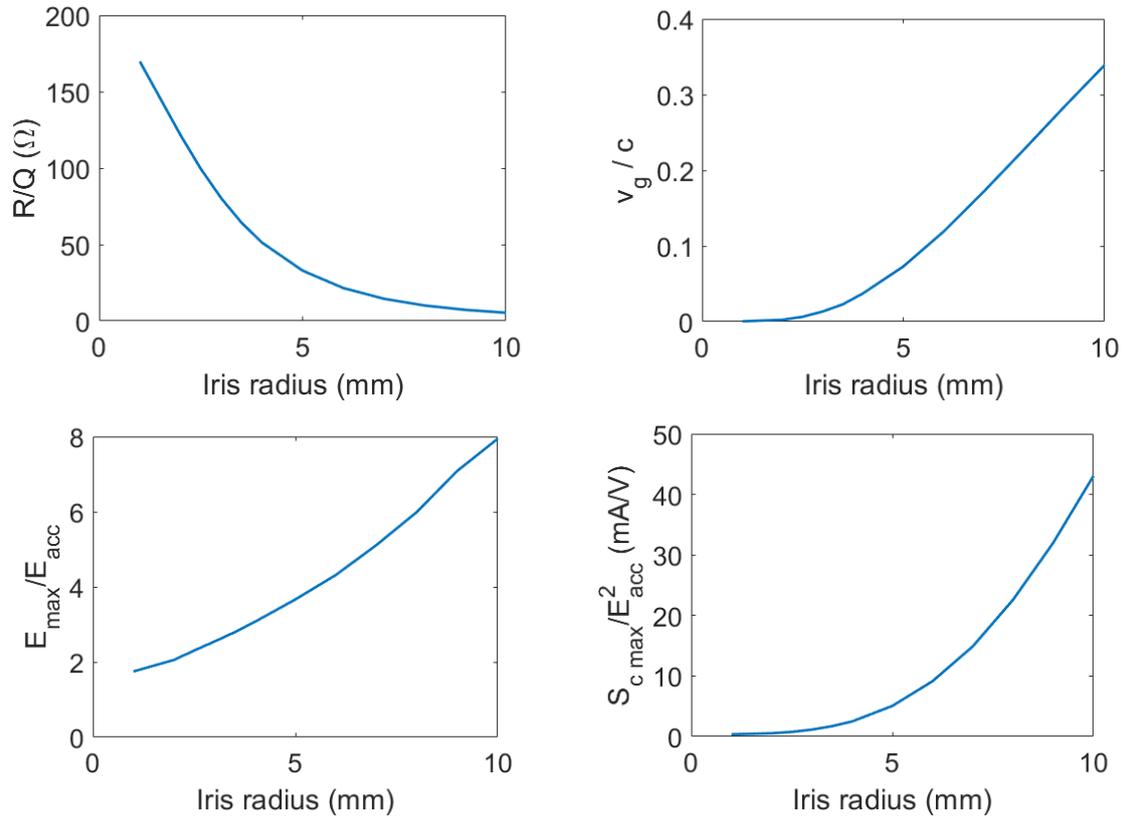


Figure 2.13: Properties of cells vs. iris radius. Top left: R/Q . Top right: Group velocity as a fraction of c . Bottom left: Maximum surface electric field at unit accelerating gradient. Bottom right: Maximum S_c at unit accelerating gradient.

A total of fifteen antennas was placed at various positions along the cell surface (Fig. 2.15) to represent possible breakdown locations. They provided good coverage

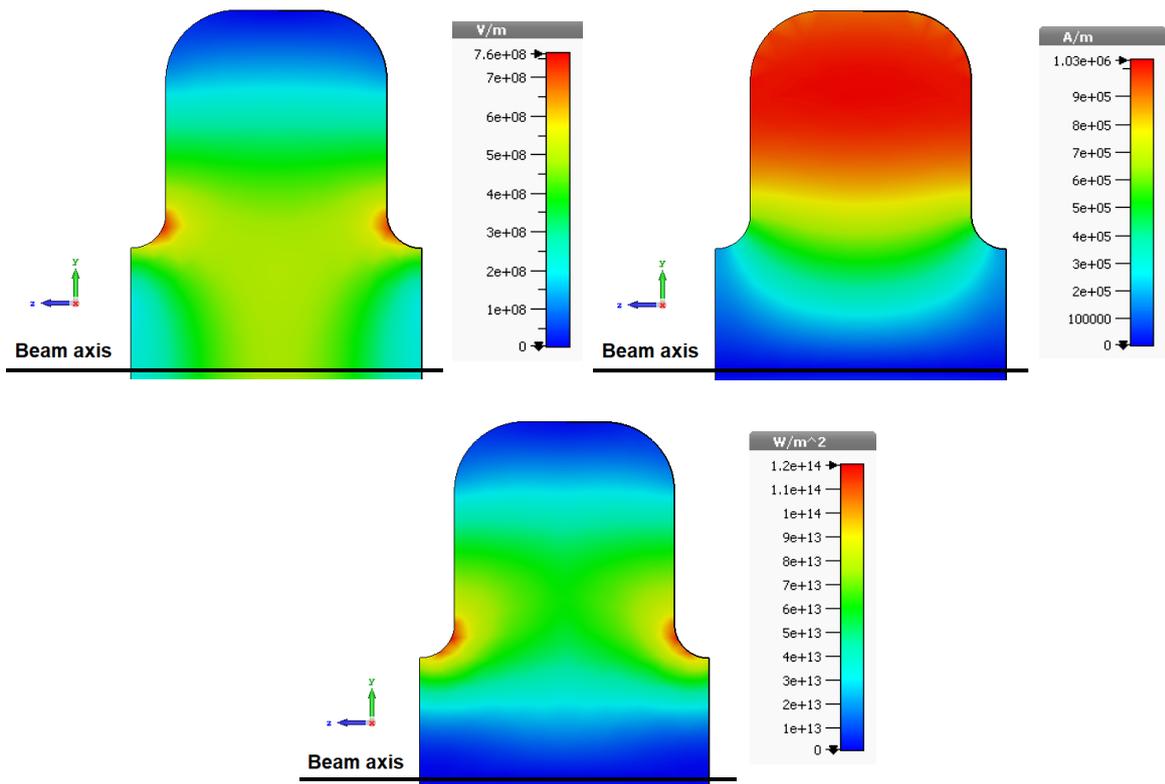


Figure 2.14: Field magnitudes vs. position on a longitudinal slice of the structure, obtained from an eigenmode solution with a 3.5 mm aperture radius. Top left: electric field. Top right: magnetic field. Bottom: S_c .

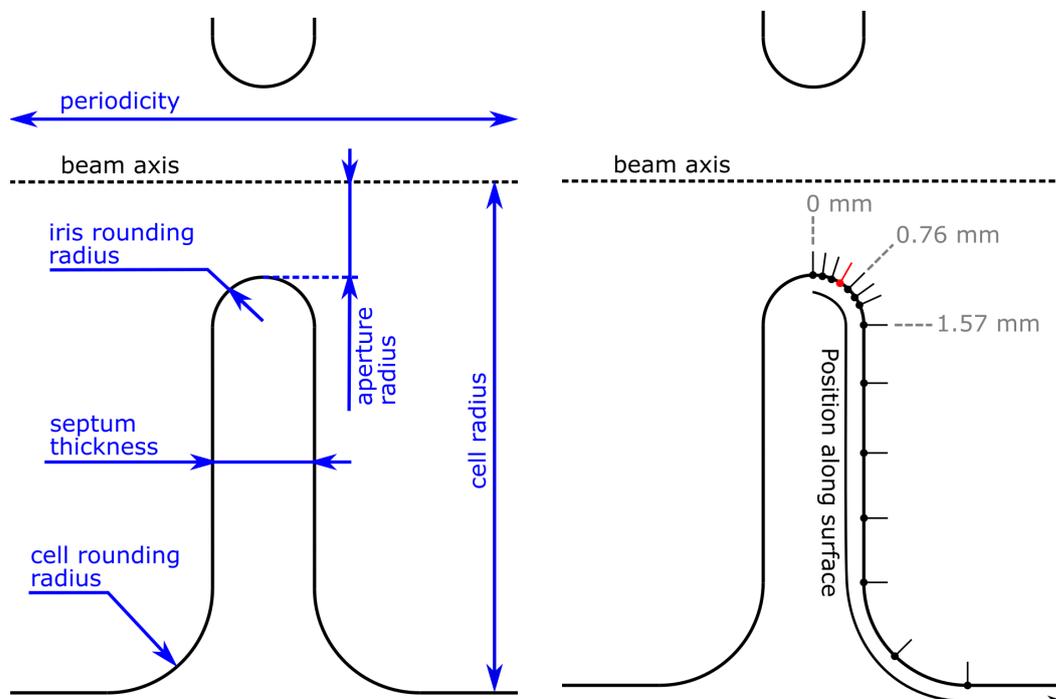


Figure 2.15: Left: dimensions of a regular cell. Right: placement of antennas for numerical R_{bd} calculations, with markings showing the distances from the iris apex in millimetres. Position 4 is highlighted in red. Antenna length not to scale.

Parameter	Value	Unit
Operating frequency	12	GHz
Phase advance	120	°
Phase velocity / c	1	-
Periodicity	8.333	mm
Iris thickness	2	mm
Iris rounding radius	1	mm
Cell rounding radius	2	mm

Table 2.2: Parameters common to all of the C3 structures simulated.

of the entire surface of the cell due to the radial and back-to-front symmetry of the simulated cells.

Structure Design

Since calculating E^* involved a consideration of power flow, an eigenmode solution was not sufficient and an S-parameter calculation was needed. Impedance matching, which plays a significant role in the power transferred between two ports, was considered crucial. Hence, it would not be useful to simulate the resonant cells alone without specifying how power was coupled into them. To this end, a travelling-wave structure was constructed out of the cells described above, coupled to circular waveguides of 12 mm radius terminated with matched waveguide ports, one being the power input and the other being the output. Fig. 2.16 shows a diagram of this setup. The input port was stimulated with TM wave. The sections of circular waveguide needed to be above cutoff for this mode at 12 GHz to allow the incident wave to propagate through them. The cutoff frequency f_c for the TM_{01} mode of a circular waveguide of radius a is given by [68]:

$$f_{c, TM_{0,1}} = \frac{c}{2\pi} \frac{\chi_{0,1}}{a}, \quad (2.35)$$

where c is the speed of light and $\chi_{0,1}$ is the first root of the 0th Bessel function of the first kind. In this case, the cutoff frequency was 9.57 GHz, well below the operating frequency.

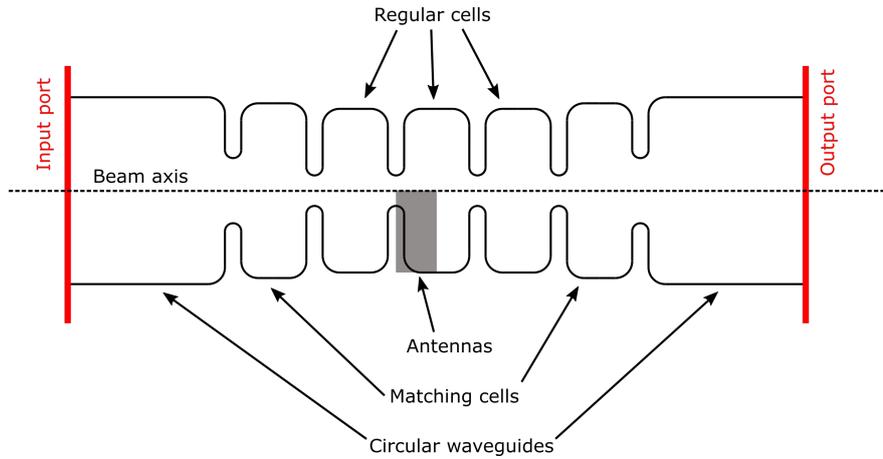


Figure 2.16: Diagram of the full structure simulated. The shaded area represents the region in which antennas were placed for calculating R_{bd} .

Iris radius [mm]	Cell radius [mm]	Matching iris radius [mm]	Matching cell radius [mm]
2.5	10.083	4.803	10.438
3	10.211	5.316	10.608
3.5	10.373	5.765	10.785
4	10.563	6.194	10.975
5	11.026	6.962	11.381
6	11.584	7.660	11.817

Table 2.3: Optimised geometrical parameters of each of the C3 structures simulated.

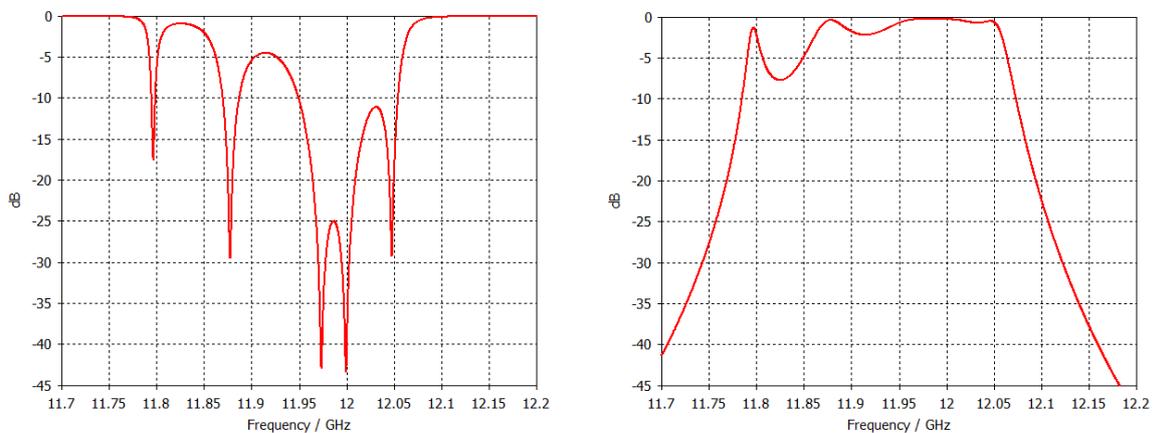


Figure 2.17: S-parameters vs. frequency in GHz of the structure with 3.5 mm aperture radius. Left: Magnitude of S_{11} . Right: Magnitude of S_{21} .

To ensure an impedance match between the structure and the circular waveguides, matching cells were incorporated between the regular cells and the input and output waveguides. The radius of the matching cells and the outermost irises could be varied independently of the parameters of the regular cells to provide minimum reflection at 12 GHz, with the optimised parameter values shown in Table 2.3. As an example, the simulated S-parameters of the 3.5 mm aperture structure are shown in Fig. 2.17. The S_{11} of the structure is very low at 12 GHz, showing that its impedance is well-matched to that of the input waveguide. In order to prevent the bleeding of fields from the matching cells into the regular cells from affecting the breakdown impedance results, the structure had three identical regular cells with the centre cell containing the antennas. In accordance with the nomenclature of CLIC structure designs, this structure was given the designation ‘C3’, indicating a constant impedance structure with three regular cells. This helped make the conditions in the centre cell more similar to those within an infinitely long constant gradient structure. Plots of the electric field along the beam axis of the full structure are shown in Fig. 2.18, demonstrating a 120° phase advance and good field flatness within the regular cells.

Results

A breakdown loading plot for one of the aperture sizes of this geometry is shown in Fig. 2.19, presenting the behaviour of each of the antennas simulated. Position 4 has the largest equilibrium E^* value of 138 MV/m, making it the most likely breakdown location in the cell. (2.8) was used as the emission function, with the parameter values $k = 5.4 \times 10^{-9} \text{ Am}^{3/2}\text{V}^{-3/2}$ and $n = \frac{3}{2}$. The value for k was manually chosen and found to produce reasonable results for both this study and simulations of tested prototype geometries, whereas $n = \frac{3}{2}$ arises from the assumption that the emission is space charge-limited, as discussed in Sec. 2.2. See Sec. 2.3.7 for details of the choice of the value of k . The antennas were chosen to be 0.1 mm long which was found to be an acceptable compromise between satisfying the condition that $l_{ant} \ll \lambda$ and ensuring a reasonable mesh size. This is also not an unreasonable size based on experimental measurements

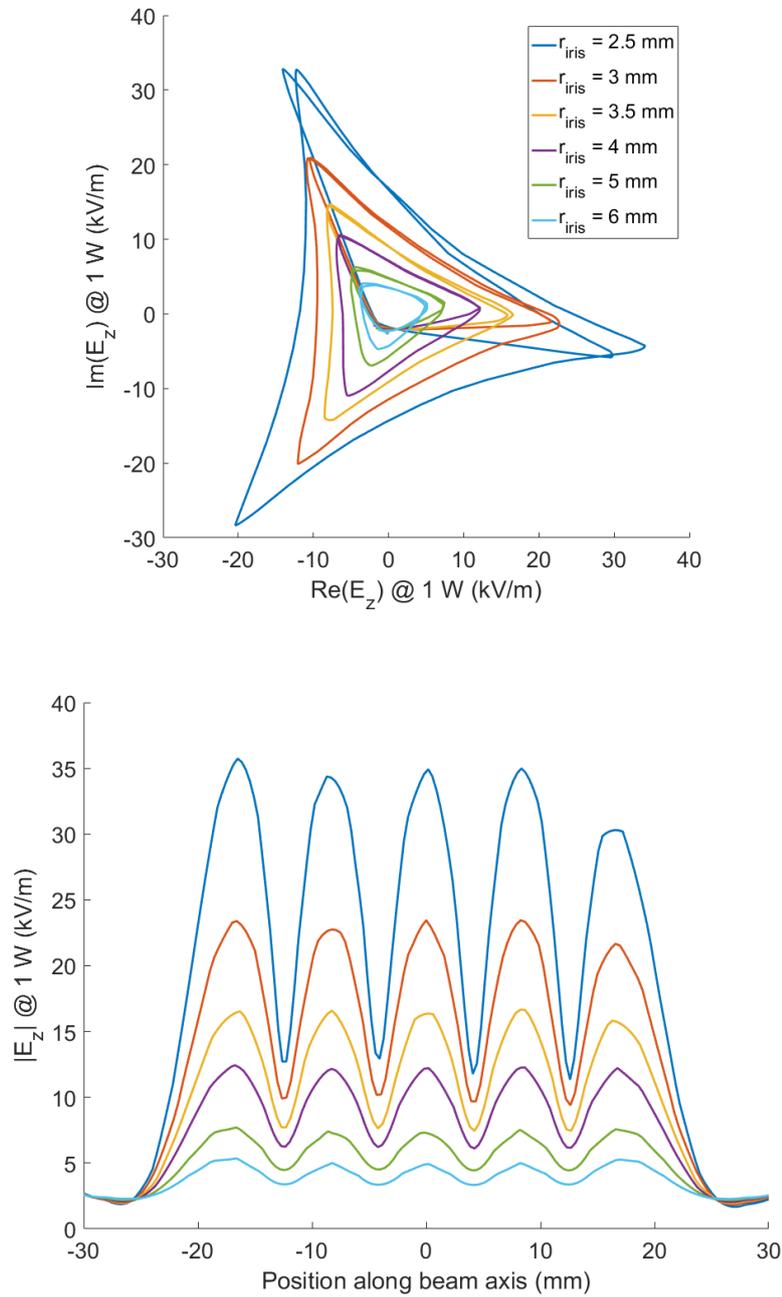


Figure 2.18: Phasor of the longitudinal component of the electric field on the beam axis E_z at an input power of 1 W. Different coloured plots represent C3 structures of different aperture sizes. Top: real part of E_z vs. imaginary part of E_z . Bottom: Magnitude of E_z vs. longitudinal position along the beam axis.

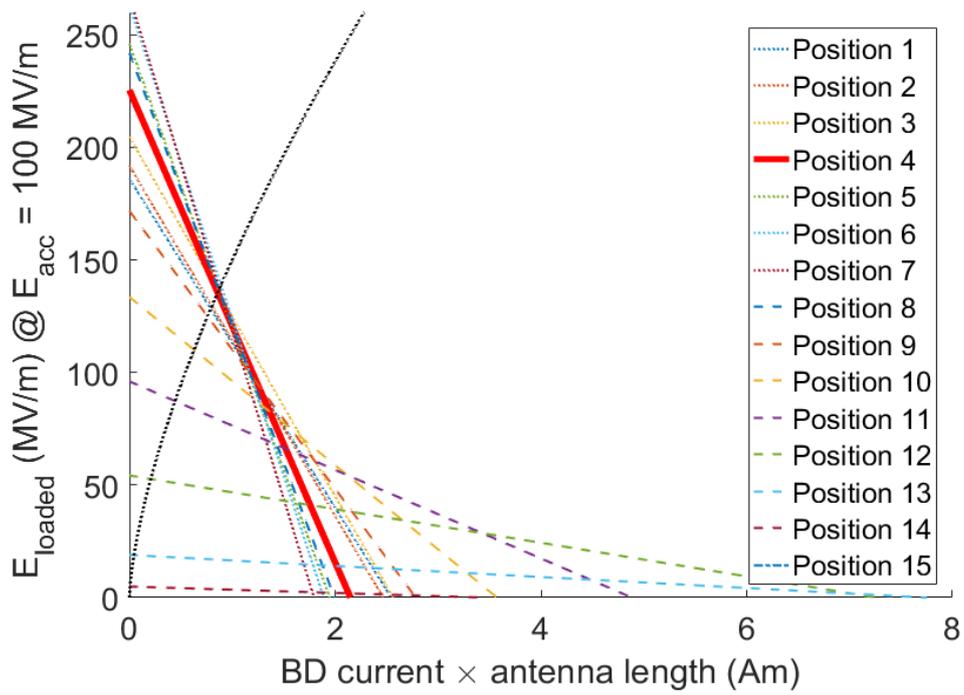


Figure 2.19: Surface electric field in MV/m vs. product of breakdown current and antenna length in Am for a 3.5 mm aperture radius and 100 MV/m accelerating gradient. Coloured lines: loading plots for different antenna locations as per Fig. 2.15, with position 1 representing the iris apex and position 15 representing the cell wall. The loading plot for position 4, which had the greatest value of E^* , is highlighted in red. Black dotted line: emission function.

of the plasma spot height. The smaller the antenna length chosen, the finer a mesh would need to be used for simulation, which is limited by available computing resources. These considerations are discussed in more detail in Sec. 2.2.1.

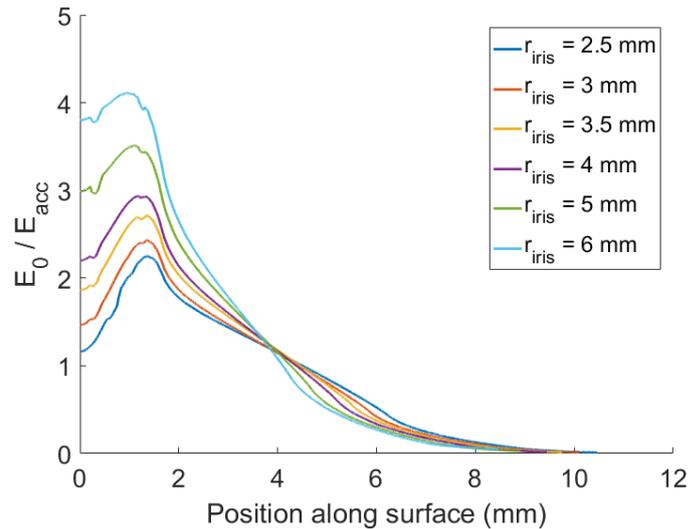


Figure 2.20: Unloaded surface electric field at unit accelerating gradient vs. position along the surface in mm. The position is defined by the arrow shown in Fig. 2.15, where 0 is the iris apex. Different coloured plots represent C3 structures of different aperture sizes. The discontinuities near the 0.5 mm and 1.5 mm marks on some of the plots are likely the result of the rounded geometry of the iris being discretised as a tetrahedral mesh in the finite-element simulation.

Figs. 2.20 and 2.22 show the distribution of E_0 and S_c respectively along the surface of the centre cell in the C3 structures simulated for various aperture sizes. The position is measured along the arrow shown in Fig. 2.15. Both quantities exhibit a peak about 1.5 mm from the apex of the iris, and are lower further away from the iris. This is consistent with typical *post-mortem* results, an example of which is shown in Fig. 2.21, which show the greatest concentration of breakdown craters in that region. It can also be seen that larger apertures result in larger peak fields for a given accelerating gradient, which is one of the reasons for making the aperture as small as possible in many practical designs from a gradient point of view.

Fig. 2.23 shows the distribution of E^* for the same geometries and locations. Due to the non-linearity of E^* with respect to the accelerating gradient, it cannot be normalised to E_{acc} in the same way that E_0 and S_c are. Instead, values are plotted for a gradient of

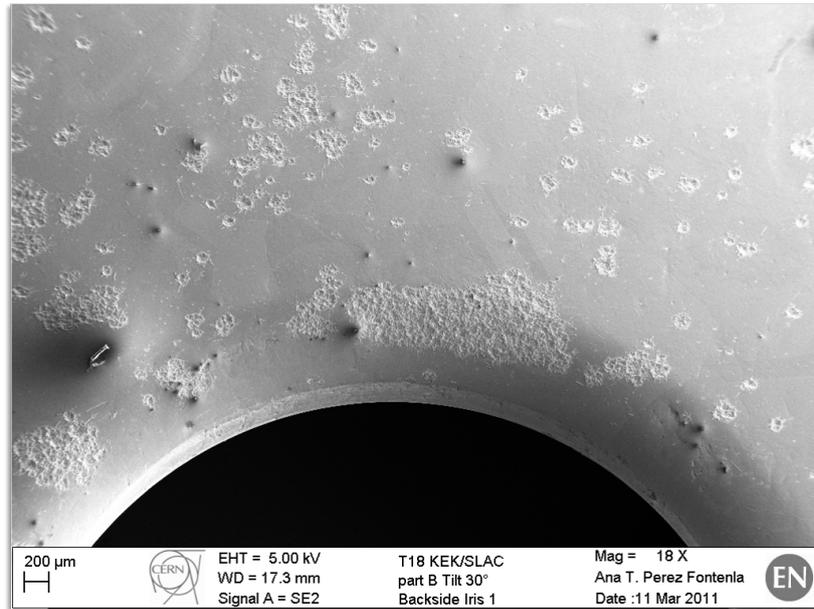


Figure 2.21: An SEM image of an iris of a T18 accelerating structure taken during *post-mortem* analysis [40]. There is a very large concentration of breakdown craters just off the iris apex.

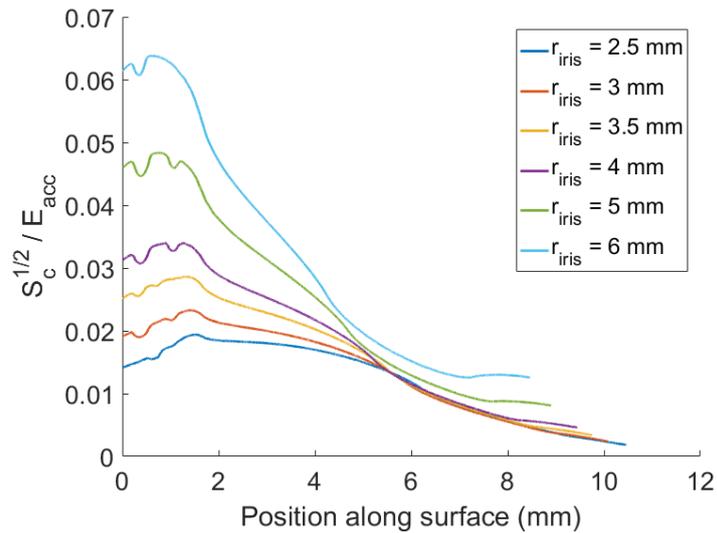


Figure 2.22: $\sqrt{S_c}$ at unit accelerating gradient vs. position along the surface in mm. The position is defined by the arrow shown in Fig. 2.15, where 0 is the iris apex. Different coloured plots represent C3 structures of different aperture sizes. The discontinuities near the 0.5 mm and 1.5 mm marks on some of the plots are likely the result of the rounded geometry of the iris being discretised as a tetrahedral mesh in the finite-element simulation.

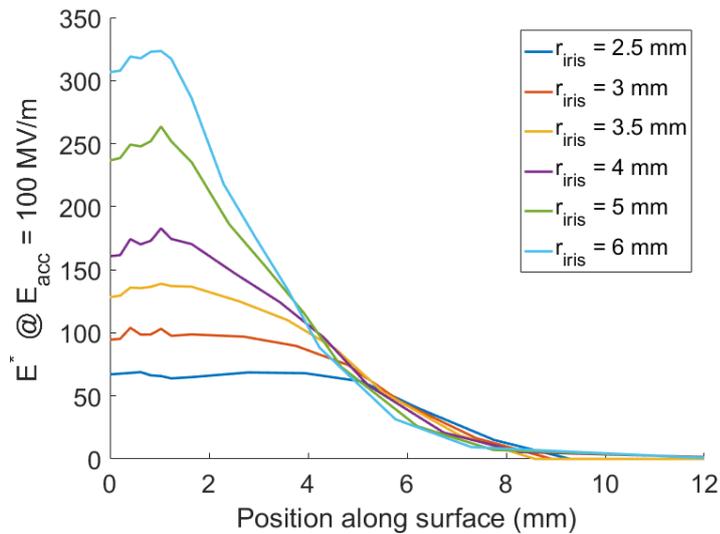


Figure 2.23: E^* at an accelerating gradient of 100 MV/m vs. position along the surface in mm. The position is defined by the arrow shown in Fig. 2.15, where 0 is the iris apex. Different coloured plots represent C3 structures of different aperture sizes. The discontinuities near the 0.5 mm and 1.5 mm marks on some of the plots are likely the result of the rounded geometry of the iris being discretised as a tetrahedral mesh in the finite-element simulation.

100 MV/m, which is the nominal average gradient for the 1.5 TeV and 3 TeV designs of CLIC [12]. E^* appears to exhibit qualitatively similar behaviour to E and S_c , showing a peak around 1.5 mm from the iris apex for iris radii 3 mm and above. The distribution of E^* is fairly flat at a radius of 2.5 mm and becomes more peaked for larger radii. This behaviour is also apparent with S_c , though a peak is still visible for the 2.5 mm case. The breakdown locations within cells operating in the TM_{010} mode thus appear to be qualitatively consistent between S_c and E^* with the chosen emission function.

A more quantitative approach for verifying this case is to compare the maximum accelerating gradient implied by various models, as shown in Fig. 2.24. In each case, a maximum allowable value for each of the three quantities P/C , S_c , and E^* was chosen. The maximum values for S_c and P/C are typical of CLIC accelerating structures [12], [42], [72]. The agreement between the three quantities is very good from a group velocity of about 1.5% of c and above. Below that value, both E^* and P/C predict much higher achievable gradients than S_c . This is believed to be due to the effect of stored

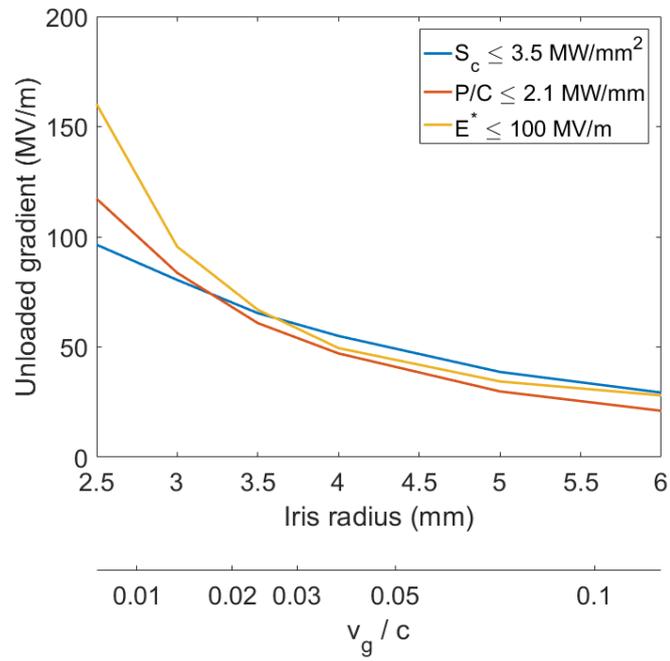


Figure 2.24: Maximum accelerating gradient without beam loading vs. aperture size and group velocity as a fraction of c . Blue curve: with a maximum permitted S_c of 3.5 MW/mm^2 . Red curve: with a maximum permitted P/C of 2.1 MW/mm . Yellow curve: with a maximum permitted E^* of 100 MV/m .

energy, which will require considerations of transient behaviour and will be discussed in Sec. 2.4.2.

It has already been established that P/C is close to zero for standing-wave structures [42], predicting unrealistically high maximum gradients (see Sec. 2.1.2). This is not the case for S_c , which does take into account reactive power flow, corresponding to energy stored in the cell. Since S_c has been shown to give much more reasonable results for standing-wave structures than P/C , it appears that it is important to consider stored energy. It seems physically justifiable to do so, since E^* is an attempt to qualitatively determine if there is enough power available to sustain the evolution of a pre-breakdown site into a full breakdown. Thus, it may be the case that even though the power flowing in from outside the structure may be insufficient for this, there may be enough energy stored locally in the cell to allow the breakdown to fully develop. Since a standing-wave structure has zero group velocity, it is reasonable to expect that the lower the group velocity of a travelling-wave structure, the greater the relative contribution of stored energy to the breakdown process. Hence, one should expect that the gradient predicted by P/C gradually deviates from that of S_c as the group velocity approaches zero. The process of calculating R_{bd} detailed in Sec. 2.2 assumes steady-state conditions, implicitly assuming that all stored energy has been depleted and the only source of power is external to the structure. A modification of the calculation of E^* to take stored energy into account will be discussed in Sec. 2.4.2.

It should be noted that structures designed with S_c in mind usually have irises with an elliptical cross-section, rather than circular as is the case here. Making the iris elliptical has the effect of lowering the peak S_c for a given gradient, improving the breakdown performance slightly. The design of the C3 structure has not been optimised for S_c , resulting in a relatively low gradient with the accepted maximum value of $S_c = 3.5 \text{ MW/mm}^2$.

Conclusion

This study of C3 structures has shown that the proposed method can be implemented in practice with numerical finite-element simulations and gives self-consistent results. The choice of the structure's design was made to be roughly representative of a high gradient travelling-wave X-band structure for a linear collider. Several aperture sizes were also simulated to take into account the range of group velocities that a practical design could have, especially relevant for tapered structures. By sampling R_{bd} and thus E^* at various locations on the wall of the centre cell, it could be verified that the prediction of both breakdown location and maximum gradient by the E^* model appear to be consistent with experimental results and S_c . From these results, one can conclude that the E^* method will most likely work well for most X-band accelerating structures operating in the TM_{010} mode. Low group velocity structures, for which some discrepancy was found, will be revisited later, as they require a consideration of the local stored energy in addition to the external power supplied.

2.3.2 Antenna Length Dependence in Numerical Simulations

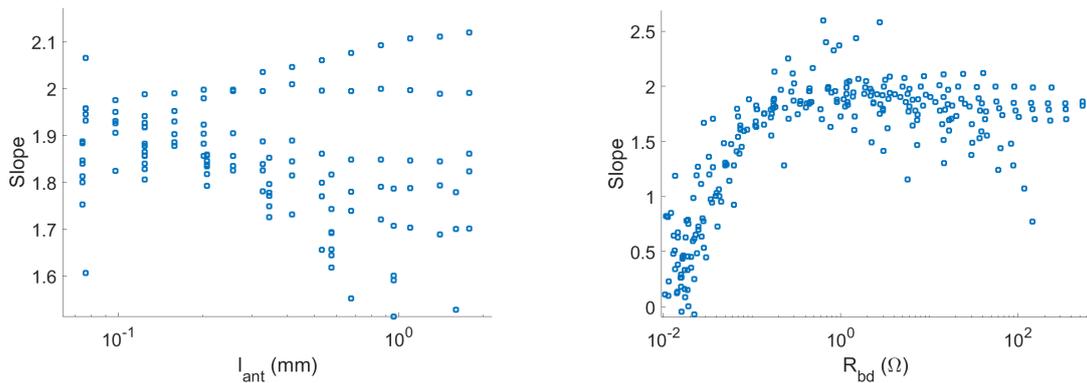


Figure 2.25: $\log(\Delta R_{bd})/\log(\Delta l_{ant})$, i.e. the slope on a logarithmic plot of R_{bd} vs. l_{ant} . A value of 2 indicates a quadratic dependence on antenna length, seen for most antennas at a length of around 0.1 mm. For very small antenna lengths, the value of R_{bd} becomes dominated by numerical errors, and for very large lengths the field can no longer be assumed to be uniform along the antenna's entire length, leading to slopes that deviate from 2. See text for more details. Left: slope vs. antenna length in mm. Slopes for antennas which had a value of $R_{bd}(0.2 \text{ mm}) < 0.5 \Omega$ are not shown. Right: slope vs. R_{bd} in Ω .

The analytical circuit models discussed in Sec. 2.2.1 featured a quadratic dependence of R_{bd} on the antenna length. In principle, any value of antenna length can be used in a simulation, with the knowledge that the result can always be re-scaled as needed using this dependence. In practice, the antenna length had to be kept within a certain range to obtain usable results.

To aid with the choice of antenna length, a set of simulations was performed in which the antenna length was scanned over a range of values in the C3 structures described in Sec. 2.3.1, with aperture sizes of 2.5, 4 and 8 mm. R_{bd} values were obtained for each of the antennas placed in the middle cell as shown in Fig. 2.15.

Fig. 2.25 shows the quantity $\log(\Delta R_{bd})/\log(\Delta l_{ant})$, (i.e.: the slope on a logarithmic plot of R_{bd} vs. l_{ant}) as a function of antenna length and R_{bd} for all of these cases. For most of the antenna locations over the three aperture sizes considered, there is a central region in which the breakdown resistance has a quadratic dependence on antenna length, i.e. $R_{bd} \propto l_{ant}^2$, determined using the slope of R_{bd} as a function of l_{ant} on logarithmic axes. Deviations from the quadratic dependence on antenna length can be seen for very long and very short antennas.

As the antenna length was increased, some of the curves deviate from $R_{bd} \propto l_{ant}^2$. The left plot in Fig. 2.25 shows this more clearly. This could be due to the assumption made in Sec. 2.2 that $l_{ant} \ll \lambda$. As an antenna is made longer in a region of non-uniform electric field, the integrated field across it may no longer depend linearly on the antenna length. In the presence of field-enhancing features, the field changes over a distance on the order of the size of the feature, which may be much smaller than the free-space wavelength (25 mm for a frequency of 12 GHz). The iris thickness of this structure is 2 mm, hence it is reasonable to expect significant changes in field over that distance and it is advisable to make the antennas significantly shorter than that. Such variations in field can be observed in the eigenmode field maps shown in Fig. 2.14, where there are regions of high E field about 1 mm in size close to the iris apex. This also explains why the slope of the plots increases in some cases and decreases in others

as the antenna length increases, as the effect depends on the location of the antenna and the specific field distribution at that location.

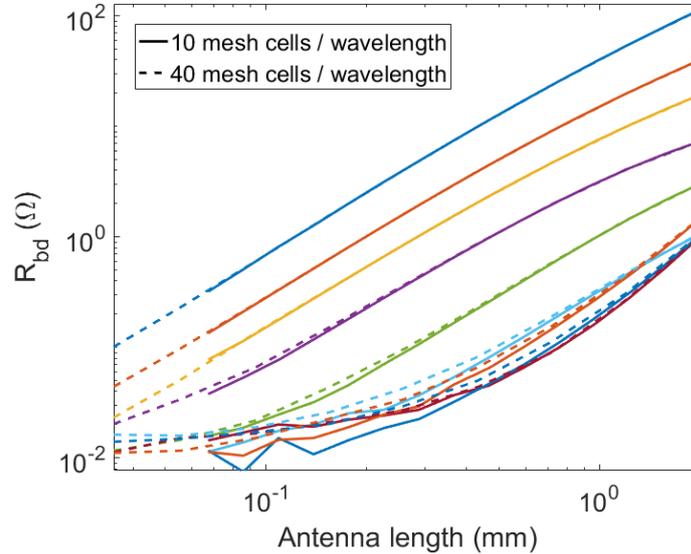


Figure 2.26: Calculated R_{bd} in Ω vs. antenna length in mm. Different colours represent different antenna locations in a C3 structure of 8 mm aperture radius. Solid curves: results obtained with a mesh size of 10 cells per wavelength. Dashed curves: results obtained with a mesh size of 40 cells per wavelength.

For very small antenna sizes, numerical errors seemed to dominate the R_{bd} values. None of the antennas seemed to have an R_{bd} value lower than $10\text{ m}\Omega$, which was believed to be an issue with the implementation of the simulation software. This caused the expected relation $R_{bd} \propto l_{ant}^2$ to break down, as can be seen for small values of R_{bd} in the right plot in Fig. 2.25. To test the influence of the mesh size on this error, the results of two otherwise identical simulations using different mesh sizes are compared in Fig. 2.26. In this figure, it can be seen that the $10\text{ m}\Omega$ lower limit did not seem to change with the size of the mesh, but the finer mesh did allow for smaller antenna sizes to be used.

From the results of this investigation, it should follow that the antenna length should be kept within a certain range of values for optimal results. The dependence $R_{bd} \propto l_{ant}^2$, expected from the theoretical calculations from Sec. 2.2.1, was used as a way of verifying the simulation results. Since errors on the order of $10\text{ m}\Omega$ were observed, the R_{bd} values should be made as large as possible to keep the relative error small. This can be done by

making the antennas long, which also reduces the need for a very fine mesh. However, making the antennas too long would also reduce the accuracy of the results, as the assumption that the electric field is constant along the length of the antenna would no longer hold. The maximum length that an antenna can be made depends on the size of the accelerating cell and the exact field pattern, but a length significantly smaller than the cell is advisable. Plots similar to the ones shown in Fig. 2.25 would be a useful tool in determining the exact length in any given case, as this effect can be seen quantitatively as a deviation of the slope from 2.

An antenna length of 0.1 mm was used for the simulations results of the C3 geometry presented in this thesis. This value was shown to give the smallest spread in slope values in the left plots of Fig. 2.25. Simulations of tested geometries such as the T24 (Sec. 2.3.5) or the CLIC Crab Cavity (Sec. 2.3.4) used antenna lengths of 0.5 mm instead, due to the much greater complexity of these geometries requiring the use of a coarser mesh to simulate. Although not optimal, the 0.5 mm value was still smaller than any geometrical features, and was kept constant between different geometries to allow meaningful comparisons between them.

2.3.3 Comparison of the Analytical Model with Numerical Simulations

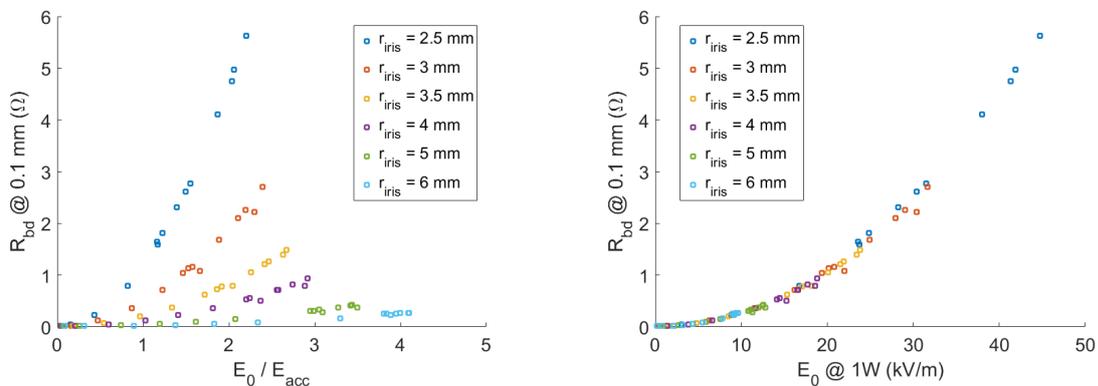


Figure 2.27: Breakdown resistance vs. unloaded surface electric field E_0 . Different coloured points represent C3 structures of different aperture sizes, as indicated on the legend. Left: E_0 values given for unit accelerating gradient. Right: E_0 values given for 1 W input power.

Additional analysis was done on the results of the numerical simulations of C3 structures discussed in Sec. 2.3.1 to check if they are consistent with the analytical model discussed in Sec. 2.2.1.

E_0 and R_{bd} are plotted against each other in Fig. 2.27 for a range of antenna locations and different aperture sizes. In this figure, it can be seen that $R_{bd} \propto E_0^2$ within any given cell, but the proportionality constant varies with aperture size. The larger the aperture, the higher the R_{bd} values, consistent with the results discussed in Sec. 2.3.1, which show that smaller apertures allow higher gradients for a given maximum E^* . If the electric fields are normalised to a constant input power, the proportionality constant becomes the same for all aperture sizes. This suggests that the relationship between accelerating gradient and input power, as in (2.2), plays an important role in defining R_{bd} . It is encouraging because it shows an explicit dependence on power flow and group velocity which have been observed to be important factors in determining the breakdown limit. From the behaviour apparent in Fig. 2.27, it is possible to deduce that:

$$R_{bd} \propto \frac{E_0^2}{P}, \quad (2.36)$$

where P is the incident RF power. By combining (2.36) with (2.2), one can obtain an expression for R_{bd} in terms of quantities that can be calculated from an eigenmode simulation of the cell:

$$R_{bd} \propto \frac{1}{v_g} \frac{R}{Q} \left(\frac{E_0}{E_{acc}} \right)^2 \quad (2.37)$$

This is consistent with the expression (2.15), confirming that the numerical results are indeed consistent with the analytical model.

For the geometry discussed in Sec. 2.3.1, the dependence of R_{bd} at the location of peak E_0 on aperture size and group velocity can be obtained using the results shown in Fig. 2.13, and are shown in Fig. 2.28. The slope of the line in the right plot in

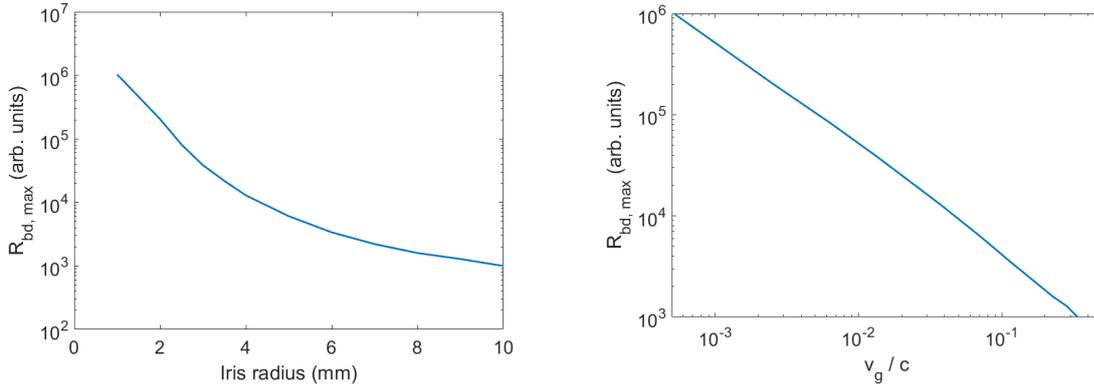


Figure 2.28: Scaling of breakdown resistance at the location of peak E_0 calculated from eigenmode results using (2.37). Left: R_{bd} vs. iris radius in mm. Right: R_{bd} vs. group velocity as a fraction of c .

this figure shows a dependence of approximately $R_{bd} \propto v_g^{-1.1}$. The deviation from the expected $R_{bd} \propto v_g^{-1}$ is due to the fact that R/Q and E_0/E_{acc} are not independent of the group velocity in this structure geometry. However, they vary relatively little compared to the variation in v_g as the aperture size is varied. This dependence on group velocity appears consistent with the use of v_g as an approximate predictor of breakdown performance as discussed in Sec. 2.1.2.

The correlation of the spatial distributions of R_{bd} and E_0 within a cell is a useful result as it means that the location of highest E^* in a given cell will, in most cases, be the location of highest E_0 , showing consistency with experiments, where breakdowns tend to happen at the location of peak surface electric field. It also helps guide the choice of antenna location in simulations, since it implies that regions of peak E_0 are of the greatest relevance for a breakdown limit. This can significantly reduce the time and effort required to perform a numerical calculation of E^* for a given geometry, as every potential breakdown site requires a separate simulation. In contrast, a quantity such as S_c which is a linear function of the unloaded \mathbf{E} and \mathbf{H} fields requires only field maps from a single simulation. Thus, knowing *a priori* that breakdowns within any given cell are most likely to occur at the location of peak E_0 , only that location needs to be considered to calculate the maximum E^* , thus characterising the breakdown performance of the entire cell.

2.3.4 CLIC Crab Cavity

The CLIC Crab Cavity is one of the structure designs discussed in this chapter that has been built and tested at high power [18]. It is a prototype for a transverse deflecting cavity for CLIC. In order to compensate for the luminosity loss due to the 20 mrad crossing angle of the beams, the use of crab cavities similar to this one was proposed as a way of introducing a time-dependent transverse deflection of the beam, causing the bunches to rotate such that they collide head-on at the Interaction Point (IP) [56].

This prototype is a backward travelling-wave structure which was designed to have a relatively large group velocity of $0.029c$ to limit the effect of beam loading on the transverse kick [56]. As the purpose of this structure is to deflect the beam rather than accelerate it, it operates in the TM_{110} mode rather than the TM_{010} mode as with most accelerating cavities. This means that the direction of the electric field on the beam axis is transverse rather than longitudinal. A cylindrical cavity has two degenerate TM_{110} modes with the same frequency but different polarisations [68]. Since only one of these modes was desired, the cells of the Crab Cavity were given a racetrack shape to detune the undesired mode in order to stop it from being excited by the beam. The spatial distribution of the field magnitude, shown in Fig. 2.29, is different from a typical accelerating cavity. The lack of axial symmetry in this mode makes it useful for a comparison of the predictions of S_c and E^* . The final design has two waveguide feeds in each of the coupling cells to ensure that the centre of the dipole mode and the geometrical centre coincide. Otherwise, with a single feed, a bunch on the beam axis will see a monopole accelerating field, which is undesirable. The prototype tested in XBox-2 only had a single feed, which results in a clearly visible asymmetry in the electric field pattern of the coupling cells.

Since this structure played a significant role in motivating the E^* approach to breakdown limits, it is important to confirm that it predicts a breakdown distribution consistent with the *post-mortem* results. This structure has a constant impedance, which was important for choosing the locations where E^* was to be calculated. Constant

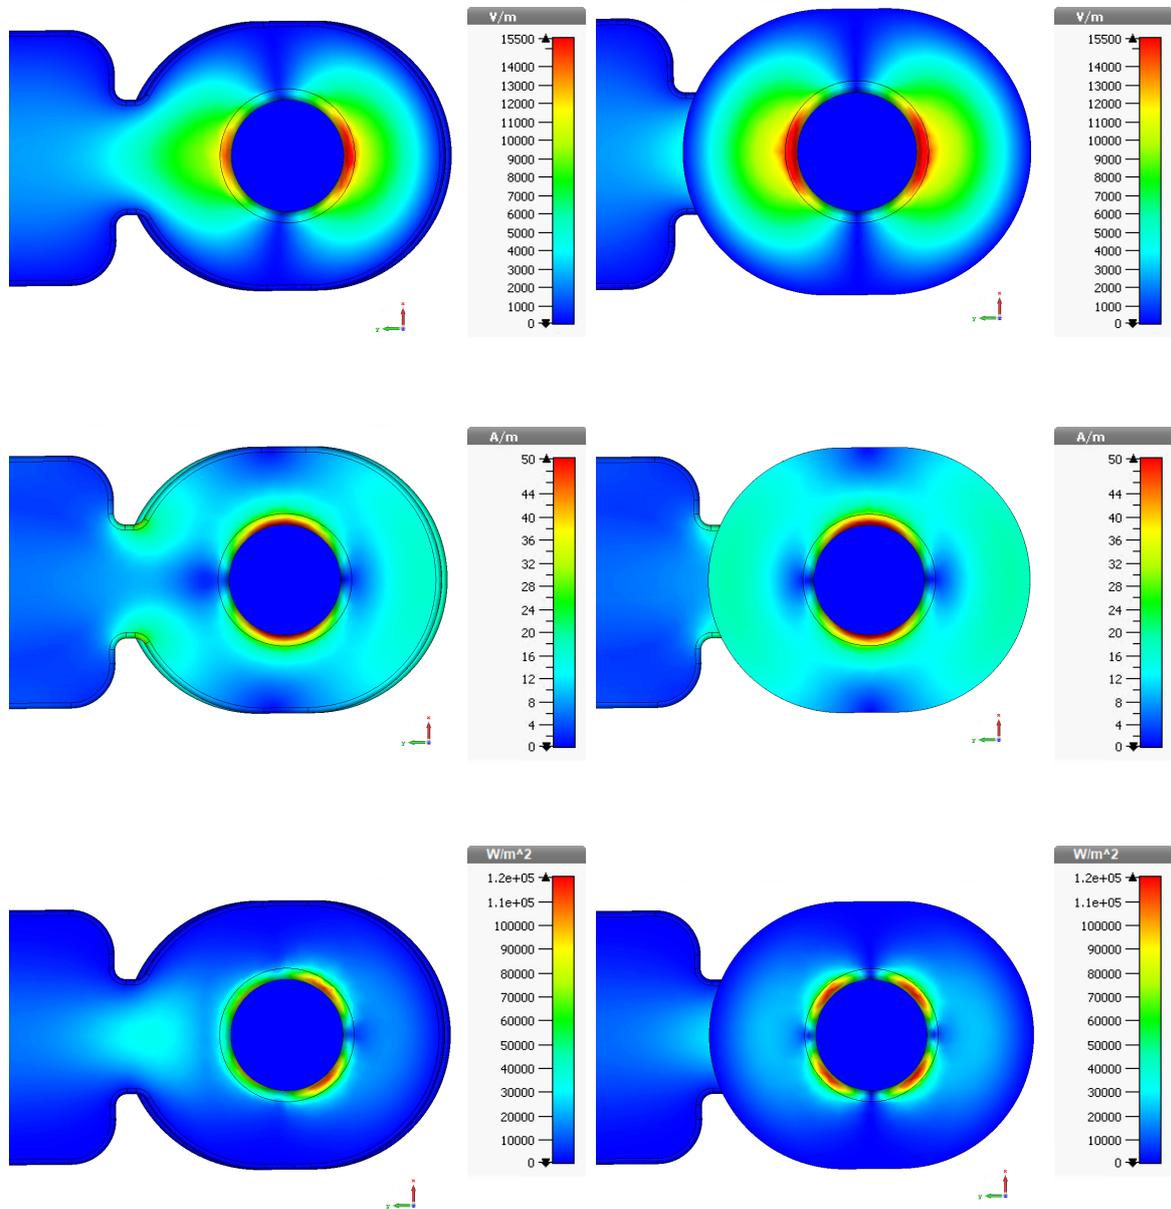


Figure 2.29: Fields vs. transverse position for the first two cells of the Crab Cavity prototype. Top left: electric field magnitude in Cell 1. Top right: electric field magnitude in Cell 2. Middle left: Magnetic field magnitude in Cell 1. Middle right: Magnetic field magnitude in Cell 2. Bottom left: S_c in Cell 1. Bottom right: S_c in Cell 2. ‘Cell 1’ refers to the input coupler cell and ‘Cell 2’ refers to the first regular cell.

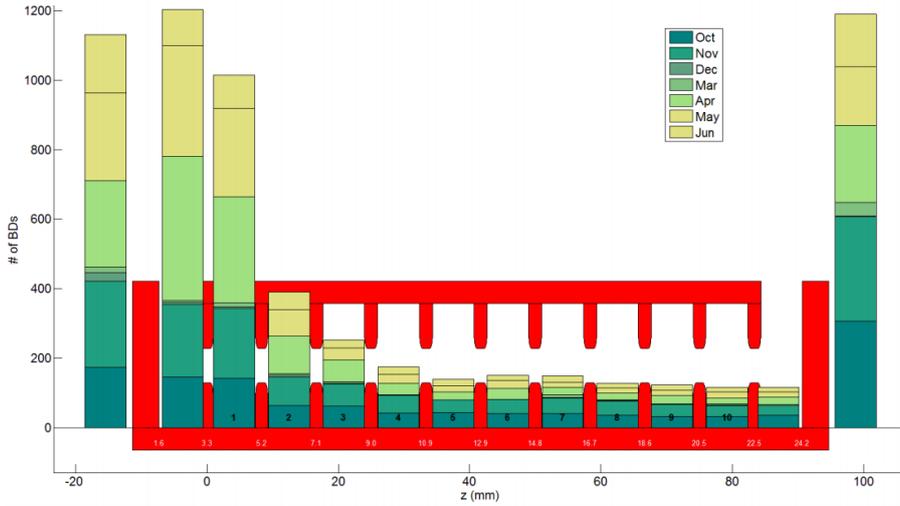


Figure 2.30: A diagram showing the longitudinal positions in mm of the cells of the CLIC Crab Cavity [18]. A histogram is overlaid showing the number of breakdowns in each cell. Each colour represents the breakdowns which occurred in each particular month of the structure’s conditioning. The bars on the extreme left and right of the plot represent breakdowns identified as having occurred outside of the structure, such as in the pulse compressor or high-power load.

impedance indicates that the regular cells have exactly the same geometry and therefore the same shunt impedance. This means that as the RF power decreases along the length of the structure due to resistive losses, the fields in each subsequent cell are smaller than in the previous cells [32]. Thus, it is expected that the majority of the breakdowns occur in the first cell. This was indeed the case when it was tested in XBox-2, as demonstrated in Fig. 2.30. Because of this, E^* was only calculated for the first two cells of the structure. Two cells were considered because the first cell is a coupling cell with a slightly different geometry and the presence of the input waveguide. The second cell is a regular cell which is identical to the other regular cells. Due to their identical geometry, the breakdown behaviour in the remaining regular cells is expected to be similar to the second cell but with slightly lower fields.

It has been established in Sec. 2.3.3 that within a given resonant cell, the location of peak E^* is expected to coincide with the location of peak E_0 . The peak E_0 and peak S_c are both located close to the apex of the iris, but they have distinctly different azimuthal positions. This simulation was performed to investigate how well the azimuthal distribution of E^* agreed with the actual breakdown locations, and to compare it with

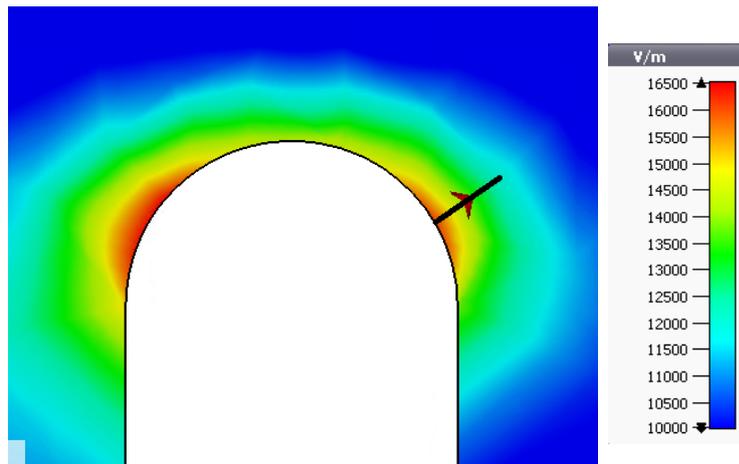


Figure 2.31: Electric field magnitude vs. position on a longitudinal slice of the CLIC Crab Cavity. The iris of the first cell can be seen and the position of one of the antennas is shown as a black line.

that of S_c . With this goal in mind, antennas were placed around one half of the iris (the structure has left-right symmetry) of the first two cells to check the azimuthal distribution of E^* around the iris. A plot of the electric field magnitude around the iris of the first cell and the position of one of the antennas relative to it is shown in Fig. 2.31. The base of the antenna was placed on the location of the local maximum of E_0 to obtain the limiting value of E^* . An antenna length of 0.5 mm was chosen due constraints from the total number of mesh cells needed to simulate the entire geometry of the structure.

The resulting breakdown loading plots for antennas placed in the first two cells of the Crab Cavity are shown in Fig. 2.33. The calculation was performed for an input power of 52 MW, which was the highest power reached during testing [18]. The equilibrium E^* values using the assumed emission function that has been used thus far are shown in Figs. 2.34 and 2.35, where they are compared with the azimuthal distribution of breakdown craters. The results appear encouraging: in both cells studied, E^* has two maxima: one at the 0° location and one at the 180° location. The distribution of E^* is asymmetric in the first cell and symmetric in the second cell, which also reflects the distribution of breakdown craters. The Crab Cavity prototype was designed to operate at 13.35 MW, and was tested up to 52 MW with a breakdown rate of $(1.10 \pm 0.25) \times 10^{-4}$

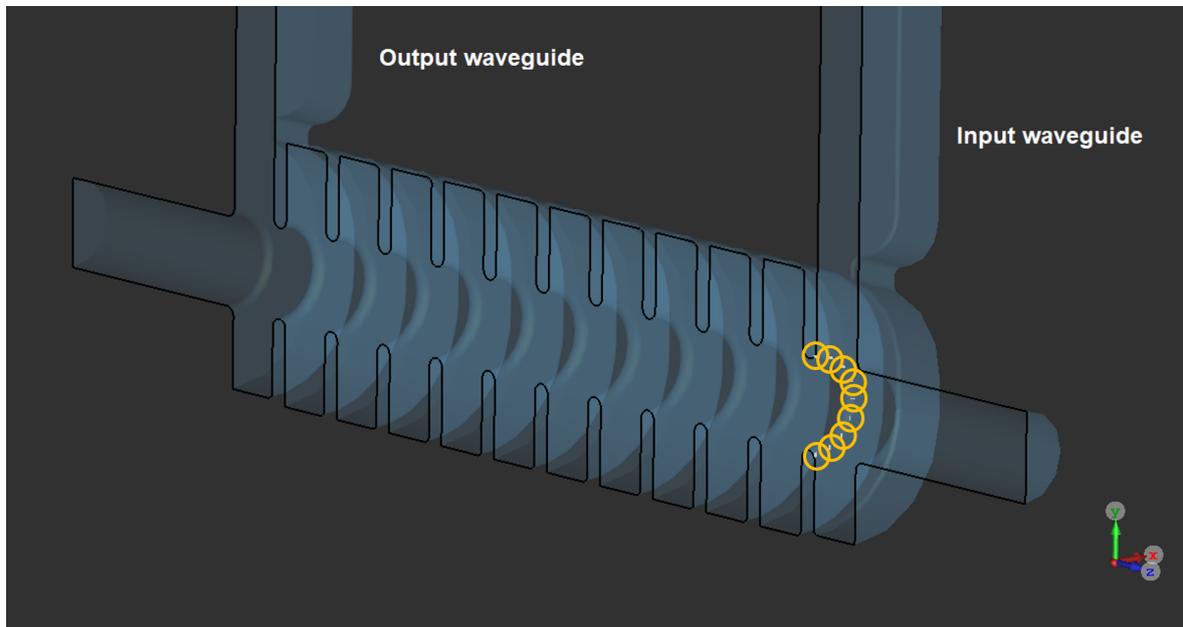


Figure 2.32: A cutaway view of the internal volume of the CLIC Crab Cavity with the locations of antennas shown as white lines with yellow circles around them.

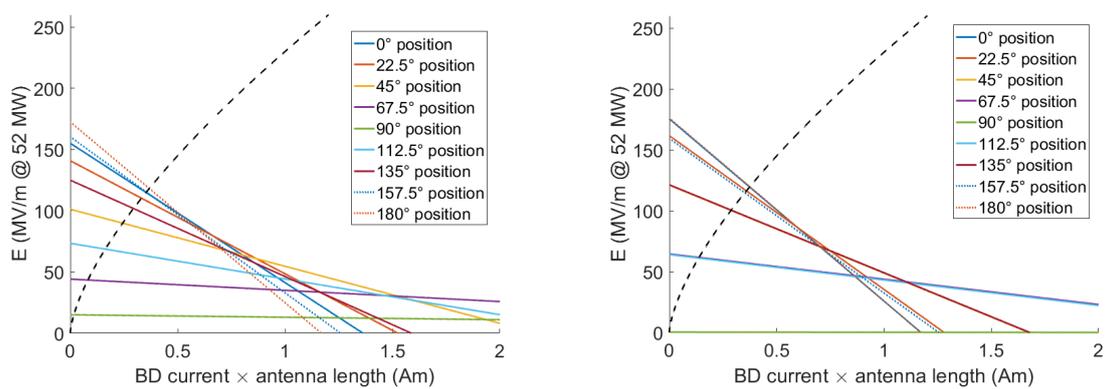


Figure 2.33: Surface electric field in MV/m vs. the product of breakdown current and antenna length in Am in the CLIC Crab Cavity at an input power of 52 MW. Coloured lines represent different antenna locations on the iris. Black dotted line: emission function. Left: Cell 1. Right: Cell 2.

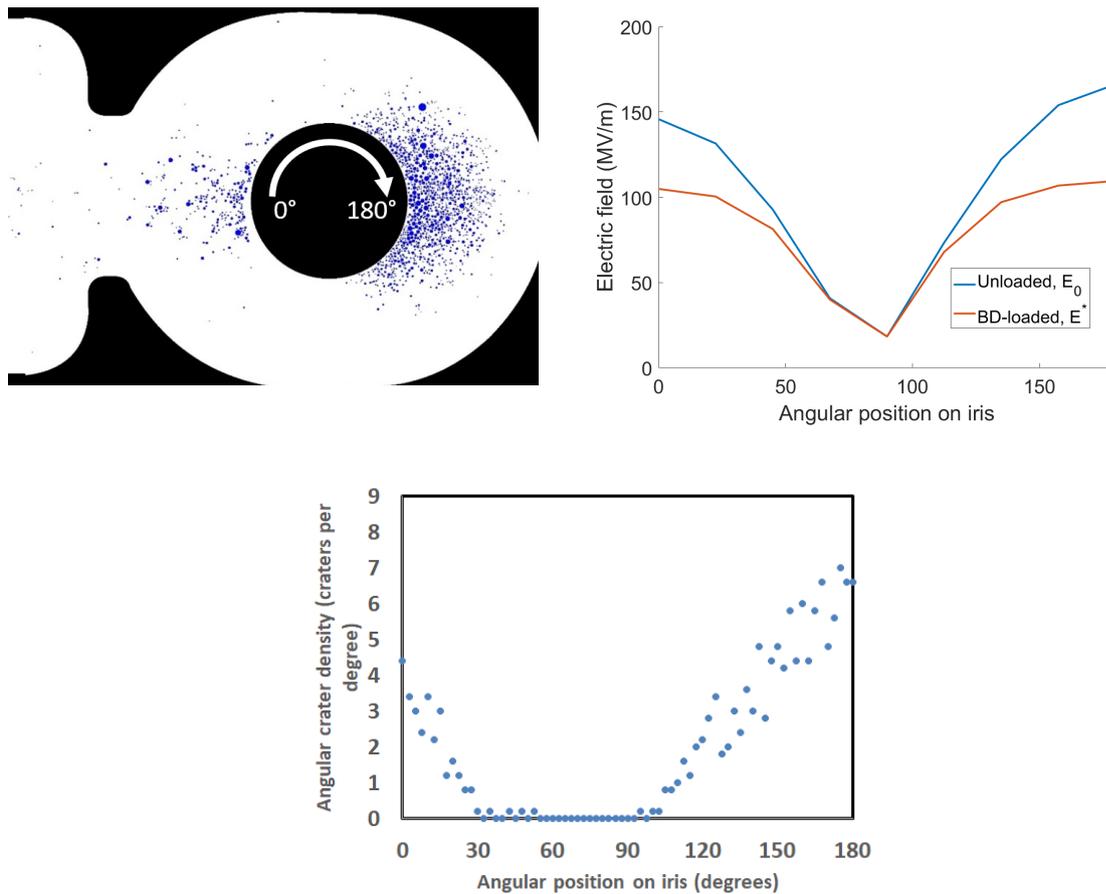


Figure 2.34: Left: breakdown locations (represented by blue dots) vs. transverse position in Cell 1 of the CLIC Crab Cavity. Right: unloaded surface electric field (blue) and E^* (red) in MV/m vs. angular position on the iris in degrees, at an input power of 52 MW. Bottom: angular density of breakdown craters in this cell vs. angular position in degrees, according to *post-mortem* studies [73]. The definition of the angular position relative to the geometry of the cell is shown in the left plot.

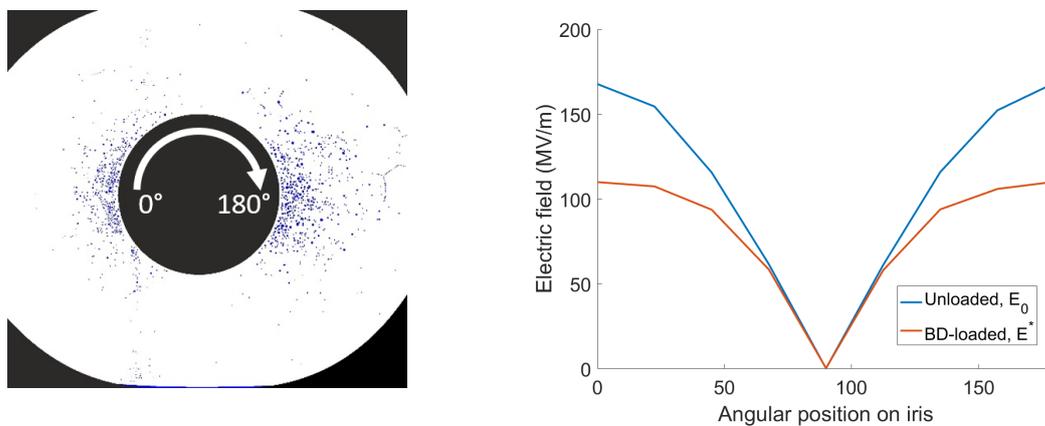


Figure 2.35: Left: breakdown locations (represented by blue dots) vs. transverse position in Cell 2 of the CLIC Crab Cavity. Right: unloaded surface electric field (blue) and E^* (red) in MV/m vs. angular position on the iris in degrees, at an input power of 52 MW. The definition of the angular position relative to the geometry of the cell is shown in the left plot.

breakdowns per pulse and a pulse length of 120 ns. The peak E^* and S_c values at the nominal and maximum power are given in Table 2.4.

2.3.5 T24 Structures

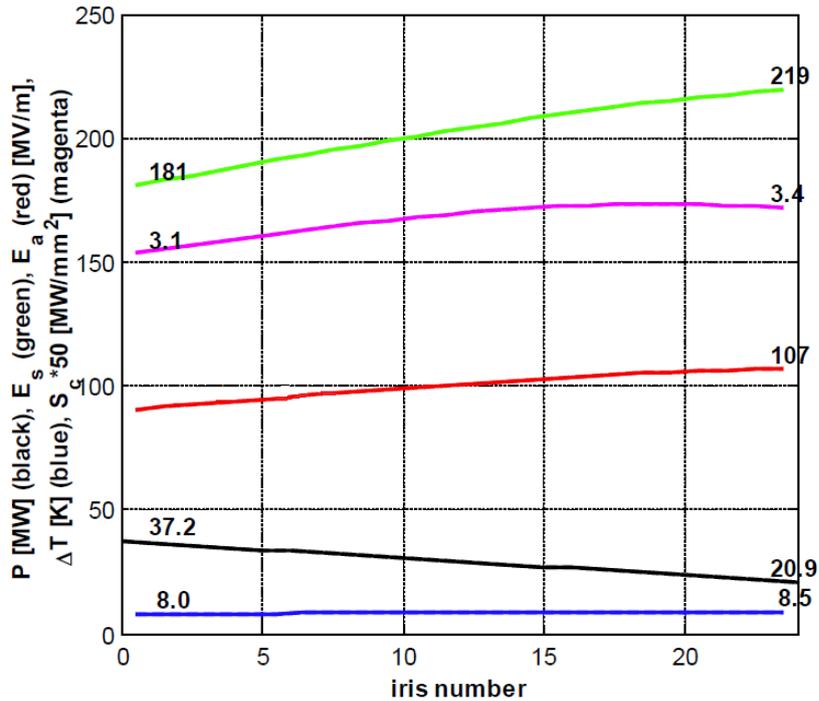


Figure 2.36: A plot demonstrating the variation of the RF power (black), surface electric field (green), accelerating gradient (red), temperature increase due to pulsed surface heating (blue), and S_c (magenta) with iris number in the T24 structure design when operated without beam at the nominal accelerating gradient of 100 MV/m [72]. Iris 1 is closest to the RF input and iris 24 is closest to the RF output.

In order to meet CLIC's specification for gradient and breakdown rate, a great deal of effort has been expended to optimise the gradient performance of CLIC accelerating structures. The T24 structure is a prototype design for a CLIC accelerating structure. It has 24 regular cells and 2 matching cells. The 'T' in its name indicates that the geometry of the structure is tapered, i.e. parameters such as aperture size and iris thickness vary along the length of the structure to control properties such as shunt impedance and peak fields in each cell. The RF power flowing through each cell decreases along the length of the structure due to resistive losses in the copper walls, and if a beam is present, beam loading as well. One of the goals of tapering is to make

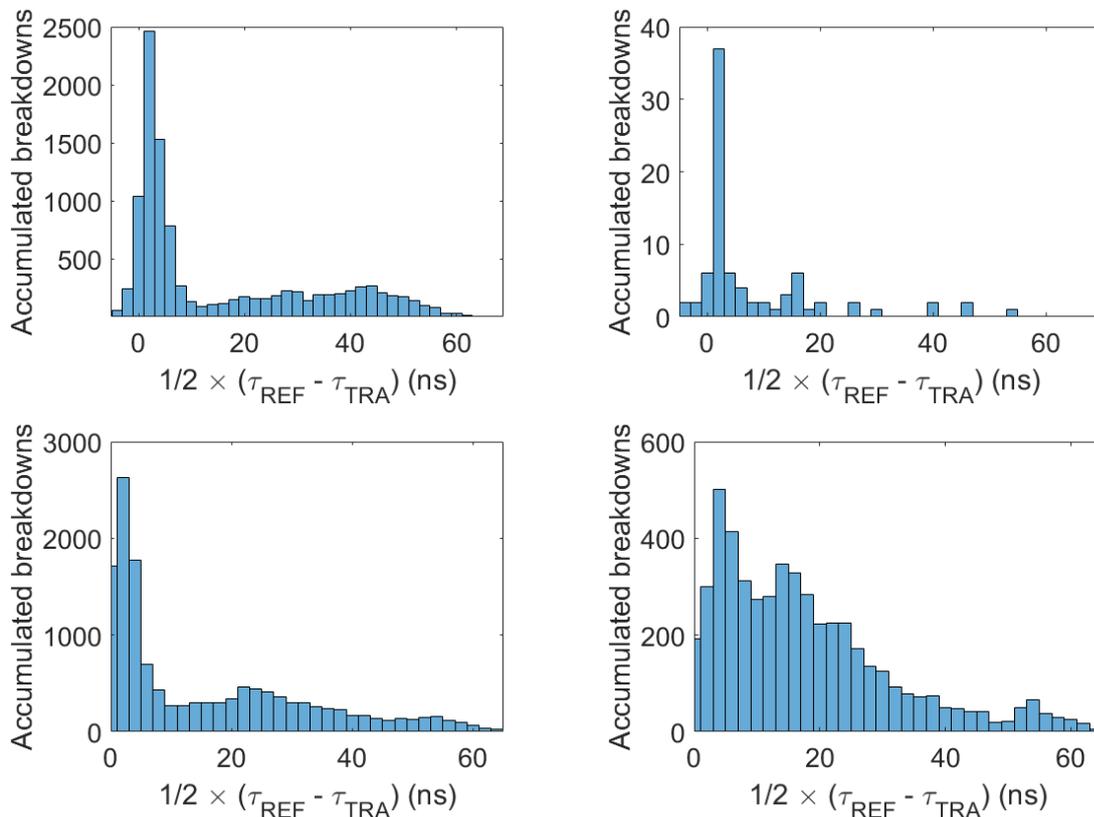


Figure 2.37: Accumulated breakdowns vs. longitudinal position in units of RF signal propagation time in ns. 0 ns represents the structure input and 65 ns represents the structure output. Top left: T24PSIN1 over its entire test in XBox 2. Top right: T24PSIN2 during a flat run at 112 MV/m gradient with 200 ns pulses after conditioning. Bottom left: T24N4 from its installation in XBox 3 until 1 June 2019. Bottom right: T24N5, from installation in XBox 3 until 1 June 2019.

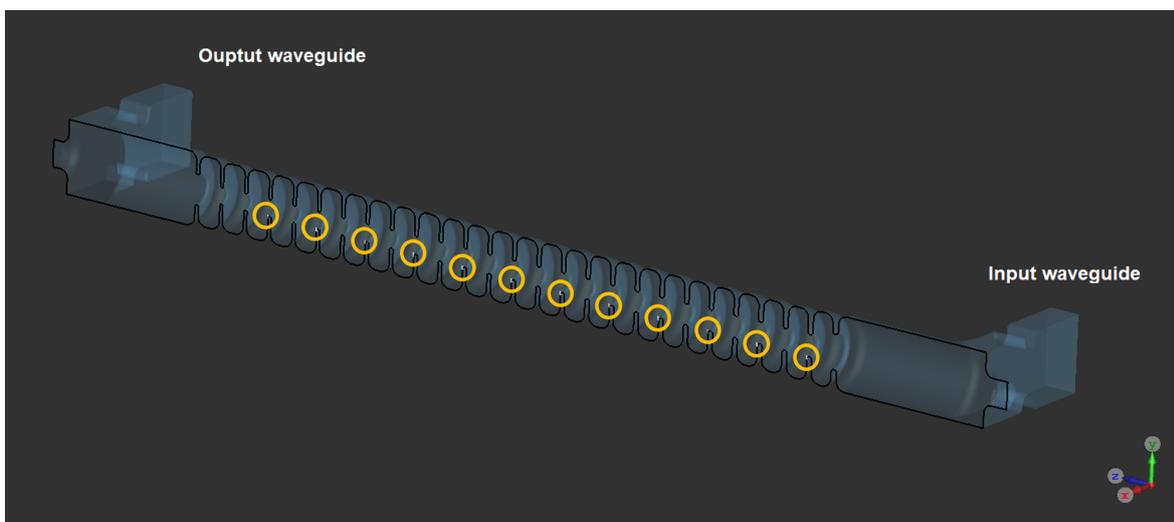


Figure 2.38: A cutaway view of the geometry of the T24 structure with the locations of antennas shown as white lines with yellow circles around them.

the gradient more uniform by compensating for the decrease in power along the length of the structure.

Tapering is important for maximising the maximum gradient of a structure, and also plays a role in long-range transverse-wakefield suppression [74]. Fig. 2.36 demonstrates the effect of tapering, showing the variation of various quantities including E_0 and S_c along the length of the structure, with no beam loading. Despite the decrease in power along the length of the structure, E_0 and S_c are both lowest in the first cell. An effort has been made to optimise the structure using S_c , ensuring that it does not exceed $3.5\text{MW}/\text{mm}^2$ at nominal power with no beam loading. Although S_c scales with the square of the field for a fixed geometry, it is closer to being flat along the length of the structure than E_0 , showing that the structure was designed with S_c rather than electric field as a breakdown limit.

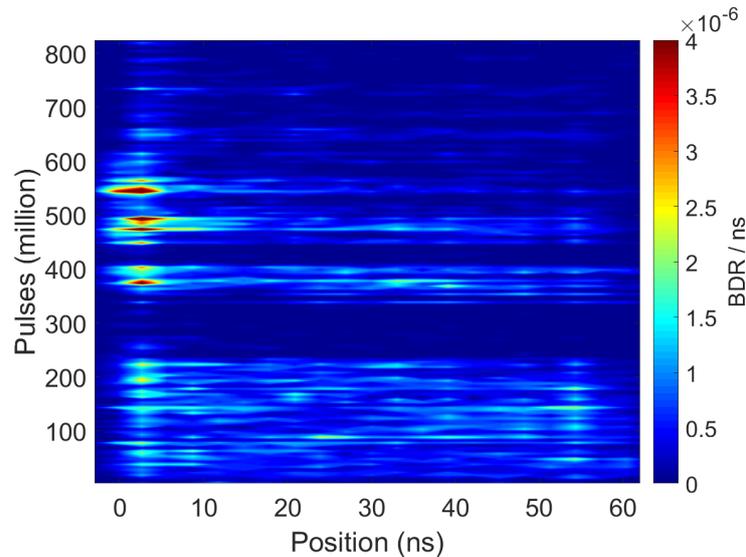


Figure 2.39: Longitudinal breakdown density in the TD26CCR05 structure in breakdowns per pulse per ns vs. longitudinal position in ns and cumulative number of RF pulses applied. Longitudinal positions are presented in units of RF signal propagation time in ns with 0 ns representing the structure input and 65 ns representing the structure output.

Despite the peak values of E_0 and S_c occurring close to the output of the structure, several tests of T24 structures at the XBox test stands have shown a large concentration of breakdowns near to the structure input, where the E and S_c are both the lowest, as shown in Fig. 2.37. This has been noticed in both this structure and other tapered

designs such as the TD26CCR05 [75] and TD24SiC [76] structures. It has been noted that typically, the breakdown locations are spread out evenly throughout the structure in the initial stages of conditioning, then get more concentrated around the structure input as conditioning progresses and reaches higher field levels. An example is shown in Fig. 2.39. An even spread of breakdowns is effectively a design objective, since specific weak spots are undesirable.

One explanation that has been proposed for this is the effect of a breakdown on the incident RF pulse: when a breakdown has fully developed, the cell is filled with plasma and free electrons [63]. This presents an impedance mismatch to the incident signal causing it to be reflected back to the input rather than propagating on to the output. The superposition of the incident and reflected signals is expected to result in a standing-wave in the section of structure between the input and the breakdown, with a peak electric field in the antinodes that is twice as large as the nominal electric field without a breakdown. Thus, these much higher fields may cause one or more secondary breakdowns upstream of the initial breakdown [77]. Assuming this is indeed the case, one should expect the most breakdown damage in the first cell, resulting in many craters with sharp, field enhancing features making it likely for another breakdown to occur there. This hypothesis is, however, contradicted by the stability of phase of the reflected RF signal [78], indicating the reflected signal does not cause secondary breakdowns upstream. An alternative explanation is provided by the E^* model, detailed below.

The exact geometry used in this simulation was that of the T24PSI structures. This design has the same iris geometry and tapering as the baseline T24 design, but both corners of the cell are rounded rather than just one, resulting in cells with back-to-front symmetry, as illustrated in Fig. 2.40.

These structures were used to test the applicability of SwissFEL's manufacturing technology to CLIC, including the use of brazing as an alternative to bonding [79]. Two prototypes of this geometry tested at CERN: T24PSIN1 which reached 55 MW in input

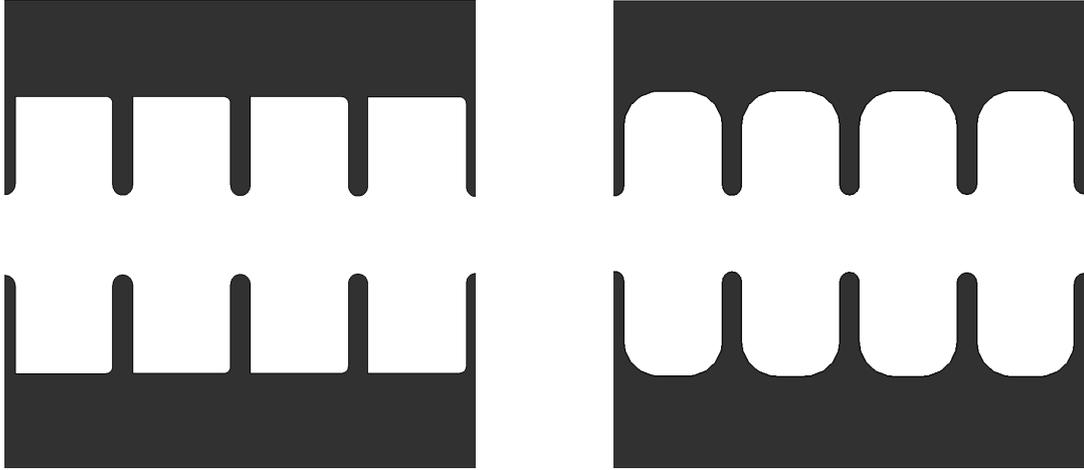


Figure 2.40: Geometry of the middle cells in the baseline T24 geometry (left) and the T24PSI geometry (right).

power and 121 MV/m in unloaded accelerating gradient, and T24PSIN2 which reached 112 MV/m in gradient after a more conservative conditioning approach. These prototypes also showed a clustering of breakdowns at the input of the structure, as shown in Fig. 2.37.

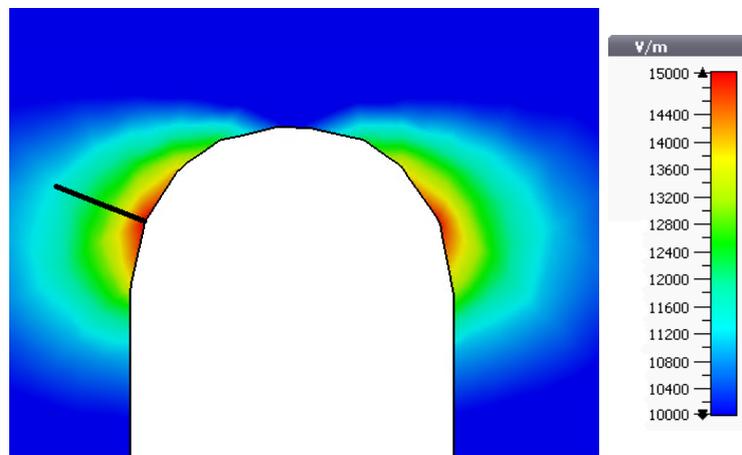


Figure 2.41: Electric field magnitude vs. position on a longitudinal slice of the T24PSI structure. The iris of a cell can be seen and the position of one of the antennas is shown as a black line.

Different cells were investigated to understand the effect of the tapering of the structure. In order to calculate the E^* in this geometry, one antenna was placed in the location of peak E_0 in every other cell (see Fig. 2.38). It was taken as a premise that peak E^* occurs in the location of peak E_0 in a given cell based on discussion in Sec. 2.3.3. The antennas were thus angled away from the apexes of each iris to place them in the

location of peak E_0 , as shown in Fig. 2.41. As with the Crab Cavity (Sec. 2.3.4), antennas of length 0.5 mm were used here.

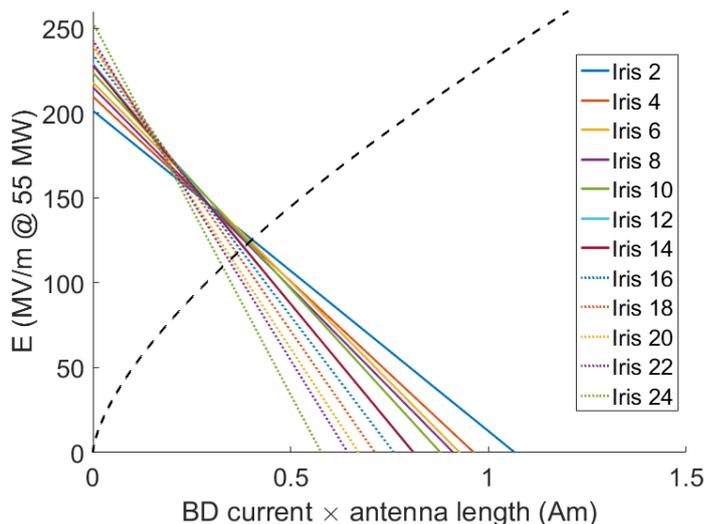


Figure 2.42: Surface electric field in MV/m vs. product of breakdown current and antenna length in Am for the T24 structure at an input power of 55 MW. Coloured lines: loading plots for different cells locations as per Fig. 2.38, with iris 1 representing the iris closest to the RF input. Black dotted line: emission function.

The breakdown loading plot resulting from this study is shown in Fig. 2.42. There is an interesting ‘reversal’ effect: at low values of breakdown current, the electric field is lowest in the cell closest to the input and highest in the cell closest to the output, as expected from the tapering profile shown in Fig. 2.36. However, at a large enough current, the lines cross over and the highest field moves to the input cell, reflecting the breakdown distributions seen in the X-Boxes. Plots of E_0 and E^* vs. position at the maximum tested power of 55 MW are shown in Fig. 2.43, where it can be seen that the peak value of E_0 is located next to the output whereas the peak value of E^* is located next to the input. This ‘reversal’ could be why breakdowns only become concentrated close to the input cell after the structure has reached relatively high power levels after some conditioning. It also implies that an optimal structure should have more tapering than a T24.

The peak E^* and S_c values in the T24PSIN1 structure are given in Table 2.4. Comparing ultimate values with the Crab Cavity, S_c differed by 25.6% between the two

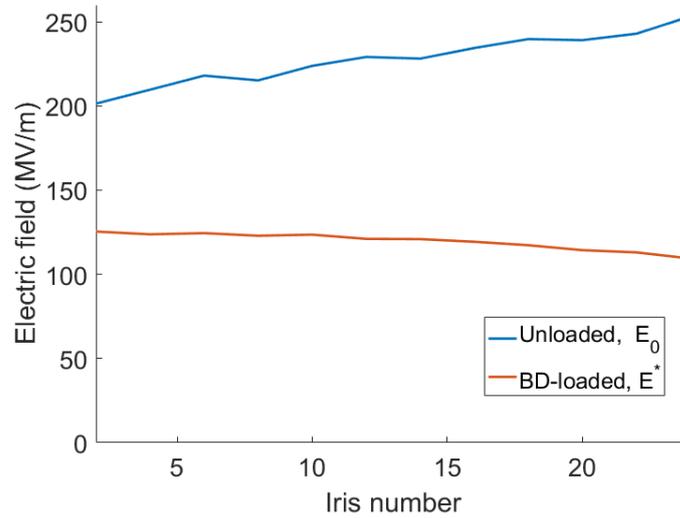


Figure 2.43: Unloaded surface electric field (blue) and E^* (red) in MV/m vs. iris number, at an input power of 55 MW.

Parameter	T24PSIN1	Crab Cavity	Unit
Nominal power	37.5	13.35	MW
Maximum power reached	55	52	MW
Peak S_c at nominal power	3.50	1.71	MW/mm ²
Peak S_c at maximum power	5.30	6.66	MW/mm ²
Peak E^* at nominal power	98.6	60.8	MV/m
Peak E^* at maximum power	115.4	109.5	MV/m

Table 2.4: Comparison of E^* and S_c results for the T24PSIN1 and Crab Cavity.

structures, implying a 12.0% discrepancy in maximum field level. On the other hand, E^* differed by 5.4%, exhibiting greater consistency as an ultimate limit.

2.3.6 S-Band Backward Travelling Wave Structure

The results of the studies of the C3 aperture scan (Sec. 2.3.1) as well as the T24PSI structures (Sec. 2.3.5) strongly suggest that the E^* model should work well for CLIC-like accelerating structures and likely cover the range of parameters over which a realistic design can vary. As the interest in high gradients goes beyond just the CLIC study [80], it would be very valuable to have a universal breakdown criterion applicable to any type of high-field device. The CLIC Crab Cavity discussed in Sec. 2.3.4 has shown one example in which E^* has performed well in a design with significant differences from an accelerating structure. Another design which E^* was tried on is an S-band Backwards Travelling Wave (BTW) accelerating structure for the TULIP medical proton linac [81]. Apart from the difference in operating frequency (2.9985 GHz rather than 11.994 GHz like in CLIC), this structure differs from CLIC structures in some key ways: it has a $7\pi/6$ phase advance per cell for synchronicity with protons with $\beta = 0.38$, and has a very small aperture which results in a low group velocity varying between $0.0038c$ and $0.0021c$. This low group velocity can be used because the nominal beam current is very small, resulting in very low beam loading. The BTW structure also features coupling holes in the septa for magnetic coupling between adjacent cells to compensate for the very low coupling via the beam aperture.

The BTW structure was designed with high gradient technology and design methodology, and was optimised using S_c . In particular, the coupling holes and the nose-cone in the first cell were designed to have the same value of S_c . This is of particular interest because the nose-cone had a large electric and small magnetic field, while the coupling holes had a small electric and a large magnetic field, making it a very good test of the comparative predictive power of S_c and E^* . The structure was tested in SBox, an S-band equivalent to the X-band test stands at CERN [82].

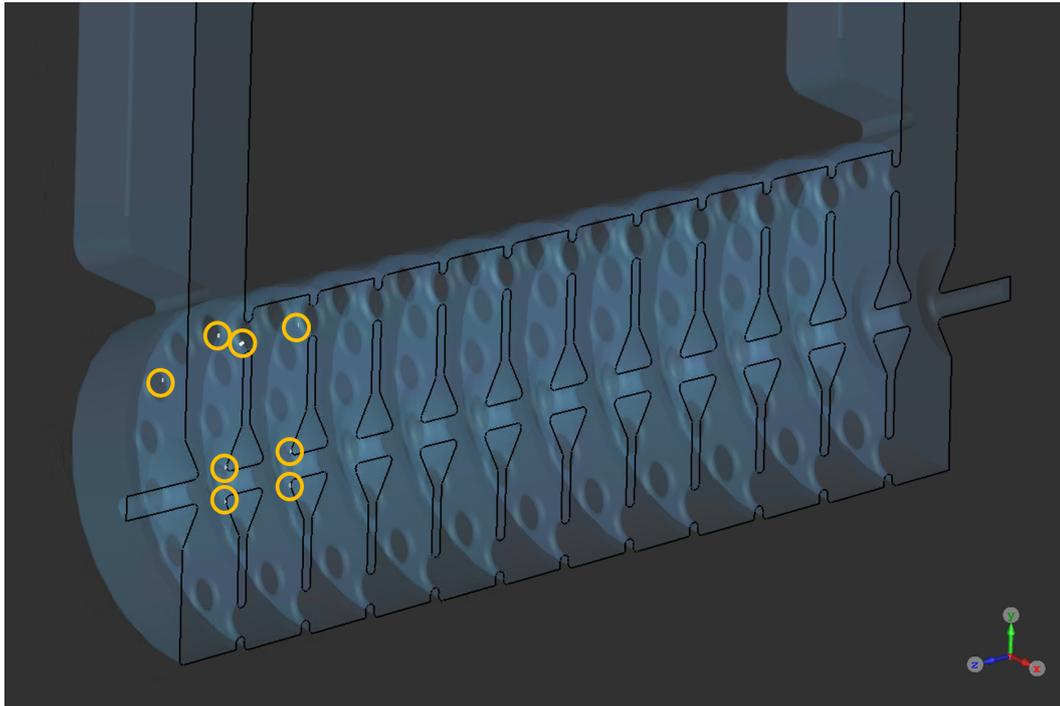


Figure 2.44: A cutaway view of the geometry of the BTW structure with the locations of antennas shown as white lines with yellow circles around them.

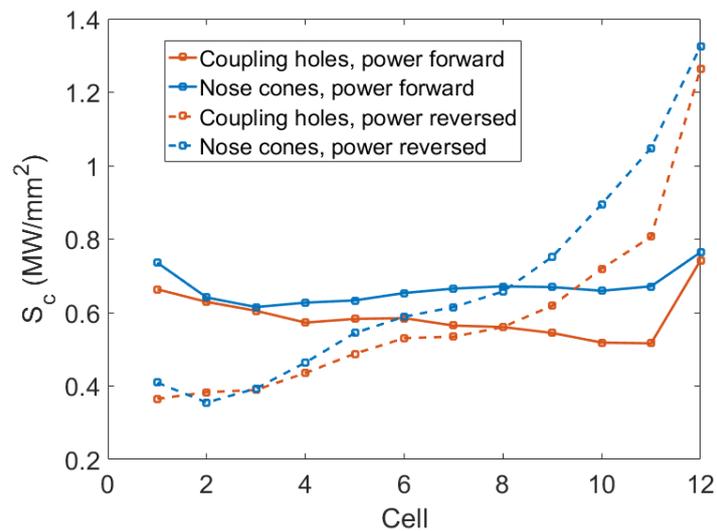


Figure 2.45: Peak S_c in MW/mm² vs. cell number in the BTW structure, at an input power of 20.16 MW. Red solid curve: in the coupling holes with power flowing in the forward direction. Blue solid curve: on the nose cones with power flowing in the forward direction. Red dashed curve: in the coupling holes with power flowing in the reverse direction. Blue dashed curve: on the nose cones with power flowing in the reverse direction.

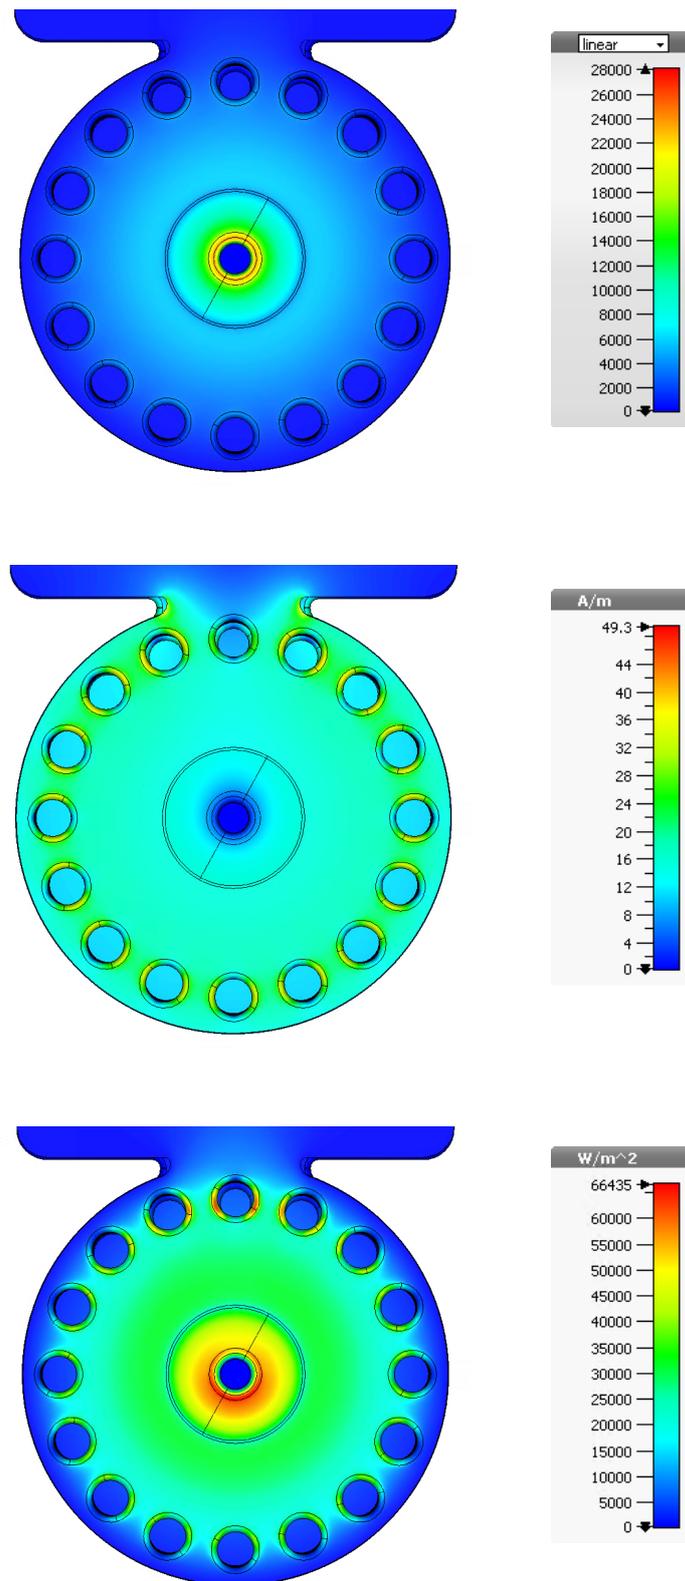


Figure 2.46: Field vs. transverse position for the output coupler cell of the BTW structure, when RF power is injected in the reverse direction. Top: electric field magnitude. Middle: magnetic field magnitude. Bottom: S_c .

Due to an error in mounting the structure for high-power testing, the structure was tested with power flowing in the opposite direction to that designed [83]. Due to the tapering of the structure's shunt impedance, this resulted in a very different field profile from the one intended, as shown in Fig. 2.45. However, the experimental results are still useful for comparing the S_c and E^* models. The test was ended because the maximum power that could be provided to the structure, 33 MW, was reached.

Fig. 2.46 shows the E , H , and S_c field distributions in the output cell, which experienced the highest power flow and field levels due to the reversed orientation of the structure. In this cell, the field distributions were not radially symmetric due to the presence of the input waveguide. The S_c was highest on the bottom of the nose-cone and is also very high on the top coupling hole, closest to the waveguide. Here, 'top' and 'bottom' refer to positions as shown in the figure. For 1 W of incident power, the peak S_c was 62.7 kW/m² in the top coupling hole, 54.0 kW/m² on the top edge of the nose cone, and 65.7 kW/m² on the bottom edge of the nose cone. In other words, the peak S_c at the top coupling hole was between that on the top of the nose-cone and that on the bottom of the nose-cone. Images from an endoscopy of the output cell, performed after the test at SBox was concluded, are shown in Fig. 2.47. Extensive breakdown damage is visible on both the top and bottom edges of the nose-cone, yet the surface of the top coupling hole is almost pristine despite the large value of S_c in this location, showing another inconsistency in which the number of breakdown craters does not appear to increase with S_c as expected.

For the E^* calculation for this structure, antennas were placed in the output cell and the last regular cell. In each of the two cells, antennas were placed on the top and bottom of the nose cone, as well as the location of peak S_c on some of the coupling holes. The antenna length was 0.5 mm. The results were assumed to scale with frequency in the manner shown in (2.33) to make them comparable to the X-band structures studied earlier. As the operating frequency of this structure is exactly a quarter of that of the X-band structures, the R_{bd} values were multiplied by 16. Another simulation was run in which the structure was installed in the correct direction, with the same arrangement

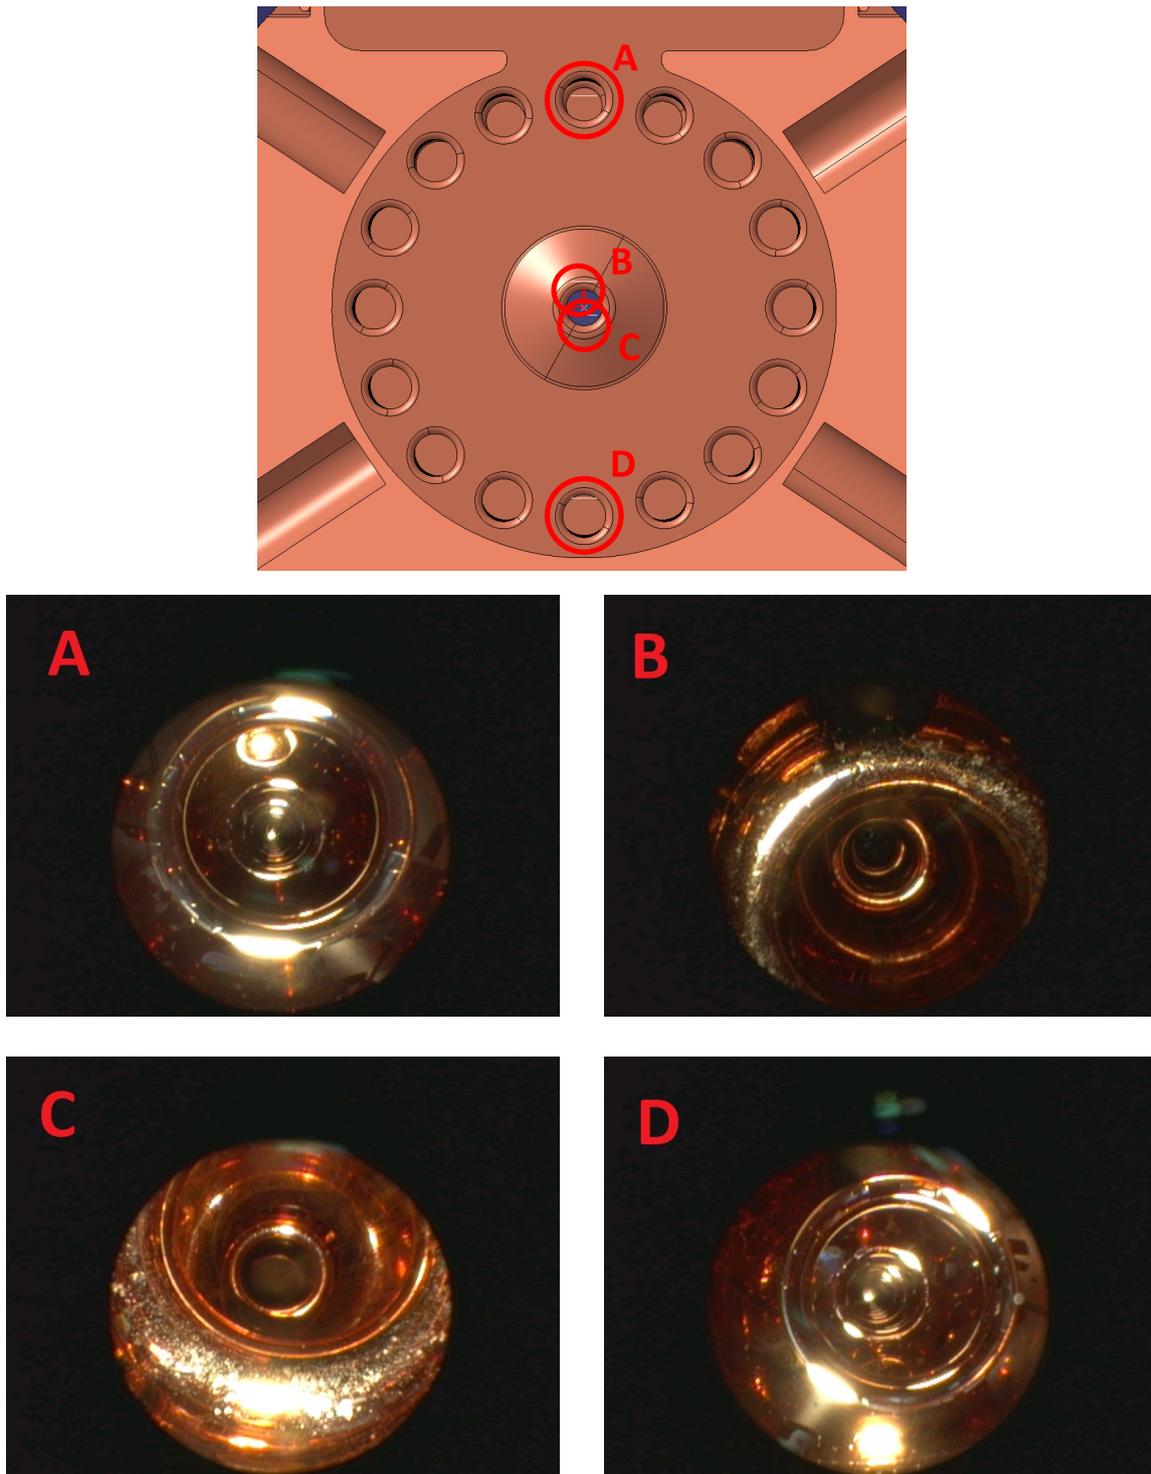


Figure 2.47: Top: locations in red circles where endoscopy images were taken in relation to the geometry of the cell. Middle left: endoscopy image of the top coupling hole, which a maximum S_c of 62.7 kW/m^2 , showing almost no breakdowns. Middle right: endoscopy image of the top of the nose-cone, with a maximum S_c of 54.0 kW/m^2 , showing some breakdown damage. Bottom left: endoscopy image of the bottom of the nose-cone, with a maximum S_c of 65.7 kW/m^2 , showing extensive breakdown damage. Bottom right: endoscopy image of the bottom coupling hole, with a maximum S_c of 41.7 kW/m^2 , showing almost no breakdowns. S_c values are given for 1 W of incident RF power [39].

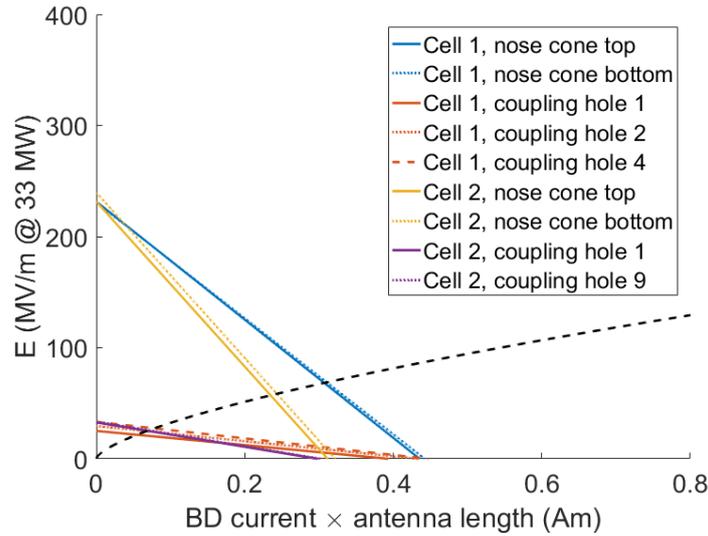


Figure 2.48: Surface electric field in MV/m vs. the product of breakdown current and antenna length in Am, in the BTW structure at an incident power of 33 MW in the forward direction. Coloured lines for different antenna locations in the first two cells as indicated on the legend. Black dashed line: emission function. ‘Top’ and ‘bottom’ refer to location as shown in Fig. 2.46. The coupling holes have been numbered starting from the top hole and increasing going anticlockwise around the cell.

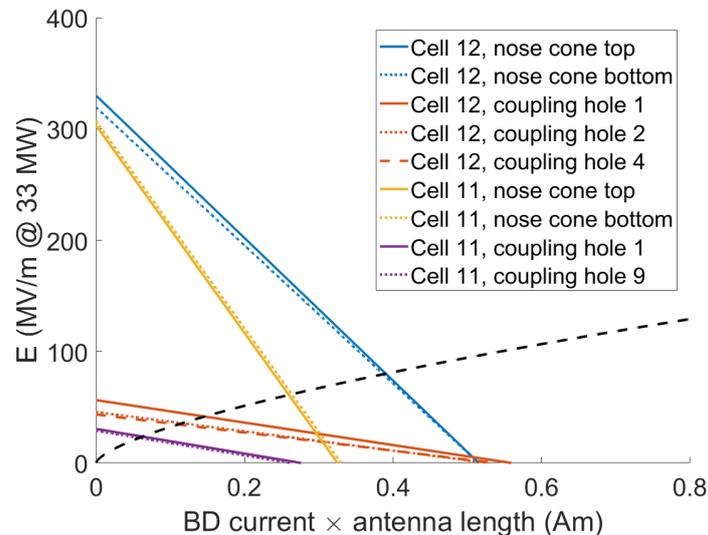


Figure 2.49: Surface electric field in MV/m vs. the product of breakdown current and antenna length in Am, in the BTW structure at an incident power of 33 MW in the reverse direction. Coloured lines for different antenna locations in the last two cells as indicated on the legend. Black dashed line: emission function. ‘Top’ and ‘bottom’ refer to location as shown in Fig. 2.46. The coupling holes have been numbered starting from the top hole and increasing going anticlockwise around the cell.

Parameter	T24PSIN1 (tested)	BTW (hypothetical)	Unit
Power	55	138	MW
Peak S_c	5.3	5.13	MW/mm ²
Peak E^*	115	115	MV/m

Table 2.5: A comparison of the peak E^* , S_c and power for the T24PSIN1 and BTW structures. The values for the T24PSIN1 correspond to the highest power reached during testing, whereas those for the BTW have been scaled to obtain the same value of E^* as the T24PSIN1.

of antennas but in the input cell and the first regular cell. Breakdown loading plots representing the results of the two cases are shown in Figs. 2.48 and 2.49.

The E^* model gives a result that is more consistent with the endoscopy images than S_c . Since E_0 was much lower on the coupling holes than on the nose cone in both cells, the resulting E^* was lower as well, explaining the lack of breakdown craters on the coupling holes. The maximum E^* on the nose-cone of the output cell was 79.8 MV/m, whereas the maximum value on the coupling holes was 41.0 MV/m, which is about half the value on the nose-cone.

As the ultimate breakdown-limited gradient of the structure could not be reached due to limitations in the maximum power available from the test stand, there is no experimental data on the limiting value of the structure. However, one can still make a comparison between the predicted maximum gradient as per S_c and E^* . Table 2.5 shows a comparison between the T24PSIN1 and BTW structures. The incident power given for the BTW structure is a hypothetical value which would be needed to achieve the same peak E^* as that reached in the T24PSIN1 during testing. At this power, the peak S_c in the BTW structure would have a value of 5.13 MW/mm², which only differs from the peak value in the T24PSIN1 by 3.3%. Thus, the predictions of maximum achievable gradient in the BTW structure according to S_c and E^* only differ by 1.6%, showing good consistency between the two when applied to the BTW structure.

This may not be an accurate prediction of the maximum power due to the BTW structure's very low group velocity. The tendency for the single-frequency E^* calculation to over-estimate the achievable gradient of low-group-velocity structures has been shown

k value	$k_c/3$	k_c	$3k_c$
Max. E^* in T24PSIN1 (MV/m)	161.9	115.4	72.3
Max. E^* in Crab Cavity (MV/m)	139.2	109.5	74.5
Discrepancy in E^* (%)	16.3	5.4	3.0

Table 2.6: A comparison of the peak E^* attained in the T24PSIN1 and CLIC Crab Cavity obtained using the single-frequency model with different values of k , as well as the discrepancy between the peak values. k_c is the final chosen value of $5.4 \times 10^{-9} \text{ Am}^{3/2}\text{V}^{-3/2}$. Larger values of k result in better agreement between the E^* values.

and discussed in Sec. 2.3.1, and is likely applicable here as well. Given the extremely low group velocity of the 12th cell of the BTW structure, $0.0021c$, one can expect a significant contribution from the stored energy of the cell, causing the true breakdown-limited maximum power to be much lower than 138 MW. Once a future experiment can conclusively test the maximum power of this geometry, more information about the frequency scaling of the absolute value of E^* can be deduced.

2.3.7 Choice of k Value

The value of k , the proportionality constant in the characteristic emission function that represents the current-field characteristic of a nascent breakdown in a well-conditioned high-gradient structure, was not determined *a priori* and had to be chosen such that the E^* results gave the best agreement with experimental results. It was found that several aspects of the E^* model depended on the exact value of k . These were the consistency in the maximum E^* reached in tested X-band structures as well as consistency between the spatial location of the peak value of E^* and breakdown locations observed in high-gradient structure tests. Due to the qualitative nature of some of these criteria, it was difficult to algorithmically find the optimum value of k . Instead, a value was manually chosen that was deemed to give a good compromise between all the criteria. For each of the criteria mentioned, examples of results obtained using different values of k will be shown.

The first criterion to be discussed is consistency between maximum E^* values in tested X-band structures. The T24PSI and CLIC Crab Cavity were a useful comparison

because they have significantly different field patterns and parameters, yet operated at the same frequency, were made out of the same material, and were tested in the same test stand. One of the starting assumptions of the E^* model is that the breakdown rate should depend only on the peak value of E^* , meaning that the maximum E^* should agree between the two structures. Table 2.6 shows the discrepancy in maximum E^* between the two structures for different values of k , demonstrating that larger values of k result in better agreement.

Another important criterion is that the spatial distribution of E^* should match typical spatial distributions of breakdown craters seen in experiments. For example, cells operating in the TM_{010} mode, usually have a concentration of breakdown craters just off the apex of the iris, as shown in Fig. 2.21. This means that within such a cell, the E^* should have a maximum close to this location. Fig. 2.50 shows the spatial distribution of E^* along the surface of the middle cell in the C3 structures discussed in Sec. 2.3.1 for different values of k . In this geometry, E^* should have a peak around the 1.6 mm mark to be consistent with experimental results. It can be seen that for the smallest value of k , each aperture size does indeed exhibit a peak at this location. For larger values of k , however, the results for small aperture sizes instead exhibit local minima in E^* at this position. This would correspond to that area being clear of breakdown craters, which is completely at odds with experiments. It should be noted, however, that the inclusion of transient behaviour, as in Sec. 2.4.2, does appear to prevent the formation of these local minima, which means that this is no longer a constraint on the choice of k if the broadband quantity is used.

Along from the location of peak E^* within any given accelerating cell, the position of the peak across multiple cells in a tapered structure can also change with the choice of k . One of the key arguments for using E^* rather than unloaded field quantities such as S_c in was the longitudinal breakdown locations in T24 structures as discussed in Sec. 2.3.5. Tests of several T24 structures showed a strong clustering of most of the breakdowns close to the RF input, despite neither the E_0 nor the S_c being the greatest there. E^* , on the other hand, was shown to be the largest in the cell closest to the

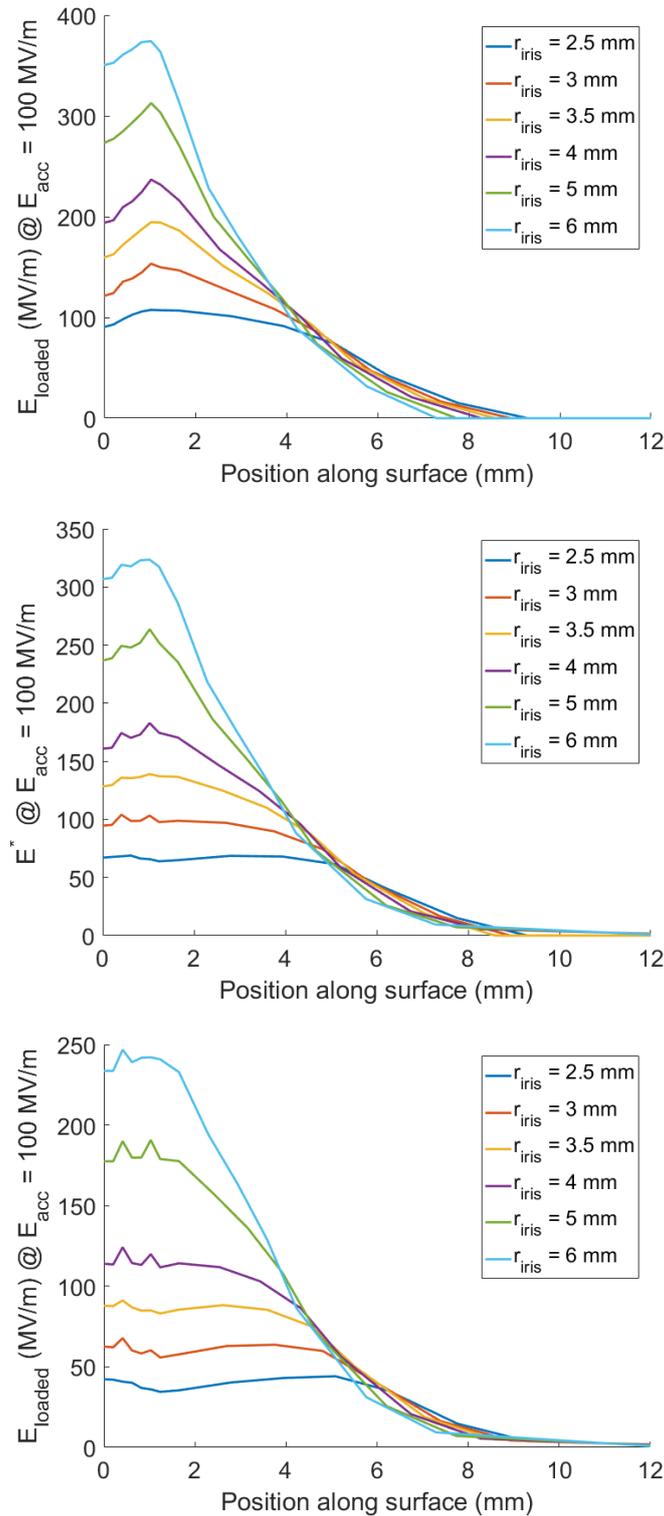


Figure 2.50: E^* in MV/m vs. position along the surface in mm. The position is defined by the arrow shown in Fig. 2.15, where 0 is the iris apex. Different coloured plots represent C3 structures of different aperture sized as indicated on the legend. Top: using $k = k_c/3$. Middle: using $k = k_c$. Bottom: using $k = 3k_c$. It can be seen that for larger values of k , the distribution of E^* for small aperture sizes exhibits a local minimum around the 1.6 mm mark, the location where breakdowns most frequently occur in experiments.

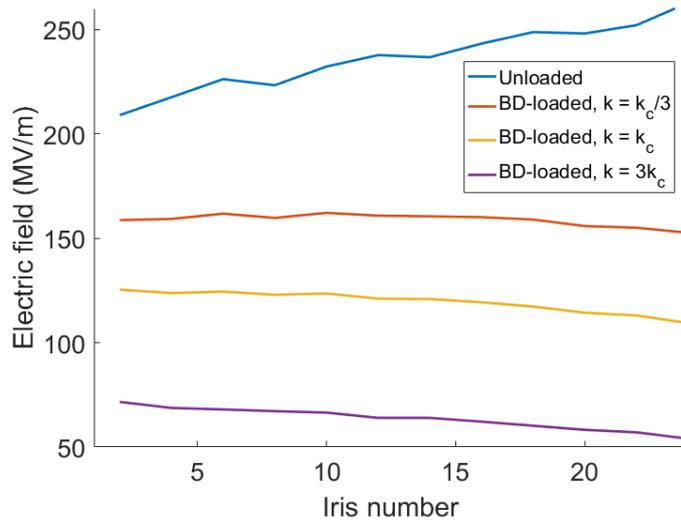


Figure 2.51: Loaded and unloaded electric field values in the T24PSI structure in MV/m vs. iris number, at an input power of 55 MW. Blue: unloaded surface electric field. Red: E^* using $k = k_c/3$. Yellow: E^* using $k = k_c$. Purple: E^* using $k = 3k_c$. In the case of $k = k_c/3$, the peak value of E^* is no longer located at the first iris.

RF input, in line with experimental results. However, the longitudinal distribution of E^* in the T24 structure does also depend on the exact value of k chosen, as shown in Fig. 2.51. It can be seen in this figure that if k is too small, the peak value moves downstream, away from the input cell. This effect does not depend on k as strongly as, for example, the position of peak E^* within the cells of the C3 structures, as in Fig. 2.50. Thus, it did not constrain the choice of k as strongly as the other factors listed above.

2.4 Improvements of the Method

The method used to obtain the results discussed thus far in Sec. 2.3 represents the simplest possible implementation of a breakdown quantity based around the idea of breakdown-loaded electric field. Some ways of modifying the method to obtain a more physically accurate model have been identified and investigated. The results of these investigations are presented in this section.

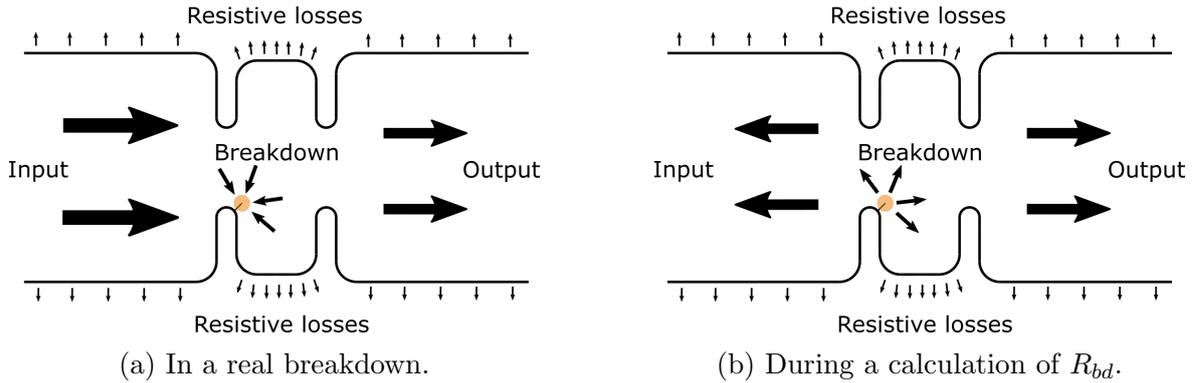


Figure 2.52: Direction of power flow during a breakdown.

2.4.1 Use of Z-Parameters

In all calculations of R_{bd} discussed so far, an antenna was placed at a potential location for a breakdown site. In the simulation, the antenna was set up to be a current source with a current of unit amplitude at the operating frequency of the structure. All the waveguide ports of the structure were connected to matched loads in the simulation. Thus, the breakdown was a source of power and the structure input was a sink. While this is still a valid calculation in principle as Maxwell's equations are time-reversal invariant [68], the situation being modelled does not exactly correspond to a breakdown in reversed time. The two situations are represented as a diagram in Fig. 2.52. What is normally the output of the structure remains a load in simulation, along with resistive losses in the copper walls of the structure. It is believed that this method has still given satisfactory results because:

- the cells in the structures being studied here are roughly symmetric in their coupling to the input and output ports, meaning that a reversal in the direction of power flow from the output port only scales the result by a constant factor, meaning that the values relative to other cells and structures are preserved.
- cells in travelling-wave structures are typically overcoupled, meaning that the resistive losses are small compared to the total power flowing through them. Thus, a reversal in the sign of this term will only cause a small error.

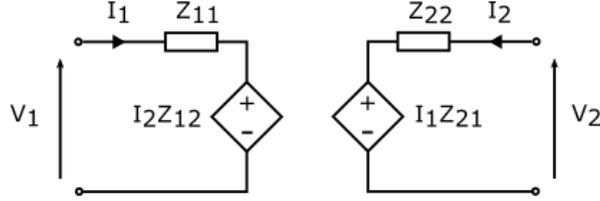


Figure 2.53: Z-parameter description of a linear two-port network.

Despite the method described above giving satisfactory results, a more physically accurate model is still likely to give better results and be less reliant on assumptions. One way to do this is to consider the matrix of impedance parameters \mathbf{Z} which describes the coupling of each port of the system to every other port. It is related to the scattering matrix \mathbf{S} as follows [68]:

$$\mathbf{Z} = \sqrt{\mathbf{Z}_0}(\mathbf{I} + \mathbf{S})(\mathbf{I} - \mathbf{S})^{-1}\sqrt{\mathbf{Z}_0}, \quad (2.38)$$

where \mathbf{I} is the identity matrix, and $\sqrt{\mathbf{Z}_0}$ is a diagonal matrix in which each element is the square root of the characteristic impedance at the corresponding port. In this case, the input and output ports, as well as each of the breakdown antennas included in the simulation, are each considered ports in this matrix. The \mathbf{Z} -matrix may be used to obtain an equivalent series resistance, analogous to the R_{bd} calculated earlier, defining the coupling between the input port and the antenna in question. For the case of a two-port network, represented schematically in Fig. 2.53, the equivalent series impedance Z_{eq} is given by:

$$Z_{eq} = Z_{22} - \frac{Z_{21}Z_{12}}{Z_0 + Z_{11}}, \quad (2.39)$$

where Z_{nm} is the element of \mathbf{Z} in the n^{th} row and m^{th} column, and Z_0 is the output impedance of the source connected to port 1. The advantage of this method is that it explicitly considers only the power flowing from the input port to the antenna. Resistive losses and other ports are correctly modelled as loads.

2.4.2 Finite-Time Broadband Quantity

A further improvement that could be implemented to the calculation of E^* is the inclusion of the effect of stored energy, which has been identified as the source of discrepancy between E^* calculation and experimental results in Sec. 2.3.1. The fact that group velocity plays an important role in determining the gradient that a structure can reach implies that transient phenomena are inherently important, and thus that the bandwidth of the structure should be considered. This can be seen from the way the complex S-parameters of an RF network contain information about the energy stored within it. For example, the Q of a resonator, which defines the bandwidth of its frequency response as $Q = \omega_0/\Delta\omega$, also allows the stored energy to be calculated as $Q = \omega_0 E_{stored}/P_{loss}$ [68]. The discharging of the stored energy in an RF cavity is a transient phenomenon, meaning that it involves a finite bandwidth rather than a single frequency point. Thus, an equivalent to R_{bd} which takes into account the coupling over the full bandwidth of the structure might result in an E^* that is sensitive to the stored energy, and thus more applicable to low group velocity and standing-wave structures, as well as making a bridge between RF and DC.

Fig. 2.54 shows the real part of Z_{eq} , denoted R_{eq} , calculated by obtaining a \mathbf{Z} matrix from a numerical simulation of the C3 structure geometry described in Sec. 2.3.1, then using the values of this matrix in (2.39). As with R_{bd} , R_{eq} has greater values for smaller aperture sizes, implying the ability to sustain a higher unloaded electric field before breakdown. The bandwidth of R_{eq} also appears to increase with aperture size, which is consistent with the increase of bandwidth of a structure as the coupling is increased. R_{eq} appears to be the largest at the edges of the passband of the structure, which is consistent with the group velocity approaching zero at these frequencies. It tends to zero far away from the operating frequency of 12 GHz, seemingly implying infinite power coupling to the breakdown at those frequencies. However, one should keep in mind that in the power delivered to the breakdown depends on not only the equivalent series impedance Z_{eq} but also the equivalent open-circuit voltage, V_{eq} , obtainable from the \mathbf{Z} -matrix as:

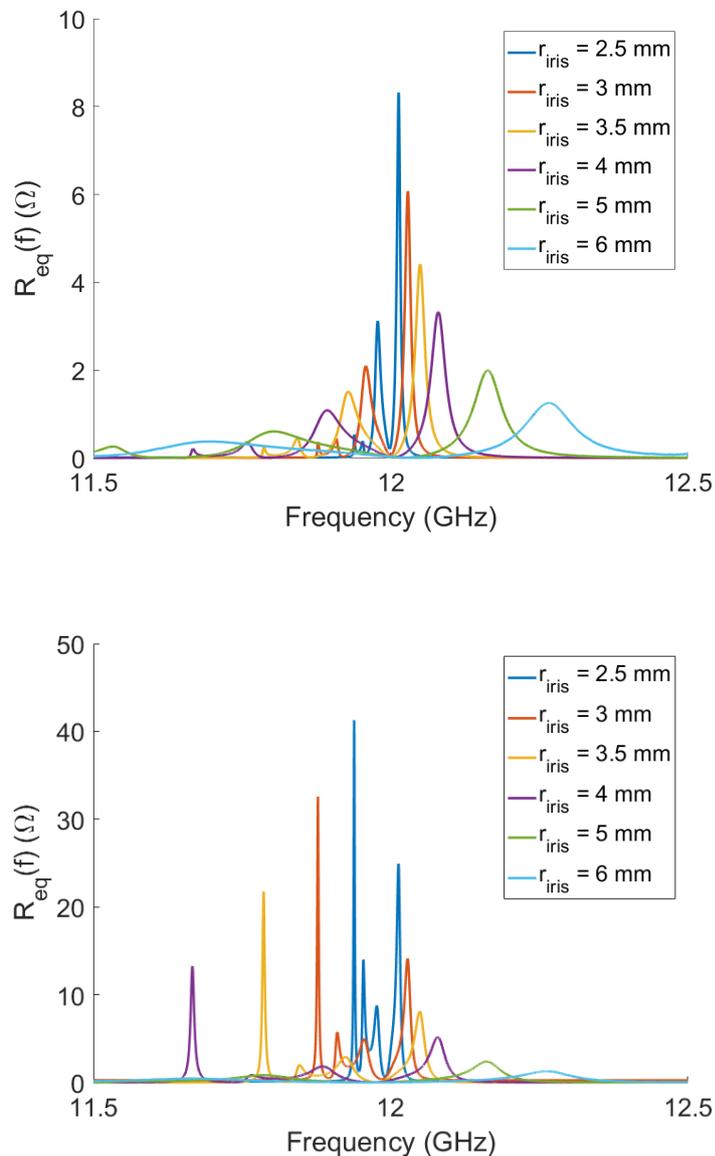


Figure 2.54: R_{eq} in Ω vs. RF frequency in GHz in the C3 structures. Different coloured plots represent C3 structures of different aperture sizes as indicated in the legend. Due to the $2\pi/3$ phase advance of the structure, the operating frequency of the structure is offset from the centre of its bandwidth, leading to asymmetry with respect to frequency in these plots. Top: At the iris apex. Bottom: At the location of peak surface electric field, a few millimetres from the iris apex. See Fig. 2.14 for the electric field distribution in the cell.

$$V_{eq} = V_1 \frac{Z_{21}}{Z_{11}} \quad (2.40)$$

In this work, V_{eq} was calculated directly by multiplying the unloaded surface electric field obtained from a finite-element simulation by the antenna length. Outside the operating bandwidth of a typical accelerating structure, the surface electric field for a given input power is very low compared to that within the bandwidth. (This is equivalent to the value of Z_{21}/Z_{11} being very low outside the operating bandwidth.) Thus, despite R_{eq} being low, very little power can be coupled to the breakdown at these frequencies. A future improvement to the method might be to take into account the variation of V_{eq} with frequency.

Considering a range of frequencies rather than just a single frequency point means that transients are now important. The problem cannot be easily solved in the frequency domain as the breakdown loading is nonlinear. Thus, a completely correct solution would require a numerical time-domain solution that takes into account both the frequency-dependent source impedance and the nonlinear load. While such a solution is possible, it was hoped that a much simpler alternative would be available by making an assumption about the transient response of this system. The emitted current as a function of time was assumed to have the same shape in every case, with the only possible difference between cases being scaling by a constant factor. This would mean that the same procedure for calculating E^* could be used as before, as long as the values of R_{eq} and the emission function were modified to take into account the assumed functional form of the emitted current.

If the emission of current during a breakdown is assumed to occur in very short, intense impulses that are much shorter in time than the response time of the structure, they can be approximated with the use of the Dirac delta function $\delta(t)$:

$$i(t) = Q\delta(t), \quad (2.41)$$

where $i(t)$ is the emitted current as a function of time and Q is the total charge emitted in such an impulse. This would mean that the voltage $v(t)$ across the antenna, which would appear as a result of the structure's response to this emission of current, would be given by:

$$v(t) = i(t) * r_{eq}(t) = Qr_{eq}(t), \quad (2.42)$$

where $r_{eq}(t) = FT^{-1}[R_{eq}(f)]$, with FT^{-1} representing the inverse Fourier transform, is the impulse response of the system as seen by this antenna. This would mean that the total energy U associated with this impulse is given by:

$$U = \int_{-\infty}^{\infty} i(t)v(t)dt = Q^2 r_{eq}(0) = Q^2 \int_{-\infty}^{\infty} R_{eq}(f)df \quad (2.43)$$

Thus, the quantity:

$$K = \int_0^{\infty} R_{eq}(f)df \quad (2.44)$$

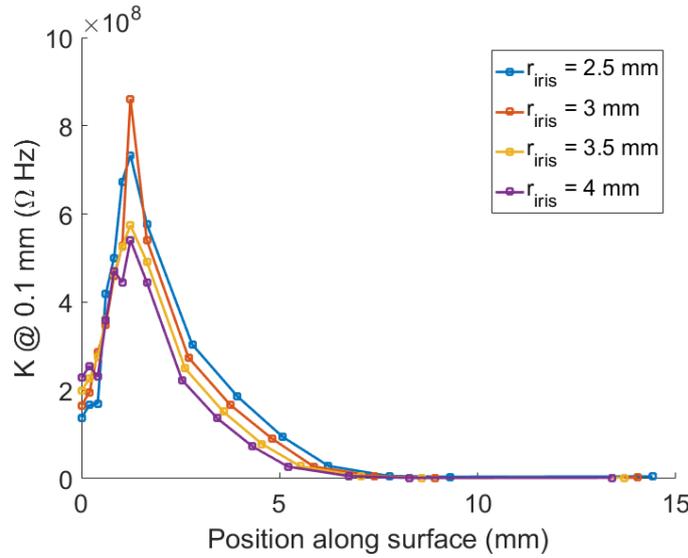


Figure 2.55: K vs. position along the surface in mm. The position is defined by the arrow shown in Fig. 2.15, where 0 is the iris apex. Different coloured plots represent C3 structures of different aperture sized as indicated on the legend.

would appear to be a good candidate for a broadband equivalent for R_{bd} with the assumption that the current is emitted over a very short period. (N.B.: As the structure being simulated is a physical system in which electric field and current take on real values only, $R_{eq}(f)$ has even symmetry around zero frequency. Thus, the integral only needs to be taken over the limits 0 and ∞ , rather than $-\infty$ and ∞)

However, as shown in Fig. 2.55, this does not appear to give results in line with expectations. For example, the peak value of K for the 3 mm iris radius is greater than the peak K for the 2.5 mm iris radius. This does not reflect the behaviour seen in Fig. 2.54, where, for each aperture size simulated, the peak value of R_{eq} over frequency decreased monotonically as the aperture size was increased. The peak in K does not seem to diminish with increasing aperture size, a behaviour seen with S_c in Fig. 2.22.

On the other hand, an integral of the square of R_{eq} , denoted L , appears to give much better results:

$$L = \sqrt{\int_0^{\infty} R_{eq}(f)^2 df} \quad (2.45)$$

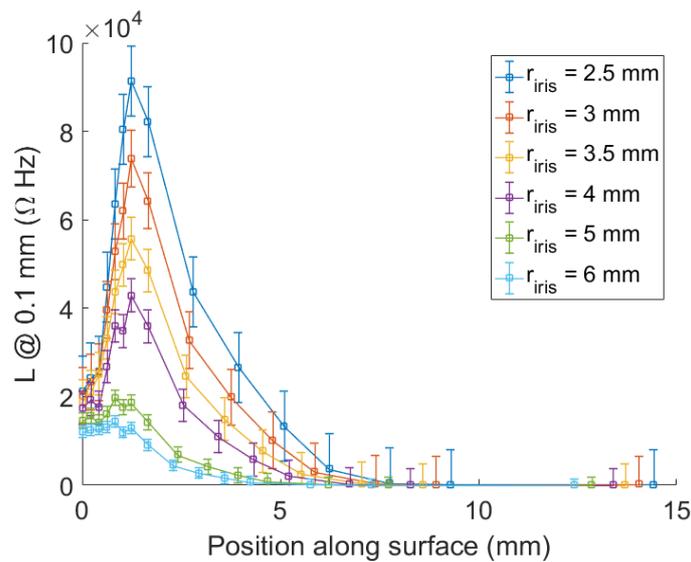


Figure 2.56: L vs. position along the surface in mm. The position is defined by the arrow shown in Fig. 2.15, where 0 is the iris apex. Different coloured plots represent C3 structures of different aperture sized as indicated on the legend.

Plots of L for various C3 geometries calculated using (2.45) are shown in Fig. 2.56. L appears to decrease monotonically with aperture size, corresponding to a lower maximum unloaded electric field, as is expected. The distribution of L as a function of position on the surface also becomes flatter as the aperture size is increased, like S_c . As was the case with R_{bd} , the value of L is greatest at the location of peak E_0 .

A physical explanation consistent with the use of L is as follows: not all of the power absorbed in the breakdown is done so by the emission site. In fact, one of the assumptions made for this model in Sec. 2.2 is that the antenna represents a stream of emitted charged particles. Power must also be provided to the emission site itself for it to heat up enough to undergo runaway. The main mechanism for this is thought to be resistive heating of the site due to the current flowing through it. However, not all of the RF power absorbed by the breakdown contributes to the heating of the breakdown site. Some of the power absorbed causes an increase in the kinetic energy of the emitted electrons, which either escape the structure or collide with the walls of the structure relatively far from the emission site. The power actually causing heating of the breakdown site is proportional to the square of the instantaneous voltage $v(t)$ across this resistance. In this case, the energy delivered to just the emission site for an impulse excitation is given by:

$$U = \int_{-\infty}^{\infty} \frac{v^2(t)}{R_l} dt = \frac{Q^2}{R_l} \int_{-\infty}^{\infty} r_{eq}^2(t) dt = \frac{Q^2}{R_l} \int_{-\infty}^{\infty} R_{eq}^2(f) df, \quad (2.46)$$

where R_l is the resistance causing the heating of the emission site. Parseval's theorem was used to move between the time and frequency domains for convenience [84]. While the value of R_l is unknown, it can be absorbed by the proportionality constant k of the emission function, as long as all candidate breakdown sites have similar properties. The square root of the integral is taken for consistency in dimensions.

Fig. 2.57 shows the correlation of the new quantity L with the local surface electric field E_0 . One can see that this quantity is no longer a simple quadratic function of E_0 , as was the case with R_{bd} , shown in Fig. 2.27. The smaller the aperture size, the

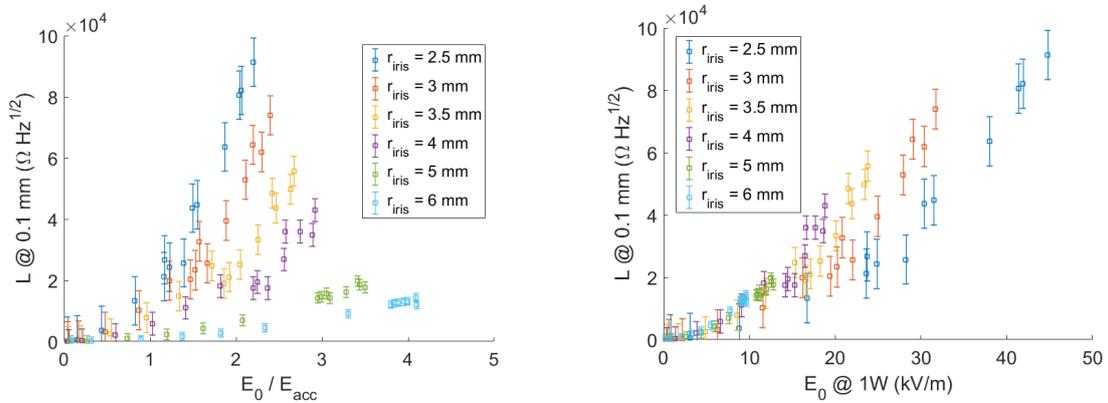


Figure 2.57: L vs. unloaded surface electric field E_0 . Different coloured points represent C3 structures of different aperture sizes, as indicated on the legend. Left: E_0 values given for unit accelerating gradient. Right: E_0 values given for 1 W input power.

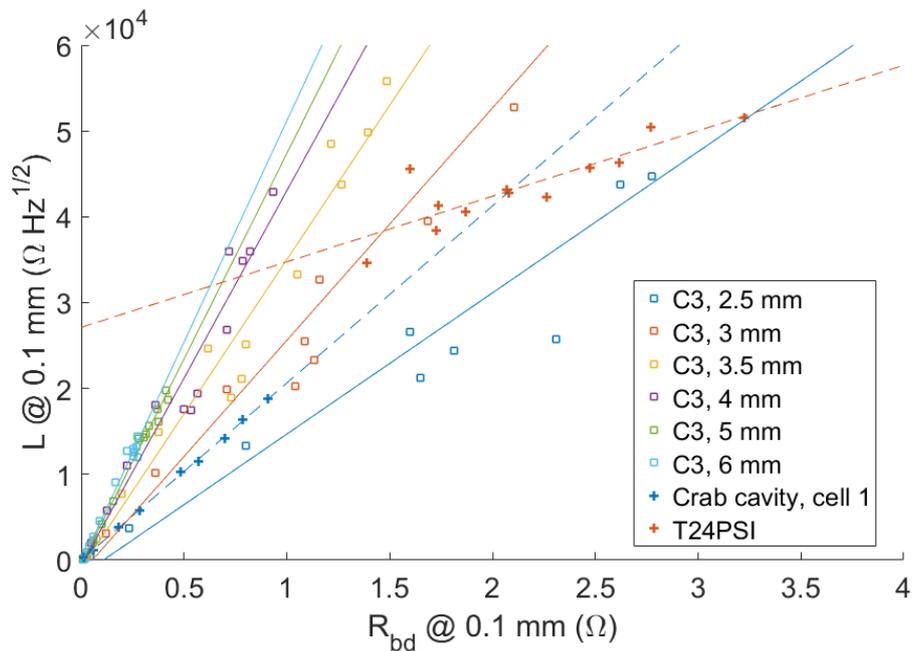


Figure 2.58: L vs. R_{bd} for C3 structures of different aperture sizes, as indicated on the legend. Results for the T24PSI and the first cell of the CLIC Crab Cavity, normalised to an antenna length 0.1 mm, are also shown. Solid lines represent linear fits to the C3 structure data, whereas dashed lines represent linear fits to the data from the Crab Cavity and T24PSI. Error bars have been omitted for clarity.

greater the deviation from the quadratic dependence $L \propto E_0^2$. This can be seen more clearly in Fig. 2.58, where L is plotted against R_{bd} for the each location and aperture size. Here it can be seen that for each aperture size there is an approximately linear relationship between L and R_{bd} . The proportionality constant appears to decrease as the aperture size becomes smaller, corresponding to lower group velocity. The value of L becomes smaller than expected, implying lower breakdown fields. This is good because R_{bd} predicted gradients which were unrealistically large for low group velocities, as discussed in Sec. 2.3.1. The slope on which the points in Fig. 2.58 lie can be interpreted as representing the relative contributions of the external power flow and local stored energy to the breakdown, with a steep slope, as with the 4, 5, and 6 mm aperture sizes for the C3 structure, corresponding to negligible stored energy, and shallower slopes corresponding to increasingly larger stored energy contributions.

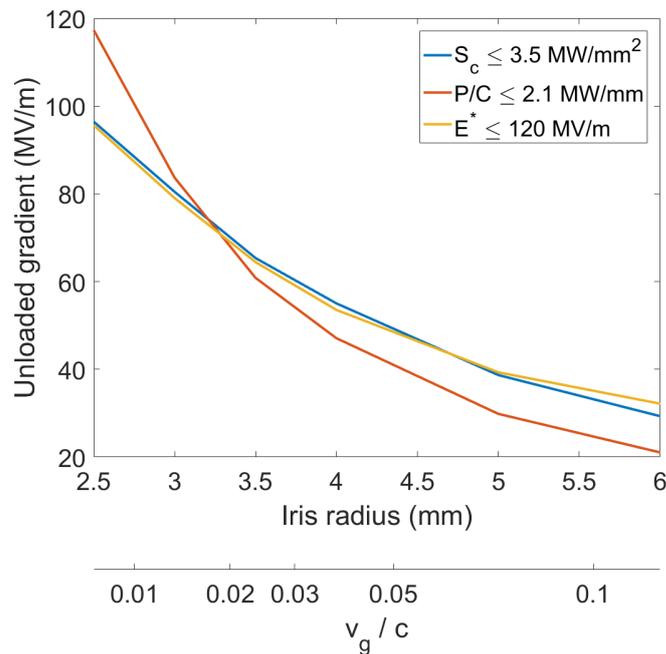


Figure 2.59: Maximum accelerating gradient without beam loading vs. aperture size and group velocity as a fraction of c . Blue curve: with a maximum permitted S_c of 3.5 MW/mm². Red curve: with a maximum permitted P/C of 2.1 MW/mm. Yellow curve: with a maximum permitted E^* of 120 MV/m.

A calculation of E^* was performed for the C3 geometry as outlined in Sec. 2.3.1, with the exception that R_{bd} was replaced with L . As L has different units to R_{bd} and a

different range of values, a different proportionality constant for the emission function (2.8) had to be chosen. In this case, a value of $7.5 \times 10^{-18} \text{ Am}^{3/2}\text{V}^{-3/2}$ was chosen for k .

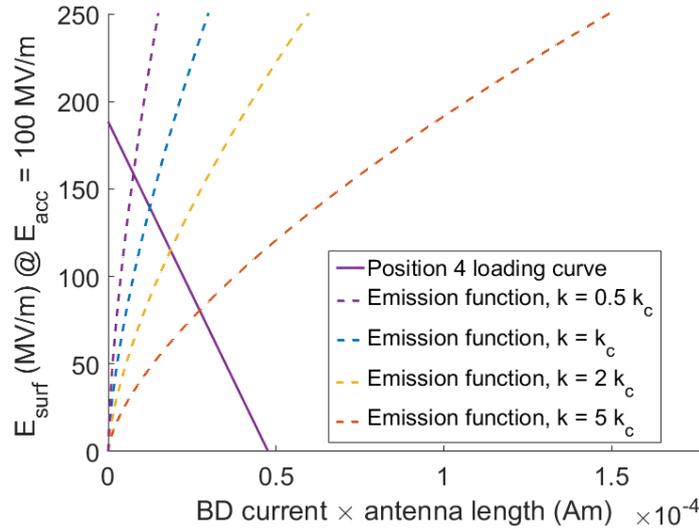


Figure 2.60: Surface electric field in MV/m vs. product of breakdown current and antenna length in Am for a 3 mm aperture radius and 100 MV/m accelerating gradient. Solid purple line: loading plot for the antenna at position 4, as per Fig. 2.15. Dotted lines: emission functions with various values of k as indicated on the legend, with $k_c = 7.5 \times 10^{-18} \text{ Am}^{3/2}\text{V}^{-3/2}$.

To obtain an estimate for the appropriate value for k , it was decided that the emission function should cross the loading curve of the antenna close to the midpoint between its x- and y-intercepts. An example using the antenna at position 4 (the location at which most breakdowns would normally occur in experiments) for the 3 mm aperture radius C3 structure is shown in Fig. 2.60. If the value of k were much larger or much smaller, the emission function would either be very close to the horizontal axis or very close to the vertical axis. This would mean that the E^* value would always be either very close to zero or E_0 respectively. Neither result would be very useful. k was then fine-tuned to give the closest match to the S_c curve in Fig. 2.59.

The exponent $n = \frac{3}{2}$ remained the same as in Sec. 2.3.1. A comparison of the maximum gradient for different aperture sizes for the C3 structure was made, using the recalculated values of E^* , S_c , and P/C , with the result shown in Fig. 2.59. In contrast

with the results obtained using R_{bd} , shown in Fig. 2.24, there is much better agreement between E^* and S_c , even for low group velocities.

A system with a finite bandwidth has a finite transient response to an excitation, with a narrower bandwidth corresponding to a larger time constant. Thus, a single-frequency model corresponds to a steady-state condition with no transient effects. Since the finite stored energy can only provide power to the breakdown for a finite amount of time, the contribution of stored energy to the evolution of a breakdown is necessarily a transient effect. Considering the impedance of the structure over a finite bandwidth appears to have achieved the desired effect of including the contribution of stored energy to the problem, resulting in predictions of maximum gradient much closer to those of S_c .

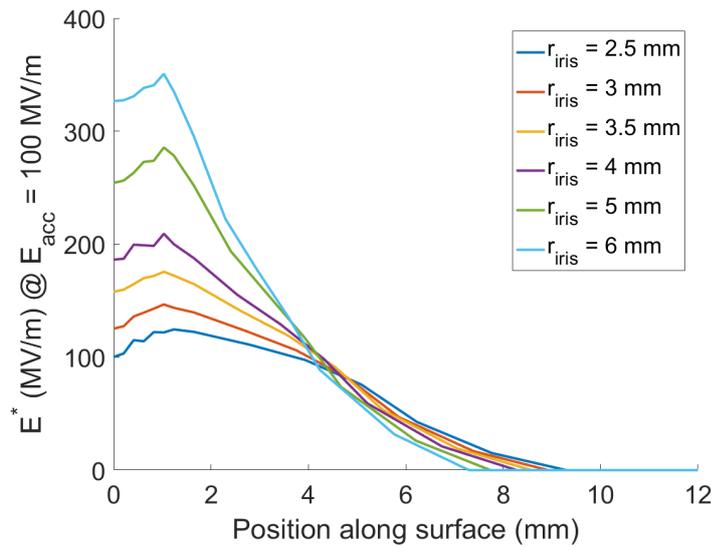


Figure 2.61: E^* in MV/m for an accelerating gradient of 100 MV/m vs. position along the surface in mm. The position is defined by the arrow shown in Fig. 2.15, where 0 is the iris apex. Different coloured plots represent C3 structures of different aperture sized as indicated on the legend.

Using L instead of R_{bd} results in a distribution of E^* within the cell as shown in Fig. 2.61. As expected, the peak value in all cases is about 1.5 mm off the iris apex, as with S_c , shown in Fig. 2.22. The general behaviour is similar to that with R_{bd} , shown in 2.23, though the peak at the 1.5 mm position remains present even for $r_{iris} = 2.5$ mm, which was a possible problem with R_{bd} that was identified earlier.

L was also calculated for two X-band geometries that have been experimentally tested, the T24PSI structure and the CLIC Crab Cavity. The antenna locations and simulation conditions were the same as described in their respective sections 2.3.5 and 2.3.4. A comparison of the L and R_{bd} values, normalised to an antenna length of 0.1 mm, for these geometries is also shown in Fig. 2.58. The resulting values are in the same order of magnitude as those of the C3 structures, reflecting the similarities in their overall dimensions. The ratio L/R_{bd} for the two tested structures also appears reasonable given their respective group velocities.

All of the sets of points, with the exception of those corresponding to the T24PSI structure, roughly form lines crossing the origin. This is due to the tapered design of this structure, which means that the properties of each individual cell, including the group velocity, vary along the length of the structure. Since the group velocity of a cell determines the ratio of power flow to stored energy, it should in turn determine L/R_{bd} . Therefore, it should follow that due to the tapering of the T24PSI structure, each cell would have a different L/R_{bd} , meaning that the points representing them in Fig. 2.58 do not form a line crossing the origin.

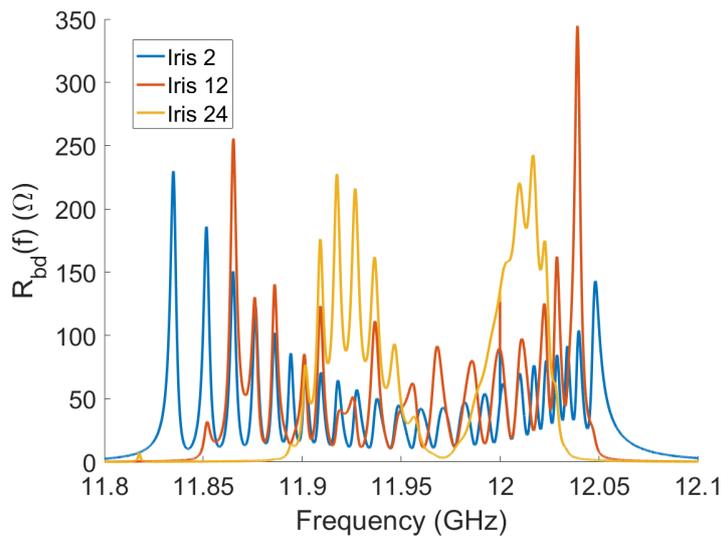


Figure 2.62: R_{eq} in Ω vs. RF frequency in GHz in the T24PSI structure. Blue: for an antenna on iris 2. Red: for an antenna on iris 12. Yellow: for an antenna on iris 24. Antenna locations are shown in Fig. 2.38, with iris 1 representing the iris closest to the RF input.

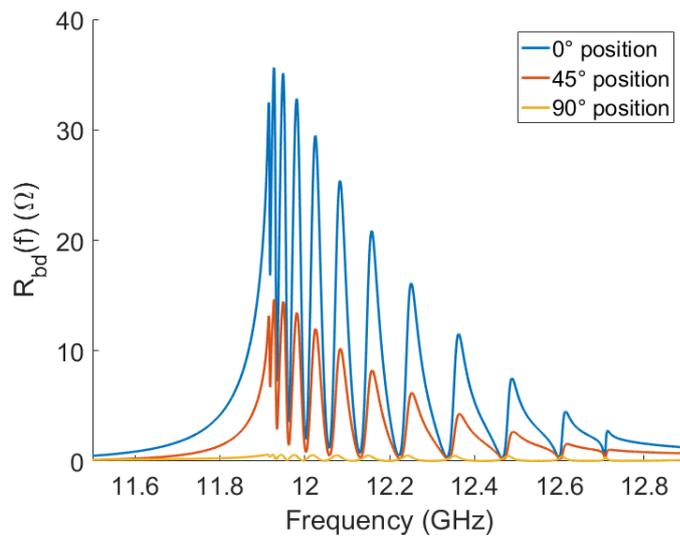


Figure 2.63: R_{eq} in Ω vs. RF frequency in GHz in the first cell of the CLIC Crab Cavity. Blue: at the 0° position. Red: at the 45° position. Yellow: at the 90° position. Positions are defined as per Fig. 2.34.

Plots of $R_{eq}(f)$ are shown for the T24PSI structure in Fig. 2.62, and for the Crab Cavity in Fig. 2.63. Of interest is the way the bandwidth decreases when moving along the T24PSI structure, which is possibly due to the tapering in the structure's group velocity. On the other hand, the plots of different locations in the first cell of the Crab Cavity resemble each other, but are scaled differently - reflecting the fact that all three antennas are located in the same cell, but have different couplings to the operating mode of the structure.

The predicted value of $R_{eq}(f)$, according to the analytical model discussed in Sec. 2.2.1, is plotted for the iris apex of the C3 structure in Fig. 2.64. While the shape of the function is very different from its equivalent obtained from the finite element simulations, as in Fig. 2.54, the way in which it is scaled down in resistance and stretched on the frequency axes with increasing aperture size shows a strong resemblance. The difference between the two is likely due to the fact that the analytical model assumes an infinitely long structure that is well impedance-matched over its entire bandwidth, whereas the structure that was simulated had a finite length and could only be impedance matched over a small bandwidth close to its operating frequency. A calculation of an analytical equivalent of L , denoted L_{an} , was performed as follows:

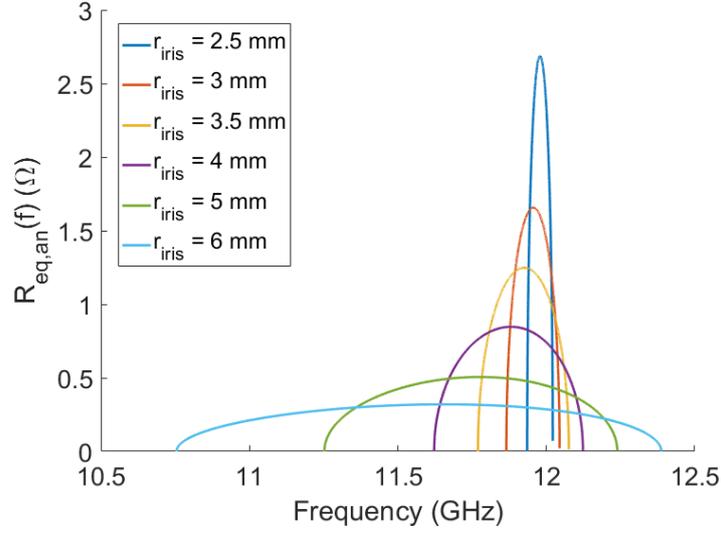


Figure 2.64: $R_{eq,an}$ in Ω vs. RF frequency in GHz, on the iris apex of C3 structures. Different coloured curves represent C3 structures of different aperture sizes, as indicated on the legend. The values were calculated using the analytical formula (2.28) and eigenmode results of the cell geometry.

$$L_{an} = \sqrt{\int_{\omega_0}^{\omega_\pi} R_{eq,an}(\omega)^2 d\omega}, \quad (2.47)$$

where $R_{eq,an}$ refers to the value of R_{eq} obtained from the analytical model as per Eq. (2.28). This results in values that are very well correlated to their equivalents calculated in finite element simulations for the same geometry and location, as demonstrated in Fig. 2.65. This means that for accelerating structures operating in the TM_{010} mode, which CLIC structures are an example of, the R/Q , spatial distribution of surface electric field, and dispersion curve (all obtainable from an eigenmode solution), are all that is required for an accurate prediction of E^* and thus the breakdown limit. Since this is also the case for the broadband model, accurate predictions for low group velocity structures should also be possible.

2.5 Extension to DC

As discussed in Sec. 2.1.4, the ideal breakdown quantity should also encompass the case of zero frequency, i.e. DC. None of the quantities discussed thus far works for DC: S_c

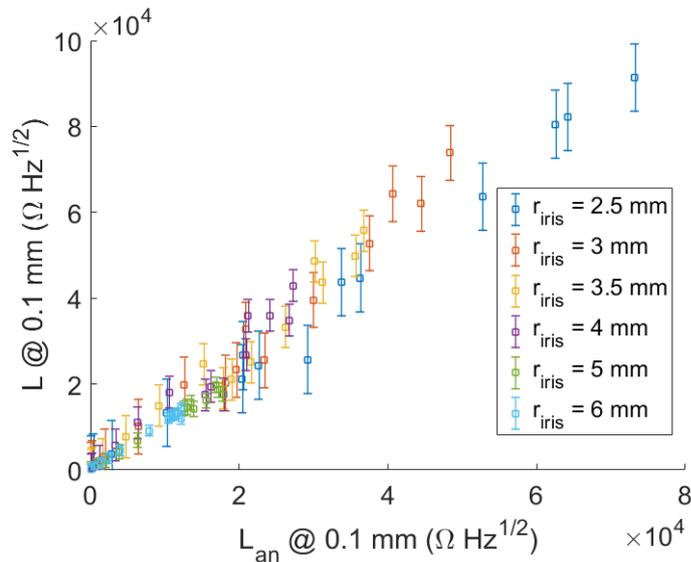


Figure 2.65: L_{an} , obtained from (2.47) and eigenmode results vs. L , calculated via finite-element simulations. Different coloured points represent C3 structures of different aperture sizes, as indicated on the legend.

and P/C both predict that a DC experiment should be able to hold an infinite surface field, since there is no power flow in the steady state, whereas Boyd's formulation of Kilpatrick's criterion discussed in Sec. 2.1.4 predicts breakdowns at zero field. A quantity for DC is needed not just for completeness, but because DC experiments are much simpler and cheaper to set up and instrument than RF tests, especially for testing candidate materials for accelerating structures, as well as investigating the fundamental science of breakdowns. The majority of the applications of breakdown science also involve DC rather than RF fields. Furthermore, a consistent theory encompassing both RF and DC results can be applied to new problems with confidence.

A number of behaviours have been observed in the LES, a DC breakdown experiment at CERN complementary to development of accelerating structures for CLIC (see Sec. 1.5) which are believed to be important in constraining the behaviour of E^* when applied to a DC experiment. They are as follows:

- The temporal distribution of breakdowns within high voltage pulses (Sec. 2.5.1).
- The spatial distribution of breakdowns in various electrode configurations (Sec. 2.5.2).

- The dependence of the maximum attainable electric field on the size of the inter-electrode gap (Sec. 2.5.3).

This is followed by a discussion of an attempt to apply E^* to a theoretical geometry qualitatively resembling that of the LES, and a comparison of the results to experimental data in Sec. 2.5.4.

2.5.1 Breakdown Times

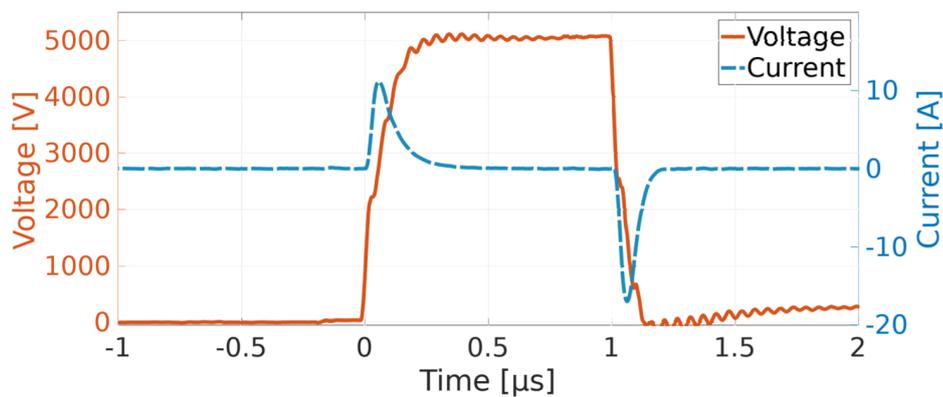


Figure 2.66: Voltage in kV (red) and current in A (blue) vs. time in μs in a typical pulse without a breakdown in the LES [85].

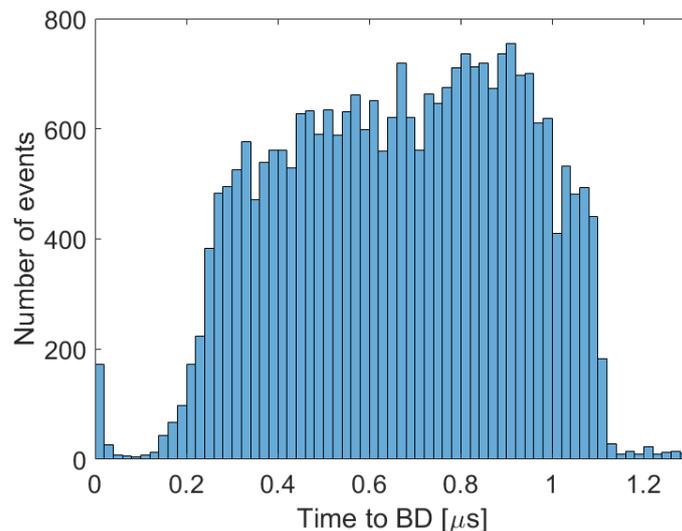


Figure 2.67: Number of accumulated breakdowns vs. time relative to the start of the high voltage pulse in μs , in a pair of electrodes in the LES [86].

It is clear that quantities based purely on power flow cannot be used unmodified at DC, as no power flows in a DC experiment in the charged, non-breakdown state. There

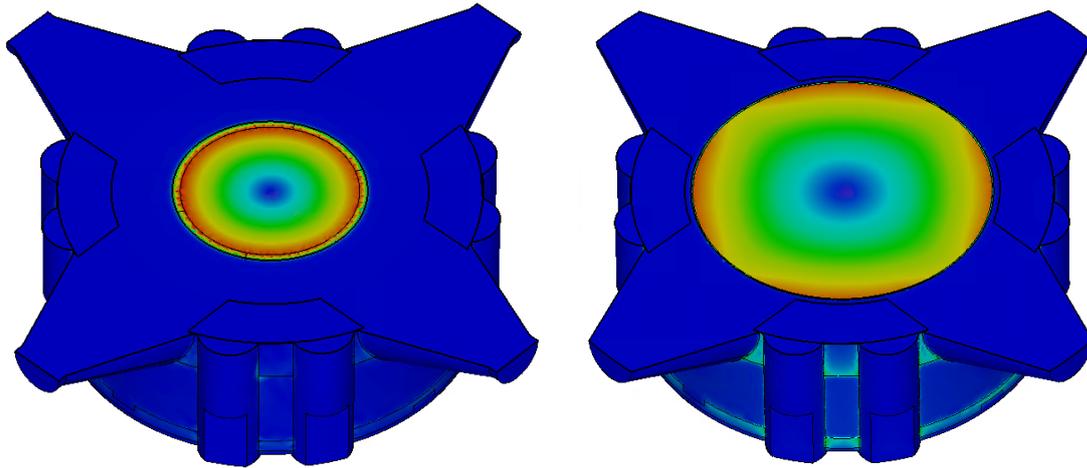


Figure 2.68: S_c , in arbitrary units (with blue representing zero and red representing the greatest value) vs. transverse position in the gap of the LES when a 100 MHz sinusoidal voltage is applied to the input. Left: With 40 mm diameter circular electrodes installed. Right: With 60 mm diameter circular electrodes installed.

is only power flow when charging and discharging the system, meaning that a quantity based on power flow would predict breakdowns to occur only at those times. Fig. 2.66 shows an example of current and voltage waveforms in the LES, demonstrating the lack of power flowing in or out of the system for most of the pulse. On the other hand, Fig. 2.67 shows the distribution of breakdown times relative to the start of the high voltage pulse for a typical run in the LES. It is clear that the timing of breakdowns is much better correlated with the voltage pulse shape rather than the power. The idea that power flow has to be considered for the system during a breakdown, separately from the case in which there is no breakdown, as per the E^* model seems a more satisfying explanation of this behaviour.

2.5.2 Breakdown Locations

Despite the inconsistency in the time of the breakdown, S_c does appear to predict the spatial distribution of breakdowns in the LES correctly. Fig. 2.68 shows the distributions of S_c for 40 mm and 60 mm diameter electrodes. Rather than a DC voltage, an RF signal at an arbitrarily chosen frequency of 100 MHz has been assumed to be incident at the input port. In both cases, the S_c is highest at the edge of the electrode, despite the E field being almost uniform within the gap. The larger S_c at



Figure 2.69: Images of 60 mm diameter DC electrodes after testing, both showing breakdown craters concentrated on the edges of the electrodes and grouped into four islands. Left: Tested at CERN, with breakdown craters appearing as dark points [87]. Right: Tested at the University of Helsinki, with breakdown craters appearing as white points [88].

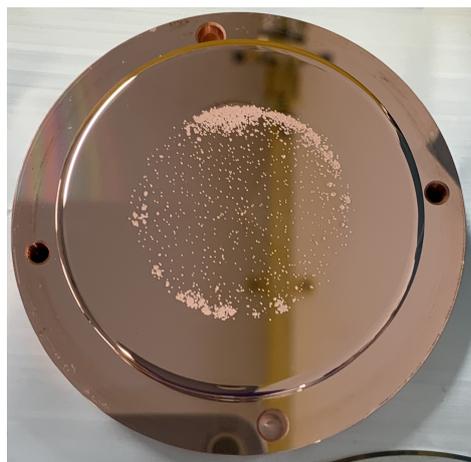


Figure 2.70: The cathode from a DC test involving a large cathode and a small anode. Breakdown craters appear as white points, concentrated in a 40 mm diameter ring [89].

the edge of the electrode appears consistent with better accessibility to incident power from the input port. A dependence on the proximity of the chamber wall can also be seen with the larger electrodes, with the appearance of four ‘islands’ of peak S_c . Images of electrodes tested in both this system and a very similar one at the University of Helsinki [85], shown in Fig. 2.69, show a distribution of breakdown craters consistent with this pattern.

It was initially suspected that the reason for the clustering of breakdowns on the edges of electrodes was field enhancement on the corners of the electrodes, resulting in a 25% higher electric field there than in the centre of the electrode [90]. Further investigations into this have not supported this hypothesis. The clustering remained when electrodes with a specially optimised edge profile, which limited the field enhancement to 3%, were used. Another test involved the use of a large-diameter cathode and a smaller-diameter anode. This configuration removed the field enhancement at the cathode completely. The clustering of breakdowns on the edge of the anode still remained (see Fig. 2.70), suggesting that either the anode plays a role in breakdown, or that the field enhancement is not the cause of the clustering. The effect of power flow during the onset of breakdown thus appears to be a better explanation.

2.5.3 Gap Size Dependence of Maximum Field

Another interesting phenomenon that has been observed in the LES is the dependence of maximum attainable electric field on the gap size. A fresh pair of electrodes was first fully conditioned, then a series of runs was performed in which ceramic spacers of different thicknesses were installed and the electrodes were re-conditioned with a gap size that varied between 20 μm and 100 μm [91]. The electric field applied was controlled by a feedback system which attempted to keep the breakdown rate at a constant value. A plot summarising these runs is shown in Fig. 2.71. In this plot, it can be seen that a higher field could be reached when the gap size was smaller, despite the otherwise identical experimental setup. It was found by Profatilova et. al. that the maximum field followed the relation $E \cdot d^{0.28} = \text{const.}$ Fig. 2.72 shows the result of

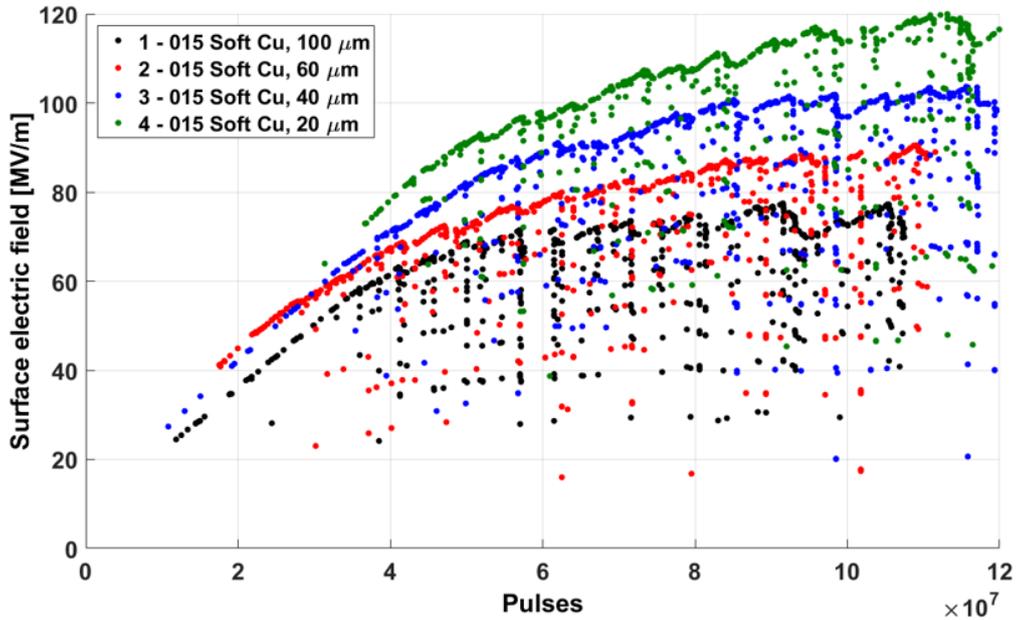


Figure 2.71: Surface electric field in MV/m vs. number of voltage pulses applied for various gap sizes in the LES [91]. Black: 100 μm gap. Red: 60 μm gap. Blue: 40 μm gap. Green: 20 μm gap.

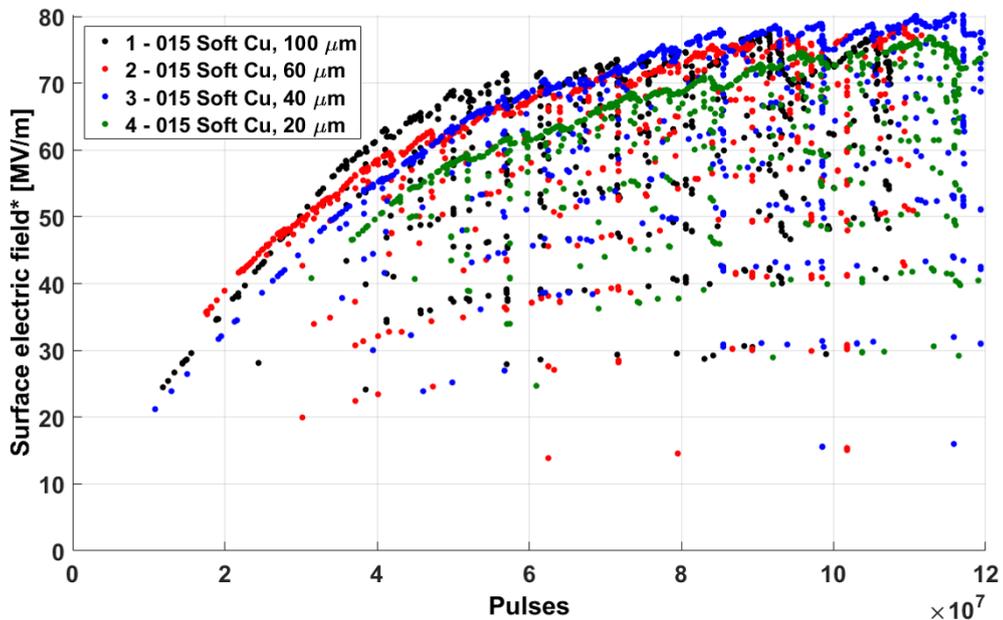


Figure 2.72: Normalised surface electric field* in MV/m vs. number of voltage pulses applied for various gap sizes in the LES [91]. The normalisation was performed by multiplying the field by $(d/d_0)^{0.28}$ has been applied to the field value to take into account the gap size dependence observed, where d is the gap size and $d_0 = 100 \mu\text{m}$. Black: 100 μm gap. Red: 60 μm gap. Blue: 40 μm gap. Green: 20 μm gap.

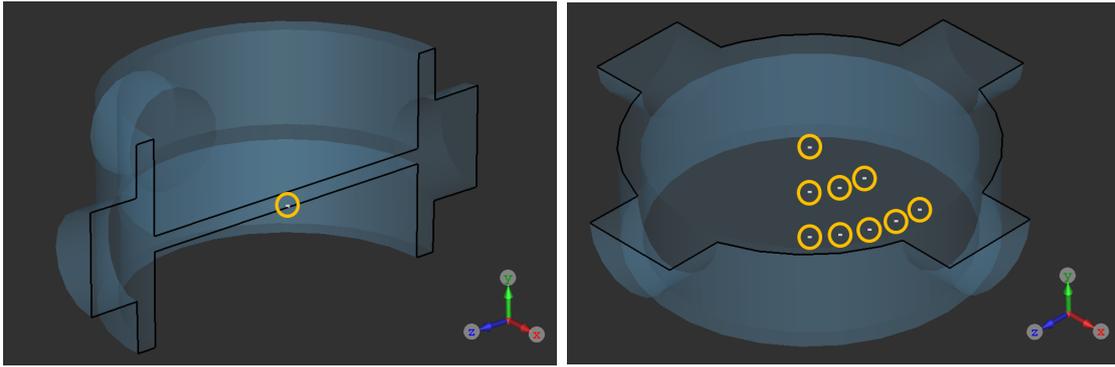


Figure 2.73: Cutaway views of the internal volume of the simplified DC system, with the locations of antennas shown as white lines with yellow circles around them. Left: cut vertically. Right: cut horizontally.

normalising the electric field using this factor of $d^{0.28}$, showing a smaller spread in the maximum value of normalised electric field between the different runs than there was with the raw electric field value.

2.5.4 E^* for a Theoretical DC Geometry

A calculation of E^* for the LES was attempted to check if the experimental observations discussed could be reproduced in simulations. Unfortunately, achieving numerical convergence in the electromagnetic simulation proved difficult due to the very small inter-electrode gap size and large and complicated geometry of the rest of the chamber. As a preliminary step towards calculating E^* for the LES, a simplified geometry was designed to check whether the qualitative behaviour of the LES seen in experiments could be reproduced in simulations.

This simplified geometry is depicted in Fig. 2.73. It features two cylindrical electrodes, 60 mm in diameter, separated by a small gap. This represents the electrodes of the LES. For easier meshing, the gap size in this model was much larger (between 0.5 mm and 4 mm) than the real LES, in which the gap could be varied between 20 μm and 100 μm (due to limitations of the maximum voltage that could be applied). To represent the windows in the LES chamber, four cylindrical protrusions were added to the simplified geometry. No ceramic spacer was included. Also, the simplified geometry did not have any feedthroughs for connecting the power supply. Instead, the entire top and bottom

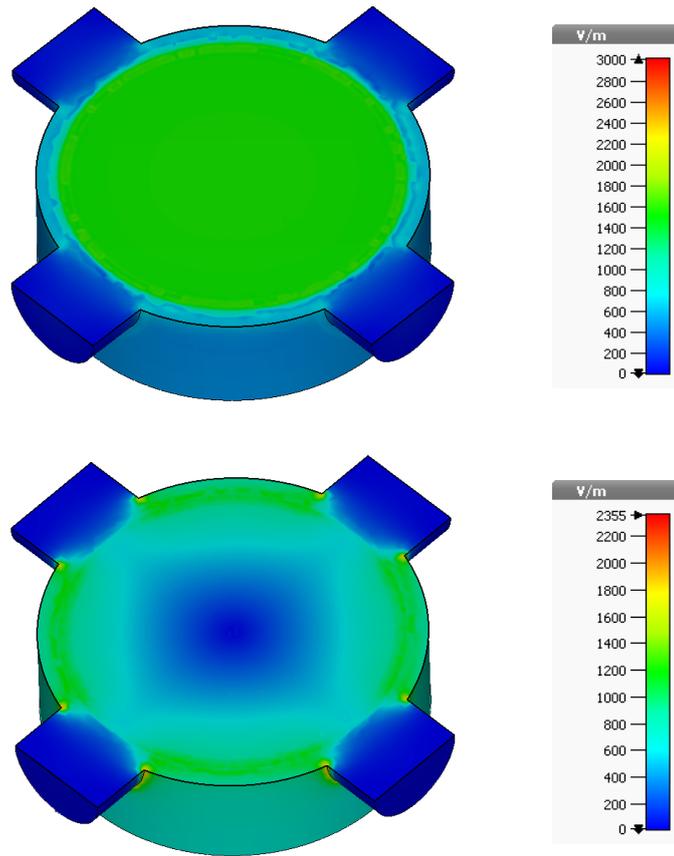


Figure 2.74: Fields vs. transverse position in the gap of the simplified DC system when a 500 MHz sinusoidal voltage is applied to the input. Top: Electric field. Bottom: S_c .

surfaces were defined as coaxial ports. Antennas were placed on a 45° sector of one of the electrodes: one in the centre of the electrode, three 13.35 mm from the centre, and five 27 mm from the centre, close to the edge. The antennas were all 0.5 mm long.

Fig. 2.74 shows the distribution of the electric field E and the S_c between the two electrodes. The S_c distribution seen in the full LES geometry, as shown in Fig. 2.68, qualitatively resembles that seen with the simplified geometry, with the maximum values being located near the edges of the electrodes, and slightly lower near the windows than away from them.

The equivalent series resistance R_{eq} for each antenna was obtained by taking the real part of Z_{eq} , which was itself calculated using Eq. 2.39. In this calculation, a source impedance value of $Z_0 = 7.5 \Omega$ was used. This is the characteristic impedance of a coaxial transmission line with a 60 mm diameter inner conductor and a 68 mm diameter outer conductor [68] - the dimensions of the simplified DC system away from the

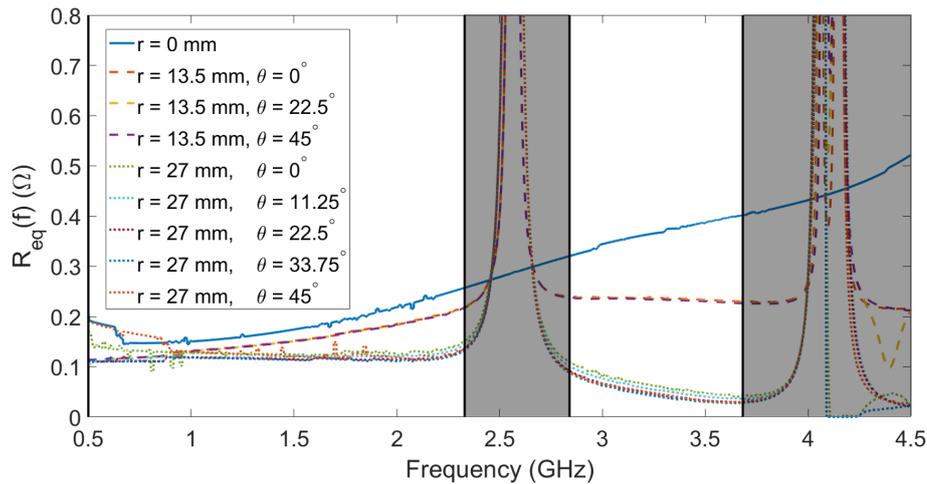


Figure 2.75: R_{eq} in Ω vs. frequency in GHz, for a 1 mm gap size in the simplified DC system. Different coloured curves indicate different antenna locations relative to the centre of the electrode, as indicated on the legend. Shaded areas indicate frequencies excluded from integration, see Fig. 2.78.

gap. Thus, there should be no impedance mismatch at the ports of the system; it is equivalent to extending the electrodes and chamber vertically to infinity. With no impedance mismatch between the inter-electrode gap and the power source, the results of this calculation should pertain to the gap and windows only, and no other features.

It should be noted, however, that this is not the case in the actual LES, which is fed with a $50\ \Omega$ coaxial cable from the high voltage-power supply, and has multiple impedance mismatches within the vacuum chamber between the high-voltage vacuum feedthrough and the gap. This impedes the flow of power from the power supply, and can result in resonances in the vacuum chamber which store energy locally. Such resonances were measured in the LES vacuum chamber and are discussed in a different context in Sec. 3.3.

A plot of $R_{eq}(f)$ for each of the antennas is shown in Fig. 2.75, for a gap size of 1 mm. In describing the location of each antenna, r represents the distance from the centre, and θ represents the polar angle such that 0° is aligned with one of the windows. The behaviour seen in these results appears generally consistent with the S_c simulation, most clearly visible roughly between 3 and 4 GHz. Excluding the peaks at around 2.6 GHz and 4.1 GHz, the value of R_{eq} is strongly correlated with the distance of the

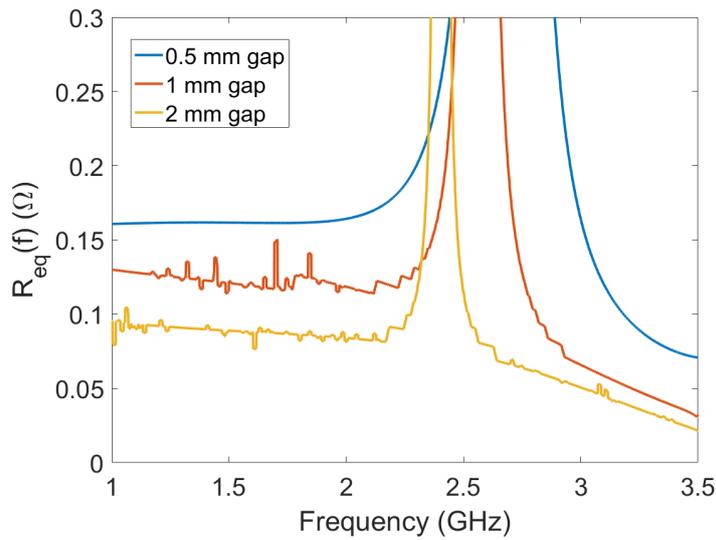


Figure 2.76: R_{eq} in Ω vs. frequency in GHz, for the $r = 27$ mm, $\theta = 45^\circ$ position. Fluctuations visible between 1 and 2 GHz are caused by numerical errors in the simulation. Blue: for a gap size of 0.5 mm. Red: for a gap size of 1 mm. Yellow: for a gap size of 2 mm.

antenna from the centre, with the different plots forming three distinct groups of $r = 0$, $r = 13.5$ mm, and $r = 27$ mm. The centre antenna has the largest R_{eq} and the antennas close to the edge of the electrode have the smallest R_{eq} . With a constant unloaded electric field over the entire electrode surface, this should correspond to breakdowns being most likely to occur closest to the edge of the electrode as seen in experiments.

Contrary to the case of a resonant accelerating cavity, a DC experiment does not have any design frequency other than zero. This precludes the use of the single-frequency method for calculating E^* , since no power can be coupled through an antenna at DC. This left the broadband method, formalised in Eq. 2.45, as the only practical way of completing the calculation. However, integrating R_{bd}^2 over frequency was problematic with the resonant peaks pointed out above, as they would dominate the result of the integration and produce results inconsistent with experiment. Specifically, it would result in E^* having the smallest value at the edges of the electrode, the location where most breakdowns occurred during testing. It is believed that the peaks correspond to resonant modes of the inter-electrode gap which do not couple well to the TEM mode of the coaxial input, and thus that they should be excluded from the calculation of L .

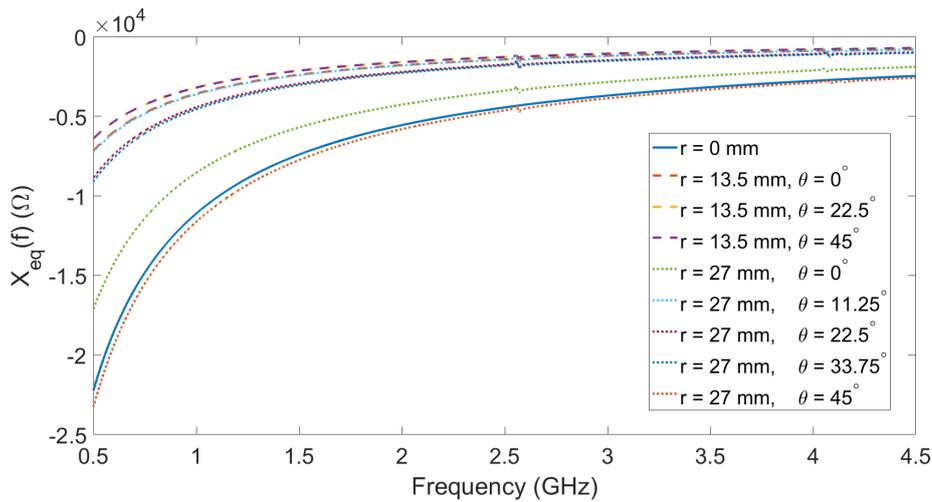


Figure 2.77: X_{eq} in Ω vs. frequency in GHz, for a 1 mm gap size in the simplified DC system. Different coloured curves indicate different antenna locations relative to the centre of the electrode, as indicated on the legend.

It was thus decided to perform the integration as normal, but over a frequency range excluding the peaks.

To do this, some consistent and quantitative way of defining the frequencies the integrate over had to be devised. This was important as the centre frequency and bandwidth of the resonant peaks depended on the gap size, as can be seen in Fig. 2.76, which meant that the frequency range had to be different for each gap size. While a future method might incorporate the transfer function which would describe how the fields in the gap couple to the fields at the input port, it was decided to use a simple, less-general method here. One method which appeared to work well used the imaginary part of Z_{eq} , X_{eq} , which is depicted for a 1 mm gap size in Fig. 2.77. In this figure, resonances can be seen at around 2.6 GHz and 4.1 GHz, corresponding to the ones seen in the real part. The resonances are superimposed onto a large capacitive reactance that is most likely the result of the capacitance of the respective antenna. The reactance X_c of a capacitor C as a function of frequency f is given by [69]:

$$X_c(f) = -\frac{1}{2\pi fC} \quad (2.48)$$

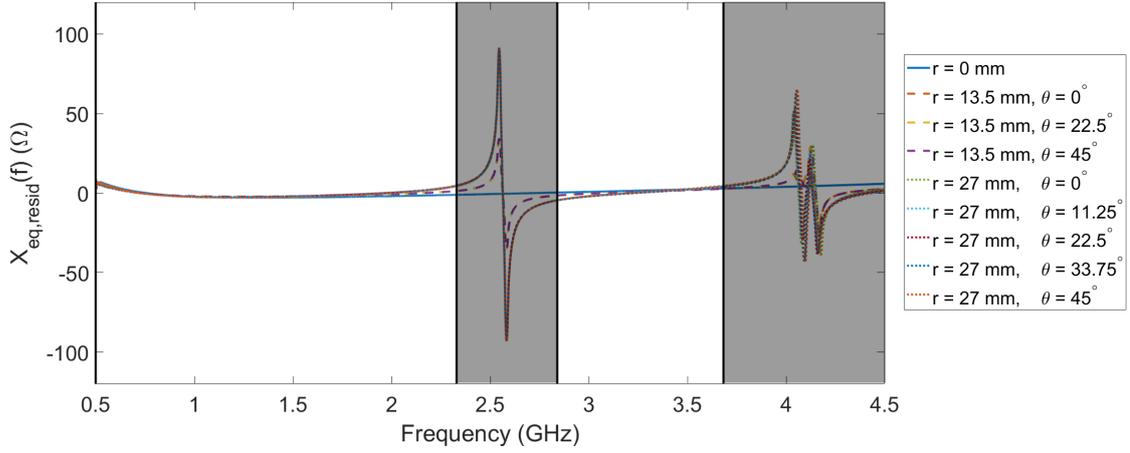


Figure 2.78: Residual X_{eq} in Ω vs. frequency in GHz, for a 1 mm gap size in the simplified DC system. Different coloured curves indicate different antenna locations relative to the centre of the electrode, as indicated on the legend. The black vertical lines show the frequencies at which $X_{eq,resid}$ crossed 5% of the positive and negative peak values. Shaded areas indicate frequencies excluded from the integration of R_{eq} .

To be able to see the resonances more clearly, $X_c(f)$ was fitted to each of the plots in Fig. 2.77. A value for C was chosen for each plot that gave the minimum mean-squared difference between $X_{eq}(f)$ and $X_c(f)$. The residual reactance $X_{eq,resid}$ for each antenna was then calculated by subtracting the fitted function $X_c(f)$ from the simulation results, i.e.:

$$X_{eq,resid}(f) = X_{eq}(f) - X_c(f) \quad (2.49)$$

The result for a 1 mm gap is shown in Fig. 2.78. The frequency limits of the integration were defined from this plot as the frequencies at which the value of $X_{eq,resid}(f)$ crossed 5% of its positive and negative peak values, and are shown in Figs. 2.75 and 2.78 as vertical black lines. In each case, there were two frequency bands over which the integration was performed. As shown in Fig. 2.75, these were 0.5 - 2.5 GHz and 2.6 - 4 GHz for the case of the 1 mm gap size.

The result of integrating $R_{bd}^2(f)$ over a frequency range defined in this way is shown for all of the antenna locations and gap sizes in Fig. 2.79. Qualitatively, the results appear consistent with behaviours seen in experiment, in that L decreases moving from

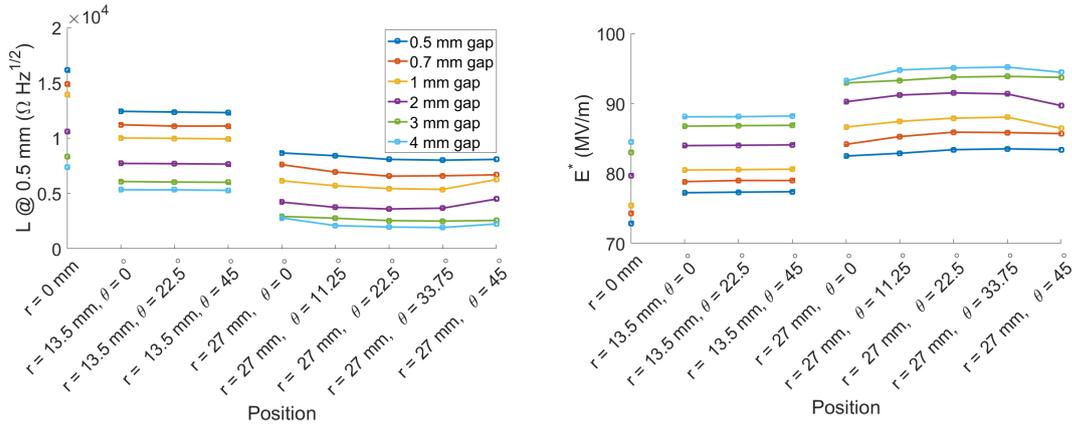


Figure 2.79: Quantities vs. position relative to the centre of the electrode in the simplified DC system. Different coloured points indicate different gap sizes. Left: L for an antenna length of 0.5 mm. Right: E^* for an unloaded gap field of 100 MV/m.

the centre of the electrode to the edge, and also decreases as the gap size is increased. For a constant unloaded electric field E_0 , a lower value of L corresponds to a higher breakdown rate. There is also a small angular dependence of L on the edge of the electrode, with the 0° position, aligned with the window, having a slightly higher L than positions not aligned with the window. This is also consistent with both the distribution of S_c and crater locations seen in experiments.

Using these values of L , and a uniform E_0 of 100 MV/m, E^* values were calculated using the same procedure as for the RF structures discussed in previous sections, also shown in Fig. 2.79. As the L values obtained for this geometry were much smaller (by about a factor of 60 when normalised to the same antenna length), a different emission function was used, with $k_0 = 6.7 \times 10^{-16} \text{ Am}^{3/2}\text{V}^{-3/2}$, which is about two orders of magnitude larger than that used for the RF structures. Since this geometry is theoretical and does not correspond to any physical experiment, there was no data to fit the value of k_0 to. Using the same value of k_0 as in the RF case, however, would have resulted in E^* being very close to E^0 for every antenna location. If this were the case, there should be no dependence on the gap size, which is also inconsistent with experimental results. Without a simulation that resembles the true geometry of the LES, or an experiment using the simplified geometry, no more information about what value k_0 should have for a DC experiment can be obtained.

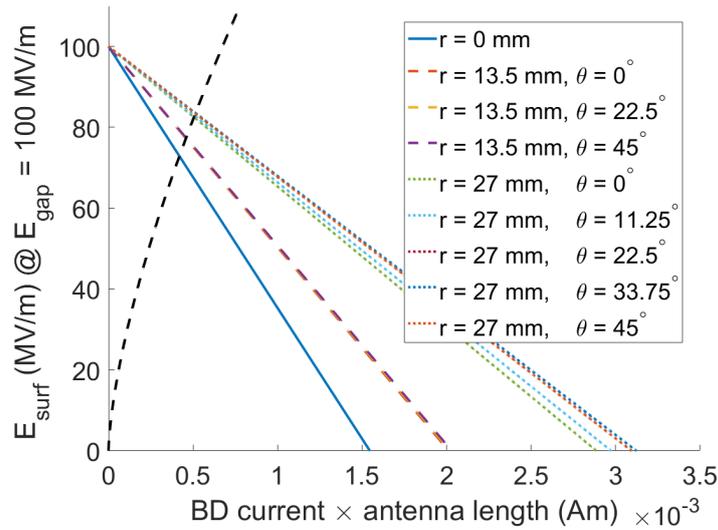


Figure 2.80: Surface electric field in MV/m vs. product of breakdown current and antenna length in Am, for a 0.5 mm gap size and an unloaded gap field of 100 MV/m in the simplified DC system. Different coloured lines represent different antenna locations relative to the centre of the electrode, as indicated on the legend. Black dashed line: emission function.

A breakdown loading plot representing this calculation for a 0.5 mm gap size is shown in Fig. 2.80. If the same value of k_0 were used as for the RF structures, the resulting E^* would be very similar for all situations, without any location or gap-size dependence. It should be noted, however, that the frequency dependence of E^* is not fully understood yet. Thus, it may be reasonable to scale either k_0 or L with frequency.

As with the experimental results from the LES, a gap-size dependence of E^* was found in this study. The dependence of the minimum value of L on the electrode for each gap size is plotted as a function of the gap size in Fig. 2.81. The location of the minimum value of L is the one at which most of the breakdowns are expected to occur, and thus defines the breakdown rate of the system as a whole. To make this result comparable with the experimental results from the LES, the value of E_0 for an arbitrarily chosen maximum allowed value of E^* of 100 MV/m for each gap size was also calculated and shown in Fig. 2.82. This reflects the way the gap-dependence study in the LES was conducted, in which the breakdown rate was kept constant whilst the maximum electric field was measured for each gap size. L appears to be proportional to the gap size to the power -0.55 ± 0.05 , with a reasonably good fit to this power-law dependence. E_0 ,

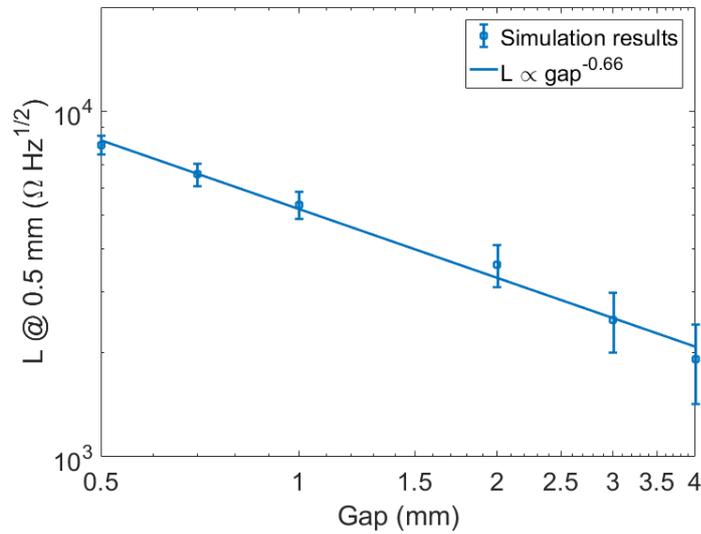


Figure 2.81: L in $\Omega \text{ Hz}^{1/2}$ vs. gap size in mm, in the simplified DC system. Each point represents the the minimum value of L over the entire electrode at a given gap size. The line represents a power-law fit of $L \propto \text{gap}^{-0.66 \pm 0.05}$.

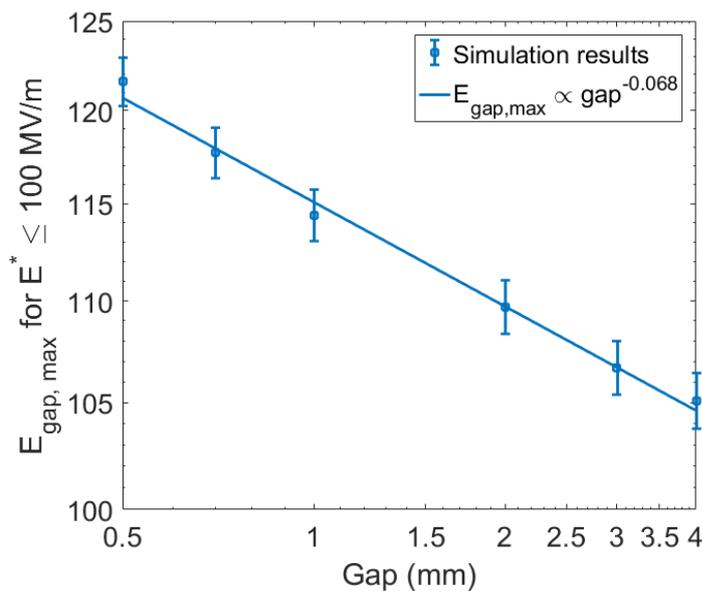


Figure 2.82: The maximum unloaded gap field in MV/m vs. gap size in mm, in the simplified DC system for a maximum permitted E^* of 100 MV/m. The points represent simulation results, and the line represents a power-law fit of $E_{\text{gap,max}} \propto \text{gap}^{-0.068 \pm 0.004}$.

has a weaker dependence on the gap size, with an exponent of -0.068 ± 0.004 . This is much lower than the -0.3 seen in the experiment, though in the correct direction, with larger gaps resulting in lower unloaded gap fields. The dependence of E_0 on the gap size should be very small for large gaps, because in this model the E^* can never be larger than E_0 . Thus, with a setup such as this one in which E_0 always decreases with increasing gap size, E_0 will asymptotically tend to the limiting value of E^* . This means that a power-law dependence cannot hold for all gap sizes in this model.

This asymptotic behaviour could potentially be a way of experimentally measuring the value of E^* , and thus help with characterising the emission function. So far in this thesis, no other way of measuring E^* during a breakdown has been identified. Only the unloaded surface field E_0 and the impedance quantities L and R_{bd} have been calculated. Such an experiment would involve progressively widening the inter-electrode gap, and for each gap size, measuring the breakdown rate as a function of surface electric field. Since, as detailed above, it is expected that E^* tends to E_0 for large gap sizes, the breakdown rate at a given E_0 should become independent of the gap size for large enough gap sizes. In this regime, it should be the case that $E_0 \approx E^*$. Since, E_0 can be easily calculated, this means that the breakdown rate can be measured as a function of E^* , i.e. $BDR(E^*)$. If it can be shown that this dependence $BDR(E^*)$ is consistent between multiple experiments with electrodes of different geometries, then it might become possible to deduce the maximum E^* in an experiment based on the current breakdown rate. This might then allow the emission function to be determined experimentally as well, since E_0 , L , and E^* would all be known under any given conditions.

The maximum voltage reachable by the LES may prove to be a limitation to such a measurement if this asymptotic behaviour only appears with very large gap sizes, thus requiring high voltages to set up the required electric field. If no experimental evidence for this asymptotic behaviour is found, this may also prompt modifications to the theory.

2.5.5 Conclusions

Due to the difficulty in simulating the actual geometry with which the experiments were performed, direct quantitative comparisons cannot be made. The simplified model studied here had much lower L values than those of the X-band RF structures, though this may be a result of the unknown scaling of E^* with frequency. It may also be due to the fact that the geometry had an extremely high group velocity of close to c , and the C3 structure study discussed in Sec. 2.3.1 showed that higher group velocities correspond to lower R_{bd} and L values.

Nevertheless, the predicted behaviour is qualitatively consistent with experimental data showing a clustering of breakdown craters on the edges of the electrodes away from the viewports, and a decrease in maximum reachable field as the gap size is increased. This shows that the idea of the breakdown rate being a function of breakdown-loaded surface electric field is compatible with experimental results from the LES. A possible avenue for probing the exact form of the emission function, along with the dependence of the breakdown rate on E^* , by investigating the gap-size dependence in more detail has been identified. Simulating the exact setup used in experiment, or simplifying the experiment in a way that would make simulations easier, would significantly increase the usefulness of the DC results in further developing the breakdown model.

2.6 Conclusions and Future Work

In this chapter, a new model quantifying the power flow during vacuum breakdowns in normal-conducting high-electric-field devices has been presented. Two options were considered for the impedance quantity. One is the single-frequency quantity R_{bd} , relating to the steady-state case, not taking stored energy into account. The other is the broadband quantity $R_{eq}(f)$, which shows the behaviour of the system over a range of frequencies, thus providing additional information about the local stored energy available to provide power to the breakdown. Using the integral $L = \int_0^\infty R_{eq}^2(f)df$ as an impedance-like quantity that can be used to determine the loaded surface field during

a breakdown gives results that work with low-group-velocity structures that have a large contribution from the local stored energy. It is possible that this quantity may not be valid in all cases, and some other function of $R_{eq}(f)$ may have to be formulated to produce a more general model.

The emission function is not known with certainty, but a possible one was found that appears to give results consistent with the X-band RF structures that were studied for this thesis. With more data, possibly from a dedicated experiment, it might be possible to define the function more accurately. It may also be possible that the limiting quantity is not the loaded electric field, but a different quantity such as the power delivered to the breakdown. Other details such as the dependence on frequency and pulse length, which are used in the design of particle accelerators by means of empirical scaling laws, might have an influence on the emission function.

An analytical circuit model for simple travelling-wave RF structures was derived and shown to exhibit the same behaviour as full-wave electromagnetic simulations, supporting the validity of the model and presenting a possible practical method of designing and optimising accelerating structures using the breakdown-loaded antenna model.

Experiments that could be useful to consider in the future to further validate and constrain the theory have been identified and are described below. In both cases, behaviours that seem consistent with the results have been identified and need to be verified quantitatively.

2.6.1 ITER Neutral Beam Injector

MITICA is a prototype of the Heating Neutral Beam (HNB) injector [92] for the ITER (International Thermonuclear Experimental Reactor) fusion reactor [93]. It is designed to output a 40 A DC neutral deuterium beam at an energy of 1 MeV, and consists of an RF ion source that produces D^- ions, an electrostatic accelerator, and a neutraliser to remove the excess charge from the accelerated ions to produce a neutral beam.

The Voltage-Holding Prediction Model (VHPM) was developed to guide the design of the high voltage DC components of the neutral beam injector, and predict the maximum voltage they can sustain [94]. It was based on the Cranberg-Slivkov model [95], [96], but adapted to work in the regime of very large inter-electrode gaps on the order of cm to m, which are uncommon in studies of vacuum breakdowns [97]. It defines a breakdown parameter W as follows:

$$W = E_c^\gamma E_a^\alpha U, \quad (2.50)$$

where E_c is the electric field on the cathode surface, E_a is the electric field on the anode surface, U is the total voltage, and α and γ are parameters fitted to experimental data. The fit was performed using data obtained from experiments on scale models and earlier operational machines. The presence of the total voltage term U implies a dependence on the gap size, which is consistent with many other DC experiments on vacuum breakdowns, including the LES at CERN.

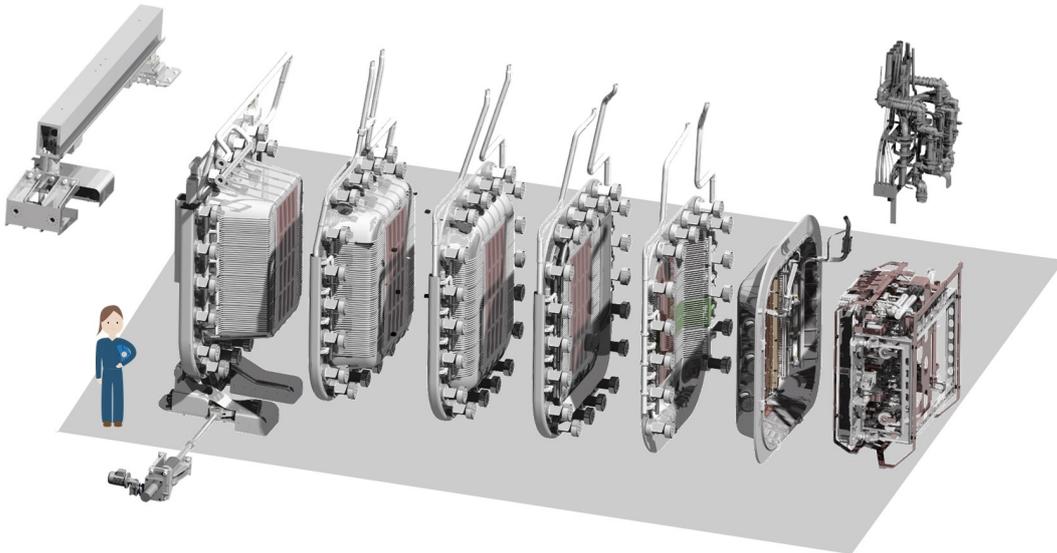


Figure 2.83: An exploded view of the ITER beam source showing the six individual accelerating grids [98].

Because of this total voltage dependence, the electrostatic acceleration stage of the injector has six separate grids, pictured in Fig. 2.83, held at successively increasing potentials. Thus, instead of a single gap with a potential difference of 1 MV, there are

five gaps with 200 kV each. This brings about the same increase in beam energy but allows a higher overall accelerating gradient, due to the smaller total voltage across any given gap.

The gap-size dependence implied by the VHPM could also be interpreted in the context of the antenna model as follows: dividing the volume of the DC accelerator into five separate regions by means of the intermediate voltage grids reduces the stored electric energy available to power a breakdown occurring within any given gap, giving a lower E^* in the gap than if there were just one gap with a potential difference of 1 MV.

It would be a worthwhile endeavour to check if the breakdown-loaded antenna model could produce results consistent with the VHPM, and thus provide an alternative explanation of the total voltage term in Eq. 2.50.

2.6.2 Choke-Mode Accelerating Structures

As part of the CLIC study, an investigation on applying the concept of choke-mode cavities to wakefield damping was performed [99]. Choke-mode cavities include a groove in the side wall containing RF absorber material to damp higher-order modes excited by the beam. A diagram is shown in Fig. 2.84. To prevent the absorber material from damping the accelerating mode, where a high quality factor is a crucial parameter, a quarter-wavelength choke is added in order to isolate the absorber from the cell. Since the choke is no longer a quarter wavelength long at the frequencies of the undesired higher-order modes, they are not stopped by it.

The radially symmetric geometry means that such a structure can be manufactured by turning only, without any milling steps, resulting in lower cost. Prototype standing-wave structures operating at 11.424 GHz that included a choke-mode cell were designed at Tsinghua University to apply this concept for a CLIC-like structure, and tested at KEK.

Part of this study involved the optimisation of the dimensions of the choke for the best high-gradient performance [102]. The geometry of the test structure is shown

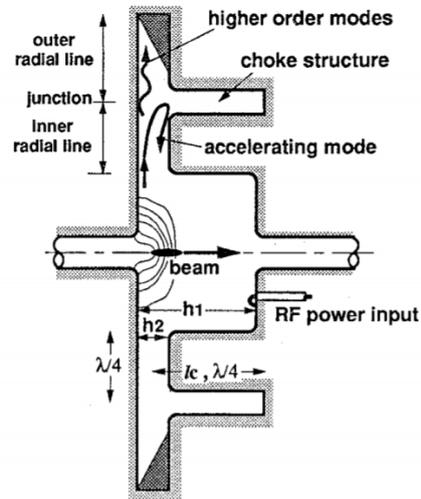


Figure 2.84: Conceptual diagram of a choke-mode cavity [100].

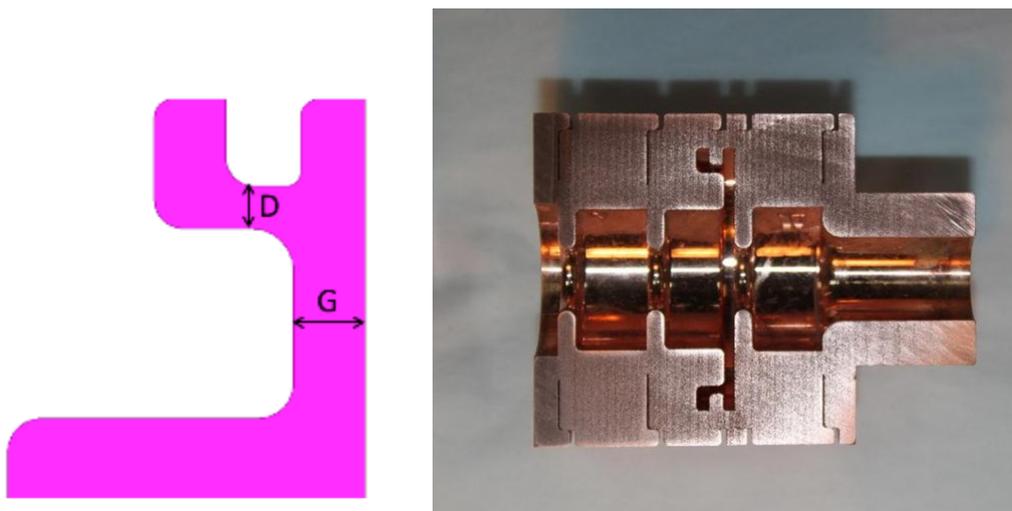


Figure 2.85: A standing-wave X-band choke-mode structure. Left: Dimensions of the choke. Right: Cutaway view of an assembled prototype [101].

in Fig. 2.85. Dimensions D and G , as indicated on the diagram were varied and the resulting prototypes were tested. The results can be summarised by a quantity CHK which is a function of the electric field within the choke and one of the geometric parameters:

$$\text{CHK} = 173 \times \left(\frac{E_{choke}^{max}}{E_{surf}^{max}} \right)^{-0.707} D^{0.711}, \quad (2.51)$$

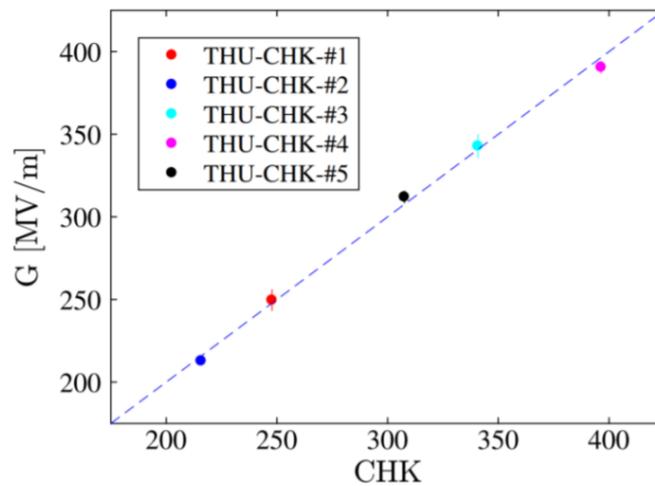


Figure 2.86: Maximum accelerating gradient G in MV/m vs. the geometric quantity CHK , defined in (2.51) [102].

where E_{choke}^{max} refers to the maximum electric field inside the choke, E_{surf}^{max} refers to the maximum surface field in the entire cavity, and D is the dimension indicated in Fig. 2.85. This quantity was found to be well-correlated with the maximum gradient achievable in the structure, as shown in Fig. 2.86. The $D^{0.711}$ term in (2.51) is reminiscent of the gap size dependence in the LES (Sec. 2.5). The choke also resembles a gap between electrodes but with an oscillating rather than DC field.

It would be worth performing E^* calculations on such a geometry, containing features seen in both RF and DC cases to verify that the observed behaviour can be reproduced.

Chapter 3

Dark Current Fluctuations

3.1 Motivation

3.1.1 Breakdown Statistics and Conditioning

The study of vacuum breakdowns under high electric fields, particularly in the context of particle accelerators, has led to the observation of a number of notable behaviours.

An important one is the stochastic nature of breakdowns in normal-conducting accelerating structures. Work done by various groups has repeatedly shown that the probability that a breakdown will occur in a given RF pulse scales with the surface electric field and pulse length [43]. One empirical scaling law used to quantify this dependence is $\text{BDR} \propto E^{30}\tau^5$, where BDR is the breakdown rate per pulse, E is the surface electric field, and τ is the pulse length [42]. Thus, there is no strict breakdown threshold as a given device can be operated at a higher field if a higher breakdown rate can be tolerated. This is in contrast to, for example, superconducting accelerating cavities which are limited by a critical magnetic field strength above which quench will inevitably occur [32].

The times and locations at which breakdowns occur are also of interest. Despite the tendency for a breakdown to produce a crater with field-enhancing sharp features,

approximately half of all the breakdowns in a given experiment have been found to occur at a location in which no breakdown has occurred previously, and after a long period (approximately 5000 pulses or more) since the last breakdown [29]. This further reinforces the implication that breakdowns are triggered by some stochastic process.

The conditioning curves of different RF structures and DC test samples become much more consistent from experiment to experiment if plotted against number of pulses as opposed to the number of breakdowns [19]. Thus, conditioning appears not to be a matter of destroying contaminants or extrinsic features, with the possible exception of very early in the conditioning process. Instead, each RF pulse seems to act to modify the structure surface to strengthen it against breakdown whether or not breakdowns actually occur.

3.1.2 Dislocation-Driven Breakdown Nucleation

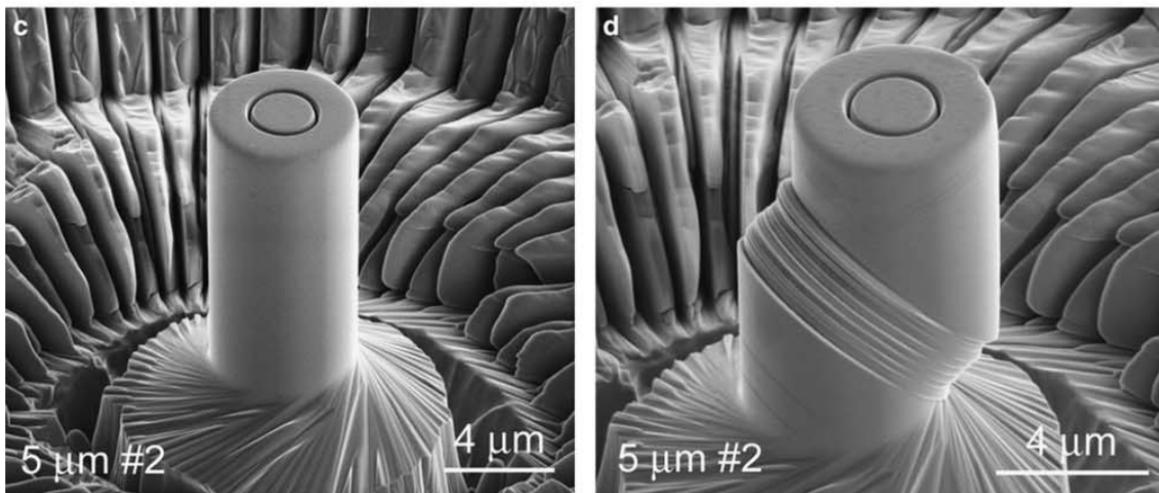


Figure 3.1: A micron-scale sample subject to a compressive force, clearly showing the slip planes along which dislocations moved during the test [103]. Left: Before. Right: After.

A proposed explanation for the statistics of breakdowns and the observed dependence of breakdown rate on electric field is given in [104], where it is suggested that the nucleation of breakdowns is driven by dislocation dynamics within the copper material. The movement of the dislocations themselves is driven by the pulsed electric stress applied to the surface. The initial stages of a breakdown are thought to involve a region of

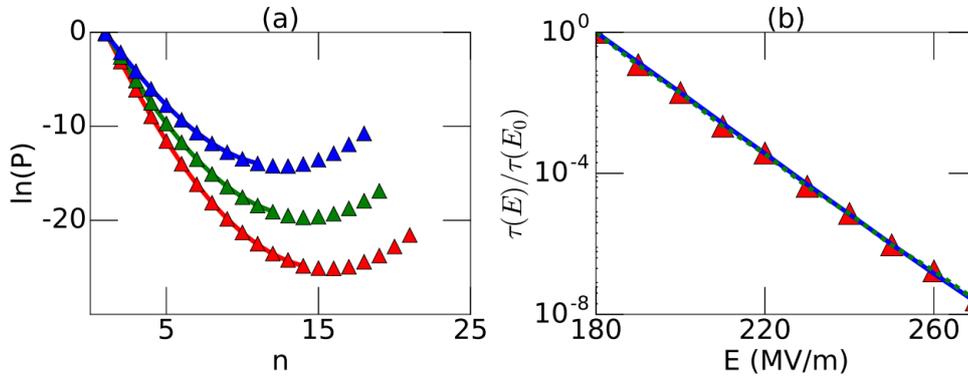


Figure 3.2: Left: Logarithm of the probability distribution vs. the number of mobile dislocations in the slip plane n [104]. Red: for a surface electric field of 200 MV/m. Green: for a surface electric field of 230 MV/m. Blue: for a surface electric field of 260 MV/m. Right: mean time to breakdown τ in arbitrary units vs. surface electric field in MV/m. Red: obtained from Monte Carlo simulations. Blue and green: calculated analytically.

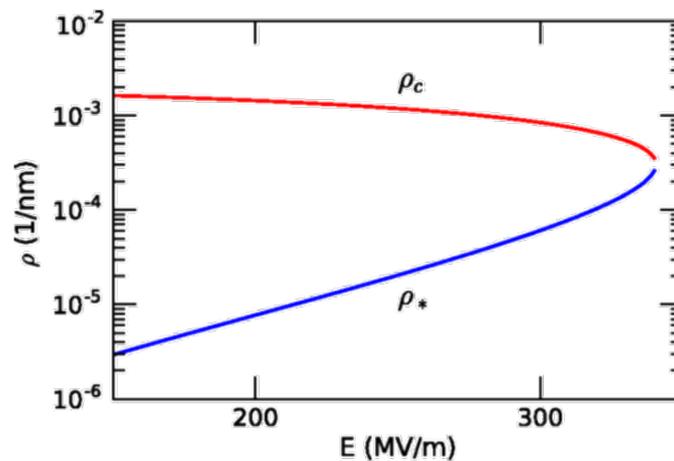


Figure 3.3: Dislocation density in 1/nm vs. surface electric field in MV/m. Red: the unstable fixed point ρ_c . Blue: the stable fixed point ρ_* [104].

some field enhancement being deformed by electric field stress, leading to higher fields and a thermal runaway process. This plastic deformation requires the motion of dislocations in the material. Plastic deformation on very small scales occurs in a step-wise and stochastic manner [105], and the kinetics of the deformation under pre-breakdown conditions are described in the model. As an illustration of the discretised behaviour of deformation on small scales, caused by the presence of dislocation slip planes, Fig. 3.1 shows a micron-scale sample before and after deformation in a compressive test.

The model describes the population of mobile dislocations in one slip plane, with their average rates of creation and depletion being denoted by ρ_+ and ρ_- respectively. The probability distribution of the number of mobile dislocations in a single slip plane is shown in Fig. 3.2. Only mobile dislocations are relevant to this model, meaning that mobile dislocations that become trapped contribute to the depletion rate. These values depend on both the current density of mobile dislocations ρ , and the surface stress σ caused by the applied electric field. There are two fixed points of ρ in this dynamical system, a stable point ρ_* and an unstable one ρ_c , where $\rho_* < \rho_c$, as illustrated in Fig. 3.3. If at any point $\rho > \rho_c$, a runaway increase in ρ occurs, corresponding to the formation of a protrusion from the surface that constitutes the first stage of a breakdown.

As this model is stochastic, the probability that a breakdown occurs is nonzero for any value of electric field, and increases with surface electric field in a manner that fits experimental data. This is shown in the plot of mean time to breakdown (directly proportional to the inverse of the breakdown rate) as a function of surface electric field in Fig. 3.2. The model also explains how conditioning can occur as a result of pulsing electric fields without breakdowns, by presenting it as a work-hardening process caused by the repeated application of stress from the electric field. In the dislocation model, the hardening of the material makes it more difficult for dislocations to move, making it more resilient to breakdown.

It should be noted that this model only attempts to explain the initial stages of the nucleation of a breakdown (see Sec. 1.6). A complete breakdown event is vastly more energetic than the dislocation motion processes considered here, requiring other self-sustaining processes to drive it from a small nucleation event. However, if the model is accurate and the field is sufficiently high and there is sufficient power available for the later stages to occur, the statistics of observable breakdowns will depend on the statistics of the nucleation events.

3.1.3 Dark Current Fluctuations

The work described in this chapter is an attempt to directly measure the variations in dislocation density described in Section 3.1.2, under stable, low-breakdown-rate conditions. Since dislocations tend to propagate to the surface as this is the lower energy state, there should be minute fluctuations in the geometry of a surface subjected to high electric fields, which should in turn result in fluctuations of the field emission characteristics of the sample.

A direct measurement of the dislocation dynamics would serve as strong evidence that the dislocation motion model is correct and could further improve understanding of the processes that lead to breakdowns, allowing better high-gradient structure designs in the future. Since the dynamics of the fluctuations are predicted to change when the system is closer to breaking down, measuring this on a structure under conditioning could also allow for an optimised conditioning process. Instead of allowing breakdowns to occur as they currently do in the conditioning process, the dark current fluctuations could be measured in real time to predict if a breakdown is imminent, and adjust the applied power and thus the surface field accordingly. Such a mode of operation could, in principle, allow structures to be conditioned to high fields without any breakdowns occurring, resulting in much less accumulated damage and potentially higher performance.

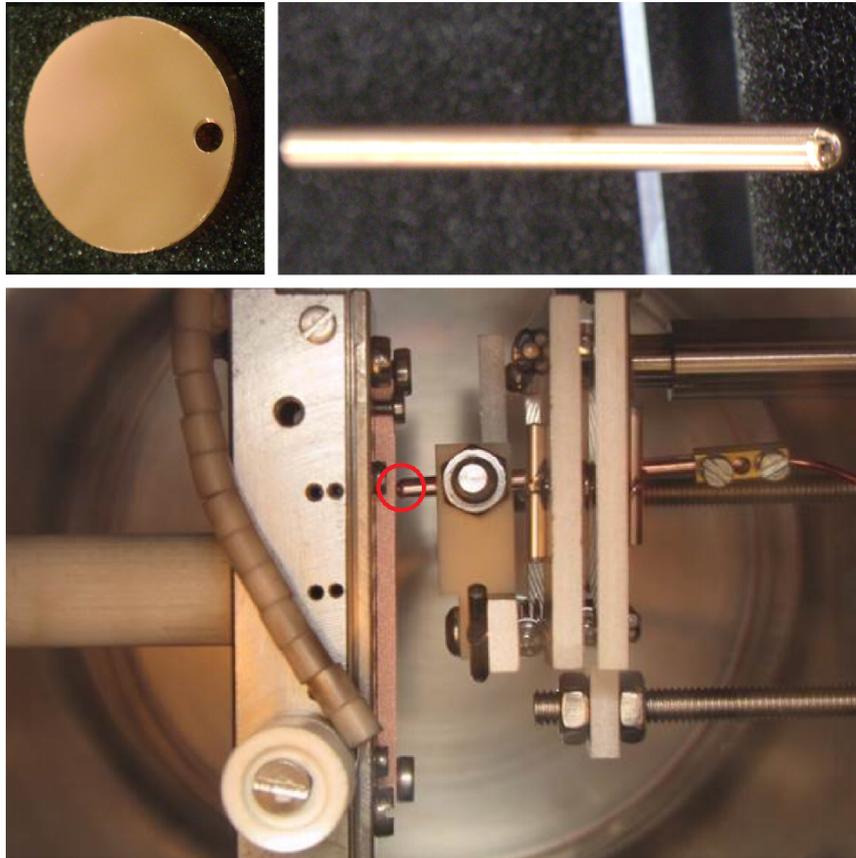


Figure 3.4: Experimental setup of the CERN DC Spark System I. Top left: flat cathode. Top right: cylindrical anode with hemispherical tip. Bottom: assembly in the vacuum chamber with a gap of about $20\ \mu\text{m}$ [22]. The location of the gap is marked with a red circle.



Figure 3.5: Current vs. time in the CERN DC Spark System I. Left: at an applied voltage of 750 V. Middle: at an applied voltage of 800 V. Right: at an applied voltage of 850 V [106].

3.1.4 Past Measurements

Some preliminary measurements of fluctuations in dark current that could be relevant for dislocation dynamics were performed in the CERN DC Spark Systems I and II [106], [107]. These systems featured a cylindrical anode with a hemispherical tip and a flat cathode, held about $20\ \mu\text{m}$ apart under ultra-high vacuum (UHV) conditions, depicted in Fig. 3.4. The measurements showed fluctuations of the dark current between several discrete levels, which rapidly increased in amplitude as the field was increased, as shown in Fig. 3.5.

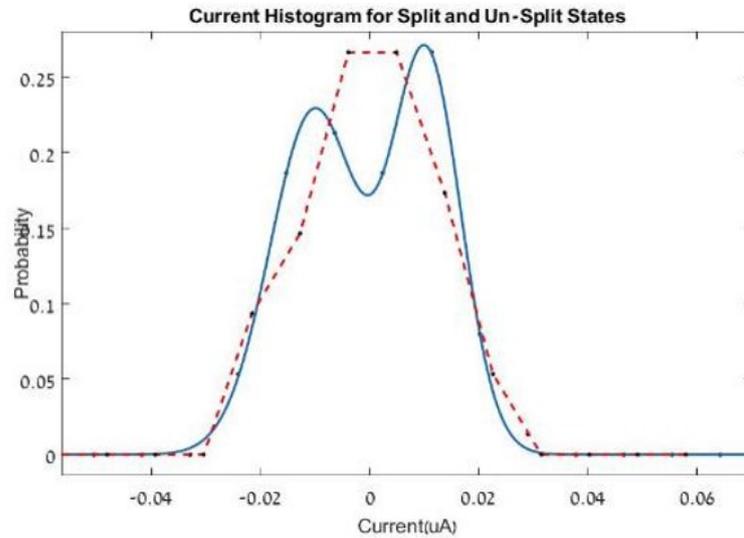


Figure 3.6: Probability distribution vs. current in μA in the CERN DC Spark System I [107]. Red dashed curve: at low voltage. Blue solid curve: at high voltage. The mean current was subtracted in each case to give a zero mean.

Fluctuations in surface geometry due to dislocation dynamics have been suggested as an explanation for this observed behaviour [108]. Fig. 3.6 shows a probability distribution of the current values from this measurement, demonstrating that the system jumps between discrete levels of current at the high electric fields. The presence of discrete levels of current supports the dislocation dynamics hypothesis, in which the changes in the geometry of field emitters are inherently quantised.

Another potentially relevant experimental observation was made with copper electrodes cooled to cryogenic temperatures [109]. In this experiment, a pair of electrodes was first conditioned at a temperature of 60 K, after which the current-voltage curve was

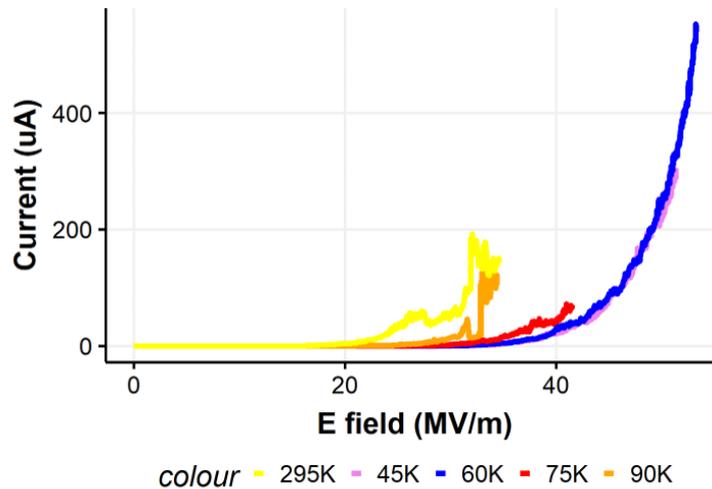


Figure 3.7: Field-emitted current in μA vs. applied electric field in MV/m in a cryogenic experiment [109]. Each curve represents a voltage scan at a given temperature in which the voltage was increased gradually until a breakdown occurred. Higher maximum voltages could be reached at cryogenic temperatures as compared to room temperature, demonstrating the dependence of temperature on breakdown rate. Yellow: at 295 K. Pink: at 45 K. Blue: at 60 K. Red: at 75 K. Orange: at 90 K.

measured at different temperatures without any further conditioning. In each case, the voltage was increased until a breakdown occurred. These current-voltage curves can be seen in Fig. 3.7. Two observations can be made from this plot: one is that at lower temperatures, a higher field could be reached before a breakdown occurred. The other is that the measurements taken at 45, 60, and 75 K appear much smoother than the ones taken at 90 and 295 K. It is yet to be determined if the smoothness of the curves is related to the fluctuations hypothesised by Engelberg et al. [104], though an argument for this is made in Sec. 3.4. However, there does appear to be a correlation between the smoothness of the current-voltage relationship and the surface electric field at which breakdown occurred. If the jaggedness of the curve is indeed a manifestation of the fluctuations being searched for, then this experimental result constitutes further evidence in support of the fluctuation hypothesis.

3.1.5 Approach

It was thought that the initial observations discussed in the previous section were worth investigating in a dedicated study. Past measurements suggest that these events are detectable by measuring dark current with time at a fixed surface field. The rate of fluctuations is expected to depend very strongly on the field, a prediction which was used in this analysis to distinguish them from noise or other phenomena.

The objectives of this study were to validate and improve on past measurements of fluctuations by others, characterise their statistics as a independent test of the dislocation model, and then measure them in RF structures as an indicator of conditioning state, and therefore to guide control the conditioning process.

This measurement was effectively a search for fluctuations in the local surface geometry of the field emitting-sites of the structure. The magnitude of field-emitted current is sensitive to both the area of the emitter and the local electric field (see Sec. 1.7). A change in geometry can change both, which could be visible as a change in the current signal when held at a fixed macroscopic electric field.

The proposed model requires that the number of mobile dislocations goes from the initial state $n = 0$ to the critical state $n = n_c$, where n_c is on the order of 10, as can be seen in Fig. 3.2. Breakdowns are known to occur when CLIC-like X-band RF structures are conditioned with high-power RF pulses of 50 ns or more in duration. Past experiments with 30 GHz cavities showed breakdowns with pulses as short as 4 ns [110]. Based on these two pieces of information, the time interval between changes in n in a given slip plane must be on the order of 0.5 ns or less, corresponding to an event rate of 2 GHz. The actual event rate in a practical experiment is likely many orders of magnitude higher, firstly due to the random walk nature of the trajectory of n to breakdown, and secondly due to the presence of a very large number of slip planes exposed to high electric fields which could potentially undergo breakdown. It is also likely, however, that not all such events are observable in practice. The practical

implication of these considerations is that the event rate is not known *a priori* and as broad a bandwidth as possible should be studied when measuring this phenomenon.

The search for dark-current fluctuations was performed in two separate experiments, one with an oscillating electric field in the RF accelerating structures in the XBoxes, and the other with a DC electric field with the LES. High electric fields and dark current are present in both systems. The results obtained from the RF measurements will be detailed in Sec. 3.2, whereas those from the LES will be discussed in Sec. 3.3.

3.2 Searches in RF Structures

3.2.1 Experimental Setup

RF structures under test in the XBoxes, described in Sec. 1.4, are known to emit dark current. A series of precision measurements of this current was performed to look for evidence of the hypothesised fluctuations.

Field emission is believed to be the mechanism behind dark current, as the measured dependence of the emitted current magnitude on the surface electric field matches that expected from field emission theory (see Fig. 3.16). As field emission is a very strong function of the local electric field, small variations in local surface field due to imperfections in the geometry are expected to dominate the current emitted by the structure's surface. These originate from a number of small emitting sites, which correspond to locations with a slightly higher surface field. Some evidence for this is shown in Sec. 3.2.3.

The structures under test in the XBoxes were subject, when fully conditioned, to high-power pulses of about 50 MW and up to 200 ns in duration [18]. This corresponds to peak surface electric fields in excess of 250 MV/m, resulting in a measurable field-emitted current from every pulse. Most of the dark current from an RF structure is likely to be emitted close to the iris, where the electric field is highest. Fig. 2.14 depicts the concentration of electric field at the iris of a typical X-band accelerating structure.

The emitted electrons are then captured and accelerated by the electric field inside the vacuum volume, and exit the structure along the beam axis, where they can be measured.

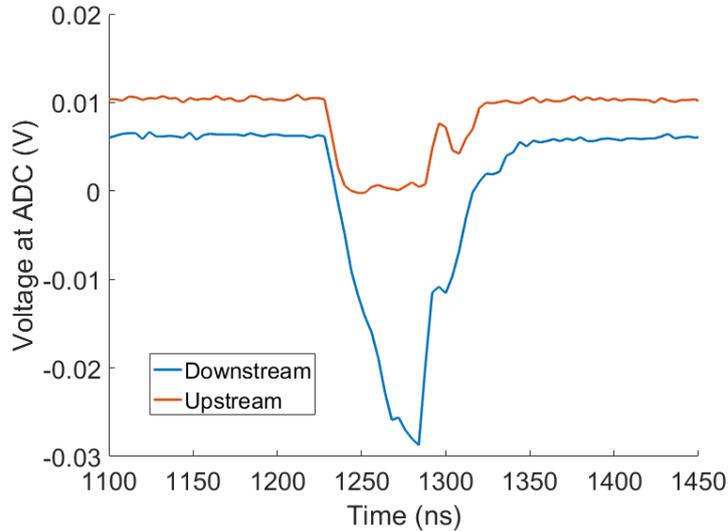


Figure 3.8: Voltage measured by the ADC in volts vs. time since the start of the pulse in ns, with the T24PSI2 structure in Xbox 2 subject to 100 ns long high-power RF pulses. Due to the SLED-type pulse compression scheme [111], peak power is attained only in the final 100 ns of the full RF pulse lasting about 1 μ s. Blue: signal from the downstream Faraday cup. Red: signal from the upstream Faraday cup.

In order for field-emitted electrons to be captured and transported to the end of a long accelerating structure, they have to be accelerated to the phase velocity v_p of the structure. For electrons of low initial energy in a structure with $v_p = c$, this can only occur above a certain minimum gradient $E_{acc} \geq 1.6/\lambda$, where λ is the RF free-space wavelength in m and E_{acc} is in units of MV/m [32]. This is 64 MV/m for an RF frequency of 12 GHz. As captured electrons will get accelerated in the downstream direction, the signal from the downstream Faraday cup is usually larger than the one from the upstream one. On the other hand, the current measured at the upstream Faraday cup is expected to be predominantly composed of electrons originating from the first cell, which have not been captured by the structure and likely have relatively low energy. Since capture is not needed for an electron to reach the upstream Faraday cup, it is also likely that a larger fraction of the electrons emitted by the first cell are measured. Typical measured signals are shown in Fig. 3.8. In this measurement, the

voltage measured by the ADC V_{meas} and the current captured by the Faraday cup I_{meas} are related by the load impedance Z_{load} (in this case $25\ \Omega$) connected to the Faraday cup as follows:

$$V_{meas} = Z_{load}I_{captured} \quad (3.1)$$

The test stands were equipped with Faraday cups on the beam axis up- and downstream of the structure, which collect the emitted charge and allow it to be measured using external electronics. They, along with their associated electronics, were installed with the intention of measuring both dark current during normal pulsing and the much larger currents produced by breakdowns. These Faraday cups were used to measure the dark current emitted by the structure to search for evidence of fluctuation.

The most attention was paid to the signal from the upstream Faraday cup rather than the downstream one. This was due to the concern that, in a structure with multiple cells, each likely containing a multitude of field emission sites, the independent fluctuations in current from each emission site might average out making it more difficult to observe individual fluctuation events. The choice to focus on the signal from the upstream Faraday cup was an attempt to restrict the problem to just the first cell.

The T24PSIN2 structure used for this experiment was conditioned first to ensure that high fields could be reached with a low probability of breakdown during the fluctuation measurement. The conditioning history of this structure is shown in Fig. 1.6. The measurement itself consisted of scanning the RF power delivered to the structure, thus varying the peak electric field inside it, and recording the signal at the upstream Faraday cup for a number of pulses at each power setting.

The pulse repetition rate was reduced from 50 Hz to 5 Hz for the duration of the measurement in order to reduce the average power dissipated by the pulse compressor, and thus reduce the temperature variation that it was subject to over the course of the power scan. Since the pulse compressor included a pair of very high-Q resonant

cavities [18], a small temperature variation on the order of 0.1 °C was enough to change the volume of the cavity enough to detune it significantly and thus modify the pulse shape. The pulse compressor cavities had temperature-controlled water cooling jackets to help maintain a constant temperature, but there was still some residual dependence of the pulse shape on the time-averaged RF power. The reduced repetition rate reduced the maximum average power by a factor of ten, also reducing the difference between the maximum and minimum average RF power delivered during the scan, in turn minimising the detuning of the cavities over the course of the scan.

3.2.2 Signal Acquisition

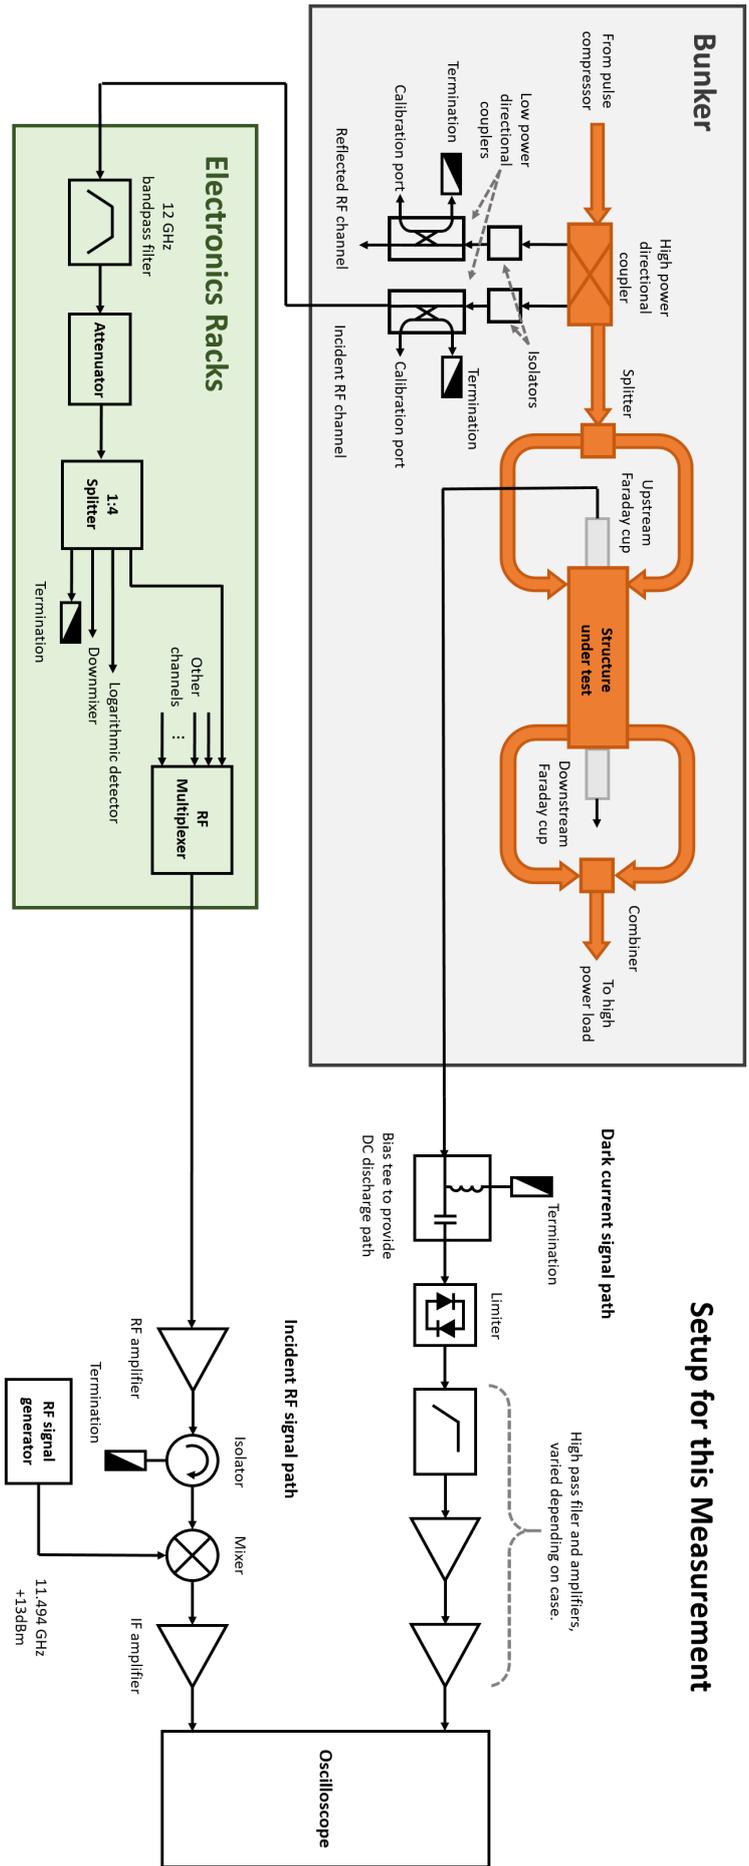
For this measurement, the hardware setup of the test stand had to be modified from its usual configuration for the conditioning of structures (see Fig. 3.9). A Tektronix MDO3104 oscilloscope [112] was used to record the relevant signals instead of the digitisers used in the normal operation of the test stand, to obtain higher resolution in both time and voltage. The use of a standalone instrument also offered better flexibility in setting up the measurement. The oscilloscope featured a sample rate of 2.5 GS/s and an analogue bandwidth of 1 GHz. The new signal paths are shown in relation to the rest of the test stand in Fig. 3.9. The signals of interest were the current captured by the upstream Faraday cup, and the incident RF signal, from which the surface electric field can be determined. The use of the same instrument to record both channels was beneficial for synchronisation between the data samples, the importance of which is discussed in Sec. 3.2.3.

Dark Current

The signals from the Faraday cups were normally read by a National Instruments 5761 digitiser card with sample rate of 250 MHz [113]. For this measurement, the signal from the Faraday Cup was re-routed as follows:

The Faraday cup was connected to the RF+DC port of a bias tee, whose DC port was terminated with a resistor. This provided a DC discharge path to ground for any

Figure 3.9: Diagram of the hardware setup for dark-current fluctuation measurements in an RF structure, showing the signal paths of both the Faraday-cup signal and the incident RF.



charge accumulated by the Faraday cup. This was necessary as some of the components further along the signal path were AC-coupled and thus were an open circuit for a DC current. Without the bias tee, the charge collected by the Faraday cup would build up over time and eventually cause a breakdown in the acquisition electronics (unrelated to vacuum breakdown or dislocation dynamics). A limiter was connected after the bias tee to protect the sensitive amplifiers and oscilloscope from the potentially very large signals which could appear on the Faraday cups in the event of a breakdown.

As no assumptions could be made about the frequency content or amplitude of the fluctuation signals, attention was paid to minimise the noise floor and maximise the bandwidth of the measurement for the best chances of a successful measurement. The internal random noise of the oscilloscope is specified in Table 3.1 for different voltage settings. At its most sensitive setting, it has a root-mean-square (RMS) value of 0.179 mV. This is much larger than the fundamental limit set by Johnson-Nyquist thermal noise [114]. Thermal noise is caused by the thermal agitation of charge carriers in a conductor and is always present, placing a fundamental limit on the precision with which any electronic measurement can be made. Its total power within a given bandwidth B is given by:

$$P_{N,th} = k_B T B, \quad (3.2)$$

where k_B is Boltzmann's constant, and T is the temperature of the system. At room temperature, and using the oscilloscope's full bandwidth of 1 GHz, this corresponds to a power of 3.8×10^{-12} W, or an RMS voltage of 13.8 μ V for a characteristic impedance of 50 Ω . The noise figure of the oscilloscope at the lowest scale setting is thus 22.2 dB, assuming the oscilloscope's input noise has a uniform Power Spectral Density (PSD). The noise figure is a quantity that describes the degradation in the signal-to-noise ratio (SNR) caused by an electronic instrument when compared to the hypothetical case in which the instrument introduces no noise apart from thermal noise, with 0 dB representing a noiseless instrument.

Vertical Scale Setting [mV/div]	RMS random noise [mV]
1	0.179
100	2.40
1000	24.7

Table 3.1: Input noise of the Tektronix MDO3104 oscilloscope [112].

A substantial improvement can be made by using a low-noise amplifier, with much lower internal noise voltage, to amplify the signal before sending it to the oscilloscope such that the oscilloscope's internal noise is no longer a significant contributor to the signal to noise ratio. The Mini-Circuits ZKL-2R5+ [115] was used as a low-noise preamplifier to this end. It has a noise figure of about 5 dB, a maximum gain of 31 dB, and a 3 dB bandwidth of 2.5 GHz.

The overall noise factor F_{tot} (the noise figure in linear units) of a cascade of amplification stages, where the i^{th} stage has noise factor F_i and gain G_i , is given by [68]:

$$F_{tot} = F_1 + \frac{F_2 - 1}{G_1} + \frac{F_3 - 1}{G_1 G_2} + \dots + \frac{F_n - 1}{G_1 G_2 \dots G_n} \quad (3.3)$$

With the oscilloscope used, using one preamplifier resulted in an overall noise figure of 9.0 dB and using two in series brought the value down to 5.3 dB. If more amplifiers were to be added, the overall noise figure would tend to the amplifier's own noise figure of 5 dB, hence bringing about diminishing returns. Either one or two preamplifiers were used depending on the case.

Apart from thermal noise, the sensitivity was also limited by the quantisation of the signal by the analogue-to-digital converter (ADC) within the oscilloscope. The ADC assigned one of a fixed number of levels to each time sample, meaning that the maximum error in voltage is half a step. To minimise this error, the signal should be amplified to make use of as much of the ADC's input voltage range as possible without saturating it. The error caused by quantisation can be modelled as another type of noise such that the digitised signal is the sum of the original signal and the quantisation noise. The power spectral density of the quantisation noise thus describes another limit

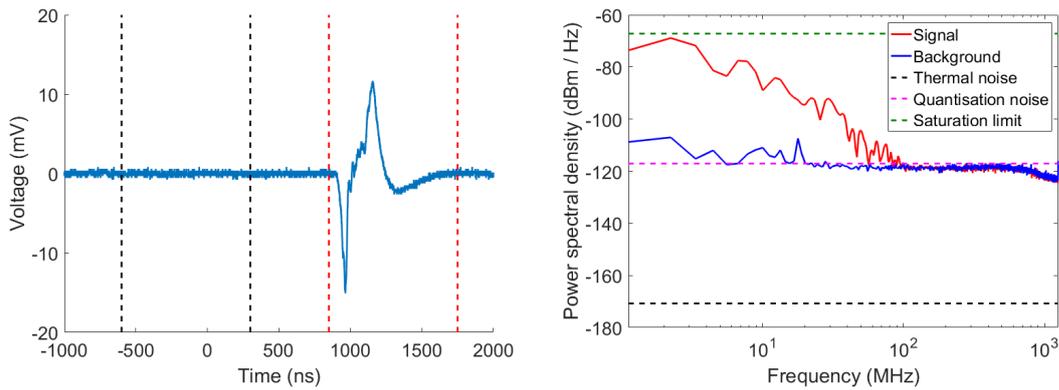
to the measurement's sensitivity. As the amount of preamplifier gain can be varied by selecting the type and number of amplifiers used, this could be either higher or lower than the thermal noise in a given measurement. The signal-to-noise ratio (SNR) for an N -level ADC whose input is a sine wave spanning the full range of the ADC is given by [116]:

$$\frac{P_{sig}}{P_N} = 3N^2/2, \quad (3.4)$$

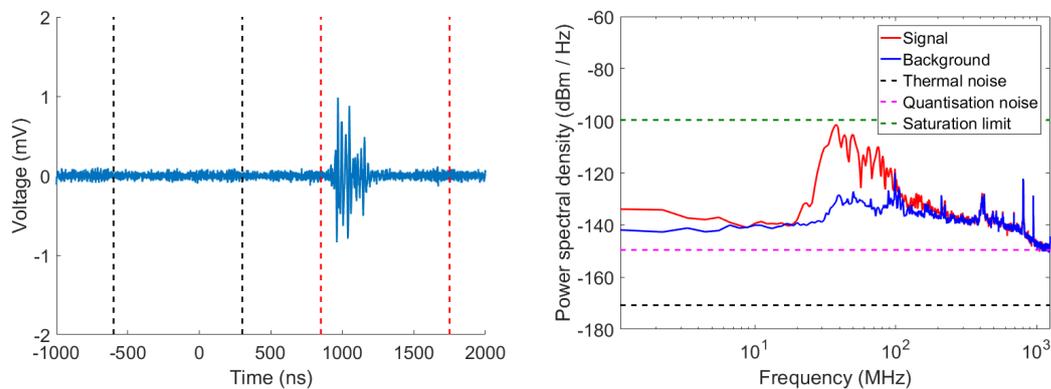
where P_{sig} and P_N denote the signal and noise power respectively. This SNR is 49.9 dB for an 8-bit ADC and 74.0 dB for a 12-bit ADC. The actual power and spectral shape of the quantisation noise depends on the exact signal being quantised. Hence, this value should be considered an estimate for guiding the choice of hardware rather than an exact value.

In itself, quantisation is not a problem. If the signal being measured is small enough, it is always possible to add more gain to the analogue signal chain before the ADC such that the thermal noise exceeds the quantisation noise, making the quantisation noise unimportant. However, too much gain can cause the ADC to saturate, losing information about the signal. Saturating the ADC results in a distorted waveform that has a frequency spectrum unrepresentative of the original waveform. Thus, an ADC with a large dynamic range (number of levels) is desirable, allowing a larger signal amplitude for a given level of quantisation noise.

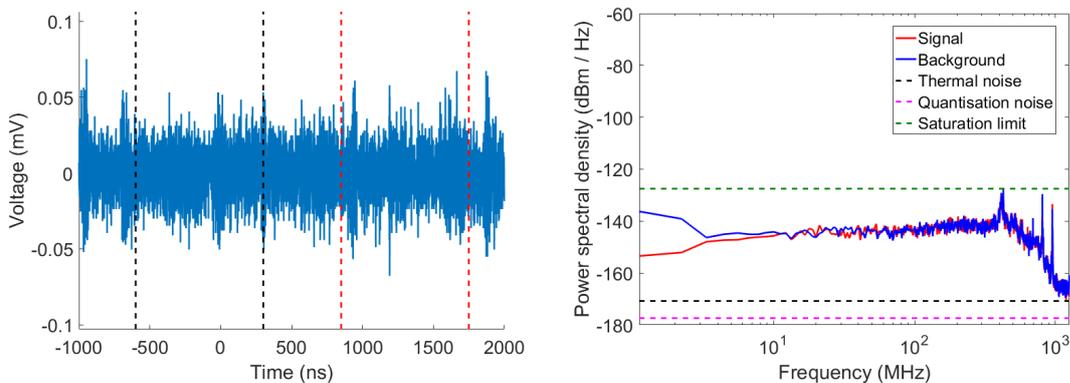
A high-pass filter was used to get around this to some extent. It was noticed that the dark-current signal was usually a relatively level pulse with some ripples. An example of this is the upstream dark current signal in Fig. 3.8. Removing the lower-frequency components of the signal removes the DC offset of the 'plateau' of this signal whilst still preserving the short-timescale features of the signal. This allowed the use of more gain without saturating the oscilloscope, thereby lowering the quantisation-noise floor and allowing the higher-frequency components of the signal to be investigated in more detail. Figure 3.10 shows signals in the time and frequency domains when



(a) 30 dB gain and no high pass filter. Some high-pass filtering still occurred as the amplifiers were AC-coupled.



(b) 50 dB gain, 50 MHz high pass filter.



(c) 60 dB gain, 400 MHz high pass filter.

Figure 3.10: Measured dark current signals, in the time and frequency domains, with different configurations of preamplifiers and filters. Rough estimates of thermal and quantisation-noise power density are shown, along with the maximum signal power before saturation. ‘Thermal noise’ includes noise from the preamplifiers but not the oscilloscope’s internal noise. The signal levels have been normalised to their estimated level at the Faraday cup. The black markers in the time-domain plot indicate the window which was used for the ‘Background’ in the frequency-domain plot, and the red markers indicate the window used for the ‘Signal’. It can be seen that removing the large low-frequency components allowed more gain to be used without saturating the oscilloscope, lowering the quantisation noise level.

sampled with different cutoff frequencies and gains, demonstrating the improvement in sensitivity with the appropriate high-pass filter.

Incident RF

As discussed above, the surface electric field within the structure had to be measured precisely alongside the current to be able to disentangle the contribution of voltage and surface fluctuations to the surface current. This was done by sampling the high-power RF pulse incident to the structure under test using a high-power directional coupler. The signal from this and other couplers was used in the normal operation of the test stand for feedback control of the power delivered to the structure, and was taken into account during routine data analysis. The path taken by the signal from the output of the high-power directional coupler to the electronics racks of the XBoxes can be seen in Fig. 3.9. The outputs from the downmixers and logarithmic detectors for each channel were sampled by data-acquisition cards in the National Instruments PXI crate controlling the test stand [18].

It was thought prudent to avoid disconnecting or interfering with the outputs of the downmixers to avoid disrupting the operation of the test stand. However, the output of the RF multiplexer shown in Fig. 3.9, normally used for calibration, was available and could be set to provide a suitably attenuated and filtered version of the incident RF signal. An isolator was placed between the external components added to this output to avoid the possibility that the signal from the local oscillator would leak back into the electronics of the test stand.

The 12 GHz RF signal from the multiplexer was not sampled directly as this would have required an analogue to digital converter with a very high time resolution, which would be prohibitively expensive. To avoid this, the signal had to be downmixed to allow a lower sample rate to be used, and demodulated to obtain the amplitude and phase of the signal. Early attempts involved the use of an RF detector diode to obtain the amplitude envelope of the RF pulses. An example of an RF pulse measured using the diode is shown in Fig. 3.12, showing a relatively noisy signal and distortion of the

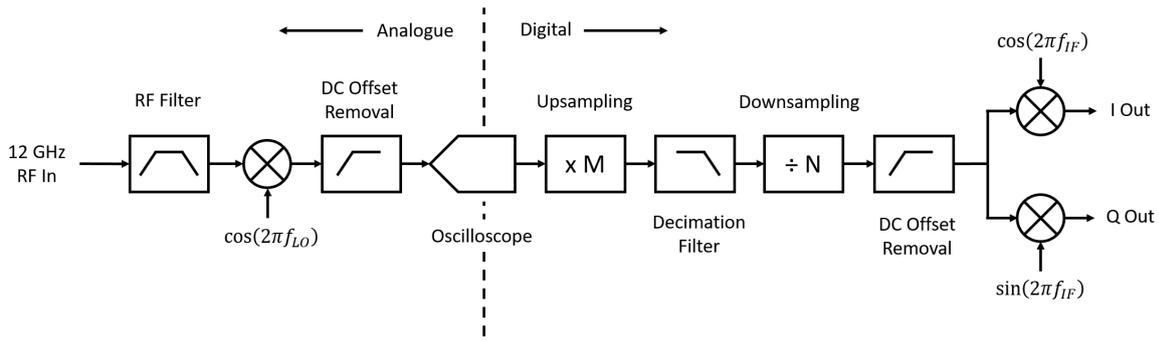


Figure 3.11: A block diagram representing the analogue and digital processing steps performed by the IQ demodulator. The blocks on the left of the oscilloscope represent analogue processing performed by physical components, whereas the blocks on the right represent mathematical operations performed on the waveforms recorded by the oscilloscope during post-processing.

pulse shape (compare with the plot of amplitude vs. time in Fig. 3.14, obtained using the improved IQ demodulator).

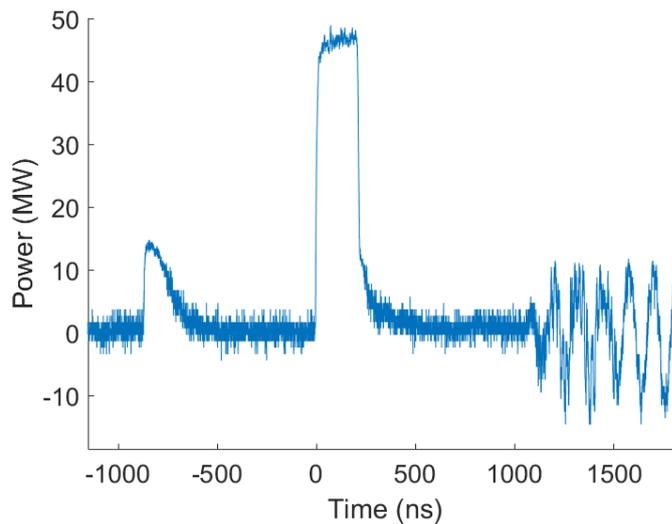


Figure 3.12: RF power indicated by the diode in MW vs. time in ns.

The diode was later replaced with an In-phase and Quadrature (IQ) demodulator, assembled out of modular off-the-shelf RF components, for better signal fidelity. It had the advantages of better linearity than the diode and the ability to provide phase information about the signal. This provided enough information for the initial 12 GHz RF signal from the directional coupler to be fully reconstructed.

A functional block diagram of the IQ demodulator implemented for this measurement is depicted in Fig. 3.11. A local oscillator (LO) frequency of 11.494 GHz was used,

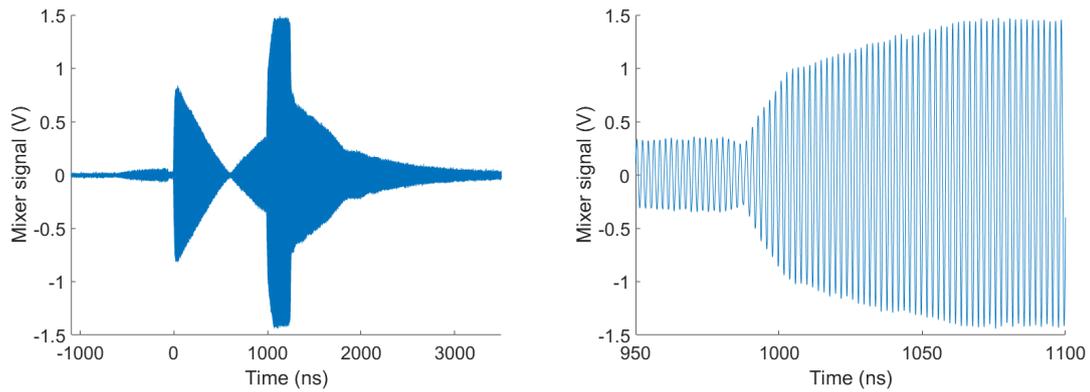


Figure 3.13: Output of the RF mixer in V vs. time in ns. Left: the entire waveform. Right: a zoom on the time axis showing the sinusoidal carrier signal.

generated by an Anritsu MG3692C signal generator [117], producing signals at 500 MHz and 23.5 GHz at the intermediate frequency (IF) port of the first mixer. The input to the RF port of the mixer was bandpass filtered to ensure that no spurious IF signals could be produced by harmonics in the incident 11.994 GHz signal. The unwanted 23.5 GHz product was removed by the input filter of the oscilloscope, requiring no additional components. The 500 MHz IF resulted in a signal well within the 1 GHz bandwidth of the oscilloscope, with ample bandwidth to accommodate the modulation of the incident RF signal. An example is shown in Fig. 3.13.

The IQ demodulation process of the downmixed IF signal can be briefly described as follows: The IF signal, represented as a series of samples in discrete time v_n , with a sample rate f_s , a time-dependent amplitude A_n and phase ϕ_n may be written as:

$$v_n = A_n \cos\left(\frac{2\pi f_0 n}{f_s} + \phi_n\right) \quad (3.5)$$

The second mixing operation can be implemented as a multiplication by cosine and sine signals at the same frequency, yielding the in-phase and quadrature products i_n and q_n respectively, i.e.:

$$\begin{aligned}
i_n &= A_n \cos\left(\frac{2\pi f_0 n}{f_s} + \phi_n\right) \cos\left(\frac{2\pi f_0 n}{f_s}\right) = \frac{A_n}{2} \left(\cos\left(\frac{4\pi n f_0}{f_s} + \phi_n\right) + \cos(\phi_n) \right) \\
q_n &= A_n \cos\left(\frac{2\pi f_0 n}{f_s} + \phi_n\right) \sin\left(\frac{2\pi f_0 n}{f_s}\right) = \frac{A_n}{2} \left(\sin\left(\frac{4\pi n f_0}{f_s} + \phi_n\right) + \sin(\phi_n) \right)
\end{aligned} \tag{3.6}$$

If these are sampled once per period of the IF carrier frequency to give new sequences i'_m and q'_m , the carrier frequency will no longer appear in the result:

$$\begin{aligned}
i'_m &= i_{m \cdot f_s / f_0} = A_m \cos(\phi_m) \\
q'_m &= q_{m \cdot f_s / f_0} = A_m \sin(\phi_m)
\end{aligned} \tag{3.7}$$

The above expressions are valid only if the sample rate is an integer multiple of the IF frequency. From i'_m and q'_m , the instantaneous amplitude A_m and phase ϕ_m can be easily determined. Picking a sample rate f_s that is exactly four times the IF carrier frequency f_0 results in the additional simplification that the local oscillator signals reduce to simple sequences as follows:

$$\begin{aligned}
\cos\left(\frac{2\pi f_0 n}{f_s}\right) &= \{1, 0, -1, 0, 1, 0, \dots\} \\
\sin\left(\frac{2\pi f_0 n}{f_s}\right) &= \{0, 1, 0, -1, 0, 1, \dots\}
\end{aligned} \tag{3.8}$$

Thus, the second mixing stage can be implemented as a simple procedure involving zeroing or negating the appropriate samples of the IF signal v_n . The sample rate of the waveforms recorded by the oscilloscope was changed from 2.5 GHz to 2 GHz by upsampling by a factor of 4, low-pass filtering to avoid aliasing, then downsampling by a factor of 5.

Hence, the modulation of the signal can be described by the complex amplitude $A(t)e^{i\phi(t)}$. The I and Q outputs of the IQ demodulator are the real and imaginary parts of this complex amplitude respectively. An example of the resulting magnitude and phase of an incident RF signal is shown in Fig. 3.14.

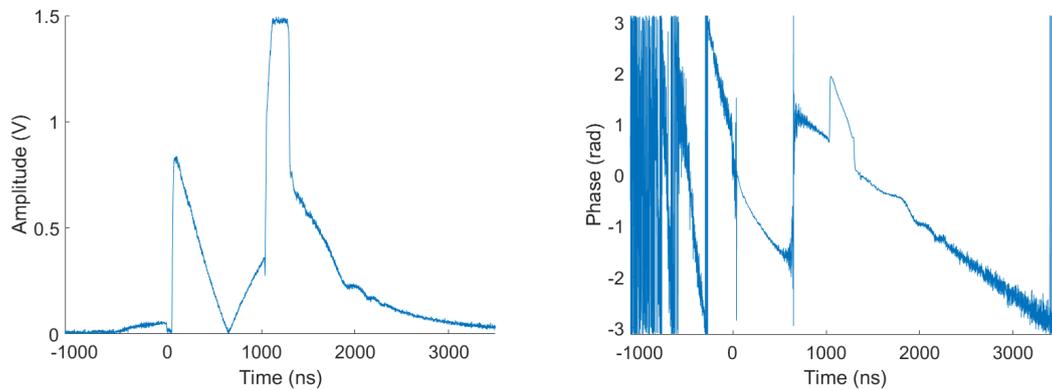


Figure 3.14: An IQ-demodulated RF pulse incident on the structure. Left: instantaneous magnitude in V vs. time in ns. Right: instantaneous phase in radians vs. time in ns.

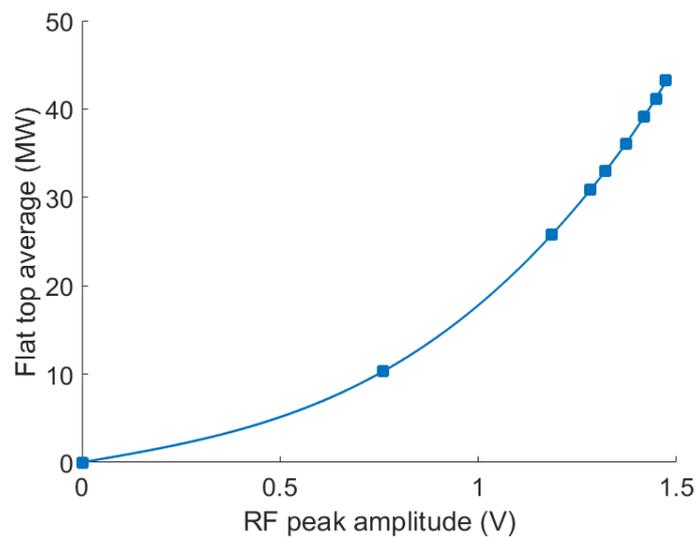


Figure 3.15: Average power of the flat-top part of the signal in MW vs. peak amplitude of the demodulated RF signal in V. Points represent measurements, while the curve is a spline interpolation of the data points.

Before each measurement, a calibration of the demodulator was performed to allow its output voltage to be related to the power incident on the structure. A new calibration was performed before each set of measurements. A calibration was performed by setting the power delivered to the structure by the test stand to a series of values, and recording the corresponding demodulated amplitude A_n , as shown in Fig. 3.15. The resulting empirical relation could then be applied to the demodulated amplitude, given in volts, to obtain the corresponding power in the structure. The incident power could then be used to determine the surface electric field within the structure. The relationship between the two could be obtained from 3D electromagnetic simulations of the structure's geometry. In the case of the T24PSI structure presented here, an input power of 37.5 MW yields an accelerating gradient of 100 MV/m and a peak surface field of 204 MV/m.

3.2.3 Data Analysis

Measurements of Field Enhancement Factor

A basic tool in the study of field emission is the Fowler-Nordheim plot, which shows the dependence of the emitted current on the applied surface field [30]. An example of the result of one of the measurements done in this study is shown, with both linear and Fowler-Nordheim axes, in Fig. 3.16. The gradient of the plot on Fowler-Nordheim axes allows the field enhancement factor β to be determined using the relation:

$$\frac{d(\log_{10} I_F/E^{2.5})}{d(1/E)} = -\frac{2.84 \times 10^9 \phi^{1.5}}{\beta}, \quad (3.9)$$

where I_F is the emitted current, E is the macroscopic surface electric field, and ϕ is the workfunction of the emitting material in eV. As the vertical axis of the plot is logarithmic, scaling the current by a constant factor does not change the slope of the fitted line. Thus, the calculated value of β is not expected to change if, for example, the Faraday cup only captures a fraction of the emitted current. The result presented in Fig. 3.16 shows a very good linear fit on the Fowler-Nordheim axes, demonstrating

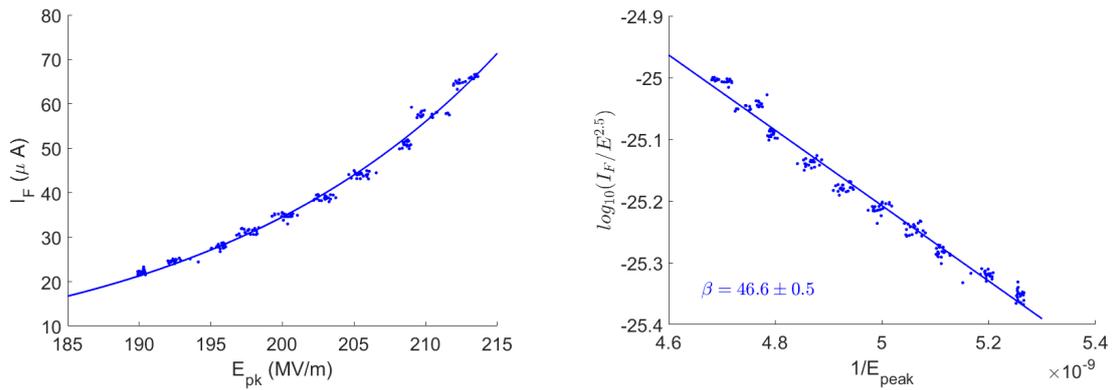


Figure 3.16: A measurement of field-emitted current from an RF structure, with a fitted value of field enhancement factor $\beta = 46.6 \pm 0.5$. Each point represents one RF pulse, with each cluster of points corresponding to one RF power level and the spread being caused by the noise in the system. Left: current in μA vs. peak surface electric field in MV/m . Right: Fowler-Nordheim axes, i.e. $\log_{10}(I_F/E^{2.5})$ vs. the inverse of surface electric field in m/V . Fits were performed by minimising the sum of the squared difference between each of the points and the fitting function.

that the functional form of the experimental current-voltage curve matches the theory of field emission very well and therefore that field emission is the main source of the current measured. The fitted value of $\beta = 46.6 \pm 0.5$ implies that there are regions on the surface of the structure where the surface field is 46 times greater than that expected from the geometry and incident power, a behaviour frequently observed in all macroscopic systems.

Figure 3.17 shows an interesting example of a power scan during which a breakdown occurred. The magnitude of the current was changed significantly by the breakdown, undergoing about a 30% reduction. As breakdowns are typically highly localised, with the resulting craters usually being on the order of $100\ \mu\text{m}$ in size (see Fig. 2.21), this must mean that a region no larger than the size of a breakdown crater must have been responsible for emitting a substantial fraction of the current emitted by the whole structure. This is one way in which this measurement could be sensitive to dislocation phenomena occurring on nanometer scales despite the total surface area subject to high fields being much larger (on the order of cm^2).

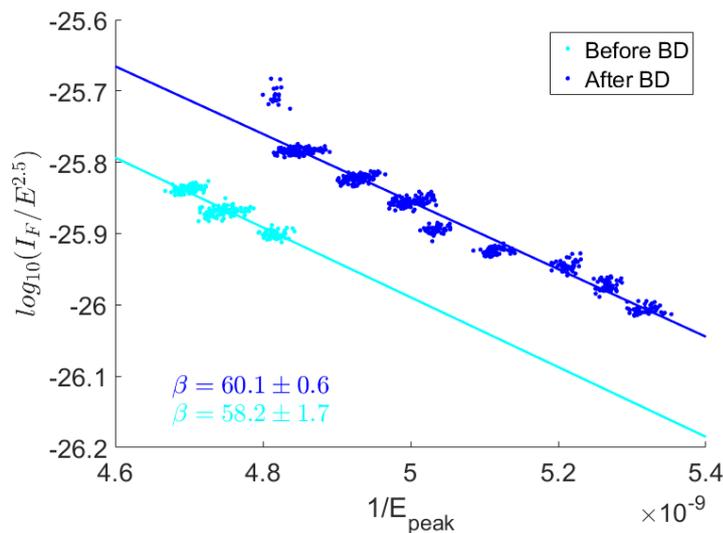


Figure 3.17: A measurement of field-emitted current from an RF structure with Fowler-Nordheim axes, i.e. $\log_{10}(I_F/E^{2.5})$ vs. the inverse of surface electric field in m/V. A breakdown occurred during the measurement. Light blue: data taken before the breakdown, with a fitted field enhancement factor $\beta = 58.2 \pm 1.7$. Dark blue: data taken after the breakdown, with a fitted field enhancement factor $\beta = 60.1 \pm 0.6$. Fits were performed by minimising χ^2 summed over all the data points, with the exclusion of the cluster of outliers at 4.8×10^{-9} m/V which were measured immediately after the breakdown.

Pulse to Pulse Variations

An example of a raw dark-current pulse measured at the upstream Faraday cup can be seen in Fig. 3.8. The shape of the dark-current pulse depends in a complex way on the RF pulse shape, propagation of the RF pulse through the structure, and propagation of the electrons from the emission site to the Faraday cup. All these effects are deterministic, and thus repeatable from pulse to pulse if the incident RF pulses remain perfectly repeatable. On the other hand, the dark-current fluctuations arising from dislocation motion are expected to be random and thus would cause the measured signal to vary from pulse to pulse even under completely repeatable conditions.

To obtain a frequency spectrum of the variations between dark-current pulses, the difference between the coherent and incoherent averages of the pulses recorded for a each power level was calculated. These will be defined mathematically below. The voltage $v_i(t)$ measured at the Faraday cup on the i^{th} pulse can be represented as the

sum of a deterministic component $d(t)$, which is identical for every pulse, and a zero-mean random component $r_i(t)$:

$$v_i(t) = d(t) + r_i(t) \quad (3.10)$$

As the random component has zero mean, averaging the measured signal over a large number of pulses leaves just the deterministic component:

$$\langle v_i(t) \rangle = \langle d(t) \rangle + \langle r_i(t) \rangle = d(t) \quad (3.11)$$

We denote the Fourier Transform of a signal by a capital letter as follows:

$$D(\omega) = FT[d(t)] \quad (3.12)$$

The power spectral density (PSD) of the deterministic component $D(\omega)D^*(\omega)$, where $D^*(\omega)$ denotes the complex conjugate of $D(\omega)$, can thus be obtained taking the Fourier transform of the deterministic component in the time domain:

$$D(\omega)D^*(\omega) = FT[\langle v_i(t) \rangle]FT[\langle v_i(t) \rangle]^* \quad (3.13)$$

This is the coherent average PSD. The PSD of each individual pulse is given by:

$$\begin{aligned} FT[v_i(t)]FT[v_i(t)]^* &= (D(\omega) + R_i(\omega))(D(\omega) + R_i^*(\omega)) \\ &= D(\omega)D^*(\omega) + D(\omega)R_i^*(\omega) + D^*(\omega)R_i(\omega) + R_i(\omega)R_i^*(\omega) \end{aligned} \quad (3.14)$$

Calculating the PSD of each of the pulses individually and then averaging the PSDs of all of the pulses gives the incoherent average PSD of the signal, i.e. :

$$\begin{aligned}
\langle FT[v_i(t)]FT[v_i(t)]^* \rangle &= \langle (D(\omega) + R_i(\omega))(D(\omega) + R_i(\omega))^* \rangle \\
&= D(\omega)D^*(\omega) + \langle R_i(\omega)R_i^*(\omega) \rangle
\end{aligned} \tag{3.15}$$

As the random component has a mean of zero, the terms $D(\omega)R_i(\omega)^*$ and $D(\omega)^*R_i(\omega)$ tend to zero when averaged. The square term $R_i(\omega)R_i^*(\omega)$, however, does not tend to zero as the variance of the random component is not necessarily zero. The difference between the incoherent and coherent averages, $\langle R_i(\omega)R_i^*(\omega) \rangle$, is effectively the PSD of the random component of the signal, i.e. that of the deviations of each individual pulse above and below the mean, and can be considered the variance of the signal at a given frequency.

To obtain an estimate of the background noise, the variance PSD of a time window placed just before the compressed RF pulse, during which no measurable dark current was emitted, was used as a reference to compare with the variance PSD of the dark current pulse.

Since the dislocation hypothesis predicts that the dark-current signals should have a variance that increases rapidly with the applied surface field, the variance PSD $\langle R_i(\omega)R_i^*(\omega) \rangle$ was studied to determine the possible presence of such a signal.

Fig. 3.18 shows a set of dark-current signals for a number of pulses as recorded by the oscilloscope with an incident power of 56.9 MW. The signal from the Faraday cup was filtered with a 50 MHz high-pass filter and amplified using one ZKL-2R5+ preamplifier. The plots of the individual pulses do not line up exactly, demonstrating a small amount of variation from pulse to pulse. The analysis method described above was used to determine if these variations could be attributed to random noise.

The frequency spectra of this measurement set are presented in Fig. 3.24. From this figure, several conclusions can be made. First of all, the ‘Background’ and ‘Background variance’ plots line up very closely, meaning that almost all of the power of the background signal is random. In other words, the deterministic component $d(t)$ in this case

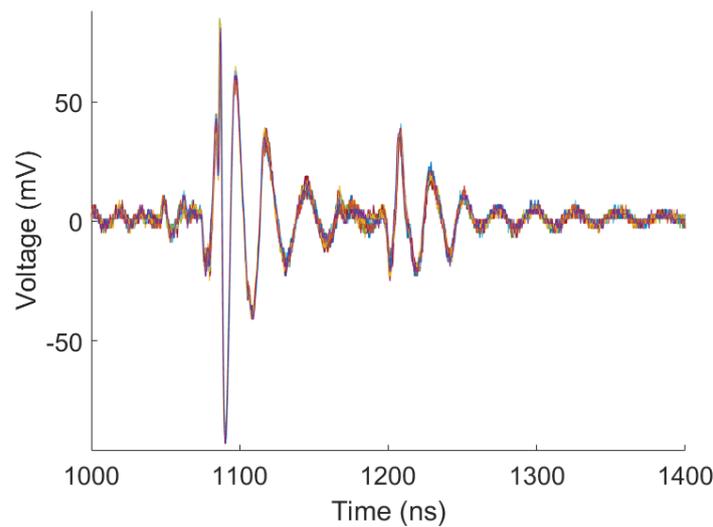


Figure 3.18: Voltage measured by the oscilloscope in mV vs. time in ns. Before being read by the oscilloscope, the signal was first sent through a 50 MHz high-pass filter. The figure shows 23 separate pulses obtained at an incident power of 56.9 MW.

is very small, consistent with what is expected to be random background noise. The level of the background signal is about 22 dB above the thermal noise level, consistent with the previous estimation of the oscilloscope's noise figure. This implies that the most significant source of the background noise is the oscilloscope itself, and that there were no other interfering signals. The dark-current signal itself does appear to have a significant amount of variance from pulse to pulse, as evidenced by the 'Signal variance' plot being well above the background noise. Thus, the background noise from the oscilloscope cannot explain the pulse to pulse differences between the dark-current signals. The 'Signal' plot is higher still, indicating that the dark current signal is the sum of a random and a deterministic component, as can be intuitively deduced from the time-domain plots in Fig. 3.18. The variance signal appears to have a peak at around 40 MHz, which will be further investigated in later sections.

Window Function

A window function was used when calculating the frequency spectra of the signals being analysed to avoid the appearance of spurious frequency components via spectral leakage. Spectral leakage is a result of the way the Discrete Fourier Transform (DFT)

works: it makes the implicit assumption that the time-domain signal being operated on is one cycle of an infinite periodic function [84]. A difference in value between the first and last samples of the time window will be treated by the DFT as a step in the assumed periodic signal, producing a broad spectrum of frequency components which were not present in the original signal and are purely a result of the way the time domain signal was sampled.

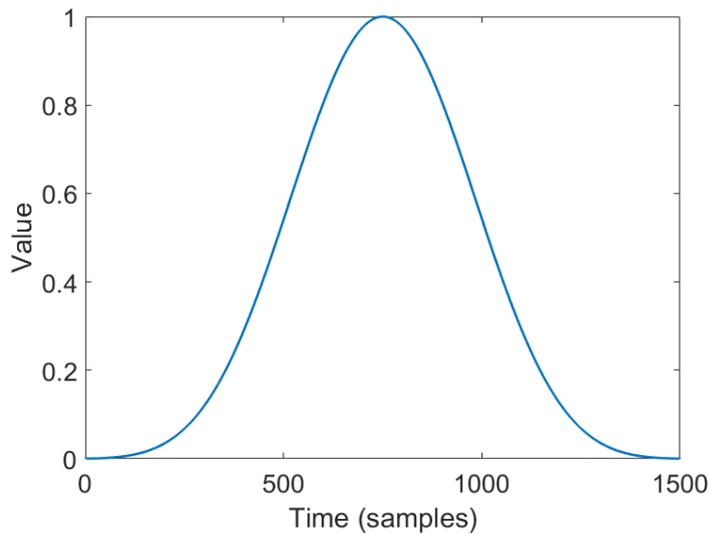


Figure 3.19: The Chebyshev window function vs. time in samples for a window length of 1500 samples.

Spectral leakage was thus an important consideration as it could significantly distort the frequency spectra of the signal and its variance. Care was taken to ensure that the time windows specified did not result in an obvious discontinuity between the first and last samples. However, due to the noise in the signal, it was impossible to guarantee no discontinuity for every pulse in a measurement. To be sure that no spectral leakage occurred in spite of this, a Chebyshev window function [118] was used. The signal was multiplied by this function to ensure that the product tapered smoothly to zero at each end of the window, thus avoiding discontinuities [119]. The Chebyshev window function is shown for a window length of 1500 samples in Fig. 3.19.

An example demonstrating the effect of this window function is presented in Fig. 3.20. Two possible input signals are shown, one with an integer number of cycles within the time window, resulting in no discontinuity, and one with a non-integer number of

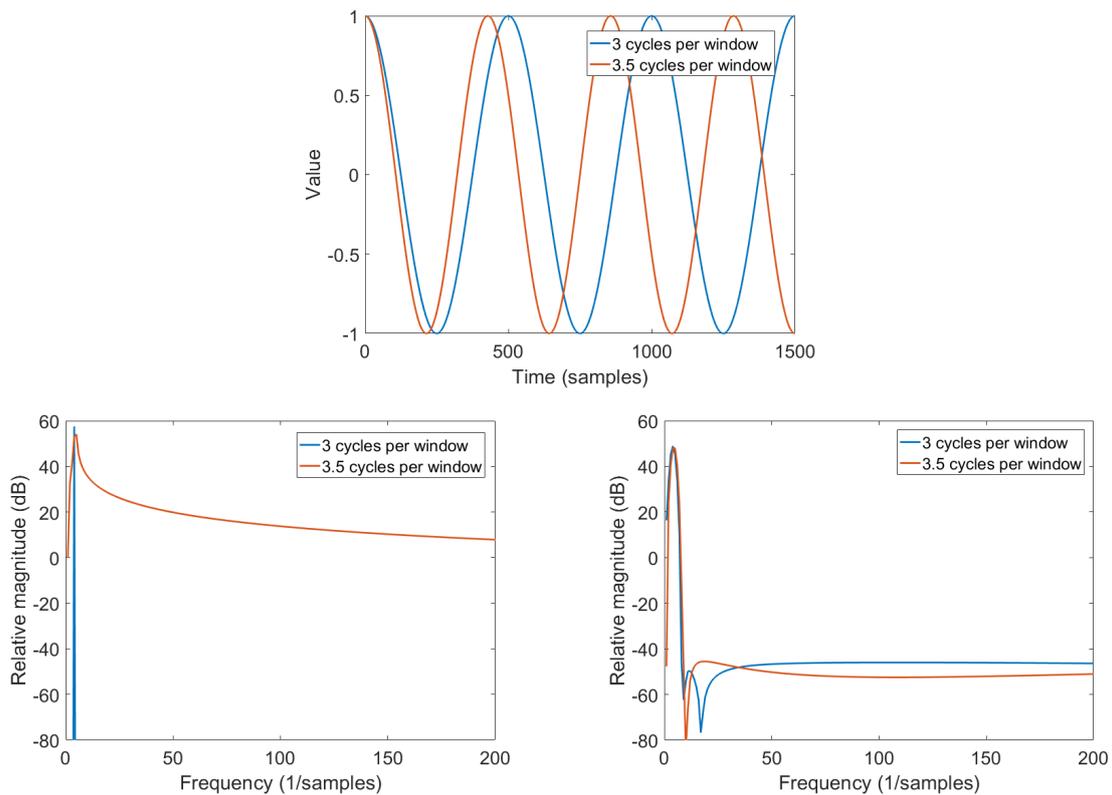


Figure 3.20: An example showing the effect of applying a window function to a signal. In each plot, the blue curve represents a signal with 3 cycles per time window while the red curve represents a signal with 3.5 cycles per time window. Top: Signals to be analysed with a DFT vs. time in samples. Bottom left: relative magnitude of the frequency spectrum in dB vs. frequency in 1/samples, without a window applied. Bottom right: relative magnitude of the frequency spectrum in dB vs. frequency in 1/samples, with the Chebyshev window applied.

cycles producing a large step. The resulting DFTs of the two signals are shown. It can be seen that without any window function, the discontinuity results in a broad spectrum of frequencies that were not present in the original signal. Applying the window function greatly reduced this spectral leakage. The disadvantage of using the window function is the introduction of some uncertainty in the frequency of the original signal, as can be seen from the broadened peaks in the corresponding frequency spectra. This was not considered a significant problem in this measurement as the objective was to determine the presence or absence of a signal and not its exact frequency. It is at any rate unlikely that the fluctuation phenomena being searched for would have had a stable enough frequency for this uncertainty to matter.

Time Alignment of Pulses

The time alignment of the sampled RF and dark-current pulses was an important source of error that needed to be considered. A misalignment in time between otherwise identical pulses appears as a spurious pulse variance, as illustrated with idealised signals in Fig. 3.21.

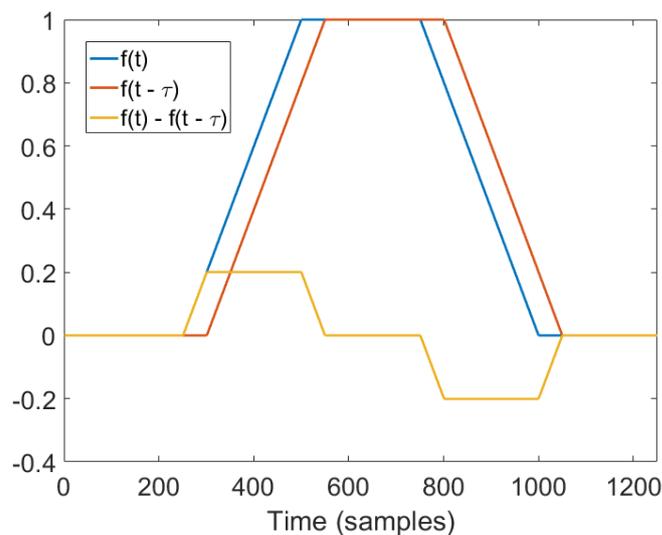


Figure 3.21: Signals vs. time in samples showing an example of the effect of time alignment on the apparent variation between pulses. Blue: an example signal. Red: a time-shifted version of the original signal. Yellow: the difference between the original and time-shifted signals.

It is thus crucial to ensure that there is as little misalignment between the pulses as possible to avoid introducing any apparent variation between the pulses. The oscilloscope used in this measurement used a threshold crossing to trigger the saving of data and synchronised each pulse to the time at which the trigger level was first exceeded. However, the implementation of this trigger appeared not to be perfect, leaving about 900 ps of root-mean-square (RMS) time jitter between the acquired pulses. With a 2.5 GHz sample rate, the sample period is 400 ps - in principle, given perfectly repeatable pulses, it should be possible to align them to less than one sample period.

In order to minimise the error resulting from misalignment, realignment in post-processing was performed. This was done by shifting each pulse in time by a number of samples n_{shift} such that the sum of squares difference between pulses was minimised, i.e.:

$$n_{\text{shift},i} = \arg \min_n \sum_j (x_{i,j-n} - x_{1,j})^2, n \in Z, \quad (3.16)$$

where $x_{i,j}$ is the j^{th} sample of the i^{th} pulse being aligned. The first pulse of the set was used as a reference to align all the other pulses. It was found that the alignment procedure worked more reliably when the smoothed magnitude was used as the input x instead of the raw waveforms. In order to prevent the minimisation method in (3.16) from overfitting to random noise in the signals, the signals were smoothed using a Locally Weighted Scatterplot Smoothing (LOWESS) filter [120] with a 30-sample window before alignment. The LOWESS filter was chosen because it provided a sharper frequency cutoff than a simple moving-average filter, whilst still being easy to adjust for best performance owing to it only having one free parameter. As the signal was AC-coupled, excessive smoothing could attenuate the signal itself, which would hinder the alignment procedure. To prevent this, the absolute values of the signals were smoothed and compared instead. Examples of raw, smoothed, and reference signals are shown in Fig. 3.22.

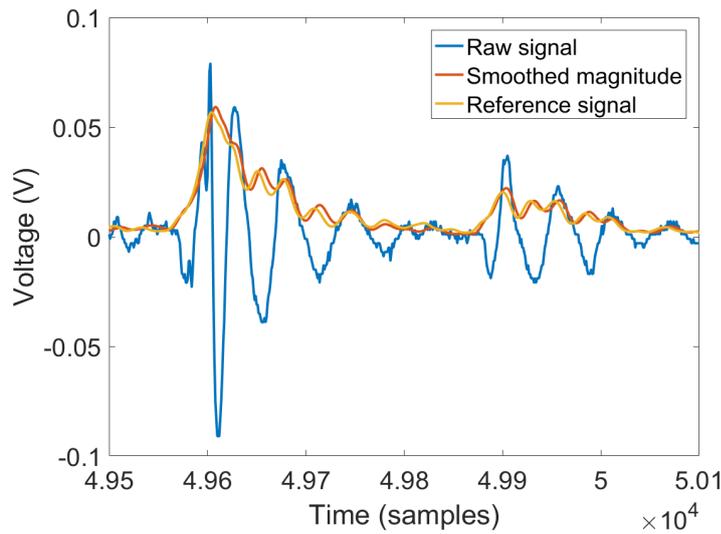


Figure 3.22: Faraday cup signal in V vs. time in samples. Blue: raw signal as measured by the oscilloscope. Red: smoothed magnitude of the raw signal. Yellow: smoothed magnitude of the reference signal. The time-alignment procedure attempts to minimise the sum-of-squares difference between the smoothed magnitude of the pulse and that of the reference signal.

Using this method to realign both the RF and dark-current pulses resulted in a significant improvement in the time jitter to about 200 ps, an example of which is shown in Fig. 3.23. This is half the sample period of the oscilloscope, representing the uncertainty with which the oscilloscope could measure time. Thus, the alignment procedure was believed to be working correctly. The residual uncertainty was thought to be a fundamental limitation caused by phase noise of the oscilloscope’s internal sampling clock, and possibly could have been reduced if a precision frequency reference were used.

The effect of the reduction in time jitter can also be seen in Fig. 3.25, where the difference in the signal variance before and after time alignment is about 10 dB over most of the signal’s bandwidth.

To be able to determine if the variance spectrum is statistically significant, an estimate of the noise spectrum resulting from the residual time jitter of 200 ps RMS was calculated to use as a reference. This spectrum would effectively quantify the uncertainty of the variance measurement as determined by the time resolution of the oscilloscope. This estimate was calculated by taking the PSD of the difference between the mean

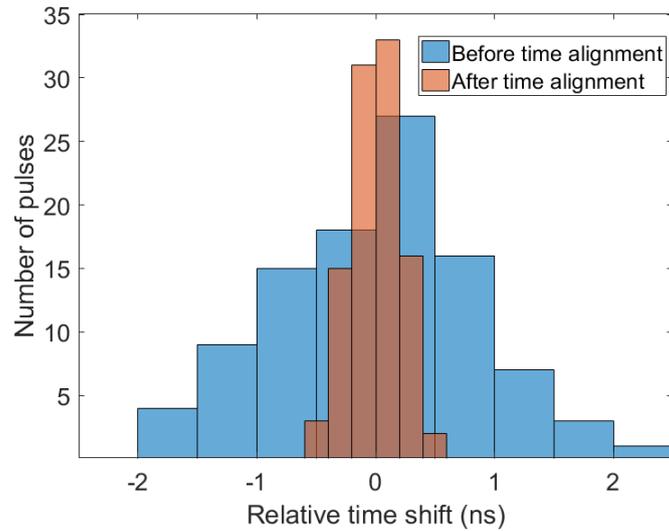


Figure 3.23: Number of pulses vs. relative time shift of the RF pulse amplitude in ns. The timing of each pulse in this plot is defined as the time at which the signal first exceeded 50% of the maximum value found in the waveform. Blue: before post-processing, showing an RMS jitter of about 900 ps. Red: after post-processing, showing an RMS jitter of about 200 ps.

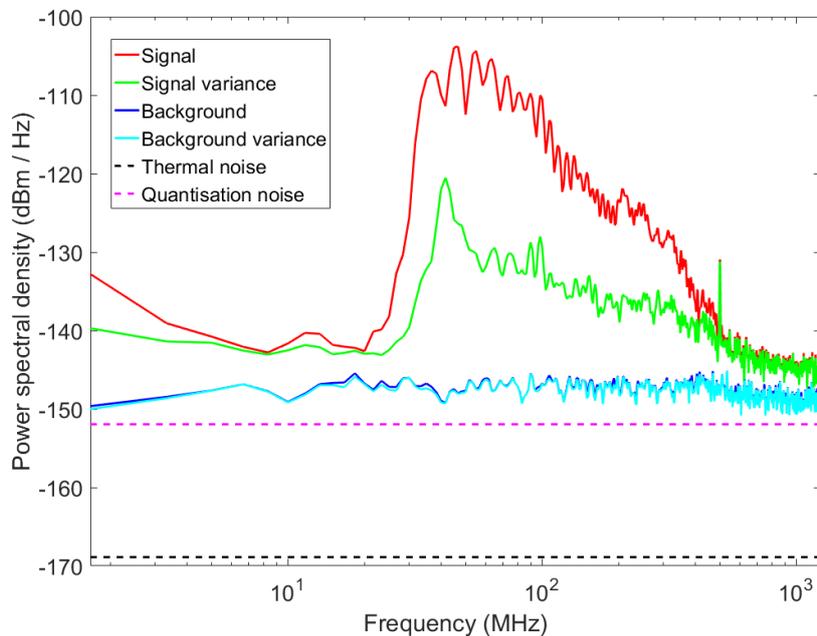


Figure 3.24: Power spectral density of the pulses shown in Fig. 3.18 in dBm/Hz vs. frequency in MHz. A 50 MHz high pass filter was used. Red: total power spectral density of the signal. Green: power spectral density of the random component of the signal. Dark blue: total power spectral density of the background noise. Light blue: power spectral density of the random component of the background noise. Black: thermal noise excluding the input noise of the oscilloscope itself. Magenta: quantisation noise.

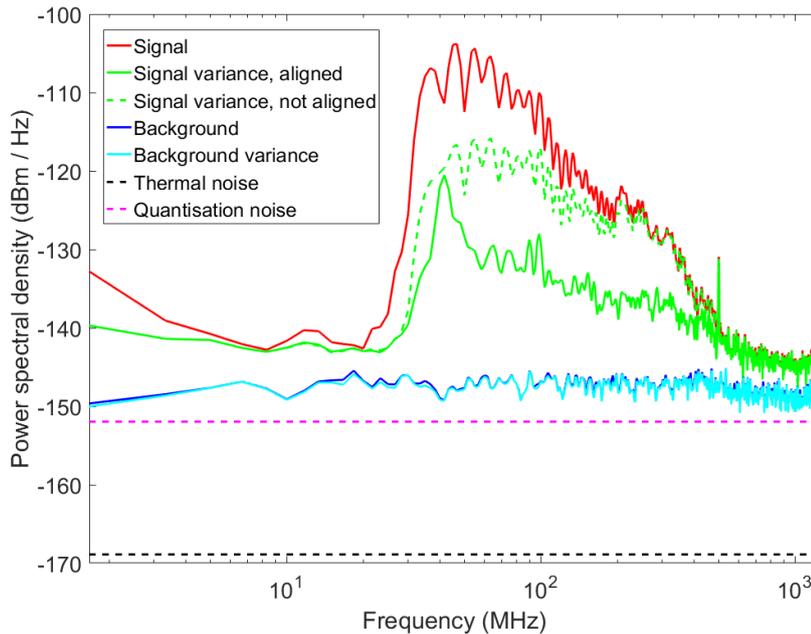


Figure 3.25: Power spectral density of the pulses shown in Fig. 3.18 in dBm/Hz vs. frequency in MHz. This is the same plot as Fig. 3.24, with the exception that the solid green curve shows the power spectral density of the random component of the signal after time alignment and the dashed green curve shows its value before time alignment.

of all the dark-current pulses and the mean shifted by the RMS jitter of 200 ps. This spectrum is shown plotted with the other spectra in Fig. 3.26.

It can be seen that the observed pulse-to-pulse variance can be attributed to the residual misalignment over most of its bandwidth. However, over the range 30 - 60 MHz, the variance PSD still exceeds the uncertainty. While this could be a candidate for the fluctuations being searched for, it is likely to be caused by another source of error which will be discussed later.

RF Modulation Effects

The analysis done thus far has assumed that each RF pulse is identical, and thus that any fluctuation in the measured dark current that exceeded the measurement noise must have been caused by a change in the properties of the emitter itself. However, the RF input has inherent uncertainty in its level and pulse shape. Such noise is expected to result in a fluctuation of the dark current, even for a perfectly stable field

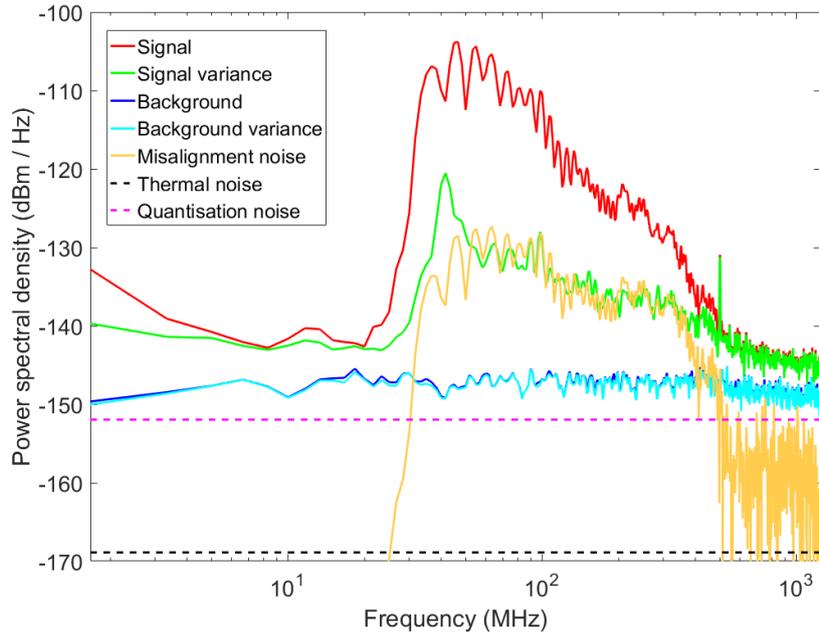


Figure 3.26: Power spectral density of the pulses shown in Fig. 3.18 in dBm/Hz vs. frequency in MHz. This is the same plot as Fig. 3.24, with the addition of the yellow curve representing the estimated noise spectrum resulting from the residual misalignment in time of about 200 ps RMS.

emitter. Thus, for any dark-current fluctuation measured to be considered significant, it would have to exceed the noise contribution from the incident RF.

An attempt was made to obtain an estimate of the contribution of the RF fluctuations to the measured dark current, though the results were ultimately not used due to shortcomings of the overly simple mathematical modelling. This estimate involved modelling the structure as a point field emitter with a Fowler-Nordheim current-field dependence. A power scan was used to empirically measure the dependence of the current captured by the Faraday cup on the incident power and thus surface electric field. Both current and the derived surface electric field were averaged over the duration of the pulse. As the properties of field emission in accelerating structures are known to vary over time, care was taken to make sure that this curve was measured immediately before the fluctuation measurement.

The peak surface electric field as a function of time $E_i(t)$ was calculated for each pulse i , based on the output of the IQ demodulator and its calibration curve. $E_i(t)$ was then

used as an argument of the empirically measured current-field dependence to obtain a predicted value for the collected current $I_i(t)$. A comparison of the dark current pulses reconstructed in this manner and the actual measured pulses is shown in Figure 3.27. The overall shape and level of the signal are similar in the two cases, but there are clearly features in the measured signal that are not accounted for by this simple model.

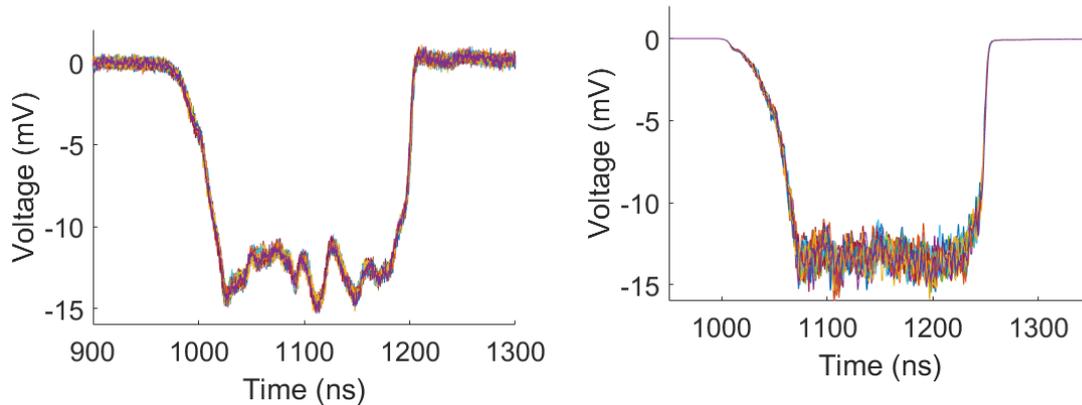


Figure 3.27: Voltage on the upstream Faraday cup in mV vs. time in ns. Different coloured curves represent different pulses at the same RF power level. Left: original signals as measured by the oscilloscope. Right: reconstructed signals based on the incident RF signal assuming a single field emitter and no electron transport effects.

An obvious limitation of this model is that it does not take into account the transport of electrons over the roughly 10 cm distance from the emission sites to the Faraday cup. The variation in time of flight may thus have a significant effect on the time structure of the current collected by the Faraday cup. This is because the electrons emitted by the upstream Faraday cup have not been captured by the accelerating structure, and thus may have a large spread of energies and thus velocities, leading to a spread in time of flight. The oscillations in the measured dark current signal in Fig. 3.27, largest just after the 1100 ns mark, appear to be very consistent from pulse to pulse, which makes it unlikely that they are the result of random dislocation motion.

Some more insight can be gained from performing a similar analysis on the reconstructed dark-current signal in the same way as the measured signals. These are shown in Fig. 3.28 for measurements made without any high-pass filtering before the oscilloscope. For frequencies above 10 MHz, the power spectral density of the actual signal

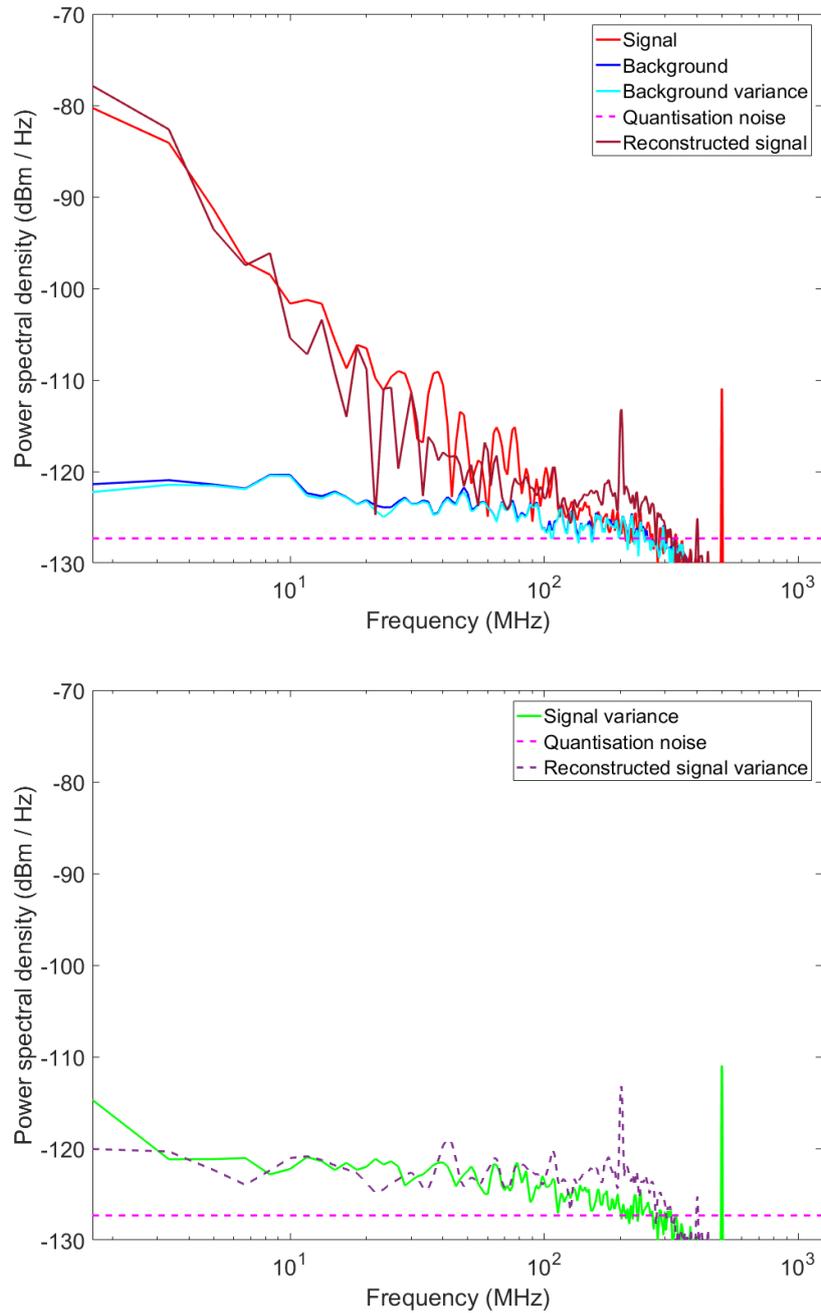


Figure 3.28: Power spectral density of dark current signals measured without a high pass filter in dBm/Hz vs. frequency in MHz. The plots are shown on two separate sets of axes for clarity. Red solid curve: total power spectral density of the measured signal. Purple solid curve: total power spectral density of the reconstructed signal. Green solid curve: power spectral density of the random component of the measured signal. Purple dashed curve: power spectral density of the random component of the reconstructed signal. Dark blue: total power spectral density of the background noise. Light blue: power spectral density of the random component of the background noise. Magenta dashed curve: quantisation noise.

exceeded that of the reconstructed one, which is consistent with the presence of oscillations in the actual signal that were not present in the reconstructed signal in Fig. 3.27.

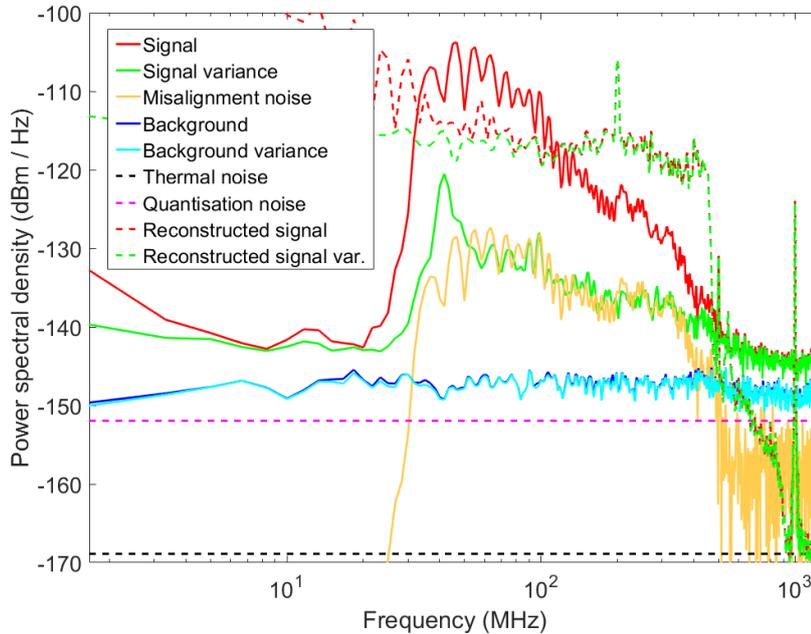


Figure 3.29: Power spectral density of dark current signals measured with a 50 MHz high-pass filter in dBm/Hz vs. frequency in MHz. Red solid curve: total power spectral density of the measured signal. Red dashed curve: total power spectral density of the reconstructed signal. Green solid curve: power spectral density of the random component of the measured signal. Green dashed curve: power spectral density of the random component of the reconstructed signal. Dark blue: total power spectral density of the background noise. Light blue: power spectral density of the random component of the background noise. Yellow: noise spectrum resulting from the residual misalignment in time. Black: thermal noise excluding the input noise of the oscilloscope itself. Magenta: quantisation noise.

More detail can be seen in Fig. 3.29, in which a similar comparison is made for measurements done with a 50 MHz high-pass filter on the Faraday cup channel. Over the bandwidth of this signal, about 20 MHz to 400 MHz, the variance of the reconstructed signal was well above that of the measured signal. Based on this model, the measured variance cannot be considered statistically significant. However, for the reasons discussed earlier, this model may not be reliable. It is also possible that the large variance in the reconstructed dark current may be a result of uncertainty in the measurement of the RF signal, rather than a real pulse-to-pulse variance, propagated to

the dark-current signal. This prompted a closer study of the RF signal in addition to the dark-current signal.

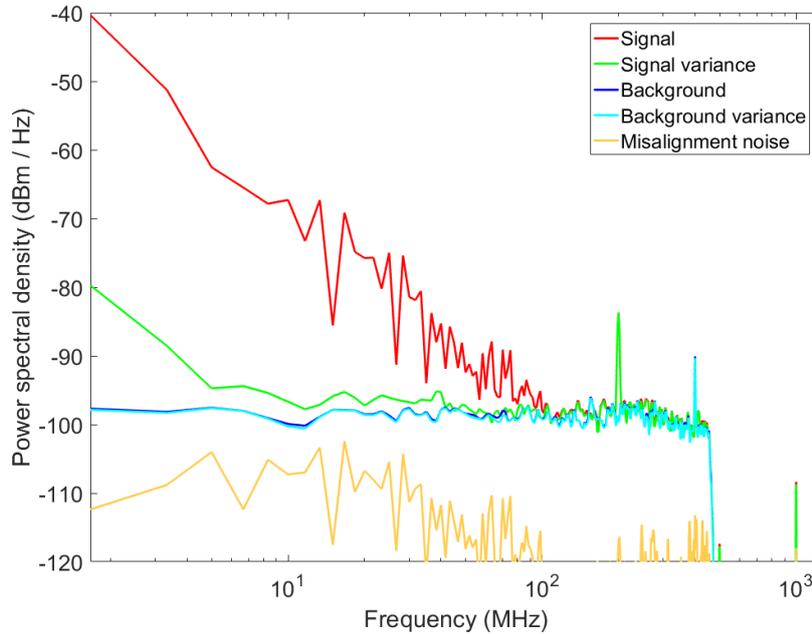


Figure 3.30: Power spectral density of the incident RF signal in dBm/Hz vs. frequency in MHz. Red: total power spectral density of the signal. Green: power spectral density of the random component of the signal. Dark blue: total power spectral density of the background noise. Light blue: power spectral density of the random component of the background noise. Yellow: noise spectrum resulting from the residual misalignment in time.

The result of performing the same analysis done on the dark current signal on the amplitude of the demodulated RF signal is shown in Fig. 3.30. It shows a pulse-to-pulse variance slightly above the background noise between 0 and 40 MHz, meaning that the variance of the RF pulses was measurable and that this should have resulted in a measurable dark-current variance, if the point field-emitter model were accurate. Due to the large difference between the observed and predicted behaviour, the point field-emitter model proposed was deemed inadequate, and no conclusions about the significance of pulse-to-pulse variance of the dark-current signal (for example as in Fig. 3.29) based on it were drawn. Without further knowledge of the behaviour of electrons travelling between the emission site and the Faraday cup, it was not thought to be feasible to obtain any more useful information from this study. A Particle-in-

Cell (PIC) simulation is one way by which electron transport in this system could be studied.

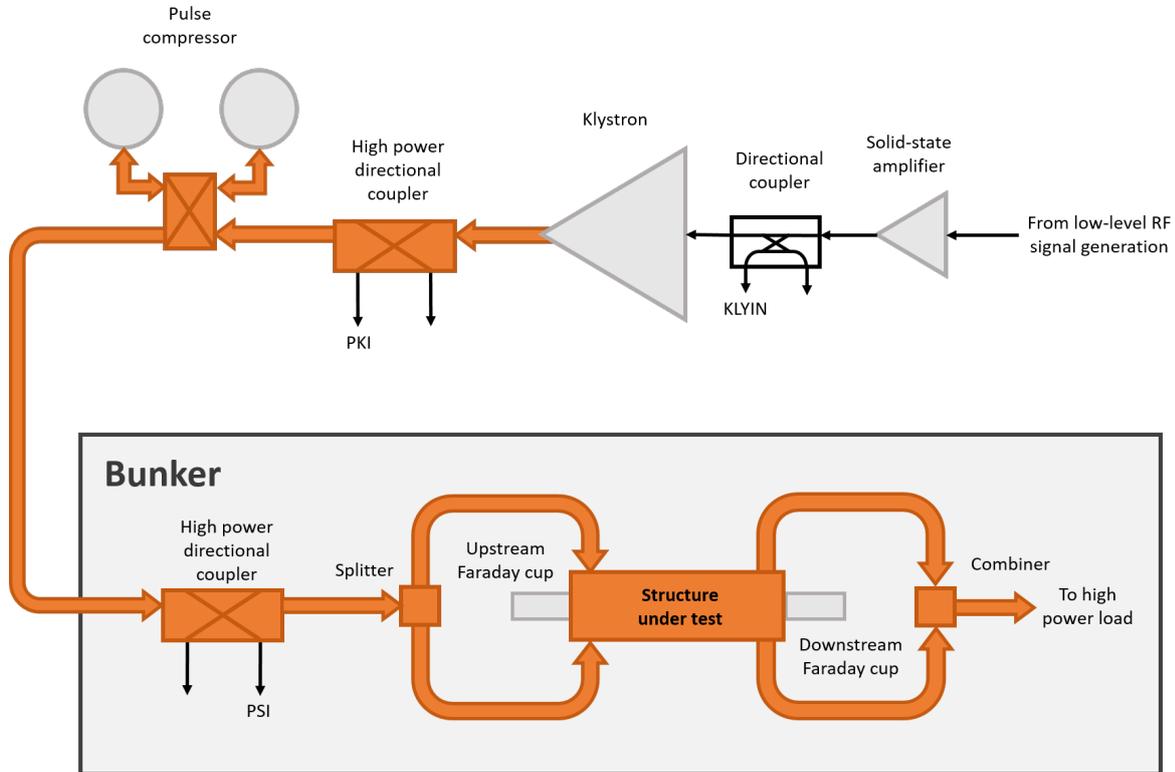


Figure 3.31: Path of the high-power RF pulse incident on the structure under test, and relevant directional coupler signals: ‘KLYIN’, the signal incident to the klystron input; ‘PKI’, the signal output from the klystron, ‘PSI’, the signal incident to the structure under test.

RF signals further upstream of the structure were also investigated to determine if a source of the 40 MHz peak in the variance of the measured dark-current signal could be identified. Fig. 3.31 shows the path of the high-power RF pulse from the solid-state klystron driver amplifier to the structure itself. Due to the use of a SLED-type pulse compressor [111] to multiply the peak power, an RF pulse of about 1.1 μs in duration was generated by the low-level RF (LLRF) electronics, which was then amplified by the solid-state amplifier and the klystron. The RF phase of the pulse was ramped over the final 200 ns of this pulse in order to cause the pulse compressor to release the energy it stored over the first 900 ns and produce a compressed pulse of high peak power lasting up to 200 ns.

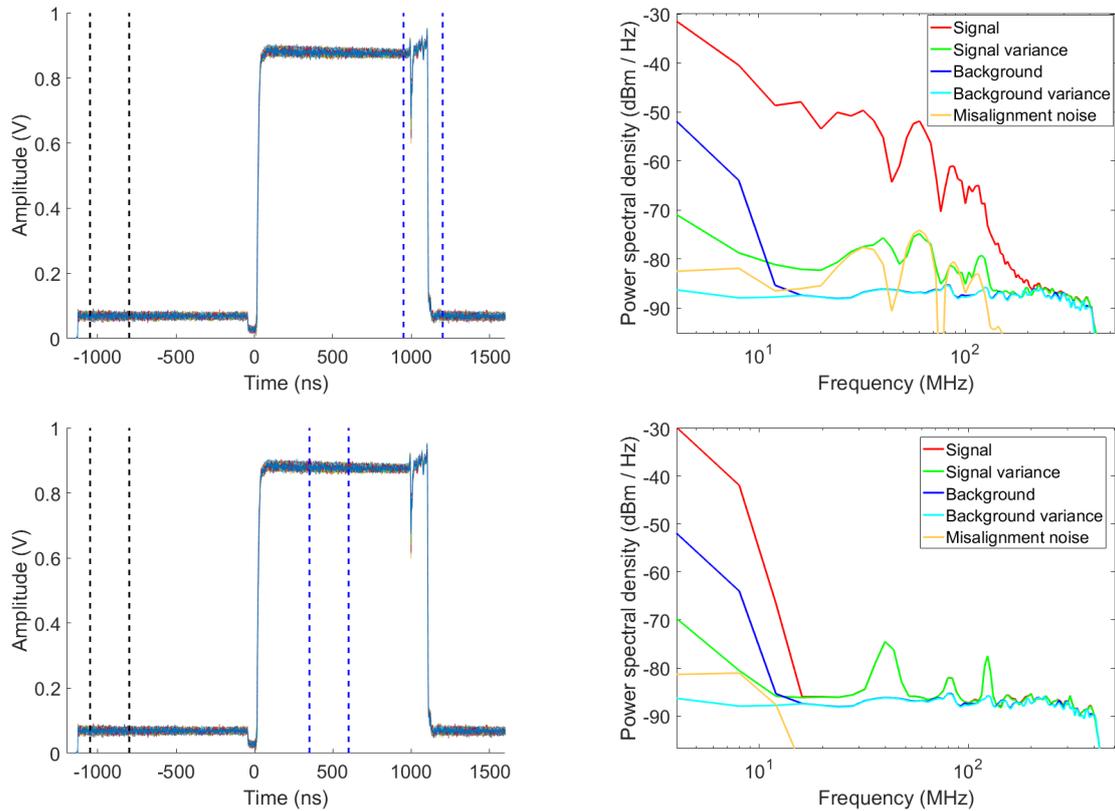


Figure 3.32: Analysis of the RF signal input to the klystron from the solid-state driver amplifier, comparing the spectral content of two different sections of the pulse. The two plots on the left show the RF signal amplitude, as measured by the IQ demodulator, in volts vs. time in ns. The time window for the DFT, show as blue dashed lines, encompasses the compressed pulse in the top left plot and the flat part of the pulse in the bottom left plot. The two plots on the right show the spectra of the signal in the respective time windows in dBm/Hz vs. frequency in MHz. In each of the spectrum plots, the red curve represents the total power spectral density of the signal, the green curve represents the power spectral density of the random component of the signal, the dark blue curve represents the total power spectral density of the background noise, the light blue curve represents the power spectral density of the random component of the background noise, and the yellow curve represents the noise spectrum resulting from the residual misalignment in time.

Fig. 3.32 shows the 1.1 μs long pulse produced by the solid-state amplifier delivered to the input of the klystron. The signal was taken from the channel labelled ‘KLYIN’ in Fig. 3.31. The limited bandwidth of the system results in variations in amplitude during the period of rapid phase change in the final 100 ns of the pulse. Spectral analysis was performed on two separate time windows: one encompassing the phase ramp used to produce the compressed pulse, and another sampling part of the flat section of the pulse. The first window has much more power at high frequencies due to the presence of a large step. The presence of large, fast edges in the compressed pulse makes it susceptible to any residual time misalignment that may be present in the recorded data. This places a limitation on the sensitivity of the variance measurement, as is evident from the elevated misalignment noise spectrum. Despite the reduced sensitivity, measurable variance exceeding the noise floor was still visible at low frequencies, indicating some pulse-to-pulse jitter in the average power level delivered to the klystron.

The second window, on the other hand, only had a significant power spectral density below 20 MHz, and less time-misalignment noise. The lower noise floor revealed three clearly visible peaks at about 40, 80 and 120 MHz, that are all above the background noise. These frequencies are all integer multiples of 40 MHz, and thus might be harmonically related to each other, suggesting that they might be the consequence of an unwanted signal coupled into the LLRF electronics responsible for generating the RF pulse. If the signal was not phase-locked with the RF pulses, it could have had a different phase in each pulse recorded and thus appeared as a pulse-to-pulse variation.

The fact that an unwanted 40 MHz signal was present in the signal incident to the klystron makes it likely that it was also present in the compressed pulse incident to the accelerating structure. There is no 40 MHz peak clearly visible in the variance spectrum of the RF incident to the structure, as shown in Fig. 3.30, though it may have been obscured by the harmonic content of other features in the signal, as with the two time windows studied in Fig. 3.32. The fact that the klystron was saturated may have reduced the amplitude of the accidental 40 MHz modulation relative to the 12 GHz carrier frequency. When driven into saturation, the differential gain (change

in output power per change in input power) is much lower than its maximum gain, meaning that any small modulation of the RF amplitude would be reduced, since the klystron cannot be considered to be a linear amplifier in this regime.

The presence of this small 40 MHz component in the surface electric field may be the reason for the 40 MHz peak in the variance spectrum of the upstream dark current. The point field-emitter model discussed in Sec. 3.2.3 also suggests that the amplitude that this component would need to have in order to explain the peak in the dark current signal is lower than the smallest amplitude measurable with the experimental setup used. Thus, it is plausible that the observed 40 MHz peak in the variance spectrum was caused by this unwanted signal in the incident RF.

Having considered the most significant sources of error, there appears to be no conclusive evidence that there was any pulse-to-pulse variation in the dark current captured by the upstream Faraday cup that could have been caused by dislocation motion under high fields. The ultimate limit to the sensitivity of the dark-current measurements performed turned out to be time misalignment from pulse to pulse, as detailed earlier in this section. One way to reduce its effect would be to avoid fast changes in the RF waveform's amplitude and phase. Fig. 3.32 demonstrates the improvement in sensitivity that this would cause. The rapid modulation is a necessary consequence of the use of a SLED-type pulse compressor in the test stand, needed to achieve the requisite peak RF power. An alternative, however, exists in the form of the Nextef test facility [121] at the National Laboratory for High Energy Physics (KEK) in Japan, in which the outputs of two 50 MW X-band klystrons are used to produce sufficient RF power without the need for a pulse compressor. The low-level RF signal driving the klystron exhibited some pulse-to-pulse jitter in its amplitude, in addition to containing unwanted components at 40, 80, and 120 MHz, which also made the desired fluctuation signal more difficult to detect. A review of the low-level RF electronics could reduce these background signals as well.

Variation of Signals with Surface Electric Field

The dislocation model predicts changes in the behaviour of the fluctuations when conditions approach breakdown, such as a change in the shape of the probability of finding the system in a given state n . Given that, a search was made for evidence of any anomalous behaviour indicative of such changes as the incident power was varied. It should be noted that although the dislocation model makes predictions of the mean time taken to move from a given state to an adjacent state, which occurs below the surface, it does not make any predictions of how the dark current signal should change in response to this. Such a prediction would require knowledge of the relationship between subsurface movement and surface changes.

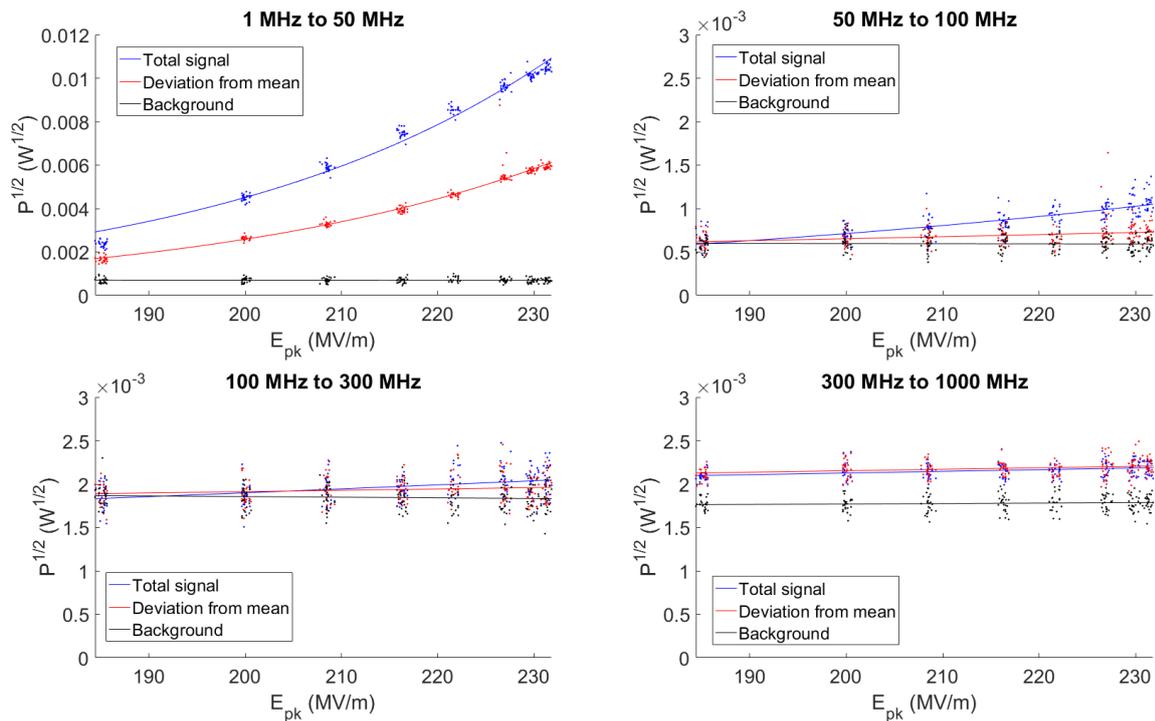


Figure 3.33: Square root of the dark current signal power in $W^{1/2}$ vs. peak surface electric field in MV/m. In each plot, the blue points represent the total signal power within the specified frequency band for each measured pulse, the red points represent the power of the deviation of each pulse from the mean of all the pulses with that nominal power level, and the black points represent the background noise. Fits were performed by minimising χ^2 summed over all the data points. Top left: between 1 and 50 MHz. top right: between 50 and 100 MHz. Bottom left: between 100 and 300 MHz. Bottom right: between 300 and 1000 MHz. Due to high-frequency switching noise from the klystron modulator, the signal in this frequency range is above the background even at low RF power levels.

In the absence of a more precise prediction, the dependence of the deviation of the signal from the mean as a function of applied electric field was studied. In this case, ‘mean’ refers to the average of all the pulses at that particular nominal RF power. To potentially gain more insight, the signal was split up into several frequency bands that were analysed separately. An example of this can be seen in Fig. 3.33. The data presented in this figure was taken without a high-pass filter, and is the same data as that in Fig. 3.24. There is a clearly visible signal in the 1 MHz to 50 MHz frequency band, which qualitatively resembles a Fowler-Nordheim curve, with a deviation from the mean that also appears to increase with the electric field in a similar manner. The signal power in the higher frequency bands is much lower, consistent with the spectra shown in Fig. 3.24.

The data was also plotted on Fowler-Nordheim axes, as in Fig. 3.34 for easier identification of field emission phenomena. A linear fit was made to each of the plots in all the frequency bands to obtain the apparent field enhancement factor β . As expected, the total signal and the deviation from the mean in the 1 MHz to 50 MHz frequency band produced good fits with positive β values, indicating consistency with Fowler-Nordheim field emission, whereas the background noise and the signals in the other frequency bands did not. The fitted β for the deviation signal was slightly higher than that of the total signal.

A similar plot for data taken with the 50 MHz high pass filter is shown in Fig. 3.35. Due to the increased sensitivity of this configuration, dark current signals are now visible in the 50 MHz to 100 MHz and 100 MHz to 300 MHz frequency bands, as evidenced by positive β values and good fits to linear functions. As before, the fitted β for the deviation signal was slightly higher than that of the total signal in each of the frequency bands.

The interpretation of these results can be aided by a simple simulation of the behaviour of a hypothetical signal with a current variance that has a known dependence on the electric field, and what this would look like on the Fowler-Nordheim plots shown above.

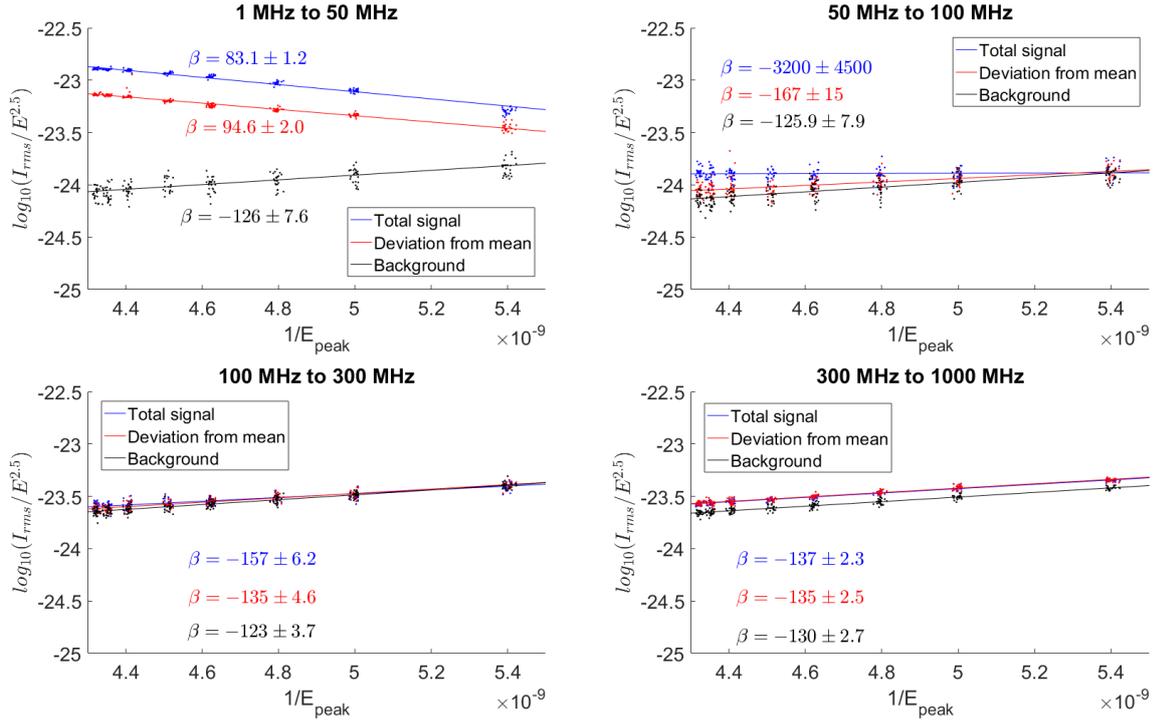


Figure 3.34: A measurement of the RMS field emitted current on Fowler-Nordheim axes, i.e. $\log_{10}(I_{rms}/E^{2.5})$ vs. the inverse of surface electric field in m/V. In each plot, the blue points represent the total signal power within the specified frequency band for each measured pulse, the red points represent the power of the deviation of each pulse from the mean of all the pulses with that nominal power level, and the black points represent the background noise. Fits of the field-enhancement factor β are shown. The fits were performed by maximising the likelihood summed over all the data points, assuming Gaussian current noise. Top left: between 1 and 50 MHz. top right: between 50 and 100 MHz. Bottom left: between 100 and 300 MHz. Bottom right: between 300 and 1000 MHz.

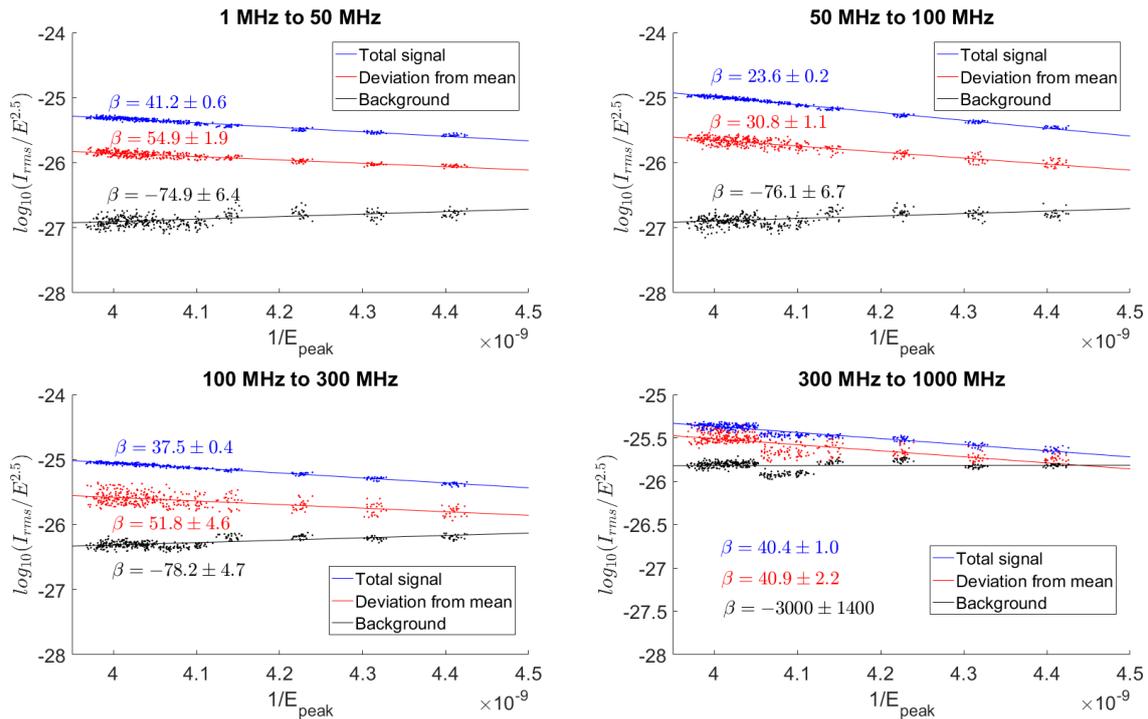


Figure 3.35: A measurement of the RMS field-emitted current, measured with a 50 MHz high-pass filter, on Fowler-Nordheim axes, i.e. $\log_{10}(I_{rms}/E^{2.5})$ vs. the inverse of surface electric field in m/V. In each plot, the blue points represent the total signal power within the specified frequency band for each measured pulse, the red points represent the power of the deviation of each pulse from the mean of all the pulses with that nominal power level, and the black points represent the background noise. Fits of the field-enhancement factor β are shown. The fits were performed by maximising the likelihood summed over all the data points, assuming Gaussian current noise. Top left: between 1 and 50 MHz. top right: between 50 and 100 MHz. Bottom left: between 100 and 300 MHz. Bottom right: between 300 and 1000 MHz.

Any consistency with a very strong dependence on electric field or of a sharp transition close to breakdown conditions would suggest the existence of the proposed dark-current fluctuations.

In this simple simulation, we define f to be the Fowler-Nordheim equation relating the emitted current I_0 to the enhanced surface field $\beta_0 E$:

$$I_0 = f(\beta_0 E), \quad (3.17)$$

where E is the surface electric field and β_0 is the field-enhancement factor. The deviation in current δI represents the pulse-to-pulse variation in dark current that has been studied in previous sections. In this model, δI is assumed to arise as a result of a deviation in surface field δE . This may be caused by a true change in the surface field, such as would be the case with a noisy power source, or a change in the local field enhancement factor by $\delta\beta = \beta/E \cdot \delta E$, which would produce the same change in current. A change in local field enhancement could be caused by a change in geometry, which could arise as a result of dislocation motion.

$$I_0 + \delta I = f(\beta_0(E + \delta E)) = f(\beta_0 E + \beta_0 \delta E), \delta E \ll E \quad (3.18)$$

The deviation in field was then assumed to have a power-law dependence on the nominal electric field E :

$$\delta E \propto E^\alpha, \quad (3.19)$$

where α is some constant. Constant, additive noise would result in a deviation that does not depend on the nominal field, yielding $\alpha = 0$, and a fixed uncertainty in the gain of an amplifying stage would result in a deviation that depends linearly on the nominal field, i.e. $\alpha = 1$. A constant deviation in β caused by a fluctuation in the geometry of the emitter independent of E , would also correspond to $\alpha = 1$. On the

other hand, the breakdown rate is proportional to E^{30} , and any process that plays a role in breakdown should have a very sudden onset. It would thus be reasonable that the occurrence of fluctuations in dark current would be a very strong function of the surface field and have a very large value of α .

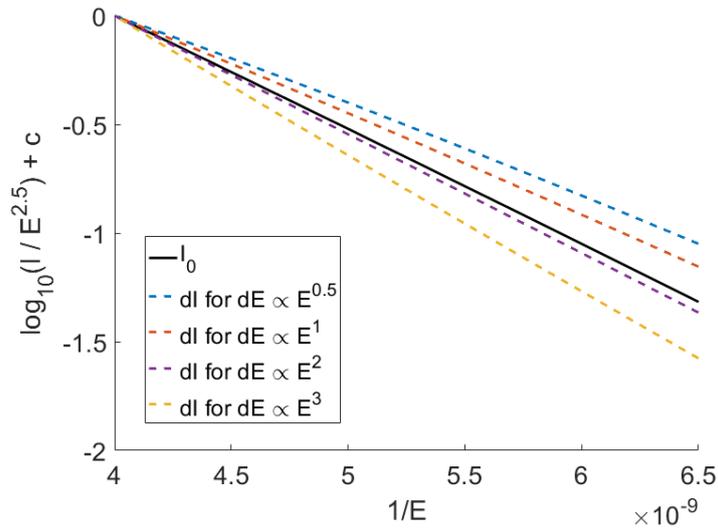


Figure 3.36: Field emitted current on Fowler-Nordheim axes, i.e. $\log_{10}(I/E^{2.5})$ vs. the inverse of surface electric field in m/V. The black solid line represents the mean field emitted current I_0 , whereas the coloured dashed lines represent the deviations in current δI arising from a deviation in field δE , which is proportional to E^α . Different colours represent different values of α , as indicated on the legend. A field-enhancement factor of 50 was assumed. A constant was added to each of the plots to ensure that they all pass through zero for easier comparison of their slopes.

Fig. 3.36 shows the result of plotting δI along with I_0 for several different values of α on Fowler-Nordheim axes, where it can be seen that the slope of the δI plot differs slightly from that of I_0 , depending on the value of α chosen.

The slope of each of the plots in Fig. 3.36 can be used to obtain a value for the apparent field enhancement β_{fit} . This was done over a range of values of α for different field enhancement factors β_0 . The relationship between the apparent β_{fit} obtained from the current deviation, and the original β_0 , is shown in Fig. 3.37. This result suggests that β_{fit} obtained from the deviation should be significantly lower than β_0 obtained from the mean current if it is caused by a very strong function of E , and higher than β_0 if the deviation depends weakly on E .

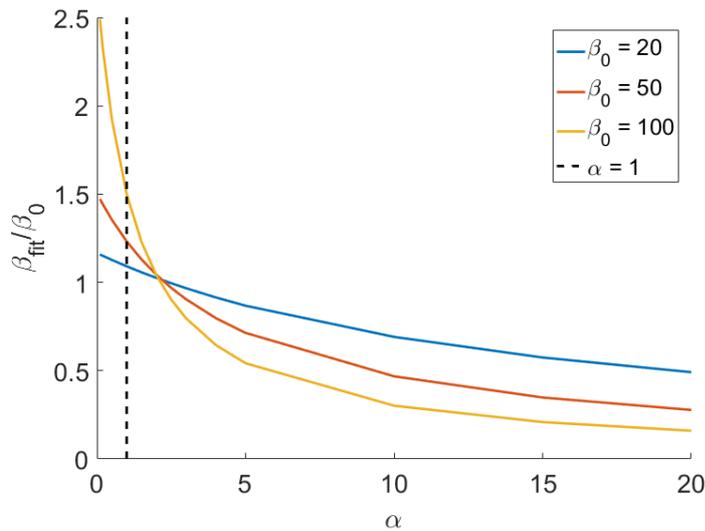


Figure 3.37: The ratio of the apparent β_{fit} of current deviations to the ‘true’ field enhancement β_0 vs. the exponent α . For each of the values of β_0 shown, a deviation with a sudden onset, and therefore a large value of α , implies $\beta_{fit} < \beta_0$.

In the experimental results shown in Figs. 3.34 and 3.35, the fitted β obtained from the deviation was higher than that obtained from the total signal by a factor between 1.1 and 1.4. By comparison with Fig. 3.37, it can be deduced that the deviations all correspond to a small α of around 1, and are not consistent with the existence of a sharp transition as predicted by the dislocation hypothesis. The measurements performed here include field levels at which a high breakdown rate was expected, meaning that the current deviation was as large as practically possible with this setup. The peak surface field reached in this data set was 247 MV/m, the highest in this series of measurements. The measurement was continued until a breakdown occurred, meaning that if any measurable precursor signals were present, they should have been observed. Despite this, no evidence supporting the dislocation hypothesis was found in studies of the dependence of the dark current signal on the incident power.

3.2.4 Conclusion

The incident RF and emitted dark current were precisely measured with maximum bandwidth and minimum noise. Measures were taken to improve the sensitivity of the measurement and analysis as much as possible. No evidence for fluctuations in

dark current caused by dislocation motion in RF accelerating structures under test in the XBox test stands have been found. This could mean that either the hypothesis is either incorrect or incomplete, or that the fluctuations were too small to be measured practically.

The measurement was hindered somewhat by the complexity of the experimental setup. The uncertainty in the surface field was made worse by the need for pulse compression to reach sufficiently high electric fields, and well as the pulsed nature of the applied field. These two constraints meant that there was large amount of background electromagnetic noise in the Faraday cup signal and the RF wave incident on the accelerating structure under test, which reduced the sensitivity of the measurement. The dynamics of electron transport between the emission site and the Faraday cups was also a notable problem. It was shown that the current reaching the Faraday cup had a time structure that did not match the expected field-emitted current given the applied surface electric field. The most likely explanation for this discrepancy is that the fraction of the emitted current that reached the Faraday cup, as well as the transit time, varied with time. These effects are also likely to have made it more difficult to measure dark-current fluctuations.

3.3 Measurement in the DC Spark System

The difficulties encountered with trying to measure dark-current fluctuations in prototype RF structures, such as the dynamics of the emitted electrons which are difficult to model with the limited amount of initial information, prompted a similar investigation on the LES. From the perspective of dislocation dynamics, the conditions in the LES are expected to be very similar to those in the RF structures, as both are subject to very high surface electric fields (about 80 MV/m for the LES, and about 220 MV/m for the RF structures), and both use the same copper alloy put through the same heat treatment. Dark current is measured in the LES as well. The LES is a much simpler system than the XBox facility and is much more amenable to modifications

in the setup. It also offers the possibility of applying an unpulsed DC voltage to the electrodes. The pulsed nature of the applied RF field in the structures tested in the XBoxes produced a significant background signal which made dark-current fluctuations more difficult to detect (see Sec. 3.2). Thus, not having a pulsed voltage is one factor which contributed to the improved sensitivity of the measurement at the LES.

Vacuum-Chamber Simulations

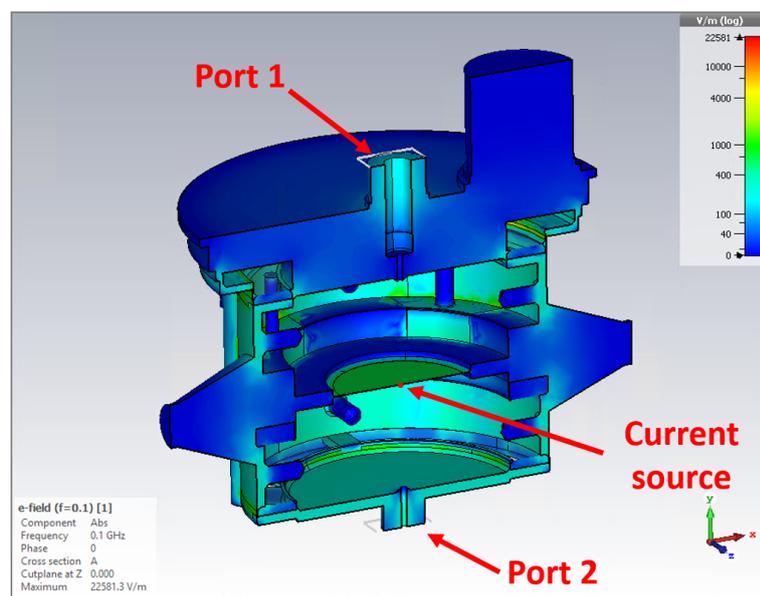


Figure 3.38: Surface electric field in the LES chamber in V/m vs. position. A 100 MHz sinusoidal current was assumed to flow across the centre of the gap. The location of the current source and the top and bottom ports of the chamber are labelled.

An electromagnetic simulation was performed on the geometry of the vacuum chamber of the LES to understand how well a fluctuation signal originating from the inter-electrode gap could propagate through the chamber to the vacuum feedthrough where it could be read. The setup of the vacuum volume and the results are shown in Fig. 3.38, with the dark current being modelled as a current flowing across the centre of the gap. The top and bottom feedthroughs were modelled as coaxial ports. A $50\ \Omega$ characteristic impedance was assigned to the coaxial ports, since this was the characteristic impedance of the coaxial cables connecting the chamber with other components.

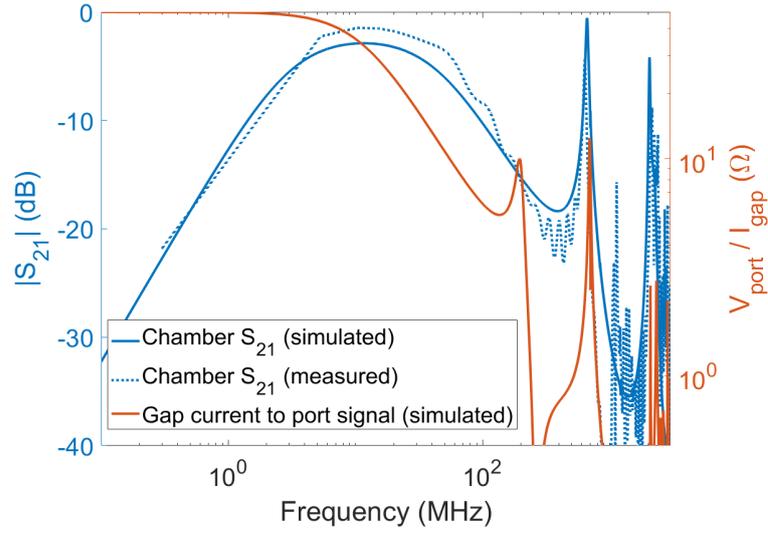


Figure 3.39: The signal propagation characteristics of the LES, with 40 mm diameter circular electrodes with a gap size of 40 μm . Red curve: simulated transfer impedance (the ratio of the voltage measured at the top feedthrough to the field-emitted current in the gap) in Ω vs frequency in MHz. Solid blue curve: simulated insertion loss $|S_{21}|$ between the top and bottom feedthroughs in dB vs. frequency in MHz. Dotted blue curve: measured $|S_{21}|$ vs. frequency in MHz.

The results of two simulations with this setup are shown in Fig. 3.39. In one of those simulations, the transfer impedance relating the current across the gap and the voltage on the top port is shown as a function of frequency. The bottom port was shorted, as it is normally done in the LES to ensure that the bottom electrode is at ground potential. A high transfer impedance is desired for maximum sensitivity to small currents. In the other simulation, the insertion loss S_{21} between the top and bottom ports was calculated. This simulation was performed to verify that the modelling of the geometry was correct, as the S_{21} of the LES chamber could be measured using a Vector Network Analyser (VNA), whereas there was no easy way of directly measuring the transfer impedance. The calculated and measured values of S_{21} , both plotted in Fig. 3.39, appear broadly consistent with each other. It is likely that the small discrepancies between the simulation result and measurement were caused by the vacuum feedthroughs, which were not modelled in the electromagnetic simulation.

The transfer impedance has a constant value of $50\ \Omega$ up to a cutoff frequency of about 3 MHz, above which it decreases as $\frac{1}{f}$. There are resonant peaks at around 200 and

700 MHz, at which the transfer impedance is significantly larger than at neighbouring frequencies, which makes these frequencies likely candidates for detecting very fast events. A possible reason for the existence of these peaks is the formation of an inductor-capacitor resonator from the capacitance of the gap and the inductance of the electrical vacuum feedthrough. Due to the existence of Johnson-Nyquist thermal noise in any electronic system [114], signals below a certain power spectral density cannot be detected. Thus, the transfer impedance is an important quantity to take note of. Using the values from this simulation, one can conclude that the thermal noise floor, assuming a $50\ \Omega$ system at room temperature, is equivalent to a current noise in the gap of $68\ \text{pA}/\sqrt{\text{Hz}}$ at 0.1 MHz, and $273\ \text{pA}/\sqrt{\text{Hz}}$ at 700 MHz. In comparison, the RF jitter in the accelerating-structure measurements detailed in Sec. 3.2 is equivalent to a current noise at the Faraday cup of about $2.5\ \text{nA}/\sqrt{\text{Hz}}$, which is more than an order of magnitude larger. As typical dark-current signals are on the order of $100\ \mu\text{A}$ in both the DC and RF setups, the LES offers the potential for much greater sensitivity than the RF experiments, especially for low-frequency signals.

High-Voltage Bias Tee

While the LES could achieve greater sensitivity than the measurements in the XBoxes, one disadvantage is that the dark current had to be measured on the input port, to which high voltages of up to 10 kV were applied to set up the required electric field in the inter-electrode gap. This made it unfeasible to directly connect any sensitive measurement equipment to the port, since this would be likely to damage it.

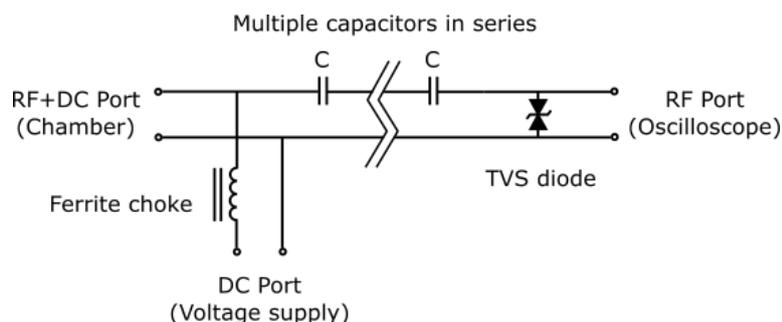


Figure 3.40: A conceptual schematic of the proposed bias tee. Several capacitors are connected in series to improve the voltage holding capability of the device.

A component which would resolve this incompatibility is a bias tee, for which a conceptual schematic is shown in Fig. 3.40. A bias tee acts as a high-pass filter between the RF+DC and RF ports, allowing AC signals to pass through but not a constant DC voltage. In this case, the RF+DC port would be connected to the vacuum chamber, the DC port connected to the high-voltage supply, and the RF port connected to the data acquisition.

Although bias tees rated up to about 100 V are frequently used in the telecommunications industry, no commercially available bias tees that could withstand a DC voltage of 10 kV could be found. Thus, a custom bias tee had to be designed for this application. The design objective of the bias tee was to ensure an impedance-matched signal path for the fluctuation signal, whilst blocking the DC high voltage. Since the frequency spectrum of the fluctuations was still not known at this point, the bandwidth was made as high as possible to maximise the chances of detecting the signal. In particular, a bandwidth encompassing the 660 MHz resonance of the vacuum chamber was considered desirable.

Real capacitors have a finite size and thus have some amount of inductance associated with the current path through the capacitor. This parasitic inductance is one mechanism which limits the bandwidth of a real device. The choice of capacitor was a compromise between voltage rating, package size (which determines the parasitic inductance), and capacitance. The capacitance C of a parallel-plate capacitor is given by [122]:

$$C = \frac{\epsilon A}{d}, \quad (3.20)$$

where ϵ is the permittivity of the dielectric, A is the electrode area, and d is the separation between the plates. Multi-layer ceramic capacitors (MLCCs) were used in this design, which consist of a stack of several parallel-plate capacitors connected in parallel in a single unit. A higher voltage rating requires a larger value of d , which in turn reduces the capacitance for a fixed electrode area. Since a capacitance that

Package type	Dimensions (mm)	Parasitic inductance (pH)
0603	1.5×0.8	360
0805	2.0×1.3	320
1206	3.0×1.5	540
1812	4.5×3.2	1130

Table 3.2: Properties of KEMET’s high-voltage surface-mount capacitors [123]. The values of parasitic inductance are given for a frequency of 1 GHz [122].

is too low will result in excessive attenuation of low frequencies, this will have to be compensated for by using physically larger capacitors with a larger area, and thus (in most cases) higher parasitic inductance. Table 3.2 shows a list of common MLCC sizes with their associated parasitic inductance.

After considering a number of commercially available options, it was decided that the best solution was to use several lower-voltage capacitors in series rather than a single capacitor rated for 10 kV. The final choice of capacitor was the KEMET C1812X471 [124], with a voltage rating of 2 kV, a capacitance of 470 pF, and C0G dielectric material - which ensures a stable capacitance with applied voltage [125]. To add safety margin, eight stages of capacitors in series were used, hypothetically giving a voltage rating of 16 kV. To increase the total capacitance, each stage consisted of two of these capacitors stacked vertically. To ensure an even split of the total voltage across each capacitor, a 10 M Ω resistor was also placed in parallel with each capacitor stage. Without these resistors, the voltage across each capacitor would be defined by the leakage resistance of the capacitor dielectric, which could vary significantly from unit to unit. As the leakage resistance is on the order of 1 G Ω or more, adding a 10 M Ω resistance in parallel made sure that the effect of the leakage resistance became negligible. The addition of the resistors produced a path between the high voltage node and ground with a resistance of 80 M Ω , which will result in a DC current of 125 μ A at the maximum voltage of 10 kV. This was not considered a hindrance as this was well within the maximum output current of the DC power supply used with the LES, and was not expected to make any difference to the fluctuation measurements.

It was found that by spacing the capacitors out with sections of transmission line, a periodic structure could be formed such that the frequency response of the bias tee was improved. Fig. 3.41 shows a more detailed schematic in which both the parallel resistors and the transmission line sections are shown; T_d refers to the group delay of each section of transmission line.

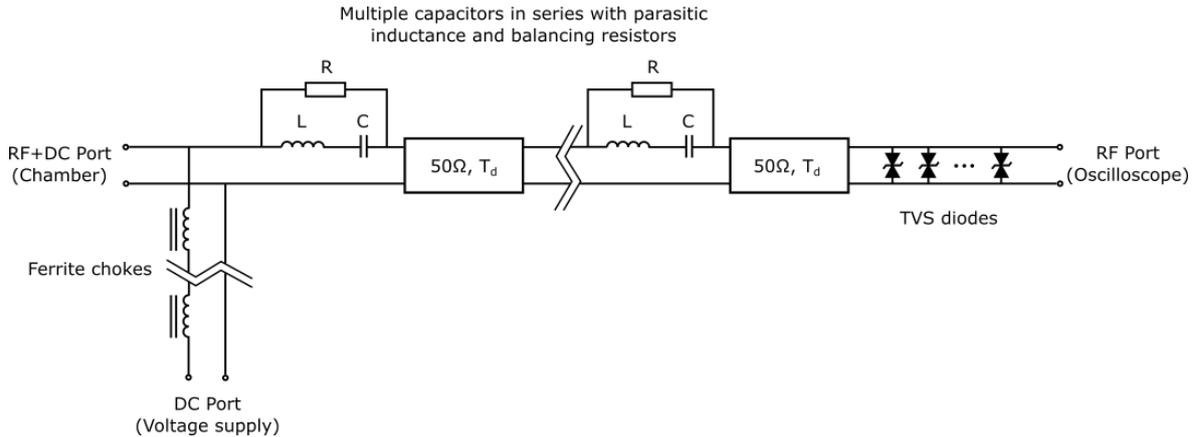


Figure 3.41: A schematic of a practical implementation of the bias-tee design including the parasitic inductance of the capacitors. Sections of $50\ \Omega$ microstripline of propagation time T_d are included between each capacitor. There is also a high-value resistor in parallel with each capacitor to ensure an even division of the total voltage between them.

Fig. 3.42 shows the results of an LTspice [126] circuit simulation of the bias-tee with different values of T_d . Here, one can see that a nonzero value of T_d brings about an improvement in the bandwidth of the device. An explanation for this improvement is as follows: the reflection from each capacitor stage, caused by the impedance mismatch due to the parasitic inductance, is shifted in phase by a different amount as each individual reflection experiences a different propagation delay. This causes the reflections to interfere destructively at the input port of the bias tee, reducing the magnitude of the sum. However, there exists a frequency at which the phase shift across one section of transmission line is equal to 180° . This results in constructive interference of the individual reflections, leading to complete reflection of the incident signal. This behaviour can be seen in Fig. 3.42 as a sharp high frequency cutoff. In the final design, the capacitors were separated by 3.9 mm long sections of microstrip, with $T_d = 24$ ps, which was found to give the greatest $|S_{21}|$ bandwidth.

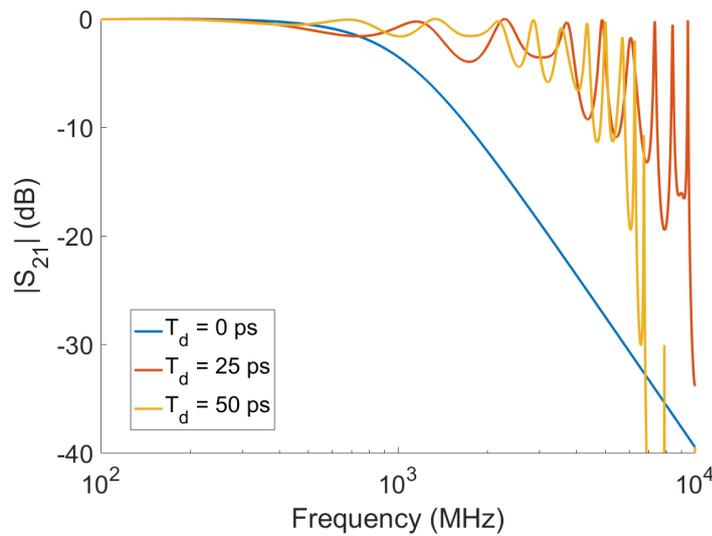
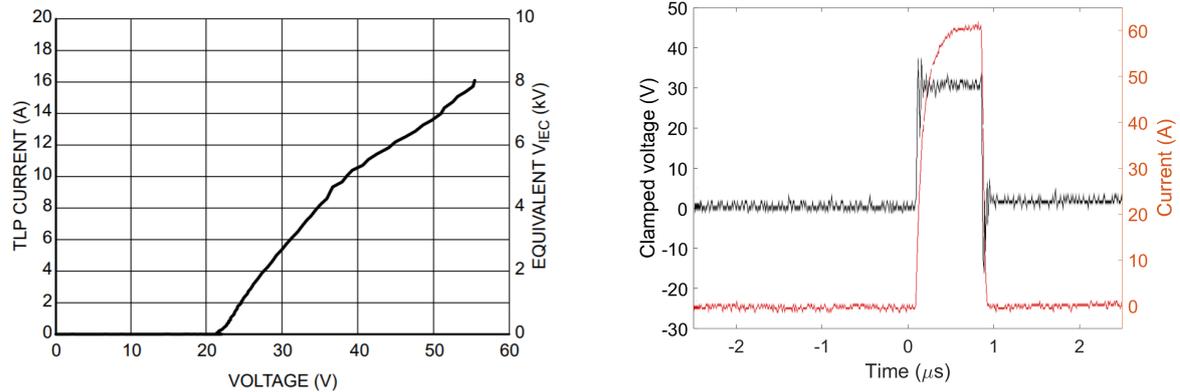


Figure 3.42: Simulated insertion loss $|S_{21}|$ of the bias tee vs. frequency in MHz. Blue: for $T_d = 0$ ps. Red: for $T_d = 25$ ps. Yellow: for $T_d = 50$ ps.

Another feature of the bias tee was the protection of sensitive electronics and human operators on the low-voltage side of the bias tee from high voltages in case of a failure of the high-voltage capacitors. A Transient Voltage Suppressor (TVS) diode [125] was introduced to clamp the voltage on the RF port in case of an arc across the capacitors. A TVS diode is normally operated in reverse bias, such that negligible current flows through it. Thus, it does not significantly affect the operation of the circuit under normal conditions. However, above a critical threshold voltage, the device undergoes avalanche breakdown and begins conducting large amounts of current. This loads the high-voltage source, causing the voltage across the diode to be limited to a safe value.

The type of TVS diode used had to be carefully chosen to avoid reducing the frequency response of the bias tee. This is due to the capacitance of the semiconductor junction [125], which can have a low enough impedance to ground to cause a significant impedance mismatch. The device chosen was the ON Semiconductor ESD7102 [127], which was specifically designed to protect high-speed data lines and thus has a very low capacitance of 0.3 pF. A single diode is rated for a maximum pulse energy of 5 mJ, while the LES chamber at a voltage of 10 kV may contain up to 30 mJ of stored energy. Thus, the bias tee was equipped with eight parallel pairs of diodes to ensure safe operation. The diodes in each pair were connected back-to-back in series, to halve the



(a) Current in A vs. voltage in V for one device, as provided by the manufacturer. Avalanche breakdown occurs at around 22 V, and the device begins conducting significant amounts of current.

(b) Voltage in V (black) and current in A (red) vs. time in μ s measured on a bank of eight pairs of devices connected in parallel when connected to a capacitor bank charged to 10 kV pulse. The voltage across the diodes was clamped to about 35 V.

Figure 3.43: Measurements of TVS diodes.

parasitic capacitance of each pair whilst only slightly increasing the clamping voltage. This gives a total capacitance of 1.2 pF to ground, which corresponds to a reflection coefficient of -14 dB at 1 GHz, still considered acceptable.

The correct functioning of this array of TVS diodes was verified by applying a 10 kV pulse from a solid-state Marx generator to this array, and measuring the current and voltage across it, as shown in the plot on the right in Fig. 3.43. The voltage was clamped to about 35 V, which, referring to the plot on the left in Fig. 3.43, corresponds to a current of about 8 A per branch, indicating a roughly equal split of the total current of 60 A. This was important to verify as diodes in parallel form a unstable system with a tendency for one diode to conduct most of the current [125]. The instability was most likely prevented by the existence of a significant series resistance within each diode, as evidenced by the linear current-voltage relationship above 20 V in the left plot in Fig. 3.43.

The assembled unit, along with a VNA measurement confirming the bandwidth of the device, is shown in Fig. 3.44. During each fluctuation measurement, the Marx generator normally used during conditioning was disconnected and replaced with a continuous DC supply, along with a range of other components as shown in Fig. 3.45. A second,

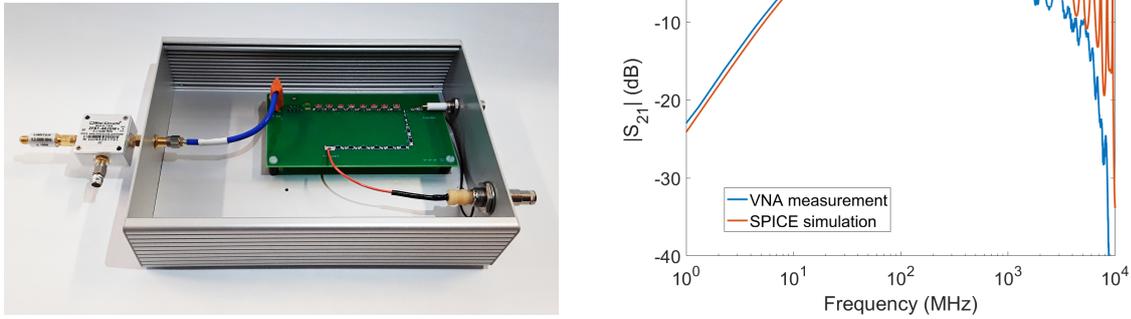


Figure 3.44: Left: photograph of the device in its enclosure. Right: Insertion loss $|S_{21}|$ in dB vs. frequency in MHz. The blue plot represents a VNA measurement while the red plot represents a simulation result with $T_d = 24$ ps. The measured $|S_{21}|$ deviated significantly from the simulation result at frequencies above about 1 GHz, which is believed to be due to dielectric loss in the FR-4 substrate, and impedance mismatches in the high-voltage connector which was not designed for high frequency signals.

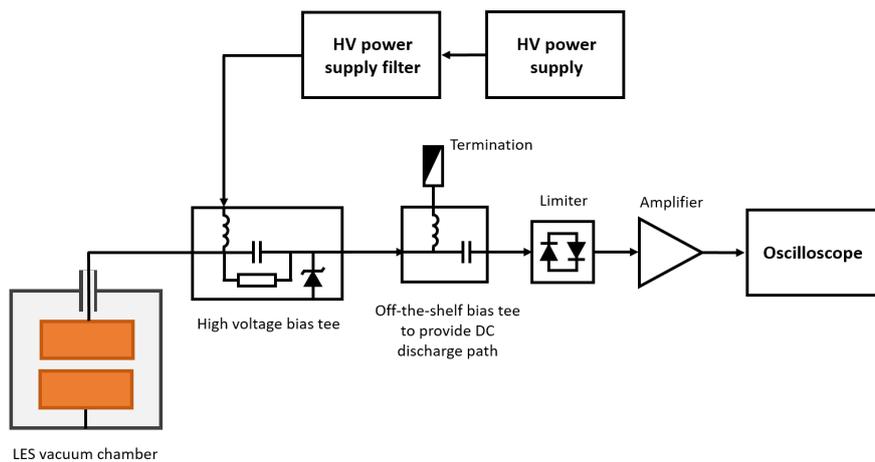


Figure 3.45: A block diagram of the experimental setup for the search for dark current fluctuations in the LES.

low voltage bias tee was connected to the RF output of the high-voltage bias tee to provide a DC discharge path to ground for the small current flowing through it. An RF limiter was also added, as the 35 V clamping voltage of the TVS diodes would have still been high enough to damage the amplifier and oscilloscope in the event of a failure. The amplifier used was a single Mini-Circuits ZKL-2R5+ [115], with a bandwidth of 10 MHz to 2.5 GHz, and a gain of about 30 dB. With this setup, any fluctuation signals originating from the vacuum chamber had a broadband impedance-matched path to the oscilloscope.

Measurements with Normal Electrodes

Initial attempts at measuring dark-current fluctuations were made with 40 mm diameter circular electrodes, with an inter-electrode gap of 40 μm . The electrodes were first conditioned up to the point at which they could reliably sustain a surface electric field of about 50 MV/m with a breakdown rate below 10^{-5} breakdowns per pulse, at which point they were run at a constant electric field for about 200 million pulses. The setup was then re-configured for fluctuation measurements as shown in Fig. 3.45. The applied voltage was then scanned while the signal measured by the oscilloscope was recorded. Care was taken to make sure no breakdowns occurred during the voltage scan.

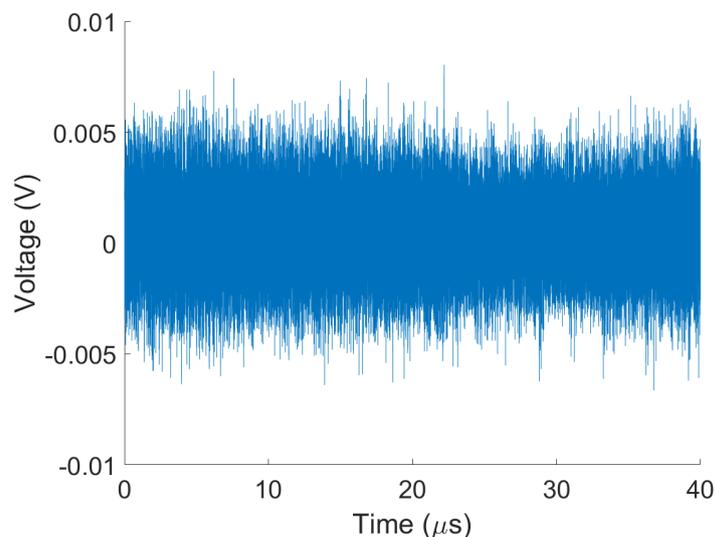


Figure 3.46: The raw signal in V vs. time in μs as measured by the oscilloscope with an applied voltage of 1800 V.

Fig. 3.46 shows an example of a raw signal waveform from this measurement at an applied voltage of 1800 V, corresponding to a surface electric field of 45 MV/m. No salient features are discernible by eye in this waveform, nor could any dependence on the applied voltage be seen. To aid with analysing this signal in more detail, a bank of matched filters [128] was created.

The impulse response of each filter was a complex exponential oscillation whose amplitude was modulated to create a short pulse. In discrete time, this can be written as a vector f_n :

$$f_n = w_n e^{2i\pi f_0 n}, \quad (3.21)$$

where w_n is a vector defining the amplitude and f_0 is the frequency of the oscillation. The filters were applied to the raw signal by calculating the cross-correlation between the raw signal s_n and the filter f_n , i.e.:

$$y_n = \sum_m s_n^* f_{m+n}, \quad (3.22)$$

where y_n is the vector representing the filtered signal, and s_n^* is the complex conjugate of s_n .

A matched filter gives a large cross-correlation value for events whose frequency content matches that of the filter, and a low value for random noise or signals with little power at those frequencies. A matched filter maximises the signal-to-noise ratio (SNR) for events which match the impulse response of the filter, i.e. $f_n = s_n$. The amplitude of the filter's impulse response should be zero outside a finite time span, so that the events detected by the filter may be resolved in time. The Chebyshev window, already discussed in Sec. 3.2.3, proved a suitable function for this task, and was thus used as w_n . A window span of 30 ns was found to work reasonably well. The impulse response of one of these filters is shown in Fig. 3.47.

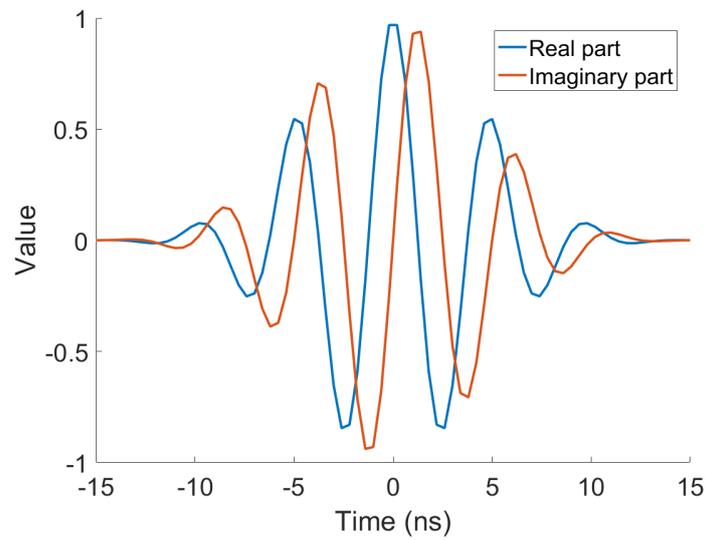


Figure 3.47: A matched filter, as defined in (3.22), with $f_0 = 200$ MHz and w_n being a 30 ns long Chebyshev window.

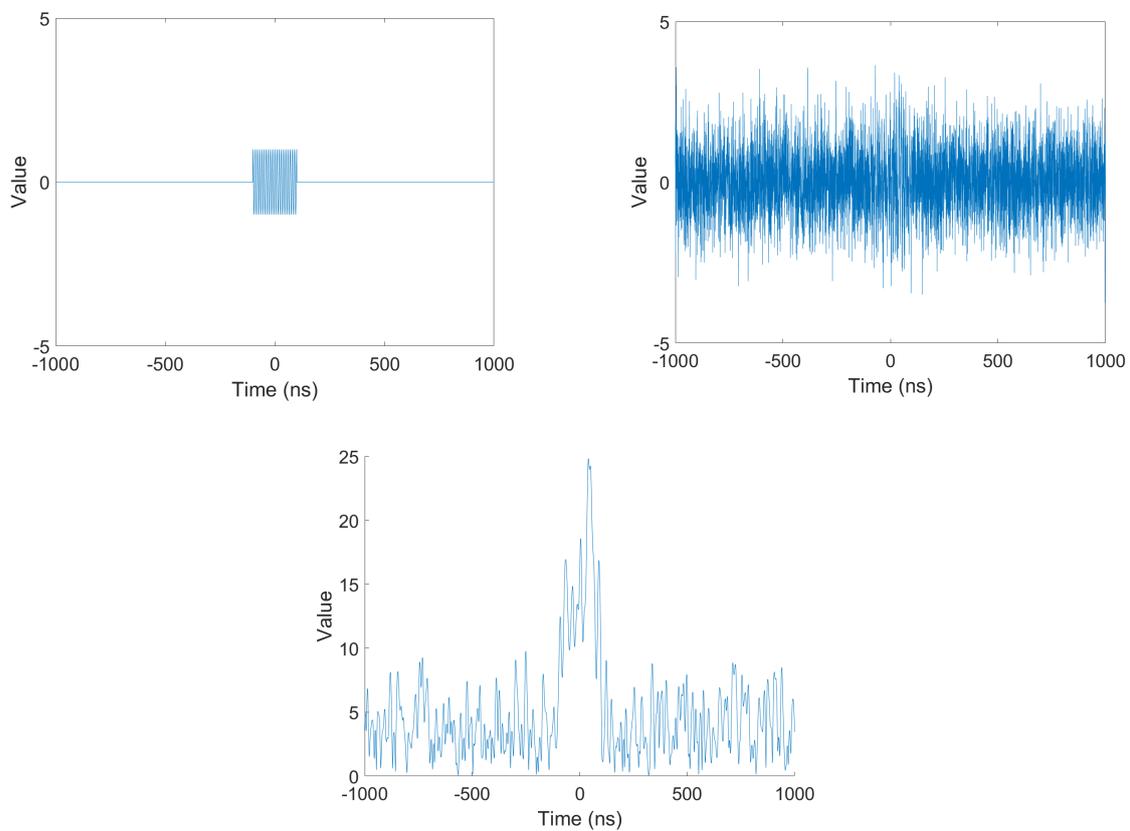


Figure 3.48: Extraction of a test signal corrupted by random noise using a matched filter. Top left: the test signal vs. time in ns. Top right: the test signal with added Gaussian noise vs. time in ns. Bottom: cross-correlation result vs. time in ns.

A test was performed on an example signal to demonstrate the effectiveness of this method, depicted in Fig. 3.48. The test signal was a 200 ns long 100 MHz sinusoidal pulse with unit amplitude. Gaussian noise of unit variance and zero mean was added to this signal to simulate a noisy environment, producing the waveform shown in the top-right plot. The signal-to-noise ratio was 0.5, and the signal was difficult to distinguish from the noise by eye. The 100 MHz filter was then cross-correlated with the noisy signal, the result of which is shown in the bottom plot. Here, the presence of the pulse can be clearly seen.

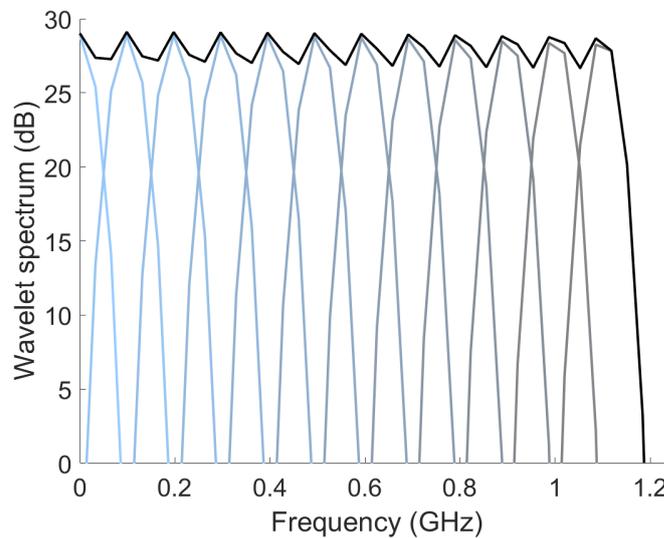


Figure 3.49: Relative magnitude of spectra in dB vs. frequency in GHz for of each of the twelve filters applied to the dark current signal. The black line represents the sum of all the spectra.

Since, as mentioned before, the specifics of the fluctuation events were not known beforehand, the optimal matched filter for these events could not be defined. Instead, a set of filters was defined to cover the entire bandwidth of the raw signal from the oscilloscope, with $f_0 \in \{0, 100, 200, \dots, 1100\}$ MHz. The frequency spectra of all the filters are shown in Fig. 3.49, demonstrating that they provide reasonably good coverage of the entire bandwidth of the signal. The frequency ripple of the sum of the frequency spectra, defined as the ratio between the highest and lowest points in the passband, was 2.35 dB, corresponding to an amplitude ratio of 1.3. This was considered small enough for the purposes of detecting the presence of a signal.

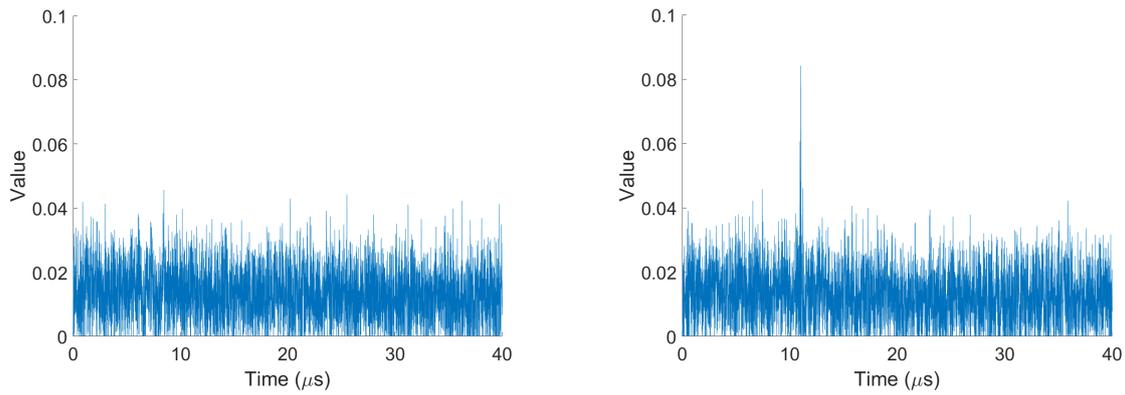


Figure 3.50: Cross-correlation result vs. time in μs using the 0 MHz filter (see text). Left: At an applied voltage of 0 V. Right: At an applied voltage of 1800 V. Occasional spikes are visible such as the one near the 10 μs mark.

The result y_n of cross-correlating the 0 MHz filter with the captured waveforms are shown in Fig. 3.50. The figure shows a typical waveform with no voltage across the gap and one with a voltage of 1800 V. There is a possible event near the 10 μs mark in the 1800 V plot. No such spikes were seen with any of the data recorded at 0 V, whereas some of the time windows at higher voltages showed more spikes such as this one. No interesting features were seen with any of the other filters. Since a key component of experimentally verifying the dislocation model was measuring the distribution of time intervals between events and its dependence on the applied electric field, it was decided that future measurements should focus on this. In the measurement thus far, the waveforms were captured in 40 μs long windows, which turned out to be too short to provide any useful information about the event interval.

Although no evidence of fluctuations was visible by eye in any of the cross-correlation results apart from those from the 0 MHz filter, it was decided to verify this quantitatively. This was done by plotting the median correlation value as a function of applied voltage for each of the filters, as shown in Fig. 3.51. All the of the plots were normalised by dividing each value by the value obtain with that filter at 0 V. In this plot, it can be seen that the median correlation with the 0 MHz filter consistently increased with voltage between 1000 V and 1800 V. No such trend was visible with any of the other filters.

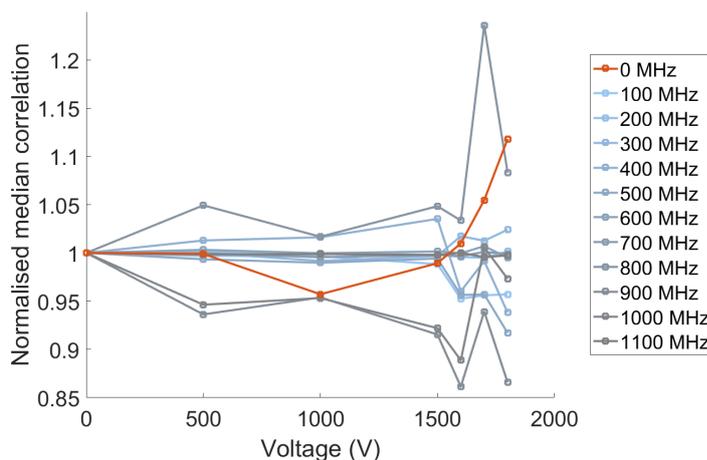


Figure 3.51: Median cross-correlation result using each of the filters vs. applied voltage in V. Different coloured curves represent different filters, as indicated on the legend. The cross-correlation values for each filter have been divided by the corresponding value at 0 V. A strong dependence on voltage is visible for the 0 MHz filter, plotted as a red curve.

Measurements with Ridged Electrodes

In an attempt to obtain a better signal, an opportunity was taken to replace the electrodes used to make the measurement with ones featuring a machined ‘ridge’ on their surface, as pictured in Fig. 3.52. These electrodes were left over unused from a different experiment [129]. The ridge on each electrode had a flat top surface that was 10 mm long and 1 mm wide.

This decision was motivated by a consideration of the effect of the capacitance of the inter-electrode gap on the transfer impedance between the current in the gap and the resulting voltage at the feedthrough that this current produces. The transfer impedance can be modelled by putting the characteristic impedance of the cable connected to the feedthrough, Z_0 , in parallel with the impedance of the inter-electrode gap, which has a capacitance C_{gap} , yielding an equivalent impedance Z_{eq} :

$$Z_{eq}(\omega) = \left(\frac{1}{Z_0} + i\omega C_{gap} \right)^{-1} \quad (3.23)$$

where ω is the signal frequency. Given the gap spacing of 40 μm , diameter of 40 mm, and the permittivity of free space, the gap capacitance was calculated using (3.20) to

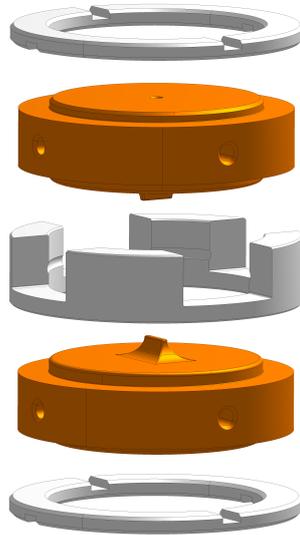


Figure 3.52: A diagram of the ridged electrodes and ceramic spacers [73].

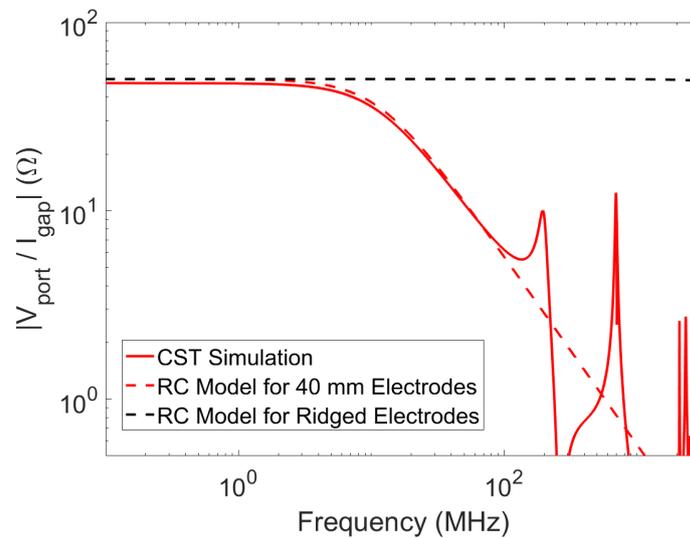


Figure 3.53: Transfer impedance between a current in the gap and the voltage at the chamber feedthrough in Ω vs. frequency in MHz. Solid red line: finite-element simulation result with 40 mm-diameter circular electrodes. Dashed red line: resistor-capacitor model of the 40 mm-diameter circular electrodes. Dashed black line: resistor-capacitor model of the ridged electrodes.

be 278 pF. This capacitance value was used along with $Z_0 = 50 \Omega$ in (3.23) to obtain a value for the transfer impedance which is shown in Fig. 3.53. In this figure, it can be seen that the model matches the finite-element simulation result remarkably well up to about 50 MHz, which is surprising given its simplicity. Above this frequency, the resonant modes of the chamber begin playing a role, and the model no longer holds. However, the model does show that if the capacitance of the gap could be reduced, the transfer impedance could be increased, resulting in larger voltage signals for a given magnitude of current fluctuation. To yield the minimum possible gap capacitance, the two ridged electrodes were rotated 90° with respect to each other, such that the ridges were crossed, resulting in an area exposed to high electric fields of about 1 mm^2 in area. Neglecting the effect of fringing fields, this yielded a capacitance of 0.22 pF, thus significantly reducing the effect of the capacitance on the transfer impedance. This is demonstrated by the black dashed line in Fig. 3.53, which shows a substantial improvement in the transfer impedance compared to the 40 mm circular electrodes.

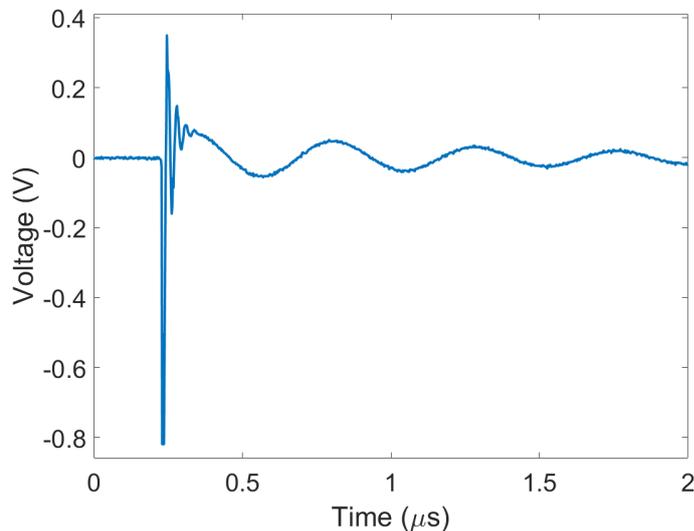


Figure 3.54: Signal in V vs. time in μs for a voltage spike and subsequent ringing observed with the ridged electrodes, at an applied voltage of 2000 V.

The use of the new electrodes appeared to bring about the desired effect, with spikes becoming clearly visible in the raw signal without any need for filtering. An example is shown in Fig. 3.54. In this plot, the superposition of two decaying sinusoids can be seen: one with a frequency of 25 MHz which decays within about 200 ns, and another with

a frequency of 2.5 MHz that takes several microseconds to decay. These are thought to be the responses of various parts of the system, to an impulse of dark current in the gap. As no resonances at either 2.5 or 25 MHz were observed when the LES was measured with a VNA, these might possibly be the response of other parts of the setup, such as the high-voltage power supply.

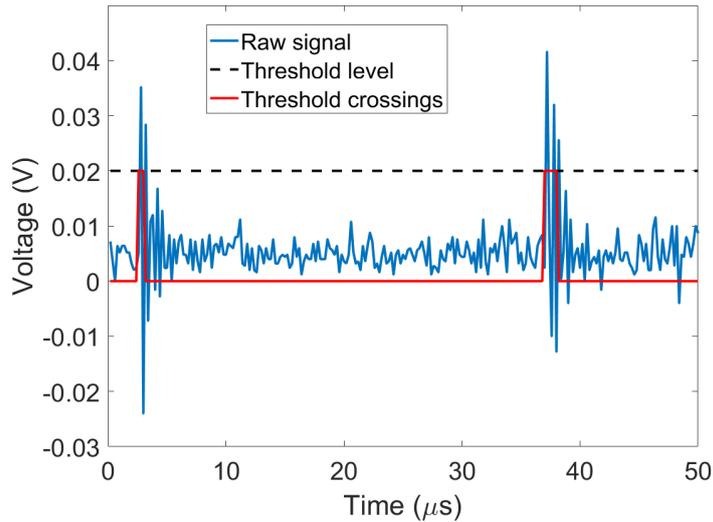


Figure 3.55: Blue curve: measured signal voltage in V vs. time in μs . Black dashed line: threshold level used to identify voltage spikes. Red curve: a representation of the output of the triggering logic, showing a nonzero value whenever the signal was above the threshold and a zero value otherwise.

Apart from the poor signal-to-noise ratio, the other issue identified in the first measurement was the short time windows over which the signal was recorded, making it difficult to extract any useful information about the time intervals between successive events. The length of the time window was limited by the memory capacity of the oscilloscope, as well as the time needed to save each waveform to disk. A window length of 10 million samples was found to be the practical limit, and each waveform of this size took several tens of seconds to save. To maximise the number of events observed per window, the sample rate was set to 5 MS/s, the minimum needed to reliably sample the 2.5 MHz ringing after the current impulse, according to the Nyquist criterion [130]. The appearance of this decaying response was used to infer that a spike had occurred, since neither the initial spike nor the 25 MHz ringing could be seen at this sample rate. Fig. 3.55 shows an example waveform captured with this sample rate. It also shows

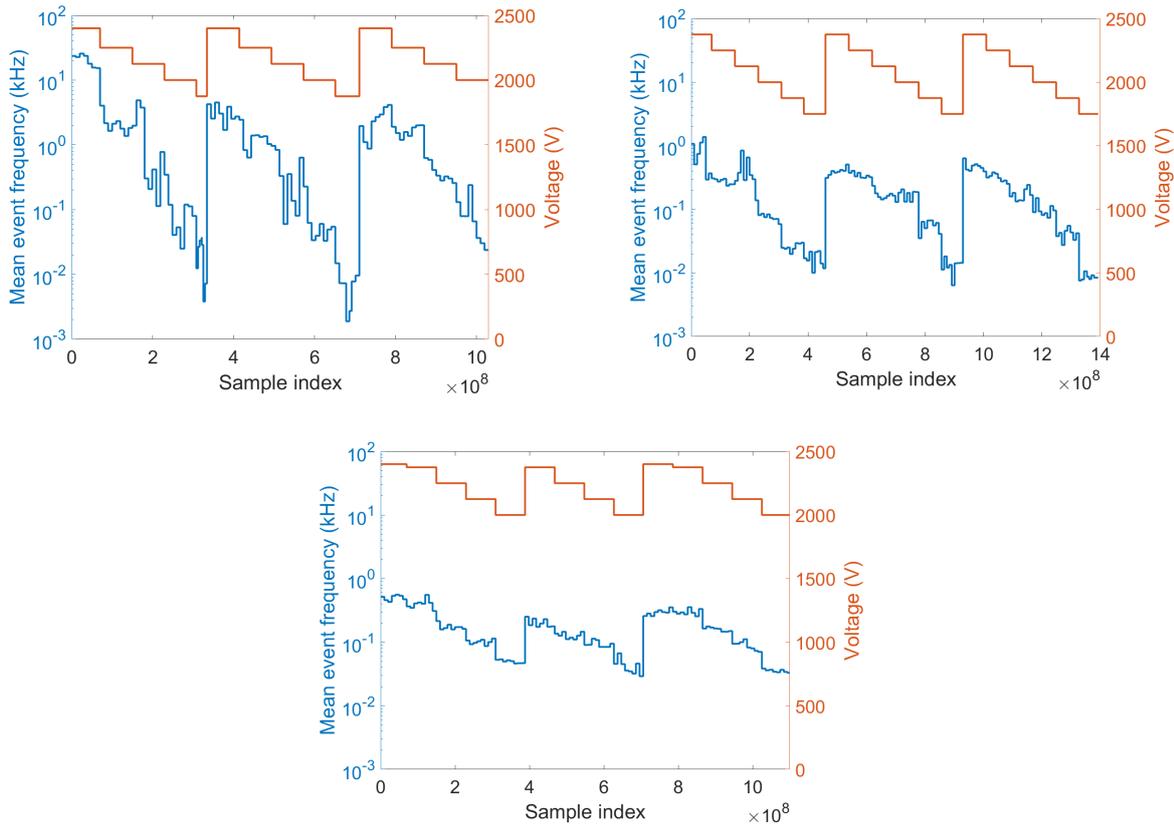


Figure 3.56: Measurements of dark-current fluctuations in the LES. In each plot, the blue curve represents the mean event frequency in kHz vs. the cumulative number of data samples taken while the red curve represents the applied voltage in V vs. the cumulative number of data samples taken. The decrease in the maximum event rate, as well as its variance, from each day to the next can be seen in this figure. Plots of fluctuation rate vs. voltage for each scan are shown in Fig. 3.57. Top left: measurement on the 29/10/2019. Top right: measurement on the 31/10/2019. Bottom: measurement on the 1/11/2019.

how the events were counted using a simple threshold-crossing trigger. The threshold level was set high enough that the background noise was unlikely to cross it, but most fluctuation events did. The time at which the threshold was first crossed during a given event was taken as the time at which that event occurred.

Several fluctuation measurements were performed over the course of a number of days to obtain an ample amount of data. Fig. 3.56 summarises a series of three consecutive measurements. On each day, three separate voltage scans were performed to check the repeatability of the fluctuation rate. At each voltage setting, eight two-second time windows were recorded at a sample rate of 5 MS/s, corresponding to 80 million sample points. In between the fluctuation measurements, the system was reconfigured for

conditioning and pulsed at a repetition rate of 2 kHz. One day of continuous running translates to about 170 million pulses.

Statistical Analysis

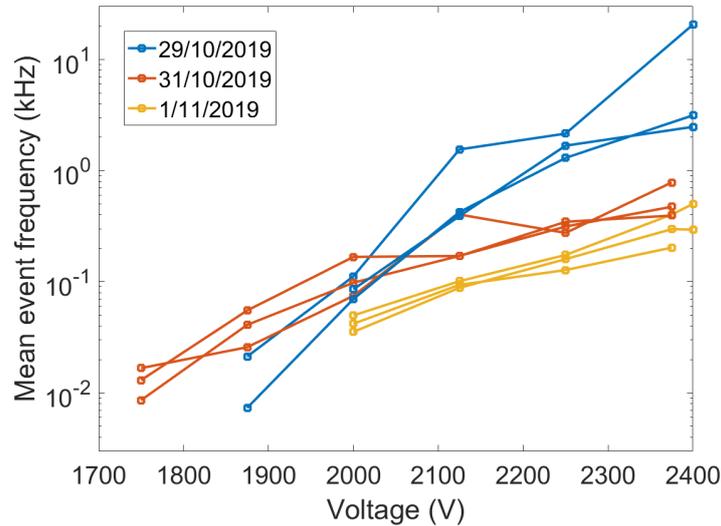


Figure 3.57: Mean event frequency in kHz vs. applied voltage in V. Each curve represents one of the voltage scans shown in Fig 3.56. Blue: measurements from the 29/10/2019. Red: measurements from the 31/10/2019. Yellow: measurements from the 1/11/2019.

Fig. 3.57 summarises the mean event rate measured over the course of the three days as a function of the applied voltage. In this figure, each of the individual voltage scans is shown as a separate plot. From this figure, it can be seen that there is a very strong dependence of the event rate on the applied voltage, as expected from the mobile dislocation model, which predicts that the rate of dislocation motion should increase as the surface stress is increased. The curves also seemed to vary somewhat between the individual scans from each day, and there was a clear difference in the rate above 2200 V between the different days.

Fig. 3.58 shows the decay in the mean event rate with time more clearly. At 2250 V, which was the highest voltage common to all the scans, the mean event rate decreased by about an order of magnitude between the first and last measurement. This is also consistent with the dislocation model, as the electrodes were conditioned in this period. The physical explanation of this is that the pulsed stress applied to the copper sur-

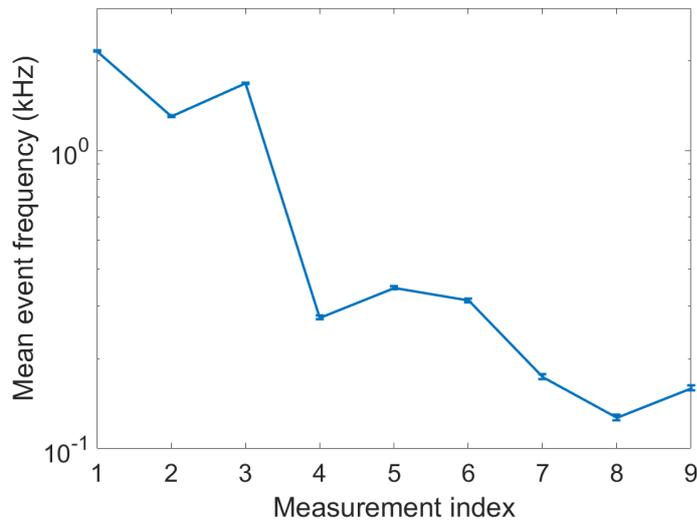


Figure 3.58: Measured event frequency at an applied voltage of 2250 V in kHz vs. voltage scan number. Each scan corresponds to one of the lines in Fig. 3.57, or one of the three voltage scans in each plot in Fig. 3.56, arranged in chronological order. Scans 1, 2 and 3 were performed on the 29/10/2019, scans 4, 5 and 6 were performed on the 31/10/2019, and scans 7, 8 and 9 were performed on the 1/11/2019. Error bars are shown based on uncertainty in measuring the rate.

face caused work hardening of the material, impeding the motion of dislocations and reducing the event rate. There was a significant variation between the rates measured on different scans on the same day. It should be noted that this is not due to the uncertainty in the measurement of the rate, which, assuming an exponential distribution of time intervals, can be expressed as:

$$\sigma_{\lambda} = \frac{\lambda_0}{\sqrt{N}}, \quad (3.24)$$

where λ_0 is the mean event rate, and N is the number of events used for that data point. The number of events measured for each data point in Fig. 3.58 ranged from 2549 to 34490, corresponding to an uncertainty ranging from 0.5% to 2.2%. The deviations in event rate were much larger than this, suggesting that the rate and possibly the probability distribution change over time as well.

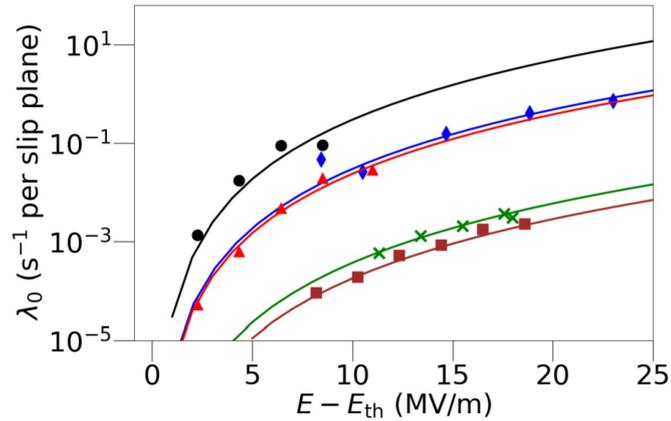


Figure 3.59: Event rate per slip plane λ_0 in s^{-1} vs. surface electric field in MV/m. A threshold E_{th} has been subtracted from the surface field value. Different colours represent measurements performed on different days. Points represent measured data while curves represent fits of the model as per Eq. (3.25) [131].

An attempt was made to fit the dependence of the mean event rate as a function of the field to predictions made by the original model. The mean birth rate of mobile dislocations per slip plane in the $n = 0$ state, λ_0 , was predicted to be:

$$\lambda_0 = \frac{25\kappa C_{tc}}{G^2 b \Delta \rho} \sigma_0^2 \exp\left(-\frac{E_a - \Omega \sigma_0}{k_B T}\right), \quad (3.25)$$

where T is the absolute temperature, and σ_0 is the stress in the case where there are no mobile dislocations, i.e. $n = 0$. The remaining parameters are constants, listed in Table 3.3. If the surface electric field E were the only source of stress, the stress would depend on E as:

$$\sigma_0 = \frac{1}{2} \epsilon_0 (\beta E)^2, \quad (3.26)$$

where β is the local geometric field-enhancement factor. However, it was found that the model and data agreed much better if a threshold field E_{th} was added such that:

$$\sigma_0 = \frac{1}{2} \epsilon_0 (\beta (E - E_{th}))^2 \quad (3.27)$$

Symbol	Value	Unit	Description
κ	0.41	-	Kinetic factor of dislocation creation
C_t	2.31	km/s	Speed of sound in copper
c	1	μm^{-1}	Density of barriers in the slip plane
G	48	GPa	Shear modulus of copper
b	0.25	nm	Magnitude of the Burger's vector of the dislocations
$\Delta\rho$	0.1	μm^{-1}	Unit change in dislocation density
E_a	0.8	eV	Activation energy of creating a mobile dislocation
Ω	5.6	eV/GPa	Activation volume for releasing new mobile dislocations
k_B	86.2	$\mu\text{eV/K}$	Boltzmann's constant

Table 3.3: Parameters of the MDDF model [131].

A possible source for this threshold field are pre-existing stresses in the material. Eq. (3.25) was fitted to different data sets by varying the value of β and E_{th} , the results of which are shown in Fig. 3.59. In this figure, each set of points corresponds to data taken on a given day. In cases where multiple voltage scans were performed on the same day, the data from all the voltage scans on that day was averaged. The value of E_{th} was found to vary between 21 and 29 MV/m, whereas the β varied between 0.6 and 3.2. The good fits between the experimental data and the prediction were considered strong evidence that the fluctuations observed were indeed caused by dislocation motion.

The exact shape of the probability distributions of the measured events was another feature that exhibited good consistency with the prediction of the dislocation model. The measured probability density function (PDF) was compared with two models. One was the exponential distribution, which is the distribution of time intervals between events that occur independently at a constant average rate [132]. Its PDF is defined as:

$$f(\tau) = \lambda_0 - e^{-\lambda_0\tau}, \quad (3.28)$$

where τ is the time interval and λ_0 is the average rate. The other was a two-parameter hypoexponential distribution [133]. In this case, the events are not independent of each other. Instead, there are two types of event, with every event of the first type

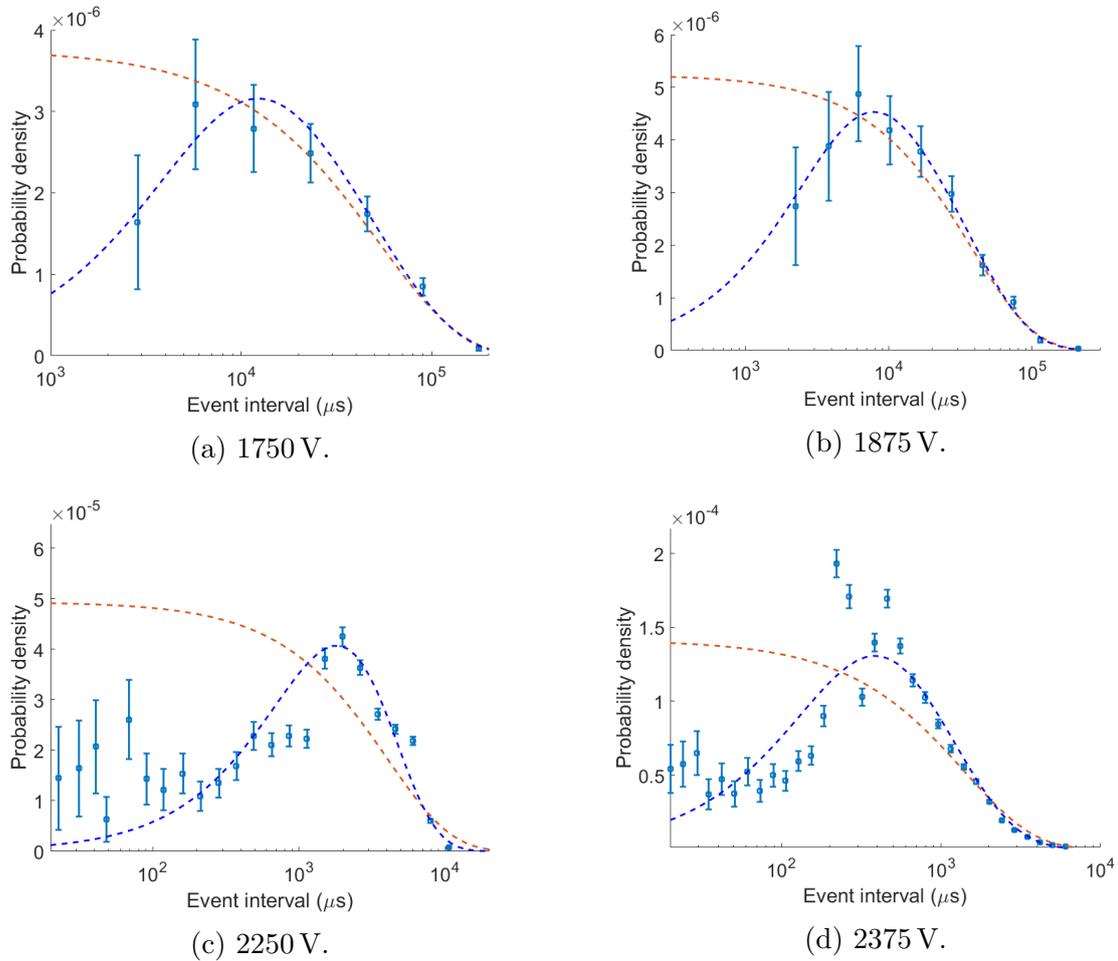


Figure 3.60: Probability density vs. time interval between successive events in μs , as measured in one of the scans on the 31/10/2019. The points represent measured data, the red dashed lines represent the best-fit exponential distribution, and the blue dashed lines represent the best-fit two-parameter hypoexponential distribution. The fits were performed by minimising χ^2 over all the data points. Top left: at 1750 V. Top right: 1875 V. Bottom left: at 2250 V. Bottom right: at 2375 V. It is believed that because the dislocation process was non-stationary and its statistics varied over the duration of the measurement, the data points for 2250 V and 2375 V have a large spread compared to the sample standard deviation indicated by the error bars.

being followed by an event of the second type. In this case, these represent the birth of a mobile dislocation (a transition from $n = 0$ to $n = 1$) and the death of a mobile dislocation (a transition from $n = 1$ to $n = 0$), with average rates λ_0 and μ_1 respectively. In principle, transitions between all the possible states should be taken into account (e.g. $n = 1$ to $n = 2$, etc.). However, since the system is most likely to be found in either the $n = 0$ or $n = 1$ states, as shown in Fig. 3.2, these two transitions should be by far the most likely to occur, and a distribution taking into account only these two transitions is expected to closely resemble the true distribution. The PDF of this two-parameter distribution is given by:

$$f(\tau) = \frac{\lambda_0 \mu_1}{\mu_1 - \lambda_0} (e^{-\lambda_0 \tau} - e^{-\mu_1 \tau}), \quad (3.29)$$

where λ_0 and μ_1 are the average rates of birth and death of mobile dislocations, with $\lambda_0 \neq \mu_1$, and τ is the time interval between successive births. Fits of both of these PDFs to experimentally measured distributions for different applied voltages are shown in Fig. 3.60. At each of the different voltages, it can be seen that the hypoexponential CDF approximates the experimental data better than the exponential CDF, further supporting a model in which there are two independent sources of events, which the dislocation model is an example of.

3.4 Discussion of Results

The measurements performed with the crossed electrodes in the LES clearly show the existence of fast fluctuations in the field-emitted current emitted by the cathode. The statistics of the time intervals between the events provide very strong evidence that they are indeed caused by the motion of dislocations within the material as predicted by Engelberg et. al.

No evidence for any such events was found in the experiments conducted on the RF structures under test in the XBoxes. One reason for this may have been the lower

effective noise floor in the LES, in part due to the lack of a pulsed electric field and thus no misalignment noise. The effect of a lower noise floor can be seen by the difference between the results from the 40 mm circular electrodes (Fig. 3.50) and the ridged electrodes (Fig. 3.54). However, the lack of observed fluctuations may also have been due to their relatively low rate. The highest rate of events measured in the LES was on the order of 10 kHz, meaning one event every 100 μ s. With the maximum pulse length typically used at the XBoxes or 200 ns, this rate of events corresponds to one event every 500 pulses. The oscilloscope used to record the data was not capable of acquiring waveforms and saving them to disk simultaneously, meaning that data could only be captured in relatively short windows several seconds apart due to the time needed to save the data to disk. Because of this, it is likely that the relatively rare pulses in which a fluctuation event could possibly be seen were not recorded at all, if they occurred while the oscilloscope was saving data to disk. The fluctuation rate measured in the LES was also seen to decrease significantly with further conditioning, which implies that fluctuation events would become even rarer in the XBoxes.

The argument could be made that with a greater surface area subject to high electric fields, one should expect a higher rate of events. However, there is experimental evidence, such as that discussed in Sec. 3.2.3, suggesting that the dark current emitted by a macroscopic device is often dominated by field emission from a very small region where the surface electric field is very slightly higher than in other regions. Since these experiments were based around measurements of this current, only geometric fluctuations that occurred in these regions would produce a measurable signal. This implies that a larger total surface area would be unlikely to result in a higher rate of fluctuations.

The direct measurement of the motion of dislocations in an RF structure undergoing high-power testing as a diagnostic technique for monitoring the progress of its conditioning may thus require further work or a different approach to become feasible. Real-time analysis of dark current waveforms for every pulse is one possible approach,

while measuring a different phenomenon linked to dislocation motion such as acoustic waves is another.

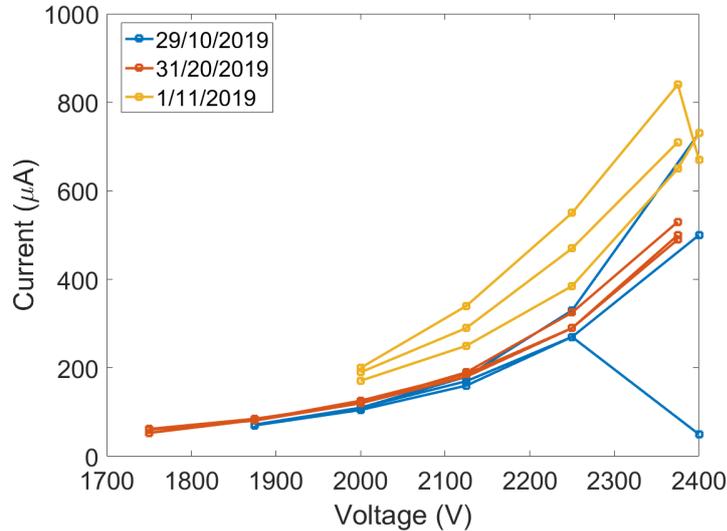


Figure 3.61: Measured field-emitted current in μA vs. applied voltage in V for each of the voltage scans shown in Fig. 3.56. Each curve represents one voltage scan. Blue: measurement on the 29/10/2019. Red: measurement on the 31/10/2019. Yellow: measurement on the 1/11/2019.

In Sec. 3.1.4, the jaggedness of current-voltage curves in field emission experiments, such as the ones shown in Fig. 3.7, was discussed. It was suggested that the jaggedness of the curves was related to the dark-current fluctuation phenomenon measured in the LES. A possible mechanism for this is that since the measurement of each point on the curve took a finite amount of time, the current could drift in a random-walk fashion between each data point, producing jagged curves. This would be consistent with the physical understanding of dislocation motion that has been put forward here, since it was reported that the current-voltage curves became much smoother after cooling to cryogenic temperatures, which is known to impede the movement of dislocations. Similar behaviour was seen in the current-voltage characteristics measured in the LES whilst the fluctuation measurements were being performed. Some of these current-voltage curves are shown in Fig. 3.61. Here, the relationship between the current and the voltage appears to have varied between individual scans on each day, and this variation was the most pronounced at the highest voltage, which was also when the greatest rate of fluctuations was observed.

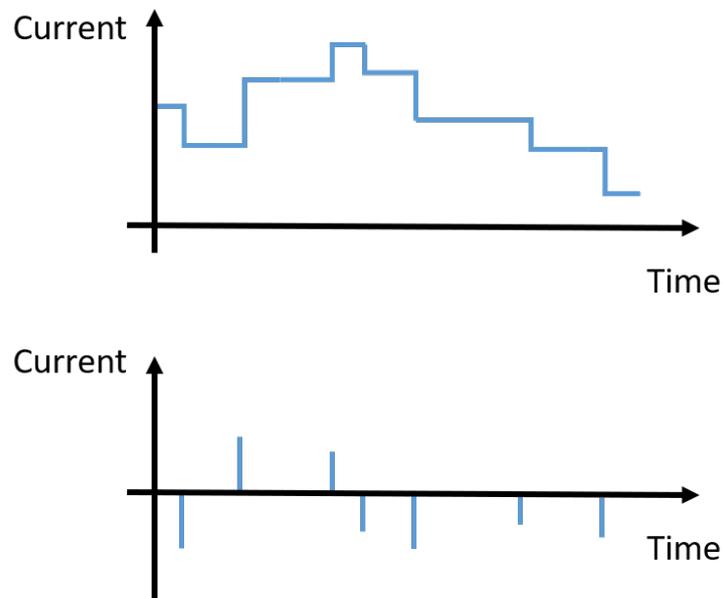


Figure 3.62: Current in arbitrary units vs. time in arbitrary units. Top: an idealised representation of the hypothesised true form of the field-emitted current in the LES. Each fluctuation event corresponds to a small random change in the total current, represented by the vertical steps. Bottom: a high-pass-filtered version of the waveform in the top plot, in which the fluctuation events are visible as spikes, but the change in DC current level is not visible. This is analogous to the output of the current setup.

These two phenomena could be related if each fluctuation event was accompanied by a step change in the total emitted current. This is a reasonable assumption, since the motion of a dislocation corresponds to a deformation, which in turn could modify the field-emission characteristics by changing the field-enhancement factor or emitting area. With fluctuation events occurring constantly, the total emitted current would then undergo a random walk of steps up or down, causing it to drift over time. A greater rate of fluctuation would thus correspond to a larger expected change in current after a fixed time interval. An idealised representation of what the field-emitted current waveform is believed to look like, compared with the signal measured in this experiment, is shown in Fig. 3.62.

Whether or not each fluctuation event indeed resulted in a change in the total emitted current could not be determined with the experimental setup used here, since the high voltage bias tee removed any DC component of the dark current. A future experiment could incorporate a DC current measurement with high time and voltage resolution to investigate this phenomenon.

Fluctuation measurements of the kind presented in this chapter could be used to identify better materials for high-gradient applications. In addition to testing the ultimate achievable surface electric field, fluctuation measurements could allow the conditioning process in different materials to be characterised more fully. Previous work has shown correlations between lower dislocation mobility and ductility with breakdown performance in DC tests [134]. Materials with low ductility such as tungsten, molybdenum and cobalt have been tried as alternatives to copper with very good breakdown performance [135]. For RF applications, however, other properties such as conductivity and ease of manufacturing are important considerations as well.

Chapter 4

Conclusions

In this thesis, two aspects of vacuum breakdowns, an important phenomenon limiting the maximum accelerating gradient of CLIC's accelerating structures, were discussed. These are the mechanism of their nucleation and their dependence on the available input power. Vacuum breakdown is a complex process involving many stages and encompassing many orders of magnitude in size, current, and power. Thus, a complete understanding of the entire process requires the study of a number of seemingly unrelated physical phenomena. Together, the two mechanisms discussed in this thesis provide the theoretical basis for three of the most important technological parameters for high-field system design: the influence of material properties on the achievable field, the origin and mechanism of conditioning, and the influence of radio-frequency design. Although this work was done in the context of the CLIC study, it has a good range of applicability to devices including high-gradient RF accelerating structures, high-field accelerating structures such as radiofrequency quadrupoles (RFQs), high-field electrostatic accelerators, and high-field vacuum electronics such as vacuum interrupters.

The work on breakdown nucleation presented in Ch. 3 was the first direct measurement of these dislocation dynamics in a system subject to surface electric-field levels typically seen in high-gradient accelerating structures. Sensitive measurements of the field-emitted current were performed in the Large Electrode System (LES), a table-top

experiment built to study vacuum breakdowns under conditions similar to those within CLIC accelerating structures. These measurements revealed brief impulses of current occurring at random intervals. The average rate of the impulses varied with the applied electric field in a manner consistent with the dislocation model, and the distribution of time intervals between impulses were fitted very well by a hypoexponential distribution, another prediction of the model. The mean rate of impulses also decreased as the sample was conditioned, supporting the idea that conditioning, the process by which structures become more resilient to breakdown after repeated exposure to high fields, is one of work hardening.

This is useful information which could be used to optimise the conditioning process. For example, it may be possible to shorten the conditioning time by hardening accelerating structures via thermal processes instead. Currently, the accelerating structures undergo a diffusion bonding process during their manufacture, which anneals and therefore softens them as a side effect. The use of hard copper electrodes in the LES has shown a decrease in conditioning time, though no hard copper accelerating structures have been tested at high power to date. Building on this, the fluctuation rate could also be used as an additional method of assessing the suitability of new materials and alloys for use in high-field applications. These experiments provide further evidence of the suitability of materials with low dislocation mobility such as molybdenum, cobalt, and tungsten for high breakdown performance.

The existence of a measurable signal that is directly coupled to the dislocation dynamics behind the nucleation of breakdowns offers the potential for an alternative way of measuring the state of conditioning of an accelerating structure to the breakdown rate. This information could, for example, allow conditioning to be done with a minimum of breakdowns and thus little degradation of the surface. It may also serve as a warning signal that a breakdown is imminent, which could be prevented by reducing the gradient temporarily. An attempt to measure these fluctuations in the dark current emitted by accelerating structures under high-power testing at the high-gradient X-band test stands at CERN has not been successful so far. Despite attention being

paid to maximise the quality of the measurement, various sources of background noise meant that the sensitivity with which the emitted dark current could be measured was fundamentally limited. Despite this, other ways of measuring dislocation dynamics in a structure in operation could include acoustic methods.

In parallel with the study of breakdown nucleation, the other aspect of vacuum breakdowns discussed in this thesis is their interaction with the power source in the onset phase. The primary objective of this study was to develop an improved quantitative limit which defines the maximum accelerating gradient at which a given structure design could operate. Such a quantity is very useful in the design phase, as it allows an optimal structure to be designed given other constraints such as RF-power limitations and its effect on beam stability.

Quantities such as Kilpatrick's criterion, which defines a maximum safe value of surface electric field as a function of operating frequency, have been used for this purpose in the past. However, evidence from high-gradient accelerating structures tested at CERN and other laboratories have shown behaviours inconsistent with this. In particular, it has been shown that different geometries at the same operating frequency could reach different maximum gradients. Currently, CLIC accelerating structures are designed using a quantity known as S_c or the modified Poynting vector, a measure of the local RF power density, with the expectation that a greater value of S_c corresponds to a greater probability of breakdown at that location. In the process of optimising the geometry of a structure, the maximum value of S_c cannot exceed a certain fixed value.

Despite the modified Poynting vector's consistency with test results from high-group-velocity accelerating structures, discrepancies have been found between the prediction and experimental results for other kinds of structures, such as deflecting structures or ones designed for accelerating low-velocity particles. Building on the idea of power flow, as opposed to the electric field, to quantify the threshold at which breakdowns are likely to occur, a new quantity was proposed. This quantity is the equilibrium electric field under breakdown loading, denoted E^* . It is to be calculated for each po-

tential breakdown site, where an emitter of charged particles is assumed to be located. Charged particles absorb energy from an electric field, meaning that the loaded field is a function of the amount of RF power that can be supplied to sustain the particle emission. The calculation of E^* requires two components: an equivalent source impedance determining the coupling of the breakdown site to the RF power source, and a function relating the emitted current to the surface field.

A number of methods of calculating the coupling of RF power to the breakdown were developed and evaluated. Among these were analytical equations derived from circuit models of travelling-wave structures, as well as numerical methods performed using a 3D finite-element solver. The method was extended to include transient behaviour, which was necessary to correctly model structures with low group velocities, in which the local stored energy may be a significant source of RF power for the breakdown. When applied to travelling-wave structures, the analytical models and their numerical equivalents gave consistent results. The availability of an analytical model means that the E^* quantity can be used in a numerical optimisation of a design, not needing computationally expensive finite-element simulations for every breakdown site and iteration of the optimisation algorithm.

The E^* quantity was compared with experimental data, giving encouraging results. The cases considered were: a T24 X-band accelerating structure, an X-band CLIC Crab Cavity prototype, and an S-band backward-travelling-wave medical proton linac structure. The locations within those structures which experienced the greatest density breakdowns during testing matched up very well to the locations which had the greatest value of E^* . In each case, the correspondence between the calculation and experiment was better for E^* than for S_c . The calculations for the two X-band structures also gave consistent values for their maximum gradients.

In another set of simulations, it was established that E^* gave similar results to S_c in terms of both breakdown location and maximum gradient for geometries similar to those of CLIC accelerating structures. This was considered further evidence of

the applicability of the E^* model, since the validity of S_c for CLIC-like accelerating structures is well-backed by experimental results.

The E^* model has performed very well in the cases to which it has been applied so far, showing the potential to become a very general model of vacuum breakdowns applicable to all kinds of high-field devices. Some areas for further work include working out a consistent way of accounting for the operating frequency, since in this work different values of the parameters of the model had to be used to fit data from X-band, S-band, and DC experiments. The behaviour of E^* for DC breakdowns is of particular interest due their broadband nature. One advantage of the LES at CERN is the relative ease with which electrodes with a custom geometries could be manufactured, meaning that dedicated tests could be done in this system.

Beyond the LES, there are many high-field and high-gradient experiments that could help further refine the E^* model. Two cases of interest that have been identified are choke-mode structures for CLIC designed at Tsinghua University, and the high-voltage components of the neutral beam injector of ITER. In both cases, attention has been paid to the breakdown performance, and an ample amount of experimental data is available along with empirical models.

Bibliography

- [1] “LHC Design Report Vol.1: The LHC Main Ring,” O. S. Bruning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole, and P. Proudlock, Eds., Jun. 2004. DOI: 10.5170/CERN-2004-003-V-1.
- [2] G. Aad *et al.*, “Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc,” *Physics Letters B*, vol. 716, no. 1, pp. 1–29, 2012, ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2012.08.020>. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S037026931200857X>.
- [3] G. Apollinari, I. Béjar Alonso, O. Brüning, M. Lamont, and L. Rossi, *High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report*, ser. CERN Yellow Reports: Monographs. Geneva: CERN, 2015. DOI: 10.5170/CERN-2015-005. [Online]. Available: <https://cds.cern.ch/record/2116337>.
- [4] R. Steerenberg, “LHC Report: Another run is over and LS2 has just begun. . .,” Dec. 2018. [Online]. Available: <https://home.cern/news/news/accelerators/lhc-report-another-run-over-and-ls2-has-just-begun>.
- [5] E. Wilson, *An Introduction to Particle Accelerators*. Oxford University Press, 2001, ISBN: 0-19-850829-8.
- [6] L. Evans and S. Michizono, “The International Linear Collider Machine Staging Report 2017,” Nov. 2017. arXiv: 1711.00568 [physics.acc-ph].
- [7] Compact Linear Collider Project, “CLIC accelerator schematic diagrams,” Jan. 2019, General Photo. [Online]. Available: <https://cds.cern.ch/record/2655160>.

- [8] “The International Linear Collider Technical Design Report - Volume 3.II: Accelerator Baseline Design,” C. Adolphsen *et al.*, Eds., 2013. arXiv: 1306.6328 [physics.acc-ph].
- [9] D. Reschke, “Performance of Superconducting Cavities for the European XFEL,” in *7th International Particle Accelerator Conference (IPAC2016)*, Busan, Korea, 2016.
- [10] “SwissFEL - Conceptual Design Report,” R. Ganter, Ed., 2010.
- [11] M. J. Stuart, J. Osborne, R. Corsini, D. Schulte, M. Draper, P. N. Burrows, S. Stapnes, and N. Catalán Lasheras, “The Compact Linear Collider (CLIC) – Project Implementation Plan,” Dec. 2018. DOI: 10.23731/CYRM-2018-004.
- [12] A. Grudiev and W. Wuensch, “Design of the CLIC main Linac accelerating structure for CLIC conceptual design report,” no. EuCARD-CON-2010-073. CERN-ATS-2010-212, 3 p, Sep. 2010. [Online]. Available: <https://cds.cern.ch/record/1346987>.
- [13] N. C. Lasheras, “CLIC: Key technology developments for the CLIC accelerator,” in *CERN Academic Training Lectures*, Geneva, Switzerland, 2018.
- [14] T. G. Lucas, T. Argyopolous, M. J. Boland, N. Catalan-Lasheras, R. P. Rassool, C. Serpico, M. Volpi, and W. Wuensch, “Dependency of the capture of field emitted electron on the phase velocity of a high-frequency accelerating structure,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 914, pp. 46–52, 2019. DOI: 10.1016/j.nima.2018.10.166.
- [15] J. Sekutowicz *et al.*, “Design of a low loss SRF cavity for the ILC,” *Conf. Proc. C*, vol. 0505161, C. Horak, Ed., p. 3342, 2005.
- [16] M. Aicheler, P. Burrows, M. Draper, T. Garvey, P. Lebrun, K. Peach, N. Phinney, H. Schmickler, D. Schulte, and N. Toge, *A Multi-TeV Linear Collider Based on CLIC Technology: CLIC Conceptual Design Report*, ser. CERN Yellow Reports: Monographs. Geneva: CERN, 2012. DOI: 10.5170/CERN-2012-007. [Online]. Available: <https://cds.cern.ch/record/1500095>.

- [17] M. Volpi, private communication.
- [18] B. J. Woolley, “High Power X-band RF Test Stand Development and High Power Testing of the CLIC Crab Cavity,” PhD thesis, 2015.
- [19] A. Degiovanni, W. Wuensch, and J. Giner Navarro, “Comparison of the conditioning of high gradient accelerating structures,” *Physical Review Accelerators and Beams*, vol. 19, no. 3, p. 032001, Mar. 2016. DOI: 10.1103/PhysRevAccelBeams.19.032001.
- [20] N. Catalan-Lasheras, A. Degiovanni, S. Doebert, W. Farabolini, J. Kovermann, G. McMonagle, S. Rey, I. Syratchev, L. Timeo, W. Wuensch, B. Woolley, and J. Tagg, “Experience Operating an X-band High-Power Test Stand at CERN,” no. CERN-ACC-2014-0166, Jun. 2014.
- [21] E. Senes *et al.*, “Results of the Beam-Loading Breakdown Rate Experiment at the CLIC Test Facility CTF3,” in *Proc. of International Particle Accelerator Conference (IPAC’17), Copenhagen, Denmark, 14–19 May, 2017*, (Copenhagen, Denmark), ser. International Particle Accelerator Conference, Geneva, Switzerland: JACoW, May 2017, pp. 1348–1351, ISBN: 978-3-95450-182-3.
- [22] N. Shipman, “Experimental Study of DC Vacuum Breakdown and Application to High-Gradient Accelerating Structures for CLIC,” PhD thesis, 2014.
- [23] L. M. Redondo, A. Kandratsyev, M. J. Barnes, S. Calatroni, and W. Wuensch, “Solid-state Marx generator for the Compact Linear Collider breakdown studies,” in *2016 IEEE Int. Power Modul. High Volt. Conf.*, IEEE, Jul. 2016, pp. 187–192. DOI: 10.1109/IPMHVC.2016.8012824.
- [24] I. Profatlova, X. Stragier, S. Calatroni, A. Kandratsyev, E. Castro, and W. Wuensch, “Breakdown localisation in a pulsed DC electrode system,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 953, 2020. DOI: 10.1016/j.nima.2019.163079.
- [25] H. Timko, F. Djurabekova, K. Nordlund, L. Costelle, K. Matyash, R. Schneider, A. Toerklep, G. Arnau-Izquierdo, A. Descoeurdes, S. Calatroni, M. Taborelli,

- and W. Wuensch, “Mechanism of Surface Modification in the Plasma-Surface Interaction in Electrical Arcs,” *Physical Review B*, vol. 81, no. 18, p. 184109, May 2010. DOI: 10.1103/PhysRevB.81.184109.
- [26] J. Kovermann, “Comparative studies of high-gradient RF and DC breakdowns,” PhD thesis, 2010.
- [27] R. Rajamaki, “Vacuum arc localization in CLIC prototype radio frequency accelerating structures,” Master’s thesis, Apr. 2016.
- [28] W. Wuensch, “Advances in the Understanding of the Physical Processes of Vacuum Breakdown,” in *High Gradient Accelerating Structure*, World Scientific, Dec. 2014, pp. 31–50, ISBN: 978-981-4602-09-9. DOI: 10.1142/9789814602105_0003.
- [29] W. Wuensch, A. Degiovanni, S. Calatroni, A. Korsbäck, F. Djurabekova, R. Rajamäki, and J. Giner-Navarro, “Statistics of vacuum breakdown in the high-gradient and low-rate regime,” *Phys. Rev. Accel. Beams*, vol. 20, p. 011007, 1 Jan. 2017. DOI: 10.1103/PhysRevAccelBeams.20.011007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.20.011007>.
- [30] J. W. Wang and G. A. Loew, “Field emission and RF breakdown in high-gradient room-temperature linac structures,” Tech. Rep., 1997.
- [31] F. Djurabekova, “Field emission from the first principles: Effect of point defects on the value of the workfunction,” in *Mini MeV Arc Meeting*, 2018.
- [32] T. P. Wangler, *RF Linear Accelerators - second, completely revised and enlarged edition*. Wiley-VCH, 2008, p. 450.
- [33] A. Latina *et al.*, “CompactLight design study,” in *60th ICFA Advanced Beam Dynamics Workshop on Future Light Sources*, 2018, pp. 85–88. DOI: 10.18429/JACoW-FLS2018-WEP1WC02.
- [34] O. Luiten, “SMART*LIGHT: a Dutch table-top alternative for synchrotron light sources,” English, KNAW brainstorm Toekomstige Grootchalige Onderzoeksfaciliteiten ; Conference date: 18-03-2015 Through 18-03-2015, 2015.

- [35] W. Kilpatrick, "Criterion for vacuum sparking designed to include both RF and DC," *Rev. Sci. Instrum.*, vol. 28, pp. 824–826, 1957. DOI: 10.1063/1.1715731.
- [36] T. J. Boyd Jr., "Kilpatrick's Criterion," *Los Alamos Group AT-1 Report*, vol. AT-1:82-28, Feb. 1982.
- [37] W. Wuensch, I. Wilson, H.-H. Braun, and M. Luong, "A very high gradient test of a 30-GHz single cell cavity," in *7th European Particle Accelerator Conference (EPAC 2000)*, Aug. 2000, pp. 271–273.
- [38] E. R. Castro, "Microscopy at CERN," in *Mini MeV Arc meeting*, Tartu, Estonia, 2017.
- [39] K. Szypula, "HG TW Proton Linac endoscope inspection," Tech. Rep., 2019.
- [40] A. P. Fontenla, "Post-mortem analysis: SEM imaging review," in *CLIC Workshop*, Geneva, Switzerland, 2015.
- [41] W. Wuensch, "The Scaling of the Traveling-Wave RF Breakdown Limit," CERN, Geneva, Tech. Rep. CERN-AB-2006-013. CLIC-Note-649, Jan. 2006. [Online]. Available: <http://cds.cern.ch/record/932674>.
- [42] A. Grudiev, S. Calatroni, and W. Wuensch, "New local field quantity describing the high gradient limit of accelerating structures," *Physical Review Special Topics - Accelerators and Beams*, vol. 12, no. 102001, 2009. DOI: 10.1103/PhysRevSTAB.12.102001.
- [43] S. Dohert, C. Adolphsen, G. Bowden, D. Burke, J. Chan, V. Dolgashev, J. Frisch, K. Jobe, R. Jones, J. Lewandowski, R. Kirby, Z. Li, D. McCormick, R. Miller, C. Nantista, J. Nelson, C. Pearson, M. Ross, D. Schultz, T. Smith, S. Tantawi, J. Wang, T. Arkan, C. Boffo, H. Carter, I. Gonin, T. Khabiboulline, S. Mishra, G. Romanov, N. Solyak, Y. Funahashi, H. Hayano, N. Higashi, Y. Higashi, T. Higo, H. Kawamata, T. Kume, Y. Morozumi, K. Takata, T. Takatomi, N. Toge, K. Ueno, and Y. Watanabe, "High gradient performance of NLC/GLC X-band accelerating structures," in *Proceedings of the 2005 Particle Accelerator Conference*, 2005, pp. 372–374.

- [44] S. Döbert, R. Fandos, A. Grudiev, S. Heikkinen, J. A. Rodriguez, M. Taborelli, W. Wuensch, C. Adolphsen, and L. Laurent, “High power test of an X-band slotted-iris accelerator structure at NLCTA,” in *2007 Particle Accelerator Conference (PAC2007)*, Albuquerque, New Mexico, USA, 2007, p. 2191. [Online]. Available: <http://cds.cern.ch/record/1096164>.
- [45] J. W. Wang, G. A. Loew, R. J. Loewen, R. D. Ruth, A. E. Vlieks, I. Wilson, and W. Wuensch, “SLAC/CERN High Gradient Tests of an X-Band Accelerating Section,” in *Proceedings of the 1995 Particle Accelerator Conference (PAC95)*, Dallas, Texas, USA, 1995, p. 653.
- [46] S. Döbert, A. Grudiev, G. Riddone, M. Taborelli, W. Wuensch, R. Zennaro, S. Fukuda, Y. Higashi, T. Higo, S. Matsumoto, K. Ueno, K. Yokoyama, C. Adolphsen, V. Dolgashev, L. Laurent, J. Lewandowski, S. Tantawi, F. Wang, and J. W. Wang, “High power test of a loq group velocity X-band accelerator structure for CLIC,” in *Proceedings of the 24th Linear Accelerator Conference (LINAC08)*, Victoria, British Columbia, Canada, 2008, pp. 930–932.
- [47] S. G. Tantawi, V. Dolgashev, Y. Higashi, and T. Higo, “Status of high power tests of normal conducting single-cell structures,” in *Proceedings of the Eleventh European Particle Accelerator Conference (EPAC’08)*, Genoa, Italy, 2008, pp. 742–744.
- [48] R. Corsini, S. Döbert, R. Fandos, A. Grudiev, E. Jensen, T. Ramsvik, J. Rodriguez, J. Sladen, I. Syratchev, M. Taborelli, M. Tecker, P. Urschütz, I. Wilson, W. Wuensch, and Ö. Mete, “A high-gradient test of a 30 GHz copper accelerating structure,” in *Proceedings of LINAC 2006*, Knoxville, Tennessee, USA, 2006, pp. 761–763.
- [49] S. Döbert, in *CLIC Workshop*, Geneva, Switzerland, 2007.
- [50] J. Rodriguez, G. Arnau-Izquierdo, R. Corsin, S. Döbert, R. Fandos, A. Grudiev, I. Syratchev, M. Taborelli, F. Tecker, P. Urschütz, W. Wuensch, O. Mete, H. Aksakal, Z. Nergiz, and M. Johnson, “30 GHz High-Gradient Accelerating Struc-

- ture Test Results,” in *2007 Particle Accelerator Conference (PAC2007)*, Albuquerque, New Mexico, USA, 2007, p. 3818.
- [51] W. Wuensch, C. Achard, H.-H. Braun, G. Carron, R. Corsini, A. Grudiev, S. Heikkinen, D. Schulte, J. Sladen, I. Syratchev, F. Tecker, and I. Wilson, “30 GHz Power Production in CTF3,” in *Proceedings of the 2005 Particle Accelerator Conference (PAC05)*, Knoxville, Tennessee, USA, 2005, pp. 1695–1697.
- [52] C. Adolphsen, “Advances in normal conducting accelerator technology from the X-band linear collider program,” Jun. 2005, pp. 204–208, ISBN: 0-7803-8859-3. DOI: 10.1109/PAC.2005.1590396.
- [53] V. Dolgashev, “Experiments on gradient limits for normal conducting accelerators,” in *21st International Linear Accelerator Conference*, Aug. 2002, pp. 266–270.
- [54] J. Paszkiewicz, “Update on high-power testing of X-band RF structures at CERN,” in *International Workshop on Breakdown Science and High Gradient Technology (HG2018)*, Shanghai, China, 2018.
- [55] T. Higo, “Update of Nextef activity and TD26CC high-gradient result,” in *CLIC Workshop*, Geneva, Switzerland, 2019.
- [56] G. Burt *et al.*, “CLIC crab cavity final report,” Tech. Rep. EuCARD-REP-2013-028, 2013. [Online]. Available: <http://cds.cern.ch/record/1710313>.
- [57] B. Woolley, G. Burt, A. C. Dexter, R. Peacock, N. Catalan Lasheras, A. Degiovanni, A. Grudiev, G. McMonagle, I. Syratchev, W. Wuensch, and E. Rodriguez Castro, “High-gradient behaviour of a dipole-mode RF structure,” unpublished, 2020.
- [58] L. Millar, “High gradient test results,” in *CLIC Workshop*, Geneva, Switzerland, 2019.
- [59] A. Korsback, F. Djurabekova, L. M. Morales, I. Profatilova, E. R. Castro, W. Wuensch, S. Calatroni, and T. Ahlgren, “Vacuum electrical breakdown conditioning study in a parallel plate electrode pulsed dc system,” *Phys. Rev. Accel. Beams*, vol. 23, p. 033102, 3 Mar. 2020. DOI: 10.1103/PhysRevAccelBeams.

- 23.033102. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.23.033102>.
- [60] K. Eimre, S. Parviainen, A. Aabloo, F. Djurabekova, and V. Zadin, “Application of the general thermal field model to simulate the behaviour of nanoscale Cu field emitters,” *Journal of Applied Physics*, vol. 118, no. 3, p. 033303, 2015. DOI: 10.1063/1.4926490.
- [61] M. Veske, A. Kyritsakis, K. N. Sjobak, V. Zadin, A. Aabloo, and F. Djurabekova, “Dynamic coupling between particle-in-cell and atomistic simulations,” *Phys. Rev. E*, vol. 101, no. arXiv:1906.08125. 5, 053307. 16 p, May 2019, 16 pages, 12 figures. DOI: 10.1103/PhysRevE.101.053307. [Online]. Available: <http://cds.cern.ch/record/2720240>.
- [62] F. M. Charbonnier, R. W. Strayer, L. W. Swanson, and E. E. Martin, “Nottingham effect in field and $T - F$ emission: heating and cooling domains, and inversion temperature,” *Phys. Rev. Lett.*, vol. 13, pp. 397–401, 13 Sep. 1964. DOI: 10.1103/PhysRevLett.13.397. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.13.397>.
- [63] V. A. Dolgashev and S. G. Tantawi, “Simulations of currents in X-band accelerator structures using 2D and 3D particle-in-cell code,” in *2001 Particle Accelerator Conference (PAC2001)*.
- [64] W. Wunsch, “High gradient breakdown in normal conducting RF cavities,” in *8th European Particle Accelerator Conference (EPAC 2002)*, Jun. 2002, pp. 134–138.
- [65] K. N. Sjøbæk, “Avoiding vacuum arcs in high gradient normal conducting RF structures,” PhD thesis, 2016.
- [66] P. B. Wilson, “A plasma model for RF breakdown in accelerator structures,” *eConf*, vol. C00082, A. Chao, Ed., TUE05, 2000. arXiv: [physics/0010002](https://arxiv.org/abs/physics/0010002).
- [67] I. Langmuir, “The effect of space charge and residual gases on thermionic currents in high vacuum,” *Phys. Rev.*, vol. 2, pp. 450–486, 6 Dec. 1913. DOI: 10.

- 1103/PhysRev.2.450. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.2.450>.
- [68] D. M. Pozar, *Microwave Engineering, 4th Edition*. John Wiley & Sons, Inc, 2011, ISBN: 978-0-470-63155-3.
- [69] M. A. Salam and Q. M. Rahman, *Fundamentals of Electrical Circuit Analysis*. Springer, Singapore, 2018, ISBN: 978-981-10-8623-6. DOI: 10.1007/978-981-10-8624-3.
- [70] K. F. Lee, *Principles of antenna theory*. 1984, ISBN: 978-047-19-0167-9.
- [71] *CST Microwave Studio*, 2018. [Online]. Available: <http://www.cst.com/products/cstmws>.
- [72] A. Grudiev, “T24_vg1.8_disk 12WNSDVG18 CLIC_G undamped 12 GHz,” Tech. Rep., 2008.
- [73] R. Peacock, private communication.
- [74] K. L. F. Bane and R. L. Gluckstern, “The transverse wakefield of a detuned X-band accelerator structure,” *Part. Accel.*, vol. 42, no. SLAC-PUB-5783, 123–169. 59 p, Mar. 1992. [Online]. Available: <https://cds.cern.ch/record/237311>.
- [75] T. G. Lucas, M. Volpi, P. J. Giansiracusa, R. P. Rassool, T. Argyropoulos, H. Bursali, N. Catalan-Lasheras, A. Grudiev, G. Mcmonagle, I. Syrachev, R. Wegner, B. Woolley, W. Wuensch, D. Esperante-Pereira, and M. J. Boland, “High power testing of a prototype clic structure: Td26cc r05 n3,” CERN, Geneva, Tech. Rep. CERN-ACC-2018-0030. CLIC-Note-1080, Sep. 2018. [Online]. Available: <http://cds.cern.ch/record/2642425>.
- [76] A. Vnuchenko, “TD24 R05 SiC N1 high power test in XBox3,” Tech. Rep., 2017.
- [77] A. Degiovanni, “Breakdown position analysis,” in *International Workshop on Breakdown Science and High Gradient Technology (HG2015)*, Beijing, China, 2015.
- [78] A. Vnuchenko, “High-gradient results at CERN I,” in *CLIC Workshop*, Geneva, Switzerland, 2018.

- [79] R. Zennaro *et al.*, “High Power Tests of a Prototype X-Band Accelerating Structure for CLIC,” in *8th International Particle Accelerator Conference*, 2017, TH-PIK097. DOI: 10.18429/JACoW-IPAC2017-THPIK097.
- [80] W. Wuensch, “CLIC: Overview of applications using high-gradient acceleration, from photon sources to medical physics,” in *CERN Academic Training Lecture Regular Programme*, 2018.
- [81] S. Benedetti, “High-gradient and high-efficiency linear accelerators for hadron therapy,” PhD thesis, IPHYS, Lausanne, 2018, p. 256. DOI: 10.5075/epfl-thesis-8246. [Online]. Available: <http://infoscience.epfl.ch/record/253063>.
- [82] A. Vnuchenko *et al.*, “High gradient performance of an S-band backward traveling wave accelerating structure for medical hadron therapy accelerators,” in *9th International Particle Accelerator Conference*, 2018, MOPML043. DOI: 10.18429/JACoW-IPAC2018-MOPML043.
- [83] A. Vnuchenko, “High-gradient issues in S-band RF acceleration structure for hadrontherapy accelerators and radio frequency quadrupoles,” PhD thesis, U. Valencia (main), 2020.
- [84] E. Kreyszig, *Advanced Engineering Mathematics, 10th Edition*. John Wiley & Sons, Inc, 2011, ISBN: 978-0-470-91361-1.
- [85] A. Saessalo, I. Profatilova, A. Kyritsakis, J. Paszkiewicz, S. Calatroni, W. Wuensch, and F. Djurabekova, “Classification of vacuum arc breakdowns in a pulsed DC system,” *Phys. Rev. Accel. Beams*, vol. 23, no. 2, p. 023101, 2020. DOI: 10.1103/PhysRevAccelBeams.23.023101. arXiv: 1911.05192 [physics.ins-det].
- [86] I. Profatilova and N.-M. Pienimaki, “Pulsed DC System News,” Tech. Rep., 2017.
- [87] E. R. Castro, private communication.

- [88] A. Saessalo, “Characteristics and statistics of breakdowns on copper electrodes,” in *7th International Workshop on Mechanisms of Vacuum Arcs (MeVArc 2018)*, San Juan, Puerto Rico, 2018.
- [89] I. Profatilova, private communication.
- [90] —, “CERN pulsed DC system: setup, measurements and results on breakdown,” in *Mini MeVArc meeting*, Tartu, Estonia, 2017.
- [91] —, “Recent progress at pulsed DC systems,” in *8th International Workshop on Mechanisms of Vacuum Arcs (MeVArc 2019)*, Padova, Italy, 2019.
- [92] “Design Description Document (DDD 5.3) Neutral Beam Heating & Current Drive (NB & CD) System,” Tech. Rep., 2003.
- [93] ITER. [Online]. Available: <http://www.iter.org>.
- [94] T. Patton, N. Pilan, P. Bettini, G. Chitarin, A. De Lorenzi, D. Marcuzzi, E. Sartori, M. Siragusa, and L. Trevisan, “MITICA intermediate electrostatic shield: Concept design, development, and first experimental tests identification,” *AIP Conference Proceedings*, vol. 2052, no. 1, p. 030 002, 2018. DOI: 10.1063/1.5083730. [Online]. Available: <https://aip.scitation.org/doi/abs/10.1063/1.5083730>.
- [95] L. Cranberg, “The initiation of electrical breakdown in vacuum,” *Journal of Applied Physics*, vol. 23, no. 5, pp. 518–522, 1952. DOI: 10.1063/1.1702243. [Online]. Available: <https://doi.org/10.1063/1.1702243>.
- [96] I. Slivkov, “Mechanism for electrical discharge in vacuum,” *Soviet Physics – Technical Physics 2*, 1957.
- [97] A. De Lorenzi, N. Pilan, A. Pesce, and E. Spada, “Validation progresses of the voltage holding prediction model at the high voltage Padova test facility HVPTF,” in *2012 25th International Symposium on Discharges and Electrical Insulation in Vacuum (ISDEIV)*, 2012, pp. 9–13.
- [98] F. Geli, “All eyes on the neutral beam source,” Jul. 2017. [Online]. Available: <https://www.iter.org/newsline/-/2764>.

- [99] X. Wu, H. Zha, J. Shi, H. Chen, T. Abe, T. Higo, and S. Matsumoto, “Design, fabrication, and high-gradient testing of X-band choke-mode damped structures,” *Phys. Rev. Accel. Beams*, vol. 22, p. 031001, 3 Mar. 2019. DOI: 10.1103/PhysRevAccelBeams.22.031001. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.031001>.
- [100] T. Shintake, “The choke mode cavity,” *Japanese Journal of Applied Physics*, vol. 31, no. Part 2, No. 11A, pp. L1567–L1570, Nov. 1992. DOI: 10.1143/jjap.31.11567. [Online]. Available: <https://doi.org/10.1143%2Fjjap.31.11567>.
- [101] J. S. Xiaowei Wu, “High-gradient performance of X-band Choke-mode damped accelerating structure,” in *International Workshop on Breakdown Science and High Gradient Technology (HG2018)*, Shanghai, China, 2018.
- [102] X. Wu, T. Abe, D. Cao, H. Chen, T. Higo, S. Matsumoto, J. Shi, and H. Zha, “High-gradient performance of X-band choke-mode structures,” in *9th International Particle Accelerator Conference*, 2018, THPAL154. DOI: 10.18429/JACoW-IPAC2018-THPAL154.
- [103] D. Dimiduk, M. Uchic, and T. Parthasarathy, “Size-affected single-slip behavior of pure nickel microcrystals,” *Acta Materialia*, vol. 53, no. 15, pp. 4065–4077, 2005, ISSN: 1359-6454. DOI: <https://doi.org/10.1016/j.actamat.2005.05.023>. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1359645405002880>.
- [104] E. Z. Engelberg, Y. Ashkenazy, and M. Assaf, “Stochastic model of breakdown nucleation under intense electric fields,” *Physical Review Letters*, vol. 120, no. 12, p. 124801, Mar. 2018. DOI: 10.1103/PhysRevLett.120.124801.
- [105] D. M. Dimiduk, C. Woodward, R. LeSar, and M. D. Uchic, “Scale-free intermittent flow in crystal plasticity,” *Science*, vol. 312, no. 5777, pp. 1188–1190, May 2006, ISSN: 0036-8075. DOI: 10.1126/SCIENCE.1123889.
- [106] T. Muranaka, “Field emission measurements with CERN DC-spark system,” in *CLIC Workshop*, 2015.

- [107] S. I. Lachmann, “Dark current fluctuations in pre-breakdown conditions,” in *6th International Workshop on Mechanisms of Vacuum Arcs (MeVArc 2017)*, Jerusalem, Israel, 2017.
- [108] Y. Ashkenazy, private communication.
- [109] M. Jacewicz, “Preliminary results from the cryogenic pulsed DC system,” in *8th International Workshop on Mechanisms of Vacuum Arcs (MeVArc 2019)*, Padova, Italy, 2019.
- [110] H. H. Braun, S. Döbert, L. Groening, I. H. Wilson, and W. Wuensch, “Status of CLIC high-gradient studies,” Jul. 2001. [Online]. Available: <http://cds.cern.ch/record/507274>.
- [111] I. Syratchev, “Status of high-power X-band RF systems development at CERN,” in *International Workshop on Future Linear Colliders 2014 (LCWS14)*, Belgrade, Serbia, 2014.
- [112] “Mixed Domain Oscilloscopes - MDO3000 Series Datasheet,” Tech. Rep. [Online]. Available: <https://www.tek.com/datasheet/mixed-domain-oscilloscopes>.
- [113] “NI 5761R User Guide and Specifications,” Tech. Rep. [Online]. Available: <https://www.ni.com/pdf/manuals/375509a.pdf>.
- [114] J. B. Johnson, “Thermal agitation of electricity in conductors,” *Physical Review*, vol. 32, no. 1, pp. 97–109, Jul. 1928. DOI: 10.1103/PhysRev.32.97.
- [115] “ZKL_2R5+ Coaxial amplifier,” Tech. Rep. [Online]. Available: <https://www.minicircuits.com/pdfs/ZKL-2R5.pdf>.
- [116] W. R. Bennett, “Spectra of quantized signals,” *Bell System Technical Journal*, vol. 27, no. 3, pp. 446–472, Jul. 1948. DOI: 10.1002/j.1538-7305.1948.tb01340.x.
- [117] “RF/Microwave Signal Generators - MG3690C,” Tech. Rep. [Online]. Available: <https://www.anritsu.com/en-us/test-measurement/products/mg3690c>.
- [118] C. L. Dolph, “A current distribution for broadside arrays which optimizes the relationship between beam width and side-lobe level,” in *Proceedings of the IRE*, vol. 34, 1946, pp. 335–348.

- [119] Julius O. Smith, *Spectral Audio Signal Processing*. 2011.
- [120] W. S. Cleveland, “Robust Locally Weighted Regression and Smoothing Scatterplots,” Tech. Rep., 1979, pp. 829–836.
- [121] S. Matsumoto, M. Akemoto, S. Fukuda, T. Higo, N. Kudoh, H. Matsushita, H. Nakajima, T. Shidara, K. Yokoyama, and M. Yoshida, “Nextef: The 100MW X-band Test Facility in KEK,” 3 p, 2008. [Online]. Available: <https://cds.cern.ch/record/1182725>.
- [122] “Choosing and using bypass capacitors,” Tech. Rep. [Online]. Available: <https://www.renesas.com/eu/en/doc/application-note/an1325.pdf>.
- [123] “KSIM,” Tech. Rep. [Online]. Available: <http://ksim.kemet.com/>.
- [124] “Surface Mount Multilayer Ceramic Chip Capacitors (SMD MLCCs) High Voltage C0G Dielectric, 500 – 10,000 VDC,” Tech. Rep. [Online]. Available: https://content.kemet.com/datasheets/KEM_C1009_C0G_HV_SMD.pdf.
- [125] W. H. Paul Horowitz, *The Art of Electronics (Third Edition)*. Cambridge University Press, 2015, p. 1220, ISBN: 978-0-521-80926-9.
- [126] *LTspice*, 2020. [Online]. Available: <https://www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html>.
- [127] “ESD7102, SZESD7102 ESD Protection Diodes,” Tech. Rep. [Online]. Available: <https://www.onsemi.com/pub/Collateral/ESD7102-D.PDF>.
- [128] G. Turin, “An introduction to matched filters,” *IRE Transactions on Information Theory*, vol. 6, no. 3, pp. 311–329, 1960.
- [129] R. R. Antti Marilainen, “Update on dynamic vacuum and neutral Cu measurement,” Tech. Rep., 2015.
- [130] C. E. Shannon, “Communication in the presence of noise,” *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, 1949.
- [131] E. Z. Engelberg, private communication.
- [132] J. Kingman, *Poisson Processes*, ser. Oxford Studies in Probability. Clarendon Press, 1992, ISBN: 9780191591242. [Online]. Available: <https://books.google.ch/books?id=VEiM-0twDHkC>.

- [133] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi, *Queueing Networks and Markov Chains*. John Wiley & Sons, Ltd, 2001, ISBN: 9780471200581. DOI: 10.1002/0471200581. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/0471200581>.
- [134] A. Descoeurdes, F. Djurabekova, and K. Nordlund, “DC breakdown experiments with cobalt electrodes,” CERN, Geneva, Tech. Rep. CERN-OPEN-2011-029. CLIC-Note-875, Jun. 2009. [Online]. Available: <http://cds.cern.ch/record/1355401>.
- [135] C. Adolphsen, S. Döbert, A. Grudiev, S. T. Heikkinen, L. Laurent, J. A. Rodríguez, I. V. Syratchev, M. Taborelli, and W. Wuensch, “High-Gradient Test of a Tungsten-Iris X-Band Accelerator Structure at NLCTA,” CERN, Geneva, Tech. Rep. CERN-OPEN-2006-072. CLIC-Note-699, Nov. 2006. [Online]. Available: <https://cds.cern.ch/record/999491>.