Toward Cascading Failure Mitigation in High Voltage Power System Capacitors

A thesis presented in fulfilment of the requirement for the degree of

Doctor of Philosophy

to the

Institute of Energy and Environment

Department of Electronic and Electrical Engineering

University of Strathclyde

Calum J. Mackinnon

28^{th} of February, 2022

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

Acknowledgements

An ancient truism states 'it takes a village to raise a child': a PhD is no exception to this sentiment, albeit as an incomparably fleeting task and even if a participant benefits from life experience beyond its object; as an individual effort this work might not strictly be collaborative, but I am lucky to have been able to undertake it amid a community of direct and indirect contributors at the University of Strathclyde and further afield, who I hope know who they are. I thank Professor Brian Stewart for his enthusiasm, guidance, support, tact, and patience throughout the course of this research project. Secondly, thanks go to you the reader for any investment you make of your time into this work, as its value would otherwise be significantly limited. University and industry colleagues including: Andrew Carlin; Alan Davenport; Dr. Martin Given; Prof. Scott McMeekin; Kevin Butter; Alastair Ferguson; David Monkhouse; and Dr. Antonios Tzimas contributed practical support, technical insights, and facilitated site visits. Dr. Vic Catterson, Dr. Michael Gilroy, Dr. Phil Bagley, Dr. Sotirios Terzis, Dr. Andrew Roscoe, Prof. Stephen McArthur, Dr. James Irvine, Prof. Graeme Burt, and many others have been inspirational to me. Their example, enthusiasm, and openness have been instrumental in encouraging my interest in a PhD. I thank my family and friends for their patience, support, and understanding. If a dissertation can be dedicated then this is to my family, and in memory of Jim McShane: his inquisitive nature inspired my interest in technology at an early age. This work builds on a wealth of academic and industrial effort, in recognition that literature is only a subset of the understanding, know-how, and appreciation for a subject matter on the whole.

Contents

Abstract

As electrical power networks adapt to new challenges, advances in high voltage direct current interconnection offer one means to reinforce alternating current networks with flexibility and control, accordingly improving diversity to become a present-day, viable alternative to network flexibility and energy storage measures. High voltage capacitors support these links and offer simple means of voltage support, harmonic filtering, and are inherent to established and emerging converter designs. Where research literature predominantly explores use of modern dielectrics in efforts toward improved capacitor technologies, but reveals little about: existing capacitor designs; associated failure modes or statistics; or avenues in monitoring or maintenance, simulation modelling equips engineers with an approach to pre-emptively anticipate probable incipient fault locations toward improving designs for systems yet to be commissioned. This Dissertation presents a high-voltage capacitor simulation model, before exploring two questions about these hermetically sealed, highly modular assets: where are incipient faults most likely to arise; and how can internal faults be externally located? Nonlinear voltage distributions are found within each and among connected units, induced through parasitic effects with housings supported at rack potential. Consequent implications are considered on: stresses within unit dielectrics, susceptibility to cascading failure, and an ability to locate internal faults. Corroboration of fault detection and location is additionally found possible using unit housing temperatures. A model is presented, developed to be scalable, configurable, and extensible, and made available for posterity. Opportunities in asset design, modelling, manufacture, and monitoring are proffered toward improvements not only in operational longevity, but in understanding and early awareness of incipient faults as they develop.

List of Publications

1. C. J. Mackinnon, B. G. Stewart, "Simulation of High Voltage Direct Current Filters", 53^{rd} *International Universities' Power Engineering Conference*, September 2018.

2. C. J. Mackinnon, B. G. Stewart, "A High Voltage Capacitor Element Model", 37^{th} *IEEE Electrical Insulation Conference*, June 2019.

3. C. J. Mackinnon, B. G. Stewart, "Thermal Profiles of High-voltage Capacitor Units", *Annual Conference on Electrical Insulation and Dielectric Phenomena*, October 2019.

4. C. J. Mackinnon, B. G. Stewart, "Regions of Electrical Stress in High Voltage Capacitor Units", *COMSOL Conference Europe*, September 2019.

5. C. J. Mackinnon, B. G. Stewart, "A Capacitor Bank Simulation Model", 38^{th} *IEEE Electrical Insulation Conference*, June 2020.

List of Conferences

The conferences attended during the course of this research programme are listed in Table 1, which uses the symbol * to denote those at which presentations and publications were given.

Table 1: Conferences

Conference	Year	Month	Dates	Location
LCNI	2016	October	11 - 13	Manchester, UK
UHVNet	2017	January	18 - 19	Glasgow, UK
AllEnergy	2017	May	10 - 11	Glasgow, UK
CIRED	2017	June	12 - 15	Glasgow, UK
UHVNet*	2018	January	15 - 16	Winchester, UK
Futurewind	2018	March	22	Glasgow, UK
AllEnergy	2018	May	2 - 3	Glasgow, UK
UPEC*	2018	September	4 - 7	Glasgow, UK
SmartFuturES*	2018	September	19	Glasgow, UK
UHVNet*	2019	January	15 - 16	Manchester, UK
EIC*	2019	June	16 - 20	Calgary, Canada
COMSOL*	2019	September	24 - 26	Cambridge, UK
CEIDP*	2019	October	20 - 23	Washington, USA
UHVNet*	2020	January	15 - 16	Glasgow, UK
EIC*	2020	June	7 - 11	Online

Chapter 1

Introduction

This Chapter introduces the premise, aims, approaches[1], and structure of this work in Section 1.1, before in more detail introducing its wider subject domain in Section 1.2 for context.

1.1 Premise

In accordance with one of Shigeo Shingo's principles (for value to be drawn according to a basis of need), an open-minded approach to identifying a domain and direction for this research project has been adopted. High voltage direct current (HVDC) is an area of interest articulated by industrial partners of the centre for doctoral training by which this project has been funded.

While some electrical transmission assets (such as transformers) are sufficiently well established and numerous as to have monitoring techniques and statistics collated on reliability, HVDC remains a relatively recent set of technologies in terms of diagnostics and prognostics. Thus, while a prospective contributor might begin a research project by exploring literature to identify a research niche, it can be daunting to instead stumble upon a metaphorical crevasse.

Principles of reliability centred maintenance invite manufacturers to include measurement points and design countermeasures in anticipation of eventual faults within each assets' design. Where vulnerabilities are understood, this knowledge informs iterative design refinements to aid detectability, mitigate consequences, and reduce the likelihood of failure modes toward improved reliability and longevity from a 'whole lifecycle' perspective. As increasing numbers of HVDC systems are set to be planned and deployed to help decarbonise the global economy, information on failure modes of assets critical to these links could prove timely. Rather than await failure or reliability data, this research adopts existing design information as a basis to anticipate failure modes by returning to the fundamentals of component physics and operation.

1.1.1 Hypothesis and Aim

This project hypothesises that:

> as yet unidentified signals betray failure modes in HVDC converter stations.

This project consequently aims to identify those *signals which arise due to faults* in HVDC converter stations, and appropriate associated measurands. It therefore explores both electrical and thermal effects associated with one particular asset critical to HVDC function: capacitors.

[1] Material given in this introduction will, naturally, be treated more comprehensively in subsequent Chapters.

1.1.2 Approaches

Motivation

An increasing use of power electronics means capacitors are becoming more common as power systems evolve to meet modern challenges; unit designs also vary as new topologies – such as modular multilevel converter (MMC) – and (composite and dry dielectric) materials emerge.

High-voltage capacitors are a suitable subject of study given the large numbers deployed at converter stations for conventional HVDC. Although some online monitoring solutions exist, capacitor banks' structure makes them opportune assets for unit-specific condition monitoring, facilitated by ubiquitous information technology, in addition to the benefits of informed design. This is particularly given units' vulnerability to signal artefacts when connected at high-voltage in converter stations and their increasing criticality in alternating current (AC) power systems.

Research Landscape

Internally fused and fuseless unit designs offer no visual indication of internal dielectric breakdown, and any existing element failures can lead to further susceptibility (faults 'cascade'). It is accordingly imperative to detect and protect against capacitor unit faults, which is typically achieved using counters and imbalance protection; recent advances can identify the string of units on which each fault occurs for efficient maintenance by narrowing the search space for faulted units. Such techniques use existing sensors, but remain limited by them. In future, distributable measurements could: 1. detect *incipient* faults; 2. significantly reduce the search for detected faults; and 3. take account of operational temperatures which are critical to capacitor longevity (every 8 to 10 °C increase in core temperature approximately halves expected life [1]). Improved understanding alone can identify vulnerabilities to inform maintenance and design.

A Model-based Approach

Modelling discussed in this Dissertation focusses on electrical and thermal effects, as capacitors are electrical assets with dielectrics primary to their function, and where operating temperatures are known to significantly influence unit health and longevity. Simulation modelling allows:

1. estimations to be made on (a) *dynamic* behaviour of assets which can only be safely measured offline, and (b) the *internal* physics of hermetically sealed units;
2. models to be (a) adapted and developed in tandem with understanding of internal designs, or conversely to be (b) configurable so as to accommodate design uncertainty and emulate a variety of unit types, and to be (c) agnostic of any one supplier;
3. data to be generated with (a) a minimum of physical waste, (b) fewer safety concerns relative to lab experiments (with (c) resilience to limits on practical work); and allows
4. support for repeatability and further work through reference to a digital model file.

Disadvantages of modelling include: 1. the time taken to build (familiarity with an environment and) a model of sufficient detail; and 2. an inevitable simplicity inherent to any model, which typically omits noise, manufacturing defects, and ambient conditions (physical assets will always be more nuanced in reality). Nonetheless, by emulating the effects of an applied electrical signal, associated electric field stresses, dielectric heating, and heat from incipient faults, modelling allows insights into asset behaviour, vulnerabilities, and failure modes to be obtained.

Dissertation Structure

This Dissertation is structured with seven Chapters, each contributing as follows.

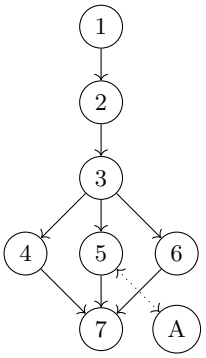

Figure 1.1: Chapter dependencies

1. The remainder of this **Chapter 1 introduces the subject area**, including drivers for change in electrical power systems, to identify technologies for which monitoring could be opportune. HVDC is a large subject, but as these systems advance so too does a gulf between expectations for reliable operation and means of assessing the failure modes which will inevitably, eventually arise in forms potentially unique to emerging technologies.

2. **Chapter 2 explores suitable research directions** and returns to the fundamentals of capacitor unit design before offering an insight into approaches used for fault detection and location in high-voltage power system capacitors, as discussed in current literature.

3. **Chapter 3 presents a configurable high voltage capacitor unit model** to aid further study in this area. This is made available online for the study of different capacitor designs, and has been designed to be configurable to that end. Good practise techniques are discussed and improvements which could be made to the simulation environment to support these techniques have been highlighted to COMSOL® directly.

4. **Chapter 4 explores electrical effects** in pursuit of probable fault locations and possible diagnostics. Unit housing is shown to provoke parasitic electric field effects, which invoke a nonlinear electrical distribution within each unit. Solvability of a corresponding electrical circuit and the influence of internal faults on terminal characteristics are discussed, and regions of high electrical field stress associated with aspects of a capacitor unit design are considered. This is found to extend throughout multiple units on a rack.

5. **Chapter 5 considers thermal behaviour** and how this is heavily influenced by element winding alignment within each unit, since foils both contain and direct thermal energy. This can be used to corroborate electrical diagnostics, which is an area of significant importance to capacitor asset longevity but with limited consideration thus far.

6. **Chapter 6 covers preparation of a physical prototype capacitor**. It outlines the design, build, and measurement of this scale replica unit in a laboratory setting, and finds practical measurements to accord with the principles of operation found by simulation.

7. **Chapter 7 collates conclusions of the dissertation** from discussion within prior Chapters. It considers learning outcomes, and possible avenues and opportunities for further research to support future work.

The author hopes that by offering background understanding, suggestions for further investigation, and a baseline for modelling, future exploration can be encouraged in this area at an intersection of energy systems, asset management, information technology, and sustainability. Chapter dependencies are represented graphically in Figure 1.1 to support selective reading.

Conventions

References to this Dissertation, the Chapters, Sections, Figures, and Tables it contains are capitalised. COMSOL requests the ® symbol accompanies references to the software environment. References to components of the modelling environment use verbatim typesetting (and in the index are capitalised) to delineate their use as a reference to terminology specific to COMSOL® rather than more general use of the term: all capacitors have geometry, but a computer model has a `Geometry` (as a node within the model tree). An algebraic convention to encapsulate newly substituted terms in brackets is followed to reduce ambiguity in algebraic working.

1.2 Context

Electrical networks underpin modern lifestyles and serve a critical (if often unnoticed) role in advanced economies; they are a cornerstone of advances in sustainable development goal 7 [2], and the nature of these networks is changing in three notable ways.

1. Power systems are becoming increasingly sustainable. Traditionally, power systems have operated by conveying power from synchronous generators, which are controllable, few in number, and centrally located on a network, but which tend to be carbon-intensive, presently nearing the end of their useful life, or are simply challenged by economic competition with modern alternatives. In efforts to reduce greenhouse gas emissions, such forms of centrally-owned historical generation are increasingly becoming superfluous to geographically disparate, more numerous, and less controllable *renewable sources*, which are necessarily far from centres of demand, connected at all voltage levels, and connected by converter interfaces which do not offer equivalent inertia to support network stability.

2. Networks should be cost-effective. These systems are expected to (a) operate at *minimal cost* to consumers [3], (b) ensure a competitive market for trading energy and its associated services, in part to (c) offer competition to other available forms of domestic energy, and to (d) further a form of social justice to permit everyone access to the energy we need.

3. Energy must be dependable. High availability ('uptime') is desired for electrical supplies to be relied upon as a technology. Not only are renewable sources less reliable, but the manner in which a network must operate to connect large proportions of converter-based renewable generation can leave it vulnerable to frequency and voltage deviations in response to sudden changes, whether faults or unanticipated variation in intermittent generation output. Thus, transitioning to a modern generation portfolio poses a risk to *security* (the expectation for energy to be *available* when needed).

1.2.1 Energy Trilemma

These three aspects of a changing energy landscape are captured by the 'energy trilemma' index [4], which attempts to outline the scale of desired change. As illustrated in Figure 1.2, research endeavour in this field revolves around these aims, and broadly conforms to three approaches to their resolution: intermittency (using bulk or distributed storage mechanisms, or shifting intermittent loads, to better match instantaneous supply and demand), interoperability (pursuit of the 'smart grid' paradigm and network flexibility), and interconnection (increasing natural supply and demand diversity through physical integration of established systems).

'Smart grid' paradigms and similar solutions designed to achieve flexibility (such as active network management and dynamic thermal ratings) unlock these capabilities with information and communication technology (ICT) stacks. Alongside ICT, power electronics offers a means of introducing controllability by decomposing and reconstituting alternating current (AC) waveforms, to thus be used together with ICT to provide the interoperability, intermittency, and interconnection needed to address the changing nature of modern networks.

Another aspect of this transition is ownership: large, established market participants which own and control conventional generation are giving way to a broader array of owners concordant with the distributable nature of renewables, with both economic and operational implications.

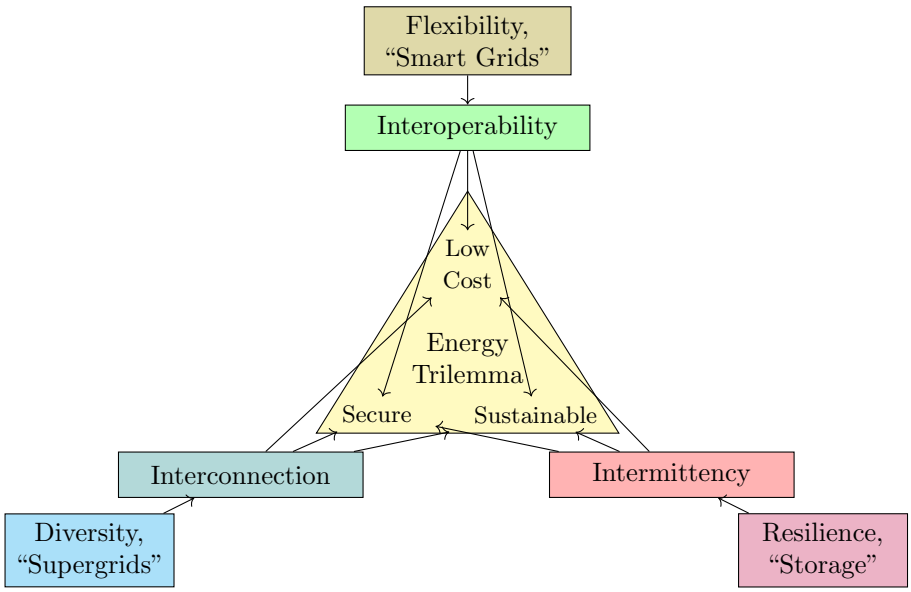

Figure 1.2: A representation of the power systems research landscape

Consumer Engagement

Over recent years, consumer engagement has been suggested as a fourth aim for electrical networks [5]. While proponents of this argument suggest the use of smart metering, awareness of carbon footprint, and the importance of cultivating responsible energy use such as through home insulation, choice of tariff, and even choice of car, these arguments are often made by those with an established interest in electrical technology or in the industry. These efforts are well intentioned, but overlook a principle of ubiquitous technology first outlined in 1991 [6]:

> "the most profound technologies are those that dissappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it."

When viewed as a ubiquitous technology, the electrical power system *should* be sufficiently pervasive, dependable, cheap, and conscionable as to be ignorable. Thus, while energy trilemma aims accord with one another, they better lend themselves toward the aims of a *ubiquitous* technology rather than toward consumer engagement per se. However, the array of advances the energy trilemma index demands for power systems globally might nevertheless draw technical interest and could increase engagement of consumers as a by-product of these core ambitions.

1.3 Interconnection

While any one of these three approaches (storage, 'supergrids', or 'smart grids') taken in extreme could suffice to address the challenges posed by the energy trilemma, in practise a combination is necessary. Energy storage is viewed as a hallmark need for electrical energy systems given its ability to offset (and hence match) instantaneous demand and supply. However, reinforcing and connecting existing networks can create the diversity between generation and load necessary to improve the scale at which any balancing must occur. Similarly, while 'smart grids' can make use of distributable ICT to co-ordinate supply and demand around constraints, interconnection offers a comparatively mature way of alleviating constraints to begin with, through diversity.

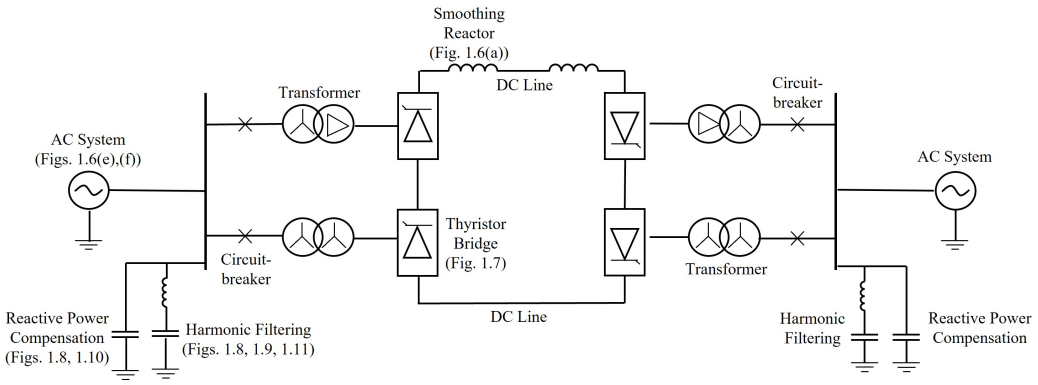

Figure 1.3: Line diagram of an LCC HVDC link [8]

In short, interconnection and preparations toward a 'pan-European' network [7] is a modern approach to conventional reinforcement: a natural next step to the historic development of network infrastructure through increasing voltage levels. It solves key aims of the energy trilemma predominantly by improving the diversity of supply and demand on a network it strengthens:

1. costs are shared (and therefore reduced) over a wider area between newly connected networks to unlock economics of scale;

2. network security is improved as a result of controllability, and since a better connected network can be made more resilient to faults; and interconnection

3. permits higher proportions of renewables by connecting different geographies, allowing each area to make the most of its natural resources.

Accordingly, the European network of transmission system operators for electricity (ENTSO-E) calls for 50 GW of European cross-border capacity by 2030, beyond projects already planned [7].

1.3.1 High-Voltage Direct Current

As an abstract concept, interconnection can be realised in practise through established AC technology, but to achieve greater capacities (where networks are already mature as is broadly true in developed economies) it is becoming increasingly appropriate to use high voltage direct current (HVDC) systems. HVDC links use power electronic converters, which although prone to losses allow efficiency gains to be made overall. As such, power electronic conversion provides a modern means of conventional power system reinforcement. A line diagram of an line commutated converter (LCC) HVDC system is shown in Figure 1.3.

1.3.2 Converter Topologies

Developments in semiconductor technology have enabled new forms of power electronics, which in turn have allowed high-voltage applications. In the 1980s mercury arc valves were used for the first HVDC link to Gotland, Sweden. Since then the underlying semiconductor technology has advanced further to respectively use thyristors and insulated gate bipolar transistors (IGBTs) for conventional line commutated converter (LCC) and voltage source converter (VSC) topologies. The main constraint with advances in these core semiconductors is heat extraction,

where modern switching elements are typically cooled using deionised water [9]. An ability to control the firing angle of switching elements (e.g. thyristors) offers converters a significant advantage over conventional AC networks, as these links are *controllable*.

Advantages of LCC HVDC

Specific advantages of LCC HVDC over conventional AC connections are listed here (numbers 1 to 6), before VSC technology is discussed as a way to expand these advantages with (7 to 8).

1. Primarily, advantages are in the *capacity* and *distance* that HVDC solutions cost-effectively offer. Although conversions between AC and direct current (DC) power are lossy, DC electrical transmission sees savings in the need for fewer conductors with more efficient use of each cross-sectional area, and by operating at higher voltages, is liberated from some losses for the length of the line (associated with capacitive components of its transfer function). Thus, for a given cost, DC links are capable of greater distances and capacities.

2. In connecting renewables, the principle for using converter-based power electronic interfaces is that generation (turbine) speeds are decoupled from that of the grid. This same principle of *decoupling frequency* in connections allows integration of systems which might otherwise remain independent: asynchronous networks can be connected via an intermediate DC stage to form 'supergrids' on continental scales.

3. The *compact footprint* of converter stations relative to the electrical power capacities they convey is another strong advantage, in that it allows reduced converter station footprints which is an important factor in connecting: offshore installations; where land is expensive; and to form high-capacity connections to urban areas [8].

4. Reduced asset footprints are also true of the associated transmission assets, due to HVDC links' need for *fewer conductors*, leading to further cost savings over long distance connections. Additionally, this capability allows existing AC circuits to be converted to higher-capacity DC connections, should converter stations be positioned at either end of existing conductors.

5. HVDC allows connections to be made *subsea* which would otherwise be limited by the capacitive nature of a saltwater environment. This is important for the viability of offshore wind at distances far from the coast, and is crucial to providing an option to circumnavigate the need for overhead lines and associated wayleaves for connections by land.

6. HVDC can even be used as a form of network *resilience* to weather: in Canada, in response to ice loading on overhead lines, a bipole link can be isolated, connected together at one end, and used to pass direct current for de-icing, before being reinstated to the network.

Advantages of VSC HVDC

As semiconductor technologies have advanced, newer HVDC topologies have been developed: VSC options allow for a different set of strengths and weaknesses but in general are more technically capable at greater cost. In addition to firing angle (α), IGBTs are self-commutating and allow VSC converters to control commutation angle (μ) (p123, [8]). Combined with IGBTs' high-frequency switching capabilities, controllable commutation angles represent a significant

advantage by permitting use of pulse width modulation: a converter can be switched to mimic an AC waveform which extends control over active power (α) with independent control of reactive power (μ). Thus, it is possible with VSC to connect to low-voltage and completely passive networks, meaning this technology can support a greater variety of network locations, renewable generation[2], multiterminal HVDC, and blackstart.

7. VSC technologies offer improved connectivity through the possibility of *multiterminal* HVDC systems (MTDC), which can allow a resource (such as large offshore wind arrays proposed for the North Sea [8]) to be shared by networks in multiple countries.

8. A main advantage is *controllability* of power flows. Back-to-back HVDC links can be used to mitigate risk of frequency events spreading and thus can provide 'firewalls' within a network. Converters can be set to transfer only the power required, and thereby: stabilise networks in terms of active and reactive power; to dampen subsynchronous modes; and can be used to mitigate the short-circuit current contribution of connected generators.

Drawbacks

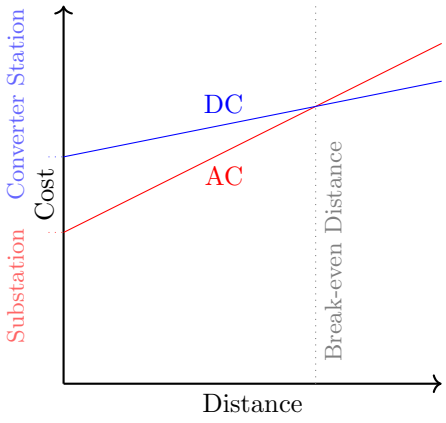

Figure 1.4: Cost-distance relationship for AC and DC connections (p320, [10])

The drawbacks of HVDC systems are their higher costs (when considered independently of capacity and distance – Figure 1.4[3]) relative to conventional AC connections, and possibly limitations with supply chains faced with global demand. Although conversion to DC facilitates subsea connections, such routes are necessarily inaccessible with significant associated costs for installation (and maintenance should a fault arise).

Increasing proportions of converter-connected generation presents challenges for network operation, particularly in terms of inertia. This is also true where demand is reduced on networks designed to be loaded: parts of the network risk becoming loosely controlled.

LCC topologies are vulnerable to disturbances and signal artefacts on the AC network, so require electrically strong network locations to connect to, in the interests of preventing commutation failure. VSC technologies, on the other hand, can connect to entirely passive networks and are able to cope with AC signal artefacts, but are instead vulnerable to DC line faults by virtue of large DC capacitors shunt connected at each station.

Moreover, given their nature as a relatively recent technological development, HVDC systems present a challenge to operators for their operation and maintenance. That new designs and emerging topologies are used for new developments, there is also a risk that the technology could carry associated risks to asset health over the course of its lifetime. For instance, the Nelson River project experienced a series of coincidental transformer alarms indicative of challenges associated with an early system design.

[2]This simple progression from (firing angle controlled) thyristors to (firing and commutation angle controlled) IGBTs is the means by which opportunities in renewables and offshore technologies are unlocked.

[3]Break-even distances are stated as around 800 km for overhead line and 50 km for subsea connections [8], [10].

Voltage-source Converters

VSC systems are a more modern and capable form of HVDC first introduced in 1997. Their main advantage over thyristor-based LCC topologies is *greater controllability* by using IGBTs, which allow control of both firing and blocking. Combined with higher switching frequencies, this allows more sophisticated behaviour by permitting techniques such as pulse width modulation and third harmonic injection to reduce characteristic harmonics and "synthesise a fully controlled AC voltage, which enables precise control of active and reactive powers" [8].

One of the most important strengths of VSC technology is that these systems are liberated from a need to connect only to strong networks. Instead, they have an ability to connect weak systems and even to network locations which are entirely passive (thus making such links suitable for blackstart capability and in connecting generation).

This high level of controllability means that characteristic harmonics associated with the conversion between AC and DC can be deliberately mitigated. And as each VSC converter can control active and reactive power *independently* of one another, each converter station can act as a static compensator (STATCOM) to support the local AC network.

A reduction in characteristic harmonics of VSC systems as well as removing a need for voltage support at the point of connection means that VSC systems require *less filtering*, which leads to the possibility of further reductions in the required footprint of VSC converter stations. Such filtering arrangements are also simpler, as they can be permanently connected to the converter and need not be switched in and out according to load profiles (page 136, [8]), thereby further removing concern for longevity of switchgear designated in LCC arrangements for this task in meeting daily load profile variation.

In LCC systems, STATCOMs bridge any difference between the reactive power required and that supplied by the filter banks connected at a given point in time. As VSC topologies can instead control this reactive power, the need for a separate supporting STATCOM diminishes.

VSC converters permit *smooth reversal* of power flow, such that they can participate more fluidly within modern networks. However, appropriate choice of dielectric must be made for cables in order for this capability to be realised in practise. This is an important ability since VSC links can transfer just the power required at any given point of time, and the ability to reverse the direction of power flow allows such systems to play a role in accommodating daily and seasonal variation in network use, and even should network power flows change over time as new connections are made.

VSC systems can be *multi-terminal*, which means they could be used to connect a common resource to multiple possible beneficiaries, such as in the case of proposals for islands of offshore wind connected to a number of surrounding countries.

Modular Multilevel Converters

Modular multilevel converter (MMC) topologies [11] are designed around half-bridge 'submodule' switching elements [12] which lead to more complex and less technologically mature designs, but offer much greater controllability. This in turn allows harmonic output to be minimised, significantly reducing the need for filtering and reactive power support and thus facilitating even further reduced station footprints.

1.3.3 Early Experience of HVDC: the Nelson River Transformers

One of the first HVDC systems brought online experienced seemingly coincidental failures [13], which suggests something intrinsic to this early link had not been fully understood at the time. An HVDC system was commissioned to connect Winnipeg in Manitoba, Canada to a hydro power station on the Nelson River, and was brought online in 1979. This link was commissioned with mercury arc valves (it would later be upgraded with Alstom's H300 thyristors in 1990 [9] and then with further thyristors in 1993 [13]), was expanded in its capacity in 1983 and 1984, and had new generation connected between 1990 and 1992. It has one ⅄ and one Δ transformer for a 12-pulse valve group on each of its two bipoles, which, when duplicated at both rectifier and inverter, sums to a total of 8 converter transformers first put into service in 1978 [14].

During these transformers' operation, a series of failures plagued the project (in particular where ethylene gassing occurred within the oil) as listed in Table 1.1 [14][4]. Two transformers were returned to their manufacturer in 1978, and others later had remedial modifications.

Table 1.1: Nelson River HVDC converter transformer (Tx.) failures [14]

Month	Tx. Location	Protection	Fault Location
Feb 1994	Inverter (⅄)	Gas Alarm → Manual Block	Valve winding
May 1998	Rectifier (⅄)	Ground Differential Relay	Line winding
Aug 1998	Rectifier (Δ)	Pressure Relief	Valve winding
Nov 1999	Inverter (Δ)	Pressure Relief	Valve winding
Feb 2000	Inverter (⅄)	Sudden Pressure Alarm	Valve winding
Sept 2000	Rectifier (Δ)	Gas Alarm → Pressure Relief	Valve winding

These provoke two noteworthy considerations. Firstly, the similar timescales of these failures is interesting in the context of their approximately twenty-year service at time of their respective failures. Second is the similar nature of the failures in predominantly being associated with the first few winding turns, particularly on valve windings [13], [14]. Importantly, on inspection, there is strong evidence of breakdown *between adjacent winding turns* on the first handful of valve winding turns where paper has in some cases been seen to be "virtually incinerated" [14].

Secondly, while alarms associated with these failures were sudden and unanticipated (even where online gas monitoring had been installed), all transformers analysed had signs of prior aging or of having suffered gradual deterioration over a sustained duration. With the exclusion of the September 2000 failure these were respectively: • returned to factory in 1979 with concern around metal particulates; • similarly returned in 1979 due to ethylene gassing; • returned in 1991 due to ethylene gassing; • evidence on post-mortem of reduced insulation resistance on winding turns; and • with paper wrapping and enamel destroyed on the first turns on all legs.

Possible causes of this breakdown mechanism are harmonics (characteristic or otherwise) introduced by conversion between AC and DC, and the potential for voltage impulses [15]. Since the reactance of an inductive winding $X_L = j\omega L$, where $\omega = 2\pi f$, is proportional to the signal frequency f, transients at high frequencies see each transformer winding as a "large reactance" [15], such that high frequency signals reach only the first few turns. Accordingly, this leads to voltage discrepancies between adjacent turns and puts stress on intermediate dielectrics.

Conventional converter station topologies feature transformers as the *first* assets seen on the AC side of converter valves, *next to which is filtering equipment* designed to attenuate harmonic signals before an interface to an AC system [9].

[4]The author offers his condolences to those privileged to have known the late first author of these papers.

Nowadays, converter transformers are designed to be more robust to such transients, by increasing the strength of insulation for at least the first winding turns seen by such transient signal artefacts (harmonics and commutation impulses) where "the thickness of the paper tape layer varies with the local stress" (p. 589 [16]). However, it is not yet apparent from the literature that similar precautions are adopted for *other* assets in the vicinity of a converter station valve hall beyond high-criticality converter transformers. As filter banks are the target destination for a converter station's characteristic harmonics, their assets are those electrically next subjected to high-frequency signal components from transformers. Should such evidently stressful signal artefacts be dissipated by local impedances and dielectrics, this highlights these assets' suitability for further investigation with a view to unlocking potential improvements in their design, possibly analogous to those now adopted as a common practise for transformers.

1.3.4 Criticality, Reliability, and Availability

Possible options for converter switching controllability can additionally enhance network capacity not only by interconnection, but through power factor correction, series compensation, and connection of renewable sources dispersed across a network in terms of both connection voltage and location. Thus, the introduction of HVDC links to a network has potential to unlock benefits and offer *combined capabilities* greater than the sum of their parts.

Since these technologies are responsible for greater capacities, they become more critical to network operation and as a component of power electronic interfaces and standalone power factor correction, shunt capacitor banks correspondingly become "increasingly important" for reliable power systems [17], across voltage levels. Since interconnection projects are typically for large capacities, their availability is increasingly critical [18] for: • network diversity; • providing access to bulk storage; and • allowing placement and connection of geographically disparate renewable resources. As this trend toward greater penetrations of renewable generation continues, and as network inertia wanes, power systems will become increasingly reliant on converter-connected systems. Hence, to maximise benefits to networks and encourage development of similar projects [19], availability is paramount for converter-based systems.

Risks

Despite an increasing prevalence in electrical power infrastructure due to a growth of HVDC and other power electronic applications (such as reactive power compensation, for instance), and despite accordingly becoming more critical to network reliability, at present fault location and monitoring for high voltage capacitors has received "very little" attention in the literature [20].

Conventional LCC converter topologies offer high capacities, but introduce characteristic and non-characteristic harmonics, and where these systems are deployed, there is a consequent risk posed to existing and potentially aged assets. Moreover, since many of these systems are recent, bespoke developments and employ technologies which continue to be refined, reliability information can be limited, commercially sensitive, and in any case is not readily available.

Rather than await statistical information on asset reliability, it is preferable to anticipate degradation and failure of certain assets. Representative models can simulate stresses and breakdown processes to aid understanding of how assets degrade over time. Simulation can also be used to investigate the influence of assets' operational environment and dynamic characteristics, which could each deviate from measurements acquired offline under static conditions.

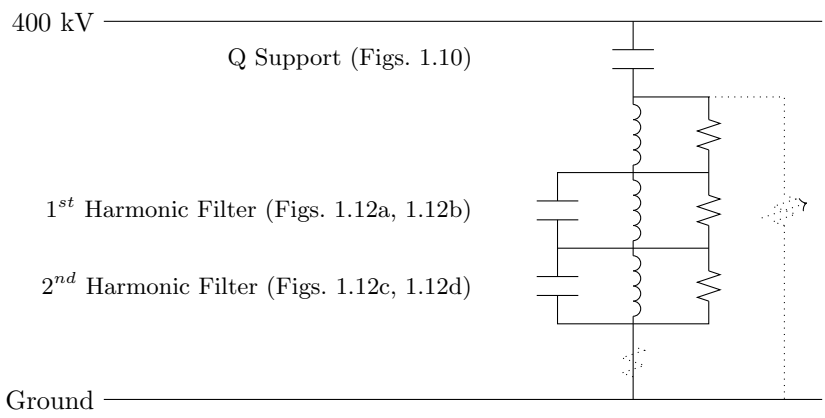

Figure 1.5: A filter bay line diagram (p108, [9])

Harmonic Filtering and Voltage Support

Characteristic harmonics result from conversion between AC and DC power, by nature of the converter (valve) topology: for example, on a 12 pulse topology the AC side would have a filter for the 11^{th} and 13^{th} harmonics with the possibility of a 12^{th} harmonic filter on the DC side. Figure 1.5[5] shows filters which attempt to keep harmonics within limits and protect other transmission assets from these high frequencies. To accommodate multiple harmonics, and to offer a degree of redundancy, there are ordinarily multiple such filter bays connected in parallel.

Voltage support at a point of connection is similarly important. LCC systems require a strong network location to mitigate risk of commutation failure, so are often sited electrically near synchronous generation (Figure 1.6e) and reactive power ('Q') support is critical.

Filtering assets are therefore central to safeguarding the electrical conditions necessary for converter technologies to contribute to a power network without degrading established (and, depending on where a connection is made, potentially critical) electrical infrastructure.

High voltage capacitor banks are one such asset, which perform a dual function within LCC converter stations: 1. as part of tuned filtering arrangements; and 2. in providing reactive power to support an interface between a converter station and the AC network it connects to.

1.3.5 Design Informed by Reliability Centred Maintenance

Although modern converter topologies are increasingly sought, conventional alternatives remain commercially viable for high-capacity connections to strong networks. Hence, it is possible that any design refinements could be considered for systems yet to be commissioned [19]. Moreover, HVDC and other converter interfaces have a large role to play in providing an element of controllability to existing networks under emerging use conditions.

Being able to build sensors into equipment during manufacture offers the ability to apply unit testing to equipment, for use throughout its development during the integration stages, acceptance testing, and then throughout its operational life. Being able to predict the remaining useful life of assets allows for longer lead times on replacement assets, in turn meaning that assets can be manufactured from further afield, aiding globalisation and economics of scale.

[5]Different grounding arrangements are reported in the literature, where some filter bays are grounded via a resistance and others use variable resistances as in [9].

1.4 Commercial Considerations

The Western Link HVDC link is a 420 km long, 2.2 GW 600 kV[6] bipole LCC connection between Hunterston in Ayrshire, Scotland, and Deeside, North Wales [21]. This link was motivated by a need to improve network capacity so as to accommodate a growth in renewable energy in Scotland requiring transmission to demand centres in England and Wales.

While HVDC projects such as the Western Link [22] can be controlled to actively attenuate subsynchronous modes[7] and hence improve capacity of existing network assets, control algorithms are not typically shared, meaning operators are unable to consider converter control as an option. Instead, operators prefer to isolate a link should oscillations pose problems [23].

Scottish Power Energy Networks (SPEN) even suggested that HVDC links can exacerbate the presence of subsynchronous modes. Such links introduce large capacitances to the network, so there is potential to significantly alter its resonant frequency and thereby increase the likelihood that an existing subsynchronous resonance could cause problems and damage assets already connected and in safe operation for years. There is scope for improved understanding among industry participants, which in this case could mean greater openness about how technologies work such that, not only could problems in the introduction of new technologies be mitigated, but benefits to the long-term operation of the network could be unlocked.

1.4.1 Scottish Power Energy Networks' Converter Station

Figures 1.6, 1.7, 1.8, 1.9, 1.10, and 1.12 show a selection of photographs[8] from a visit to Hunterston converter station, kindly facilitated by industry colleagues as this research project began. Although these photographs relate to one specific converter station, they are shown in this section to provide a graphical introduction to the type of equipment typical for conventional converter stations generally. From the capacitor bank arrangements in Figure 1.10, and the system voltage, typical voltages across a individual unit can be inferred to approximately 10 kV.

To illustrate, Figure 1.10b highlights jumper connections between series-connected units in yellow, and Figure 1.11 outlines in more detail the connections evident in Figures 1.10c and 1.10d for a large capacitor bank, including connections to a current transformer (CT), which divides this two-sided bank into an 'H-bridge' arrangement consisting of four series '*strings*' of units.

Points of Discussion

SPEN share ownership of the Western Link with National Grid, but the contractor (Siemens) is responsible for commissioning and operation of the link for the first five years of service. An early stage of this project involved meeting with Alastair Ferguson and Kevin Butter from SPEN in an effort to gauge areas of concern SPEN might have in anticipation of assuming responsibility for operation of this LCC connection and its converter station. Some of the discussion points raised are listed as follows.

[6] This project was the first subsea connection at this voltage [21].

[7] Subsynchronous resonances (p198, [10]) are oscillations below nominal network frequency as a consequence of series capacitive and inductive effects on long AC network connections, and can effectively limit the maximum capacities practical for critical network links: the 'B6' boundary between Scotland and England in Great Britain is an example. As such oscillations can pose a risk to the mechanical integrity of network-critical synchronous generators, they must be evaluated and network operations accordingly limited to protect valuable generation.

[8] Access and photograph permissions were obtained courtesy of SPEN but copyright remains with the author.

(a) DC reactor

(b) Converter valves

(c) Heat extraction

(d) Spare converter transformer

(e) Hunterston B nuclear facility

(f) AC network connection

Figure 1.6: Main components of Hunterston HVDC converter station

Figure 1.7: Valve hall

Figure 1.8: Filter hall

(a) Filter bay (looking south)

(b) View along filter bays (looking westward)

(c) CT for 400 kV bank

(d) Other RL assets in the filter bay

Figure 1.9: Hunterston HVDC converter station filter bay assets

(a) Eastern elevation of a capacitor bank

(b) Western elevation (with jumpers in yellow)

(c) Large capacitor bank (eastern elevation)

(d) Large capacitor bank (western elevation)

Figure 1.10: Varied unit orientations at Hunterston HVDC converter station

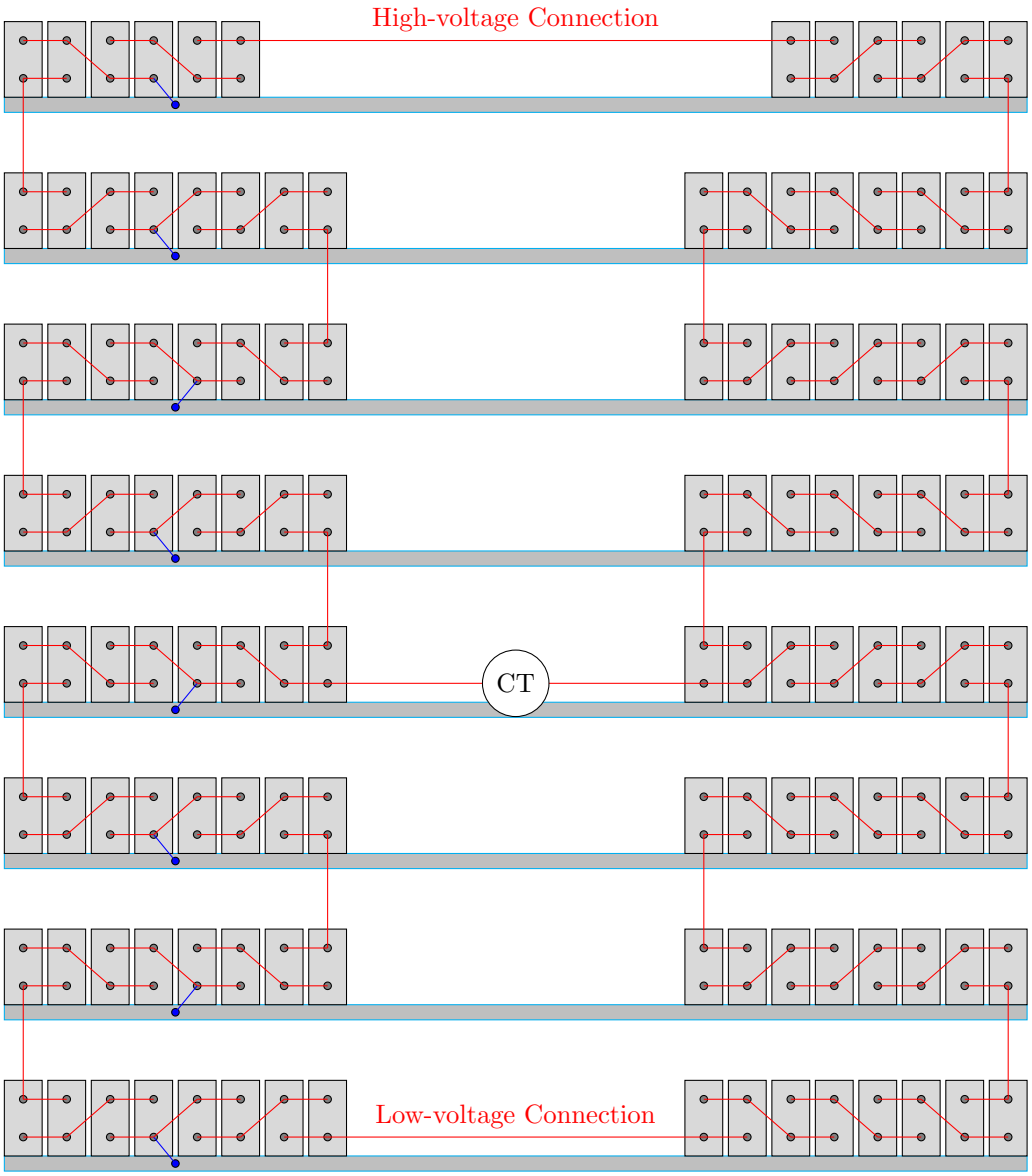

Figure 1.11: Jumper (red), rack tie (blue), and CT connections for both sides of a large bank

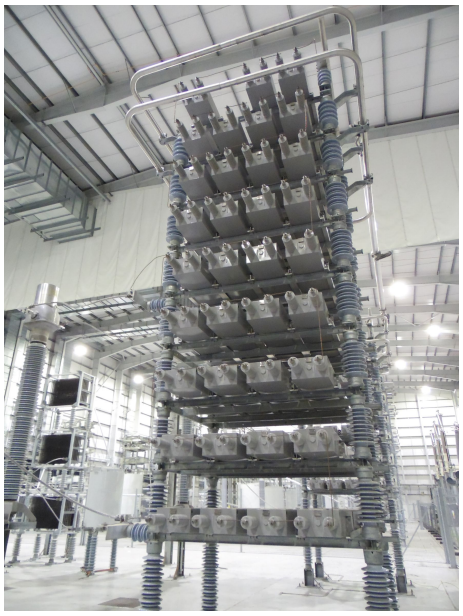

(a) Elevation of a second capacitor bank

(b) Opposite elevation of a second bank

(c) Small capacitor bank (southern elevation)

(d) Small capacitor bank (northern elevation)

Figure 1.12: Capacitor banks at Hunterston HVDC converter station

- Monitoring for the STATCOM. This is a single point of failure for both poles of the converter station and it is not possible to enter the room while it is online, and hard to see in the door, such that a means of online monitoring for this particular asset would be valuable to SPEN as a future operator of this site.
- Monitoring for the cable. The cable is the single largest investment in the link, covers a large distance, and is not easily accessible. It is measured to be around only 4 Ω which is remarkable given that it spans a distance in excess of 400 km.
- Temperature. More than 500 MW of heat need to be extracted from the converter station under full load, meaning there is a significant challenge to extract heat from the facility and there are also concerns for the health of electrical assets and considerations for personnel safety given the unsuitability of the valve and filter halls as a working environment under conditions where these buildings could reach 40 °C. This is partly due to the need to build the station indoors to protect it from coastal air.
- For operational reasons, filter hall doors must be opened once monthly, possibly only to maintain them. Given implications for a need to isolate nearby filter bays to do this, a form of monitoring could save effort needed to repeat this monthly task.
- To comply with safety rules, significant access restrictions are in place for when the system is operational which likely prevent use of thermal imaging as a monitoring technique.
- HVDC converter station switchgear has a heavy duty cycle. Ordinarily switchgear (circuit-breakers) are used for rare switching events for network reconfiguration or in the event of faults (also, ideally, rare). On HVDC converter stations these must make and break current to connect capacitor banks according to load profiles which vary on a daily basis. Manufacturers' expected lifetimes are based on typical switching duties, rather than those associated with an HVDC link. There is a question about duty cycle but also concerns raised about the unusual currents these assets might have to break or make onto.
- Similarly, transformer tap changers are required for voltage control (and therefore operability of this LCC link) so could be put under additional switching duty. This unusual switching duty could invalidate expectations of asset longevity, which is of particular interest for critical assets such as switchgear and transformers.
- As a regulated asset, no ground return path is permitted for this bipole Western Link. If either pole is out of service there is no return path and as such this link does not have the redundancy often associated with bipole arrangements.
- Thyristors could be prone to failures which could be hard to detect. Each individual cell should be reliable, but there are 28 thyristors per module, 3 modules per valve, and 1 redundant thyristor (which gives per phase 84 overall) where failures can be tolerated on up to 3 thyristors on any valve but at most 1 per module.
- Discharge resistors are designed to dissipate static charge before earthing can be applied, but do not have monitoring. It is thus not presently possible to verify a discharge, and is considered best practise simply to wait for the appropriate time before applying an earth.
- DC-side surge arrestors and DC reactors are considered critical assets, but not as critical as the cable which is comparatively inaccessible and a significant expenditure.

Filtering Capacitors

Following this dialogue, filtering capacitors were noted as assets facing operational challenges by the: 1. heightened temperatures they will experience over operational lifetimes; 2. electrical proximity to repetitive switching operations; 3. large number on site; 4. need for filter banks to switch in line with daily changes in loading on the link; and 5. filter hall access restrictions.

Temperature

LCC HVDC links require strong points of connection on existing networks in order to facilitate commutation. Frequency pertubations on the AC side of a converter station can provoke commutation failure. The Western Link was therefore sited between Hunterston, at the site of a nuclear facility, and Deeside, at a strong node on the transmission network: both are proximal to the coast and are vulnerable to salt concentrations in coastal air.

The Hunterston converter station has electrical assets housed indoors: vast filter and valve halls were constructed for the purposes of protecting electrical assets from the coastal environment and providing a homogenous, controllable operating environment. Assets exposed to sunlight, for instance, can be susceptible to imbalanced lifetimes (and operational performance) when one side is exposed more than another. The trade-off is not only capital cost, but also operational conditions over the long term since infrastructure cannot as readily dissipate heat when it operates indoors. Significant cooling apparatus is necessary on site to extract heat generated by the converter station and regulate the internal environment.

1.4.2 A Changing Network

The Western Link was proposed prior to the closure of Longannet coal power station. Hunterston and Deeside were chosen as connection points for being electrically strong points on the transmission network and close to the coast for a cable route to be predominantly subsea.

In time, however, the nature of demand and generation connecting to the network has changed. Coal plants such as Longannet were removed from service, making Scotland one of the first European countries to have completely coal-free generation. As an LCC system, the Western Link takes time to accommodate power reversal, but as synchronous plant in Scotland is being disconnected, this link could prove useful in occasionally transferring power northward beyond its primary envisaged role in supporting export of renewable power to demand centres.

Moreover, one of the risks associated with HVDC converter stations is the nature of characteristic harmonics they can propagate onto electrical networks. One question could concern the effects of increased harmonic output on assets connected electrically nearby, such as transformers associated with the nuclear facility at Hunterston, as seen from the converter station building by the view in Figure 1.6e. At time of writing it now seems certain that Hunterston B will soon be disconnected and eventually decomissioned [24]. While this implies that: (a) the LCC Western Link could stand to benefit from additional voltage support to strengthen this network location; and that (b) this might relieve the aging but critical assets which connect this 1970's nuclear station from characteristic and non-characteristic harmonics induced by a new converter station within close electrical proximity (as highlighted in Section 1.3.3); it remains the case that (c) LCC systems generally connected at 'strong' network locations can nonetheless threaten critical, and potentially aged assets already in situ, such as generation transformers.

Chapter 2

Background

Literature associated with capacitor assets mimics that of HVDC more generally, as introduced in Chapter 1, in that a significant portion of research endeavour understandably focusses on improvements in future designs, and comparatively little considers failure mechanisms of existing technology. Human factors possibly contribute to a dynamic in which self-motivated research best lends itself toward 'building' new technologies (which stand to be more readily seen and appreciated [25]) rather than protection of existing ones, possibly if "no one wants to dwell on the inevitable demise" of important infrastructure [26]. This Chapter introduces the research direction pursued, offers an overview of capacitor unit designs, bank topologies, and briefly discusses failure modes: research directions and motivations are outlined in Section 2.1; and Section 2.2 details the contents of a literature review into capacitor unit designs.

2.1 Research Avenues

More recent HVDC projects are based on modern converter topologies. While such designs are interesting in their own right, this project primarily focusses on routes toward a means of condition monitoring, and as such the most recent project designs are not necessarily of greatest interest. Rather, those topologies which are potentially the most liable to faults and to which operators have already made substantial long-term commitments stand to benefit most from any means of monitoring or mitigation which could result from present-day research. Moreover, learning from existing systems is a potentially fruitful source of improvements for enhancements which can be incorporated in emerging system designs. This project therefore has a predominant focus on more conventional, line commutated converter (LCC) HVDC systems.

2.1.1 Capacitors

Capacitors are deployed as primary or auxiliary assets in support of a variety of applications within power systems, most notably in power factor correction, harmonic filtering, signal smoothing, and voltage support for modern electrical power systems. Capacitors connected in series are effective in mitigating voltage drop (such as on long spans of network used to reach distant demand or generation), and are additionally especially effective in correcting power factors where there is inherent line reactance [10]. Series capacitors can also smoothen transitions which occur where demand or generation fluctuate, as is characteristic of renewable resources.

Shunt capacitors are introduced where line current is limited by thermal considerations.

Capacitors are critical for conventional, modern, and emerging HVDC system topologies. In the context of power electronic converters, capacitors fulfil three important roles:

1. *harmonic filtering* by accommodating and smoothly dissipating the heightened energies in high-frequency signal components within tuned filters;

2. *voltage support*, which for conventional (LCC) arrangements crucially enables smooth converter operation by mitigating against network disturbances and commutation failure;

3. *facilitating power conversion* directly as controllable components in modern topologies.

In filtering arrangements (Figure 1.8) specifically, capacitor banks attenuate characteristic harmonics generated by the conversion process: switching elements introduce discontinuities as they fire (or, in VSC, block) a sinusoidal waveform at any point other than its zeros.

A filter is tuned through selection of capacitive and inductive asset values to join a set which attenuate these $6n \pm 1$ characteristic harmonics for some small value of $n \in \mathbb{N}$.

Shunt-connected capacitances support voltage at points of connection, and on DC lines to smooth ripples. LCC systems are vulnerable to commutation failure which can occur as a result of network disturbances. Capacitors play a role in addressing this by supporting voltage and generating reactive power at important network positions: primarily where a converter station connects to an AC network[1]; and while VSC systems can independently control active and reactive power such that they have less need for AC voltage support, the large capacitances required to support and balance voltage on each pole of a DC link instead present a vulnerability as sources for any potential DC-side faults.

Lastly, more recent multilevel and 'flying capacitor' topologies [27] use capacitances *within* converters, directly controllable by switching elements. This approach requires highly sophisticated converter control, but by varying the energy available to each individual switching element 'cell', it can eliminate the sharp discontinuities which are inevitable in more established systems. In this manner, harmonics can be almost entirely eliminated and much greater controllability can theoretically be attained with these emerging converter station design options.

Of these varieties, those capacitances connected on the AC side of a converter interface are more relevant to conventional forms of HVDC, and have advantages in being similar in form to those used for many applications throughout AC power networks. Motivations for the choice of these assets as a subject of study are presented in Section 2.1.2.

[1] LCC converter stations are often sited close (Figure 1.6e) to synchronous generation also for this reason.

2.1.2 Motivations for Fault Mitigation

Some failure rates (for units and banks respectively) are given in Table 2.1[2] [28] and Table 2.2 [17]. Naturally, the larger the system, the more costly an associated failure becomes: "the cost of a capacitor failure in a small system is significant. The cost of a failure in a large system is massive" [29]. Costs of resolving a fault additionally increase with the fault's progression in time; early detection of incipient faults can offer the opportunity to avoid, resolve, or mitigate the impact of a fault which might otherwise go unnoticed until an inevitable outage, thereby improving overall availability, power quality, safety, and minimising maintenance costs [30]. In the context of Section 2.1.1, it is important to minimise risk of faults on capacitor assets in particular, as:

1. filters on LCC and VSC systems both (a) attenuate unwanted harmonic frequencies and also (b) provide reactive power support, meaning assets in these filters are a common point of failure for both of these functions;

2. there are few examples of capacitor monitoring at present, either in the literature or as commercial options, and those which are reported have been designed around (and thereby constrained in their design by) existing points of measurement;

3. in high-voltage large-capacity converter stations, capacitor units are numerous, such that

   (a) even low failure rates risk being compounded by large asset numbers,

   (b) any known failures are difficult to locate and identify, adding to the time required for maintenance outages, and

   (c) advantages gained through improvements in monitoring, operation, or design can be borne out across a large asset population;

4. as intimated in Section 2.1.1, any learning outcomes could similarly apply to capacitors used in other LCC and VSC filters, DC voltage balancing on VSC systems, and in MMC valves, as well as on flexible AC transmission systems (FACTS) and STATCOMs (such as that also installed at Hunterston).

A focus on failure modes specific to capacitor banks is additionally advantageous, since:

5. these remain only constituent assets within much larger systems, where (a) impulses from converter switching elements could deteriorate capacitor dielectrics, (b) signals from a converter could in turn offer a means of monitoring switching elements' health (such as by trying to detect slow breakdown for static charge layers in thyristor or IGBT modules);

6. a level of technical challenge accompanies a need to monitor (or mitigate) internal faults, since (a) capacitor banks are tiered in racks which increase in voltage, necessitating distributable monitoring, (b) the prerequisite underlying technologies have only relatively recently become commercially available, (c) if any of a redundant set of capacitive elements fail, the rest experience greater voltage so become more likely to fail too;

7. units tend to be designed for operating temperatures T of around $-50\,^{\circ}\mathrm{C} \leq T \leq 55\,^{\circ}\mathrm{C}$ [31], [32] so enclosed filter halls that approach this upper limit could aggravate deterioration;

[2] SC can be inferred to stand for static (reactive power) compensation.

Table 2.1: Example capacitor failure rates [28]

Project (Installed)	Application	Number of Units	Failures (annual %)
Itaipu (1985-1993)	HVDC	15224	0.03
IPP (1986-1991)	HVDC	13786	0.05
Pacific (1987-1994)	Filter	1008	0.15
Itaipu (1990-1994)	SC	5040	0.04
Rihand-Delhi	HVDC	6103	0.004

Table 2.2: Failure counts in Egyptian capacitor banks [17]

Vendor	1	2	3	4	5	Total
Number of Banks	183	83	76	60	55	457
Total MVA_r	886	416	400	324	296	2383
Commissioning (from)	1985	1990	1990	2006	1996	1985
Commissioning (to)	2006	2008	2008	2007	1996	2008
Lifetime	23	18	18	2	12	23
Total	160	62	70	6	11	309

8. with adequate prognostics, acquiring spares and planning maintenance can be arranged according to need, in line with a 'pull-value' principle of continuous improvement; and

9. future designs, maintenance, and operation can be refined by relevant learning outcomes.

Significantly, high voltage capacitors represent a unique hazard in high-voltage testing laboratories: beyond 15 incidents before 1953, a further 9 are documented in [33] from a 30 year duration preceding 2009. These documented cases are unlikely to be an exhaustive list and sadly include 4 fatalities. Capacitor units' ability to store charge poses particular threat to anyone working within their vicinity. Without an ability to detect developing faults within hermetically sealed units (or those associated with discharge resistors critical to dissipate charge following isolation), capacitors will always pose a risk to safety (of debatable acceptability).

As capacitor designs have evolved from the use of external fuses to modern designs where fuses are positioned internally (or omitted entirely) in pursuit of greater efficiencies, locating a faulted element within a bank is no longer a simple case of visual identification. Such assets are susceptible to 'cascading' degradation [1] and comprise many individual units, making them difficult to maintain. Knowledge of fault locations could therefore: aid maintenance, to allow intervening where these cascading faults occur, while additionally informing future designs.

Importantly, capacitor banks are: • modular in nature; • large (in that they each comprise many units), particularly when used to support reactive power at high voltages; • subject to "very little" coverage in the public literature [20]; • the 'target' for harmonic content and signal artefacts; and • susceptible to cascading failure. Together, these attributes make capacitor banks a suitable candidate asset for further study toward eventual measurement and monitoring.

2.1.3 Health Assessment

[34] presents an offline, high-voltage test to facilitate end-of-life assessment of individual capacitor units. This test acoustically detects partial discharge for a specified duration based on component temperature as obtained by a thermal-imaging camera. Virtual measurements, where readings from modular hardware is fashioned into more meaningful measurements by configurable software, is introduced as the approach taken and [34] further suggests that the

proposed test augments those in IEC 60871-1, which are indicated to exist primarily to check asset condition at commissioning.

A risk is ordinarily assessed in terms of its probability of occurring and potential impact, but while IEC 60812 [35] considers risk in terms of these criteria, it additionally includes anticipation and detectability of an occurrence as a related factor. Any complete breakdown of a dielectric, for instance, is usually preceded (so could in theory be anticipated) by incipient dielectric discharges. It is worthwhile to consider the benefits of condition monitoring in these terms; the threat posed by an undesirable event is greatly increased when an operator considers that it could occur unexpectedly, and have broader repercussions should it go undetected.

Routine health assessment through individual testing or online condition monitoring is also a precursor to operational efficiencies; good prognostics can minimise the inventory waste associated with spares, and could allow preferential operating conditions to be afforded to critical assets. With data to support justification, maintenance routines can prioritise a system's least reliable assets, and systems could in principle be operated with asset health in mind.

2.1.4 Condition Monitoring

As introduced in Section 2.1.2, the sooner an incipient fault can be detected, and interventions (if any) applied, the lower the overall costs of the fault. As in [34], platforms exist to measure small numbers of (or individual) units [36], [37], but such approaches appear to be necessarily conducted offline in a static environment, and for which the regularity of these techniques being applied does not appear to be reported; whereas online approaches are available (most notably for protection) but have resolution to at best one leg of a capacitor bank bridge arrangement rather than being able to locate faults to individual units [38], [39], [40]. A distinction between online and offline monitoring is necessary given the different (ambient and electrical) conditions assets experience when under load.

A monitoring technique which could operate online but resolve faults to a smaller subset of units would help to co-ordinate replacements and reduce outage times [40]. Given knowledge capacitor unit designs, their operation, dielectric fluid, and knowledge of an ambient and electrical operating environment, a system designed to monitor a capacitor bank could measure any of the parameters indicated in Table 2.3[3]. These measurands might each 1. have different challenges associated with sensing and measurement, 2. or vary in relevance based on the unique characteristics of an asset's design or environment, but could 3. offer insights and allow asset conditions to be inferred, and to 4. permit correlations to be drawn between: (a) network conditions (outages, configuration, power quality); (b) converter control philosophies, operation, and switching impulses; (c) voltage and reactive power support; (d) (daily) load profiles and switching operations; (e) partial discharge and visible condition; (f) bridge imbalance and element breakdown count; (g) fault and maintenance history; and (h) weather.

2.1.5 General Electric (Alstom Power)

Antonios Tzimas contributed to this project by facilitating a meeting at General Electric (GE) premises on the 3^{rd} of July 2017, at which it was possible to discuss GE's view on HVDC and have a tour of GE facilities. Fault mode effects (and criticality) analysis (FMECA) [35], Markov techniques [41], and fault tree analysis [42] standards describe a set of techniques which can be

[3]Capacitance is listed as a periodic measurement as it can be measured offline [36], [37].

Table 2.3: Potential measurands

	Continuously	Periodically	Ignore
To Measure	Temperature	Dimensions	Pressure
	Voltage	Dielectric Fluid Level	Dielectric Strength
	Current	Capacitance	State of Charge
			Vibration
			Ambient Temperature
			Sound
			Loading (P & Q)
			Electromagnetic Signals
			SCADA Events

useful in analysing risk across a whole system. As discussed in Section 2.1.3, [35] introduces the idea that risk is not only likelihood by consequence, but is also influenced by detectability: a 'risk priority number' can sometimes be used as an attempt to capture the ability or otherwise to detect a developing risk [35]. In view of this whole-system thinking, GE posited that certain maintenance-accessible assets could be sacrificially operated to better protect the most critical or otherwise less accessible assets. However, a prerequisite is the need to be able to first monitor the state of health of the assets hypothetically placed under greatest burden.

2.1.6 Approaches in Data Analytics

There are 3 broad, non-exhaustive, and not mutually exclusive categories for approaches to data analytics, each with relative merits and drawbacks and suitable to different initial conditions.

Knowledge-based Techniques

Knowledge-based techniques are particularly useful where established knowledge or practical experience exists, even if shared only among a limited number of 'experts'. The process of knowledge elicitation is often itself beneficial in bringing to light good (or indeed suboptimal) practices which are applied through experience, as well as highlighting the rules or approaches used toward offering automated approaches, expert systems, or decision support.

The utility of knowledge elicitation techniques remains subject to 1. access to established expertise, 2. time-intensive processes required to gather such knowledge, 3. the relevance of expertise to increasingly diverse asset-bases and evolving operational circumstances, and 4. the possibility of commercial protectionism. Knowledge and associated rule-based techniques have, however, a significant advantage in being comprehensible, explainable, contextualised, and in that regard are more readily trusted to be reputable than other approaches.

Data-driven Techniques

With advances in technology data has become a commodity. Data can reveal valuable asset insights and combinations of 'white-box' and 'black-box' machine learning techniques are both available and increasingly accessible to allow as much information as data contains to be gleaned. Data-driven techniques excel where sufficient data already exists or where sophisticated or

complex behaviour could be difficult to understand unless considered on aggregate, statistically. Their successes lie in the low orders of complexity of information relating to real-world physics.

Drawbacks of these techniques often include 1. a primary need for sufficient, clean, and labelled data, 2. the inexplicability of results or patterns, and 3. a propensity toward overfitting algorithms and techniques to the data available, which risks a system's applicability to a wide and varied asset base.

In addition, with 'deep' and reinforcement learning techniques this process is itself increasingly automated. Hence it is the data itself (and its quality) which often becomes the bottleneck in the application of data-centric approaches. An ability to specify, develop, and deploy sensors to gather data to begin with could help alleviate any constraints posed by limited volumes or quality of data toward asset (or system) observability, diagnostics, and prognostics.

Model-based Techniques

Although based in digital technology, simulation or 'model-based' techniques are more a 'white-box' approach than data-driven alternatives tend to be, in that findings and insights are more accessible to developers and more readily lead to inferences and explicable insights.

Such approaches can, however, be comparatively resource-intensive given the time neccessary to develop a model of sufficient detail (quality) to generate 'synthetic' data of adequate relevance, and the model used might be particular only to a one specific design or attribute. Such models also risk making significant omissions compared to real-world assets and tend to be idealised versions of practical configurations. Aspects such as 1. sources of noise, whether (a) environmental, (b) internal, or (c) associated with measurement, 2. operating conditions 3. manfuacturing defects and 4. asset history are at risk of being omitted from computer models unless explicitly captured. Some of these information sources are relevant to the development of 'digital twins' which aim to track the condition of important assets.

A significant benefit of model-based approaches is that neither physical assets nor measurement data are strict prerequisites to the use of a model to prepare synthetic data.

2.2 Technical Background

In general, there is: 1. a significant established body of knowledge and understanding comprising engineering know-how and well-documented standards which inform present-day capacitor specifications; 2. literature detailing current designs and exploring the potential for new ones; 3. limited (publicly available) detail on unit designs used on particular projects, or measurement data obtained; and 4. information and research about monitoring, failure-mode prediction, detection, or location. This Section reports on items 2 and 4.

2.2.1 Capacitors in Power Systems

As introduced in Chapter 1, *harmonic filtering* is one application of capacitors as part of tuned arrangements (which could be applied throughout a power system but) of particular importance at converter interfaces. A risk of network resonances mean that harmonic filtering is important (to attenuate signal frequencies at which a network is resonant) and can also be a consequence of introducing capacitance in the first place. Furthermore, in *modern converter topologies* significant capacitors are also required to support voltage on the DC link, or even to form switchable modules in converters themselves [8]. Supercapacitors have "not yet" become a further role for capacitance within power systems, wherein energy can be stored for *transient control* (p23, [10]). However, a simple use of capacitance in power systems is in *power factor correction*. In handling large amounts of energy, the more efficient a system the greater the rewards which follow over a long term; even small improvements in efficiency sustained over an asset's life can lead to significant accumulated gains. Network efficiency is thus paramount to electrical power systems, for which often "the resistance of lines is negligible compared with reactance" (p70, [10]). As networks predominantly comprise inductive assets their reactances are seen by a signal at fundamental frequency, which draw reactive power, leading to power factor lag, inefficiencies, and voltage drop (p170-176, [10]). The most significant application of capacitance on power systems is therefore at high-voltage, to counter such effects and maintain system power factors close to unity, whether connected: 1. as banks of static capacitors, (a) in series or (b) in shunt; or 2. within power electronic compensators, such as (a) static VA_r compensators (SVCs), and (b) STATCOMs. Directly connected capacitors are more efficient but less controllable than compensators, and all require (some form and arrangement of) capacitor units to be installed [10].

2.2.2 Unit Design

Absolute values of capacitance required for power systems can be so significant as to necessitate subdividing a bulk capacitance into an array of connected units, in the interests of: 1. substation footprint, 2. practical manufacture, 3. transport, 4. economic scalability, and 5. heat dissipation. Practical capacitors therefore comprise collections of connected units. Incidentally, this modular approach lends itself toward resilience to faults (both through simple redundancy and by limiting fault extents) as well as possible maintenance and replacement through targeted intervention.

A line diagram of an internally fused unit is shown in Figure 2.1, where units typically comprise: 1. bushings; 2. wound elements, in turn comprising (a) terminals (whether discrete foils or a sprayed surface), (b) dielectrics (often polypropylene in modern designs); 3. element

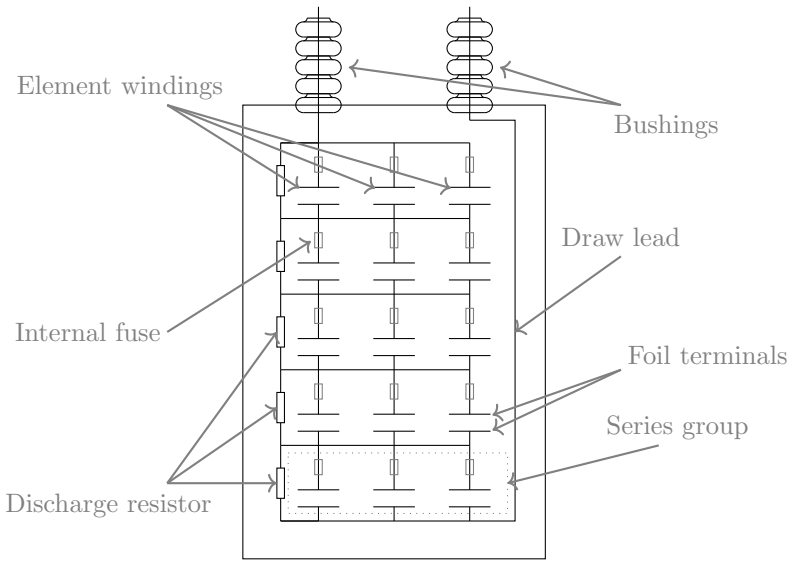

Figure 2.1: A capacitor unit line diagram [28]

groups, connected by (a) schooped ends[4] (or tabs in older designs), (b) connecting leads ('busbars', in some designs), (c) a draw lead (if required) to return a circuit to bushings on a common side; 4. discharge devices; 5. fusing (if applicable); 6. a dielectric impregnant; and 7. an outer housing (sometimes referred to as a 'can'), with (a) rack mounts, (b) a nameplate, and (c) an outer coating. Unit design and construction is detailed in [38] and [1]. The remainder of this Section explores select trade-offs in capacitor unit design.

Protection

Combinations of element size, and parallel and series connections between them can be chosen to achieve desired nameplate ratings; unit sizes designed to balance asset footprint with heat dissipation; and fusing options selected to trade tolerance to dielectric breakdown with operational efficiency [38], where:

1. *externally fused* banks remove each whole unit from service in the event of failure, and make faults easy to find although they require more space for fuses to fit;
2. *fuseless* units short-circuit each faulted series section [38], reducing reactance;
3. *internally fused* designs isolate only failed elements, increasing reactance after each element failure;
4. *'self-healing'* or metallised-oxide designs offer fault resilience and graceful degradation, but at reduced operational efficiency [44]; and
5. *unfused* designs [45][5] are possible provided other protection can be deemed sufficient.

[4]Winding ends are 'schooped' by leaving a margin of one terminal exposed throughout the winding while a margin of dielectric protects the other (Figure 2.3 illustrates the use of margins). The surface of a winding end can then be sprayed with molten metal to form a single terminal and facilitate a connection [43], [44].

[5]Fuseless capacitor bank designs feature (typically long [38]) strings of series-connected units which are in parallel but unconnected from one another; whereas unfused options feature connections between parallel strings. In both cases, breakdowns cause permanent short-circuits, where in the case of an unfused configuration the entire group is shorted as a result.

Fuses are the "most common and the generally recommended" approach to unit protection [46]. In selecting a fusing option, the protection each choice offers must be balanced with operational efficiency and the capacitance an operator is prepared to lose with any single element failure.

Discharge resistors are used to safely dissipate charge after a disconnection, and as will be further discussed in Section 3.1.2, discharge resistors are sometimes illustrated [47] as being located between bushing terminals, although this is often a mere simplification as in practise such components can be located throughout a unit [28]. They must dissipate residual charge in the event of disconnection to below 50 V in 5 minutes or less for capacitors rated above 600 V [48]. It is safety critical that discharge resistors and their connections do not fail.

In 'self-healing' designs, metallized electrodes are vaporised by the arc of a dielectric breakdown, thereby isolating the failed film and permitting most of a series element to remain active. Self-healing designs are robust, but unsuitable for high-power applications due to overheating and thermal gradients [44]. Instead, high-power applications more commonly employ internally fused designs [17], [49].

Temperature

As unit dimensions affect the ratio between volume and outermost surface area, a unit's size determines its ability to expel heat. A cube increased in size by a factor k, has volume $v = k^3$ and surface area $a = 6k^2$, such that the surface-area-to-volume ratio $\frac{a}{v} = \frac{6}{k}$. Heat dissipation is inversely related to size, as smaller units are better able to dissipate heat to an ambient environment, but larger unit sizes make more efficient use of a substation footprint. This is a critical attribute for these assets since temperature not only influences operational characteristics, and neutral (or bridge) imbalance, but crucially affects capacitor unit *longevity*.

For every 8 °C increase in hotspot temperature, useful life of a capacitor unit halves [1], [50]. Capacitors heat under load due to both resistive (potentially large) and dielectric (often negligible) losses taking place within the enclosed unit, where choice of dielectric and design of terminals (resistive losses are higher in metallized oxide designs) influence the scope of this heating [51]. Foil alignment, coupled with the high thermal conductivity of metal terminals in discrete foil designs, consequently leads to an "anisotropic thermal conductivity" [51][6].

Akin to transformers, for which there exists a widely established literature, capacitor unit hotspots influence their projected useful lifetimes [52], and are highly dependent on this anisotropic nature of heat dissipation. Thus, the ambient conditions, electrical loading, current health, and unit design all contribute to the thermal conditions which dicate dielectric aging.

Dielectrics heat under electrical load, and as these materials degrade incipient faults will arise which are in turn likely to have a thermal component and further heat local dielectrics. Thermal components of incipient faults could therefore be detected, should they exceed background levels of normal thermal activity. Thus, a progression from 1. a healthy unit (with dielectric heating under load) to 2. one with an incipient defect (a point heat source of increasing power), to 3. a unit with a failed or fused element and a change in heating for its remaining element dielectrics, could be a means of corroborating electrical measurements. This anticipated progression of changing temperature in time could be explored in future work.

Sunlight can disproportionately affect one side of a capacitor bank and cause imbalance [53], seasonal variation in ambient temperatures, and daily variation in load all contribute to a dy-

[6] 'Anisotropic' behaviour is that which varies by direction; heat travels more in some directions than in others.

namic thermal environment, in which changes in unit temperature, moreover, naturally lead to deviation in capacitance values (as typically documented in unit datasheets). It is inevitable that *statically measured* (and nameplate) unit capacitances change under operational conditions, however slight any change might be in relative terms.

These technical considerations accentuate the emphasis placed on the temperature of filter halls expected for operational conditions of the Western Link as introduced in Section 1.4.1, given that this installation is unusual for having assets housed indoors in combination with the need for cooling, as discussed in Section 1.4.1.

Dielectrics

Dielectric integrity dicates longevity of electrical assets and plays a "major role" as a primary component in high voltage capacitors [38]. Relative permittivity (ϵ_r) is a function of voltage amplitude, frequency, and temperature [54], and a critical attribute of a dielectric material.

Designs have adapted over time: historically, units made use of aluminium foil terminals and dielectrics such as askarel [55]; before Kraft paper impregnated with mineral oil as dielectric became more common [29]; paper dielectrics have gradually been replaced with polymers, most commonly polypropylene (polyethylene terephthalate, and polyvinylidene fluoride [29]) films, with a transition too from PCB-containing dielectric fluids to environmentally safer but flammable alternatives [26], toward modern, synthetic, environmentally-friendly dielectric fluid choices [31], [56], [57], [58]; and for modern converter applications even entirely without liquid dielectric [59]. More recently, research investigates nanocomposites' ability to enhance dielectrics [60], [61], [62].

Discrete foil designs are more common in high-capacity converter stations due to their efficiency, since discrete layers of aluminium foil have high conductivity by comparison to metallized-oxide designs which instead offer a resilient alternative in which dielectric layers are sprayed with metallised vapor to create terminals which vaporise in the event of a dielectric breakdown, to thereby protect the vicinity of the dielectric fault afterwards. In such 'self-healing' designs, internal arcing 'vaporises' part of the terminal and thus allows the unit to degrade gracefully [43]. Polypropylene layers [63] and aluminium foils can be 12 μm thick, and multiple layers of dielectric material can be combined to a thickness of 20 μm [44] (dielectric layers are typically combined to limit the effects of manufacturing defects and improve reliability), and some commercially available units have an outer housing only 1.5 mm thick.

Dielectrics also vary, and while polypropylene film is commonly used within element windings, fluid filled ('wet') unit designs make use of a dielectric fluid to permeate the structure and occupy any voids within windings or winding groups. Modern dielectric fluids are often proprietary synthetics, and selected for fire and environmental safety: SAS and C101 are two categories of liquid dielectric, which have sufficiently similar properties as to be interchangeable for practical purposes [64]. Chemical properties are known for some trade-name dielectrics such as Midel 7131 [65], and since the energies and materials involved in capacitor units puts them at risk of explosion [66], fire safety (as well as environmental reasons) can be assumed to motivate the choice of these fluids irrespective of their particular attributes.

Windings

Capacitances in parallel support reactive power, and those in series accommodate voltage. Windings, groups, and capacitor units can therefore be connected in series or parallel, so that together an asset can be respectively configured to accommodate these requirements. Figure 2.2 illustrates that element windings are typically formed by layering strips of aluminium foil between a polypropylene dielectric, which are wound to maximise surface area between the terminals, while facilitating a straightforward manufacturing process and preparing an element in a compact form [29]. Provided a mandrel can be securely removed from the centre of a winding, the larger the mandrel the less the turn radius of a resulting winding. Note that, for the purposes of this Dissertation, winding 'alignment' refers to a vector orthogonal to flat foil surfaces which result from stacking windings within a unit (Figure 2.2), as a unit vector describes a 2-dimensional plane.

An area of metallic foil terminals remains exposed at either end while layers are wound to form each capacitive element, before windings are stacked and compressed to fit the housing. Adjacent rolls are connected by the exposed foil terminals on each, firstly in parallel into groups termed 'series elements' or "series sections" [1] (element groups), which in turn are connected in series and stacked throughout a unit. A unit's voltage rating and capacity for reactive power support are respectively determined by the total capacitance connected in series and in parallel in this manner, and which vary by design.

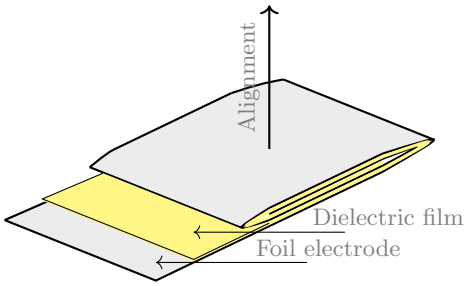

Figure 2.2: A simple winding

Relative to dielectric layers, each terminal layer is slightly extended at one end of the winding and recessed at the other such that an electrical connection can only be made at either end. Thus, a winding of two aluminium foil terminals enables each terminal to connect at either end: a simple margin of dielectric protects each from connecting to the alternate terminal. Slightly more nuanced designs can have multiple foil terminals for each layer by introducing margins at intermediate points along a winding in addition to those at each end: this permits a greater number of interleaving series capacitances to form between layers in particular to accommodate high voltage designs, as shown in Figure 2.3.

Extended-foil winding designs, in which part of a terminal is exposed from the dielectric at either end, allow connections to be formed by "schooping" where solder is applied to each extended terminal, or alternatively, connections can be made with foil tabs soldered to individual terminal layers. By these means: element windings connect in groups; these groups connect as a stack; and in turn, the stack of winding groups connects to a unit's terminal bushings (typically at one end of the housing). In either case jump leads can be used to connect winding groups and if a design has an uneven number of winding groups a draw lead can complete an internal circuit by connecting a terminal at the far end of a unit to the remaining bushing.

Connections are a point of potential failure, so care is due [17] since foil edges (and tabs) can aggravate electrical stresses [67] and these connections are also prone to resistive heating as current is conveyed between elements, even where such connections are electrically sound.

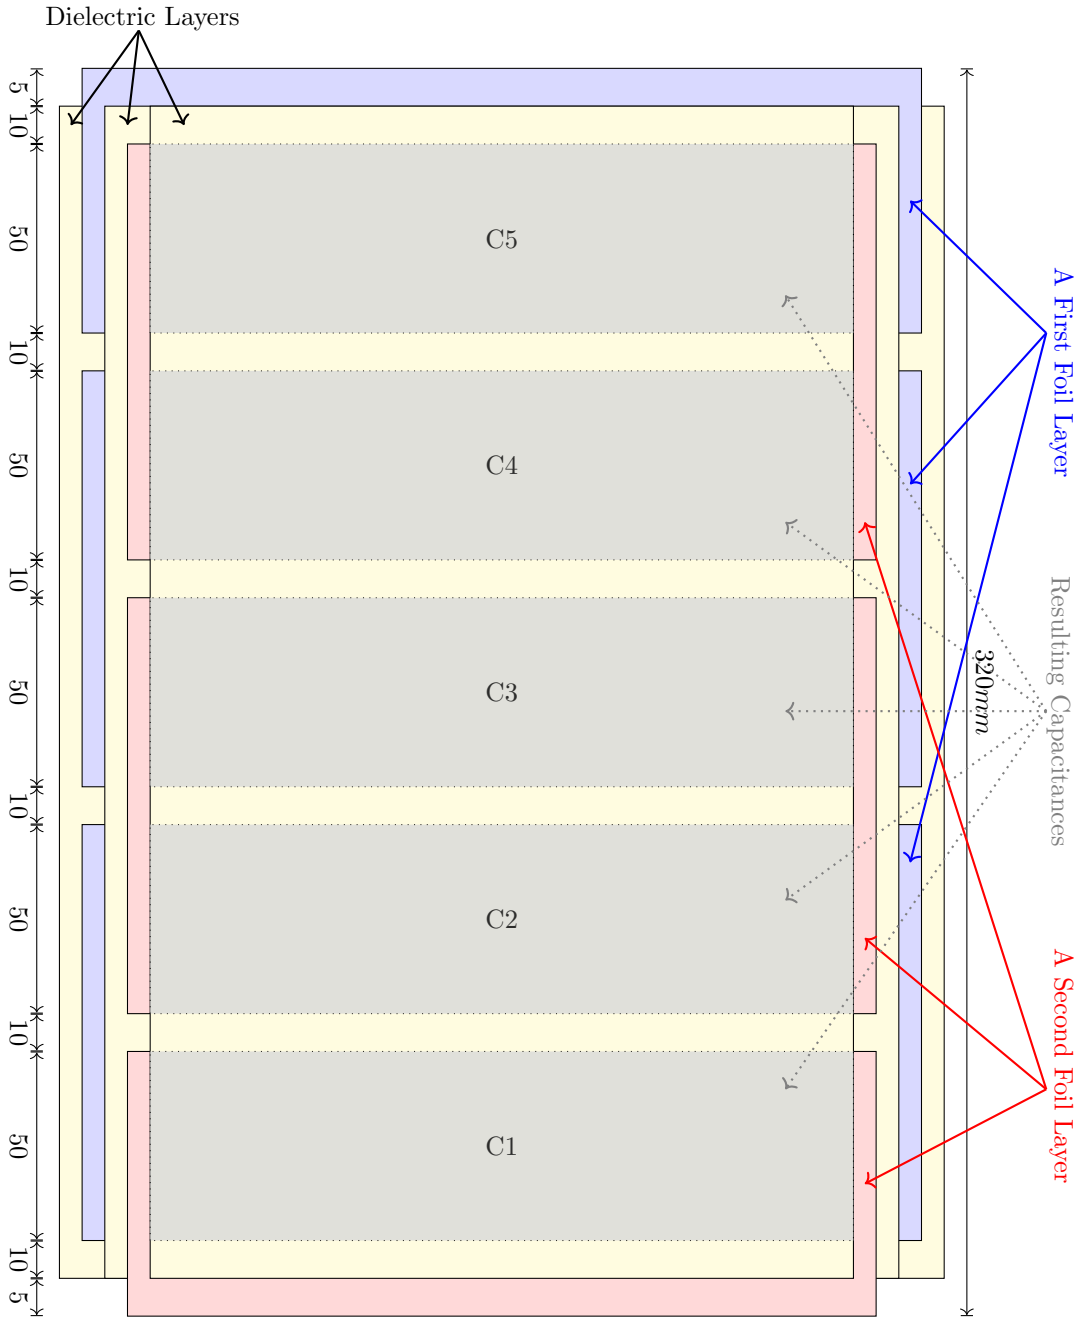

Figure 2.3: Design of interleaved windings

2.2.3 Capacitor Banks

As seen in Figure 2.4[7], a capacitor bank consists of a number of individual units connected in series or parallel: the amount of capacitance connected in parallel permits support for reactive power, whereas the capacitance connected in series determines the maximum voltage a capacitor bank is able to withstand [38]. Adjacent racks of a capacitor bank are at different electrical potentials (to reduce the voltage step from the asset structure and the signal it carries) and hence are insulated from one another.

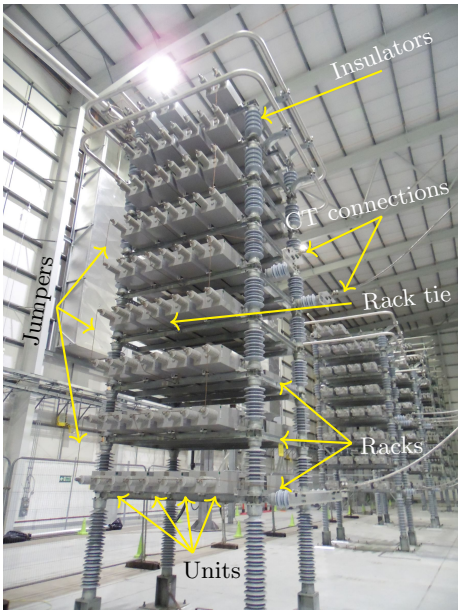

Figure 2.4: A large capacitor bank

In addition to varied fusing arrangements, capacitor banks can be arranged in: 1. ⅄, 2. Δ, 3. double-⅄, or 4. H-bridge topologies; and in each case can be either • grounded, or • ungrounded [45]. There is therefore a high degree of variability between designs, although similar principles apply in each case. It is in part due to this variability in the design of both units and banks overall that configurability is an aim of the model to be introduced in Section 3.1.

In HVDC filtering, the voltage rating of high-voltage capacitor banks is usually the "most significant factor in determining the total cost of the AC filters" [68], and as reactive power and voltage requirements are respectively met through connections in parallel and series[8], this could in part be due to practical limits on how many units can be supported by any one rack. Shunt capacitor banks designed to connect at high voltage must necessarily have many units among which to distribute voltage; with only a small number of units supported by each rack, it follows that the number of racks must necessarily increase in turn (alongside insulators, jumper connections, rack materials, and associated space and access requirements), raising the height of the bank and escalating overall costs.

To connect to line voltages within a limited substation footprint, racks increase in voltage as they increase in height. Rack mounts allow units to adopt different mounting arrangements and conduct voltage to secure unit housings at the potential of a rack tie. Each unit has a metal housing and sits on a rack: since racks are at voltage, so too is this housing. Only a limited number of capacitor units can therefore be used on each rack in order to protect element windings from rack voltage, where each capacitor relies on the dielectric fluid it contains to insulate elements from housing. 'Jumper' connections are made between units in series to form 'strings', which in turn are connected in parallel to support a common phase.

Rack faults pose a particular risk in such circumstances due to being difficult for protection to detect [69]. Where the same phase is connected by multiple strings, faults which occur between them (whether between jumpers, or where unit housings are affected by virtue of being tied to one of these strings), conventional approaches to protection can take a long time to detect and isolate such faults, which can cause "major equipment damage" [69] over sustained durations.

Ordinarily, in a bridge or neutral arrangement the intention is for parallel strings of units to

[7] A current transformer (CT) is used for protection but is not shown.
[8] This is similarly the case with the configuration of elements within a single unit.

Table 2.4: Table of consequence of failure by fusing type [38]

Fusing	e_{Series}	$e_{Parallel}$	SC/OC	C_{lost}	Q_{lost}	V_{Series}	$V_{Parallel}$
External	many	few	SC	small	small	small	0
Internal		many	SC	small		no change	small
Fuseless	many	few	SC			$\leq 10\%$	0
Unfused			SC			$\leq 10\%$	0
Self-healing			OC	minimal	minimal		

balance so that minimal current and voltage is seen on a connection between their midpoints. This allows a protection relay to discern any change in operation which leads to a sufficient imbalance to cause a trip, but this traditional approach is typically insensitive to degradation on parallel branches, where a bridge can remain balanced despite 'offsetting' faults [1].

In addition, fault conditions must be discerned from inherent tolerances, failure of redundant elements, and environmental factors such as changing ambient temperature or sunlight heating some strings more [70]. Capacitor bank failures vary [66], and under ordinary circumstances a bank might trip on average less than once each year [71], but where it remains possible for internal breakdowns to occur within designed tolerances.

The more evenly capacitor elements are matched, the more evenly are stresses on dielectric shared, and as increased voltage stresses accelerate dielectric aging [1], consequently the less likely any single element is to experience disproportionate stresses such that it would be at risk of premature failure. Accelerated aging further predisposes a bank to cascading element failure: where a failure occurs, dielectrics in the fewer remaining elements (whether those in series with a failed series section in fuseless designs, or those in parallel with a failed element in internally fused designs [70]) are subjected to greater voltages; and effects on asset reactance can moreover expose a weakened bank to greater voltages in practical settings.

2.2.4 Cascading Failure

In fuseless designs (featuring long strings of series-connected elements which fail to short-circuit), element failure can short an entire series section and leaves fewer elements to share the electrical load [1]. In turn, "as more series sections fail, the voltage stress on the remaining series elements increases, and the time between failures is reduced" [1]. In internally fused designs (where elements fail to open-circuit on account of fuses), faults reduce the series-section capacitance and place greater stresses on elements remaining in parallel with a failed one.

Moreover, where capacitance (C) is reduced, and in some scenarios[9] (with a strong network connection) increased reactance $X_C = \frac{1}{j\omega C}$ can consequently attract greater voltage to be shared by the fewer remaining elements.

These effects of element failures heighten a bank's susceptibility to further degradation and, should they be left unchecked, will provoke an element failure '*cascade*' [47], [49], (p10, [45]), where fused designs are "possibly more likely than fuseless to experience cascading failure" [70].

For each fusing arrangement, Table 2.4 [38] considers the implication on the remaining healthy elements and the bank overall. Thus, where one component (whether an element or unit) fails within a parallel group, it is likely that failure will recur within that same group. Where a sufficient number of parallel elements experience faults, or otherwise (as is the case for externally fused units) where any one element fault inserts a short-circuit to an entire

[9]This is less likely to apply where these assets support network voltage on long transmission lines.

element group, voltages on other series-connected elements increase, which in turn subjects them to an additional share of the stress and "accelerates the aging of the dielectric" [1]. In addition, Moreover, as voltage stresses increase each time, mean times between subsequent failures reduce, and it is worth considering that in fused arrangements the number of operated fuses alone is "not necessarily a reliable indication of the actual condition of the capacitor units in the bank" [72].

Potential for knock-on consequences necessitates that protection relays discriminate between faults which leave a bank within its tolerance and those which potentially place it at risk of overvoltage (and further breakdown) due to subsequent signal artefacts, switching transients [48], [30], and transient inrush currents [17]. "System transient overvoltages" are a particular risk to capacitor banks [45] as they can lead to additional breakdown by exploiting a prevailing weakness. This is not always straightforward, partly since failures at different locations can mask one another from protection relay oversight [1].

Although high voltage capacitor units can in general be considered reasonably reliable, the significant numbers of units and elements which comprise a large capacitor bank compound the likelihood of even improbable faults. Thus, should an early fault occur on one phase of a capacitor arrangement, the likelihood of further such faults is aggravated toward that same phase, accentuating both the mean time to failure following a fault and the resulting phase imbalance. In time, this cascading effect can encourage early onset of faults, but designed element redundancy offers time in which an operator can be forewarned to schedule maintenance interventions as a mitigation. Such intervention is only possible, however, should it be possible to detect, track, and preferrably locate incipient faults or early failures in a timely manner.

Of course, that is *if there is* sufficient tolerance; in the event that an 110% overvoltage threshold is reached (which on a small bank could be due to a single element failure [45]) then the speed at which faults can be located matters all the more. Furthermore, where standards (and literature) stipulate a minimum of 10 units or elements be connected together in series to provide for this overvoltage requirement, this is based on the assumption of a *linear* relationship between the number of healthy units or elements and the voltage seen by any one.

2.2.5 Approaches to Failure Detection

Beyond intentional motivations listed at the start of Section 2.2.2, in terms of detecting, locating, and allowing for possible interventions to faults, it is "convenient" [46] that banks are structured in such a modular fashion. Despite this opportunity, it remains the case that "very little research" on fault detection methods in capacitor banks has been reported in the literature [20]. What literature there is discusses:

1. *offline measurements* applied to each unit [34], which "may or may not be a part of the regular maintenance schedule" [45];

2. *imbalance techniques* wherein a fault affects the current and voltage distribution, which in turn affects the imbalance seen by the protection relay, and a change in phase angle of this imbalance current and voltage can identify the phase and section of a fault using existing protection relays [38], [47], [73], and which can even detect concurrent failures [74], in general allowing failures to be located to within one arm of a balanced arrangement;

3. *reactance techniques* to identify transitions from a healthy to a faulted state with existing measurement points [69], [75] (including 'superimposed reactance' [39], [20], [76], [77]);

4. high-fidelity *switching transient measurements* to detect problems with capacitor faults [78] or associated switchgear [30], [79]; and

5. measurement of *polarisation current* and *dielectric loss* [63].

For the most part, these techniques benefit from existing measurement systems, but as a requirement this also limits this direction of research as a limited number of measurement points is "less than ideal" [39], especially as banks in high-voltage contexts grow large enough that a single arm of a bank can consist of as many as 30 individual units (Figure 1.10). [73] notes a further benefit of tracking faults: if no more than a tolerable number of faults are detected (within designed redundancy), scheduled maintenance could even be made optional or avoided.

Detecting faults is one challenge, but locating them is another altogether. Should a fault arise in a bank (without external fuses as a visual indicator), an ability to identify the phase and 'string' of series connected units in a bank can help by defining a minimum search space within which engineers are to manually locate faulted units during an outage. Given capacitors' role in voltage support, it can be important that "outages be as short as possible" [38], but invariably require a "time-consuming process" wherein "the majority of the outage time" [47] is spent locating faults. Accordingly, this search prevents proactive interventions being made frequently, and increases outage durations in response to faults which place a bank outwith tolerances (which in turn can have implications for a wider system, such as loading of conventional HVDC).

Increasingly, information and embedded system ubiquity opens up new fields in computer science and electronics (from 'cloud' to 'fog', and more recently 'edge' computing), such that there could yet be opportunities in alleviating such measurement restrictions by applying distributable means of monitoring and detection. Should additional measurements be possible, not only might it be possible to locate element failures to a much narrower collection of units as a less unweildy search space (thereby reducing outage times), but it could allow for some level of monitoring of *incipient faults*, rather than exclusively those resulting from elements which have already failed.

Moreover, even in view of these reported state-of-the-art techniques for failure detection, an improved understanding of where fault locations are *likely* to be within an otherwise unweildy array of capacitor units could be used alongside a recent count of such failures to inform which units or parts of a capacitor asset are best prioritised to be searched first. The literature reported in this Section predominantly considers system-level studies and simulations in order to consider the signals which can be detected from existing transducers.

This Dissertation instead attempts to return to fundamentals of the assets in question in order to improve understanding of likely fault locations, so as to: 1. inform which units are best prioritised in a search; 2. explore promising approaches to more ubiquitous measurement for future online monitoring; which could in turn 3. facilitate online monitoring and location of incipient faults in addition to element failures.

Chapter 3

Modelling

This Chapter: in Section 3.1, introduces COMSOL Multiphysics® as a modelling environment and presents a simulation model developed to study capacitor units; before Section 3.2 explores elements of good practise and approaches found to be useful to extensible model development.

3.1 Modelling in COMSOL Multiphysics®

COMSOL Multiphysics® is a simulation environment with strengths in combining the simulation of concurrent physical phenomena. It has been chosen for this project due to: this ability, its versatility, and for having an active community.

Model files in COMSOL® adhere to a consistent structure in a 'Model Builder' window, which contains a tree of nodes which specify each simulation model. This deliberate user interface design promotes a common workflow, to improve legibility and thus portability of model files between users and across application domains [80]. The primary nodes in this structure are: 1. Global Definitions; 2. Component with (a) Definitions, (b) Geometry, (c) Materials, (d) Physics, and (e) Mesh; 3. Study; and 4. Results.In addition to a tab associated with each of these items, the ribbon within the environment window also contains a 'Developer' tab from which an 'Application Builder' can be accessed. This feature allows studies to be packaged such that they can be run from a simplified user interface, and shared to allow others to explore simulation configurations first-hand, and secondly this tab permits the user to script aspects of the model file to improve the repeatability of common or tedious tasks, for instance. This feature provides functionality to support sharing model files and even user-friendly 'applications' using COMSOL Server®. Model files are in a proprietary format and a named single user licence has been obtained for this project.

The capacitor model illustrated in Figure 3.1 is developed as a basis for simulations. This `Geometry` is intended to be simple to permit any problems in its development to be easily uncovered and reduce simulation times, while being amenable to accommodating element numbers representative of real assets and allowing general principles of operation to be explored. It has undergone a number of iterations over the duration of this project, but is intended to be configurable, extensible, and reasonably scalable.

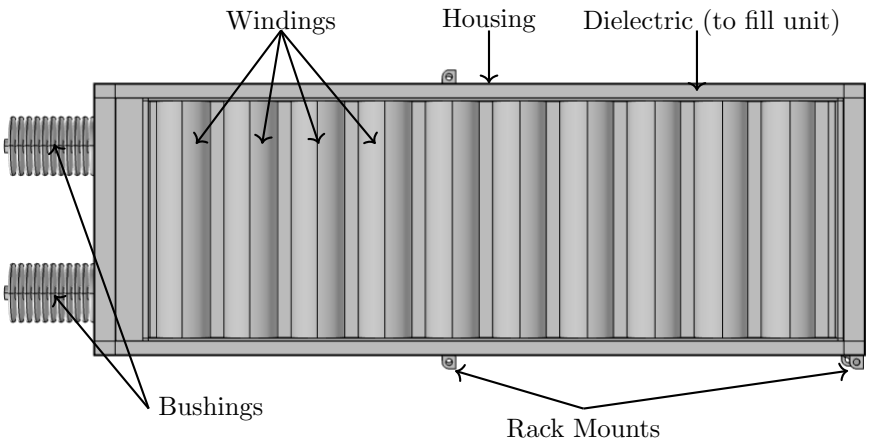

Figure 3.1: A baseline model geometry

3.1.1 Model File

A model file (Figure 3.2), adapted and extended over the course of this project, has been stored for posterity in the University of Strathclyde's PURE repository under a "Capacitor Unit Model File" title and a Creative Commons BY-SA 4.0 license.

Its DOI is: https://doi.org/10.15129/a03a8415-43c0-4e19-9fee-748b67f5c0e1.

This file is neither perfect nor a finished product but is stored for reference as a possible starting point for future development, or as a means to gain familiarity with COMSOL® as a platform, in conjunction with associated documentation.

As discussed in Section 4.1.3, modelling has potential to calculate basic characteristics for a given unit winding design, such that these can potentially be used to develop diagnostics and thereby improve manufacturers' offering. Therefore, a CC BY-SA 4.0 licence is employed to encourage models used as the basis of future diagnostics to be shared: for sanity checking, but also to promote competition by decoupling a diagnostic offering from manufacturers' pre-requisite knowledge. Unless this baseline model is shared in turn, operators could risk becoming beholden to an asset's original manufacturer for diagnostic and maintenance for the entire duration of its operational life.

Versioning

Versioning is conducted manually with version numbers assigned to filenames such that specific files and results sets can be referenced: a primary number is assigned to iterations on the main workflow and secondary numbers assigned to those on workflow branches. Brief descriptors for each iteration are included as part of this naming convention. Table 3.1 lists file versions used to generate the datasets used in each of the main studies in Chapter 4 and in Chapter 5.

Rather than store each model file containing each set of results, at various stages of development, a single model file is stored in the repository with which sets of results can be generated by reapplying the configuration used in each `Study`. This permits the stored file to collate enhancements introduced to the model between each set of generated results (such enhancements include the use of 2D boundaries, improved winding details, configurability through parameters and `Definitions`, and support for multiple geometrically similar units).

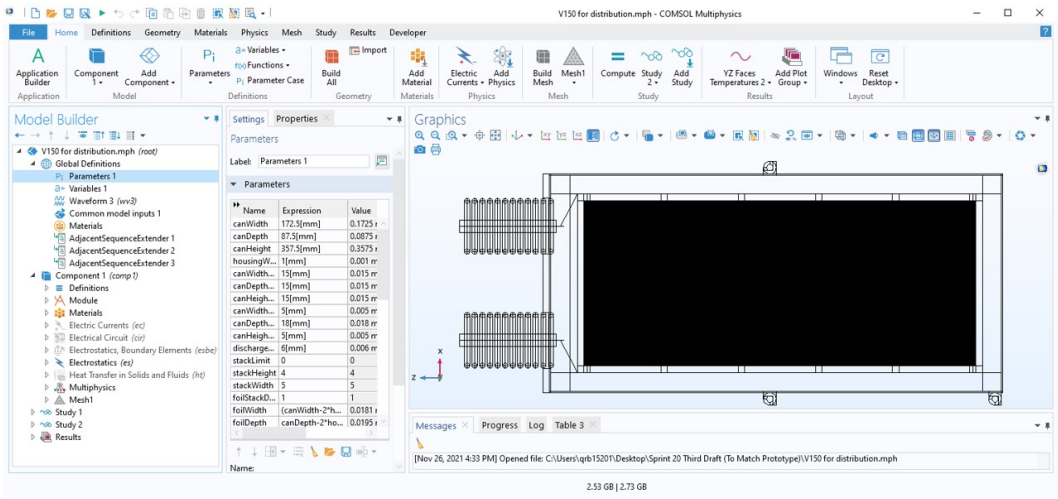

Figure 3.2: The COMSOL® environment, showing a model tree, parameters, and graphics

Table 3.1: Subset of model file version history

Publication	Version number	Relevant Data Tables	Sprint
EIC2019	V40-6	`Probe Table 28`	5
CEIDP2019	V81-5; V81-6-4; V81-7	(1D and 3D `Plot Groups`)	9
COMSOL2019	V93-C1	`Probe Table 28`	12
EIC2020	V105-11	`Table 5` & `Table 6`	15
Section 4.1.4	V105-11-4	`Probe Table 1` & `Table 2`	18
Section 6.3.1	V150	`Table 2` and `Table 3`	20

3.1.2 Geometry

In COMSOL®, a `Geometry` forms the environment in which 2D and 3D model components are available, where components comprise points, boundaries, domains, and objects: these distinctions are necessary to specify references when extending each main part of the model.

While it is possible to explore and learn about a given technology or application as it is developed, it is preferrable to begin with a design which includes an understanding of some of the main changeable attributes associated with studies to be run. Thus, features of a model design (and these attributes in particular) can inform a set of parameters which can be defined and used to specify Geometry components, such that the resulting Geometry can be made *readable*. In turn, readability improves a model file's: • *configurability*, as a rebuilt geometry will adapt to parameter changes; • *maintainability*, as it is easy to debug or modify a geometry in relation to pre-defined, human readable quantities; • *portability*, as readability allows others to more readily uncover the intended meaning behind a developed geometry; and • *extensibility*, in that subsequent adaptations are easier to make when an initial model can be understood.

Axis Orientation

Parameters also largely remove the need to provide specific points in relation to an origin and model axes, which provide a point of common reference for the model `Geometry`. Components of the `Geometry`, `Mesh`, and in postprocessing results have been named (for ease of reference) according to their axial alignment: foils are described as being stacked in alignment with either the 'X', 'Y', or 'Z' axes, for example, and planes chosen for defining `Mapped` meshing according

to the axes along which a boundary lies.

Capacitor Design Options

For capacitors used in power systems, capacitive elements are wound, but there remain a variety of design options. Specifics of winding designs are largely left to interpretation, as limited detail of winding designs is given in publicly available literature, which are further discussed in Chapter 6. Models developed in this project make use of winding designs where foil edges are at the centre of the winding, although a variety of sizes, alignments, interleaved designs, stacking arrangements and number of turns are possible. In its present form, this model file uses conditional (`If`, `Else`, `End If`) nodes[1] to accommodate foil alignments in either of 3 directions but does not also allow winding turns to be positioned on either of the remaining axes, as a default is chosen for each, and it offers only limited support for multiple stack designs. In other words, the model offers 3 of 6 possible combinations (shown in Figure 3.3, in which yellow planar foil boundaries are bounded by purple surfaces to indicate the opposing sides on which winding turns are made) of foil and foil-turn alignments, and accommodates multiple winding stacks only divided perpendicular to foil turns.

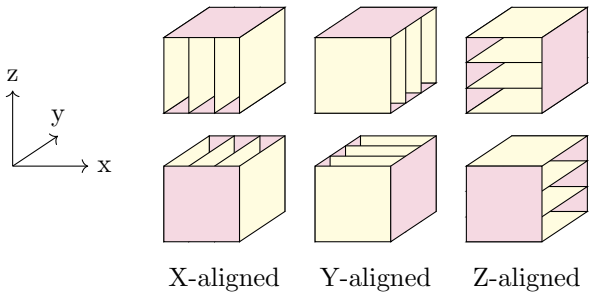

Figure 3.3: 6 alignment options

Boundaries as Foils

Early iterations of the model file used in this project initially involved 3D capacitive 'foils': this effort was misguided and resulted from a limited familiarity with COMSOL® at the time. Boundaries alone can be established within the `Geometry` as a way to model (extremely thin) foil terminals, which is a preferrable approach since: 1. the resulting `Geometry` is cleaner; 2. the `Mesh` is simpler and more efficient, as it must only cover dielectric rather than domains of both dielectric and foil; 3. an improved representation of reality, since (a) real foils are thin, and (b) as a greater number of turns can be included such that windings themselves are more true to those used in real assets. However, this stands to further be improved upon by allowing layered materials to be used for winding definitions in future model revisions, as will be discussed further in Section 7.3.

Peripheral Components

Initial models include space for a single discharge resistor to serve as a single device between terminals for the whole unit (based loosely on simplified literature [58]). An improved understanding of discharge resistor positioning developed over the course of this project, but `Geometry` and `Mesh` sequence nodes to support this early option remain within the model file, using conditional `If` and `End If` nodes in the sequence. Conditional sequence nodes allow a model to be configurable: where conditional statements evaluate a parameter, the sequence enabled for a model's `Geometry` or `Mesh` can be set to change; and hence peripheral components

[1] Incidentally, so too does this Dissertation adapt its styling and content with if, else, and fi nodes in LaTeX.

can be included in or omitted from a common model file. Discharge resistors are usually placed throughout practical units and are associated with each element group rather than necessarily with the unit as a whole. Future refinements of this model could therefore include a set of options for discharge resistor location, or omit them entirely, given the limited volume each resistor might realistically require within a unit. Bushings and rack mounts are additionally provided on the `Geometry` of each unit for entirely aesthetic purposes. To emulate tab connections to each winding and internal connections between element windings and bushings (including any draw lead), subsequent selections can be facilitated by specifying `Domain`s to intersect those encapsulating foils in a layered approach, from which suitable boundaries can then be chosen.

Layered Approach

Some details of commercially available capacitor units are available online (for example in datasheets) and some details are not publically available. To make use of available information about unit size, and housing design, while also allowing winding designs to be reconfigured within a model, it is necessary to decouple internal aspects of unit design from the extremeties.

Additionally, capacitor models can be designed to scale to represent windings in greater detail (with potentially very thin layers of material), provided a `Mesh` sequence can also be scaled in tandem. For these reasons, a layered approach has been adopted in the developed model, whereby intermediate layers act as an interface intended to decouple an external layer (which can be specified according to available design information and have meshing appropriate to its size) from internal layers (which can be designed to change – for example, to test different winding arrangements – and which in detailed models must accommodate a narrow mesh).

In model provided in Section 3.1.1, `Domain`s are used to configure units with a set of four concentric layers, where a naming convention is chosen such that these can be referenced from `Geometry`, `Definitions`, and `Mesh` sequences in a reasonably consistent manner: 1. foilLayer; 2. growthGap; 3. dielectricFill; and 4. outerHousing. Naming conventions improve readability by creating 'slices' of a model that persist across otherwise distinct nodes of a `Component` tree. Thus, the source of an error can be located and remedied with comparative ease, irrespective of whether it is best addressed by changes in `Geometry`, `Mesh`, or another model component.

3.1.3 Mesh Sequences

A `Mesh` sequence is applied to the model `Geometry`, wherein a layered approach introduced in Section 3.1.2 aids efficient meshing. The way a `Mesh` is applied depends on how a geometry is finalised: every mesh respects `Domain` boundaries (`Mesh` points fall on these boundaries), but a boundary can also either delineate between a separate `Mesh` for the `Domain` on each side, or a `Mesh` can continue from one `Domain` to another (the same points on a boundary serve both adjacent domain meshes). `Form Union` and `Form Assembly` both allow objects to be combined into a single new object, where two overlapping objects would form a new object of three `Domain`s, the difference being in whether a mesh is continuous (`Form Union`) at `Domain` boundaries or discontinuous (`Form Assembly`). `Form Assembly` is appropriate for models with moving parts, for example. Capacitor models considered in this project use `Form Union` as only electrical and thermal effects are studied. Meshing is manually specified for model *scalability* by allowing a reasonably efficient mesh despite the possibility of narrowly spaced foils.

Table 3.2: Count of mesh elements by layer

	V120	V150
foilVolume	32019	17766
growthGap	159499	293414
dielectricFill	273743	268773
housing	3694	4575

Initial Structure

Thin components of the model `Geometry` (foils, and a thin outer housing) would lead to an unnecessarily inefficient number of `Mesh` elements should an automatic `Free Tetrahedral Mesh` be applied. To improve the quality of `Mesh` elements, reduce their number, and thereby reduce simulation execution times, a more efficient `Mesh` can instead be specified manually. While simplified geometries are relatively impervious to longer execution times imposed by less efficient meshing, it becomes important to use an efficient `Mesh` when a model scales to be representative of real assets with their associated detail. Hence, the `Mesh` for this capacitor model is designed to fit each of the four concentric `Geometry` layers:

1. foilVolume: an innermost section applies a `Mapped Mesh` to foil surfaces, which is `Swept` to fill `Domain`s between them to create a central array of anisotropic elements, sized according to user specified parameters;

2. growthGap: to expand a `Mesh` from narrowly-spaced elements on foils, to more isotropic and reasonably sized elements, a `Mapped Mesh` is defined with a scaling factor appropriate to the `Mesh` defined on foils, and is swept across each outward face of this layer;

3. housing: an outer housing is also thin (at 1.5 mm), and as such is conducive to a `Mapped` and `Swept Mesh`, similar to that used for foils to maintain overall `Mesh` efficiency;

4. dielectricFill: a `Free Tetrahedral Mesh` joins the `Swept` meshing on the growth layer to that on the outer housing.

Combined, these four `Mesh` 'layers' create a reasonable quality, *scalable* `Mesh` on the `Geometry`'s outer housing and central foils, between which a limited volume of free meshing is cultivated to provide continuity. The count of `Mesh` elements for two differently configured model versions (including V150 – the model file as published in Section 3.1.1) are given in Table 3.2, in which the influence of a `Swept Mesh` is apparent on reducing the number of elements in the foilVolume layer in relative terms. In these models, a `Swept Mesh` is similarly applied to the thin housing with an additional associated reduction in relative `Mesh` element numbers and hence model complexity. This approach is intended to reduce simulation time, permit increased numbers of foil elements, and adapt to user preferences for `Mesh` fidelity. Where possible:

1. `Mesh` nodes are named according to their layer to maintain readability; and

2. `Box Selections` are defined within the `Definitions` node based on the same parameters as the `Geometry` such that they can scale and adapt with possible parameter changes to resiliently select boundaries for mapping and `Domain`s for `Swept` or free meshing.

Table 3.3: Model materials

Component	Material	ϵ_r	Conductivity (Sm^{-1})
Dielectric	`Transformer Oil`	3.9	1×10^{-14}
Housing (& rack mounts)	`Steel AISI 4340`	1	4×10^{6}
Foils	`Aluminium 6063-T83`	1	3.030×10^{7}
Bushings	`Porcelain`	6	1×10^{-15}

Discussion

`Mesh` sequences do not support conditional nodes, so while such nodes can be used to change the characteristic nature of a `Geometry` sequence in response to select parameters, a manually specified `Mesh` sequence requires manual intervention in response to any such changes. Without support being included within COMSOL®, one intermediate solution could be to have sets of nodes within a sequence which can be enabled or disabled as necessary to accommodate each possible fundamental change in `Geometry`, as occurs for this model with different foil alignments.

It is good practise to repeat a `Study` with multiple meshes to verify conformance and ensure the arrangement of a particular `Mesh` (being a necessary intermediate stage of a process) does not itself have a significant bearing on results.

3.1.4 Materials

Capacitance is an accumulation of electric field over a dielectric (the primary element in any capacitor), so it follows that the most important materials in capacitor assets are those used as dielectrics. Strong dielectrics can withstand greater electrical field stresses to form thinner layers (allowing capacitor terminals to be placed closer together) or allow potential differences to be increased, and consequently unlock greater capacitance for a given unit rating and volume. Modern capacitor units use polypropylene films[2] (multiple films are combined according to design), and are permeated by a dielectric liquid to insulate capacitor terminals throughout. Mineral oil has historically been used to impregnate Kraft paper dielectrics [29] before these have been replaced with modern synthetic alternatives as materials advance.

However, available details on the chemistry of modern proprietary dielectric fluids seem to be limited, aside their attributes as being environmentally safer and less hazardous than conventional mineral oil dielectrics. In defining dielectrics used for simulation in this project, it is therefore assumed to be sufficient to approximate the relative permittivities available for synthetic-ester based dielectric fluids[3], in acknowledging that these vary (and can be readily amended in the model file). This is one area where the model stands to benefit from improved detail should characteristics representative of modern dielectric fluids be known.

Where available, materials are selected from built-in options and are summarised in Table 3.3. Steel is used for unit housing and rack mounts [58]; foils are aluminium; bushings porcelain; and the most critical material to any capacitor is the dielectric for which a relative permittivity $\epsilon_r = 3.9$ is applied (albeit using mineral oil as an in-built material option), and which can be amended via a user parameter to facilitate `Parametric Sweep` study types.

[2]Polypropylene film has relative permittivity $\epsilon_r = 2.2$ [67].
[3]As examples, dielectric fluid tradenames include Midel 7131 [81], Dielectrol [82], and Faradol [77].

3.1.5 Physics Interfaces

Once a model has a `Geometry`, `Mesh`, and `Materials` specified, physical effects and their interactions can be explored. COMSOL® offers a suite of modules which extend the core package to emulate various physical phenomena, which can be selected according to each user's interests. From the 'ACDC' module, `Electrostatics` (es), `Electrostatics with Boundary Elements` (esbe), `Electric Currents` (ec), and `Electric Circuit` (cir) interfaces are used in this project, as is the 'Heat Transfer' module's `Heat Transfer in Solids` interface for point heat sources.

Electric Circuit

`Electric Circuit` (cir) is a clean and straightforward interface which stands independent of the `Geometry`. It allows: 1. circuits to be specified from a range of electrical components; 2. signals to be generated within those circuits (including complex signals, such as the presence of harmonics); 3. thoses signals to be coupled to a `Geometry` using a selection of nodes discussed in Section 3.1.6; and importantly 4. measurement of electrical parameters at terminals as seen by the circuit (rather than position points on a `Geometry` to be measured less directly by a geometry-orientated `Physics` interface).

This interface (as illustrated by Figure 3.4) sufficies for developing simple circuits[4] which can be configured to couple with an established `Geometry` and thereby apply electrical conditions to geometrical models.

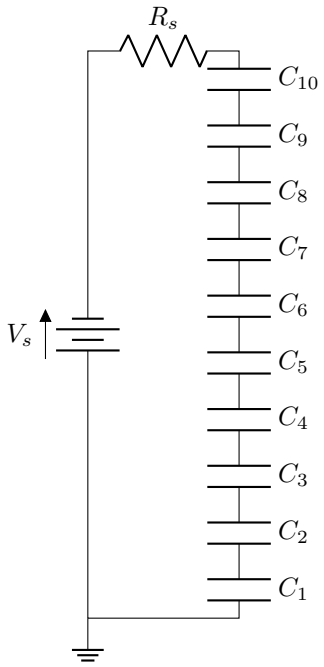

Figure 3.4: Coupling an electric circuit interface with a 10-element model

Electric Currents

An `Electric Currents` (ec) interface has been applied to the `Geometry` in early stages of model development to determine the distribution of voltage among the modelled windings subjected to an electrical signal applied via a coupled `Electric Circuit`. Equations for this interface are given as static cases of the equation of continuity, one form of Ohm's law, and how electric fields relate to conventional electrical potential (V), respectively as Equations 3.1, 3.2, and 3.3, for which $\mathbf{J}$ is current density (and $\mathbf{J}_e$ the current density of an external source), Q_j represents a current source, $\mathbf{E}$ is the electric field, and σ the electrical conductivity [83].

$$\nabla \cdot \mathbf{J} = Q_{j,v} \tag{3.1}$$

$$\mathbf{J} = \sigma\mathbf{E} + \mathbf{J}_e \tag{3.2}$$

$$\mathbf{E} = -\nabla V \tag{3.3}$$

[4]This interface could be more developer-friendly as connections must be manually specified between terminals of circuit components as numbered attributes, so it helps to draw the intended circuit out on paper in advance.

Electrostatics

`Electrostatics` (es) is the most appropriate interface for modelling capacitors as it is capable of modelling `Stationary` and changing currents as discussed in Section 7.3.2. It makes use of Equations 3.4 and 3.5, where ρ relates the space charge density, and $\mathbf{D}$ the electric displacement field, to determine the nature of electric fields in the model, again in the presence of coupling with an `Electric Circuit` interface via `Terminal` and `Floating Potential` nodes applied to each foil and the model housing. `Electrostatics` offers a measure of capacitance to be calculated by a `Global Evaluation` node in the results based on simulated electric fields, which further makes it an appropriate choice for the types of studies undertaken in this project.

$$\nabla \cdot \mathbf{D} = \rho_V \tag{3.4}$$

$$\mathbf{E} = -\nabla V \tag{3.5}$$

Electrostatics with Boundary Elements

`Electrostatics with Boundary Elements` (esbe) offers an alternative to Electrostatics which gives continuous rather than discretised results, although it generally takes longer to simulate. This offers greater detail for subtle electrical field effects, without the need to rerun a more precise `Study`, but at the expense of a greater computational burden, and otherwise allows simulation of the same electrical properties over a `Geometry` and its `Materials`.

Similar to `Mesh` conformance introduced in Section 3.1.3, having both `ES` and `ESBE` interfaces as options can allow a model to be studied with each of these interfaces to verify conformance in results. Equation 3.6 indicates how ESBE solves for electric field effects through a dielectric (defined by the permittivity of free space and relative permittivity of the dielectric, $\epsilon_0\epsilon_r$).

$$\nabla \cdot (-\epsilon_0\epsilon_r \nabla V) = 0 \tag{3.6}$$

Waveform Signals

Simple sinusoidal alternating current signals can be applied to the model by choosing the `Sine source` option within a `Voltage Source` node in the `Electric Circuit` interface and specifying appropriate parameters. In this manner, such sources can then be combined in this interface to synthesise harmonics, provided these signals remain consistent in time. However, AC signals are not limited to sinusoids and a variety of waveform types can be included by first establishing them as a `Waveform` function added to the `Global Definitions` node. These can then be used within a `Voltage Source` by setting it up as an `AC Source` and specifying the waveform and a time parameter (for example `wv1(t)`). It is necessary to use an appropriate `Study` time to make use of signals which change in time.

Heat Transfer

Heat Transfer interfaces allow temperatures (T) and heat flux ($\mathbf{q}$) to be simulated given material properties (density ρ, thermal conductivity k, and specific heat capacity at constant pressure C_p), convection ($\mathbf{u}$), and thermal powers (Q) including a thermoelastic damping heat source (Q_{ted}) within a model [84]. Thermal conductivity and heat transfer coefficients can be defined, particularly on outermost boundaries as part of a `Heat Flux` node to represent heat dissipation

from the model. Heat generation can be applied throughout a domain (to represent dielectric losses) and as `Point Source`s which are used in Chapter 5 to represent incipient breakdown such as partial discharge. `Heat Transfer in Solids` solves for Equations 3.7 and 3.8, while a similar interface, `Heat Transfer in Solids and Fluids`, solves the same equations but omits the thermoelastic damping term Q_{ted}.

$$\rho C_p \mathbf{u} \cdot \nabla T + \nabla \cdot \mathbf{q} = Q + Q_{ted} \tag{3.7}$$

$$\mathbf{q} = -k\nabla T \tag{3.8}$$

3.1.6 Means of Coupling

The `Electric Circuit` interface offers a suitable way to monitor the electrical behaviour of a model, establish values of a power supply, and to acquire measurements into tables of results. To facilitate interaction with a 3D model, three nodes in the `Electric Circuit` interface permit coupling with `Electrostatics` or `Electric Currents` interfaces [83], illustrated in Figure 3.5 in which solid lines indicate explicitly specified associations, and dotted lines (and brackets) used to denote unspecified associations which are instead inferred by interfaces, and are useful to keep in mind to limit confusion in coupling interfaces together.

External I vs. U

Each `External I vs. U` node acts as a voltage source in the electric circuit. It takes two nodes in the circuit and acts as a voltage source between them, reflecting the value of an arbitrary voltage connected from elsewhere (such as a voltage measured in another physics interface). It thus can be thought of as having one explicit input (a measured voltage), one implicit input (the ground implied by that voltage), and two outputs (the points it connects into a circuit).

External U vs. I

Similarly, each `External U vs. I` node allows coupling between interfaces, but acts as a current source between two points specified in the electric circuit, reflecting an arbitrary value of current

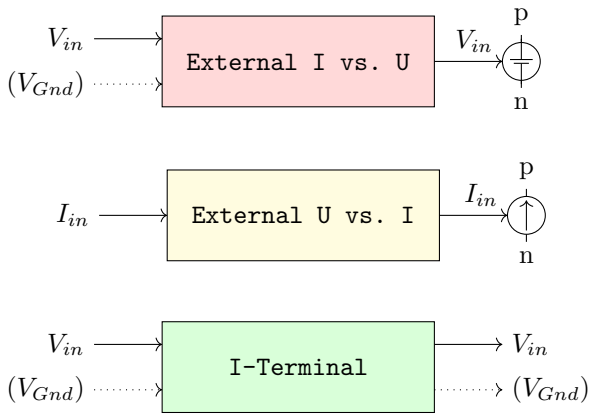

Figure 3.5: Nodes to couple geometrical (left) to Electric Circuit (right) physics

which represents one measured elsewhere (such as a current in another physics interface). Such nodes therefore have one input (a current), and two outputs (points it connects in a circuit).

I-Terminal

Alternatively, `I-Terminal`s take only one explicit input (a potential difference), an implicit input (a ground the potential is relative to), and one output (a node in the circuit) to which the voltage of a `Terminal` in another physics interface is coupled as a voltage in the circuit.

3.1.7 Study Configurations

`Study` nodes are associated with the *root* node of a model tree and are used to configure solvers and simulation approaches. As a means of control over simulation approaches, these allow techniques which are efficient for a particular problem to be specified, to therefore lessen the time required for each study. Otherwise, if a solution can be estimated, setting this estimate as `Initial Values` within physics interfaces can significantly reduce simulation times. 'Sweep' investigations can also be configured here to explore relationships between results and changes to model parameters, material properties, and signals applied. In this project, `Stationary` studies are predominantly used, as effects explored in Chapters 4 and 5 relate primarily to asset designs and geometry rather than changes in time. These offer straightforward `Study` steps and are a sensible way to sanity-check a model before more extensive and potentially time-consuming studies are performed. `Time Dependent` studies simulate models which change in time. This ability can be useful for future studies, which in particular could explore behaviour related to harmonics or other signal artefacts. `Iterative` and `Direct` solvers can be enabled to best suit the problem: `Iterative` solvers place less of a burden on computational resources but are less robust to nuanced solution spaces than are `Direct` solvers. Particular considerations (such as choice of solver algorithm and solution tolerance) can be configured for each. `Study` nodes store solutions on completion, so that they can be used to present results through postprocessing.

3.1.8 Post-processing

Results can be graphically displayed by post-processing, wherein straighforward `Cut Plane` and `Cut Line` visualisations are mostly used for voltage and thermal distributions in this project. Various visualisation techniques are available to display solutions, however, as Figure 4.12 illustrates. `Time Dependent` study results can also be post-processed into animations.

3.2 Discussion and Modelling Trade-offs

3.2.1 Parameterisation

Further to readability as discussed in Section 3.1.2, a more readily apparent benefit to model parameterisation is *configurability*. Tables 3.4[5], 3.5, and 3.6 respectively list parameter descriptions, set values, and expressions for model files used. These Tables reference file versions introduced in Table 3.1, and Table 3.6 exhibits duplication where parameters have been deprecated in exchange for others as the model evolved. Parameters allow a capacitor model to adapt to match different unit sizes and foil configurations, alongside mesh and dielectric conditions, all from a common point within the environment user interface. This ease of reconfigurability can in turn facilitate comparisons between unit designs (both in terms of size and internal winding arrangements), material properties, and meshing. On the other hand, while parameters allow a model to be more easily specified, references to each parameter must permeate the file for this to be effective. For instance, a `Geometry` can be recalibrated by amending a unit's height parameter, but only where all appropriate parts of a `Geometry` make equivalent reference to a common 'height' parameter. This is true not only for a `Geometry` but for other model aspects: a `Mesh` sequence must also refer to the same parameter in order to appropriately mesh a changed geometry without need for manual intervention following each change.

Furthermore, COMSOL® has recently introduced an 'Application Builder' feature which allows a model to be packaged and shared, where parameters can be amended and a study repeated, beyond the host (licensed) computer. Hence, `Definitions` are an appropriate way to incorporate parameters for reference by other parts of a model file.

3.2.2 Definitions

While not strictly necessary, `Definitions` are similar to parameters as a form of good-practise. Although model `Definitions` take time to set up, in general they: • improve model readability; • allow the model to be readily adapted, reconfigured or adjusted; and • are extensible.

Definitions used for the model introduced in this Section 3.1 attempt to follow consistent naming conventions (based on those in Section 3.1.2), and to some extent are based on the same parameters as the model `Geometry`, such that these `Definitions` (and in turn the `Materials`, `Mesh` sequence nodes, and physics interfaces which use them) can conform to any change in parameters and hence make the model more easily configurable. Increasingly representative models (and the number of windings and winding turns which comprise them) necessitate the use of `Definitions` to collect related boundaries for winding groups in an efficient manner. The remainder of this Section 3.2.2 discusses techniques used to this end.

A Repeated Geometry Sequence

`Array` nodes can be used to repeat a `Boundary`, `Domain`, or `Object` within a `Geometry`, but cannot change attributes with each repeated instance. To create a series of geometrically similar winding turns, but with increased turn radii each time, it is instead necessary to use a common formula to specify each individually. A lengthy sequence of conditional statements (as shown in Figure 3.6) can instead test for each increment of a parameter `turns`, with support (presently) for 40- and 10-turn maxima for windings respectively aligned with the X- and Z-axes.

[5] 'Can' is an alternative term for a single capacitor unit.

Table 3.4: Parameters descriptions

Parameter	Description
canWidth	The width of the can (x axis).
canDepth	The depth of the can (y axis).
canHeight	The height of the can (z axis), selected from the ABB Capacitor's Buyer's Guide.
housingWidth	The width of material used for the outermost housing cannister (in reality this is part of the overall dimensions).
canWidthOffset	
canDepthOffset	
canHeightOffset	The offset of foils from the base of the can.
canWidthHalfGap	The gap between foils widthwise (x axis).
canDepthHalfGap	This was 44.8[mm] for alignment 0.
canHeightHalfGap	
dischargeResistorHeightOffset	A height specified to allocate space for wiring and a discharge resistor between the terminals.
stackLimit	The maximum size of the stack.
stackHeight	
stackWidth	The number of elements stacked widthwise (along the x axis).
foilStackDirectionZ0X1Y2	This decides the direction of foils (stacked horizontally, vertically, or depthwise).
foilWidth	The width of each foil.
foilDepth	
turns	This is the number of times a foil is turned for each element.
gapsPerElement	This is the number of dielectric gaps for each element.
foilThicknessZ	The thickness of foils as calculated by dividing the available space by the number of elements specified in stackHeight.
foilThicknessX	
foilThicknessY	
foilThickness_static	The thickness of foils used for stackLimited geometries (quick, simple simulations).
foilCornerFoldRadius	Was: (canDepth+canDepthOffset*-2)/6
bushingDistance	The distance between the centres of each bushing.
rackMountHeight	The height of the handles on the outside of the can, according to the ABB Capacitor's Buyer's Guide.
rackMountDepth	The width of each handle on the outer casing.
foilMeshDepthwise	The number of mesh points depthwise on flat faces of each foil.
foilMeshWidthwise	The number of mesh elements widthwise on each flat foil element.
voltageSourceAmplitude	The magnitude of the voltage source.
eRelativePerm	Relative Permittivity for purposes of parametric sweep.
limitedModuleGrowthGapDivisions	The number of segments that the growth gap has in limited-element models.
foilHeight	
connectorThickness	This is always best to be less than half the distance of a foil along the X Y or Z direction depending on its alignment.
connectorGroupCount	How many elements are in each series section?
tempNominalDivisor	to divide foilWidth for diagonal connectors.
rackMountBase	
connectionDivision	The number of times to divide the central volume to create paths for connections.
foilZPosition	
foilBaseZ	
minor	A minor offset so box definitions encapsulate the designated space.

Table 3.5: Parameter values used in versions of the model file

Parameter	V40-6	V81-{5, 6-4, 7}	V93-C1	V105-11	V150
canWidth [mm]	403	403	172.5	172.5	172.5
canDepth [mm]	178	178	87.5	87.5	87.5
canHeight [mm]	1140	1140	357.5	357.5	357.5
housingWidth [mm]	1.5	1.5	0.75	0.75	1
canWidthOffset [mm]	20	20	10	10	15
canDepthOffset [mm]	15	12	10	10	15
canHeightOffset [mm]	30	30	10	10	15
canWidthHalfGap [mm]	4	6	5	5	5
canDepthHalfGap [mm]	44.8	25	15	15	18
canHeightHalfGap [mm]	8	8	5	5	5
dischargeResistorHeightOffset [mm]	40	60.6	0	6	6
stackLimit	0	0	0	0	0
stackHeight	10	18	5	10	4
stackWidth	1	1	5	1	5
foilStackDirectionZ0X1Y2	–	{0, 1, 2}	1	0	1
turns	–	4	2	4	40
foilThickness_static [mm]	0.5	0.5	0.5	0.5	0.5
foilCornerFoldRadius	0	0	0	0	0
bushingDistance [mm]	220	220	90	90	90
foilMeshDepthwise	2	3	3	3	3
foilMeshWidthwise	3	3	6	6	6
voltageSourceAmplitude	100	100	100	100	100
eRelativePerm	3.9	3.9	3.9	4	1.4
limitedModuleGrowthGapDivisions	8	1	1	1	1
connectorGroupCount	–	6	5	5	5
tempNominalDivisor [m^2]	–	0.1	1	1	1
rackMountBase [mm]	–	–	10	10	10
connectionDivision	–	–	3	3	3
minor [m]	–	–	–	–	0.00001

Table 3.6: Calculated parameters used throughout model development

Parameter	Expression
foilWidth	(canWidth-2*housingWidth-2*canWidthOffset-2*canWidthHalfGap*stackWidth)/stackWidth
foilDepth	canDepth-2*housingWidth-2*canDepthOffset-2*canDepthHalfGap
foilHeight	canHeight-2*(housingWidth+canHeightOffset+canHeightHalfGap)-dischargeResistorHeightOffset
rackMountHeight	canHeight/2
connectorThickness	(canDepth-2*housingWidth-2*canDepthOffset-2*canDepthHalfGap)/connectionDivision
gapsPerElement	2*(turns+1)
foilThickness	(canHeight-2*housingWidth-2*canHeightOffset-2*canHeightHalfGap-dischargeResistorHeightOffset)/(stackHeight*10+3)
foilThicknessZ	(canHeight-2*(housingWidth+canHeightOffset+canHeightHalfGap)-dischargeResistorHeightOffset)/(stackHeight*2*gapsPerElement)
foilThicknessX	(canWidth-2*housingWidth-2*canWidthOffset-2*canWidthHalfGap)/(stackHeight*2*gapsPerElement)
foilThicknessY	(canDepth-2*housingWidth-2*canDepthOffset-2*canDepthHalfGap)/(stackHeight*2*gapsPerElement)
rackMountDepth	canDepth-2*housingWidth-2*canDepthOffset
foilZPosition	housingWidth+canHeightOffset
foilBaseZ	housingWidth+canHeightOffset+canHeightHalfGap

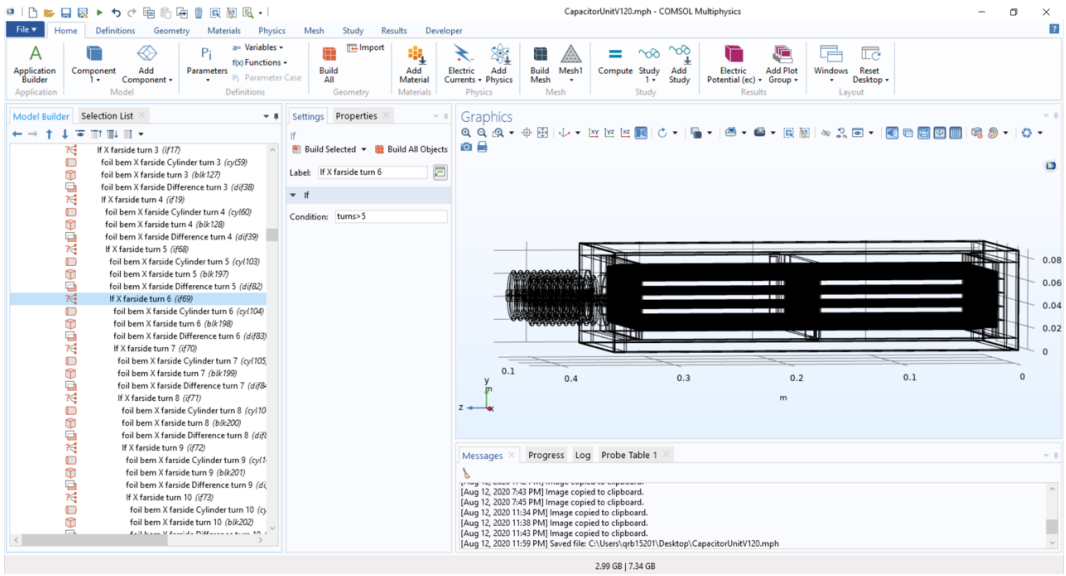

Figure 3.6: A sequence of repeated conditional geometry nodes

Choosing an Innermost Terminal Layer

For this approach to work, even when different axes and arrangements are used, it is best to understand at least how to choose a first boundary of a winding terminal, and then the rest of the element winding can be selected easily through a `Definition` sequence for that element. Figure 3.7 further illustrates how this works for a changed foil alignment, with windings aligned to the Z-axis. This `Box Selection` configuration will select foil 20, as indicated in Figure 3.7.

- The simplest axis is the one that the foil turns extend on. For the image in Figure 3.7, this is the y axis, which uses a Y minimum of `housingWidth+canDepthOffset` and a maximum of `canDepth-(housingWidth+canDepthOffset)`. As foil turns occupy the 'growthGap' layer, it is enough to limit the box selection to layers beyond this one provided an (`Entity` or `All vertices`) `inside box` option is chosen for the `Output Entities` option, since any turn designed to connect between two layers will extend beyond any `Box Selection` that selects a first layer.

- For the axis that terminal boundaries are stacked along (e.g. the Z axis in Figure 3.7) it is possible to pick a single boundary using a minimum of: housingWidth+canHeightOffset +canHeightHalfGap+(gapsPerElement)*19*foilThicknessZ and a maximum of housing-Width+canHeightOffset+canHeightHalfGap +(gapsPerElement+1)*1*foilThicknessZ. An integer between `gapsPerElement` and `foilThickness` can be used to select the n^{th} winding. `gapsPerElement = 2×turns+1` gives a count of the number of spaces between layers on a given winding.

- The remaining (X) axis is more complicated due to the possibility of multiple stacks:

  - x minimum = housingWidth+canWidthOffset+canWidthHalfGap*1 + 0*(foilWidth-0*canWidthHalfGap)/stackWidth

  - x maximum = housingWidth+canWidthOffset+canWidthHalfGap*1 + 1*(foilWidth-0*canWidthHalfGap)/stackWidth.

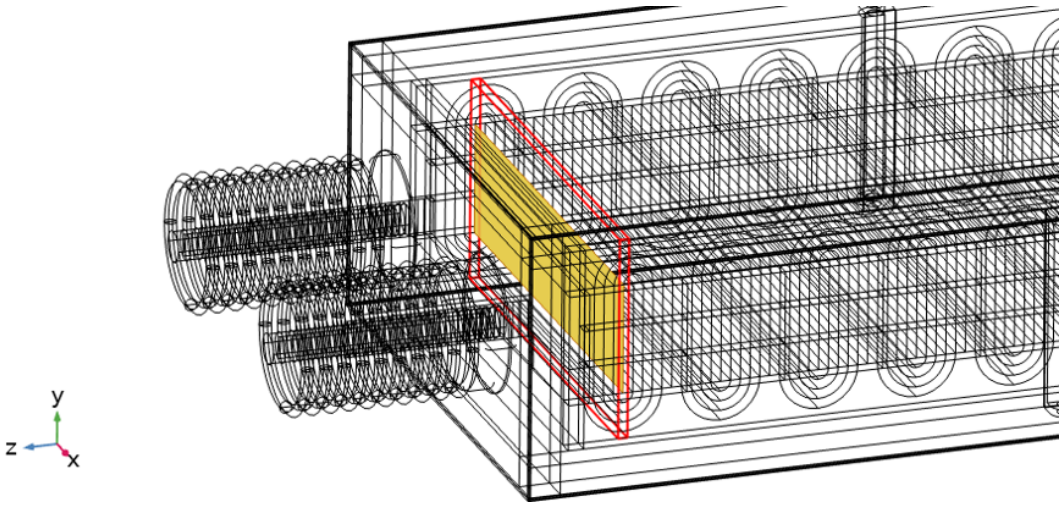

Figure 3.7: A single boundary selected from the specified settings

Figures 3.8 and 3.9 further illustrate how an initial boundary is chosen for X-axis aligned windings, where Figure 3.9 shows a `Definition` stop shy of a margin in an interleaved design, and the context of making this initial selection for the first node of an element `Definition` sequence. In all cases the objective is to encapsulate exactly one boundary with a `Box Selection` (red), such that remaining boundaries in the same element will be selected by reference from the `Definition` sequence. This minimises manual effort in defining all boundaries in a winding.

In Figure 3.7, winding terminals are defined with boundaries which are distinct from one another, so an automated process can easily differentiate between windings. However, in practise adjacent elements are often electrically connected, and as such, both boundaries which are to be electrically connected can be selected by the same `Box Selection` as the first of a definitions sequence to reduce its overall length in defining a connected pair of adjacent windings.

The dimension used to select an innermost boundary depends on winding alignment, and where an expression in terms of the `gapsPerElement` parameter is added to the maximum and minimum fields for the X, Y, or Z parameter to correspond with windings which are stacked in that direction as specified by parameter `foilStackDirectionZ0X1Y2`. The multiplier, k, used in this term is $k = 2n - 1$ for the n^{th} winding in the chosen direction.

In more general terms, while the model is not entirely configurable (as dictated by bounds on the capabilities of COMSOL®), it is easier to adapt at least in terms of the `Geometry` and `Mesh`. Errors and warnings do arise in compiling a `Mesh`, but since this part of the model broadly depends on `Definitions`, such issues are often resolved through amending nodes in the `Definitions` subtree of the `Model Builder`. It is therefore necessary to check the `Definitions` a model uses following a reconfigured `Geometry`, and while it is helpful to have a `Mesh`, `Materials`, and other model aspects which adapt with updates made to `Definitions`, this remains a cumbersome process inherent to each reconfiguration.

Spacers Direct Expansion to Adjacent Boundaries

After selecting an initial winding boundary, remaining boundaries can be selected by repeatedly:

1. selecting those adjacent to the current selection with `Adjacent` selections; and
2. subtracting any not part of the chosen winding using `Difference` nodes.

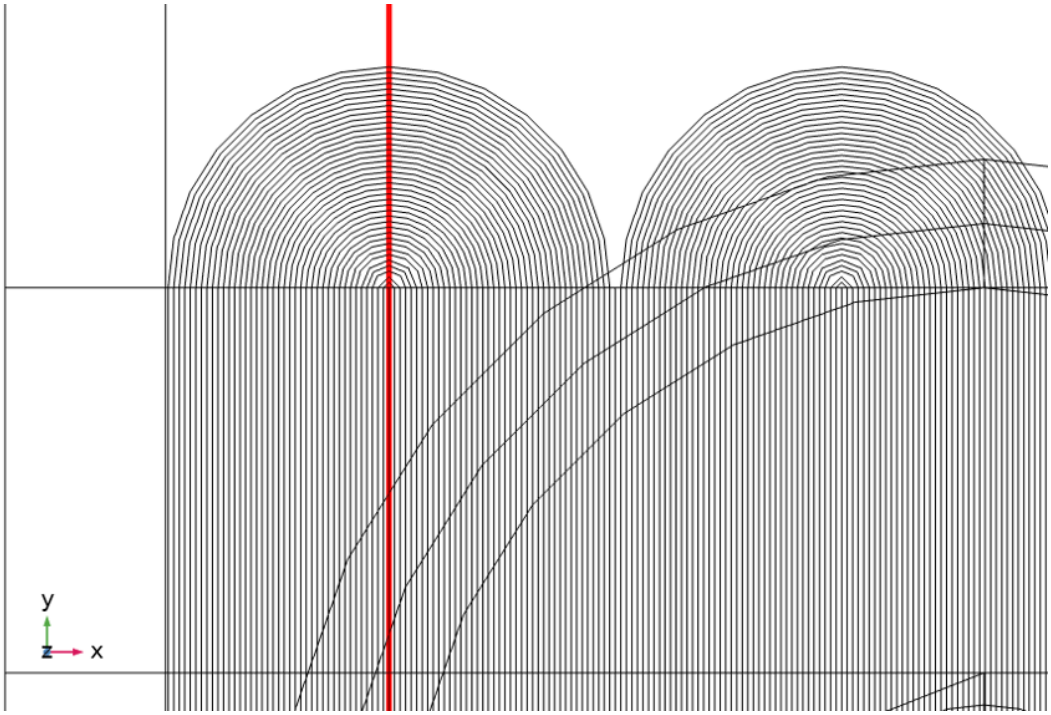

Figure 3.8: Selecting an innermost winding boundary

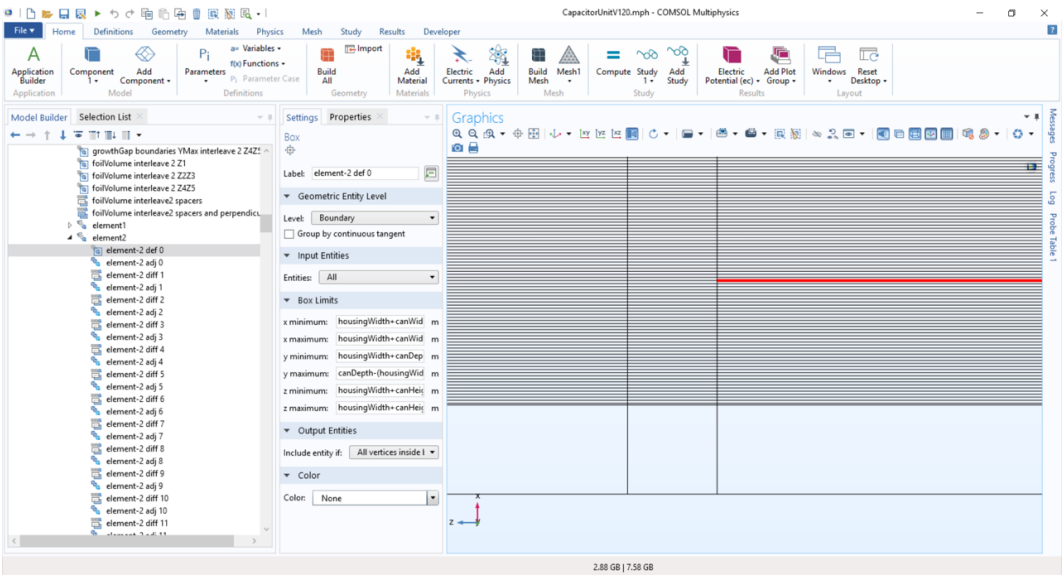

Figure 3.9: Selecting an innermost winding boundary with a box selection

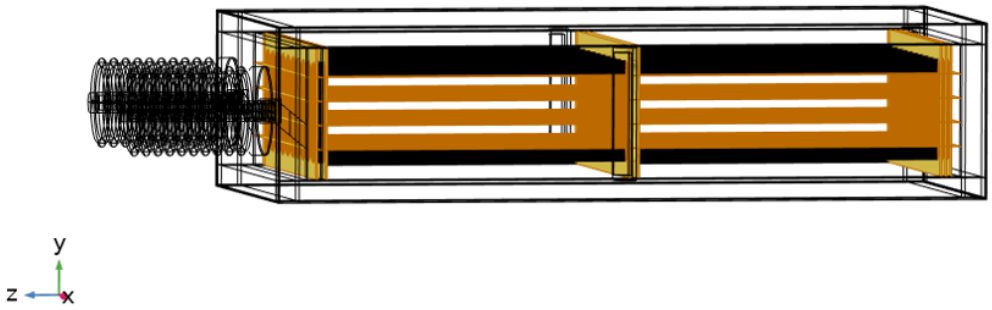

Figure 3.10: Boundaries defined as spacers (shown in yellow)

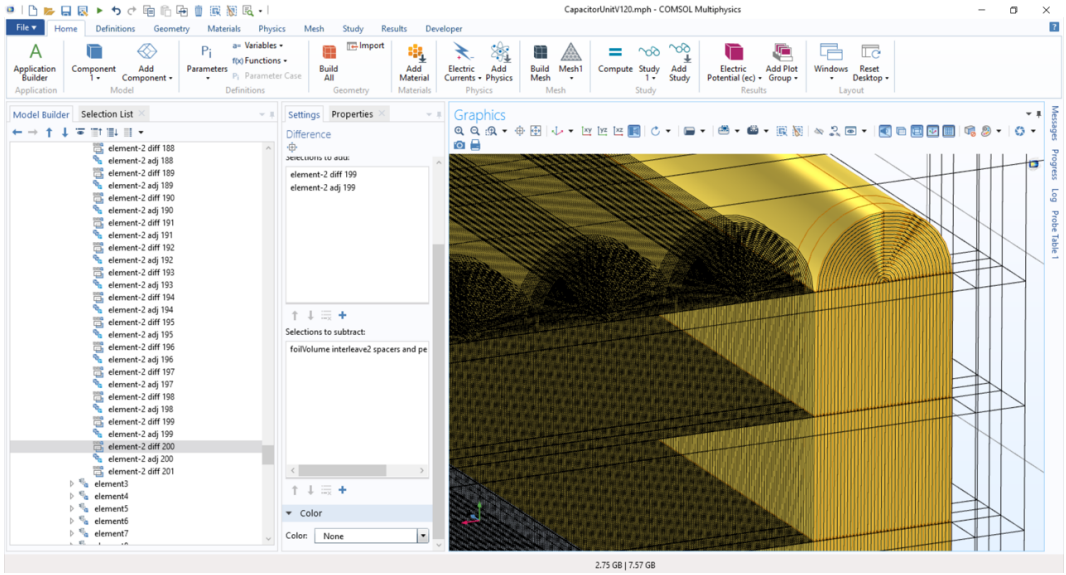

Figure 3.11: A defined winding

This second item requires 'spacers' to be defined as boundaries explicitly *not* a possible part of a winding, such that these can be subtracted in point 2. In yellow, Figure 3.10 shows spacers chosen for an X-axis foil alignment: all boundaries adjacent but perpendicular to foil layers.

A Resulting Foil Definition

Figure 3.11 shows the collection of boundaries contained within a single `Difference` definition at the end of a repeated sequence. Such long sequences appear cumbersome within the simulation environment, but once established these can be amended or removed in the knowledge that they can be recreated with relative ease from a model method. This is a time-efficient and readily-reconfigurable way to accurately define all boundaries in a winding of many turns.

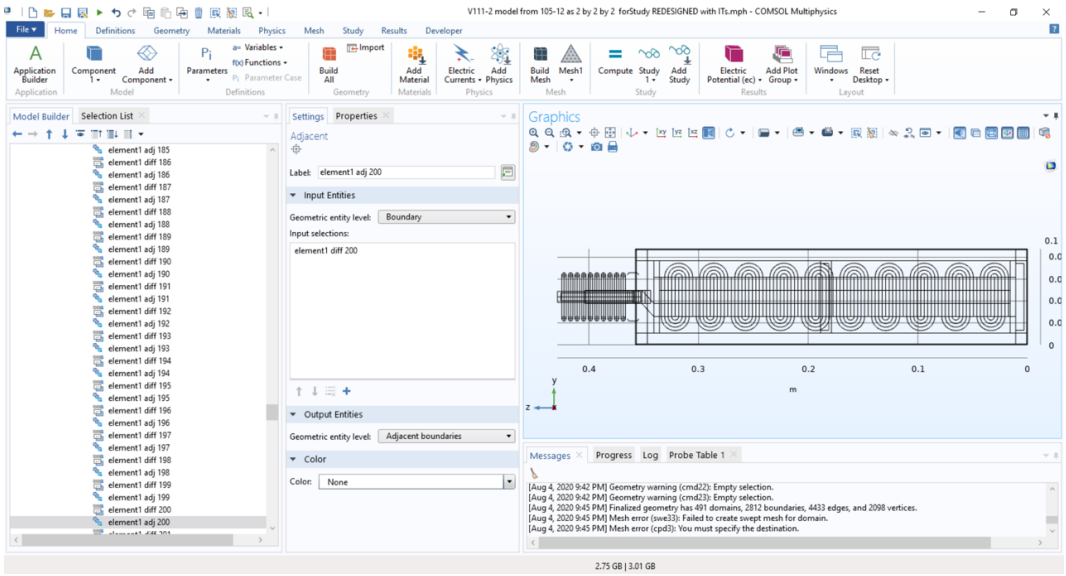

Figure 3.12: Screenshot of a definition sequence generated by a method in subsection 3.2.3

3.2.3 Model Methods

Methods are a powerful feature within COMSOL® and which can be used to automate routine tasks as a form of scripting using Java. As introduced in Section 3.2.2, `Definitions` allow a model to be reconfigured with relative ease; it can take a long time to manually select every boundary comprising a winding of more than only a handful of turns; and to establish `Definitions` for each foil layer from adjacent boundaries, it is possible to use a combination of `Adjacent` and `Difference` nodes.

Figure 3.12 shows a sequence of such nodes which has been generated by the Java code below. This short code makes it extremely easy to repeat a pair of nodes in the `Definitions` branch of the model, such that alternate `Adjacent` and `Difference` nodes can be chained together, each referencing the node before, in order to begin from a single boundary specification and repeatedly extend it to all adjacent boundaries aside those defined as being explicitly outwith the element winding.

This code simply copies existing nodes, so an initial `Box` definition, `Adjacent`, and `Difference` node is needed for each new element sequence (preferrably grouped to improve navigation). Specific index numbers for the `Adjacent` and `Difference` nodes to be repeatedly duplicated can be obtained from the 'Tag:' field of the 'Properties' tab for these nodes. The 'AdjacentSequenceExtender' method can then be edited from the 'Edit Method' option in the 'Method Call' settings window, so that can be specified within the source code before it is run: 1. a label (a) name and (b) number; 2. an index for each of the two original (a) `Adjacent` and (b) `Difference` nodes in the sequence; and 3. the number of total desired repetitions.

Method Code for Extended Element Sequences

```
/* This script generates long sequences of adjacent boundaries, designed to allow
 * foil elements to be automatically defined as sets of (many) adjacent boundaries
 * given one boundary is manually configured with a box selection to begin the set. */
```

```
/* It's important to update these 5 parameters before we run this method each time. */
int index = 2; //This is the label number of the last difference selection node.
int adjIndex = 5839; //This is the tag number of the last adjacent selection node.
int difIndex = 5877; //This is the tag number of the last difference selection node.
int repetitions = 200; //The number of repetitions (how long the sequence should be).
String naming = "element10"; //Something to differentiate the sequence from others.

for (; index <= repetitions;)
{
  model.component("comp1").selection().duplicate("adj"+(adjIndex+1), "adj"+adjIndex);
  model.component("comp1").selection("adj"+ ++adjIndex).label(naming+" adj "+index);

  with(model.component("comp1").selection("adj"+adjIndex));
    set("input", new String[]{"dif"+difIndex});
  endwith();

  index++;

  model.component("comp1").selection().duplicate("dif"+(difIndex+1), "dif"+difIndex);
  model.component("comp1").selection("dif"+ ++difIndex).label(naming+" dif "+index);

  with(model.component("comp1").selection("dif"+difIndex));
    set("add", new String[]{"dif"+(difIndex-1), "adj"+(adjIndex)});
  endwith();
}
```

Discussion

This model remains far from being truly easy to reconfigure as it requires manual intervention to recalibrate `Definitions` in response to changed parameters (and hence changes in a `Geometry`). The efforts thus far are simply an attempt to limit the amount of intervention necessary, and in the hope of eventually being able to further minimise any such manual configuration. Further work could make more comprehensive use of `Box` selections (which in turn refer to parameters common to the `Geometry` elements being targeted) such that reconfigurability can be improved.

3.2.4 Indicative Configurations

Figure 3.13 illustrates a model specified by the the parameters in Table 3.7, and with windings of 20 turns. The configuration of turns can be more readily seen with a view parallel to the winding turns, as for this model in Figure 3.14, and in Figure 3.15 for a model with 40 turns. As provided in Section 3.1, the model file can accommodate winding alignments in each direction and for each alignment, a number of winding turns $t \leq 40$.

Table 3.7: Model Configurations

Parameter	Value
`canHeight` =	357.5 mm
`canWidth` =	172.5 mm
`canDepth` =	87.5 mm
`housingWidth` =	0.75 mm
dielectric offset (all axes) =	10 mm
`foilStackDirectionZ0X1Y2` =	0 (Z)

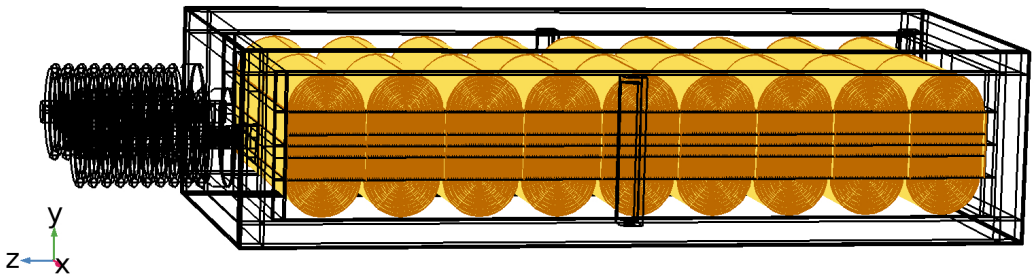

Figure 3.13: 20 turn model wireframe

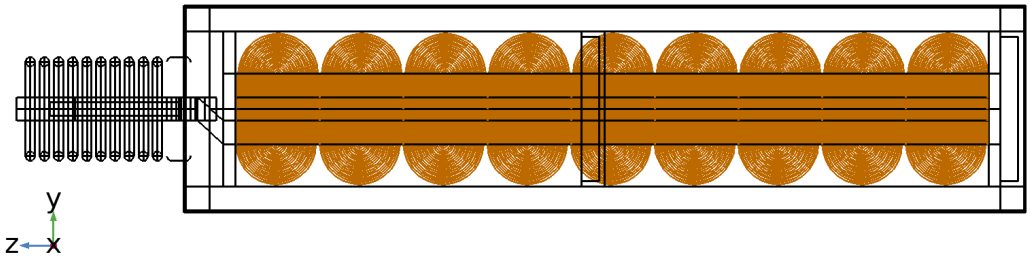

Figure 3.14: 20 turn model

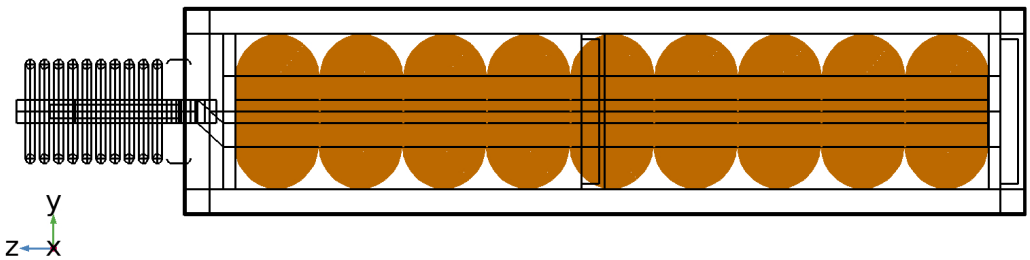

Figure 3.15: 40 turn model

Chapter 4

Electrical Effects

The model presented in Section 3.1 is used to generate results which in this Chapter are presented and discussed in terms of electrical effects. This model is reconfigured to evolve in complexity for each Section. A simple 10-element model forms a basis to explore the distribution of voltage within a unit's internal windings, and how series element failures affect this, in Section 4.1; a more sophisticated 25-element interleaved design reveals, in Section 4.2, where electric field stresses arise; before Section 4.3 extends the model to a more complex case of multiple units on a shared rack, to explore the relationship between unit position and potential differences.

4.1 Effects of Failure on Voltage Distribution

This Section considers the impacts of series element failure in fuseless capacitor units, where an element fault forms a short-circuit which removes its group from the set of elements in service.

Where this Section considers faults, it explores locations for individual elements faults. In other words, only a first fault to arise on a previously healthy capacitor unit is considered: the effects of a subsequent failure occurring on a unit which has already experienced a fault are not. Similarly, fault conditions outwith capacitive elements remain unexplored, such as where a fault might arise between terminals of adjacent windings, or to a draw lead or housing.

4.1.1 Simulation Studies and Results

The model presented in Figure 3.1, with ten 'windings' of 3D foils arranged in series along the Z axis of the model `Geometry`, has two boundary conditions configured in support of simulation studies: one represents electrical connectivity between a signal source and elements; and another simulates the behaviour of the asset model given its specified `Geometry` and `Materials`. A 50 Hz, 10 kV supply signal at 1 Ω impedance is applied[1], which provides initial results: voltages throughout the `Geometry` (and at each foil) are shown in Figure 4.1 and Figure 4.2[2], and relative to a linear distribution in Figure 4.3. Table 4.1 lists voltages at each foil, for 11 scenarios to

[1] These arbitrary, nominal values are chosen to be representative of units sized for HVDC link AC filtering.

[2] In Figure 4.2, 'Arc length (m)' is the term used to illustrate the length of a cut line running from one end of a unit to the other, straight through its centre in the same direction as in Figure 4.1. Synthetic noise on this line graph resembles gear 'teeth', which partly show variation in electric potential between electrodes of each capacitive element, and are in part due to these terminals having non-zero thickness in this simple early model.

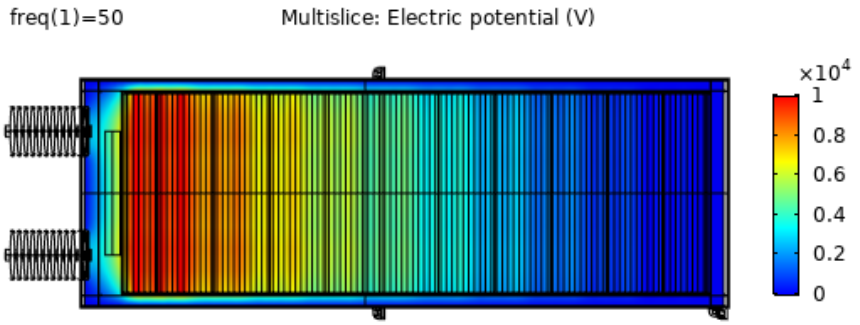

(a) Voltage over an XZ-plane cross-section

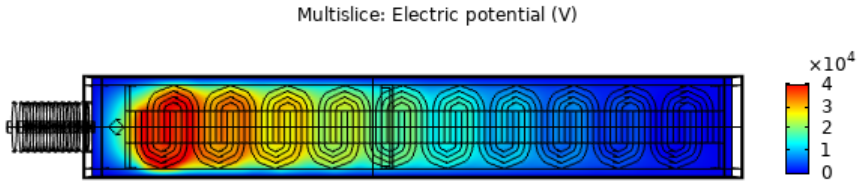

(b) Voltage over a YZ-plane cross-section

Figure 4.1: A distribution of voltage over simple unit cross-sections

show how any short-circuit between foils can affect voltage. `Terminal`s 0 and 10 (T0, T10) consistently hold 0 V and 10 kV, respectively.

4.1.2 Discussion and Implications

Each row of Table 4.1 corresponds to a single short-circuit scenario, beginning with an absence of any short-circuit and where each pair of adjacent foils is short-circuited in turn. Each column corresponds to voltages at each of ten discrete foils in the model. The results are provided graphically in Figure 4.1 and Figure 4.2, which show the distribution of voltage within a unit without a short-circuit to be nonlinear, as is more clearly illustrated in Figure 4.3.

Since there is variation neither in the `Geometry` between any two foils nor in the dielectric

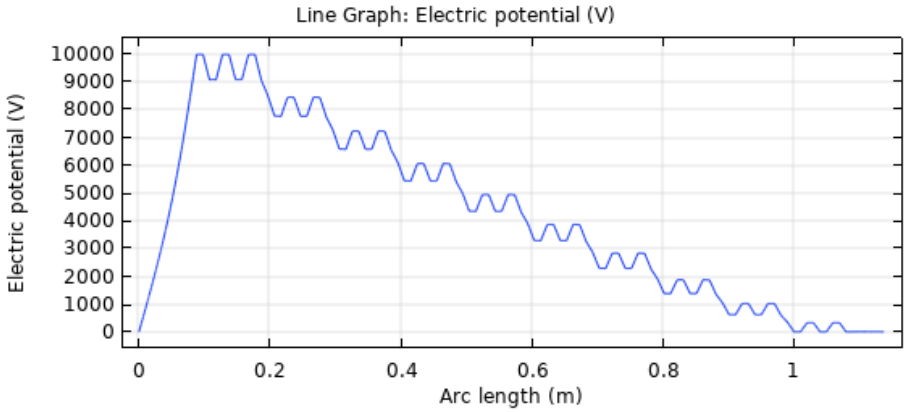

Figure 4.2: Voltage over a cross-section and cut line of a simple unit

Table 4.1: Short-circuit foil voltages at 50 Hz

Short Circuit	T1	T2	T3	T4	T5	T6	T7	T8	T9
None	454.39	1176.5	2062.3	3045.1	4089.3	5172.6	6299.3	7477.6	8703.8
T0-T1	0	504.84	1317.1	2313.8	3424.9	4607.6	5855.4	7170.6	8544.8
T1-T2	451.04	451.04	1238.6	2228.3	3343.1	4536.0	5798.2	7130.7	8524.0
T2-T3	468.74	1171.1	1171.1	2122.7	3228.0	4428.7	5709.7	7067.7	8490.9
T3-T4	496.99	1254.9	2119.8	2119.8	3183.8	4370.9	5654.9	7026.1	8468.3
T4-T5	521.86	1329.6	2276.6	3247.2	3247.2	4386.3	5649.2	7015.0	8460.6
T5-T6	539.91	1384.0	2391.0	3456.4	4497.6	4497.6	5705.8	7041.7	8471.1
T6-T7	553.67	1424.9	2476.0	3610.1	4758.6	5854.0	5854.0	7121.3	8506.1
T7-T8	564.07	1454.9	2536.2	3715.3	4931.6	6131.2	7274.2	7274.2	8574.7
T8-T9	570.69	1473.6	2573.0	3778.1	5032.3	6289.0	7521.1	8687.0	8687.0
T9-T10	573.45	1481.5	2589.0	3805.8	5077.8	6361.7	7637.0	8872.6	10000

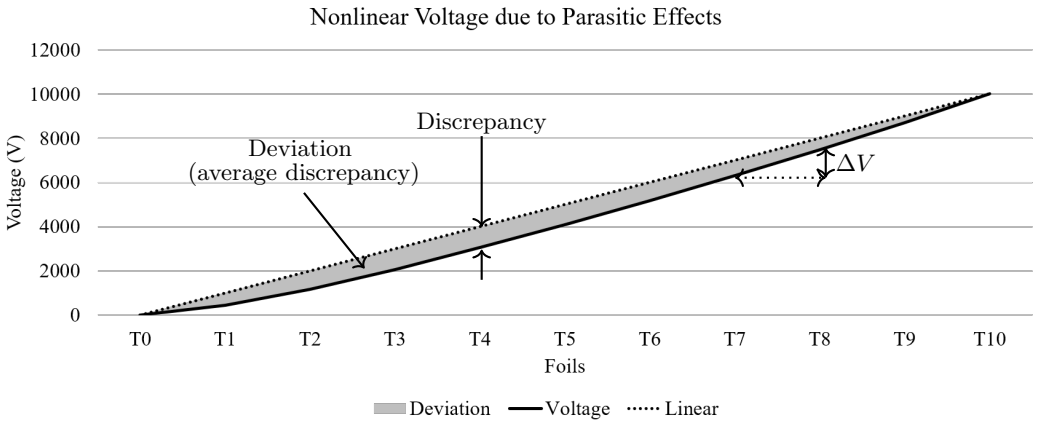

Figure 4.3: Nonlinear voltage due to parasitic effects

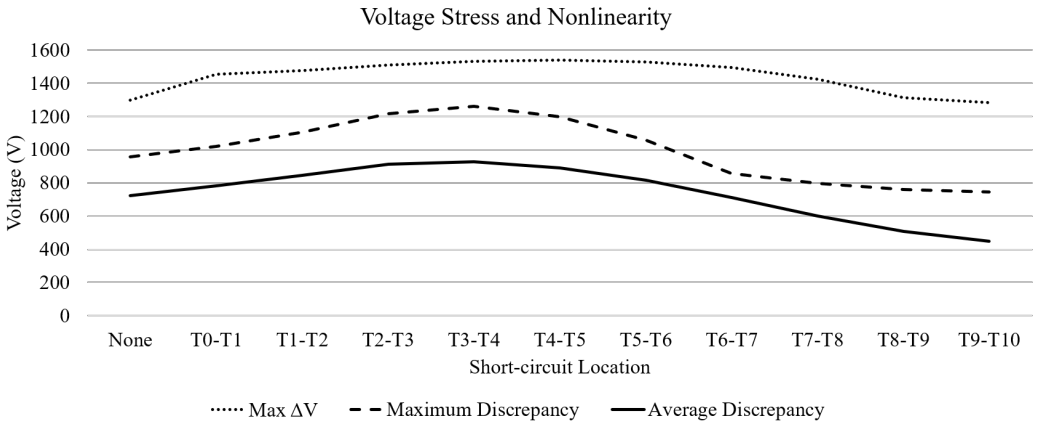

Figure 4.4: Voltage stresses between foils and deviation from a straight line

permittivity throughout the model, this nonlinearity must intrinsically result from the capacitor unit geometry. While capacitances arise between adjacent foils, as expected, a potential difference between each foil and a 'grounded' unit housing also invites a parasitic capacitance: one that forms between each terminal and ground in parallel with all lower-voltage foils and their own associated stray capacitances which accumulate as element groups are combined in series; this cumulative effect appears to provoke a subtle nonlinearity in the distribution of voltage shared due to the effect on overall capacitance.

The extent that a voltage distribution deviates from a linear relationship, for each short-circuit simulation scenario listed in Table 4.1, can be evaluated by the integral between them, indicated by the area highlighted between the curve and line of Figure 4.3, or otherwise compared by average and maximum discrepancies from a linear distribution and the associated maximum voltage stress (ΔV) between any two adjacent foils as indicated in Figure 4.4. Any nonlinear distribution suggests that some elements experience greater voltage stresses than others, with consequent predisposition to degradation and eventual failure, and implying that parameters such as these therefore describe a unit's ability to evenly distribute voltage and mitigate the chance of further accelerated dielectric aging which can lead to cascading series element failures [1]: the lower they are; the more resilient the unit to this failure mode.

Figure 4.4 shows that a short-circuit increases maximum voltage stress (Max ΔV) between any remaining pair of adjacent foils compared to a healthy asset, save for failure of the highest-voltage element. In this anomalous circumstance the voltage distribution throughout a unit is more linear, as is the case for any short-circuit occurring above the sixth foil. Short-circuits below this foil lead to greater maximum and average discrepancies from a linear voltage distribution, as respectively illustrated by the broken and solid lines.

These results might offer a mild reassurance, given that healthy assets ordinarily experience greatest potential differences between their highest voltage foils (Figure 4.3). However, caution is due since in this model, an applied voltage remains unaltered between scenarios and this study does not therefore consider the influence of changes in overall unit (or bank) reactance which would occur in the event of an element failure on a real power system: greater total reactance is known to draw greater voltages, which would further change the potential differences in units which have blown a fuse in response to an internal fault [38].

The `Geometry` itself is similarly simplistic since it contains only a limited number of foil elements, and omits: elements connected in parallel; effects of temperature; and the effects of a capacitor bank structure, each of which would influence a real system. A change in reactance from series element failure could in turn lead to filter detuning, and compromise capacitor banks' ability to support reactive power and correct power factors.

Furthermore, a dielectric's response to an applied signal varies with frequency and temperature, as its characteristic eigen frequency[3] influences how its relative permittivity ϵ_r varies as seen by frequency components of the applied signal [54]. In turn, this will affect capacitances and voltage distributions within a unit and throughout a bank. Such considerations are an opportunity for more detailed study, particularly where harmonics characteristic of common converter topologies can be investigated, or where a focus can be given to modern dielectrics.

[3]Dielectric eigen frequency is a polarisation effect ($\omega_e = \frac{1}{\tau_p}$), where τ_p is the dipole relaxation time. ω_e is a function of a material's binding energy, and thus a function of temperature, and which "may vary considerably for different dielectrics" (p290-291 [54]).

4.1.3 Parasitic and Stray Capacitances

Unit housing affects capacitance, as this structure is proximal to and at some potential difference from (most, if not all) internal elements. Stray capacitances therefore form between elements and housing, on aggregate effect a subtle nonlinearity in voltage distributed among connected elements (or element groups), and have a slight influence on units' terminal capacitance.

Parasitic effects do not occur evenly for all foils: those at either end of a stack of windings have greater capacitance due to their increased exposure to housing. Such effects will accordingly differ with unit design: the use of interleaved winding designs; the presence of draw leads; the positioning of fuses and discharge devices; the arrangement of element groups within units; and regions where dielectrics are subject to incipient faults or variation in temperature, all serve to affect stray capacitances. Hence, while a unit can be designed to approximate a specified nameplate capacitance between terminals, this rating does not exist in isolation from its environment: influences on the dielectric between foils and housing will also lead to subtle changes in overall unit terminal capacitance.

In practise, relative to those of the windings, parasitic couplings are comparatively insignificant (having, on average, an influence equivalent to 9.29% of the size of each element capacitance by impedance, as calculated by Table 4.2) and difficult if not impossible to measure in practise for real assets and especially in practical environments. Moreover, couplings are likely tempered by inductive effects due to: winding design; long jumpers between racks; and peripheral components within capacitor units. However, parasitic capacitive effects are responsible for the nonlinearities inherent to capacitor energy distributions, result from unit topology, and if estimated through simulation or measurement could offer a unique understanding of each unit design, and could therefore prove useful in series element failure diagnosis and location.

Using a model to find voltages on intermediate foils within each unit, it might be possible to estimate these subtle effects (for an albeit sanitised simulation model, but which are impossible to obtain in practise), such that comparative changes in terminal capacitance ratings can be used to infer which of a set of series element winding groups has experienced a fault. As illustrated in Figure 4.4, knowledge of fault locations has implications for subsequent dielectric stresses, and thus also on consequent disposition to further faults.

Solvability of a Simple 2-group Unit

For a simplistic circuit representing a potential unit of only 2 element groups in series, as in the circuit diagram of Figure 4.5, the presence of stray capacitances C_3 and C_4 to housing could be inferred for each node. Initial equations can be derived from a mesh analysis and can be equivalently written in matrix form as in Equation 4.1. Conventional mesh analysis problems offer component values such that voltages and currents can be found to solve the circuit, but in this case voltages can instead be assumed to be known to solve instead for values of the capacitive circuit components.

To be solvable, each of the 9 unknown values in Figure 4.5 must have either a known value or an associated equation.If at least: two parameters can be evaluated (say, both voltages found through modelling, or more practically, the current and voltage at unit terminals); impedances related based on their circuit connections in the form of Equation 4.2 (which evaluate the reciprocal of z_T in terms of series and parallel arrangement of capacitances in the model); and five mesh analysis equations derived as in Equation 4.1, then 8 values or relationships are

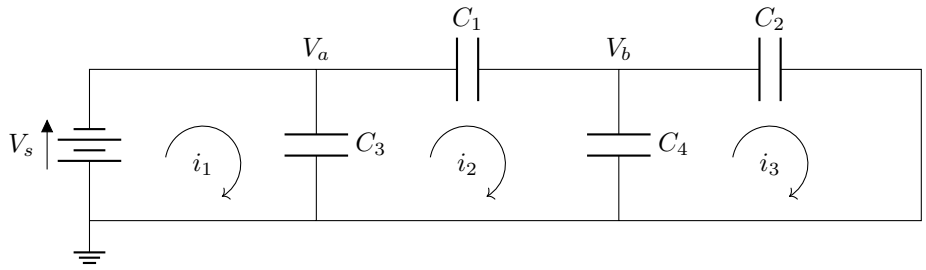

Figure 4.5: A simplified equivalent circuit with stray capacitances C_3 and C_4

available to solve a 9-variable problem (which is insufficient).

As discussed in Section 4.1.2 (which states that – for a sanitised, idealistic model – each pair of foils forming an element is similarly spaced with equivalent dielectric properties), if a similar geometry in the model for the winding capacitances can permit their equivalence to be assumed ($C_1 = C_2$), it is enough to solve for parasitic capacitances in this simple case. Since any extension in the number of series-connected element groups is accompanied by an expansion of Equation 4.1, including this equivalence assumption to form Equation 4.3 can theoretically allow a circuit of *any* number of series element groups to be solved. Hence, it is conceivable that an ability to measure terminal characteristics such as the current and voltage across a unit's terminals could be enough to solve for internal electrical characteristics experienced by a unit's windings, given a baseline understanding of the initial element and unit design.

$$\begin{pmatrix} v_a \\ v_a - v_b \\ v_b \\ v_a \\ v_b \end{pmatrix} = \begin{pmatrix} i_1 & 0 & 0 & 0 & 0 \\ 0 & i_2 & 0 & 0 & 0 \\ 0 & 0 & i_3 & 0 & 0 \\ 0 & 0 & 0 & i_1 - i_2 & 0 \\ 0 & 0 & 0 & 0 & i_2 - i_3 \end{pmatrix} \begin{pmatrix} z_t \\ z_1 \\ z_2 \\ z_3 \\ z_4 \end{pmatrix} \tag{4.1}$$

$$\frac{1}{z_T} = \frac{1}{z_3} + \frac{1}{z_1 + \frac{1}{\frac{1}{z_2} + \frac{1}{z_4}}} \tag{4.2}$$

$$z_1 = z_2 \implies \frac{1}{z_T} = \frac{1}{z_3} + \frac{1}{z_1 + \frac{1}{\frac{1}{(z_1)} + \frac{1}{z_4}}} \tag{4.3}$$

5 Series Group Design

A more practical unit design with 5 groups of 2 windings each is shown in Figure 4.6 with one terminal (shared by 2 windings and a draw lead) highlighted in red. Its circuit diagram in Figure 4.7 can be used to derive a total impedance as described by Equation 4.4. If each pair of element groups in the model can again be assumed to be sufficiently similar then $z_1 = z_2 = z_3 = z_4 = z_5$ and z_1 alone suffices as in Equation 4.5, which can in turn be rearranged in terms of an overall terminal impedance z_T, given as Equation 4.6.

$$z_T = \cfrac{1}{\frac{1}{z_a} + \cfrac{1}{z_1 + \cfrac{1}{\frac{1}{z_b} + \cfrac{1}{z_2 + \cfrac{1}{\frac{1}{z_c} + \cfrac{1}{z_3 + \cfrac{1}{\frac{1}{z_d} + \cfrac{1}{z_4 + \cfrac{1}{\frac{1}{z_e} + \frac{1}{z_5}}}}}}}}}} \tag{4.4}$$

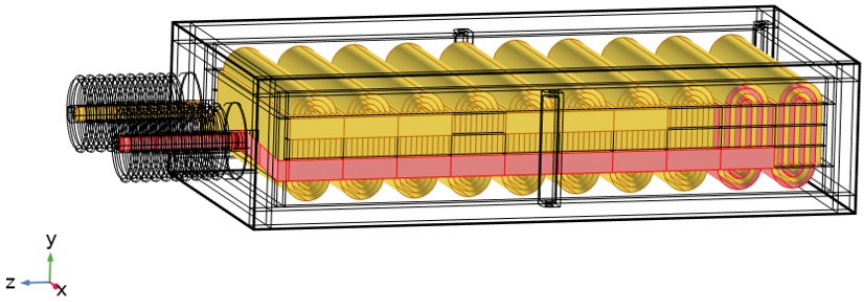

Figure 4.6: Unit with 10 windings connected as 5 series element groups

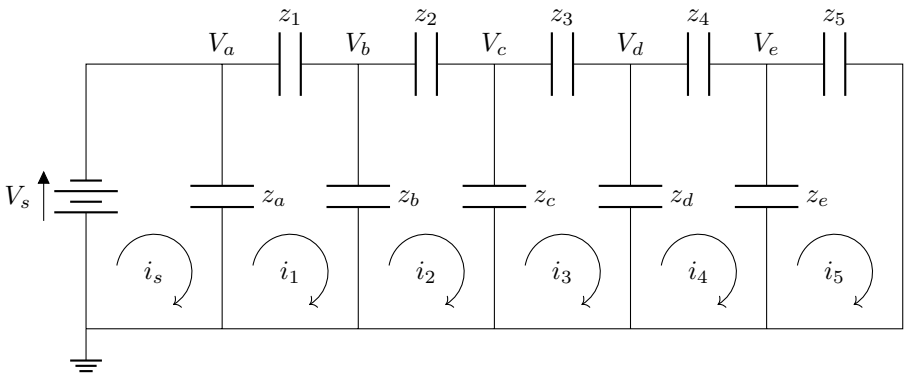

Figure 4.7: Circuit diagram equivalent to the unit in Figure 4.6.

$$z_T = \cfrac{1}{\frac{1}{z_a} + \cfrac{1}{z_1 + \cfrac{1}{\frac{1}{z_b} + \cfrac{1}{(z_1) + \cfrac{1}{\frac{1}{z_c} + \cfrac{1}{(z_1) + \cfrac{1}{\frac{1}{z_d} + \cfrac{1}{(z_1) + \cfrac{1}{\frac{1}{z_e} + \frac{1}{(z_1)}}}}}}}}}} \tag{4.5}$$

$$\begin{aligned}
z_T = z_a\big(&z_1^5 \\
&+z_1^4(z_b + 2(z_c + z_d + z_e) + 1) \\
&+z_1^3(z_c + 2z_bz_c + 2z_bz_d + 2z_bz_e + 2z_d + 3z_cz_d + 4z_cz_e + 3z_dz_e + 2z_e) \\
&+z_1^2(z_cz_d + 2z_cz_e + 3z_dz_e + 4z_cz_dz_e + 3z_bz_cz_d + 4z_bz_cz_e + 3z_bz_dz_e) \\
&+z_1(z_cz_dz_e + 4z_bz_cz_dz_e)\big) \\
\hline
&z_1^5 \\
&+z_1^4(z_b + 2(z_c + z_d + z_e + 1)) \\
&+z_1^3(3z_c + 4z_d + 4z_e + 1 + 3z_cz_d + 4z_cz_e + 3z_dz_e + 2z_bz_c + 2z_bz_d + 2z_bz_e) \\
&+z_1^2(z_c + 2(z_d + z_e) + 6(z_cz_e + z_dz_e) + 4(z_cz_d(1 + z_e) + z_bz_cz_e) + 3(z_bz_cz_d + z_bz_dz_e)) \\
&+z_1(z_cz_dz_e + 4z_bz_cz_dz_e + z_cz_d + 2z_cz_e + 3z_dz_e + 4z_cz_dz_e) \\
&+z_cz_dz_e
\end{aligned} \tag{4.6}$$

The circuit illustrated in Figure 4.7 shows the element windings and parasitic capacitances to housing, which is in this case assumed to be at the potential of the lowest voltage bushing, such that it has 6 electrical nodes. For this already highly-simplified capacitor design, $n = 5$ and $k = 2$ such that $\binom{n+1}{k} = 15$ theoretical capacitances exist of which 5 are element groups; 4 are parasitic to ground (the lowest voltage element group capacitance contains the final parasitic capacitance); and 6 represent theoretical stray electric field couplings between non-adjacent terminals.

Combinatorics are needed for a comprehensive understanding of stray and parasitic effects, but for greater numbers of element groups (in higher voltage designs) such analyses become intractable due to *combinatoric expansion*. Automated analysis would thus be necessary to fully quantify internal impedances, and even this remains limited to units with manageable numbers of element groups. Matlab's support for symbolic variables (`syms`) offers one approach to computationally solve more complex equivalent circuits.

However, element winding groups are, intuitively, by far the largest capacitances in a unit such that stray effects to unit housing can be considered comparatively minute: accounting for one between each terminal and ground leaves 5 (in the case of Figure 4.6) remaining theoretical capacitances which occur between non-adjacent electrical terminals and between elements and the draw lead. Such capacitances are thought to be largely shielded due to intermediate elements, sufficiently small that they can be considered negligible for this analysis, and in any case are likely more greatly affected by misrepresentations and simplifications in a simulation model than necessarily representative of effects present in practical designs, although they will have some influence.

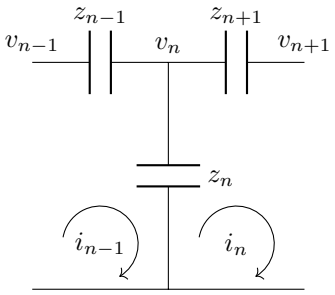

Figure 4.8: Stray capacitance in the general case

Additional effects present in practical settings include winding inductance, which could counteract stray capacitances. Moreover, where designs include a draw lead, it is considered good practise for this to connect to the lower voltage terminal, and in the case of Figure 4.6 where a low voltage terminal is grounded (where a unit sits in the centre of a rack on the side of a rack tie), any effects due to a draw lead are already incorporated within results.

Combinatoric Patterns

Terms can be grouped as a polynomial in z_1 when solving algebraically for a given unit design, as in Equation 4.6: in this form coefficients appear to follow patterns of binomial coefficients. Considering the general case, the impedance of a single stray capacitance z_n can be found relative to that of each element z_1, as in Equation 4.7, and as illustrated in Figure 4.8.

In Equation 4.7, a binomial relationship is apparent (denominator coefficients match those of the 2^{nd} row of Pascal's triangle – as 2 element groups are considered). These parasitic effects accumulate throughout a unit of series-connected element groups and therefore have a combined influence on voltages of adjacent foil windings. In turn, this explains the pattern of higher-order binomial coefficients which arise in algebraic working when solving for more complex designs.

$$z_1 = z_{n-1} = z_{n+1} \implies \begin{aligned} v_n &= (i_{n-1} - i_n) z_n \\ \frac{v_n}{z_n} &= \left(\left(\frac{v_{n-1} - v_n}{z_{n-1}} \right) - \left(\frac{v_n - v_{n+1}}{z_{n+1}} \right) \right) \\ \frac{v_n}{z_n} &= \frac{v_{n-1}}{(z_1)} - \frac{v_n}{(z_1)} - \frac{v_n}{(z_1)} + \frac{v_{n+1}}{(z_1)} \\ \frac{z_1}{z_n} &= \frac{v_{n-1} - v_n - v_n + v_{n+1}}{v_n} \end{aligned}$$

$$\frac{z_n}{z_1} = \frac{v_n}{v_{n-1} - 2v_n + v_{n+1}} \tag{4.7}$$

Equation 4.7 is applied to the voltage distribution of the healthy scenario in Table 4.1 to find the comparable influence of parasitic effects (C_n) for each of 10 capacitor elements (C_1) in the simple model it considers, given in Table 4.2. Proportional influences decrease as the parasitic effect associated with each terminal is electrically in parallel to a greater combined capacitance of winding elements and their own stray effects, save for an increase for T9 at the high-voltage end of the unit (possibly due to greater exposure and proximity to unit housing).

Table 4.2: Relative influence of stray capacitance by comparison with elements

Terminals	T1	T2	T3	T4	T5	T6	T7	T8	T9
Voltage (V)	454.39	1176.5	2062.3	3045.1	4089.3	5172.6	6299.3	7477.6	8703.8
$\frac{C_n}{C_1}$ (%)	58.92	13.91	4.70	2.02	0.96	0.84	0.82	0.64	0.80

4.1.4 Changed Terminal Characteristics

In the event of a short-circuit between terminals (a series element group failure), a change to a unit's overall capacitance is inevitable. It is less apparent how such a change can be interpreted from the current and voltage seen at a unit's terminals, or if any such detected change can be used to infer the winding foils involved in a dielectric breakdown.

For the general case of an n-group unit, an $n \times 1$ matrix can describe these capacitances, as in the case of the right-most matrix in Equation 4.1. As presented in Section 4.1.3, the equations in Equation 4.1 are solvable provided an assumption is made about the similarity of winding capacitances ($z_1 = z_2$ in this case), and at least two circuit quantities are known, which could be unit terminal voltage and current.

5-group Unit Studies

Using the model in Figure 4.6, simulations are run using a `Direct` solver for `Stationary` studies of each possible series element short circuit failure, as in Section 4.1.1 but for this more recent model. A further improvement is the use of `I-Terminal` couplings between an `Electric Circuit` and both `Electric Currents` and `Electrostatics` interfaces, which in turn respectively use `Terminal` and `Floating Potential` boundary conditions. A 10 kV DC voltage is applied through a 0.1 Ω supply resistance, and simulating with `Electric Currents` allows the change in terminal current to be compared with a healthy scenario, as shown in Figure 4.9. Tables 4.3 and 4.4 provide results which indicate that there are minimal differences between the outcomes of each interface, and moreover, Table 4.3 shows that there *are* subtle differences in terminal current change due to the *location* of series element failure short-circuits.

Figure 4.9 illustrates the change in current drawn from a 10 kV, 0.1 Ω supply corresponding with each fault location. Notably, these changes are only sufficiently significant for faults between higher-voltage element groups, as similar values are seen for the `T3-T2`, `T2-T1`, and `T1-T0` scenarios, but high-voltage fault detections alone could prove useful as these winding groups are most prone to failures anyway. Faults between windings and housing are not simulated during these brief studies, but could be an appropriate avenue for future work.

At first glance, such subtle changes in unit impedance would require high fidelity measurement to discern or infer failure locations. Fortunately, these small values result from a highly simplified simulation model employing minimal numbers of windings and winding turns, and wide spacing between foils. Practical units have much larger capacitance, and therefore reactances low enough to admit larger currents, so it would be entirely speculative to gauge the suitability of measurement systems without more detailed and representative modelling.

Results in Tables 4.3 and 4.4 are in keeping with those in Sections 4.1.1 and 4.3, such that `I-Terminal` nodes appear to have little influence on results, even if these offer more efficient means of simulation with reduced compuation times and less constrained solution spaces, as further discussed in Section 4.3.

Table 4.3: Electric currents modelling

Scenario	$I_{Terminals}$	V_0 (V)	V_1 (V)	V_2 (V)	V_3 (V)	V_4 (V)	V_5 (kV)
Healthy	1.3231 (nA)	0	1653.5	3456.7	5275.2	7477.6	10
$T_5 = T_4$	1.7206 (nA)	0	2208.6	4616.9	7045.8	10 (kV)	10
$T_4 = T_3$	1.5987 (nA)	0	2148.1	4490.4	6862.6	6862.6	10
$T_3 = T_2$	1.5023 (nA)	0	2084.2	4365.5	4365.5	7046.3	10
$T_2 = T_1$	1.5298 (nA)	0	2157.1	2157.1	4428.6	7072.4	10
$T_1 = T_0$	1.4885 (nA)	0	0	2296.4	4514.6	7113.6	10

Table 4.4: Electrostatics modelling

Scenario	$I_{Terminals}$	V_0 (V)	V_1 (V)	V_2 (V)	V_3 (V)	V_4 (V)	V_5 (kV)
Healthy	0 (nA)	0	1653.3	3456.2	5274.5	7476.6	10
$T_5 = T_4$	0 (nA)	0	2208.6	4616.9	7045.8	10 (kV)	10
$T_4 = T_3$	0 (nA)	0	2147.8	4489.9	6861.8	6861.8	10
$T_3 = T_2$	0 (nA)	0	2084	4365	4365	7045.4	10
$T_2 = T_1$	0 (nA)	0	2156.8	2156.8	4428.1	7071.5	10
$T_1 = T_0$	0 (nA)	0	0	2296.1	4514	7112.7	10

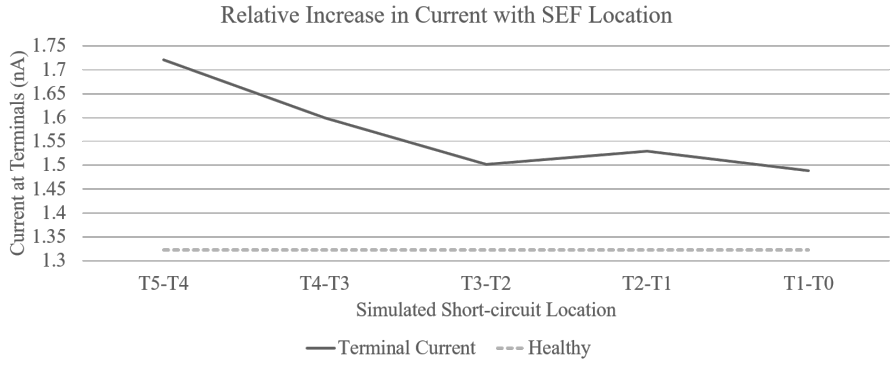

Figure 4.9: Relative change in terminal current for each fault location

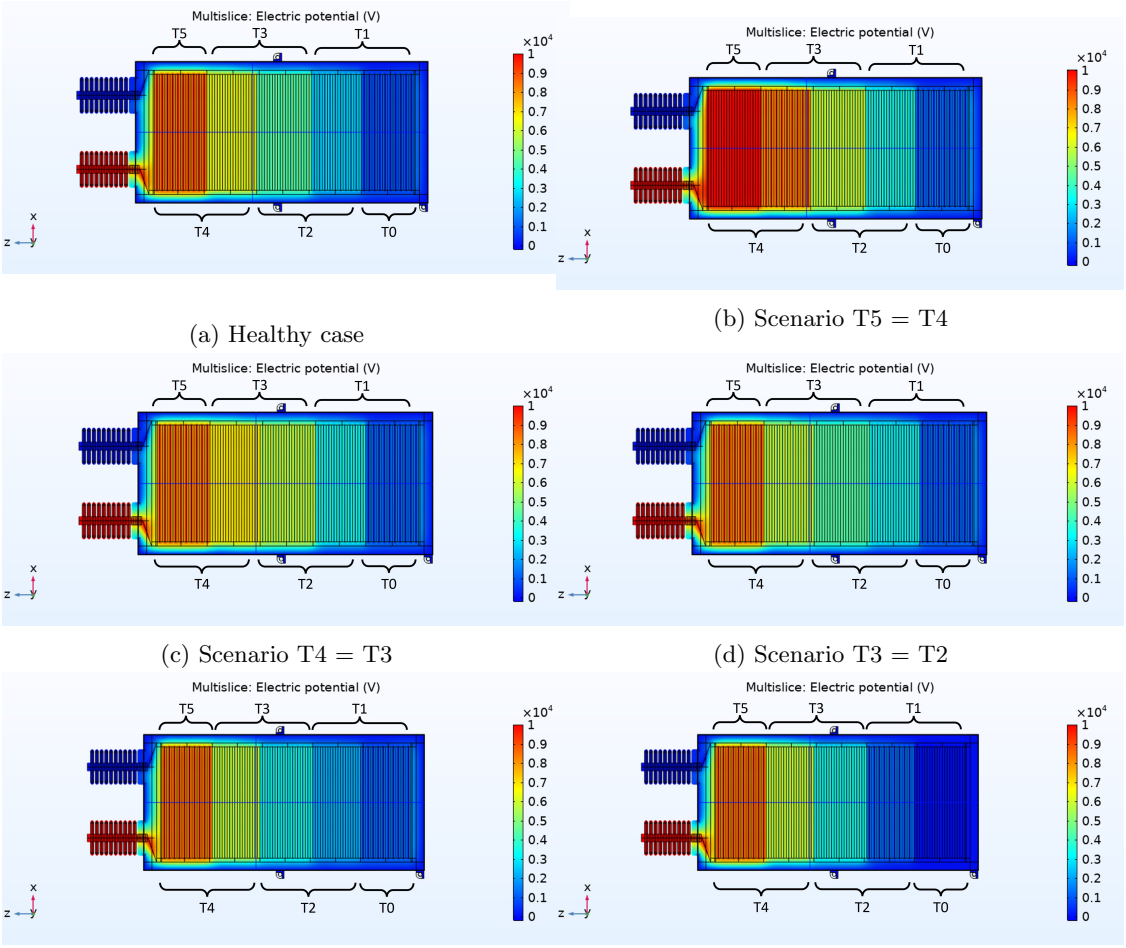

(a) Healthy case

(b) Scenario T5 = T4

(c) Scenario T4 = T3

(d) Scenario T3 = T2

(e) Scenario T2 = T1

(f) Scenario T1 = T0

Figure 4.10: Electric potential in an XZ cut plane

Fault Location Inference

Should it be possible to use terminal current and voltage to infer or estimate the solution to an equivalent circuit for a given unit design, it could be possible to interpret changes in unit characteristics due to degradation and failures over the course of an asset's operational lifetime. This is the benefit of simulation modelling: it is possible to resolve (at least an approximate) *characteristic matrix* for a given capacitor unit design, which can in turn be used to interpret changes whether as part of knowledge based diagnostics or as a feature set useful for data analytic approaches should routine measurement data be available.

Eventually, if change in terminal capacitance varies with failure location, this measure could be used to identify a failed element group. This ability would allow engineers to gain:

1. **improved insight from routine measurements** – should individual units be measured during an outage, not only can their overall value be compared to the last inspection but the *change* in value can be used to infer the affected group within each unit;

2. **evidence for future design enhancements** – failure locations would reveal where improvement to unit designs are most appropriate; and

3. **information about the electrical environment** – if a bank is subject to routine impulses, failure locations could reveal how far through a bank such impulses propagate.

Accordingly, fault location inference could be used to corroborate alternative fault detection and location approaches, as to be further discussed in Chapter 5.

4.2 Electric Field Effects in a Capacitor Unit

This Section investigates electrical field stresses within a high voltage capacitor unit, and aims to better understand these internal electric fields, not least since electric fields form capacitance, but also to find where these electrical stresses arise beyond designed capacitor elements, where dielectrics are subject to high field strengths, and hence to better inform future unit design.

A discrete foil unit is modelled with interleaved [29] windings as a high-voltage design. This is an approach more likely to be used in high voltage applications, but the chosen configuration is also better suited to exploring electric field effects due to its more involved structure: interleaved windings each have multiple elements; series-connected windings conduct current in opposing directions; and an odd number of windings necessitates a draw lead.

Due to their role and position within converter stations, capacitor banks are subject to impulses from being switched in and out of service, as well as from converter operation and its associated harmonics. Improved understanding of where units' internal components could be subject to stresses from these high-frequency signal artefacts is necessary to understanding the conditions which predispose a unit to fault, the risk a given cause might pose (to inform a post-mortem); and offer insight into internal locations susceptible to incipient faults.

4.2.1 Setup

This study uses a discrete-foil, high-voltage capacitor model as illustrated in Figure 4.11. It makes use of model parameters to resize the `Geometry` and orient foils along the Z axis, before margins are introduced to allow selective `Definitions` to create an interleaved winding design: windings and foil turns occupy almost the entire volume of a capacitor unit, and `Domains` are specified to create perpendicular divisions. These divisions mean `Definitions` can be selectively applied to each winding to specify each as either: a single contiguous capacitive element; or as an interleaved series-connected group of elements spaced with dielectric margins, as indicated in Figures 2.3 and 6.3 [29]. Foils are 2D boundaries, and as dielectric layers are omitted from the `Geometry` (for simplicity), the remaining volume approximates such layers as a dielectric volume of relative permittivity $\epsilon_r = 3.9$ to mimic mineral oil as introduced in Section 3.1.4.

This model accordingly has five interleaved windings, each with five series capacitances, and aligned with the Z-axis of the model `Geometry`. An odd number of windings necessitates use of a draw lead to connect all windings in series and make the 'ground' bushing connection [1], included to demonstrate its role in shaping electrical stresses within a unit, albeit with two-dimensional foil boundaries to maintain a simplistic `Geometry`. Windings connect in series at alternate ends of the unit, with a draw lead to make the low-voltage connection along the module's length to the low-voltage bushing, where the draw lead, its associated foil terminal, and the housing are all at ground potential. Additional features such as rack mounts, bushings, and connecting leads are included to finalise the `Geometry` and to provide the option for detailed effects (of a draw lead, for example) to be explored when boundary conditions are applied.

4.2.2 Simulation Studies

A two-step process is followed: firstly, an `Electrical Circuit` interface is configured and coupled with one for `Electric Currents` to ascertain the distribution of voltages across foils

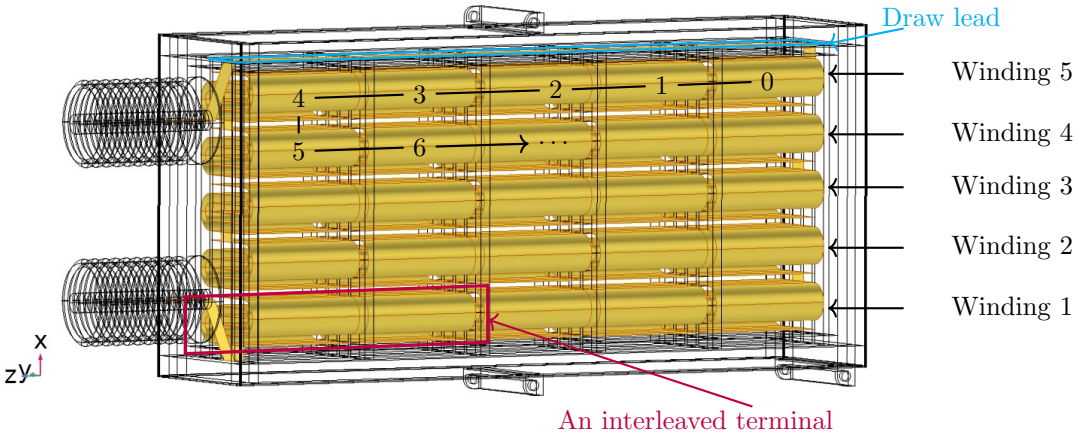

Figure 4.11: A unit with interleaved foil windings in yellow and numbered elements

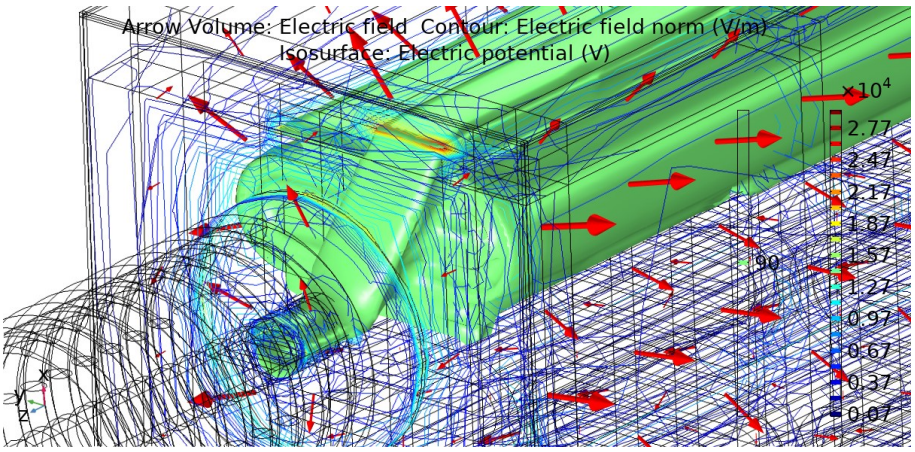

Figure 4.12: Visualisation techniques in COMSOL®

within the unit. Secondly, these voltages are again specified in an `Electrostatics` interface to study electric field stress. With this process, the effects of a ground connected draw lead and application of a 10 kV source are investigated using a steady-state study.

4.2.3 Results

Figure 4.12 illustrates postprocessing options available in COMSOL® by showing electrical effects within the model: an `Isosurface` is set to 90% of the maximum voltage to encapsulate the high voltage winding[4] and draw lead; and an `Arrow Volume` and `Contour` lines reveal electric field direction and strength. Figure 4.13 illustrates another visualisation technique available in COMSOL®, where `Streamlines` are used to show the electric field. With `Streamlines` coloured according to voltage, this reveals: how a draw lead (to the right of this diagram) distorts the field; some fringing effects between adjacent windings; and the potential for unintended parasitic effects to arise between adjacent foils and between foils and housing.

[4]This particular model was initially configured with Winding 5 connected to the high voltage terminal via a draw lead, but was reconfigured for other studies in this Section to allow for the fact that draw leads are vulnerable components [1] which should instead be connected to the low voltage terminal.

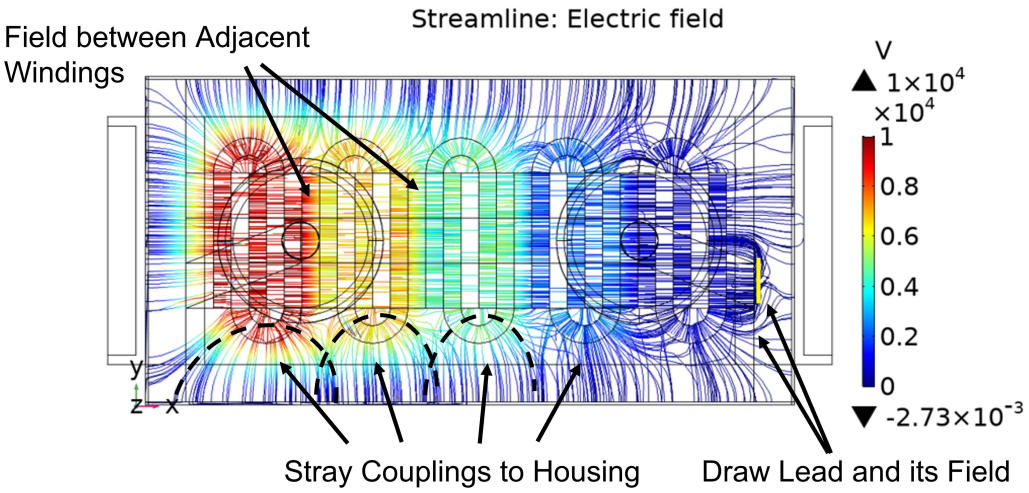

Figure 4.13: A streamline of electric field shows the nature of stray effects

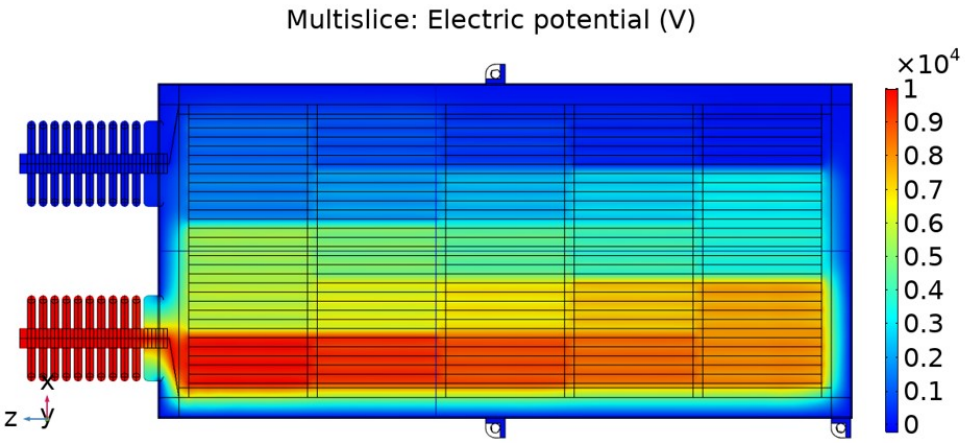

Figure 4.14: A cut plane shows voltage to increase within each interleaved winding

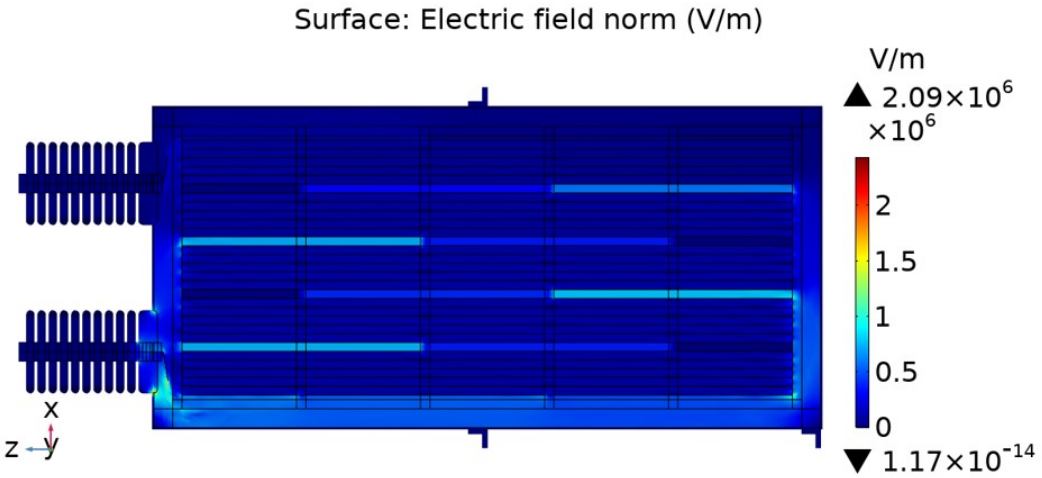

Figure 4.15: An XZ cut plane shows where field stresses arise between adjacent windings

Table 4.5: Voltages across foil windings (V)

Winding	Foil 5	Foil 4	Foil 3	Foil 2	Foil 1	Foil 0
5	1367.7	1010.8	701.0	402.2	171.3	0.0
4	1367.7	1761.3	2153.2	2555.5	2932.4	3342.9
3	5613.0	5151.1	4695.1	4232.4	3781.6	3342.9
2	5613.0	6078.3	6548.6	7010.8	7461.1	7892.9
1	10000.0	9575.3	9144.6	8730.1	8303.8	7892.9

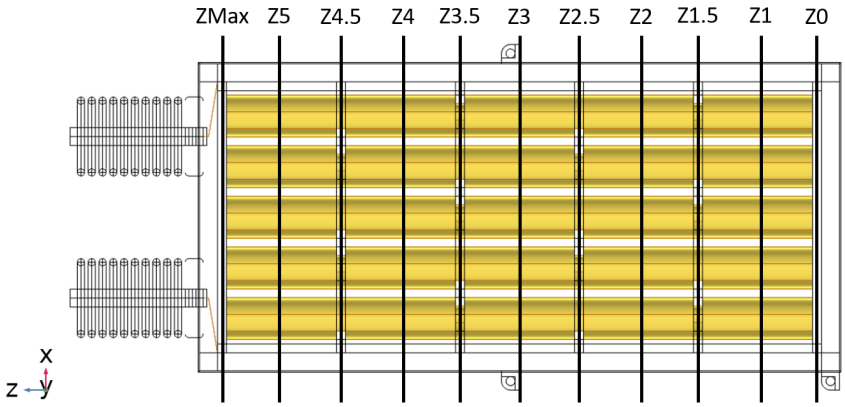

Figure 4.16: Indexing for cut planes

Figure 4.14 shows the distribution of voltage in the capacitor unit graphically, where values are presented in Table 4.5 (in an order to match the orientation of Figure 4.14), and Figure 4.15 highlights how these voltages incite regions of electric field stress within the model due to its series-connected interleaved winding design.

Figure 4.16 indicates the placement of `Cut Plane`s along the Z axis of the unit `Geometry`. `Cut Plane` diagrams are included in Figures 4.17, 4.18, and 4.19, to show where electric field stresses arise throughout a model cross-section. Colour legends are common to all Figures here aside the latter two diagrams of Figure 4.17, which respectively show field stress on the inner and outer sides of a draw lead, while the topmost diagram in Figure 4.17 shows field strength between windings 4 and 5. Figure 4.18 shows electric field stresses between windings and the housing on an XZ plane, where heightened stress can be seen to accord with winding voltages. Figure 4.19 comprises XY `Cut Plane` diagrams positioned at Z2, Z2.5, and at ZMax as indexed by Figure 4.16, which respectively reveal electrical field stresses bisecting foil overlaps, overlaps with dielectric margins, and the bushing end, of each of the set of five capacitive windings.

4.2.4 Discussion

Results introduced in Section 4.2.3 are discussed before conclusions are drawn in Section 4.4.2. The `Geometry` used in this `Study` is a simplistic representation of commercially available assets, and therefore electrical stresses arising in practical settings could vary significantly from presented results. In practise, 1. foil corners and edges are typically folded to mitigate stresses on surrounding dielectrics; 2. units can be configured with even numbers of windings to mitigate need for a draw lead [1]; and 3. winding ends are 'schooped' [43]; although 4. representative winding turn radii are much less than those used in these models; and crucially 5. layers of dielectrics used in real windings will be significantly thinner than simulated.

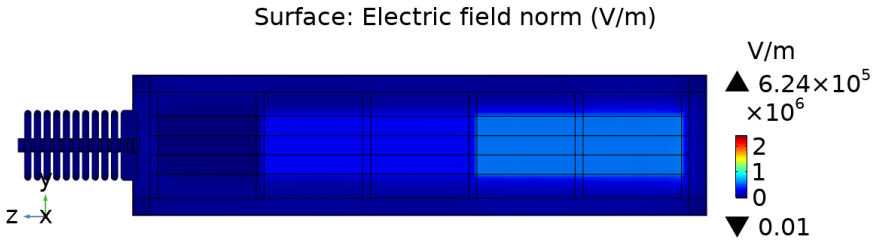

(a) A cut plane in YZ showing the distribution of electric field between adjacent interleaved windings

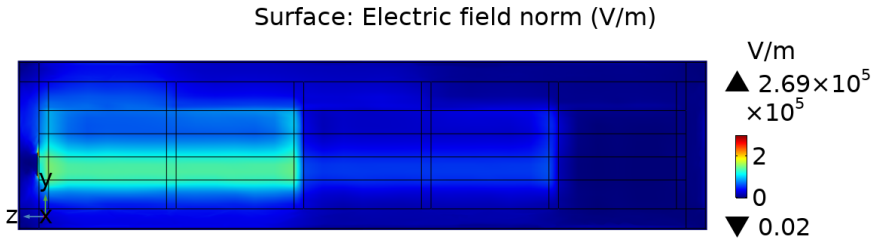

(b) A cut plane in YZ showing heightened electric field between a winding and a draw lead

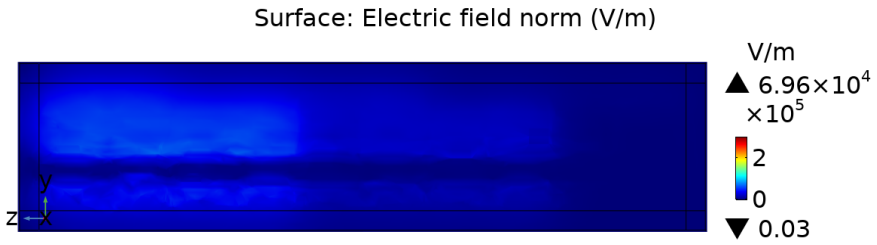

(c) A cut plane in YZ showing minimal electric field between a draw lead and housing

Figure 4.17: YZ distributions of electric field stress

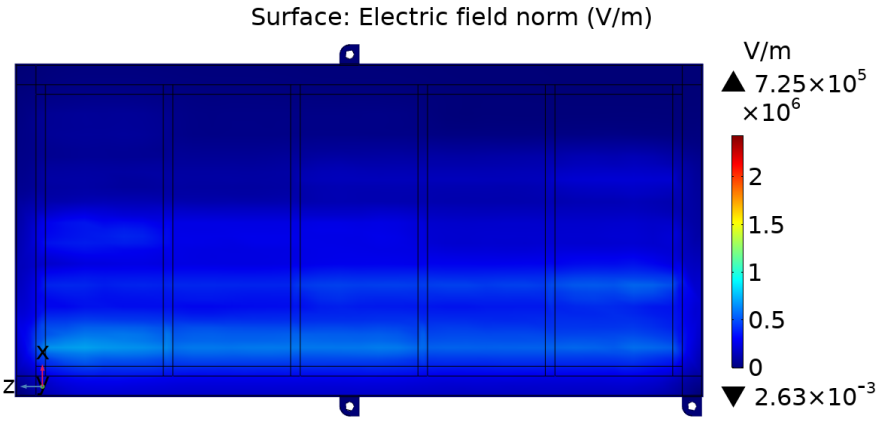

Figure 4.18: A cut plane in XZ showing electric field effects of winding turn radius

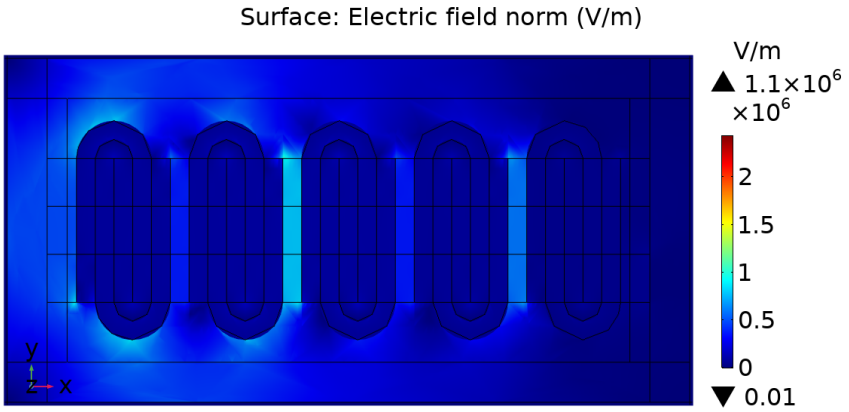

(a) A cut plane in XY at Z2 showing electric field effects at the centre of 5 elements

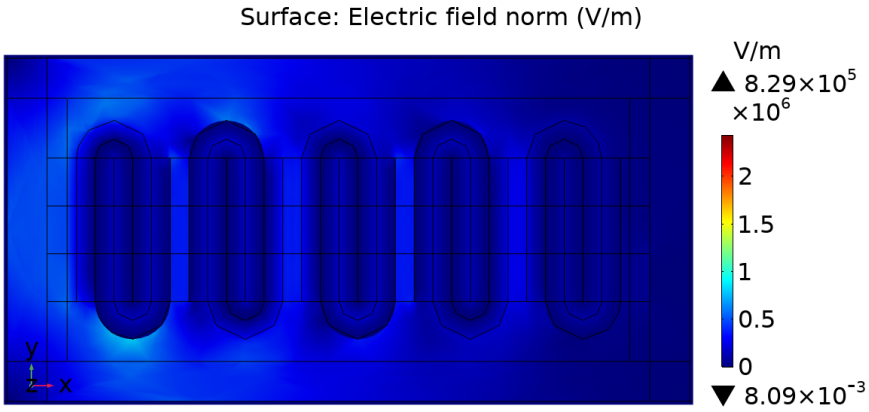

(b) A cut plane in XY at Z2.5 showing electric field at element margins

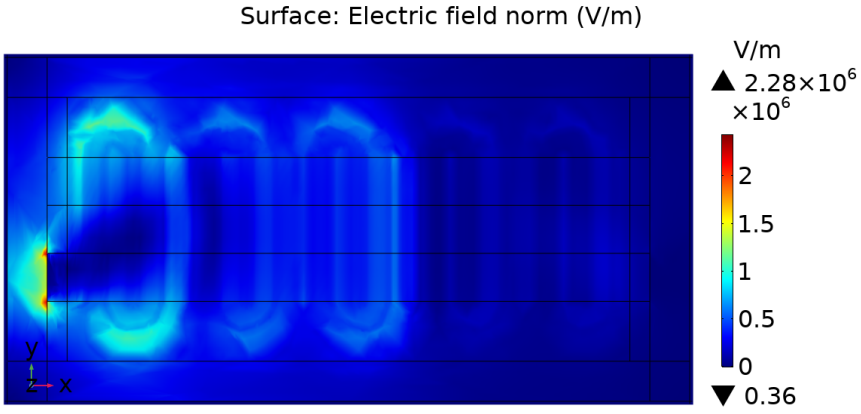

(c) A cut plane in XY at ZMax showing electric field effects near unit bushings

Figure 4.19: XY distributions of electric field stress

Results

Electric field stresses are strongest where connections are made between the unit terminals and capacitive elements, and where bushings meet unit housing. This can be seen in Figure 4.15 and Figure 4.19, and these field strength maxima establish an upper bound on the normalised colour legend for diagrams in Figures 4.15, 4.19, and 4.18.

Table 4.5 and Figure 4.14 illustrate voltage distribution within the model. Between adjacent windings, stresses are strongest between neighbouring foils with greatest voltage differences and are minimised between connected foils, so the pattern of electric stresses between windings shown in Figure 4.15 appears sensible in context of the interleaved winding arrangement when voltages are considered. Four areas have stresses induced between adjacent foils of greatest potential difference, although Figure 4.15 shows that strengths of electric field in these areas are not equal. Beyond aspects of the unit design, and as discussed in Section 4.1, when a nonlinear voltage distribution is considered it can explain why some of these areas experience greater stresses than others (for example, the electric fields shown between windings 2 and 3 in Figure 4.15 are stronger relative to those between 1 and 2, 3 and 4, and which are certainly greater than those between windings 4 and 5). Therefore, a nonlinear voltage distribution would suggest electric fields are also, consequently, unevenly distributed within the model.

Good practice in capacitor manufacture is to fold foil edges and corners before forming each element winding. The model simulated for this study is simplistic and does not include folded corners or edges on element foils, and the sharp points which result can be seen to contribute to electrical stresses in Figures 4.15, 4.17, and 4.19, especially at winding ends.

Figure 4.18 shows a `Cut Plane` close to the tips of foil turns, and highlights field stresses which arise along the length of each winding fold. It shows that these stresses are proportionate to foil voltages, as in Figure 4.15, but also suggests winding shape has a role in cultivating electric field stresses: an implication which is further supported by regions of electric field stress which form around foil folds as seen in Figure 4.19. Foil bend radius, therefore, could affect the strength of field in these regions. If similar capacitance can be achieved using either broad or narrow element windings (large or small bend radius), then from an electric field perspective the former is preferred in any unit design. Where one foil is exposed at a dielectric margin, electric field stresses increase around the underlying foil, as illustrated by Figure 4.19b.

A draw lead on the right of Figure 4.13 has a distinct effect on the electric field, heightening stress on dielectric between it and the nearby winding. Since such leads must cross the entire length of the unit, good design practice would have them connect to the terminal closest in voltage to unit housing, as in this `Study`. Moreover, it might be beneficial to position draw leads on whichever side of a unit is lower to mitigate against effects of temperature on proximal dielectrics [1], and draw leads can be avoided entirely should an even number of element windings be chosen, where connections need only be made at one end of the unit [1].

Electrical field lines illustrated in Figure 4.13 indicate areas of increased electric field stresses within the unit, illustrate the presence of stray capacitances, and reveal fringing effects which form between elements (such as at margins). Accordingly, Figure 4.13 implies parasitic capacitances (between elements and housing) to depend on element position: elements in the lowest and topmost windings are exposed to a greater surface area of housing, so appear to be more electrically proximal in terms of greater coupling and hence stronger stray capacitance effects.

Table 4.6: Simulated materials

	Al	Steel	Mineral oil	Porcelain
Component	Foils	Housing	Dielectric	Bushings
Conductivity (Sm^{-1})	$30.30e^6$	$4.032e^6$	$10e^{-15}$	$1e^{-15}$
Relative Permittivity	1	1	4	6

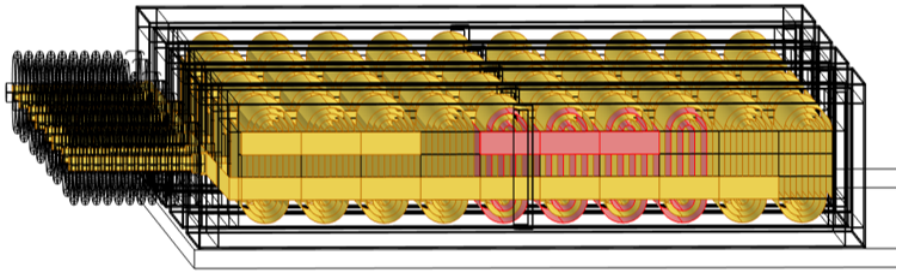

Figure 4.20: Foils in a capacitor unit with 5 series groups of 2 elements each

4.3 Modelling Multiple Units

As discussed in Section 2.2.3, individual capacitor units are modular subcomponents of larger assets. Capacitor banks are structured in racks for efficient footprints in addition to stresses placed on dielectrics from tiered voltages.

Figure 4.20 shows a simplistic design for a set of 4 adjacent units: each comprising 10 discrete-foil elements of 4 turns of aluminium foil, paired to form 5 series element groups amid a relative permittivity $\epsilon_r = 4$ mineral oil dielectric. It reveals foil turns, a draw lead, and a terminal shared by adjacent element groups (in red) for one side of a rack of 8 units in Figure 4.21. The units modelled omit fuses and use materials in Table 4.6. Figure 4.21 also illustrates series jumper connections between units on each side (A and B), which connect in parallel to a 40 kV source voltage with a rack tie on Side A.

4.3.1 Simulation and Results

Two studies are undertaken: firstly, a single pair of units is modelled in series with a 22 kV source voltage to represent two units between the rack tie and voltage source on Side A; before a study is conducted of the complete rack shown in Figure 4.21 (showing the Side A rack tie).

Connections between units are respectively shown for each study in Figure 4.22 and in Figure 4.21, where an electrical circuit interface in COMSOL Multiphysics® connects: 1. in the first study, (a) unit housings to ground, and (b) both units in series; and 2. for the second, (c) units on each side ('A' and 'B') in series, (d) both sides in parallel to share a common ground, measurement nodes, and source voltage and resistance, and (e) a rack tie from a jumper between Side A's central units which secures housings of all units on the rack at a shared potential.

ΔV is used to approximate the stresses placed on dielectrics, where $\Delta V_i = V_i - V_{i-1}$ for $V_{-1} = 0$ and for $0 \leq i \leq n$, where $i \in \mathbb{Z}$ within the set of n foils. This 'rate of change' in voltage serves to illustrate regions of electric stress placed on dielectrics used within individual elements, for a simplistic comparison (e.g. it takes no account of stresses outwith windings).

The first study mimics one quarter of a rack. A resulting cut-plane voltage distribution is shown in Figure 4.22, with results provided in Figure 4.23, both showing relative change in voltage (ΔV) as a means to approximate dielectric stress.

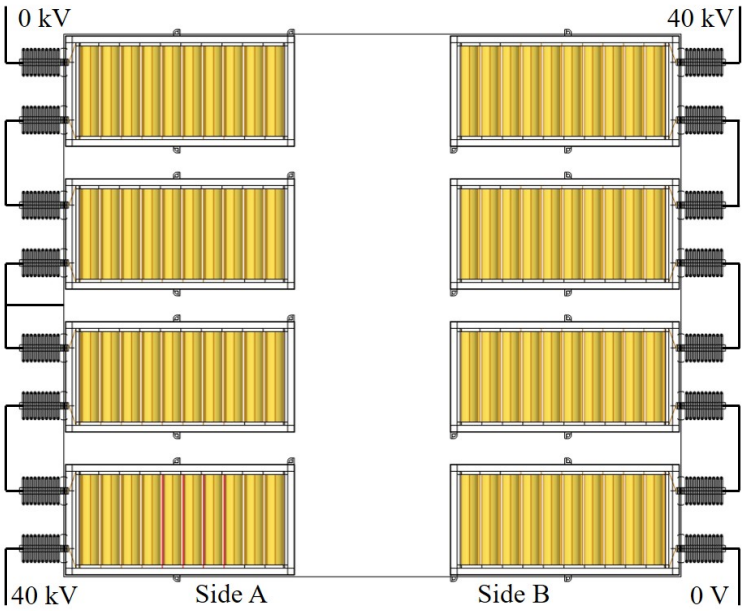

Figure 4.21: Foils in a capacitor rack of 8 units

Table 4.7: Foil voltages on rack (kV)

	A1	**A2**	**A3**	**A4**	**B1**	**B2**	**B3**	**B4**
0	0.000	6.801	11.91	22.13	0.000	5.593	15.67	27.41
1	0.683	8.482	12.83	25.35	0.489	7.385	17.93	29.86
2	1.878	9.982	14.47	28.83	1.379	9.329	20.26	32.37
3	3.385	11.18	16.63	32.46	2.560	11.37	22.62	34.89
4	5.071	11.87	19.22	36.19	3.981	13.49	25.00	37.43
5	6.801	11.91	22.13	40.00	5.593	15.67	27.41	40.00

Secondly, a rack of 8 individual units (Figure 4.21) is studied with the same interfaces to provide results outlined in Table 4.7 and illustrated in Figure 4.24 and Figure 4.25.

4.3.2 Discussion: First Study

Colour variation evident in Figure 4.22 at a glance reveals an uneven share of stresses, and corresponds with the dotted line in Figure 4.23. These dielectric stresses result from the relative increase in voltage from one foil to the next, suggesting that the unit connected to voltage is more likely to suffer series element failure, or failure of dielectric between foils and unit housing, than is the case for the unit on the left to which rack voltage is tied. Thus, failures are more likely on units furthest from a rack tie, so numbers of capacitor units on each rack are limited by such disparities between voltages on unit foils and housing, considering strengths of dielectric materials used.

Nonlinear Voltage

Figure 4.23 shows that nonlinear voltage distributions exist beyond any single unit and are present across all series-connected units. In the case of the first study, where there are no complicating factors, this nonlinearity is straightforward.

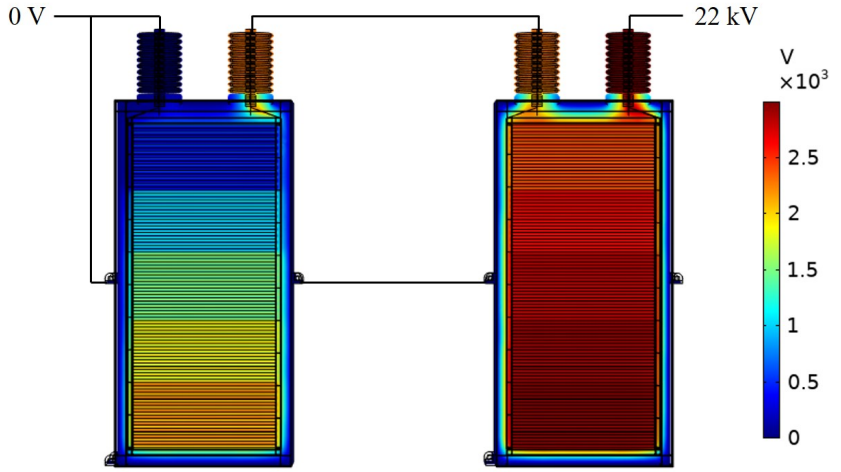

Figure 4.22: Voltage shared by units connected to a rack (left) and to 22 kV (right)

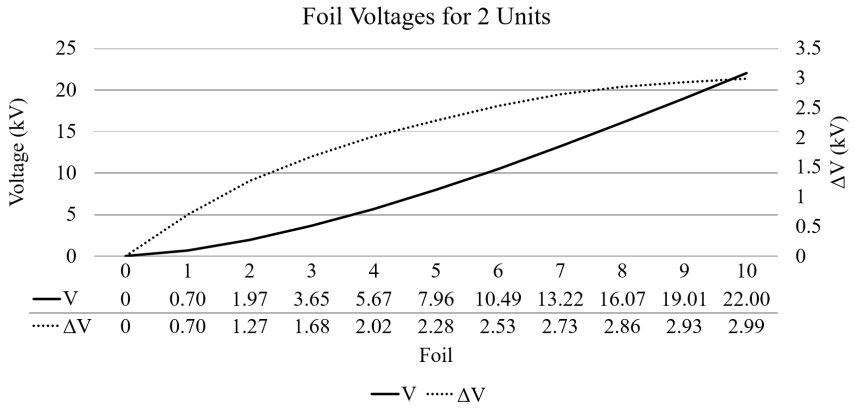

Figure 4.23: Voltages shared by adjacent units

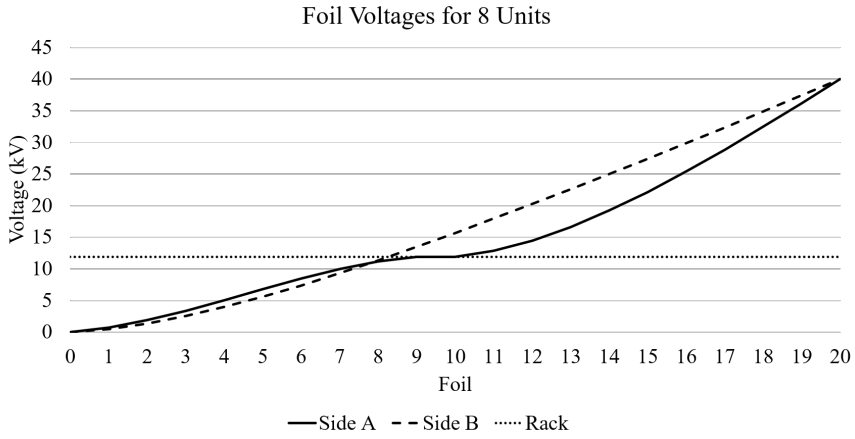

Figure 4.24: Foil voltages

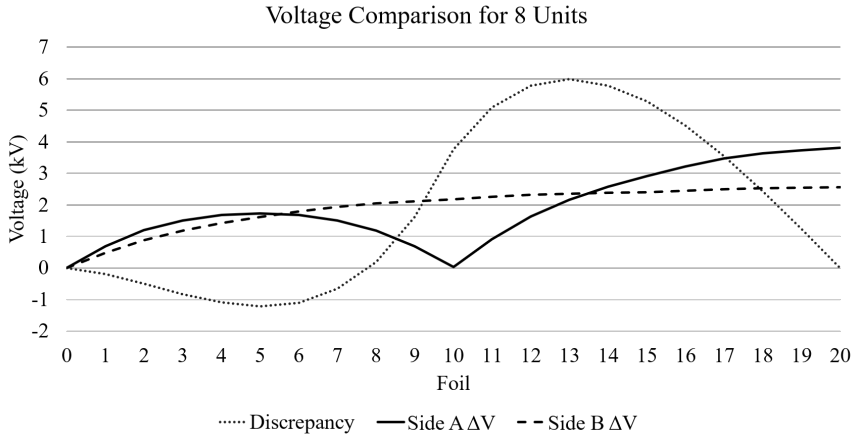

Figure 4.25: Foil voltages relative to rack potential

4.3.3 Discussion: Second Study

For the model in Figure 4.21 with results in Figure 4.24, a similar nonlinear distribution is evident in the broken line representing foils on Side B, where there is no rack tie. However, Figure 4.24 also shows Side A has an inflection point such that its slopes are steeper and dielectrics therefore more stressed, and hence more susceptible to element failure than Side B.

Rack voltage (the dotted line) is around only 30% of the applied voltage due to this nonlinearity. Potential differences between winding foils and that of the rack (and unit housings) test dielectrics most severely in corner units, and thus especially those connecting to high voltage.

Effects of a Rack Tie

As only one side ordinarily supports rack voltage, discrepancies are inevitable. Extending the results of Figure 4.24, Figure 4.25 highlights ΔV for each Side A and B: the aforementioned rack tie provokes a point of inflection in Figure 4.24 which corresponds with a local minimum on Side A's ΔV at foil 10; whereas ΔV for Side B (the broken line) gradually increases toward the high voltage connection, as in Fig. 4.24. Side A sees greater stress for supporting rack voltage.

Foils $\{0, 5, 10, 15, 20\}$ represent terminals with jumper connections (as in Figure 4.21), so a local maximum in ΔV for Side A on foil 5 occurs where a jumper connects units A1 and A2. Further work would be required (using more comprehensive models) to make certain, but this suggests such connections to be sites of heightened stress throughout a bank, where midpoints between rack ties – jumpers which connect between racks (or rather, the dielectrics immediately adjacent to these connections) – are likely to experience higher stresses.

Figure 4.25 shows a voltage discrepancy ($\Delta V'$) between equivalent foils on each parallel side as a dotted line, where: $\Delta V'_i = V_{i,\ Side\ B} - V_{i,\ Side\ A}$ for i as defined in Section 4.3.1. Such a discrepancy might be of limited interest for the design considered in Figure 4.21, but implies a heightened risk of rack faults in capacitor bank designs where separate strings of units have proximal connections. For instance, should distinct unit strings share the same face of a bank, jumpers between racks would interleave between inner and outer units (as illustrated in [1]), such that this discrepancy due to rack tie distribution could predispose a bank to rack faults.

Manufacturers might prefer to alternate sides which support rack potential, as this could help balance voltage distributions, a bridge or neutral, and also aid fault finding in future.

Unit Design

All units modelled have draw leads (Figure 4.20) in a common configuration, but which can be avoided in practical designs by using even numbers of series groups. Where draw leads *are* used, Figure 4.25 reveals the most stressed terminals on each unit not always to be those at higher voltage, suggesting units are best positioned cognisant of voltage and dielectric stress distributions, with draw leads on the terminal electrically closest to the nearest rack tie, as well as with heat in mind [1]. While voltages age dielectrics, failures ultimately result from electrical field stresses too: bushing connections are a particular risk for field stress, for instance; so units can be designed cognisant of such influences to improve their resilience [85].

Cascading Failure

Where one string of capacitor units on a bank determines rack voltages, it follows that a series element failure on the opposing side alters its voltage distribution, so can accentuate the disparity between parallel unit strings. This suggests that not only can series element failure predispose a bank to further breakdown primarily as a result of there being fewer remaining elements and with changed characteristics (such as increased reactance drawing greater voltage), but also that discrepancies between strings are aggravated, heightening risk of rack fault. Were it possible to measure, changes in rack voltages over a bank's operational life might be one means of determining the extent of failures, but only for the supporting string of units.

Use of COMSOL®

In coupling `Electrical Circuit` and `Electric Currents` interfaces, solution spaces are sensitive: an ability to find a solution depends on the source voltage (V_{sc}), a value of source resistance (R) connected between the source and parallel sides of the capacitor rack, relative permittivity for the dielectric (ϵ_r), and solver error tolerance. For large models, it can take an unpredictable amount of time to reach a solution for a given problem. For the same `Geometry`, `Mesh`, selections, `Materials`, and solver (with relative error 0.005), solutions are found to be at (V_{sc}, R, ϵ_r) = (40 kV, 1 $k\Omega$, 4), and (22 kV, 0.22 Ω, 4).

4.4 Conclusions

4.4.1 Series Element Failure and Voltage Distribution

The primary outcome of this Chapter is an ability to adapt and study the model presented in Chapter 2 to simulate electrical effects present in capacitor assets. This model is scalable through meshing techniques; adaptable through parameters; and while imperfect, offers a means of simulation toward understanding dynamic asset behaviour under operational conditions.

Results presented in Section 4.1.1 suggest that:

- voltages are distributed *nonlinearly* throughout a capacitor unit, as a consequence of parasitic capacitances induced by the unit geometry;
- following a short-circuit, the voltage distributed across remaining foils is also nonlinear, and this distribution varies as a function of the location of each failed capacitor element;
- higher voltage foils experience the greatest voltage stresses, but simulation results suggest that subsequent voltage stress increases (and therefore further propensity for dielectric degradation and breakdown) are overall most exacerbated when lower voltage elements fail.

This model is limited by its simplicity, but it offers a basis on which such considerations can be explored in future work, for which the capacitor model can be scaled to better represent real assets, and for which different geometries, electrical frequencies, and thermal constraints can be considered.

Security, diversity, and cost of an electrical power system depend on the health, reliability, and longevity of its constituent assets. An improved understanding of capacitor behaviour can help inform operation and maintenance decisions in the context of a wider electrical power system.

Results obtained by simulation indicate voltages between internal foils and an external housing effect nonlinear voltage distributions in a high voltage capacitor element model.

4.4.2 Electrical Stress

Although the model employed in this study is highly simplistic, some indicative conclusions are inferred:

- field stresses are strongest around terminal connections where elements connect to the bushings, and where bushings meet the unit housing;
- uneven stresses form between windings in response to voltage distributions across the elements in an interleaved, series-connected winding design, and are also affected by a nonlinear voltage distribution across capacitive elements, further contributing to their unevenness;
- foil edges and sharp points give rise to stress, suggesting support for a practise of folding foils at edges and corners;
- the size of windings can influence electric field stresses within a unit by virtue of the turn radius as element turns contribute to field stress, which could be reduced by increased fold bend radius;

- dielectric margins where winding foils interleave show increased electric field stress around the underlying foil, and these sites also constitute edges of each terminal;
- draw leads cultivate stress along the entire length of a unit, so designs with an even number of windings are preferred to avoid the need for one; and
- electric field visualisations suggest parasitic capacitances vary according to element position within a unit.

Coupling between `Electrical Circuit` and `Electric Currents` interfaces assume voltages are with respect to ground, so `Voltmeter` values are therefore first converted to an `Electrostatics` interface to be better understood in a geometrical context. Should future models adopt more direct couplings between these interfaces, they could be used directly. *Configurability* is central to iterative development approaches since changes are common where new information or ideas become available, or where prototypes are constructed which a simulation study should match. Parameterisation of a model `Geometry` can facilitate configurability, and make changes easier, particularly where parameters apply also to `Definitions` and where those `Definitions` form points of reference to apply `Materials`, a `Mesh`, and boundary condition interfaces.

At present, COMSOL® allows parameters to be used in *conditional* statements which aids flexibility in a model, but should *iteration* be additionally supported in creating a `Geometry` much more detailed and sophisticated models could be built. Furthermore, were control flow to be supported not only for `Geometry` sequences but also for `Definitions` in the form of both conditional and iterative declarations (analogous to '`if`' and '`loop`' statements in software), then parameters can be used to influence the shape and nature of a `Geometry` and its applicable `Definitions`, and therefore also the manner in which `Materials`, `Mesh` sequences, and boundary condition interfaces are applied. This would allow the possibility of developing detailed models resilient to change, which can be made easily configurable through parameters and hence of greater value to COMSOL®'s `Application Builder` functionality, too.

4.4.3 Multiple Units

Two studies presented in Section 4.3 highlight heightened stresses on dielectrics in capacitor units on the edge of a rack, due to nonlinear voltage distributions which extend throughout series connected units in addition to greater potential differences between internal foils and a common rack voltage shared by parallel capacitor unit strings. This has implications for asset design and an understanding of likely failure locations.

The following conclusions can be drawn:

- nonlinear distributions extend throughout multiple series connected units, such that element stresses are therefore unevenly distributed among units too;
- voltage distributions are influenced by a point of inflection due to presence of a rack tie, where each side of a rack then experiences a difference in distributed voltage;
- this disparity between otherwise similar unit strings exacerbates element stresses on the side hosting a rack tie;

- rack ties and bushing connections are likely to see increased propensity for stress, not only due to their proximity to housing and internal connections but since rack ties form a mid-point where stresses on element dielectrics are large; and
- corner units (those furthest in voltage from the rack) and units which connect to high voltage terminals are most susceptible to dielectric degradation and faults.

4.4.4 Further Work

Layered Materials

COMSOL® 5.5 supports layered materials, which could be an efficient way to model capacitor windings in future, in terms of: ease of model configuration; enhanced detail; and possibly more efficient computation. This approach would define foils and dielectric layers between them, to create a layered material for use as coherent domains in the `Geometry`, for improved model representativeness, meshing, and which might furthermore allow both solid and fluid dielectrics to be modelled. With more efficient modelling, or otherwise greater computational resources, expanding models to include multiple racks of a bank can allow the distribution of energy between racks to be explored.

Larger Models

Work remains to quantify voltage distributions for more detailed simulation models, and to explore the relationship between element failures on one side of a bridge arrangement and subsequent increased stresses on the other due to tiered rack voltages shared by both sides.

Moreover, stresses are found to be greatest on corner units of a bank, but where it is assumed that voltages distributed on racks themselves are similar in a healthy case: it might instead be the case that some rack levels experience greater overall voltages. Should it be practical to improve simulation efficiency or employ greater computational resources, more of a bank structure can be modelled in detail in future work.

Model Simplifications

The models studied in this Chapter are highly simplistic and do not account for many aspects of operational capacitor units. Should more representative model designs be considered for additional investigations, further work could build on improved models and in turn consider variation of a given design within manufacturer tolerances.

Chapter 5

Thermal Effects

This Chapter explores thermal properties and behaviours of capacitors through modelling to complement electrical measurands in support of diagnostics, given the critical influence of temperature on capacitor dielectric longevity. It investigates detectability of an internal heat source from distributions of temperature on a capacitor unit's outer housing. A high voltage capacitor unit model simulates propagation of thermal energy from heat sources (which emulate internal incipient faults at select positions within a unit) to a steel enclosure. Resultant heat distributions are shown to be readily influenced by foil arrangement and heat source location within a capacitor, since foils both contain and direct heat.

It considers the influence of winding alignment in Section 5.3, the influence of fault position in Section 5.4, and investigates the limits of fault detectability according to source power in Section 5.5, before proposing future research directions in Section 5.6. Discussion and conclusions are then respectively offered in Sections 5.7 and 5.8.

5.1 Introduction

Capacitor units are generally considered relatively reliable components within a wider system, but even modest failure rates could become cumbersome when scaled to match the number of components employed for high-capacity network applications. Since units are unrepairable [17], early detection and tracking for developing faults and asset degradation is preferred. As discussed in Section 2.2.2, capacitor banks' inherent modularity is practical for manufacture and construction, but additionally lends itself to targeted maintenance should fault locations be known. Accordingly, this Chapter explores how temperature associated with partial discharge and series element faults behaves, and in particular how the location of a point heat source influences thermal profiles of a unit's external housing, by employing a simulated model of a high voltage capacitor unit (as illustrated by Figure 5.1 which highlights internal elements).

Prior to fuse operation, insulation defects can invite partial discharge to deteriorate a dielectric. Partial discharges act as heat sources, and while they can occur anywhere within the dielectric of a unit are more probable under points of high electrical stress (such as foil corners, edges, leads, and points of connection).

For every 8 to 10 °C increase in temperature, the expected useful life of a capacitor unit halves [86], [1]. This applies to capacitor units operating in a 40 to 65 °C operating range [87], and may vary according to internal design, unit size, surface and environmental factors. Ca-

pacitor dielectrics generate heat under normal conditions: [88] estimates the temperature rise which occurs throughout the dielectric volume of a capacitor unit, and shows it relates to the value of Tan δ; and [50] considers the transfer of heat within individual self-healing capacitor elements used in pulsed power applications. Hotspots arise within units as a consequence of sealed designs [52], and due to heating under electrical load. Furthermore, by generating heat, these dielectric losses could mask any signal from a heat source of particular interest within a capacitor unit, such as sites of partial discharge or weak electrical connections.

An ability to detect incipient faults and series element failure or locate them within a bank of capacitor units could allow degradation to be monitored, would reveal faults on parallel arms of a bridge arrangement, and could inform operation and maintenance decisions, principally to allow select units to be prioritised for replacement based on an inference of their health. As discussed in Chapter 4, timely interventions informed by an awareness of unit health could alleviate stresses and thereby offer preventative means of fault mitigation.

5.2 Simulation Studies

As discussed in Chapter 3, a configurable capacitor unit model (Figure 5.1) has been developed in COMSOL Multiphysics® and can be set to simulate foils stacked according to either an X, Y, or Z axis[1]. For studies in Section 5.3 and Section 5.4, it is configured with 3 groups of 6 capacitive elements comprising discrete aluminium foils, each wound 4 turns with dielectric of relative permittivity $\epsilon_r = 3.9$ to mimic mineral oil and layers of Kraft paper. The model is additionally configured with:

- a boundary condition to emulate `Heat Transfer in Solids` specifies `Thermal Insulation` for unit materials (including element foils);
- a 10.45 $Wm^{-2}K^{-1}$ heat transfer coefficient within a `Heat Flux` node allows unit housing to cool in a room temperature environment;
- 50 Wm^{-3} of heat is applied with a volumetric `Heat Source` throughout dielectrics to emulate dielectric losses; and
- a 0.2 W `Point Heat Source` is specified at each of a set of distinct locations.

For differing foil configurations, `Cut Line`s are specified on each face of the unit housing to act as a common point of reference and to allow heat profiles to be graphed. These are illustrated in Figure 5.2, which names housing faces and `Cut Line`s relative to model axes. Another `Cut Line` (dotted, in Figure 5.2) centrally bisects all foils in the direction of the X-axis, as this foil alignment is chosen for studies in Section 5.4. This capacitor model is simplistic in many ways, but is primarily intended to illustrate how heat is distributed on outer housing (where it might then be possible to detect patterns or changes in these distributions).

[1] As outlined in Figure 2.2, 'alignment' is described by the axis orthogonal to flat winding surfaces.

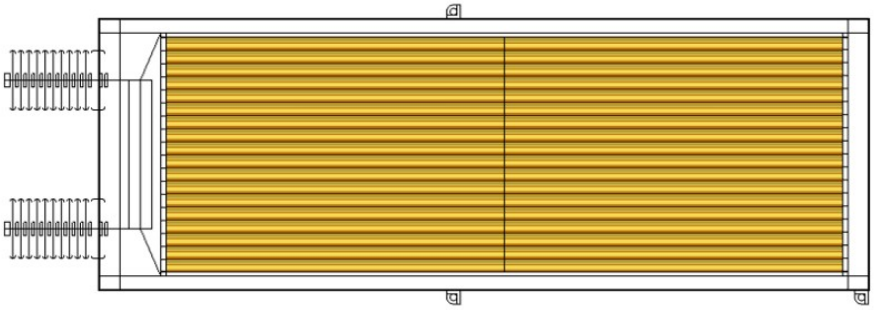

Figure 5.1: A capacitor model with foils aligned to the X-axis

5.2.1 Studied Design Aspects

This Chapter presents three studies, each exploring effects of a different *model or design aspect*.

1. Firstly, each of the three *alignment configurations* is studied for a `Point Heat Source` in the centre of the unit. Using model parameters to reconfigure the model for foils aligned to each of the X-, Y-, and Z-axes, which are provided in Figure 5.3, respectively as Figure 5.3a, 5.3b, and 5.3c. Results are then generated for a `Point Heat Source` in the centre of a central winding for each configuration, which are presented in Figure 5.4.

2. Secondly, to avoid repetition, a unit with elements stacked along the X axis (vertically, as a unit would typically sit on a rack, and as shown in Figure 5.1) is simulated to explore the propagation of heat from *different source locations* within the unit, as in Section 5.4.1. This involves testing three such `Point Heat Source` locations:

   (a) a foil tab connection;

   (b) a second-from-last foil in an element roll; and

   (c) a centremost foil edge,

   where all apply to a central (9^{th}) element. These scenarios do not directly correspond to particular types of fault, but to a limited extent are intended to mimic partial discharge.

3. Thirdly, detectable temperature changes are explored as a result of *varied output powers* of a `Point Heat Source` at each location, to compare housing temperatures. This third study uses a different model configuration from the first and second: it considers a model with 9 element windings stacked in alignment with the Z axis, rather than 18 stacked in alignment with the X axis, and uses porcelain for bushings as detailed in Section 5.2.2.

5.2.2 Materials

Where `Porcelain` is used on bushings (as in Section 5.5), the following parameters are specified. • Thermal Conductivity: 1.5 $Wm^{-1}K^{-1}$. • Density: 2.4 $g\ cm^{-1} = 2400\ Kg\ m^{-3}$. • Specific heat capacity at constant pressure: 1070 $J\ Kg^{-1}K^{-1}$. • As a solid, the ratio of specific heats for porcelain is 1. The ratio of specific heat for `Transformer oil` is also assumed to be 1 for the purposes of these simplified studies of assets with constrained volume.

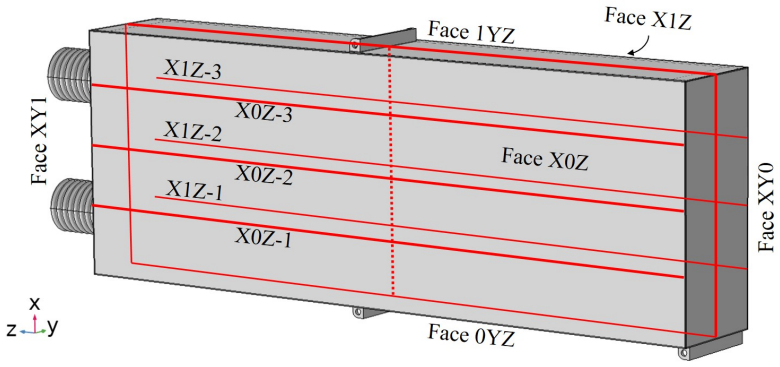

Figure 5.2: Cut lines specified in the model for postprocessing

5.3 The Influence of Winding Alignment

Figure 5.3 firstly illustrates the different winding alignment options achieved through modifying the `foilStackDirectionZ0X1Y2` parameter, which is designed to behave as an 'enumerated type'[2]. A 0.2 W source power[3] is used in each case, and result diagrams are introduced in Section 5.3.1 before results are presented and discussed in Section 5.3.2.

5.3.1 Diagrams

Values shown above colour legends in Figures 5.4, 5.6, 5.7, and 5.8 are maxima for the whole 3D model volume in each case, not only those temperatures displayed on the outermost surface: the hotspot of any unit occurs internally, and this maximum temperature is denoted at the top of the legend although each colour legend is itself normalised to a range from 293.15 to 294 K.

5.3.2 Results

Firstly, as seen in Figure 5.8, the 50 Wm^{-3} of volumetric heat (to represent dielectric losses as introduced in Section 5.2) effects a baseline symmetry in the temperature of a unit's housing[4]. This pattern is more clearly seen when a `Point Heat Source` is included at a similar position on a centremost foil of a central (9^{th}) element for a similar capacitor unit with differing internal winding alignments: Figure 5.4a, Figure 5.4b, and Figure 5.4c show foil alignment to significantly alter heat propagation characteristics, but these distinct cases share commonality. In each: • two faces (perpendicular to both foils' planar surfaces and turns) expend a lot of heat; • two (being parallel with turns but perpendicular to flat layers) are of medium heat; and • two (being in parallel with foil faces) remain comparatively cool. Thermally conductive aluminium foils direct thermal energy according to their alignment, toward the faces of a capacitor unit parallel to neither flat foil surfaces nor turns made as foils fold into elements.

As thermal energy follows winding alignment, should heat sinks be applied to reduce capacitor units' operating temperatures they would be best placed on the housing faces most proximal to internal winding terminals to have greatest effect in cooling element dielectrics. Hence, as foils direct heat, in all cases prior knowledge of a capacitor unit's internal design could significantly aid the interpretation of thermal measurements on a unit's housing.

[2] A type of variable supported by high-level programming languages
[3] 0.2 W is an arbitrary value as an initial baseline, where source powers are varied for studies in Section 5.5.
[4] This is further discussed in Section 5.5.2.

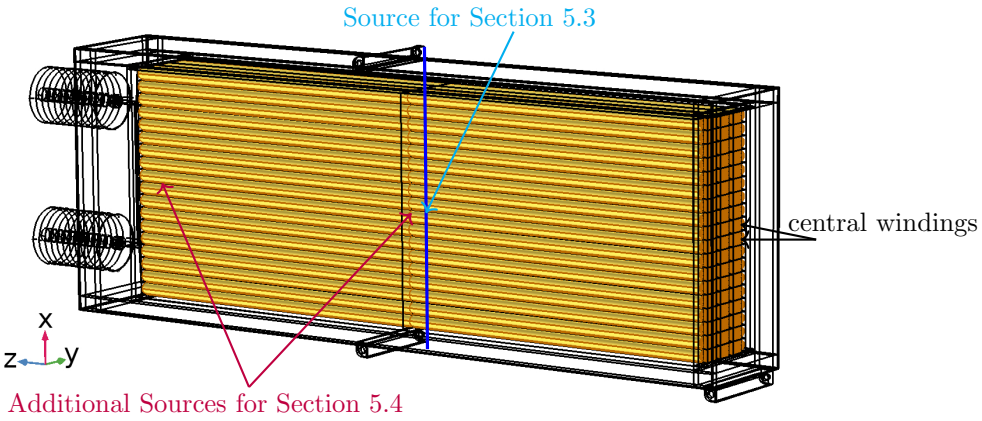

(a) Foils aligned to the X axis

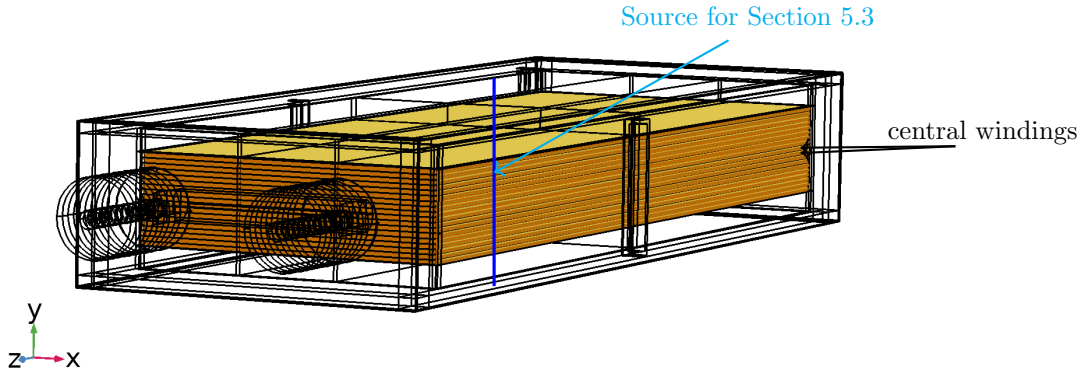

(b) Foils aligned to the Y axis (rotated to show winding turns)

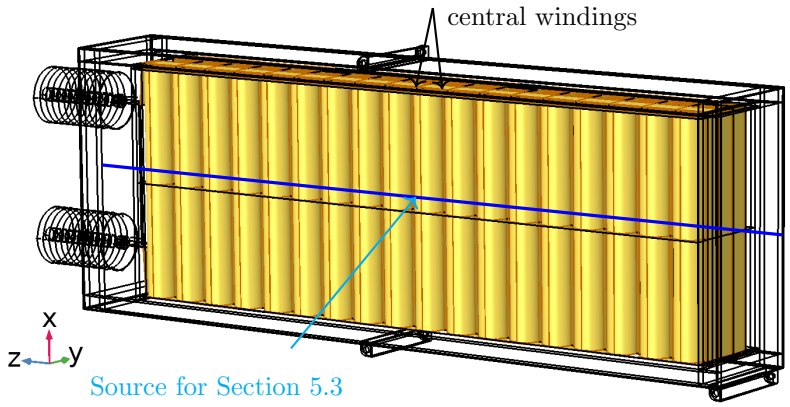

(c) Foils aligned to the Z axis

Figure 5.3: Three options for foil alignments (each showing a central bisector in blue)

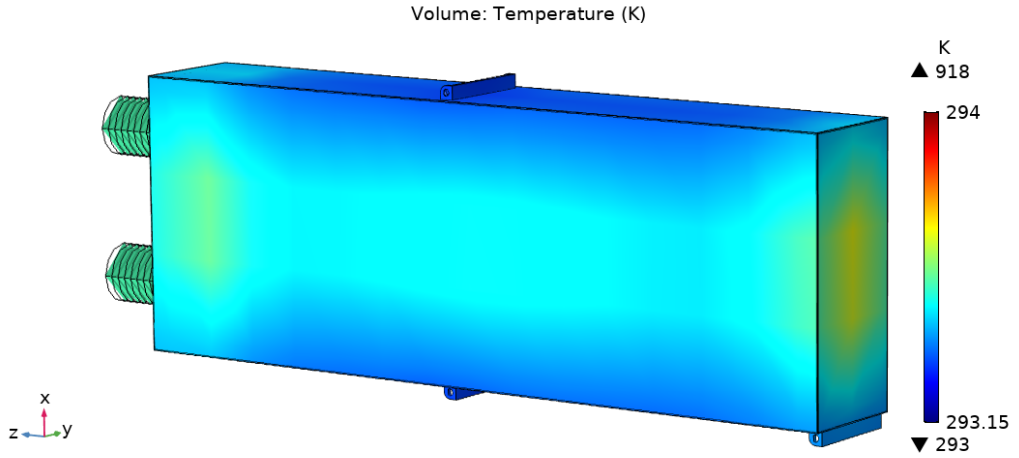

(a) X stacked - central heat source profile

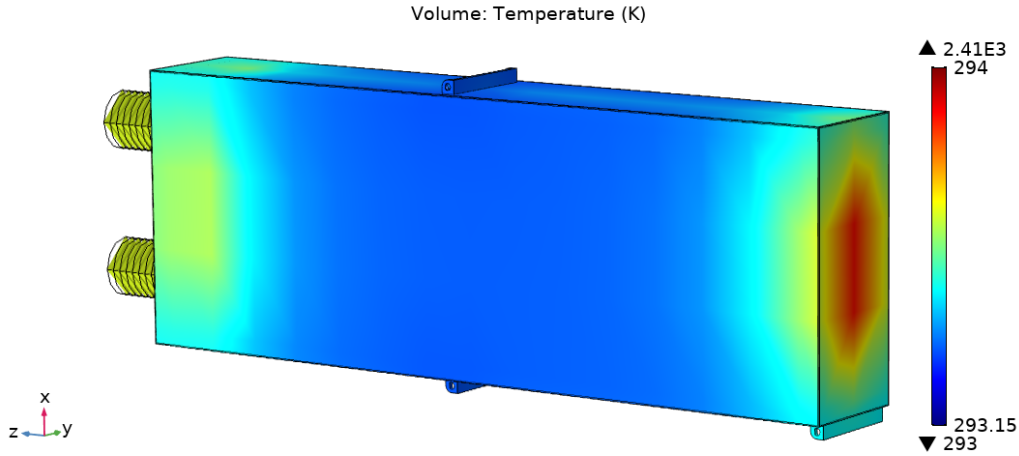

(b) Y stacked - central heat source profile

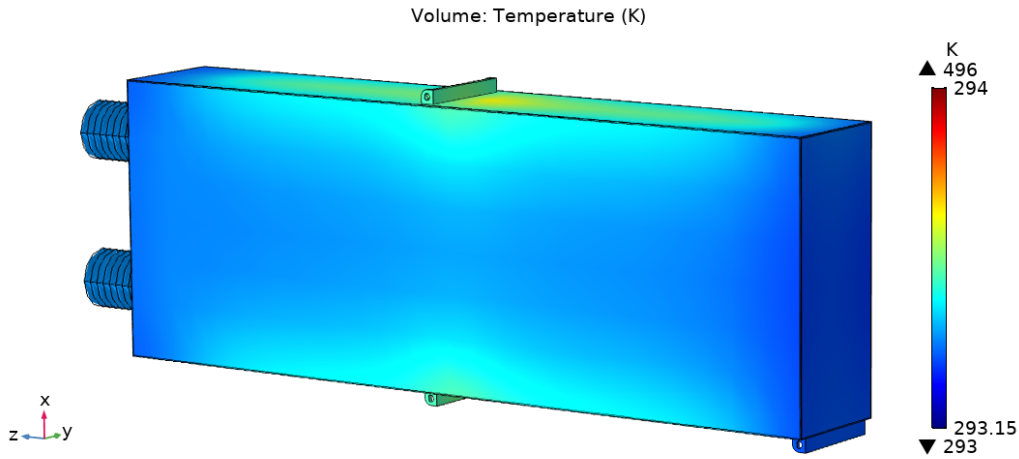

(c) Z stacked - central heat source profile

Figure 5.4: Thermal profiles differ with foil alignment

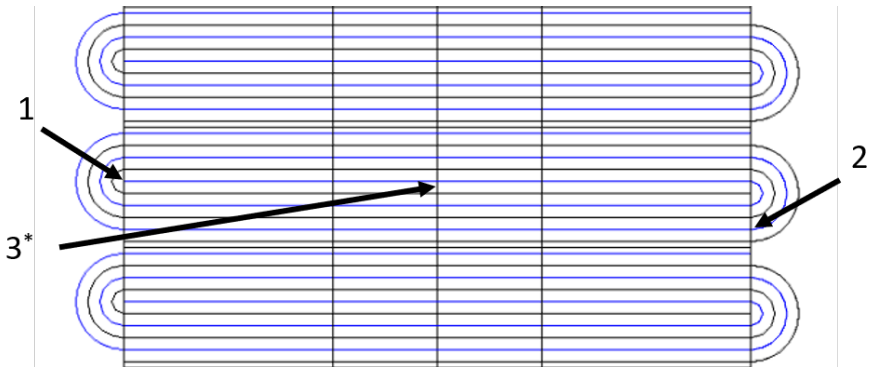

Figure 5.5: The positions of heat sources

5.4 The Influence of Fault Location

As explored in Section 5.3, winding alignment has a significant bearing on the variation of temperature on units' outer housing, all else being equal. This Section explores how such effects are influenced by the location of a heat source within a unit of fixed winding configuration.

5.4.1 Heat Source Positions

Figure 5.5 outlines three locations for point heat sources used, which give thermal profiles outlined in Figures 5.4a, 5.6, and 5.7, where Figure 5.8 is provided for comparison to show the effects of dielectric heating alone. These `Point Heat Source` locations[5] are:

1. Figure 5.4a – the edge of a (9^{th}) centre foil;
2. Figure 5.6 – a corner of a penultimate foil (near the outer edge of a winding); and
3. Figure 5.7 – a corner of a tab connector (at one end of a winding and the unit).

These are chosen so as to make use of existing points within the model `Geometry` as it is configured. It would otherwise be possible to amend the model `Geometry` to include additional `Cut Planes`, which could be positioned to intersect at desired points.

Keeping foil alignment constant, temperature is found to vary according to fault location. These visual representations are more clearly quantified in Figures 5.9, 5.10, and 5.11 according to the `Cut Lines` indicated in Figure 5.2. Figure 5.9, Figure 5.10, Figure 5.11 each include four graphs, which respectively show measurements across `Cut Lines` for a 0.2 W `Point Heat Source` on an edge of: 1. an innermost edge of a central winding; 2. a penultimate turn of a central winding; and 3. a tab connector, of an X-aligned foil unit.

5.4.2 Discussion

By comparison with Figure 5.4a, Figure 5.8 shows the case where no point heat source is present (to represent a 'healthy' scenario); Figure 5.7 highlights the asymmetry resulting from a heat source on a tab-connector at one end of a unit; and Figure 5.6 increased heat transfer from medium-temperature faces as a result of a penultimate-foil heat source location. It is also

[5]Locations identified in this list are numbered in accordance with those indicated in Figure 5.5.

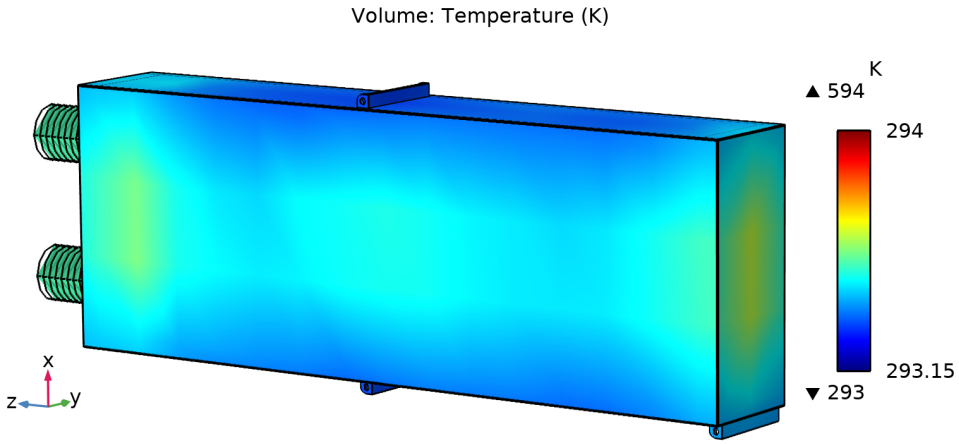

Figure 5.6: Point heat source on a penultimate foil

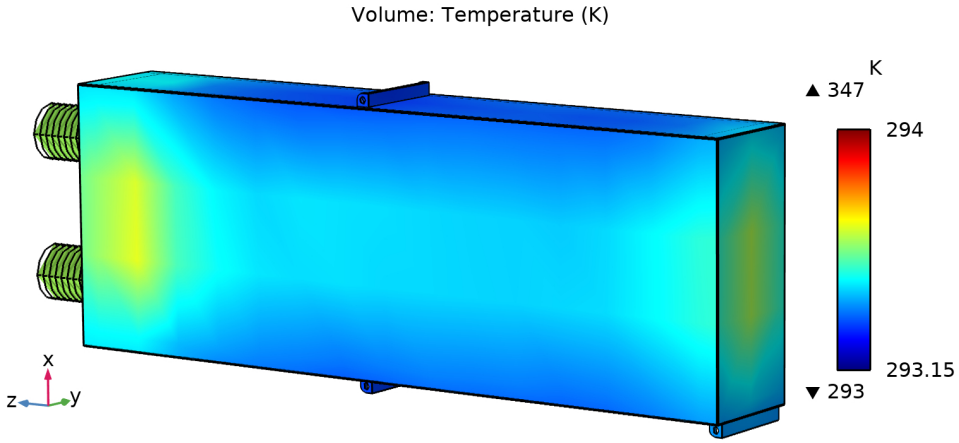

Figure 5.7: Point heat source on a tab connector

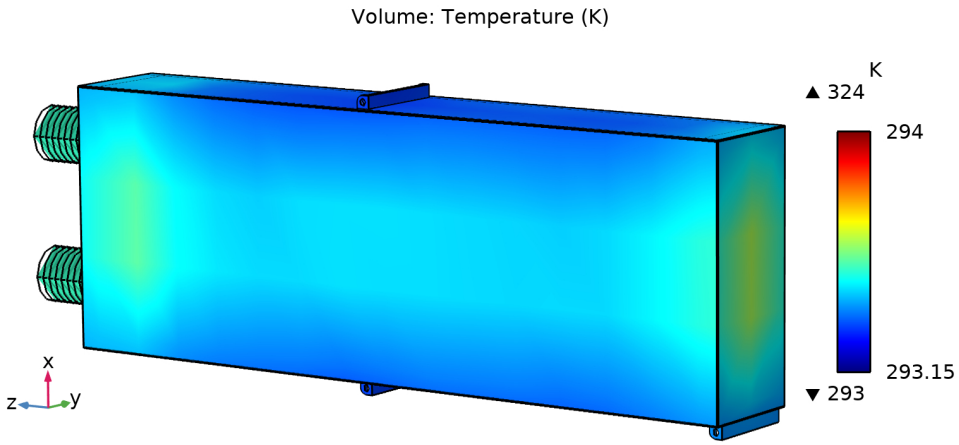

Figure 5.8: No point heat source

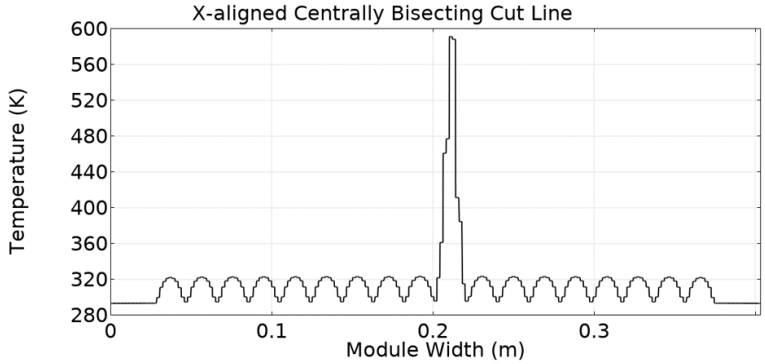

(a) Temperature along a centrally bisecting cut line for a central heat source

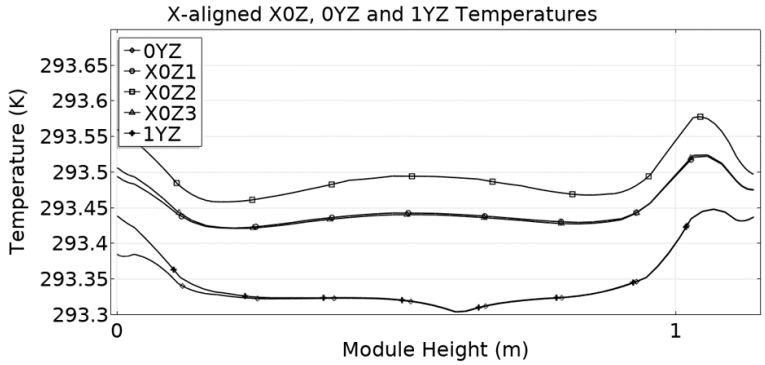

(b) Temperature along X0Z, 0YZ, and 1YZ cut lines for a central heat source

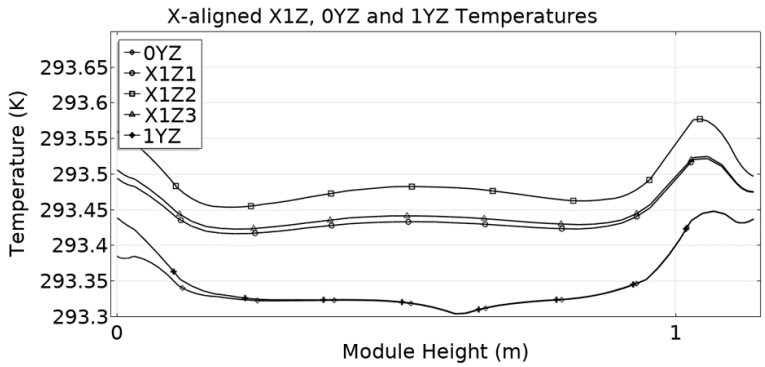

(c) Temperature along X1Z, 0YZ, and 1YZ cut lines for a central heat source

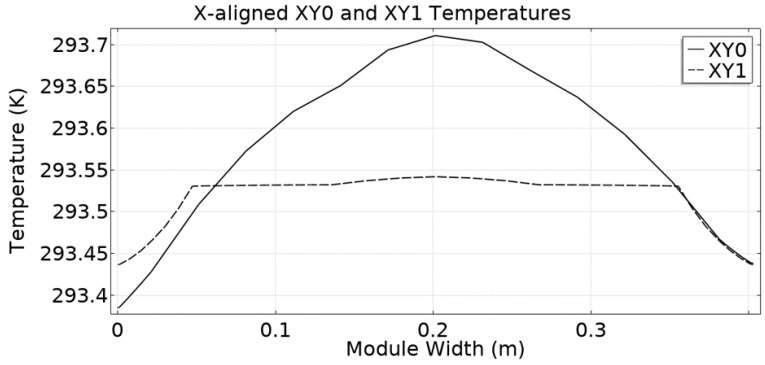

(d) Temperature along XY0 and XY1 cut lines for a central heat source

Figure 5.9: Central heat source profiles

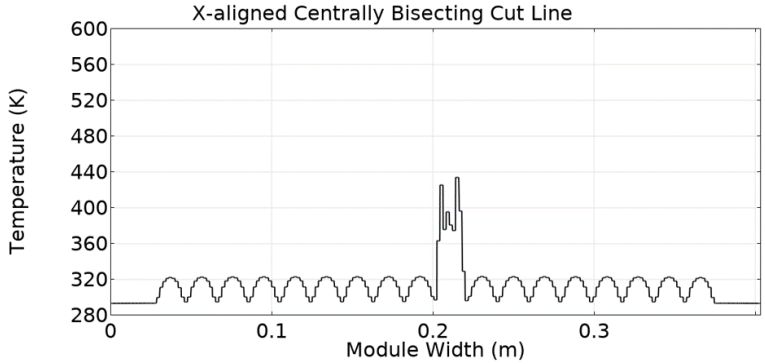

(a) Temperature along a centrally bisecting cut line for a penultimate foil heat source

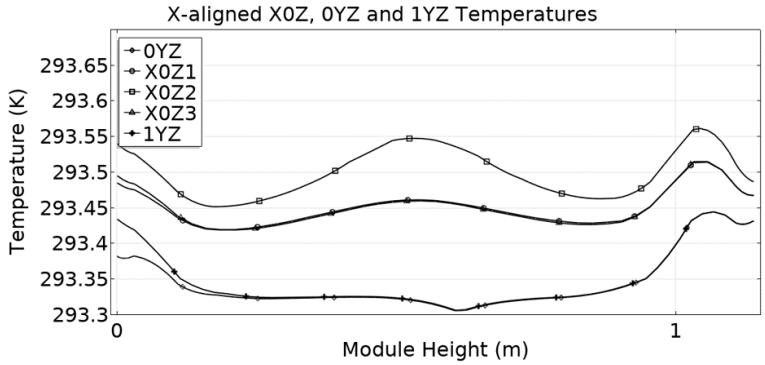

(b) Temperature along X0Z, 0YZ, and 1YZ cut lines for a penultimate foil heat source

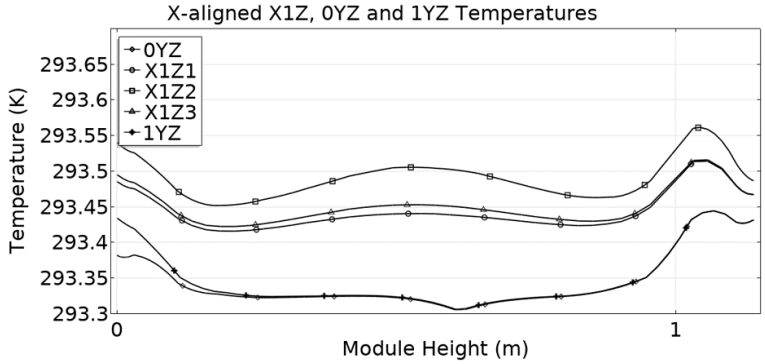

(c) Temperature along X1Z, 0YZ, and 1YZ cut lines for a penultimate foil heat source

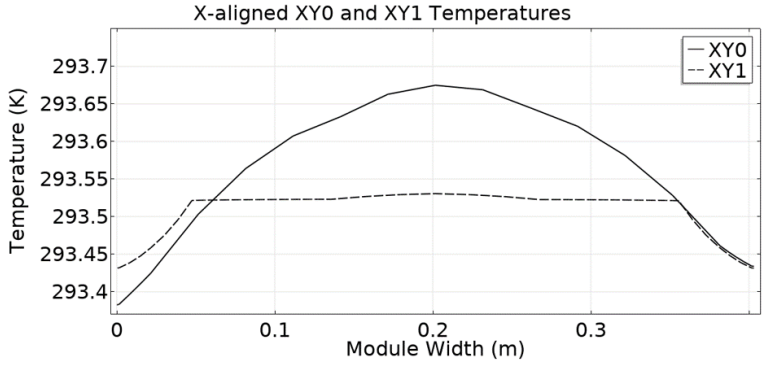

(d) Temperature along XY0 and XY1 cut lines for a penultimate foil heat source

Figure 5.10: Penultimate foil heat source profiles

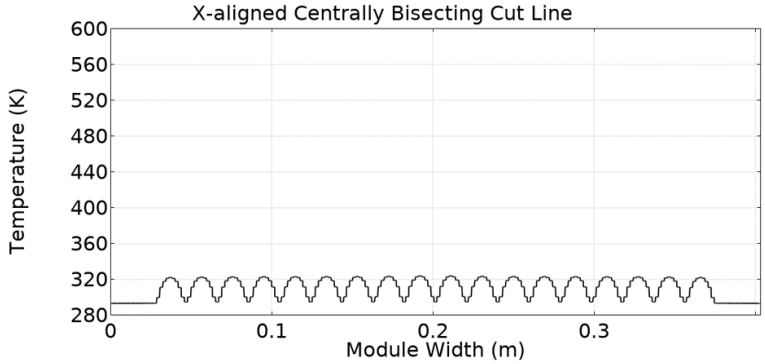

(a) Temperature along a centrally bisecting cut line for heat on a tab connector

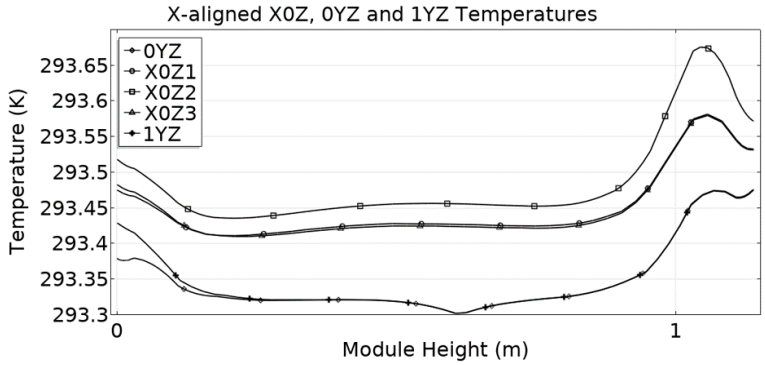

(b) Temperature along X0Z, 0YZ, and 1YZ cut lines for heat on a tab connector

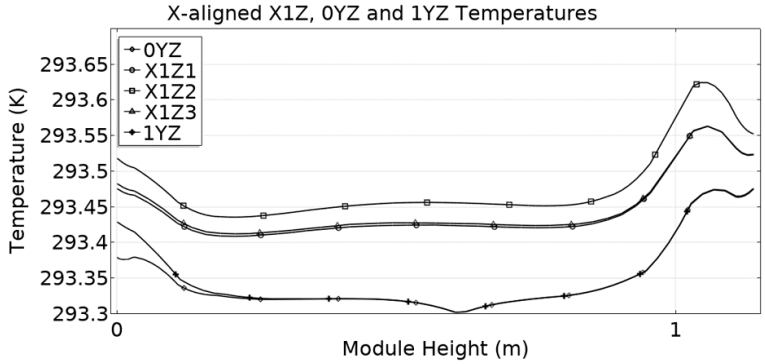

(c) Temperature along X1Z, 0YZ, and 1YZ cut lines for heat on a tab connector

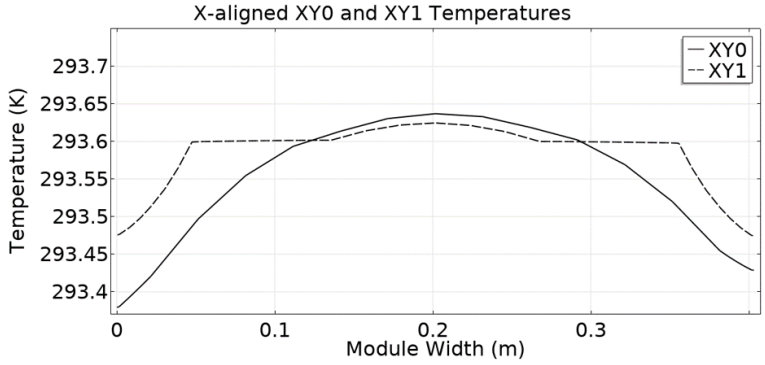

(d) Temperature along XY0 and XY1 cut lines for heat on a tab connector

Figure 5.11: Tab connector heat source profiles

notable that the greatest hotspot temperature occurs, intuitively, for a thermal fault in the centre of a winding, and the least occurs on a tab-connector outwith element windings.

This is also evident in Figure 5.9a (a centre foil) when compared with Figure 5.10a (a penultimate foil), where the maximum temperature in the former is more than twice that of the latter. The lack of an equivalent spike for Figure 5.11a suggests these peaks occur since foils *contain* heat within windings due to the low thermal conductivity of the dielectric layers relative to the higher thermal conductivity of electrodes in discrete foil capacitor designs.

Central and penultimate foil positions are indicative of two extremes of a continuous scale of source position within a winding. These results indicate that the more central a source of thermal energy, the greater the number of dielectric layers (and equivalently the more maximal the length of conductor) through which this energy must dissipate; conversely, heat generated at penultimate foil locations must traverse a minimum of winding layers, short of positions on the extremeties or external to windings entirely (indicated by those on tab-connectors). Hence, temperatures are higher the more central the heat source (fault) is within a winding.

Temperature increases are also less pronounced on faces `X0Z` and `X1Z` in Figure 5.9. For a central heat source, more energy is contained by foils such that it cannot as easily spread throughout an element, as indicated by a pronounced spike in Figure 5.10, and less readily reaches the housing. Thus, thermally insulating aluminium foils contain heat within elements.

5.5 Identifying Limits

The model's parameterised configurability as introduced in Chapter 2, is used to investigate the influence of the number of winding turns used in a 9-winding (10-terminal) discrete foil unit with windings aligned with the Z axis (Figure 5.12). The following Figures (Figure 5.13 and those in Appendix A) offer results of studying this reconfigured model design.

5.5.1 Results

Figure 5.13 provides a set of results for a unit with 10.45 $Wm^{-2}K^{-1}$ external heat flux, 50 Wm^{-1} dielectric heating, and without any active `Point Heat Source`s, to resemble a 'healthy case' scenario. Studies repeat for varied 1. model turns, 2. heat source (a) power, and (b) location.

Results in Figure 5.13 offer a point of reference and a means of comparison to the associated result sets provided in Appendix A which are indexed in Table 5.2 according to the number of winding turns in each model and the position of thermal heat sources[6]. To facilitate comparison, most of the axes are ranged similarly for a set of different scenarios, as outlined in Table 5.1.

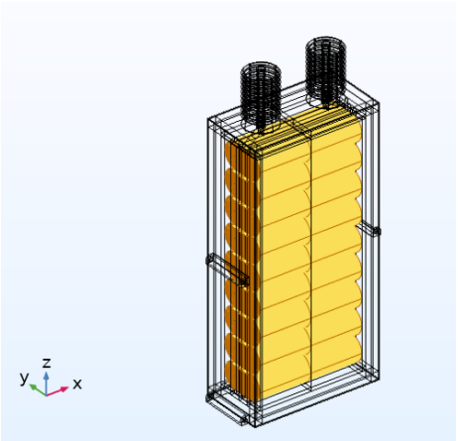

Figure 5.12: A 9-winding, 40-turn model

[6]Each Figure in the Appendix refers to Table 5.2 for ease of navigation between results sets and this Table in the digital version of this document, since all Table and Figure references link the reader to their subject.

5.5.2 Discussion

The most interesting subfigure in Figure 5.13 is that on the bottom right, which shows the temperature along a line through the centre of the unit: this temperature increases within each foil winding due to the combination of applied dielectric heating (in an attempt to represent dielectric losses) and the foil being specified as a form of `Thermal Insulation`. Since, in the model used for the studies in this Section 5.5, only dielectrics within windings have a heat source applied to represent dielectric losses (stray capacitances discussed in Chapter 4 are present, but are not considered sufficiently significant to be modelled in this `Study`), it is natural that temperatures within each winding are greater than those between the windings and housing. However, there is also an attenuation in the winding temperatures within the outer layers of each winding and also in windings at either end of the stack (those at the top and bottom of the unit). While temperatures along this central bisector are attenuated toward either end (due to increased surface area to volume ratio at top and bottom of the unit), the reason for increases at the extremes of cut lines `XY0` and `XY1` (which are parallel to windings) are due to the lateral propagation of heat as directed by wound foils. Thus, these two findings – that foil windings both *contain* and *direct* heat – are both evident even in the muted results associated with a healthy scenario. Moreover, temperatures along the sides of a unit's outer housing increase toward the ends of the unit in this scenario which only models dielectric losses, and which is indicative of the improved heat dissipation (due to increased surface area to volume ratio) at the top and bottom of a healthy unit.

From results in Appendix A, it can be seen that changing from a 20-turn and a 40-turn model leads to a substantial increase in the hotspot temperatures (shown as maxima at the top of colour legends in (a) Figures), and through the centre of the unit as seen in (e) Figures.

The penultimate heat source leads to a broader heat distribution on XY0 and XY1 faces, and as in the previous sets of studies also produces two peaks in `Cut Line`s through units' centre as heat propagates around some of the outermost layers of a winding. There is also a difference between Figures (c) and (d) for each penultimate-turn point heat source scenario, where these Figures broadly resemble one another for other source location scenarios. This is brought about by the location being on a *penultimate* winding turn, and therefore more readily facilitating thermal dissipation outward from where the winding dielectric ends on the opposite side of the unit from the applied heat source. To an extent, this effect is a result of the winding design used in this model, but the principle that heat will escape more readily to one side of a unit than to the other is likely to apply wherever a unit has winding layers which end on alternate sides of the winding. Manufacturers could take advantage of this principle by designing windings which intentionally terminate dielectric and foil layers in opposite directions in order to exaggerate this thermal difference, since doing so could also make temperature differences on outer housings noticeable enough to be useful for future diagnostics and monitoring.

In the case of thermal faults occurring outwith any winding (such as modelled here on a connection between windings) temperatures through a unit's centre line are again low enough that dielectric losses can be seen as an influence on winding temperatures as in (e) Figures. It is notable that in such cases, the outermost turns of windings adjacent to those with a fault see a subtle increase which is abnormal in comparison to central-source scenarios, indicating that the outermost turns of windings can be affected by excess heat from faults nearby but external to them. Scenarios where a connector fault is simulated are most obviously characterised by

Table 5.1: Format description for Figure 5.13

Index	Position	Description	Range (K)
(a)	Top Left	Outer housing	293 to 299
(b)	Top Right	`XY0` and `XY1` Faces	293.16 to 293.44
(c)	Centre Left	`X0Z` Face	293.1 to 295.2
(d)	Centre Right	`X1Z` Face	293.1 to 295.2
(e)	Lower Left	`0YZ` and `1YZ` Faces	293 to 298.5
(f)	Lower Right	Central bisecting line	Variable

Table 5.2: Results of winding turns, source power, and position

Position	Power (W)	20 Turn Model	40 Turn Model
Centre	0.2	Figure A.1	Figure A.13
	0.6	Figure A.2	Figure A.14
	1	Figure A.3	Figure A.15
	1.4	Figure A.4	Figure A.16
Penultimate	0.2	Figure A.5	Figure A.17
	0.6	Figure A.6	Figure A.18
	1	Figure A.7	Figure A.19
	1.4	Figure A.8	Figure A.20
Outer	0.2	Figure A.9	Figure A.21
	0.6	Figure A.10	Figure A.22
	1	Figure A.11	Figure A.23
	1.4	Figure A.12	Figure A.24

asymmetry in (a) and (b) Figures and by much greater outer housing temperatures, which is intuitive as most heat can readily reach housing and does so closest to the fault itself.

Temperature sensors such as the DS1920 [89][7] can take measurements at a resolution of 0.5 K, which implies that even for cases where `Point Heat Source` output powers of 0.2 W are detectable, but only if sensors are located where peaks on the housing temperature arise. Hence while 0.2 W sources are the minimum simulated for studies in this Section, these might in practise not be a minimum limit depending on fault locations and the resolution of temperature sensors employed. Results in Appendix A might nevertheless offer insight into thermal distributions to inform future work or otherwise permit inferences made by measurements beyond minimum detection thresholds. The greater the thermal source power of an internal fault, the more readily it could be detected in the case where there is a more general distribution of such sensors. However, should designs be engineered to cultivate areas of probable temperature rise in which sensors can be located, the more sensitive such detection techniques will be.

5.6 Further Work

In view of the simplistic nature of the present model (discussed further in Section 5.7), the potentially complex nature of capacitor units' operating environments, and although initial results suggest temperature to be a useful measurand to detect and even locate incipient faults, further work is required to assess these signals' detectability in realistic ambient environments.

[7]The DS1920 is chosen as an example in part due to its support for the '1-Wire' protocol. This capability simplifies configuration of a series of these devices for distributed measurements.

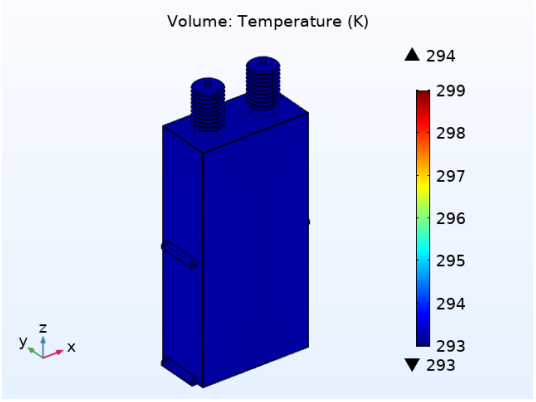

(a) Housing temperature

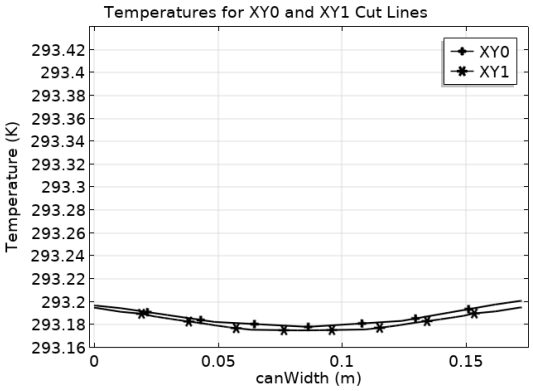

(b) XY0 and XY1 cut lines

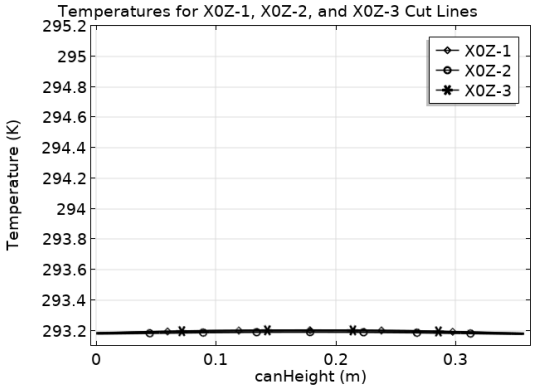

(c) X0Z cut lines

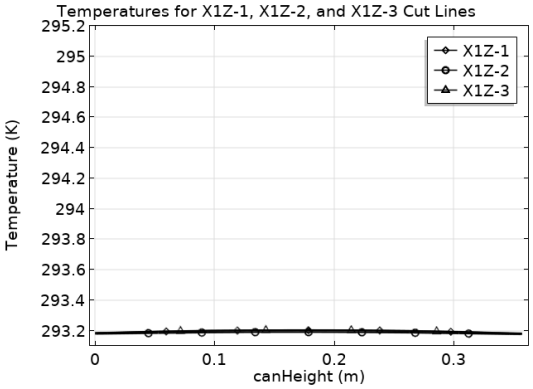

(d) X1Z cut lines

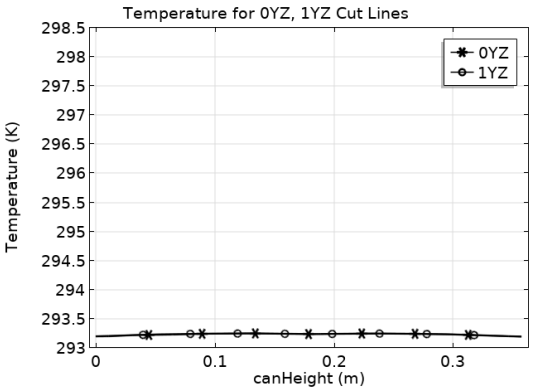

(e) 0YZ and 1YZ cut lines

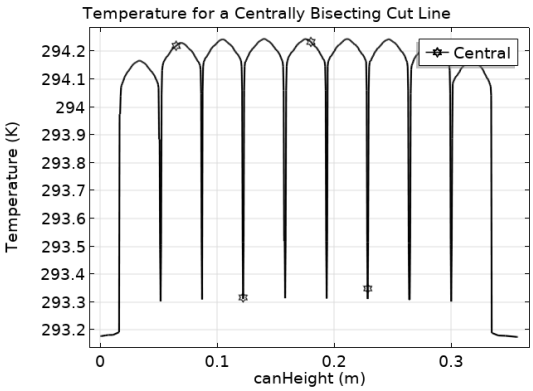

(f) A centrally bisecting cut line

Figure 5.13: 40T 0W of a set in Table 5.2

5.6.1 Simulation

Future work could consider how readily a heat source can be discerned from measurements affected by: 1. temperature sources from incipient faults; 2. ambient sources of heat; and 3. those intrinsic to the role of an asset within an electrical circuit, as (a) electrical loading and (b) the incidence of harmonics could both influence temperatures of healthy capacitor units.

Where the time and computational resources permit, it would be appropriate to use a more detailed simulation model for future investigations. With such a model, it would be interesting to quantify, through a series of repeated studies, the extent to which opposing faces differ in temperature in accordance with the relative position of a heat source on the winding. At one extreme (shown by penultimate foil studies) there is a clear disparity in temperature, and at the other extreme (shown by centre foil edge studies) any such disparity is minimal.

While these results indicate that there is a difference, it would be preferrable to explore in more detail the relationship for a greater resolution of points between these extremes.

Discharge Resistor Action

Discharge resistors are included in capacitor units as a safety device. A high proportion of safety incidents in high voltage research laboratories have occurred with capacitor units, as highlighted in Section 2.1.2. As resistors, discharging will result in resistive heating, which could be used to verify each discharge as another heat source internal to capacitor units.

Partitioned Capacitors for Heat Dissipation

In broad terms, capacitor unit housings are simple structures which contain an internal winding structure; provide a hermetic seal to contain dielectric fluid; and offer elements a degree of protection from the external environment. Unit housing is also the primary means of heat dissipation where the dielectric produces heat whenever an asset is online. Should the housing consist of protrusions into the internal dielectric, it could subdivide assets' internal elements (element groups), but in so doing also provide a more efficient means of heat extraction. Thus, capacitor housings could be designed to facilitate heat dissipation by maximising heat transfer using internal segmentation, and which could additionally influence convection of internal dielectric fluid such that fluid is contained to a particular group or region rather than permeate throughout a unit. As mentioned in Section 5.3, it is appropriate for any such partitions to be designed in alignment with units' natural heat dissipation. Such benefits could be in addition to improved resilience to high-frequency impulses and signal components.

Design for Detectability

Manufacturers could design units to improve fault detectability, such as by using internal partitions or alternating winding direction within units. As discussed for faults occurring on a winding's penultimate turn, manufacturers can make use of the disparity in consequent heat dissipation to facilitate diagnostics and monitoring by having different layers of winding dielectric terminate in opposing directions. There is potential to further improve fault locability by alternating the direction (clockwise or anti-clockwise) of adjacent windings which do terminate in opposing directions, such that capacitor units designed in this manner will have areas on

their outer housing for which a temperature rise is probable should a fault occur in either of a pair of such adjacent windings.

Like transformers, capacitor failures have potential to do considerable damage. Where inflammable dielectrics are used, undetected incipient faults have the potential to lead to catastrophic failures. To ensure safe energy dissipation in the event of sudden disconnection, discharge resistors within sealed units are sized to expend residual charge within 5 minutes of disconnection. Without sufficient safeguards, it would therefore be prudent to sanity-check energy dissipation through such internal discharge resistors were a means of detection available.

The thermal measurements discussed in this Chapter thus far provide one means of offering this capability: in view of the locations of discharge resistors, a drop in heat due to dielectric losses would be combined with an increased heat through discharge resistors, which would appear as point heat sources on the outer edge of internal windings and result in an asymmetrical thermal profiles in the immediate aftermath of electrical disconnection.

New Dielectrics

Natural-ester based dielectrics are increasingly sought for improved environmental impact over conventional mineral-oil based alternatives. As fluid dielectrics likely experience some natural convection effects due to temperature, these dielectrics could have an influence on unit heating both through differences in thermal conductivity and in viscosity (which affects any natural convection), which remains to be explored. In addition, research into dry dielectrics could alternatively result in designs which could be easier to handle, without risk of leak or the need to process fluids in manufacture and maintenance.

This would lead to an improved understanding of fault detectability from unit enclosure temperature patterns. It would provide insight into discernability of internal faults amid environmental and electrical influences, validate the principle of thermal propagation aligned with element foils using a physical model, and would suggest a preliminary sampling sensitivity and frequency necessary for a monitoring system to detect subtle incipient faults.

This research would further the reported state of the art, by advancing research toward a means of identifying incipient faults and dielectric failures in individual capacitor units. Thermal monitoring could allow capacitance to be dynamically evaluated under operational conditions, more accurate fault location, and could improve estimates of remaining life for these assets. Ultimately, an ability to identify faults on specific units could allow a set of choice (faulted) units to be replaced during routine planned outages, providing a means of early intervention to mitigate cascading failures.

5.7 Discussion

If a heat source occurs within an element, but not in the centre, the winding direction permits thermal energy to readily reach one intermediately heated face than it does the other. As such, two faces of a unit container affected by a point heat source will be intermediately heated to differing extents (as seen by different temperature increases in Figures 5.10b and 5.10c). A temperature increase on one intermediately warm face relative to the other suggests an approximate position of a heat source within an element, given prior knowledge of turn direction. Were manufacturers to alternate the foil fold direction of adjacent elements, it could aid accuracy

of location estimates which employ such information.

It is interesting to note the variation in hotspot temperatures shown from results of Section 5.3, which appear to increase for arrangements in which dielectrics are most narrow.

Compared with the scenarios in Figure 5.9a and in Figure 5.10a, Figure 5.11a shows no distinct temperature spike, suggesting elements remain comparatively cool. A markedly higher temperature is found across face `XY1` of the housing (where bushings are) and a correspondingly high temperature results in a pronounced signal toward the right of `X0Z` and `X1Z` line graphs. Hence, a heat source on a foil tab increases temperature at only one end of the capacitor: any temperature rise is asymmetrical; and no two opposing faces are heated.

It could be possible to locate a source of heat to: a capacitive element; a dielectric layer within the element; and to a location on the foil layer. Correspondingly, heat-emitting faults such as partial discharge could be detected (and to some extent located) within a capacitor unit should thermal profiles be measurable, such as with sensitive infra-red imaging, distributed temperature sensors, or heat-sensitive stickers on a unit housing. Localised dielectric heating can affect relative permittivity, and thereby alter the capacitance of an element within the unit. As such, any means of detection could also be corroborated electrically.

Simplifications

Relative to real, commercially available capacitor modules, this preliminary steady-state study makes simplifications: dielectric losses and environmental influences are omitted; a unit model is a simplistic approximation; and time taken for heat to propagate to the outermost enclosure is ignored but matters for intermittent heat sources. Instead, commercially available units vary in their: size and structure; internal design; dielectric materials; environment; and electrical loading. A homogenous dielectric is modelled, whereas in reality a dielectric would have solid and liquid components (which in turn could experience subtle convection effects), not to mention layers of material which could vary in direction of anisotropic properties from one layer to another. Results of this study are accordingly subject to more nuanced and realistic observations which require validation in a laboratory setting.

Convection

Fluid dielectrics in capacitor units might be less able to convect within the unit than it would be in transformers, which could lead to the possibility that warm dielectric fluid could accumulate at the top of a module and predispose certain elements to breakdown.

5.8 Conclusions

Thermal profiles across a unit housing are readily influenced by internal foil arrangement. Incipient faults such as partial discharge can be thermally detected, but knowledge of internal designs is prerequisite to identifying preferable measurement locations and data interpretation: manufacturers may therefore be best positioned to detect and locate unit faults toward reliability centred maintenance. As heat propagates according to foil alignment, an element containing a heat source can best be identified by temperature distributions on warmest housing faces; a foil surface (depth) within the element can be approximated from temperatures of intermediately heated faces, relative to those of the warmest faces and to one another; and heat sources

outwith foil elements can be recognised by overall asymmetric thermal profiles. While the simplistic studies presented in this Chapter omit realistic sources of thermal noise, they illustrate principles by which location of a subtle fault such as partial discharge can be inferred. Such inference would use knowledge of internal unit design and an understanding heat propagation to encode the location of its origin in thermal profiles of high-voltage capacitor units.

Chapter 6

Prototype

The practical work detailed in this Chapter has taken place in the David Tedford High Voltage laboratory at the University of Strathclyde. Initially, attempts were made to obtain a commercial capacitor unit with which to test the principles of operation and explore if electrical effects introduced in earlier Chapters are borne out in reality (through measurement), but commercial sensitivities around unit designs (and possibly the nature of an investigation into failure modes) precluded the availability of such an asset for this study.

To circumnavigate these understandable reservations, the practical work detailed in this Chapter has instead aimed to develop a prototype discrete-foil capacitor unit, at $\frac{1}{2}$ scale ($\equiv \frac{1}{8}^{th}$ volume) of a $735 \times 345 \times 175$ mm housing design[1] to match parameters typical of a Vishay unit [32] but with internal windings described in [29] (Figure 6). This practical work is motivated by a desire to demonstrate principles outlined in earlier Chapters (especially the role of parasitic field effects in provoking unique changes in terminal characteristics illustrated in Chapter 4) in a practical, laboratory environment, and furthermore acts as a basis for improved, albeit rudimentary, understanding of approaches applicable to component design and manufacture[2].

Outline designs are presented in Section 6.1, before Section 6.2 documents steps taken toward construction and brief learning from the process adopted. Section 6.4 then lists ambitions for continuation of the work and Section 6.5 draws preliminary conclusions.

6.1 Scale Model Design

Figure 6 in [29] shows an internal winding design with interleaved terminals. The unit is initially designed as a set of discrete components: • windings • an outer housing • a lid (including terminal connections) • internal connections • oil as dielectric, and • measurement devices.

6.1.1 Dimensions

Dimensions for a scale model of a capacitor unit can be selected from the formulae used to define the design of the chosen element. It should be noted that if a scale model is required

[1] The prototype design initially allowed for an additional 10 mm to accommodate internal bushing connections, but this was reduced to simplify the build leading to an eventual $715 \times 345 \times 175$ mm approximation. Similarly, housing widths have not been scaled down (from 1.5 mm steel [32], [58] to 0.75 mm steel) in preference of 1.5 mm aluminium due to the practical availability of this material during the build.

[2] This Chapter describes work undertaken in late 2019, early 2020, and in 2021 either side of a closure of university premises as a pandemic countermeasure.

to be reduced in size by a particular factor of an original design, it is worth considering the ability for an element design to scale by that factor: a 2-layer dielectric element would more easily scale in half than would a 3-layer dielectric design, for instance.

Terminology

Terminology used for designs in this Chapter (of which some are illustrated in Figure 6.4) are:

- *terminal* is one electrically continuous component to which a signal can be connected;
- *element* a single pair of foil terminals which form a capacitance;
- *series-element* a group of elements connected together in parallel (for such groups to in turn be connected in series);
- *capacitor unit* is a set of elements or element groups in shared housing, complete with bushings;
- *concentric* designs have foil terminals which are initially aligned along the same direction such that they can be wound from one end;
- *continued* designs are wound from the centre without central edges (the kind ultimately used in Section 6.2);
- *counterroll* describes a design in which foil terminals are initially aligned in opposite directions such that they can be wound from an overlap in the centre;
- *graded* refers to elements wherein dielectric layer thicknesses increase with turns;
- n is the number of turns taken in rolling;
- h_e is the height of an element winding;
- d_e is the overall depth of an element winding;
- h_f is the height of a foil layer, typically $12\mu m$;
- h_d is the height of a dielectric layer;
- d_f is the depth of the (flat) foil between turns;
- l_a is the length of a layer (of dielectric); and
- l_i is the length of a layer (of foil terminal).

Equations which define the lengths of Kraft paper and aluminium foil used and dimensions of each resulting element are deduced for each design, and dimensions are thereby delineated for a scale model of a capacitor unit.

6.1.2 Housing Design

Figures below illustrate designs for a unit housing and the overarching unit design. The next Section then considers winding designs in greater detail.

To make use of available materials: layers of 1.5 mm thick aluminium and 10×10 mm polyvinyl chloride (PVC), a design for unit housing and structure was developed as illustrated in Figure 6.1 and in Figure 6.2.

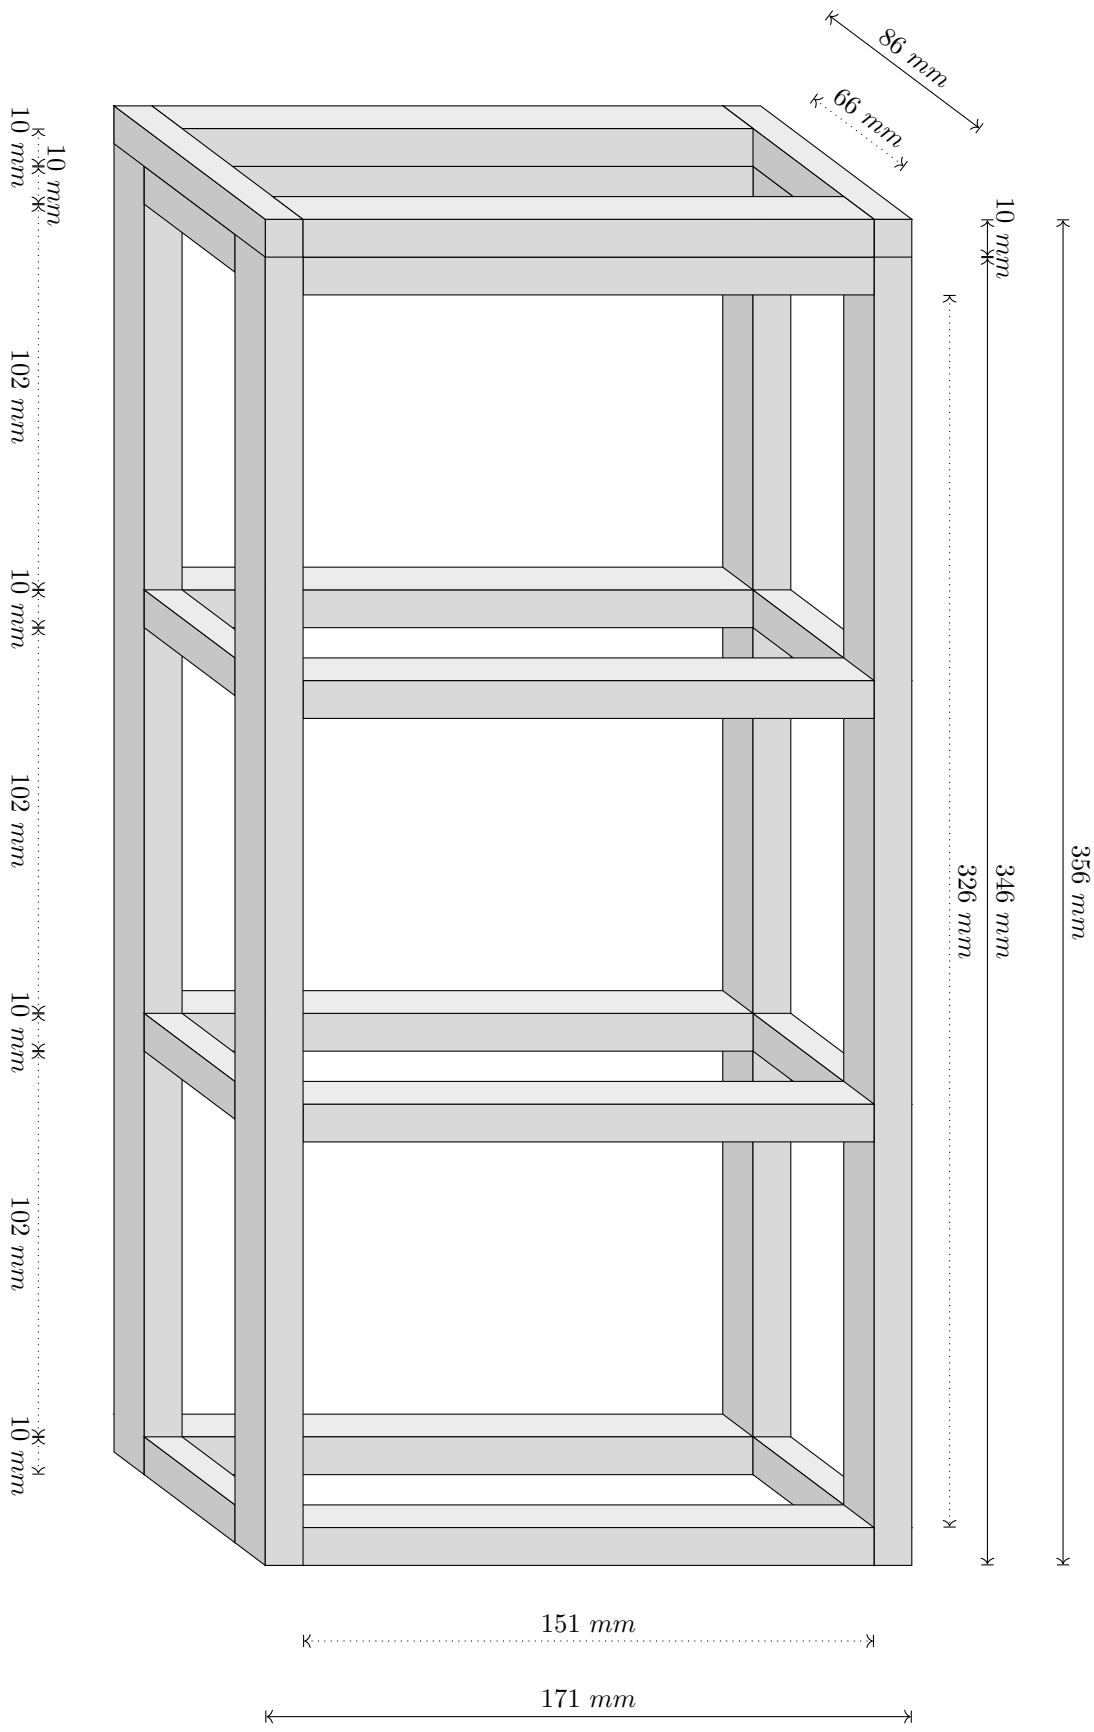

Figure 6.1: Frame for a scale model

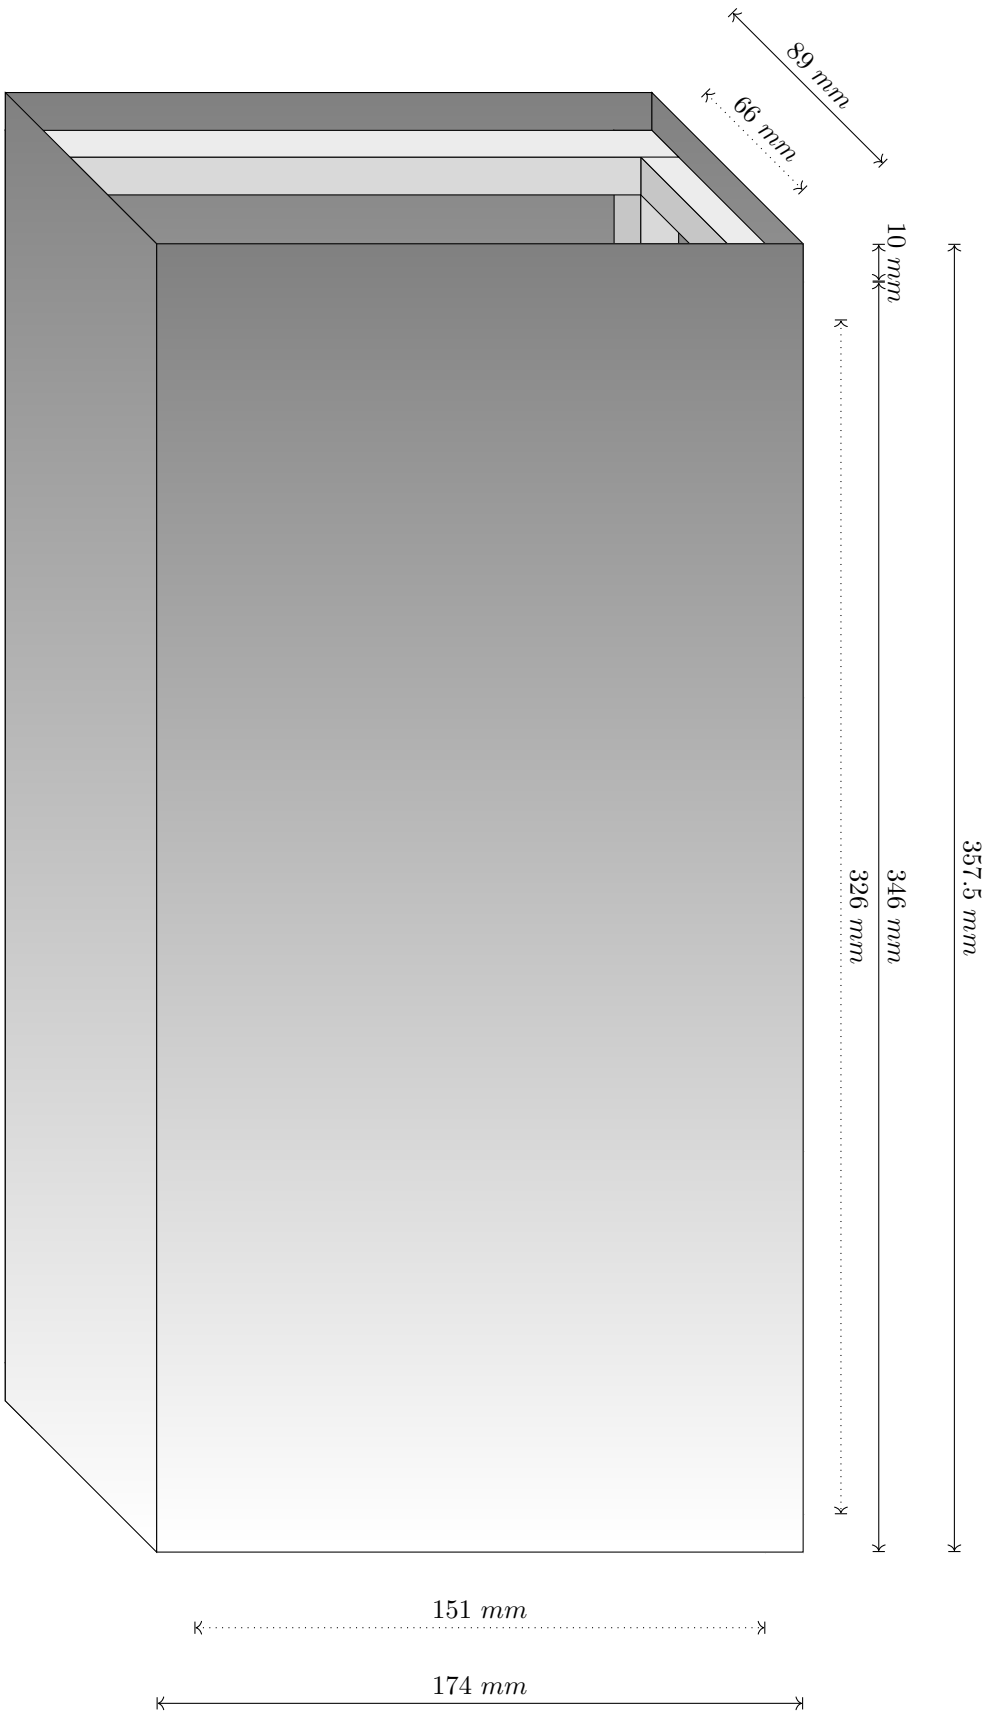

Figure 6.2: Scale model frame with enclosed sides

6.1.3 Winding Design Options

This Section explores a various winding design options, before reasons for those actually chosen for the prototype model are discussed in Section 6.1.4.

Windings can have designs which are discrete, where terminals are contiguous across each layer, or interleaved, where a terminal layer is broken by margins to form multiple smaller capacitances across the same dielectric. Interleaving accommodates higher voltages, and an example of an interleaved design is shown as a flattened cross-section in Figure 6.3, where 5 series capacitances are formed within a single winding from two layers of interleaved terminals.

Winding designs are presented in this Section with: 1. an illustrated cross-section of its constituent foils and dielectric before and after rolling (where arrowheads indicate a material continues); 2. equations which describe the width, height, depth of the winding; and 3. length of its constituent components in terms of the number of turns given.

For the purposes of this Section, all designs (excepting Figure 6.5 and graded designs) are illustrated 1. with two turns, where (a) 'height' refers to the vertical dimension on the page, (b) 'depth' to the horizontal, and where (c) 'width' is shown 'into the page', and is independent of other aspects of winding design so accordingly is omitted. The Equations included in this Section account for thickness of dielectric layers only, for which lengths are estimated according to an intermediate line cut straight through the centre of each. Maximum bend radii are assumed not to influence winding designs, but could in practise affect real dimensions. Terms used in these Equations are introduced in Section 6.1.1, and illustrated within Figure 6.4.

Capacitor unit windings are for the most part insulated from one another and are connected through • foil tabs, • schooped winding ends, or • single layers of exposed foil on adjacent windings. All are presumed to occur at lateral ends of a winding or otherwise not to affect the element cross-section. In the interests of comprehensiveness, some designs indicate windings with a foil layer exposed at either edge to allow windings to be series connected by stacking.

A 3-layer Concentric Element

Figure 6.4 shows a 3-layer concentric winding, which is also shown sized for the prototype design pursued in Section 6.2 and as a flattened form in Figure 6.5 and Figure 6.6 respectively. This element is described by Equations 6.1 to 6.7 in terms of n, where layers of the winding are labelled by vowels (i and o are innermost and outermost foils) as intimated in Figure 6.4.

$$h_e = (3h_d + 2h_f)(2n + 1) \tag{6.1}$$

$$d_e = d_f + (3h_d + 2h_f)\Big(2n + \frac{1}{2}\Big) \tag{6.2}$$

$$l_a = (2n + 1)d_f + \Big(\frac{6n^2 - n}{2}h_d + 2n(n - 1)h_f\Big)\pi \tag{6.3}$$

$$l_i = (2n + 1)d_f + \Big(\frac{6n^2 + n}{2}h_d + 2n(n - 1)h_f\Big)\pi \tag{6.4}$$

$$l_e = (2n + 1)d_f + \Big(\frac{6n^2 + 3n}{2}h_d + 2n^2h_f\Big)\pi \tag{6.5}$$

$$l_o = (2n + 1)d_f + \Big(\frac{6n^2 + 5n}{2}h_d + 2n(n + 1)h_f\Big)\pi \tag{6.6}$$

$$l_u = (2n + 1)d_f + \Big(\frac{6n^2 + 7n}{2}h_d + 2n(n + 1)h_f\Big)\pi \tag{6.7}$$

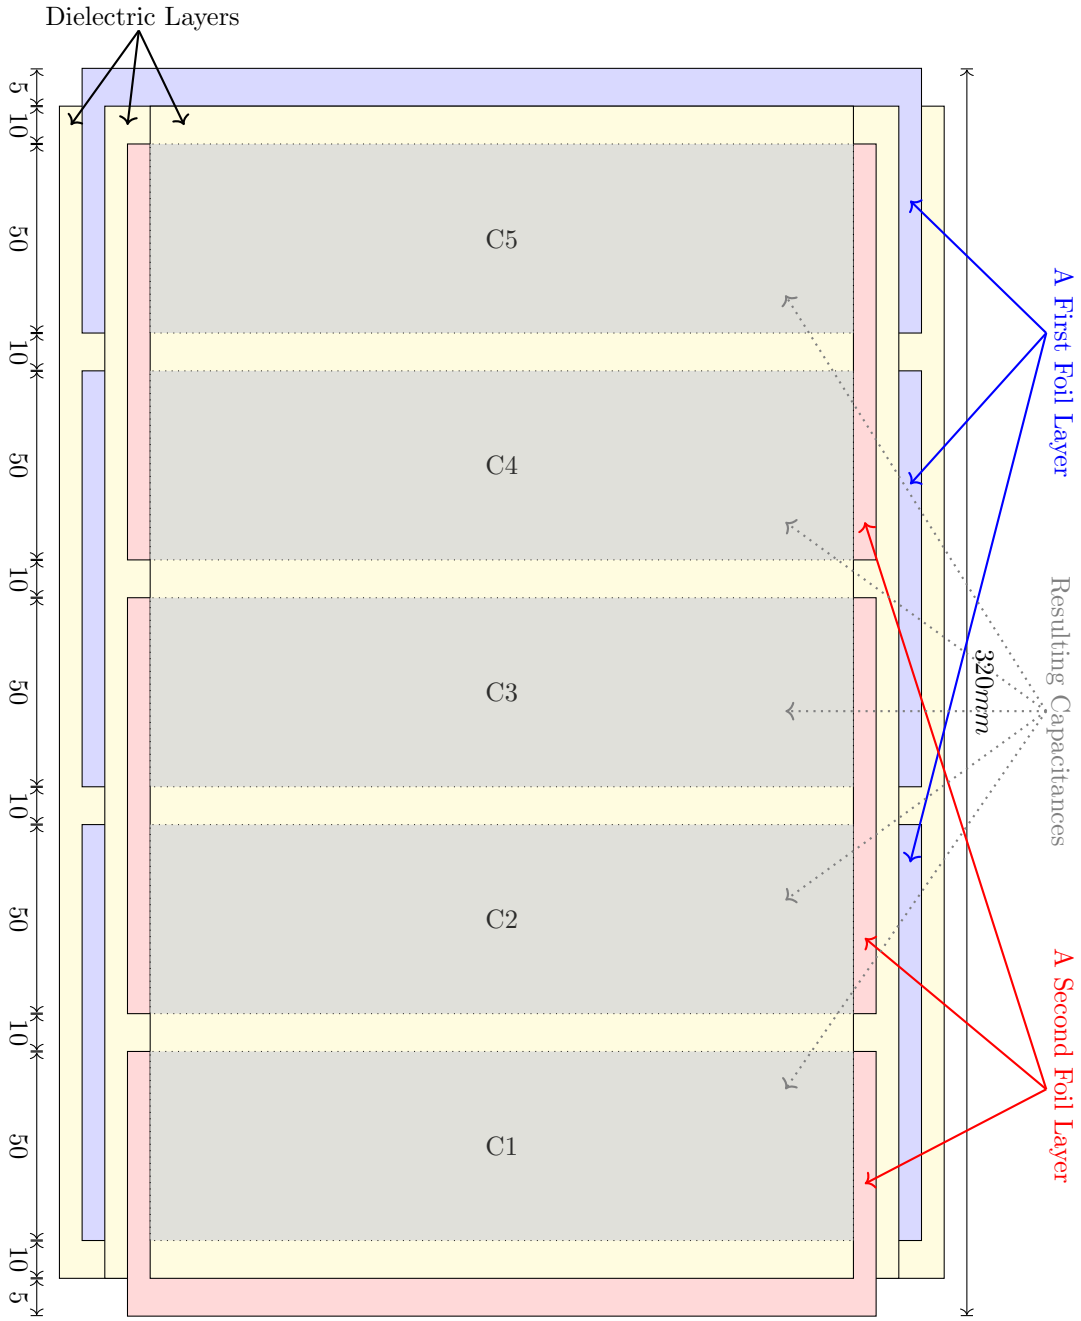

Figure 6.3: Design of interleaved windings

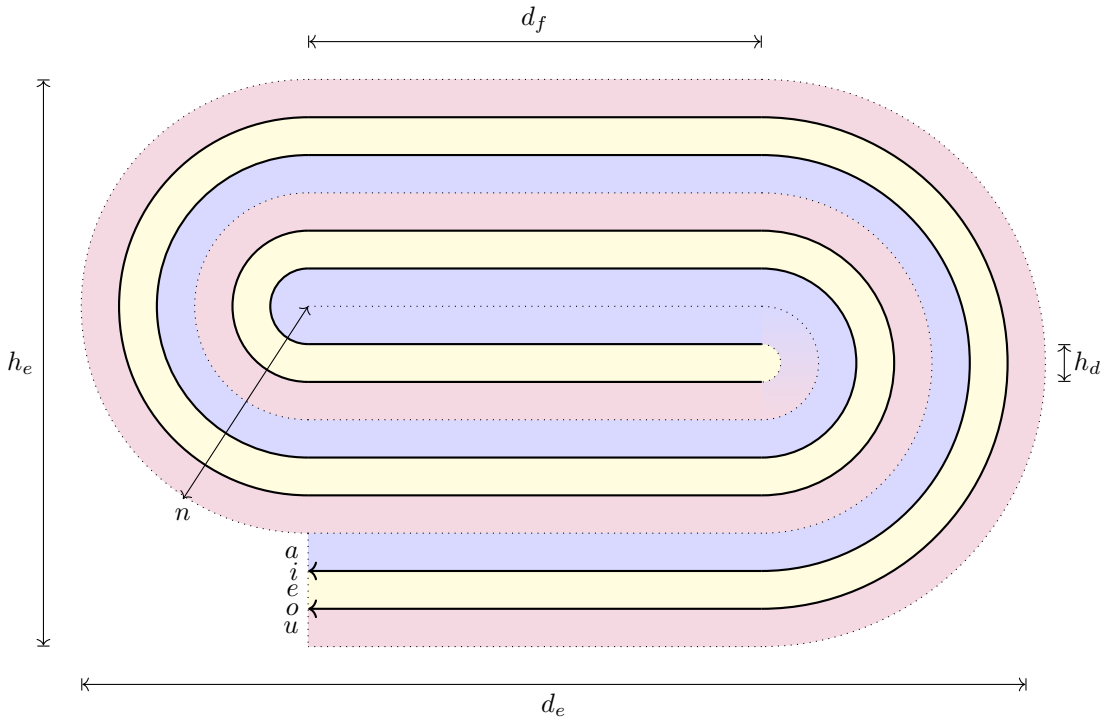

Figure 6.4: A concentric element with 3 dielectric layers

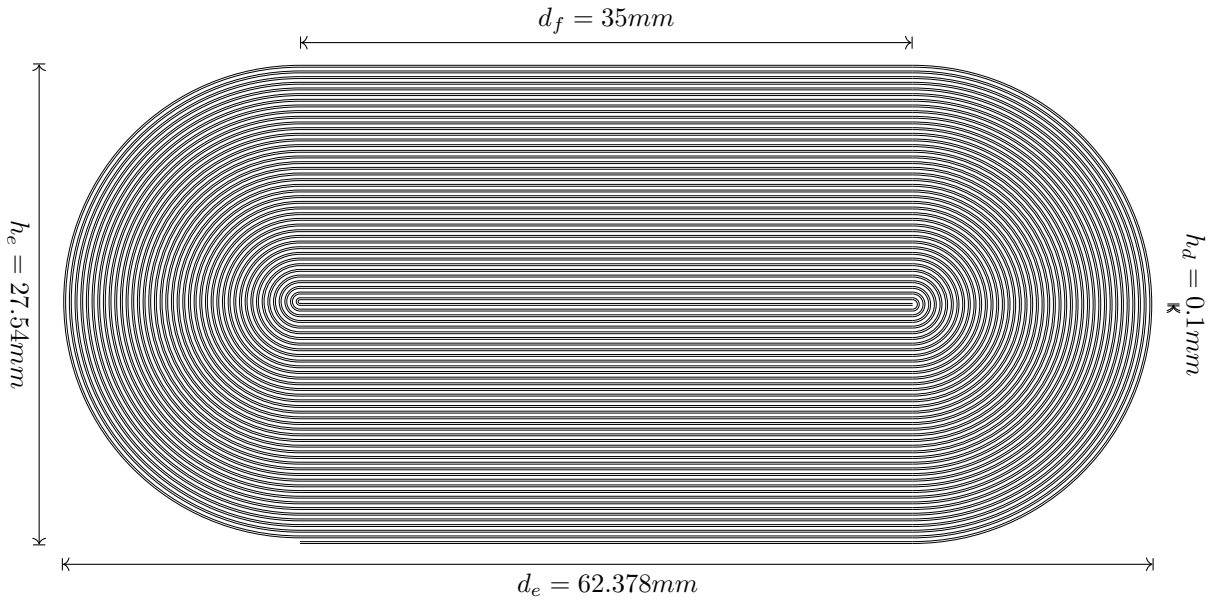

Figure 6.5: A second concentric element with 3 dielectric layers

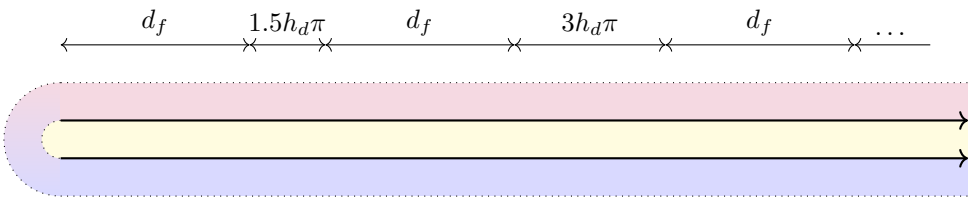

Figure 6.6: A flattened concentric element with 3 dielectric layers

Dielectric Layer Reduction with Counterroll Windings

'Counterroll' designs can offer higher capacitances for a given overall winding volume by reducing the number of dielectric layers necessary, but at the likely expense of greater manufacturing difficult due to a need to 1. arrange winding layers and prepare a winding precisely at a centre point and 2. accommodate greater winding distances (and an associated increased propensity for splay) due to winding in both directions. Although such a design still requires (usually folded) electrical terminal edges to be at potentially warm centre of a winding, in a counterroll design such edges of opposing terminals are at least separated from one another. Each terminal in this design category also boast an even distance from opposing terminals in either direction, which offers a more even distribution of electric field among the dielectric layers used.

A counterroll element with 2 dielectric layers is illustrated in Figures 6.7 and 6.8, and can be described by Equations 6.8 to 6.11, which measure a distance outward from a winding's centre as an alternative to evaluating the entire length of dielectric as in Equation 6.15. Foil lengths, l_i in Figure 6.7, is equivalent to Equation 6.14 but where h_d is substituted for $2h_d$.

$$h_e = (4h_d + 2h_f)(n+1) \tag{6.8}$$

$$d_e = d_f + (4n+2)h_d + 2nh_f \tag{6.9}$$

$$l_a = (n+1)d_f + \Big(\frac{n(2n+3)}{2}h_d + \frac{n(n+1)}{2}h_f\Big)\pi \tag{6.10}$$

$$l_u = (n+1)d_f + \Big(\frac{n(2n+1)}{2}h_d + \frac{n(n-1)}{2}h_f\Big)\pi \tag{6.11}$$

Similarly, a 1-layer winding of the same design is possible as shown in Figure 6.9, flattened in Figure 6.10, and is described by Equations 6.12 to 6.15. This further increases the benefits of using a counterroll design to improve the capacitance-to-volume ratio, but again at a cost of manufacturing ease and with increased risk of dielectric breakdown as a result of fewer dielectric layers combined with a risk of still having terminal edges at vulnerable winding centre points.

$$h_e = (2n+3)h_d + 2(n+1)h_f \tag{6.12}$$

$$d_e = d_f + 2(n+1)h_d + 2nh_f \tag{6.13}$$

$$l_i = l_o = (n+1)d_f + \Big(\frac{n(n+1)}{2}h_d + \frac{n(n-1)}{2}h_f\Big)\pi \tag{6.14}$$

$$l_a = (2n+3)d_f + \Big((n+1)^2h_d + n(n+1)h_f\Big)\pi \tag{6.15}$$

Continued Windings

Figure 6.11 shows a 'continued' winding design, which benefits from having no central edges. Figure 6.12 intimates a further benefit, where this approach allows a foil to be wound from a centre point of an array layers which is more accommodating of manufacturing tolerances, but with an electric field distribution less balanced than counterroll options. Continued designs are

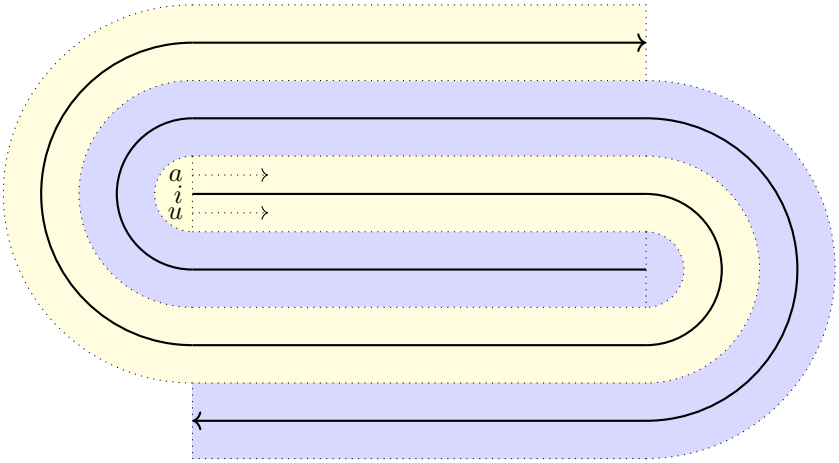

Figure 6.7: A counterroll element with 2 dielectric layers

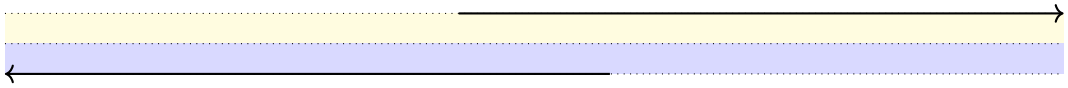

Figure 6.8: A flattened counterroll element with 2 dielectric layers

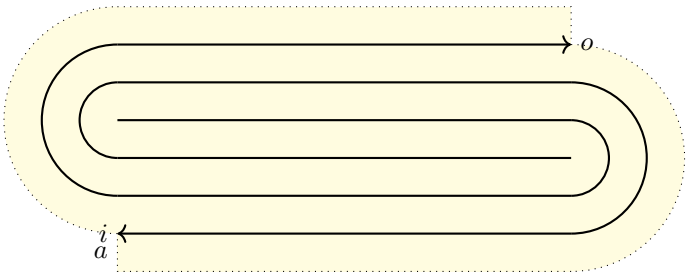

Figure 6.9: A counterroll element with 1 dielectric layer

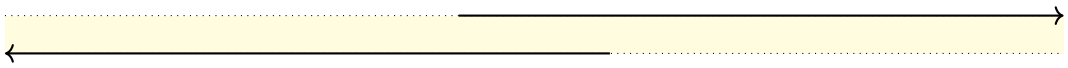

Figure 6.10: A flattened counterroll element with 1 dielectric layer

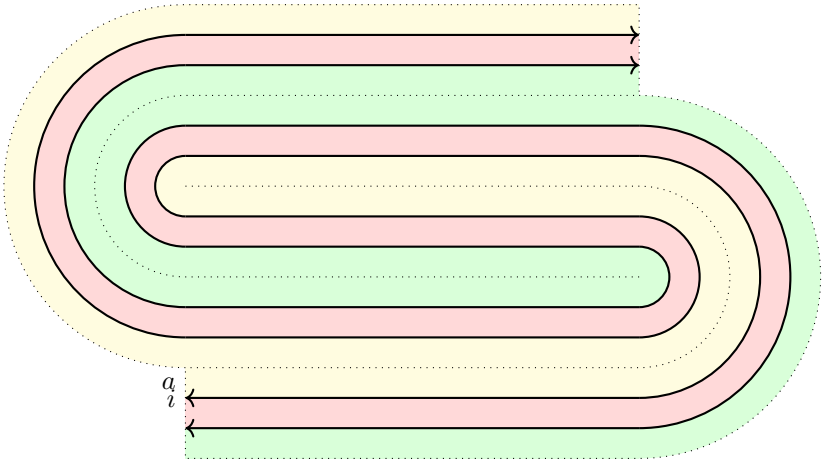

Figure 6.11: A continued element with 3 dielectric layers

described by Equations 6.16 to 6.19, where all dielectrics are naturally of equivalent length l_a.

$$h_e = (2n+1)(3h_d + 2h_f) \tag{6.16}$$

$$d_e = d_f + 2n(3h_d + 2h_f) \tag{6.17}$$

$$l_a = (2n+1)d_f + \big(3n^2 h_d + 2n^2 h_f\big)\pi \tag{6.18}$$

$$l_i = (2n+1)d_f + \big(3n^2 h_d + n(2n-1)h_f\big)\pi \tag{6.19}$$

Insulation Grading

An illustration is given of a winding element with insulation grading in Figure 6.13, which is expanded upon in Figure 6.14 where an unwound array of constituent layers is shown, and with 5 layers in Figures 6.15 and 6.16. These are of counterroll design although this is not exclusive.

Such designs are likely in practise to be difficult to manufacture, and while an increase in dielectric layers can add to overall winding volume, tactical use of this approach could offer improved reliability for certain elements, ameliorate the severity of thermal containment, and crucially could provide some resilience by reflecting a style of countermeasure used in other assets in the face of challenges intrinsic to the electrical environment as discussed in Section 1.3.3.

6.1.4 Winding Specification

Having explored a set of potential winding designs, all design options were attempted:

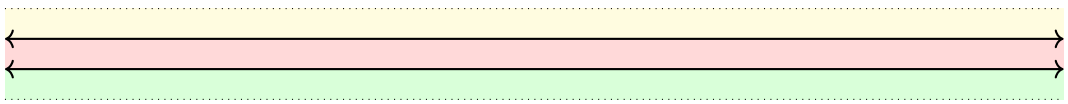

Figure 6.12: A flattened continued element with 3 dielectric layers

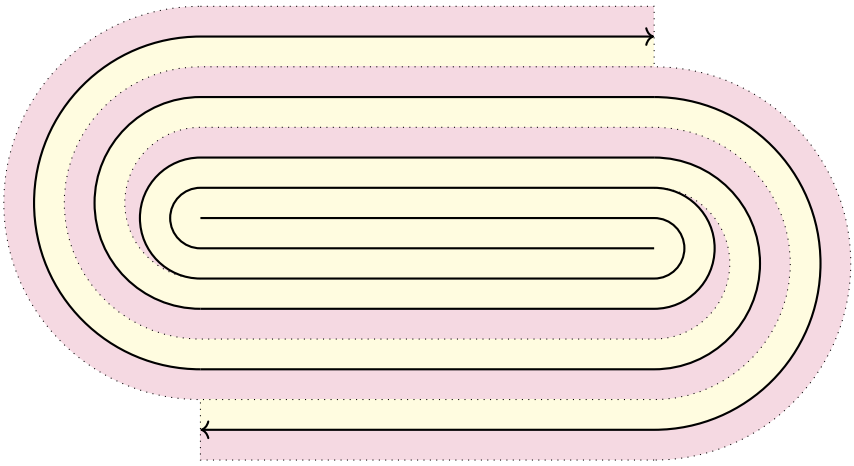

Figure 6.13: A graded counterroll element with 3 dielectric layers

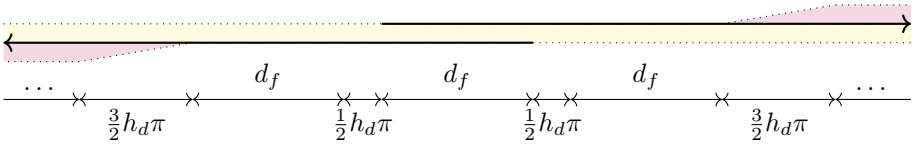

Figure 6.14: A flattened counterroll element with 3 graded dielectric layers

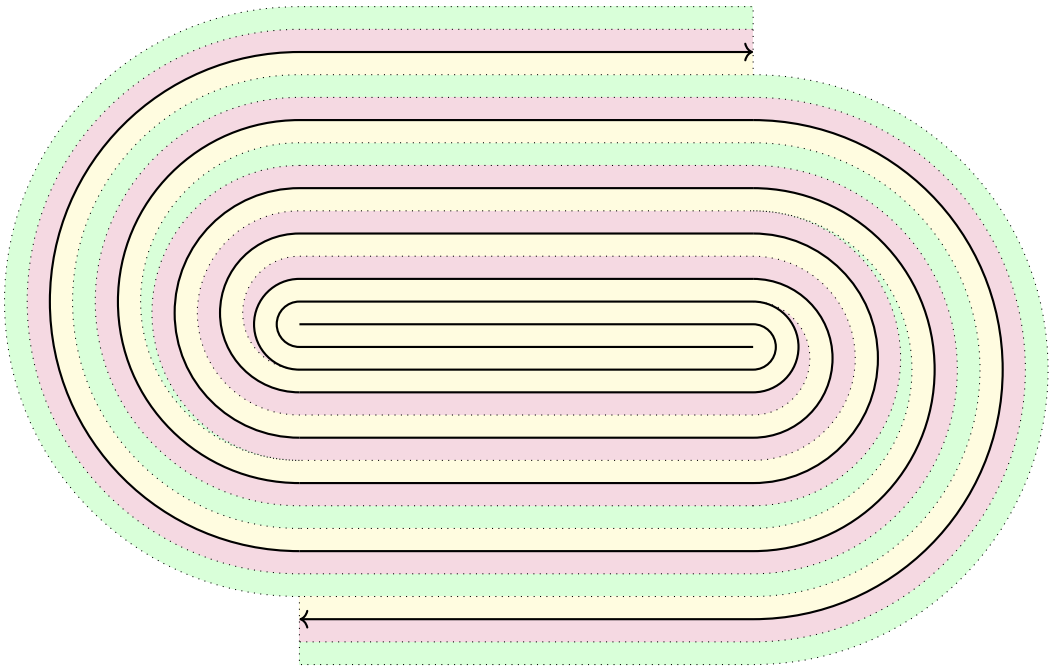

Figure 6.15: A graded counterroll element with 5 dielectric layers

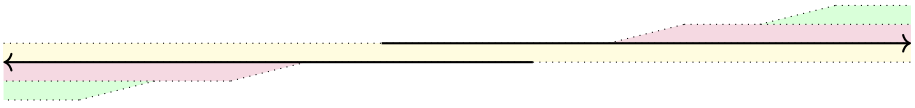

Figure 6.16: A flattened counterroll element with 5 graded dielectric layers

1. a 'counterroll' winding design was initially motivated by the space-efficiency a single-layer dielectric counterroll winding could offer, however, to prepare a winding layer arrangement (set out on adjoining tables), this required twice the preparation space of other design options due to a need to accommodate layers being extended in opposing directions;

2. 'concentric' windings were then attempted, but as Section 6.2 discusses, winding layers splay when rolled which prevented this approach from being practical; and consequently

3. a 'continued' design was instead chosen.

Since a *continued* design involves positioning a mandrel in the centre of the winding layer arrangement to begin the process of winding outward from a central point, this both reduces the overall preparation space required and limits the extent to which layers are affected by splay during preparation (by again reducing the distance to be wound in each direction). It additionally means there are no aluminium foil edges in the centre of a winding.

To specify a winding design for a 3-layer element to fit the overall capacitor unit dimensions in Section 6.1.2, the constraints (in mm) are: $h_e = 28, h_d = 0.1, h_f = 0.012, d_f = 37, d_e = 66, w_e = 330$. When reconfiguring applicable equations presented in this Section for a 3-layer discrete foil counterroll winding, the permissable number of turns which satisfy the constraints in each dimension are found to be $n_h = 42.7099$ and $n_d = 44.5031$, such that $\therefore n = 42$.

These formulae and calculations are theoretical and are difficult to translate into a winding design unless an already well refined process is used. These designs have been passed to Andy Carlin who constructed the frame and outer housing for a physical prototype.

6.2 Prototype Construction

With respect to principles of continuous improvement, where a thorough understanding and description of a process is crucial for consistency and prerequisite to improvements in quality, automation, and efficiency, the process adopted to prepare a prototype capacitor unit is described in this Section. Winding preparation consists of six stages: 1. set and clean a workspace; 2. layer acquisition; 3. paper cutting (Section 6.2.1); 4. foil cutting (Section 6.2.2); 5. setting out; 6. winding rolling (Section 6.2.3); and 7. measurement (Section 6.2.6).

6.2.1 Paper Roll Cutting

Having acquired rolls of parcel paper and 30 $cm \times 10$ m rolls of aluminium foil from local shops, paper was run through the draw saw to cut it to rolls of 310 mm width. This is a quick way to cut many rolls of paper to the required width. A similar attempt was made to size foil with this saw, but foil edges cut this way could not be cleanly unwrapped, and as such a more manual process has been required to cut foil layers to size.

While length of foils necessary for the winding dimensions were calculated to be less than 5 m, efficient use of aluminium foil rolls sold for domestic settings was impractical, as these rolls: were frequently uneven and crumpled toward the centre of the roll; had glue and uneven edges; and were almost universally less than the 10 m advertised.

6.2.2 Foil Cutting

Steps in a process of cutting foils to size are illustrated in Figure 6.17. Figure 6.17a shows a workspace established in the laboratory, where a set of four tables have been positioned and cleaned to create a workspace at least 5 m long, along which a single layer of foil can be unwound as in Figure 6.17b. With a 1 m rule, the edge of this foil layer is then marked at 1 m intervals from one end, for 5 m (Figure 6.17c), at which point additional marks are made at 11 cm, and 22 cm from one edge and 7 cm from the other (Figure 6.17d). These intervals are chosen so that existing edges are used on two of the foils to reduce the impact of inevitable errors where a blade catches on the foil. A ruler (or rule) is then used to continue these lateral width markings at periodic intervals (no more than the length of the shortest rule apart) from the 5 m mark back to the beginning of the winding (Figure 6.17e) so that these widths can be maintained each time a rule is repositioned along the foil when making a cut, before the rule and ruler are positioned together along these markings for a guillotine blade to be run between them to make each cut (Figure 6.17f). Since foil is fragile, both rule and ruler can be used in combination to distribute applied pressure on either side of the foil where a cut is made. This pressure is important when cuts are finally made, as it minimises slack in the foil, and thereby reduces likelihood of the foil becoming caught, which results in a ragged edge. Lastly, foil ends are cut in the same manner at the 5 m mark in Figure 6.17d, before resulting foil strips can be tidied of any ragged edges. Including setting out and measurement, it almost takes 2 hours in the lab to (carefully) cut the set of 6 foils needed for each interleaved winding.

6.2.3 Winding Rolling

Materials

Five materials are used to prepare each winding: • 2 × architect ruler • 3 × paper layer • 4 × 11 cm wide foil strip • 2 × 7 cm wide foil strip • cellotape (for completion).

Outline Process

As in Figure 6.18, 5 layers are individually positioned with appropriate spacing for foil margins (Figure 6.18a) along a 5 m workspace to form an 'array' (Figure 6.18b). A midpoint is found and logbooks and books positioned as paperweights[3] (Figure 6.18c) before two architect's rulers are placed on either side to clamp the winding layers to form a makeshift mandrel (Figure 6.18d).

The process of rolling these layers into a winding then occurs in stages. Books are slid along the array for approximately half a metre in either direction in an attempt to: • iron out any slack; • provide tension to layers being wound; and • bound each distance to be wound at any one time. Once books have been extended to new positions in either direction along the array, layers are wound (preferably with tension in each direction) until this process returns books to meet at the winding centre point (Figure 6.18e), at which point books can be positioned to replace manual pressure on the winding. This pause point is illustrated by Figure 6.18f and allows 1. a lengthy procedure to be subdivided, 2. an array to be repositioned on the workspace if necessary, and 3. winding layers to be realigned using corrective measures.

[3]There are many benefits to a thorough literature review.

(a) 5 m long workspace

(b) Unwound aluminium foil

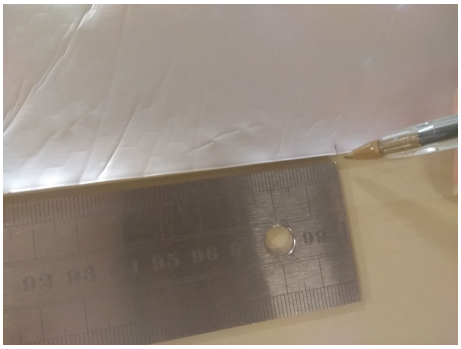

(c) 1 m interval marks on one edge

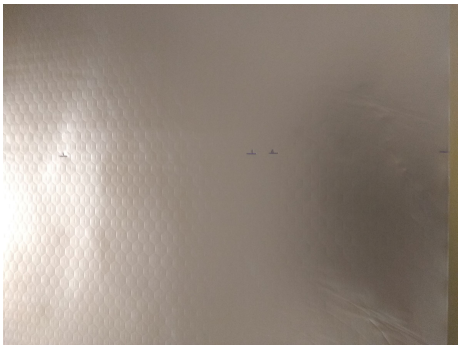

(d) Width marks at 5 m point

(e) Width marks continue along length

(f) Cut with pressure on each side

Figure 6.17: Foils being cut to size

(a) Setting out foil and paper layers

(b) An array set out along the workspace

(c) Books on either side of a centre point

(d) Architect rulers forming a mandrel

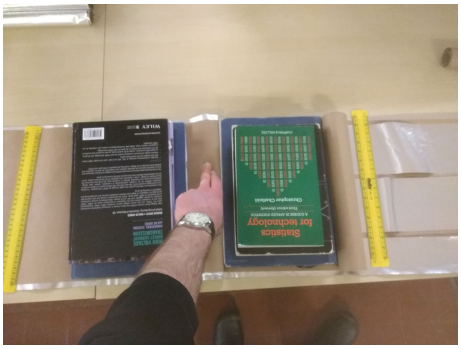

(e) Pause point during a winding process

(f) Using books to secure a pause point

Figure 6.18: Winding process 1

6.2.4 Corrective Measures

'Concentric' winding designs would be wound from one end around a perpendicular mandrel for the 5 m length of each layer. However, winding layers naturally splay (Figure 6.19) during this process, with a relationship which appears to be stronger than linear; the longer a winding is to be rolled, the greater the maximum splay when no interventions are employed. This is likely due to an uneven turn radius across the width of each layer at the point of being wound around the mandrel: one which is not exactly perpendicular; or is buckled; or applies uneven tension as it is turned, can result in inconsistent slack in each layer being wound. Even if the mandrel is perfectly configured, layers are subjected to different tensions simply due to their position: those which form inner turns will extend further outward from the winding than those forming outer turns merely as a result of a tighter bend radius.

Splay not only increases overall winding width but also reduces the effective capacitance, so a variety of static and dynamic measures can be used to minimise this. Static measures include winding outward from a centre point (rather than begin at one end), and preferrably to choose simple designs such that windings can be cut to size *after* being wound. To initiate a winding from a centre point, two rules are clamped together (albeit manually) to form a mandrel which can then be turned to initiate a winding from the centre: halving the effective length of an array susceptible to splaying in each direction.

Beyond static means of initiating each winding, it is appropriate to dynamically control the mandrel as it is rolled to correct for inevitable imperfections. A design with multiple foils per layer places added importance on this because of the set of foils in intermediate layers: it is desirable for these foils to maintain a near-constant spacing throughout the winding. A handful of mechanisms aim to combat splay and keep the layers of the winding array aligned.

Ironing slack out prevents gaps. This is done first manually but must be regularly repeated to accommodate the varying extents to which turn radius affects the amount of slack on each layer. At each pause point slack is manually ironed out by applying tension to each layer beyond the ruler in either direction, and books subsequently placed and pushed outward along the array, and preferrably not too far in either direction, to make doubly sure there is a minimum of slack between a winding and where books provide tension (Figure 6.20).

Paperweights balance tension in each direction. The closer each paperweight is to the mandrel, the greater the tension on layers which can be wound. Figure 6.21 attempts to illustrate that books acting to provide this tension can therefore be positioned to achieve as equitable a distribution in either direction as practical, and, time permitting, at limited distance from the mandrel at any one time in an attempt to increase the tension which can be applied. This is not necessarily an exactly equal horizontal distance, as layers in one direction begin at a height relative to the other, and (at least in this case) paperweights can be of unequal weight.

Splay is affected by lateral position. The lateral position of the mandrel perpendicular to the direction of the winding can be used to mitigate splay, as indicated in Figure 6.22. Initial turns of a winding are crucial: if early interventions reduce divergence on these turns, benefits are seen for the remainder of the winding.

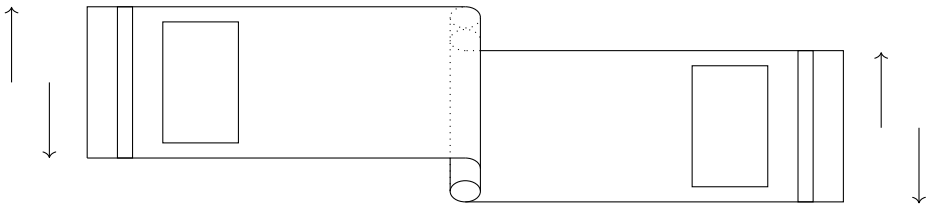

Figure 6.19: Natural splay

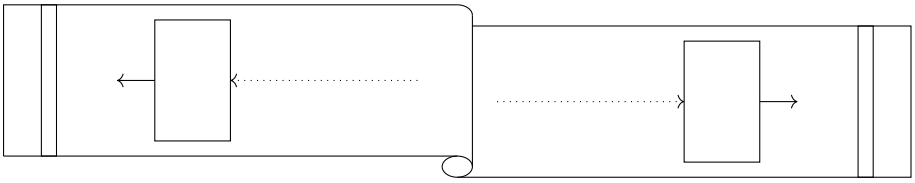

Figure 6.20: Iron out slack

Former angle influences splay. Figure 6.23 attempts to highlight the influence of mandrel angle on subsequent splay, which can in turn cause divergence between layers. If the angle of the mandrel relative to the array is used to counteract the natural splay, then a winding alignment can be maintained within tolerances. This angle is also relative to the point on its rotation, which presents a challenge when using this approach during winding since it can be difficult to use this technique manually as it necessitates: familiarity with the angle made between the mandrel axis and incoming winding layers; matching the frequency of axial rotation with the rotation of the mandrel itself off of its axis as layers are wound.

6.2.5 Practical Correction

One alternative to corrective 'mechanisms' used during the winding process is to adopt a more manual, measured approach during process pause points. With weight applied to hold the winding in place, books can be removed from one side to allow layers to be repositioned as in Figure 6.24a. Figure 6.24b illustrates that each layer of the winding array is then taken in turn, stretched to remove any slack, and repositioned relative to the other layers (through measurement where necessary). Once layers are tidily realigned, weights are replaced (Figure 6.24c) to reinstate pressure on realigned layers, and to restore weight on the incomplete winding (Figure 6.24d). Layer realignment is then repeated for the remaining direction (Figure 6.24e), until both sides have been adjusted (Figure 6.24f), before winding continues by reapplying manual pressure to the winding (Figure 6.25a), extending books along the array, and continuing the process until the next pause point (Figure 6.25b). The general process described in this Section 6.2.3 is repeated for remaining windings (Figure 6.25c) until all are complete. Figure 6.25d shows the six windings prepared in total for this project.

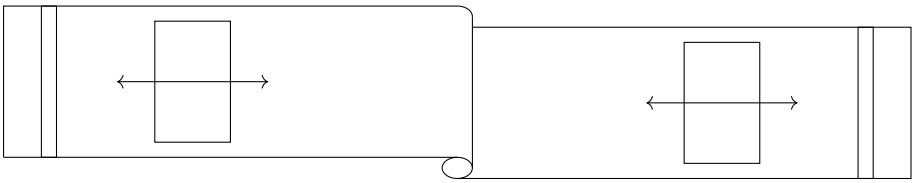

Figure 6.21: Paperweight position

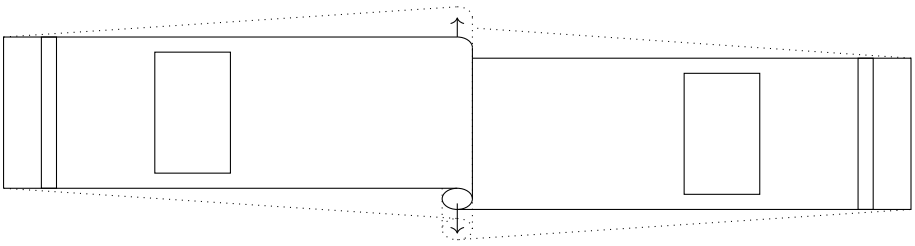

Figure 6.22: Mandrel position

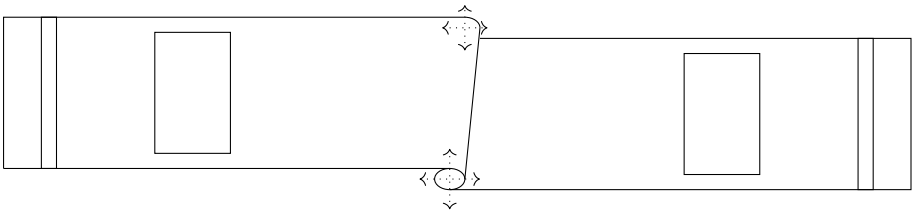

Figure 6.23: Turn angle

6.2.6 Final Windings

Figure 6.25d shows 6 final windings from which 4 have been selected to fit the housing in Figure 6.26. Fortunately, extra windings have been prepared to compensate for discrepancies, as one was measured at substantially reduced capacitance. At present this is believed to be due to a short circuit between adjacent foils within either of the two foil layers.

Corrective actions taken during the winding process have been intended to keep a winding to within its design tolerances, which in principle had originally been intended to support resultant capacitances by ensuring sufficient overlap between terminals. However, those windings which (subjectively) experienced most correction also have the most attenuated capacitances.

The first winding has highest capacitance, which corresponds with a long, almost contiguous process used in its preparation and during which corrective interventions were largely neglected and layer tension deemed paramount. One drawback, however, is that this winding extends further than its counterparts at either end. Corrective interventions could be effective at maintaining alignment if the winding is thicker in its centre, which in turn possibly results from plastic architect's rulers bending outward during the first few turns. Slack between layers would therefore be necessary to keep the winding aligned. As distance between foils is paramount to capacitance, corrections therefore reduce the capacitance. Such an effect could be further aggravated since each foil sees its counterpart terminal in both directions (inward and outward relative to the winding on both prior and subsequent turns), so increased distances between opposing terminals lead to a significant reduction in resulting measured capacitance.

Each corrective action introduces slack at winding edges which brings a high probability of inadvertently increasing distance between terminals with a noticeable depreciation in capacitance. Maintaining layer tension while winding is critical to high element capacitances and a minimum of correction is therefore preferable.

6.2.7 Internal Connections

While a total of 6 windings have been prepared, only 4 can fit the designed housing dimensions, so four windings with the most similar measured capacitances are selected to be connected as a

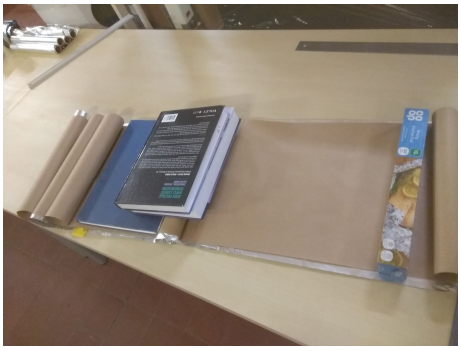

(a) Clear books from the right side

(b) Realign each layer in turn

(c) Return books to a realigned right side

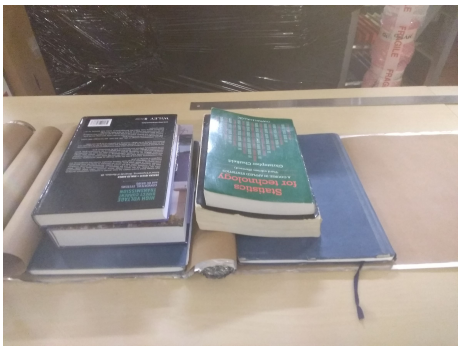

(d) Secure the pause point from the right

(e) Adjust each layer on the left

(f) An adjusted leftmost array

Figure 6.24: Winding process 2

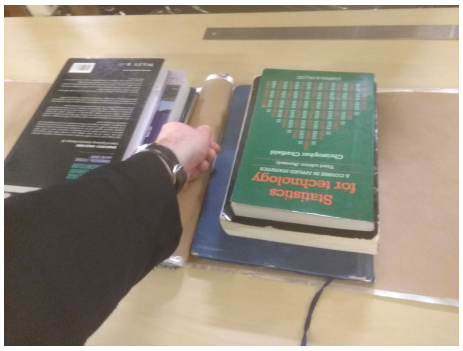

(a) Preparing to continue winding

(b) Another pause point

(c) Process repeated on winding 6

(d) All six windings

Figure 6.25: Winding process 3

Figure 6.26: Windings in housing

Table 6.1: Winding measurements

Winding	1	2	3	4	5	6
Capacitance (nF)	18.27	6.61	10.04	10.48	10.41	10.78
Connected outwith Can	discarded		5.35 nF		5.37 nF	
Connected within Can	discarded		5.24 nF		5.37 nF	
Connected in full Can	discarded		5.80 nF		5.78 nF	
All Connected in Series	discarded		2.93 nF			

series group within the unit. These have similar capacitance values but vary in physical size due to differences (varying tensions) in the winding rolling processes used between them. Although winding 1 offers greatest capacitance, omitting it from the final unit due to its anomalous capacitance also avoids a problem its increased height would pose relative to the unit housing.

After windings are individually measured, strips of aluminium foil are folded to form tab connectors (Figure 6.27a). Having cut these together at either end, the foil layers remained connected along the line of cut. As the aluminium foil has been cut cleanly, and is does not quickly oxidise, this is likely due to a form of cold-welding which explains the difficulty in separating foils cut together, whether in Section 6.2.2 or in Figure 6.27a. In future, a coating or material could be applied along a line of cut to allow multiple layers to be directly cut on any similar projects in future as a possibility for more efficient use of time.

Selected windings are then paired for connectors to be positioned at the face farthest from the completed unit's bushings. Table 6.1 gives measurements taken (in respective rows) of: 1. each winding individually; then 2. following their connection into pairs; 3. as each pair is positioned within an otherwise empty unit housing (with the lid lifted too); 4. with each pair individually measured from within a full unit; and 5. where series-connected pairs are in turn series connected to complete a unit which is measured from one terminal to the other.

This high-voltage design uses elements in series to withstand higher voltages, which leads to a reduction in overall capacitance in accordance with the formula $\frac{1}{C_{Series}} = \frac{1}{C_1} + \frac{1}{C_2}$ (by analogy, the effective distance across a dielectric is increased between terminals) which results in the reduction in measured capacitance as windings are connected in sequential rows of Table 6.1.

Measured on an open table-top, capacitance was found to increase when light manual pressure is applied: pressure reduces the distance between terminals in a loose winding, thereby increasing measured capacitance; within the unit complete with both winding pairs pressure exists due to geometrical constraints, which is believed to be the reason for a relative increase in capacitance measured in a full can compared to one containing only one winding pair.

On the other hand, a less apparent effect was also noticed where a conductive object (a human hand) was moved near to the winding without applying pressure, and which at certain locations very slightly reduced measured capacitance on an otherwise open table-top: fringing electric field effects at margins between terminals could explain this subtle effect if a nearby conductive object bridges a margin and effectively introduces an additional (if largely ineffective) series element; should unit housing or adjacent windings cause similar interference, these two effects could to some extent counteract one another and explain the negligible difference between the 5.37 nF pair when measured outwith and within unit housing.

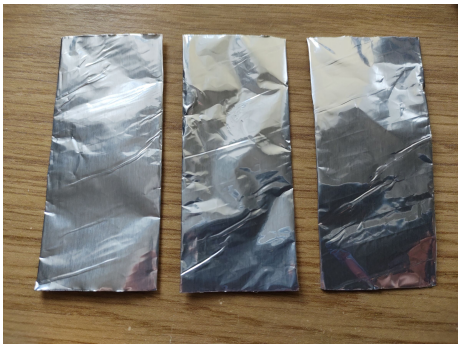

(a) Tab connectors

(b) Far tab connector

(c) Opened-top unit

(d) Completed unit

Figure 6.27: Integrating the unit windings

Figure 6.28: RCL meter

Table 6.2: Multimeter measurements

H A BC D (nF)	Healthy	A Fault	B Fault	C Fault	C2 Fault	D Fault
H A	0.16	0.16	0.16	0.16	0.15	0.16
H BC	0.16	0.16	0.16	0.16	0.15	0.16
H D	0.16	0.16	0.16	0.16	0.15	0.16
A BC	5.98	6.61	5.98	5.98	5.98	5.98
A D	3.03	3.17	3.02	3.02	3.17	3.18
BC D	5.93	5.93	5.93	5.93	6.56	6.60

Table 6.3: RCL meter measurements

H A BC D (nF)	Healthy	A Fault	B Fault	C Fault	D Fault
H A	0.144	0.144	0.144	0.1437	0.1446
H BC	0.145	0.145	0.145	0.1455	0.1461
H D	0.144	0.144	0.145	0.1455	0.1448
A BC	5.88	6.500	6.509	5.881	5.881
A D	2.94	3.086	3.084	3.088	3.101
BC D	5.83	5.832	5.833	6.463	6.506

6.3 Measurement

For the purposes of recording measurements, the windings (according to the order as shown in Figure 6.27c) are labelled A, B, C, and D: where B and C are connected via a tab connection these are also combined to form terminal 'BC' for the purposes of taking measurements; unit housing is also designated as terminal 'H'. Hence, the completed practical prototype and simulated model each have 4 terminals and $\binom{4}{2} = 6$ measurements to be made.

On initial measurements of capacitance between foils and the housing, a difference was apparent between the measurements made relative to different sides of the unit housing. The unit uses a separate panel for each face of the unit housing, which don't naturally connect to form an electrically continuous surface as in the digital model. Solder was used selectively between adjacent panels to establish electrical connectivity between all faces of the unit housing.

To assess the change in capacitance seen as a consequence of faults, short-circuits were introduced to each winding in turn by introducing a strip of aluminium foil to connect the exposed foil with the opposing layer positioned between intermediate layers of paper. In some windings (B and C) there are more than two layers of paper between turns of the same exposed terminal, such that there are multiple possible locations that the opposing foil could occupy: it was on occasion necessary to reposition a connecting strip of aluminium foil in order to reach the opposing terminal and effect a change in measured results. C Fault and C2 Fault in Table 6.2 are both taken with a short circuit introduced to different intermediate layers.

Measurements were first made using a multimeter and are presented in Table 6.2 before a more accurate RCL meter (Figure 6.28) was used to repeat the measurements to generate the results in Table 6.3. Results are in nF for both of these Tables. The final measurement in Table 6.1 of 2.93 nF corresponds with the value for the 'Healthy' case measurements between terminals A and D in Table 6.2, but for the lid and bushings having been present in the former.

The prototype unit was then subsequently measured as 3.102 nF using the PM 6303 RCL Meter in the lab, via the bushings as in Figure 6.28, before connections to and between windings were again measured under healthy and faulted scenarios to provide the results in Table 6.3.

Table 6.4: Model (with 15 mm offset) relative permittivity (ϵ_r) sweep

Permittivity	C (nF)
1	0.2903 nF
4	1.1000 nF
6	1.6380 nF
8	2.1756 nF
9	2.4443 nF
10	2.7131 nF
10.5	2.8474 nF
10.75	2.9146 nF
10.80	2.9281 nF
11	2.9818 nF
12	3.2505 nF

Table 6.5: Simulation measurements at $\epsilon_r = 1.4$

H A BC D (nF)	Healthy	A Fault	B Fault	C Fault	D Fault
H A	0.061860	0.062396	0.061966	0.061954	0.061861
H BC	0.064084	0.064085	0.064221	0.064229	0.064084
H D	0.061862	0.061863	0.061950	0.061983	0.062375
A BC	0.74263	0.83772	0.81627	0.74317	0.74263
A D	0.38283	0.40658	0.39818	0.39989	0.40466
BC D	0.74154	0.74154	0.74208	0.82288	0.82798

Table 6.6: Simulation measurements at $\epsilon_r = 10.8$

H A BC D (nF)	Healthy	A Fault	B Fault	C Fault	D Fault
H A	0.075169	0.075271	0.075189	0.075187	0.075169
H BC	0.075587	0.075587	0.075611	0.075613	0.075587
H D	0.075169	0.075170	0.075187	0.075193	0.075267
A BC	5.6756	6.3998	6.2451	5.6789	5.6756
A D	2.9091	3.0883	3.0291	3.0423	3.0734
BC D	5.6673	5.6673	5.6706	6.2964	6.3246

6.3.1 Corresponding Simulation

To assess whether the principles shown through simulation are demonstrable in reality, and to help identify discrepancies, a simulation model has been prepared to match the constructed unit and is modified: 1. to match the size of the physical prototype by (a) resizing housing parameters, (b) adding a 6 mm lid, (c) configuring windings with similar turns and orientation, but (d) keeping bushings and rack mounts as in the original model; 2. materials were changed (a) foils are set as `Aluminium`, (b) the housing as `Al 6063`, (c) dielectric removed (as the prototype is not oil-filled) from all but layers within each winding, leaving the immediate surroundings of windings as air filled; 3. offsets (the width of the PVC frame) was then increased to 15 mm in an attempt to reduce paper layer widths; 4. different permittivities for paper insulation were parsed; and 5. faults introduced similar to tests on the prototype. Table 6.4 summarises results for a variety of paper insulation relative permittivities in this arrangement.

Results

On the model with a 10 mm offset, an overall capacitance of 463 pF was found with simulation of 1.1 relative paper permittivity, and 780.79 pF with a permittivity of 2. Both of these values are far from the 2.930 nF (Table 6.1) measured on the prototype unit. One apparent reason for this large discrepancy is visually apparent in the thickness of dielectric layers: while the prototype windings have some slack and while there will be gaps between foil layers, these are unlikely to be as uniformly distant as are those in the simulation model. The simulation model was readjusted to use a 15 mm offset in an attempt to more closely replicate the narrow winding layers of the physical prototype, and different values of terminal capacitance were evaluated for varying dielectric relative permittivities, as provided in Table 6.4.

As Kraft paper is known to have a relative permittivity of 1.4 [90], this was selected for more detailed evaluation of fault scenarios alongside a purely synthetic model which uses an unrealistically high permittivity of 10.8 as an approximation of the true, measured capacitances. These models are respectively used to generate the results provided in Table 6.5 and Table 6.6.

6.3.2 Discussion

In Table 6.2, the minimal change in the numbers in the B Fault and C Fault columns relative to the healthy case suggests either that the introduced fault was ineffective or that a short-circuit already existed in the winding in each case. Windings B and C each have three paper layers between turns of the same foil terminal, so of the two possible foil layers between these layers of paper one will be a correct choice for the purposes of introducing a foil tab designed to emulate a short-circuit, and hence it is possible that the wrong intermediate layer was chosen for measurements in column B Fault of Table 6.2. That a change is seen in Table 6.3 for column B Fault with respect to the Healthy case implies that for this measurement the correct intermediate layer was chosen and a short circuit more effectively introduced, and that a second attempt (column C2 Fault) in the same Table shows results more in keeping with the B Fault column adds further evidence that this error also arose in the test setup for winding C.

While discrepancies exist between results in Table 6.3 and Table 6.6, these illustrate a similar behaviour: A 10% increase in the presence of a fault in Table 6.3 approximately correlates to the 12.7% increase seen in Table 6.6; and these increases also correspond to the location of the

fault relative to the position of the measurements. Moreover, and while errors are possible, it appears that the greatest increase in measured capacitance occurs where faults occur closest to the points of measurement, which reinforces findings in Chapter 4.

Reasons for a discrepancy between the measurements made on the prototype unit and those acquired by simulation are likely to include: 1. omission of a PVC frame from a simplistic simulation model; 2. spacing between terminal layers in the simulation model, relative to the prototype windings which are tightly contained by the unit frame; and 3. variation in windings such as (a) the number of layers of dielectric between terminals, as multiple layers of paper separate layers of foil in some windings, and (b) the model uses a 'counterroll' winding design too impractical to build due to splay; not to mention 4. the possibility[4] of short-circuit faults internal to physical windings, as each would increase the overall capacitance by a step change.

6.4 Continuation

For future work in preparing a similar prototype it would be preferrable to avoid interleaved winding designs which, although able to offer some technical advantages in a completed unit, complicate development, increase preparation time, and are particularly vulnerable to defects such as splay (which can lead to internal short circuits) and ragged edges after being cut. Sacrifices in technical capabilities of an end product can be considered an acceptable trade-off for timely and more amenable development when the aspiration is a mere prototype.

Winding tension takes precedence over alignment, although a balance is necessary in practise. If interleaved designs are avoided, in principle there could be some scope to tolerate splay by using deliberately oversized layers in the knowledge that a resulting winding can be cut cleanly[5] to size at either end once it is wound. However, this approach might preclude the use of tab-connections or schooping to form connections at winding ends; instead, areas of foil terminals could be exposed along the edge of the winding which would furthermore allow windings to connect directly in series when stacked.

Rather than attempt to verify simulated results by constructing a physical replica to match an emulated model, a more practical approach is to first prepare a suitable prototype and instead retrospectively resize a computer model to match. This methodology underscores a need for including configurability within developed model files, as discussed in Chapter 2.

Should circumstances permit, a completed prototype capacitor unit could be used to: 1. heat the unit to different temperatures and measure changes in capacitance each time; 2. use point heat sources within the unit to validate thermal behaviours studied in Chapter 5; and 3. evaluate changes in parasitic capacitance by measuring overall capacitance before and after implementing short-circuit faults between terminals.

6.5 Conclusions

This Chapter has discussed stages of preparatory design and practical implementation toward a scale model of a high voltage capacitor unit, intended as a means toward verification of

[4] Or likelihood, given terminals' proximity over a 10 *mm* margin and propensity for splay while winding.

[5] Where it is necessary first to mitigate against the possibility of cold welding when cutting multiple layers as discussed in Section 6.2.1.

simulated behaviours within a laboratory setting. Indicative conclusions can be drawn from these efforts and are outlined as follows.

1. In preparing a prototype it can be considered acceptable to pursue a simplified (if less technically interesting) asset design for the purposes of: (a) easier and more time-efficient development; and (b) learning to resolve simple problems before being presented by more challenging ones.

2. Thus, interleaved winding designs are unnecessarily impractical where a greater number of windings of simpler design could offer a more straightforward alternative, and which: (a) could be prepared with oversized layers in anticipation of splay, such that they can be cut to fit lateral dimensions after winding; (b) can be connected via areas of exposed terminal on outer layers which connect in series when windings are stacked, rather than connect at either end; and (c) where more than the necessary number are prepared to accommodate errors and achieve a more comparable capacitances.

3. Winding designs are easier to prepare in practise as 'continued' windings since: (a) a centre point used to initiate a winding does not have to be precise; (b) the distance of winding to be rolled in each direction is halved (with an associated reduction in the amount of natural splay induced); and (c) avoid the risk of having terminal edges at a vulnerable centre of a winding as is the case in 'concentric' or in potentially more space-efficient 'counterroll' designs.

4. It is beneficial to follow a clear process, which: (a) can adapt in light of new information, understanding, and from experience; (b) is documented so that it can be followed with relative consistency and is repeatable by being unambiguous to others; (c) includes break points to accommodate human factors and practical interventions in pursuit of a good-quality outcome.

5. This process should use correction mechanisms to counteract inevitable deviations from a design (a) at cost to the end result (in this case through attenuated winding capacitances), (b) in as limited a way as possible so as to prioritise winding tension over alignment to maintain desired capacitance, and (c) as early as possible when any slight deviation is first noticed, since early intervention minimises the overall corrective action necessary.

6. A sensible approach to physical prototype construction is to: (a) begin with simulation modelling to first understand engineering trade-offs in order to minimise the number of prototypes and waste involved [52]; (b) build from the centre outward, since it is more appropriate and significantly easier to size and construct simple components (unit housing) around more complicated internal windings; particularly since (c) it is better to update simulation models to match a constructed equivalent prototype than to attempt to physically replicate a simulation; which in turn (d) underscores the importance of simulation models which are adaptable to this end.

Chapter 7

Conclusions

This Chapter aims to offer brief discussion, collate conclusions from Chapters 4, 5, and 6, to proffer suggestions for further work, and to summarise some main contributions of this research.

7.1 Discussion

While Chapters 4 and 5 cover complementary subject matters, they have a common imperative in developing a basis for monitoring, detecting, and potentially locating faults within capacitors. An overarching perogative is to maximise availabilities, and by extension, the availability of HVDC converter stations and other assets within electrical proximity to them.

A Cultural Emphasis on Reliability

As outlined in Chapter 1, there is a pressing need for dependable, cost-effective, and increasingly sustainable energy systems, not least in pursuit of the United Nations' sustainable development goal 7. Interconnection, and stewardship of the assets which provide it, is an integral component of efforts to meeting international aspirations.

A model-based approach has been adopted for this project as discussed in Section 2.1.6, at least in part to investigate the physical behaviour of specific assets in detail without a need to await failure data, which: 1. take time to acquire, potentially over an asset's lifespan; 2. require a sufficient asset volume for reliability data to be of statistical value; and which 3. appear to be treated as commercially sensitive by a competitive industry of limited size. Despite the second point, anecdotal evidence is nevertheless a potentially rich source of insight into asset behaviour, operating conditions, and failure modes. In many industries, fault modes are afforded forensic root cause analysis, in part motivated by a desire to understand the specific circumstances around a particular failure, but more so as to discover wider conditions in which a fault can arise with a view to early prevention of further incidents. There is a cultural element to this, as it takes effort to explore in detail even slight defects, and within an organisational context, to be assured of the support for and time required of an investigation. Risk of 'blame' in a working culture must be put aside to create conditions for collaborative engineering dialogue sufficiently free of human bias or commercial interests, notwithstanding a difficulty in delineating responsibilities for electrical assets which tend to be influenced by a wider network, weather, and broad regulatory and commercial environments.

Reliability Collaboration

Aspirations toward open data principles are being increasing recognised within the power system industry, albeit primarily in terms of system-wide information pertaining to day-to-day activities rather than in relation to specific assets. Reliability information is arguably one area where openness matters: operators, manufacturers, maintenance teams, and asset engineers each have a unique view of a system which is in a mutual (and wider societal) interest to protect and maintain in line with principles of reliability centred maintenance. Collaboration toward this shared goal is key.

Some fault conditions and data are shared [17], but not all faults lead to the public availability of fault investigation outcomes. An ability to share understanding on the conditions and circumstances surrounding asset failure (in an aggregated, anonymised manner wherever appropriate) could allow all interested parties to learn from relevant findings, and potentially avoid placing similar assets at similar risk, to the benefit of overall availability of critical infrastructure. Ultimately, as well as financial justifications to an asset owner or industry, large infrastructure projects are further justified in pursuit of environmental, societal, and wider economic ambitions, and in the case of asset health, safety implications take natural priority. Reliability data are possibly too critical to form merely one basis of competition among many.

An Integrative Approach to Asset Health

A developing trend in modern medicine is toward integrative understanding[1]: rather than see a disease in isolation as a product of fortune, to persistently examine root causes even where these occur between interdependent systems – where social, psychological, and other environmental influences are considered. While reliability centred maintenance offers an appropriate framework for improving asset performance throughout operational lifecycles, and as will be further discussed in Section 7.4.4, there could yet be benefits to a translation of emerging thinking from a more conventional meaning of 'healthcare' to its analogy to key infrastructure within an engineering context.

Recent Examples of Interconnector Faults

Over the course of this project, the Western Link has undergone commissioning, and further HVDC links are proposed [19] for the east coast of the UK. Outages have occurred on the IFA1 interconnector[2] for which one bipole failed [18] as the other was out of service, and on the Western Link, which delayed the project [91] and constrained generation [92]. Besides the overall capacity, the risk of a *sudden* loss of a large interconnector capacity can give rise to knock-on impacts to the AC networks it connects to. Such events respectively serve to reinforce the importance of having information (and preferrably some means of early warning) about the health of high-capacity assets in operational contexts, and of the value that ambitious projects [21] can offer for future asset designs, should fault information be shared after-the-fact: this was the case for the scenarios discussed in Section 1.3.3 which lead to now commonplace design improvements (by applying reliability centred maintenance) for assets of this nature.

[1] This is characterised by work in journals such as *Psychosomatic Medicine*, the *Journal of Psychosomatic Research*, and in growing disciplines such as 'psychoneuroimmunology'.

[2] Interconnexion France-Angleterre (IFA)1 runs between Sellindge in England and Les Mandarins in France.

7.2 Collated Conclusions

The modular nature of capacitor assets presents an opportunity for monitoring. Fault detection is made easier, but a unit-by-unit basis for detection and monitoring has, as yet, neither been achieved nor proposed in current publicly available literature. Literature typically assumes a *linear* redistribution of electrical quantities following a fault, which are known to lead to *cascading* failure modes if left unchecked.

A capacitor unit model has been prepared and made available as a dataset online. This model has been designed to be a configurable starting point for future studies relating to physical phenomena and dielectric breakdown in capacitor assets. It is shared under a CC-BY-SA 4.0 license to encourage open-innovation approaches given the public nature of academic funding and the public interest in avoiding high-voltage outages.

Modelling can be improved with good-practise design, the earlier used the better. Familiarity with a simulation environment and clarity on the envisaged model can significantly improve early stages of development. Configurability can be achieved through parameters, making use of conditional statements, and applying `Definitions`, and which also improve model readability as a windfall gain. Similarly, support is there to be used: using simple models to ask brief questions can lead to a better understanding of the environment capabilities. Model methods offer a powerful tool with which to automate repetitive or tedious processes, allowing sophisticated models to be developed with relative ease.

Possible sequence control flow improvements have been suggested to COMSOL. Support for conditionality in a `Geometry` sequence could be complemented by support for iteration, which, should this also extend to `Definitions` (and therefore be additionally applicable to other aspects — a `Mesh`, `Materials`, and `Physics` interfaces), highly sophisticated and adaptable models could be easily developed in future.

Housings provoke stray capacitances which affect *nonlinear* voltage distributions. Element failures therefore affect the resulting distributions as a function of their location, and causing subtle variation in changes in terminal characteristics in each case. Higher voltage elements are subjected to greatest electrical stresses and so (all else being equal) are most likely to harbour faults. Indeed, higher-voltage element faults are preferred since they aggravate nonlinear distributions (and stresses) less than do those at lower-voltages.

Electric field stresses arise throughout a unit and depend strongly on its design. Stresses are particularly strong: around bushing connections; between windings at different voltages; at foil edges, corners, and on the outer edge of narrow foil turns; at dielectric margins (at either end of each winding, and at winding mid-points for an interleaved design); and around draw leads. These findings suggest stresses could be reduced by folding foil corners and edges, avoiding a draw lead, and possibly in a balance of winding numbers with their turn radius.

Nonlinear distributions extend throughout multiple units which connect in series. This distribution means that stresses on each element in a string of connected units are uneven,

predisposing some to faults more than others. Similarly uneven electrical stresses arise between each terminal and unit housings connected at a voltage common to all units on a rack.

For multiple units positioned on a rack, corner units experience greatest stresses. Such units are the furthest from a rack tie and hence a common rack voltage, subjecting dielectrics between housing and internal elements to large voltage disparities. Moreover, corner units are where high-voltage connections are typically made. Each of these factors place corner units under greatest electrical stresses and therefore puts them at heightened fault susceptibility.

Rack ties induce inflection points in voltage distributions on one side of the rack. Greatest element stresses are therefore seen on the same side as a rack tie, and where a midpoint between such ties is particularly prone to large voltage disparities (and therefore dielectric stresses) between adjacent terminals which coincide with where bushing and jumper connections are made. Thus, alternating the side of rack ties on adjacent racks could aid future monitoring.

Winding foils contain and direct heat, such that unit housings heat symmetrically. Faults which occur outwith element windings give rise to asymmetric thermal profiles on units' outer housing, and those within windings release heat which is directed along the plane of foil alignment (mostly in parallel to winding turns), which gives rise to: two warm faces; two intermediately warm faces; and two faces which are comparatively cool. Knowledge of internal unit designs which includes the orientation and placement of windings is therefore a prerequisite to monitoring which uses housing temperatures.

Faults with a thermal output (such as partial discharge) can be located with heat. The distribution of heat on a unit housing's warmest faces can be used to identify an element winding; should winding foils terminate on opposite sides as in the simulated model, the proportion of heat on intermediately warm faces (relative to one another and to warm faces) can allow a layer within the winding layer to be approximated as a fault location; and faults outwith elements can be identified from an asymmetrical heat profile.

Fault location could be more accurate if winding dielectrics end on opposing sides. Moreover, should winding directions then be alternated for adjacent windings, it could become easier to differentiate between the windings which contain a fault and give rise to an uneven heat dissipation on opposite sides, which could aid future fault finding based on thermal profiles.

Single-capacitance windings are much simpler to prepare than interleaved designs. Other considerations for the process of preparing windings are: the inclusion of break points to accommodate rests and interruptions; maintaining winding layers under tension (which influences the resulting capacitance); corrective measures to keep a winding on track as it is being wound; and building from the inside out by beginning with windings is preferrable.

Tolerances are inevitable, so configurable models can be adapted to include them. This is another reason reconfigurable models are useful. It is sensible to intend to build more windings than are necessary as a result of anticipating some errors and waste (unless processes are already well refined). Corrections made during a winding process allow errors to be accommodated before they grow out of hand, and would also be necessary for continuous production.

7.3 Future Modelling Options

The primary conclusion of this work is the preparation and publication of a capacitor unit model file, which although imperfect and simplistic, is configurable, extensible, permits study reproducability and can act as a starting point for future work. This Section aims to: • summarise learning outcomes from development efforts as part of this project; • outline the most recent feedback for the current iteration of the model; • note known areas where it might be improved; and • articulate requests made to COMSOL for simulation environment enhancements.

7.3.1 Modelling Practise

Structures used during early stages of model development improve model readability, but are ultimately unnecessary and in some respects could impede the model. While the intention behind a layered approach is to improve model readability, this has a consequence on `Mesh` efficiency as small `Domain`s resulting from layer boundaries at a unit's corners are modelled where materials are continuous in real assets, so the model could be simplified. Other suggestions for improvements to the model, and principles to follow during development efforts, are listed here:

- early clarity on the model to be developed (a result of not having had clarity on discharge resistor placement during early phases of this project is the model's support for a single discharge resistor fixed between terminals, which is not a realistic representation);
- build configurability, where parameters apply to a `Geometry` and to its `Definitions`;
- use `Definitions` common to `Mesh` sequences, `Materials`, and `Physics` interfaces;
- it takes time to learn about appropriate features ideally in advance of model development;
- in the case of this model, this could involve the layered shells functionality;
- scripting is powerful and facilitates sophisticated modelling, but makes use of Java; and
- complement efforts to build familiarity with the COMSOL® environment by – engaging with COMSOL support early (if available, support can help reduce barriers to the next increment of an iterative development process and help save time in the long run) and – attending a COMSOL conference will help build understanding of the environment.

7.3.2 COMSOL Support Feedback

Throughout the process of building familiarity with the simulation environment and the preparation of this model, best practise has not always been followed initially, so frequent revision of the approaches taken has been useful and specific feedback has proven necessary. Any developer can be tempted to assume that an intended behaviour has been correctly modelled, but one learning outcome gained over the course of this project is that there can be significant benefit to wider (preferrably external, and early) feedback to validate an implemented approach.

In addition to community support online, occasional support from COMSOL has clarified use of the simulation environment and shed light on best practise options available. It is best to create the simplest model possible to share with COMSOL support as a basis for an enquiry. Throughout this project, insights and suggestions have been used to improve the model, but some suggestions remain to be incorporated in further work should this model be used in future.

Aside validating some modelling approaches taken, the most recent response (of which the latter 4 remain to be included in future work) from COMSOL support included suggestions to:

1. use `I-Terminal` nodes in place of `External I vs. U` nodes;
2. choose an `Electrostatics` interface over `Electric Currents` for capacitor models since it can account for dielectric polarisation;
3. note that `Ground` nodes in coupled (ES, CIR) interfaces do not represent the same ground;
4. optionally remove bushings and rack mounts which have no bearing on internal dielectrics and are not the focus of any `Study`[3];
5. optionally merge some of the smaller `Domain`s for a cleaner model and improved meshing;
6. optionally use `Floating Potential` nodes[4] for terminals in the `Electric Currents` interface since capacitance can be calculated with `Surface Integration` of $\frac{Q}{V}$; and
7. in more detailed models, `Iterative` solvers could be used for one interface and `Direct` used on another (which could be appropriate where interfaces are coupled together).

7.3.3 Options for Further Model Development

Literature relevant to resolving fault detection and location [20] to the level of individual units within a capacitor bank is limited, despite such a capability being an opportune outcome of an inherently modular design. Stipulations in literature and industry standards for minimum numbers of capacitor units and elements connected in series (to offer redundancy and a 110% overvoltage margin [38]) appear to be based on an assumption of a *linear* increase in voltage over the healthy capacitor elements which remain. Moreover, faults which pose a threat of further breakdown 'cascading failure', in the event of transient overvoltages are of particular concern in terms of capacitor bank protection and ought to be resolved quickly.

In part due to its scalability, this model can be replicated and extended to form multiple-unit models. Such models would not only improve the simulation of individual units but also the emulation of capacitor banks composed of multiple units. Accordingly, the existing model could be extended to include studies where:

1. a greater number of units are simulated; and
2. where multiple faults are introduced (as particular combinations of early faults could result in distinct characteristics, and differing susceptibilities to further degradation).

The remainder of this Section considers further opportunities to improve or adapt the model.

Dielectrics

One of the most significant abstractions made in this model pertains to its `Materials`, and in particular the dielectrics, employed. As introduced in Sections 2.2.2 and 3.1.4, advances in modern dielectrics are crucial to improvements in capacitor unit technologies and designs: but have been abstracted in this model by considering them a single homogenous `Domain` of constant dielectric. This therefore remains a prominent potential area for future work.

[3] While this option could improve model efficiency, these aspects are primarily used for aesthetic purposes.
[4] To some extent this would improve the view of the model tree from the user interface too.

Time Dependent Studies

With the ability to simulate signals which change in time allows harmonics, as well as other signal artefacts, to be applied to a model. This presents an opportunity to investigate where stresses might arise in capacitor units from the application of selected frequencies, which could be relevant, in particular, to characteristic harmonics seen at HVDC converter interfaces. Moreover, the propagation of impulses through a capacitor bank could be determined as part of an assessment of a disproportionate susceptibility of certain units (or elements) within a bank to impulses due to frequent capacitor (or converter) switching.

Domain Merging

As noted in Section 7.3.2, an unintended consequence of having built the model as a series of concentric layers is that a number of discrete `Domain`s result even as part of the same continuous materials, some of which are small (such as at corners). Since `Mesh` elements conform to `Domain` boundaries, improvements in `Mesh` efficiency might be possible by merging such `Domain`s or rebuilding a model without this layered structure. Thus, merging smaller `Domain`s would relieve the `Mesh` sequence from having to conform to these (otherwise meaningless) internal boundaries. Conversely, maintaining planar `Domain`s makes `Mapped` and `Swept` meshing possible (which is useful on a thin outer housing), and aids readability, so a balance might suffice.

Terminal Edges

In future the model could be designed to make use of more practical winding designs. At present, each winding terminal has one edge on the outermost turn of each winding, but also an edge at the centre of each winding. While some capacitor units may employ such designs, having made practical efforts in building a laboratory prototype (discussed in Chapter 6) it is now apparent that such designs are impractical for three reasons:

1. multiple layers used in windings can splay as they are rolled around a mandrel, so beginning from a centre-point to turn in both directions minimises this splay;
2. it is practically difficult to position terminal layers such that they each begin in the centre of a winding but extend in opposite directions, whereas it is comparatively straightforward and space efficient to concurrently turn winding layers which begin and end together; and
3. foil corners and edges are prone to electric field stresses and as such (due to effects discussed in Chapters 4 and 5) are best positioned exclusively on a winding's outer turns rather than its centre where temperatures are greatest.

Moreover, were this refinement to be reflected in an updated model, and where a mid-point boundary is selected within a winding and using the same process for foil selection discussed in Section 3.2.2 is followed, an additional benefit of having foils continue throughout a winding in both directions could be that the number of boundaries which could be captured by a selection sequence of a given length would double.

Layered Materials

COMSOL® version 5.5 introduced support for layered materials, which allow a series of heterogeneous material layers to be specified as single, coherent `Domain`s in the model: this simplifies

the model `Geometry` and facilitates development while allowing modelling interfaces to solve for a series of distinct material layers. Detailed documentation on using layered materials within a model is provided by COMSOL [83], [93]. During this project, the foundation of the model had by this point been completed, but were this approach to have initially been used for the model (or adopted as a future refinement) then the `Geometry` and `Mesh` would be simpler, and hence some of the process used during this project to accommodate narrow spaces between foils could have been avoided by leveraging this improved functionality of the simulation environment.

Reconfigurable Meshing

To allow a `Mesh` sequence to be resilient to changes in the `Geometry`, it is best to apply each step to an existing `Definition`, ideally where `Definition` selections are established spatially by using `Box` and `Cylinder` selections which can be combined. In this manner, `Explicit` selections can be avoided both for `Definitions` and in specifying a `Mesh` sequence, in attempts to limit any manual reconfiguration necessary in reapplying a `Mesh` to a changed model `Geometry`.

Configurability

Configurability aids iterative development: iteration is intrinsic to a variety of development methodologies and continuous improvement principles; it is sensible for modelling to be updated as better practises and new information emerge. Moreover, model-based techniques ordinarily require validation in laboratory settings, and as such it is prudent to match any simulation model to the unique characteristics of a physical replica (such as that in development as introduced in Chapter 6). If a model-based study can be readily updated and repeated to accommodate such characteristics, its development need not wait the prior preparation of a prototype to facilitate matching one after-the-fact, and the whole investigation can be streamlined. Such repetition is made easier if each step requires only minimal intervention.

Parameters aid such configurability. `Geometry` sequences can be specified based on parameters, enabling easier development of the geometry specification, and allowing a model to adapt to changes in specified parameters (thereby preparing the model for use in COMSOL®'s '`Application Builder`' feature), and improving readability of the `Geometry` sequence.

7.3.4 Improvements Requested from COMSOL®

As highlighted in Section 3.2, and as above, parameters can be used as a common point from which to specify a model configuration: since COMSOL® supports conditional statements within a `Geometry` sequence, the nature of a model can adapt based on specified parameters.

As it stands, the model developed for this research project uses highly repetitive sequences for the `Geometry`. If not only conditional (`If`) statements but also iterative or repetitive ('`Loop`') statements were available for the `Geometry`, control flow could be made fully adaptable and more sophisticated models could be developed from a simple, readable `Geometry` sequence.

Developing detailed model geometries comprising many repeated components can be tedious. `Array`s offer a limited form of repetition: the same input node can be repeated, but not altered each time (this would therefore not allow, for instance, a series of replicated winding turns to be specified where each subsequent winding has a slightly larger turn radius). Should support for both *conditional* and *iterative* nodes be available for `Geometry` sequences, user parameters

could help establish readable, adaptable, and sophisticated model geometries which would be straightforward to develop and scale.

Once a `Geometry` has been established, a comprehensive set of `Definitions` (while not necessary) can make subsequent stages of a modelling exercise easier to configure and more resilient to changes, as `Materials`, `Mesh` sequencing, and `Physics` interfaces can each refer to `Definitions` specified in relation to the `Geometry`. In contrast, the use of `Explicit Selection`s to specify `Materials`, `Mesh` sequence, or boundary condition interface, necessitate a user checking each is correctly specified following a `Geometry` change.

Where a model uses parameters, `Definitions` specified exclusively according to these parameters can adapt in line with geometrical changes: good `Definitions` are therefore key to reconfigurability and in allowing a model to adapt to change without requiring users to engage in a time-intensive manual adjustments.

However, if parameters are used for control flow through conditional or repetition declarations in a `Geometry` sequence, but support is not similarly provided for `Definitions`, then manual intervention invariably remains a necessary response to any change which makes use of `Geometry` conditions or iteration. Were COMSOL to support control flow options not only in `Geometry` sequences, but also for `Definitions`, users could develop models which are readily configurable by a set of parameters which apply directly to both a `Geometry` and to `Definitions` on that model. `Materials`, `Mesh` sequencing, and boundary condition interfaces can therefore all be declared relative to `Definitions` and become free from a need for manual updates in response to a change in model parameters.

Hence, by creating 1. new ('`Loop`') nodes for repetition in `Geometry` sequences, and 2. extending support for both types (`If` and '`Loop`') of control flow nodes to `Definitions`, COMSOL® could support greater adaptability within model files which use parameters, and hence improve on-the-fly dynamism and usefulness of applications prepared using the `Application Builder` facility from a common model file. These aspirations were raised in person at COMSOL Europe 2019 in Cambridge with members of the COMSOL® development team and subsequently articulated via email [94].

Since studies in COMSOL® have potential to take hours or even days to run, a useful design feature might be to permit a study to be paused at potential break-points during a solver process (such as between iterations), and for the solver state to be serialised to hard disk. It would then be possible for users to work in a more agile manner, with an ability to realise which studies could take a long time, and potentially interrupt them to prioritise more meaningful work, or to resume an ambitious `Study` overnight or on weekends. This could improve productivity and expand the scale of model COMSOL® can meaningfully support, as more complex simulations could run without burdening a user during the time their attention is available, and by facilitating an approach where large problems could still be solvable with typical computing resources.

7.4 Future Research Options

To complement specific examples of future work as might be applied as amendments to the model, this Section offers more general suggestions as possible directions for further research.

7.4.1 Further Exploration of Dielectrics

Temperatures and frequencies affect capacitance, so are worth exploring in more detail. This Dissertation primarily focusses on geometries and designs of modern capacitor units and offers only a cursory exploration of the dielectrics used, but these are primary components in the system, and future work could explore how modern dielectrics might behave in view of the temperatures capacitors are likely to experience in operational contexts, alongside different signal components they are likely to see.

7.4.2 Nonlinear Voltage Distribution

Section 4.1 considers a theoretical, static assessment of a unit's relative resilience following a first element failure, but it is also true that in the best case such a unit would be derated. However, there are further practical implications in part since capacitor units are not typically derated in response to faults: it is firstly near impossible to identify a faulted unit in an operational context; and also that the whole bank would be impacted by a change rather than only one unit at a time. Thirdly, redundancy is ordinarily built into the bank. Derating any one unit might be impractical, but detection could be used to monitor faults and allow the whole bank to be derated when a certain health condition is reached. This is one reason monitoring, and fault detection and location are important capabilities.

The Distribution of Stresses throughout a Bank

In the literature, it is reported to be necessary that individual capacitor units are measured, in order to appropriately situate them to balance parallel strings of a bank, and in turn to balance parallel phases of a capacitor arrangement [1], both of which have a bearing on protection settings and sensitivities [45]. It is assumed that unique measurements from each unit are manually sorted to achieve a reliable balance between strings and between phases.

However, where a nonlinear distribution of potential differences is established, we can infer that a typical capacitor bank will see greater electrical stresses on units closer to the high-voltage connection. Figure 4.24 and Figure 4.25 further reveal how potential differences are distributed throughout a bank with parallel unit strings: as Figure 4.24 suggests, the influence of this distribution will extend to rack voltages, which each sit at some proportion of the potential difference of adjacent racks. Chapter 4 modelled only a single rack, so the manner in which multiple racks share a potential difference remains beyond the scope of this Dissertation, but it is possible to infer that a disparity could grow between strings of the same bank where one string consistently supports rack voltages (as in Section 2.2.3 and Figure 2.4, as an example).

Specifically, Figure 1.11 illustrates one capacitor bank arrangement in place at the converter station visited, where one side of the parallel arrangement supports rack potential. At this point, and in lieu of more detailed studies, an engineer might choose to assume that the same principles of Figure 4.24 apply across all racks of the bank:

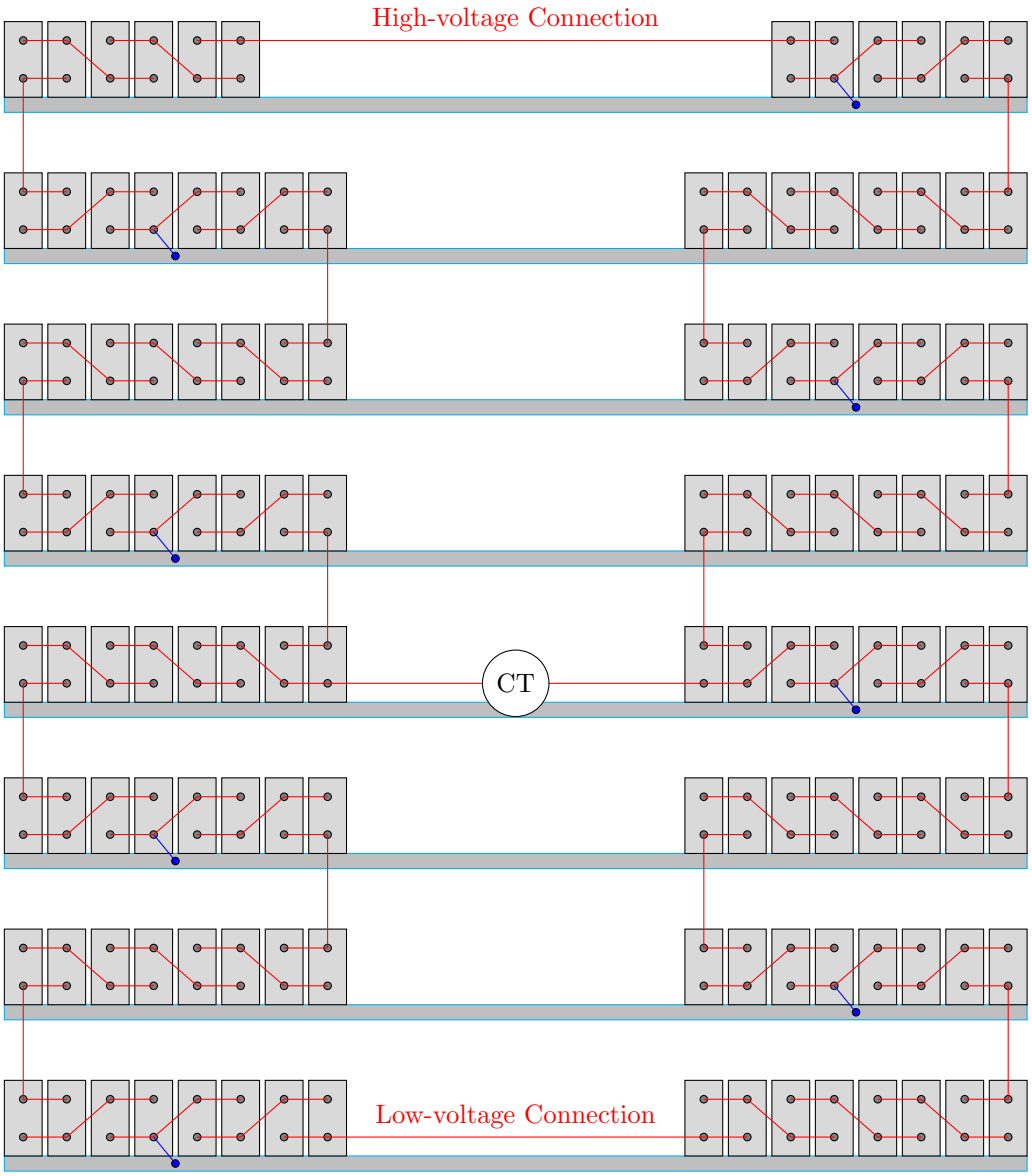

Figure 7.1: Jumper (red), rack tie (blue), and CT connections for both sides of a large bank

1. that the rack tie introduces a point of inflection (more clearly depicted by a return to $\Delta V = 0$ shown in Figure 4.25); and

2. that rack potential is some proportion less than half of the potential difference applied across units on one side of the rack.

An extrapolation of these findings to a hypothetical study of a bank (where multiple racks would be modelled), might find that an arrangement with rack ties on one side sees inflection points on only the side supporting rack voltage, and consequently

- greater stresses placed on unit dielectrics from stepwise changes in potential seen where inflection points occur (as in Figure 4.25, where ΔV for Side A exceeds the maximum ΔV for Side B); and

- risk that a discrepancy could arise between potentials on each parallel string.

Instead, should rack ties be placed on alternate sides of adjacent racks, as in Figure 7.1, it might be possible to

- share maximum stresses placed on unit dielectrics between parallel strings;

- smoothen the ΔV curve between rack ties, as these would become twice as far apart;

- reduce the likelihood of a potential difference developing between equivalent units on either string; and

- possibly even facilitate improved fault detection or location in future, as a fault on either side of a bank would then have some influence on voltage for one rack or another.

Parasitic Effects in Relation to Frequency

As capacitor banks can often be placed in proximity to power electronics and subjected to associated switching impulses over the course of their operational life, such design improvements could be considered in terms of improving unit resilience to high-frequency signals. Subtle inductive effects also arise as a result of winding turns, meaning that stray effects within capacitor units are both small and can also be overall either capacitive or inductive. Either way, while element groups considered on their own act as high-pass filtering components, parasitic shunt capacitances would pass low-frequency components and would in and of themselves help attenuate high-frequency signal artefacts (harmonics, impulses, etc.). It could be the case that the outermost dielectric between unit housing and the windings is the likely location for degradation and faults due to high-frequency signals. Thus, extending the housing internally to better segregate winding groups could serve to increase stray capacitances, and thereby improve the attenuation of high frequency signal artefacts to in turn improve resilience of these assets to routine impulses. Furthermore, by increasing parasitic capacitances in such a manner it could also exaggerate the change seen in terminal impedance characteristics in the event of element failure and consequently improve diagnostics as a secondary benefit.

Improved resilience to high frequencies could be important for capacitor units proximal to converter stations due to their exposure to impulses from the conversion process itself (i.e. thyristor and IGBT switching), but also from the increased switchgear duties required to place capacitor loads in and out of service to match reactive power support required to meet daily

variation in electrical demand, as intimated in Section 1.4.1. However, to incorporate such features in a design, it would be important to be aware of the electrical field stresses which could be exacerbated by inward protrusions from the housing.

7.4.3 Unit Allocation

As discussed in Sections 4.1.3 and 4.1.4, given knowledge of internal unit design and assuming winding capacitances to be approximately equal, it is possible to solve a circuit to infer fault locations from changes in terminal characteristics. This approach could be better informed should manufacturers measure and record winding capacitances prior to unit assembly.

While it is easier to manufacture units in batches to meet similar values of capacitance [69], manufacturing tolerances naturally mean that each unit has a unique value. In forming a set of unique units into a balanced arrangement, characteristics of individual units risk being combined in a manner which could provoke an unintentional imbalance. Accordingly, since requirements stipulate capacitor units should be individually measured and positioned to balance a bank [1], industry could benefit from a tool to keep track of unique measurements and search for the most appropriate unit positioning. By so doing, it could offer a part-automated means of sorting units by electrical characteristics[5], to:

1. balance parallel sides of a bridge arrangement (or parallel phases or strings of an overall installation) [1];
2. suggest positioning cognisant of variation in rack potential to minimise stresses due to unit positions on a rack;
3. position the units most robust to high-frequencies where harmonics and impulses are most likely to stress the bank;
4. offer subtle grading or redundancy at points of high-voltage connection to improve resilience at vulnerable points; and
5. routinely compare recent unit measurements with historical ones to infer failure locations as faults progress over time.

Such an allocation tool might alleviate engineering personnel from manually assigning measured units to rack positions, and therefore save a limited amount of time. Moreover, once uncoupled from a need for human deliberation, this task could be performed at short notice, in tandem to other work, and potentially in environments where it might otherwise be difficult to concentrate on nuances of bank design. As such, a part-automated unit allocation tool could even allow this part of a design to take place on site while units are individually measured and prior to asset integration and commissioning, thus benefiting from recent unit measurements. Such a tool could encourage consistency, reduce scope for human error, and permit allocation to be done at short notice should this become necessary. It would also facilitate comparison between measurements recorded at different points during an asset lifecycle, and hence pave the way for more frequent measurement or ‘online’ monitoring.

[5]While the first of these points has a basis in the literature, the rest are speculative.

7.4.4 Capacitor Asset Salutogenesis

Pathogenesis is the study of disease and is the predominant focus for modern medicine; salutogenesis, on the other hand, considers what is required for a system to function in good health [95]. In terms of engineering, reliability centred maintenance remains an epitome of asset management, but even this retains a focus on mitigating inevitable failure modes and is analogous to a framework for managing pathogenesis in medicine rather than necessarily promoting 'salutogenic' asset *health*.

An engineering equivalent to a salutogenic model would make use of online condition monitoring as is already occasionally employed in industrial contexts, but would move beyond occasional interventions intended to prevent deterioration to instead offer a part-automated continuous support 'harness'[6] in which assets operate more comfortably, and hence reliably. Such a harness would require means of 1. monitoring, to measure stresses and assess behaviour akin to the role of existing approaches in continuous condition monitoring; 2. control, in order to infer condition and evaluate health; and 3. actuation, to provide some (even limited) means of intervention known to limit the effects of usual stresses on asset health (dielectric integrity).

Therefore, the modular nature of capacitor banks presents an opportunity, not only for condition monitoring, fault detection, and location, but in offering controllability toward support for adaptable, salutogenic operation in which these assets are supported to operate reliably, dependably, and for greater longevity by promoting dielectric health. Given the relationship between capacitor hot-spot temperature and longevity discussed in Chapter 2 the application of the Peltier effect for thermoelectric cooling could be a suitable approach if directed to individual units which see the greatest temperatures, electrical stresses, or with prior history of internal faults. As unit temperatures influence their support for voltage, and therefore the harmonics involved, maintaining a temperature range is therefore also imperative to asset function.

Applying an array of thermoelectric modules (together with heat sinks on exposed faces) to individual capacitor units would allow for excess thermal energy to be recaptured as electrical energy, and to reinvest some of that energy into directed cooling where beneficial, all while offering a possible means of temperature monitoring in the interests of other aims presented in this Chapter. Selective cooling of individual capacitor units in a controllable manner such as this could allow: • mitigation of uneven heating as a result of sunlight, for instance where one side of a bank is at risk of sufficient imbalance that a bridge arrangement could trip; • cooling to be directed to the very surface of unit housings, therein improving its effectiveness compared to ambient cooling of an entire filter hall, for example; • improved resilience where certain units are subject to increased dielectric stresses – due to their electrical position relative to others in a bank, or – due to existing degradation or element breakdown; • to mitigate against decreased asset longevity incurred due to variation in daily and seasonal ambient temperatures; or • to retain effectiveness of units subjected to similar variation in electrical load.

7.5 Contributions

Contributions from this Dissertation are briefly outlined in the remainder of this Chapter, by firstly considering • novel contributions of this work, • early collaboration and impact opportunities, • before a short summary is provided to conclude.

[6] 'Harness' is a term borrowed from test-driven software development.

7.5.1 Novelty

This Dissertation makes the following contributions, which offer novelty in the context of present-day publicly available literature. These are further outlined in the following Sections.

1. Winding design options are presented, and discussed in terms of construction practicality.
2. Voltage distributions and electric fields within units are modelled to improve understanding of where dielectrics are most stressed, and the relationship with location of series element fault and unit terminal characteristics is explored: each fault location uniquely influences a change in unit characteristics, with high-voltage locations offering a change which is discernable and thus allows an internal fault location to be inferred.
3. Different unit configurations are used to illustrate heat dissipation from capacitor windings: this leads to principles which offer an additional means of fault location inference within hermetically sealed units, and potentially for *incipient* faults.

Winding Designs

Literature refers to windings [29], [96] but commonly focusses on the materials involved in modern capacitor unit designs. This focus is appropriate given the wealth of research underway into improved dielectrics and the potential benefits such advances could yield in terms of asset reliability, footprint, and efficiency. This Dissertation considers geometrical and design options for windings, including a discussion of relative benefits and drawbacks, to complement established knowledge. Naturally, manufacturers likely have established understanding of these considerations, but this is a contribution in the context of publicly available material.

Dielectric Stresses

Hermetically sealed capacitor units without external fuses appear as 'black boxes'[7] to the asset engineer interested in internal stresses or fault conditions. Literature offers various ways to detect and (to some extent) locate faults within capacitor banks [34], [39] which can be considered 'grey box' techniques, but even these offer limited information (in terms of a bank's operational characteristics differing from those measured in a laboratory environment, and in terms of the precision of locating a fault in one unit among many). Thus, while modelling is itself limited in its accuracy, it offers a more explicable 'white box' approach to understanding internal dynamics of capacitor units. In this Dissertation, model-based techniques are used to reveal where stresses on internal dielectrics are greatest, and thereby to glean insight into locations within a bank where capacitor units are most predisposed to incipient faults and dielectric breakdown.

Thermal Profiling

Incipient faults can be detected and *located* from housing temperatures. While these relationships have still to be explored and defined in detail, their underlying principles are outlined as a means to corroborate high-frequency detection techniques and with a view to the opportunities which increasingly capable embedded systems and 'internet of things' connectivity can afford. Moreover, possible design considerations are proposed, motivated by an intention to improve detection and to mitigate risk of failure cascades in future asset designs.

[7] Availability of information on a system's function can be termed: black = no; grey = some; & white = all.

7.5.2 Industrial Engagement

While acknowledging limitations on commercial sensitivity, the relevance of certain projects and the interest each might have in different areas, efforts have been made to broaden understanding by engaging industry colleagues outwith the University environment during this project, in:

- Scottish Power Energy Networks, as described in Section 1.4.1;
- General Electric, as outlined in Section 2.1.5.

Motivations for this project were presented [97] to the research and development director of Synaptec Ltd. which is a spin-out based out of the University of Strathclyde during a first-year project review meeting on the 6^{th} of November 2017. Despite "very little" academic interest in capacitor fault monitoring and location techniques globally [20], more recently Synaptec report developing "a capacitor monitoring solution proof-of-concept" prototype for this domain [98].

7.5.3 Summary

Rather than study the chosen network components (HVDC converter stations) from a *system* perspective, this project adopts a more specific approach by honing in on one asset type in particular and revisiting its fundamentals of operation through *physics* modelling. In so doing, an understanding of failure modes complements techniques introduced in existing literature which detect, track, and locate failures by highlighting where they are most likely to arise.

In this regard, three outcome contributions are reached:

1. development of an adaptable simulation model allowing this work to be extended;
2. identifying a nonlinear distribution of electrical energy within units which
   (a) affects the likelihood and severity of faults at different locations,
   (b) influences the subsequent changes in terminal characteristics following a fault, which
   (c) can be used to identify fault locations;
3. identifying uneven heat profiles which correspond with incipient fault location.

Through investigating these physical effects, this project uncovers where incipient faults are probable – information which is itself useful to engineers tasked with searching a bank of units in a limited outage timeframe – and accordingly highlights good practise approaches for manufacturers to implement now toward enabling timely fault detection later.

As with much fundamental research, this project began without data or initial models as a basis, but with an open-minded approach to the challenges faced by the electrical power industry as it evolves to face global sustainability objectives. Since this work focusses on a subject with limited modern-day research activity, it is hoped the contributions made can set some groundwork for further investigation and development. Much work remains, for instance, to expand on the associated findings in terms of: 1. estimating the prominence of proposed measurands in noisy environments; 2. quantifying the bounds of their utility; 3. in developing and proving (a) measurement technologies, and (b) any associated data processing techniques.

Moreover, this project demonstrates that modelling offers efficient (without physical waste) means to find routes to monitoring and facilitate reliability centred maintenance for *any* asset.

Appendix A

Detectable Thermal Behaviour

Results presented in this Appendix correspond to thermal studies documented in Chapter 5.

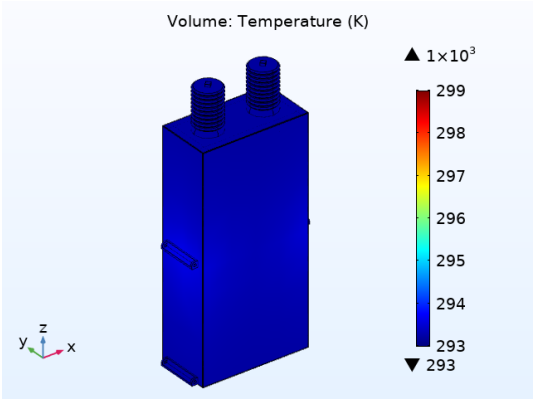

(a) Housing temperature

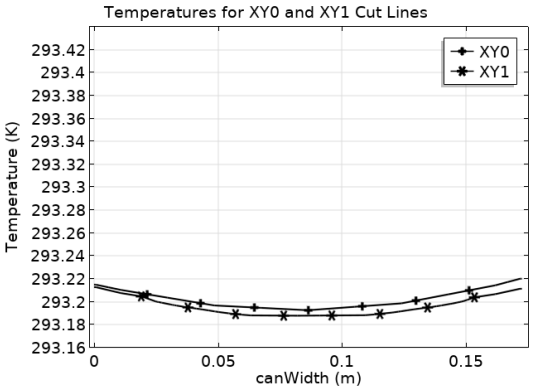

(b) XY0 and XY1 cut lines

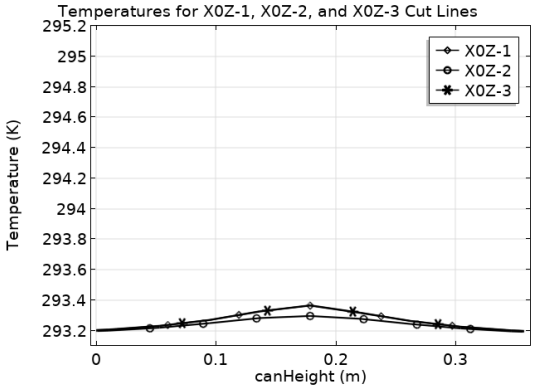

(c) X0Z cut lines

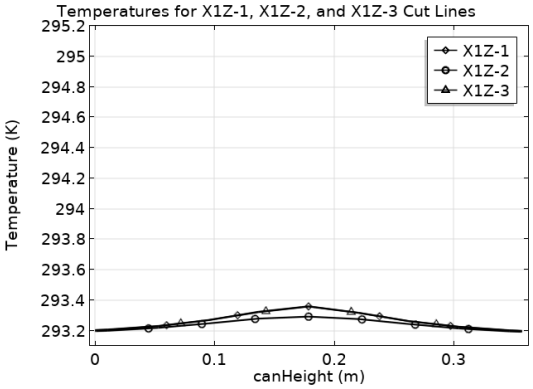

(d) X1Z cut lines

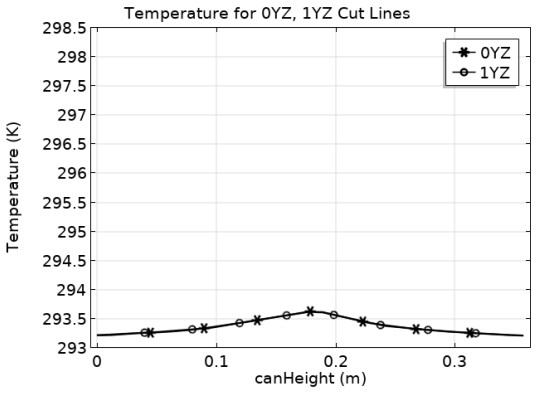

(e) 0YZ and 1YZ cut lines

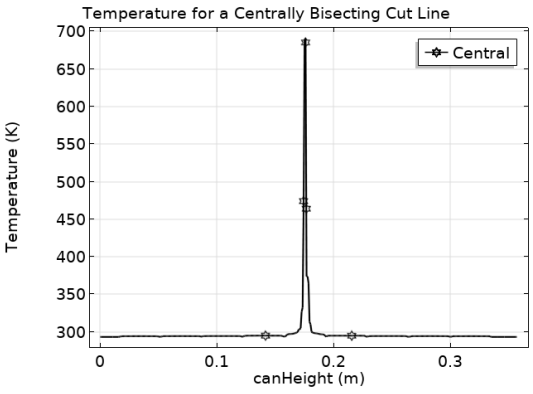

(f) A centrally bisecting cut line

Figure A.1: 20TCentre0-2W of a Set in Table 5.2

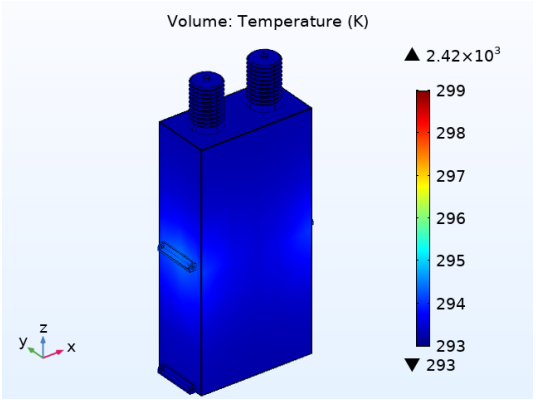

(a) Housing temperature

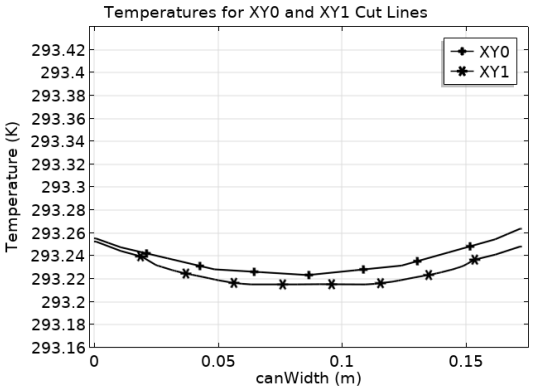

(b) XY0 and XY1 cut lines

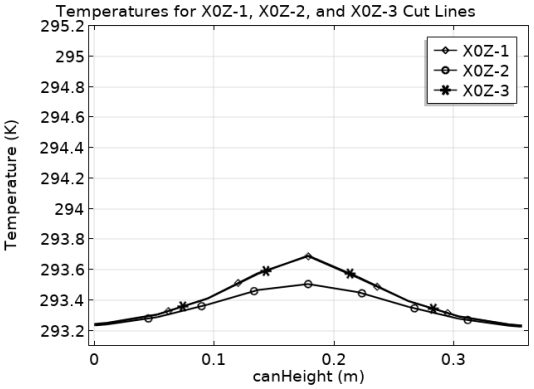

(c) X0Z cut lines

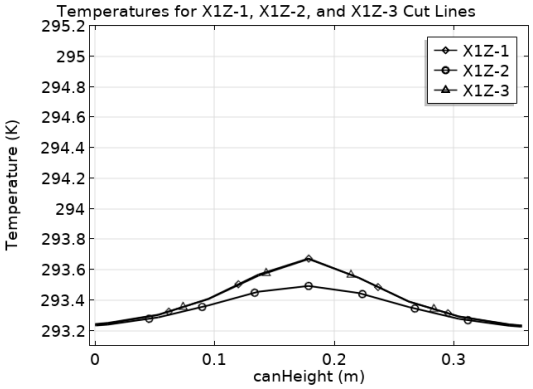

(d) X1Z cut lines

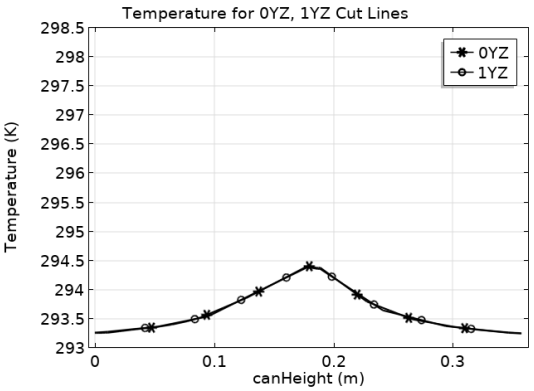

(e) 0YZ and 1YZ cut lines

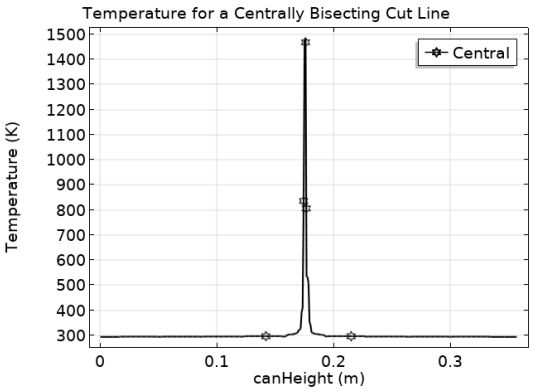

(f) A centrally bisecting cut line

Figure A.2: 20TCentre0-6W of a Set in Table 5.2

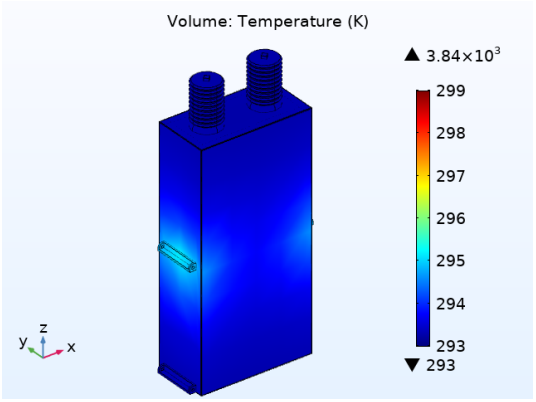

(a) Housing temperature

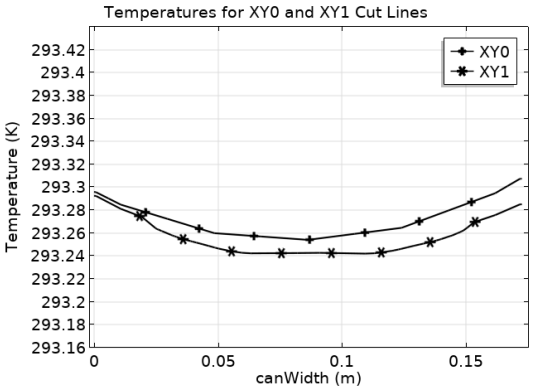

(b) XY0 and XY1 cut lines

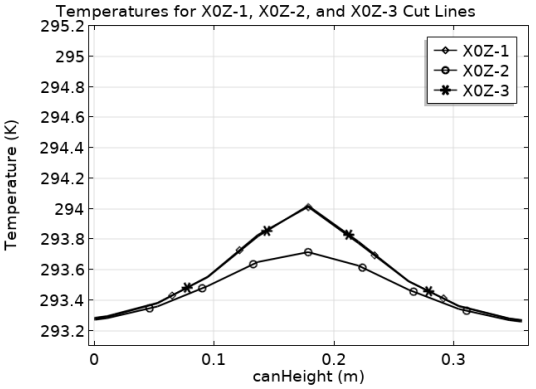

(c) X0Z cut lines

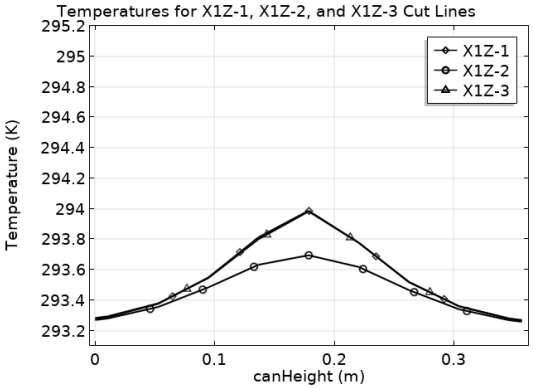

(d) X1Z cut lines

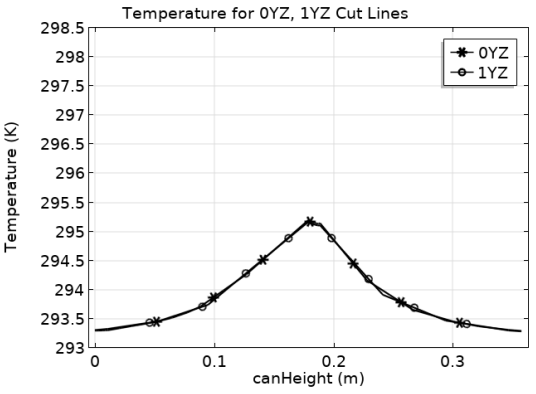

(e) 0YZ and 1YZ cut lines

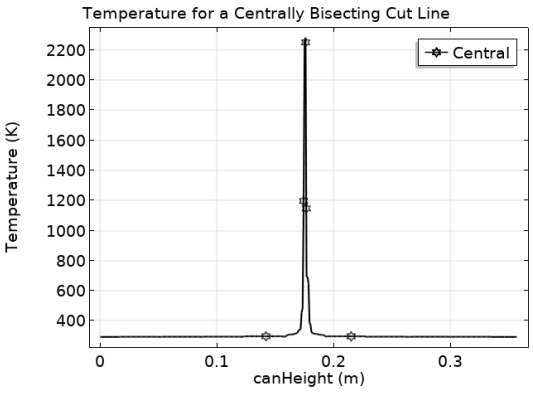

(f) A centrally bisecting cut line

Figure A.3: 20TCentre1W of a Set in Table 5.2

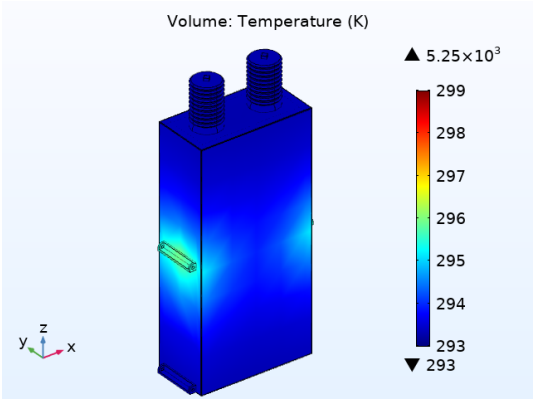

(a) Housing temperature

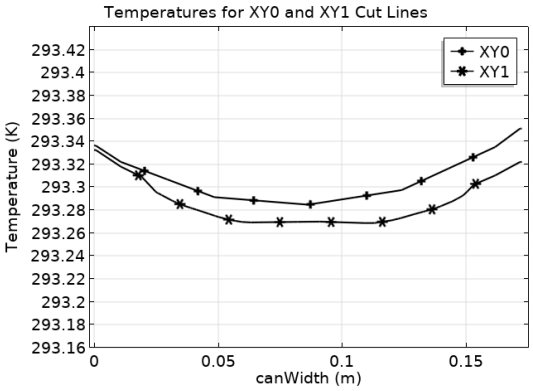

(b) XY0 and XY1 cut lines

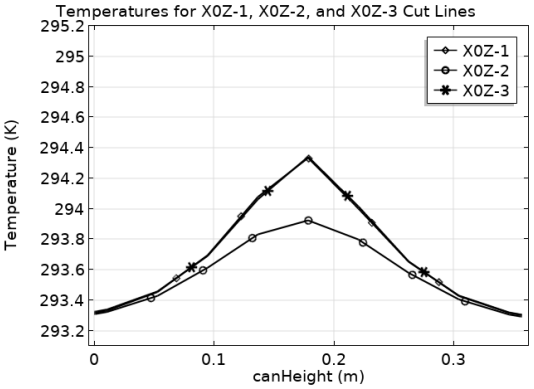

(c) X0Z cut lines

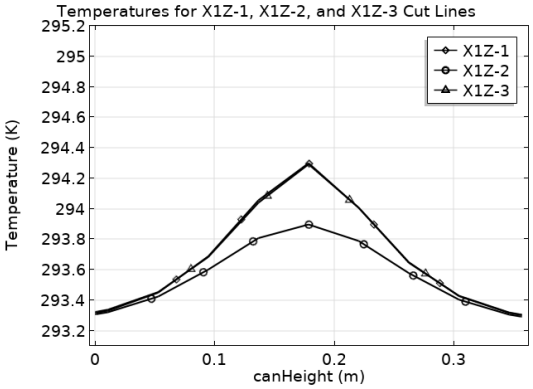

(d) X1Z cut lines

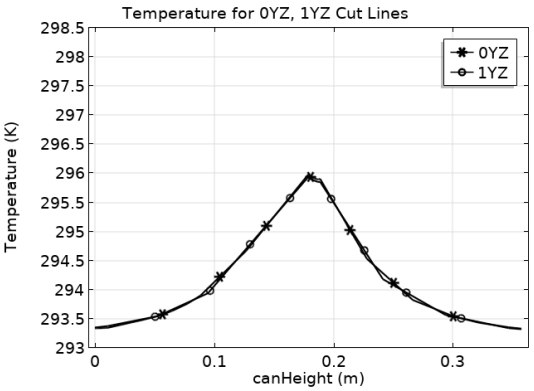

(e) 0YZ and 1YZ cut lines

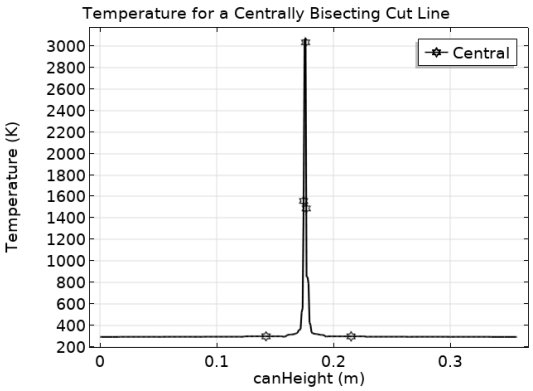

(f) A centrally bisecting cut line

Figure A.4: 20TCentre1-4W of a Set in Table 5.2

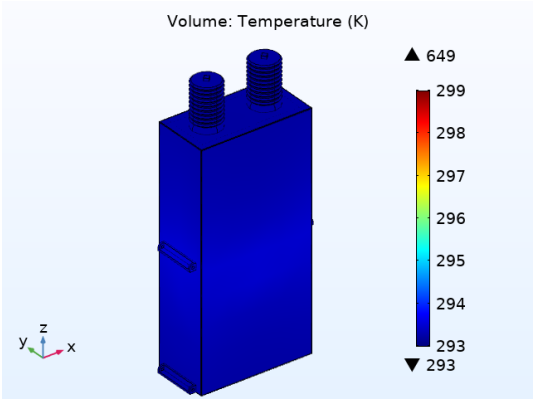

(a) Housing temperature

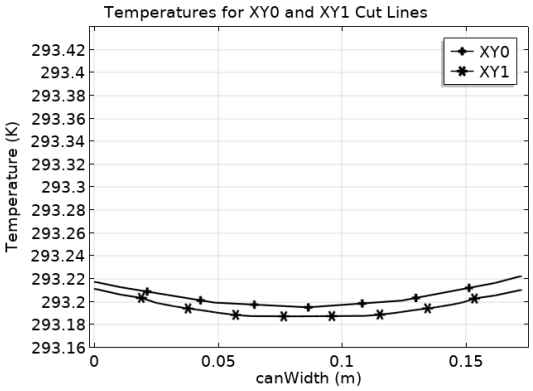

(b) XY0 and XY1 cut lines

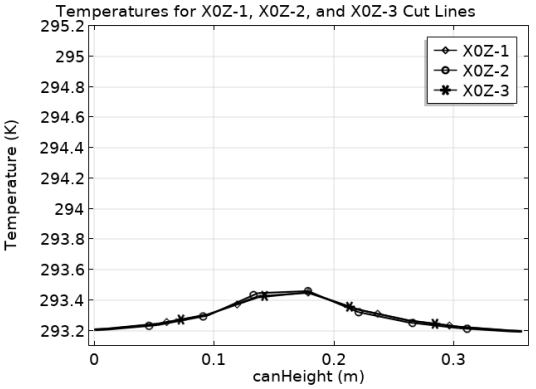

(c) X0Z cut lines

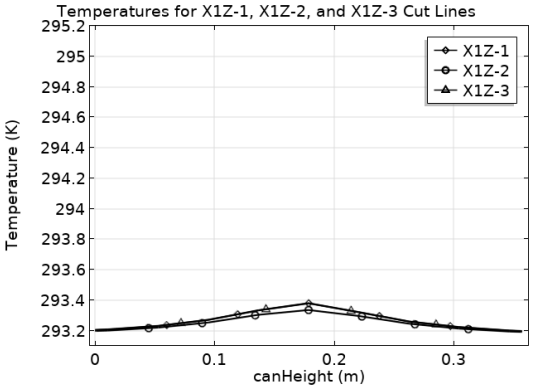

(d) X1Z cut lines

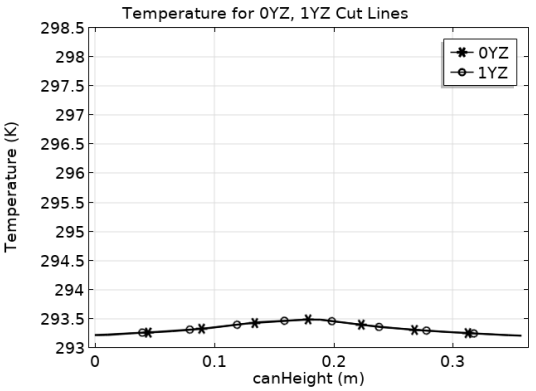

(e) 0YZ and 1YZ cut lines

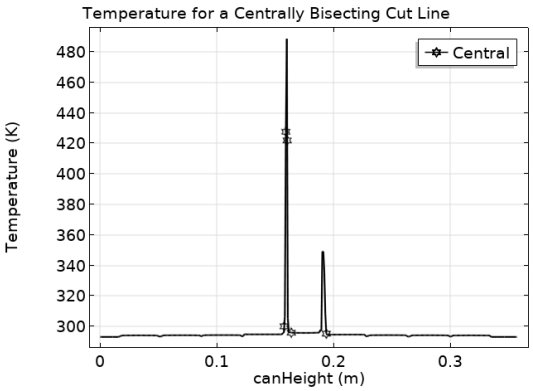

(f) A centrally bisecting cut line

Figure A.5: 20TPenultimate0-2W of a Set in Table 5.2

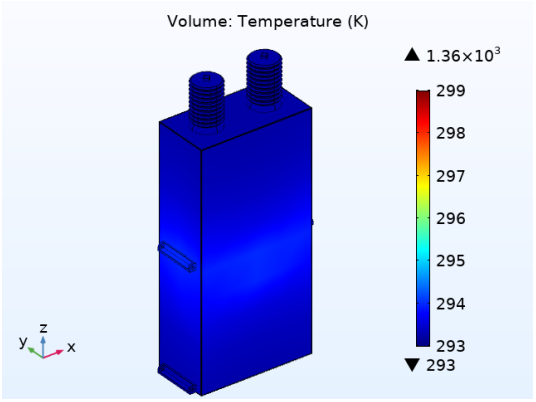

(a) Housing temperature

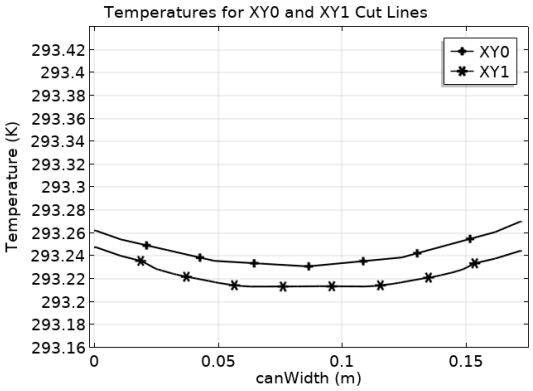

(b) XY0 and XY1 cut lines

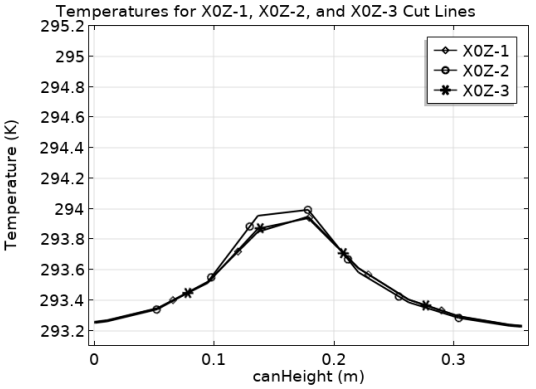

(c) X0Z cut lines

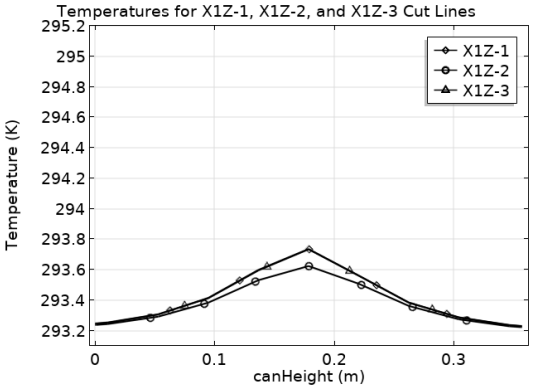

(d) X1Z cut lines

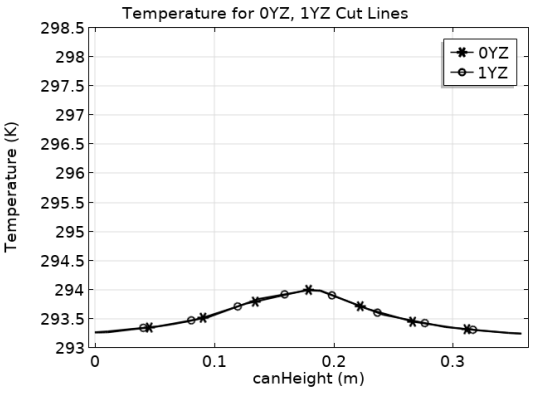

(e) 0YZ and 1YZ cut lines

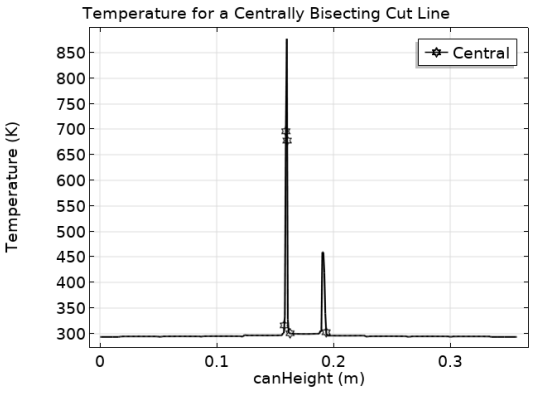

(f) A centrally bisecting cut line

Figure A.6: 20TPenultimate0-6W of a Set in Table 5.2

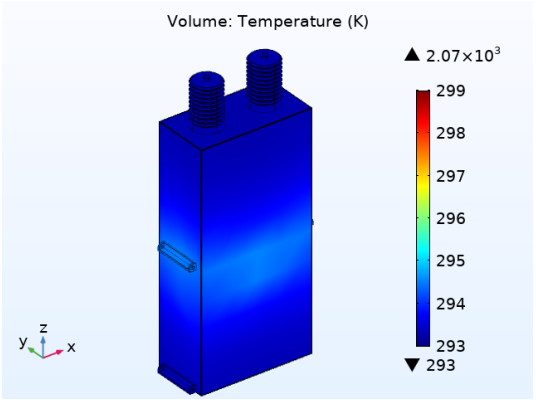

(a) Housing temperature

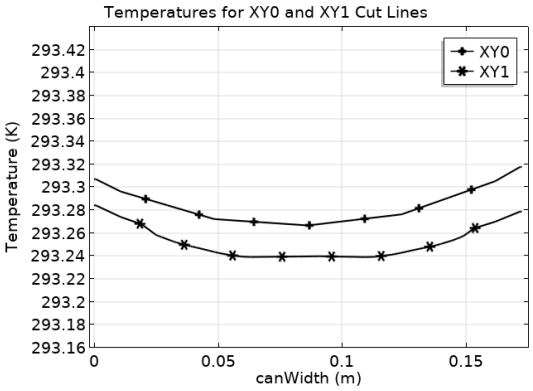

(b) XY0 and XY1 cut lines

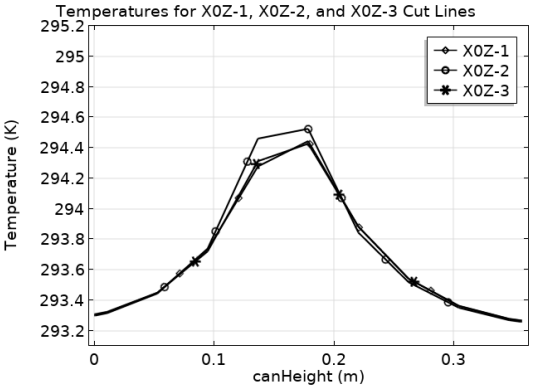

(c) X0Z cut lines

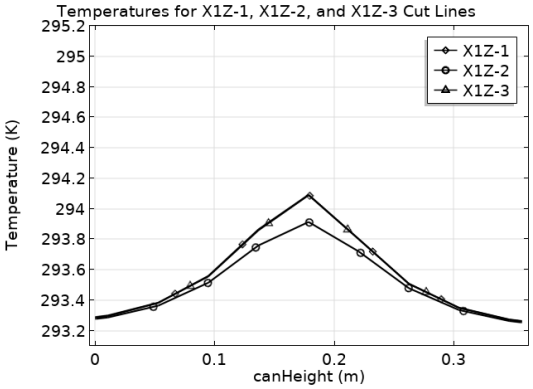

(d) X1Z cut lines

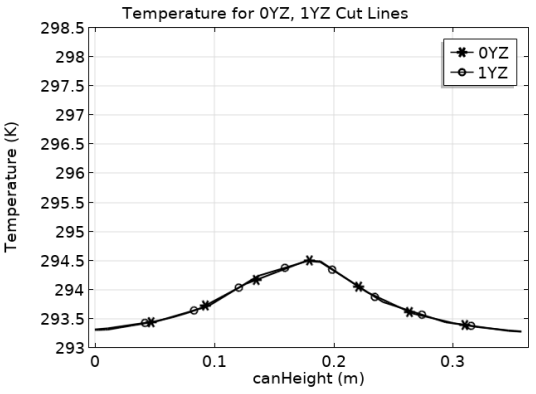

(e) 0YZ and 1YZ cut lines

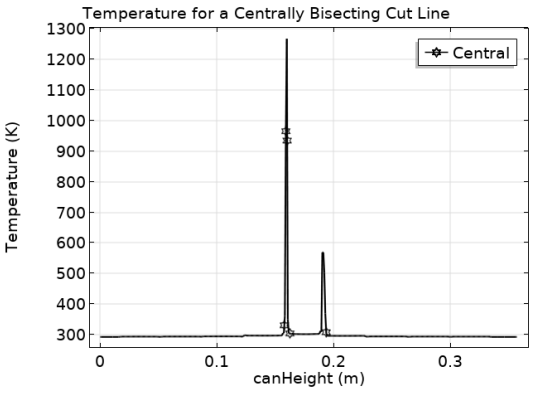

(f) A centrally bisecting cut line

Figure A.7: 20TPenultimate1W of a Set in Table 5.2

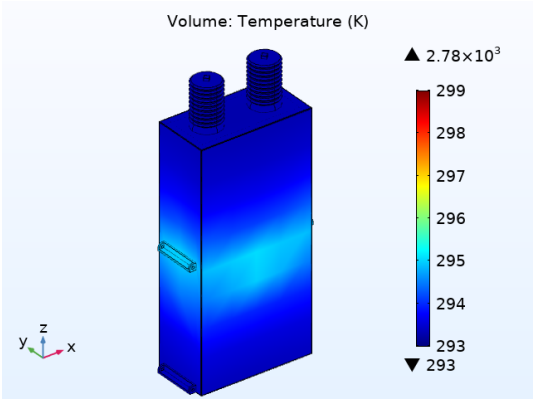

(a) Housing temperature

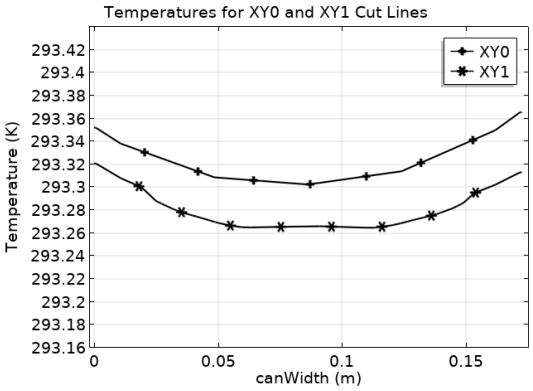

(b) XY0 and XY1 cut lines

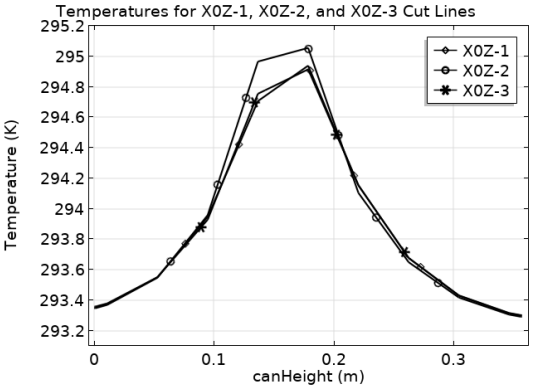

(c) X0Z cut lines

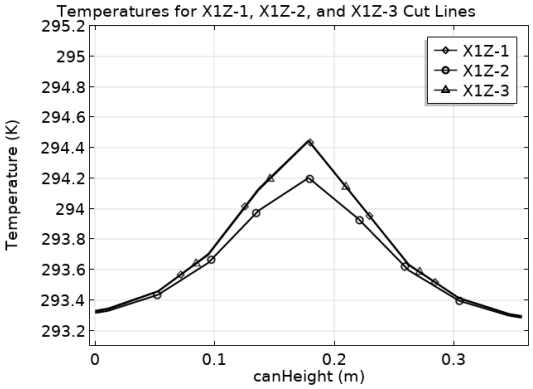

(d) X1Z cut lines

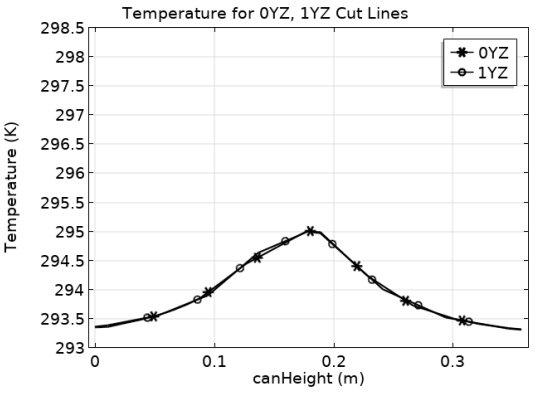

(e) 0YZ and 1YZ cut lines

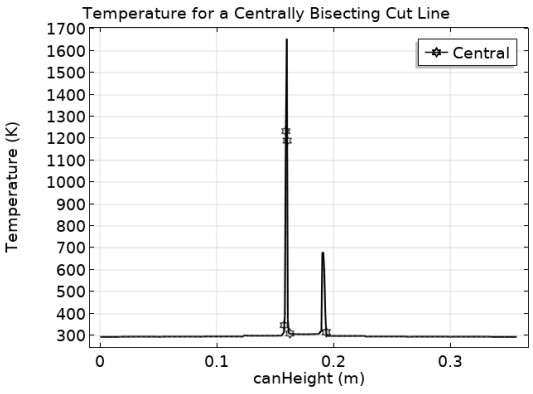

(f) A centrally bisecting cut line

Figure A.8: 20TPenultimate1-4W of a Set in Table 5.2

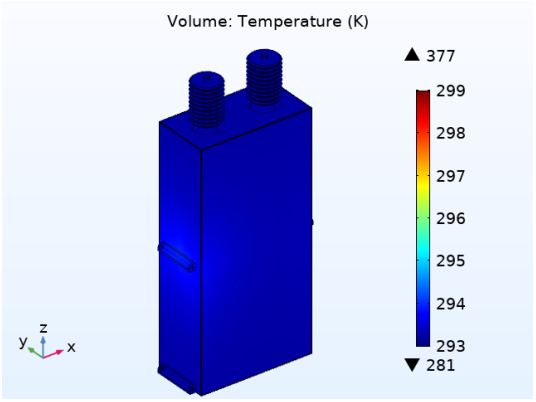

(a) Housing temperature

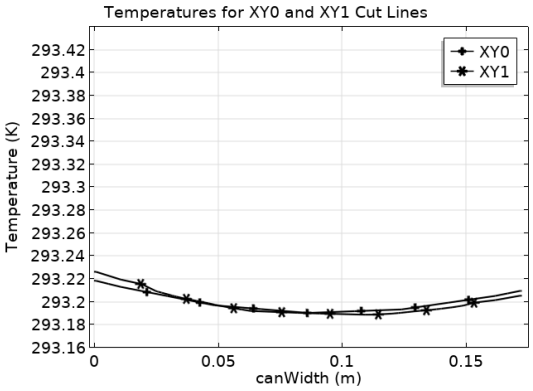

(b) XY0 and XY1 cut lines

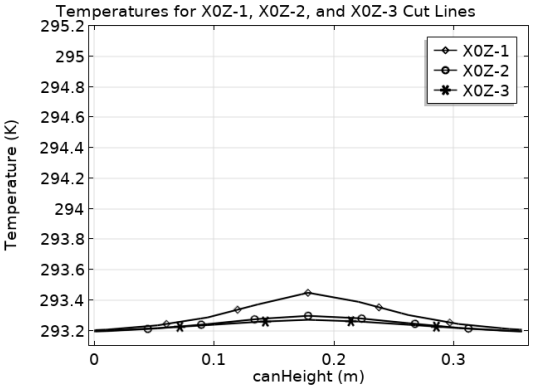

(c) X0Z cut lines

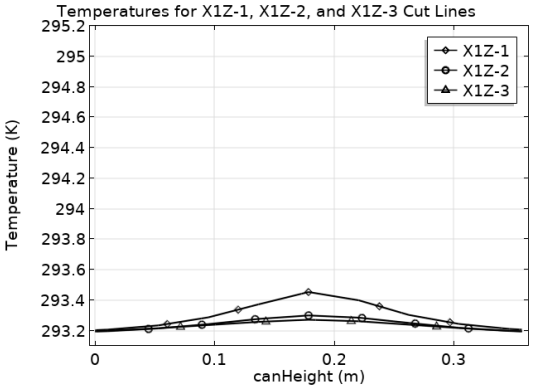

(d) X1Z cut lines

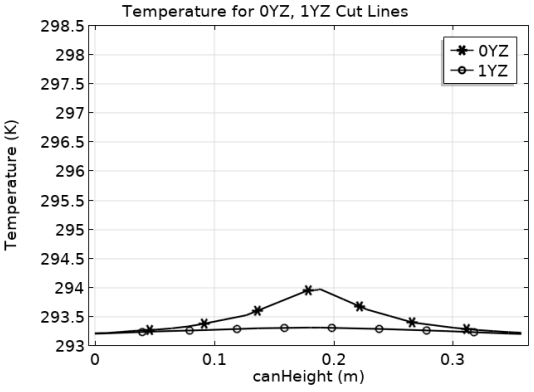

(e) 0YZ and 1YZ cut lines

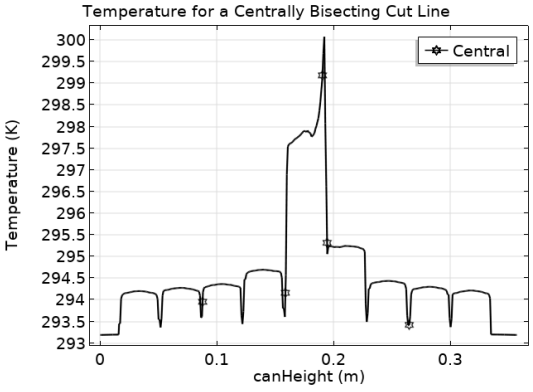

(f) A centrally bisecting cut line

Figure A.9: 20TOuter0-2W of a Set in Table 5.2

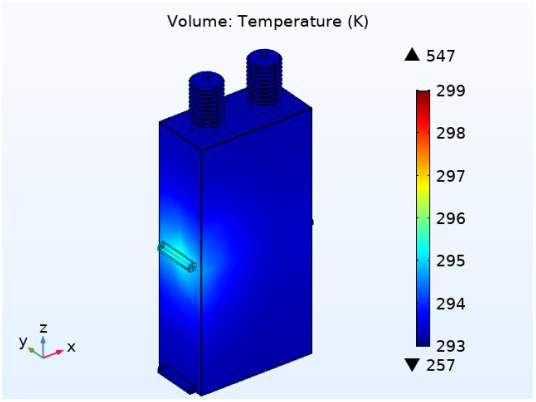

(a) Housing temperature

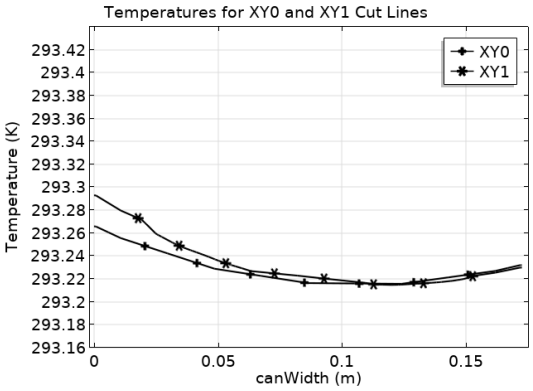

(b) XY0 and XY1 cut lines

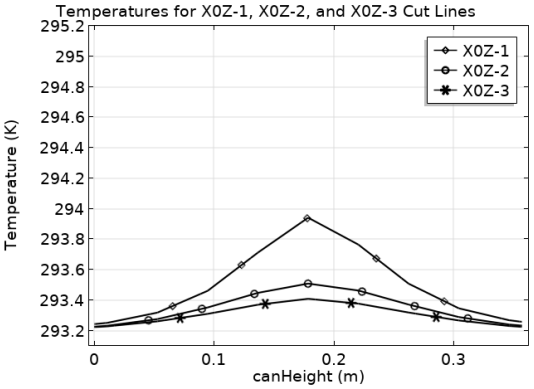

(c) X0Z cut lines

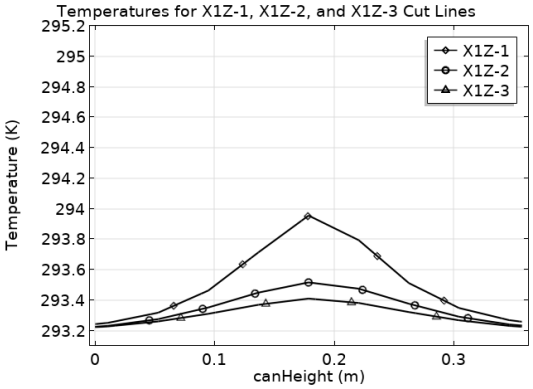

(d) X1Z cut lines

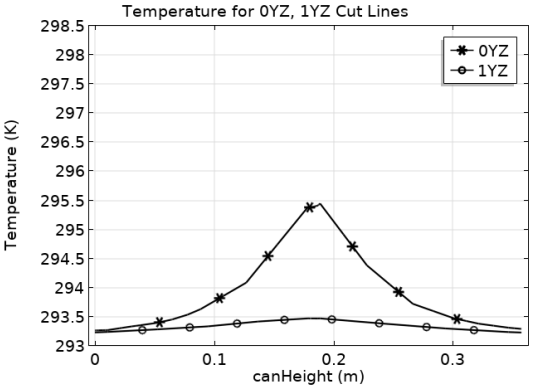

(e) 0YZ and 1YZ cut lines

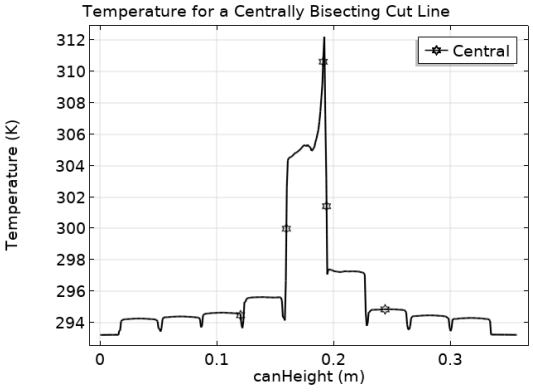

(f) A centrally bisecting cut line

Figure A.10: 20TOuter0-6W of a Set in Table 5.2

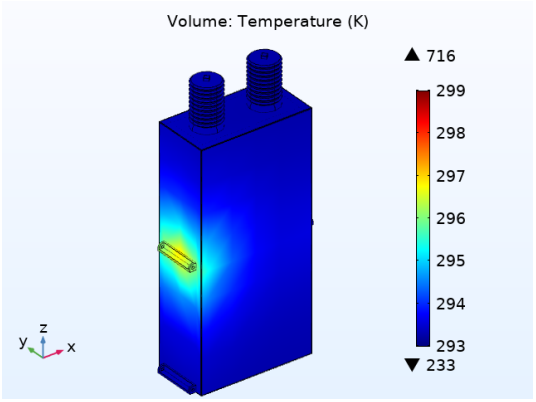

(a) Housing temperature

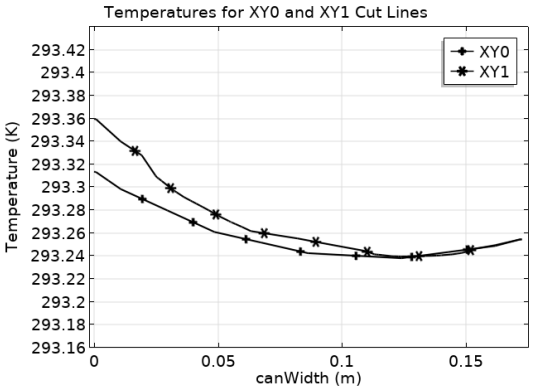

(b) XY0 and XY1 cut lines

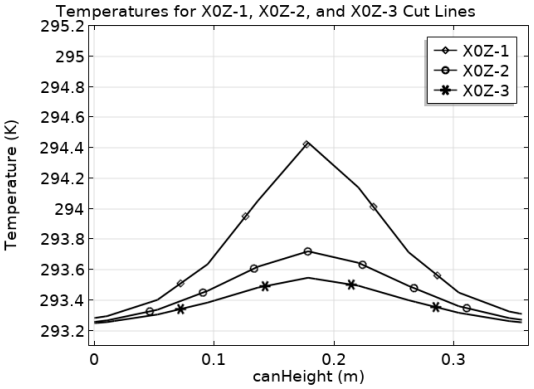

(c) X0Z cut lines

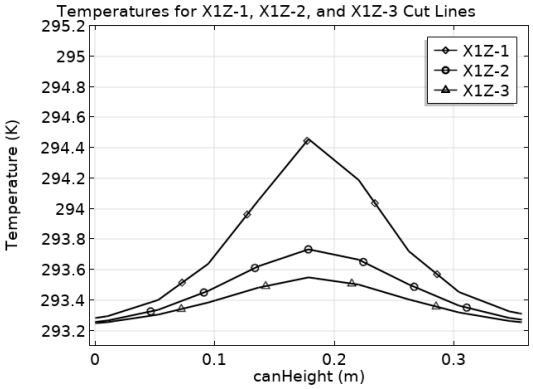

(d) X1Z cut lines

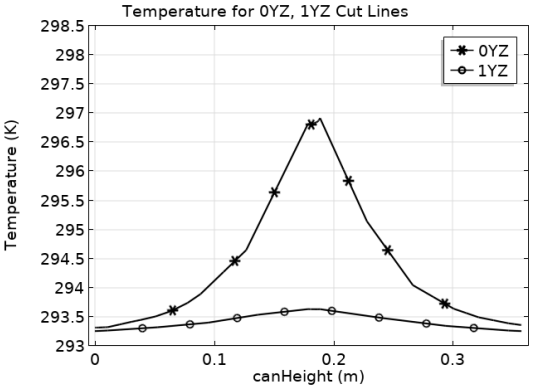

(e) 0YZ and 1YZ cut lines

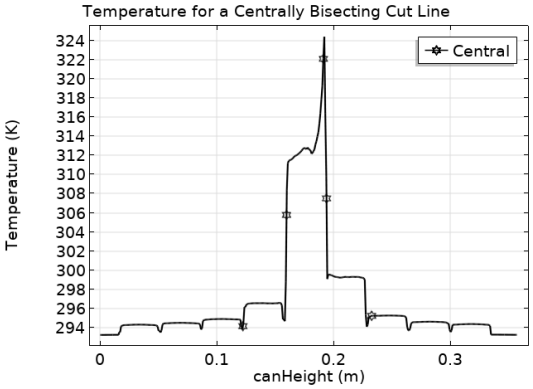

(f) A centrally bisecting cut line

Figure A.11: 20TOuter1W of a Set in Table 5.2

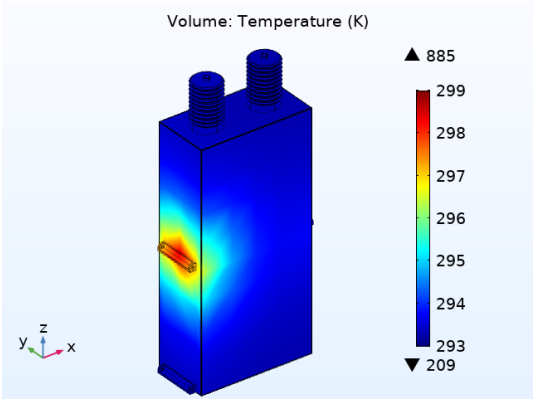

(a) Housing temperature

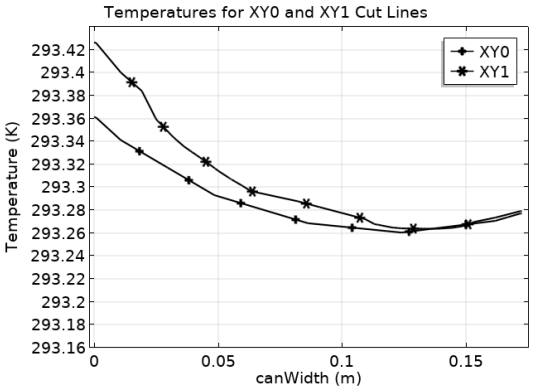

(b) XY0 and XY1 cut lines

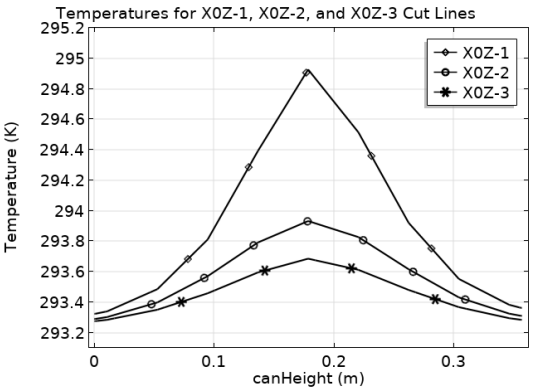

(c) X0Z cut lines

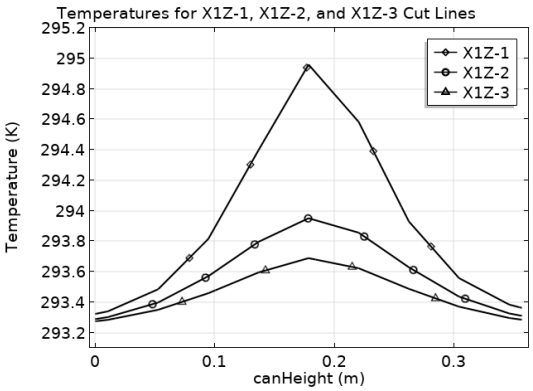

(d) X1Z cut lines

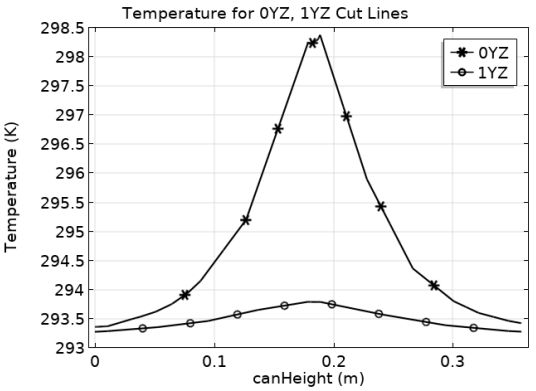

(e) 0YZ and 1YZ cut lines

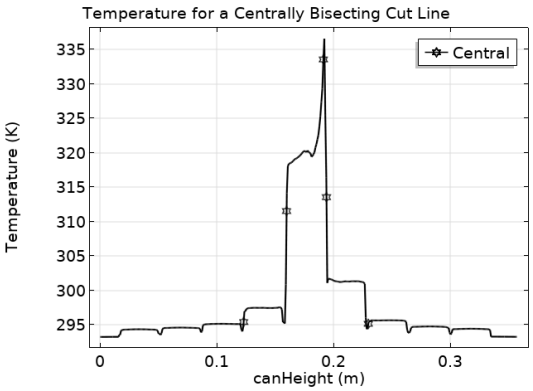

(f) A centrally bisecting cut line

Figure A.12: 20TOuter1-4W of a Set in Table 5.2

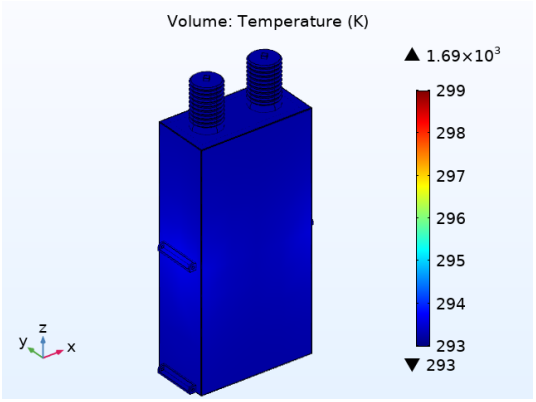

(a) Housing temperature

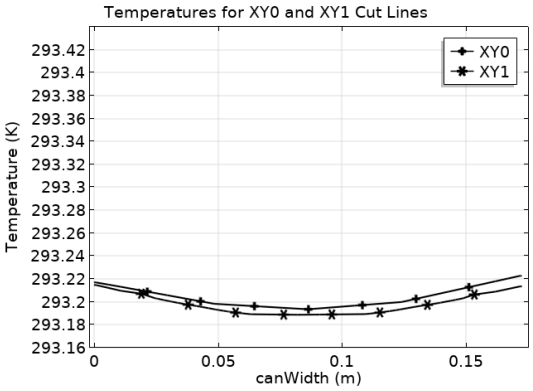

(b) XY0 and XY1 cut lines

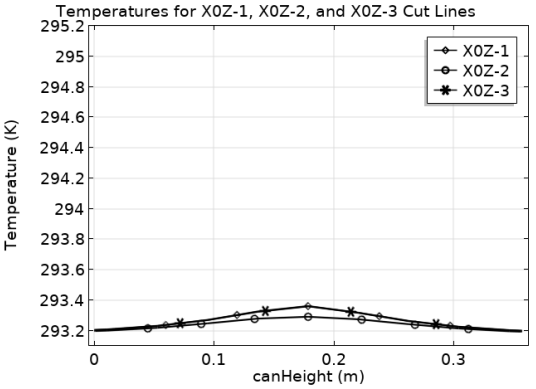

(c) X0Z cut lines

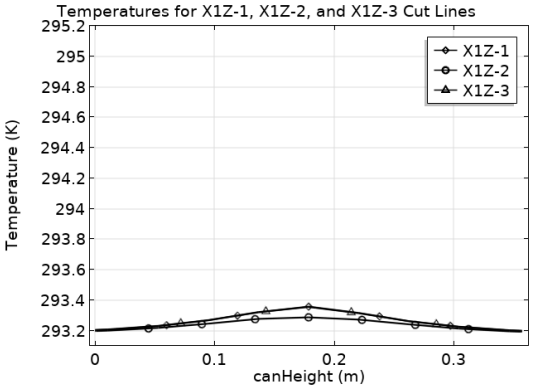

(d) X1Z cut lines

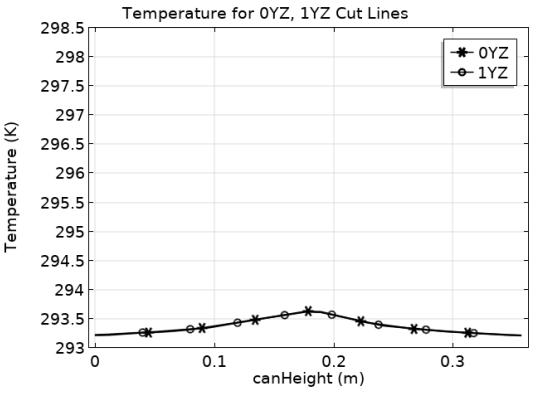

(e) 0YZ and 1YZ cut lines

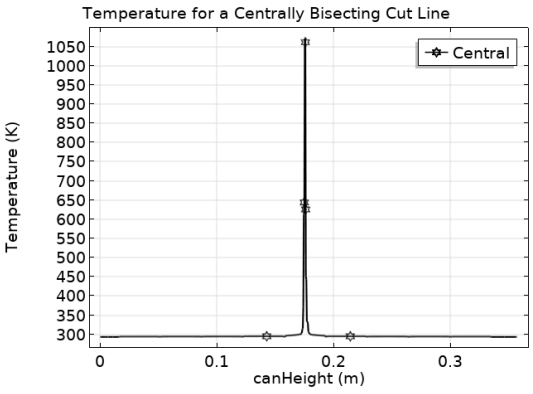

(f) A centrally bisecting cut line

Figure A.13: 40TCentre0-2W of a Set in Table 5.2

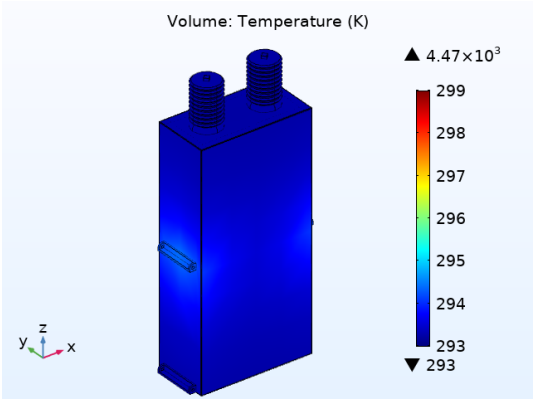

(a) Housing temperature

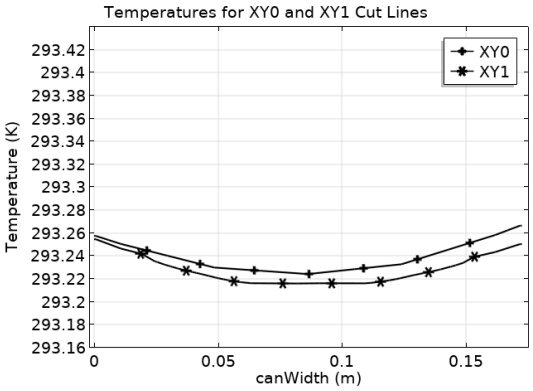

(b) XY0 and XY1 cut lines

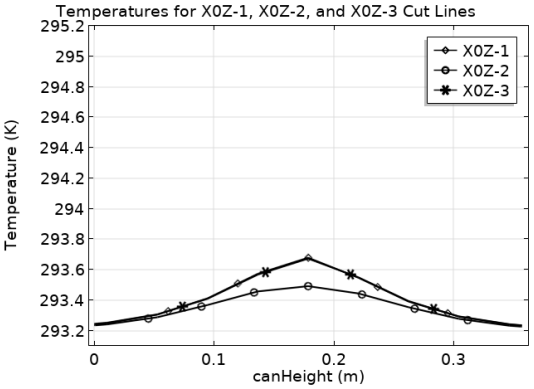

(c) X0Z cut lines

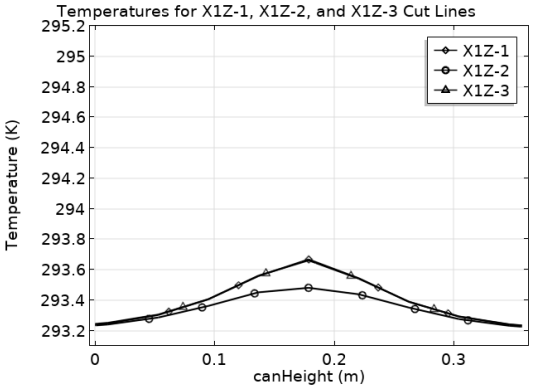

(d) X1Z cut lines

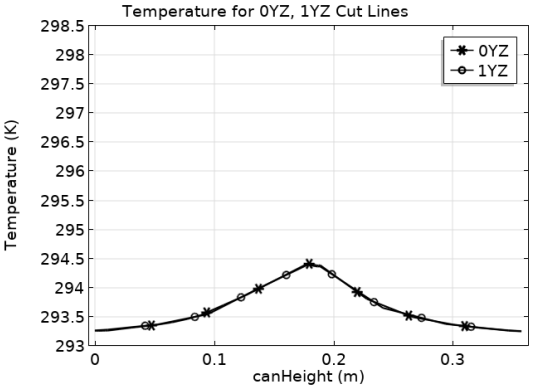

(e) 0YZ and 1YZ cut lines

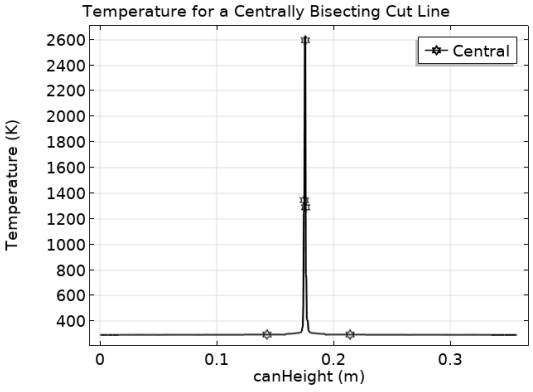

(f) A centrally bisecting cut line

Figure A.14: 40TCentre0-6W of a Set in Table 5.2

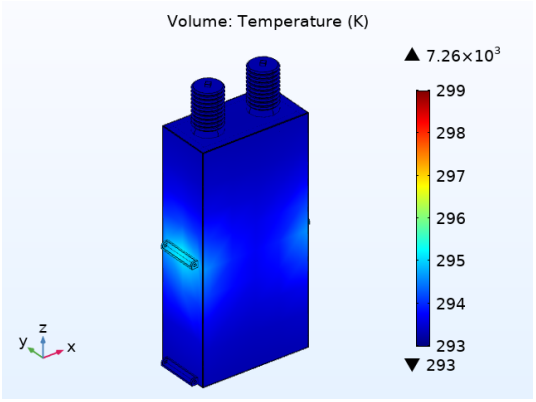

(a) Housing temperature

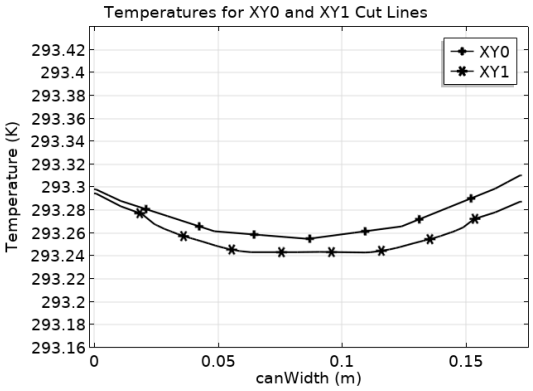

(b) XY0 and XY1 cut lines

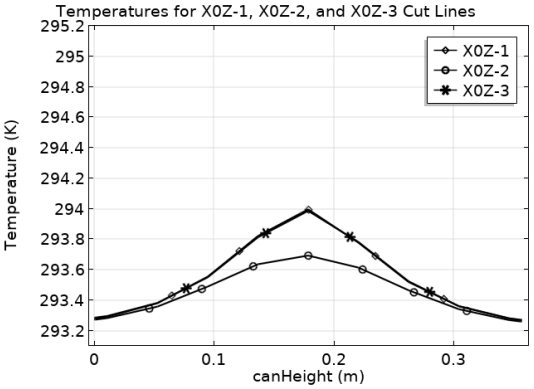

(c) X0Z cut lines

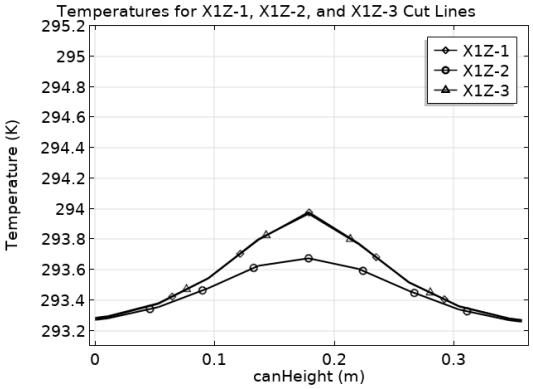

(d) X1Z cut lines

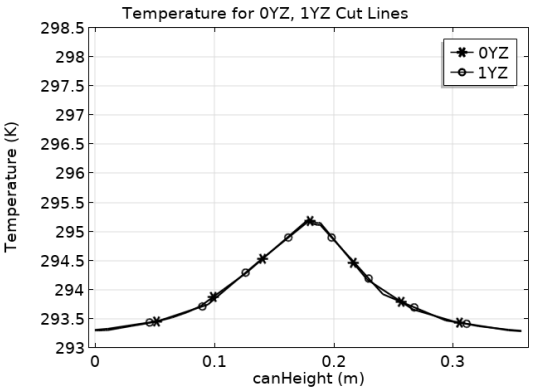

(e) 0YZ and 1YZ cut lines

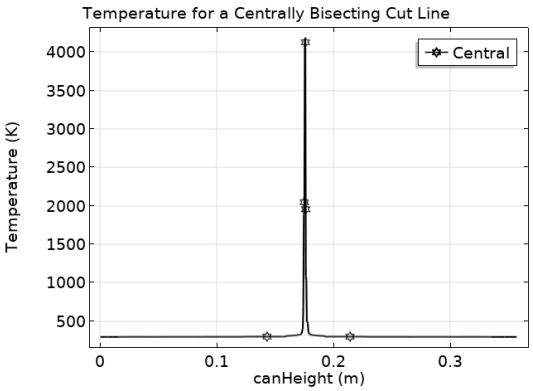

(f) A centrally bisecting cut line

Figure A.15: 40TCentre1W of a Set in Table 5.2

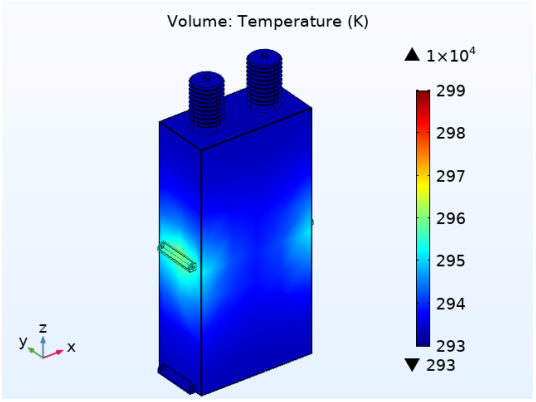

(a) Housing temperature

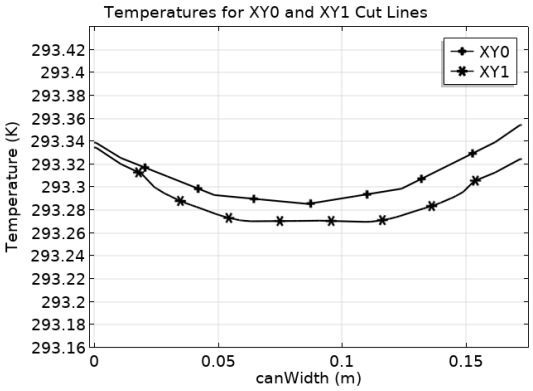

(b) XY0 and XY1 cut lines

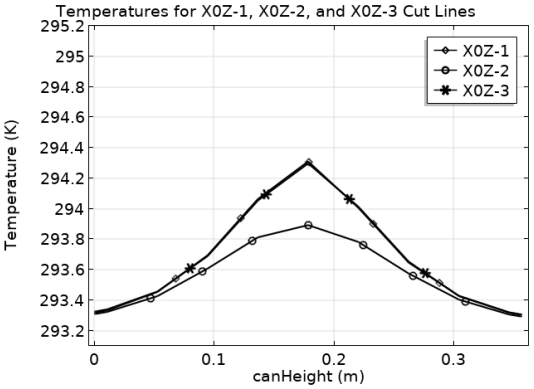

(c) X0Z cut lines

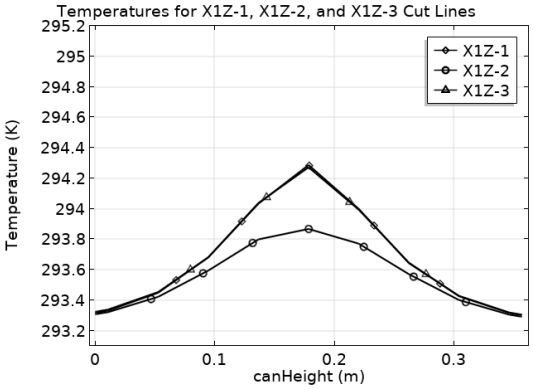

(d) X1Z cut lines

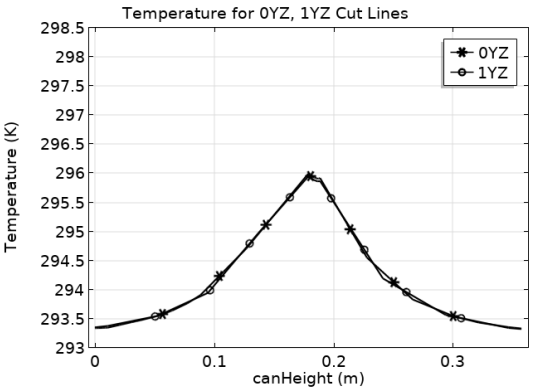

(e) 0YZ and 1YZ cut lines

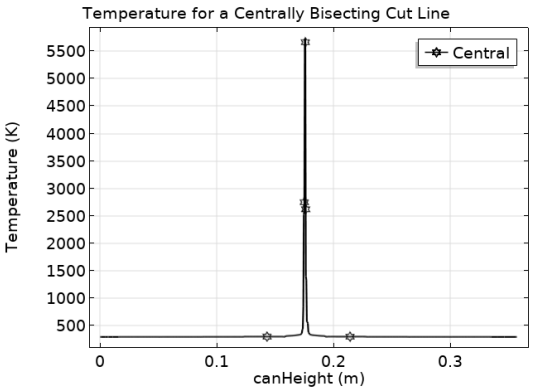

(f) A centrally bisecting cut line

Figure A.16: 40TCentre1-4W of a Set in Table 5.2

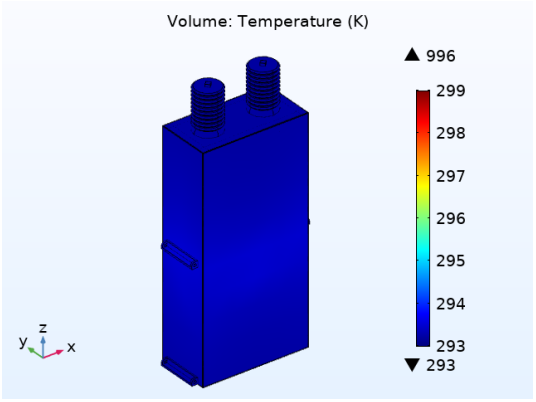

(a) Housing temperature

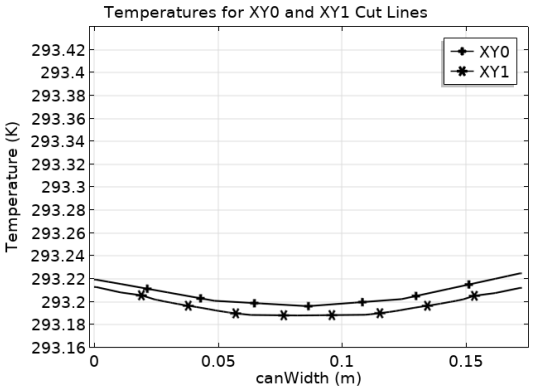

(b) XY0 and XY1 cut lines

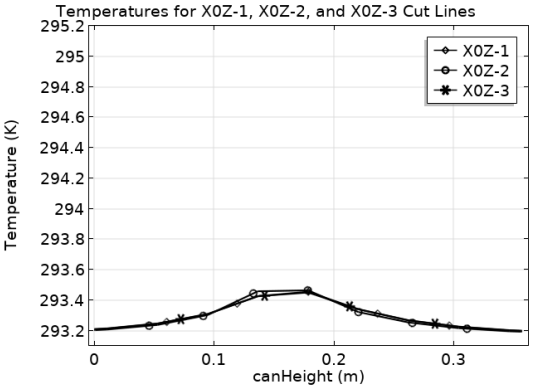

(c) X0Z cut lines

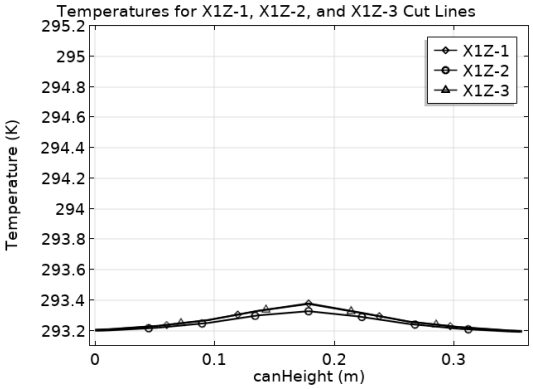

(d) X1Z cut lines

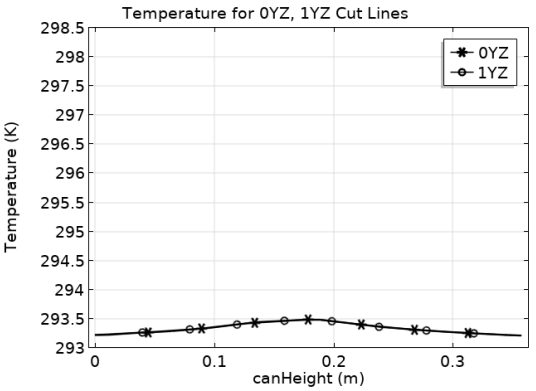

(e) 0YZ and 1YZ cut lines

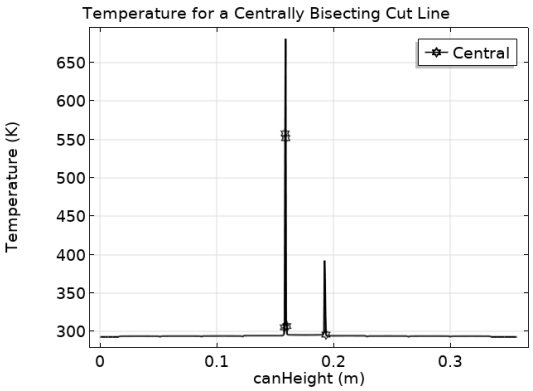

(f) A centrally bisecting cut line

Figure A.17: 40TPenultimate0-2W of a Set in Table 5.2

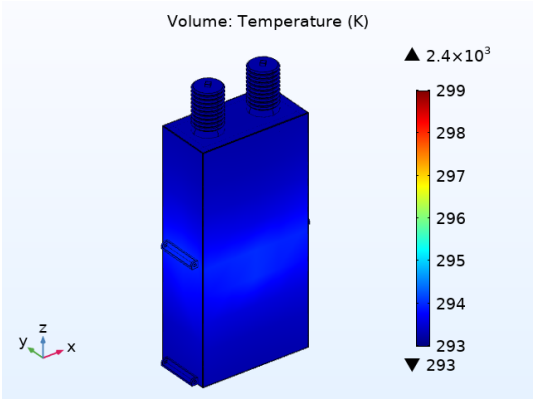

(a) Housing temperature

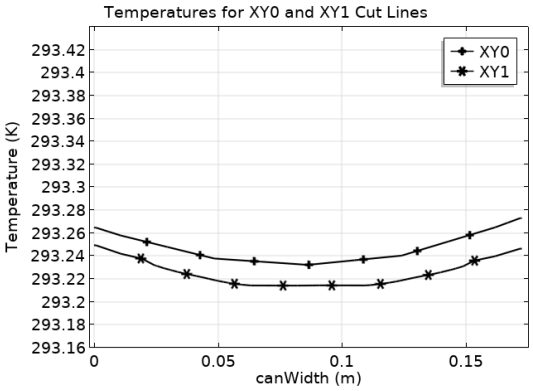

(b) XY0 and XY1 cut lines

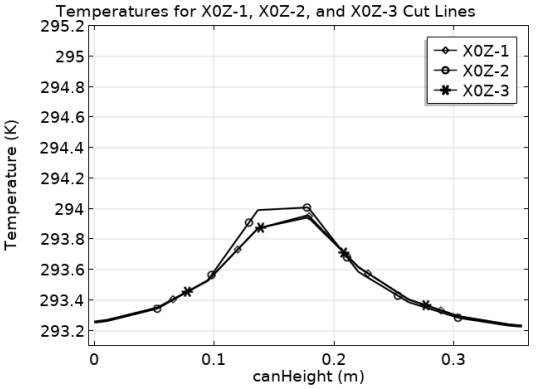

(c) X0Z cut lines

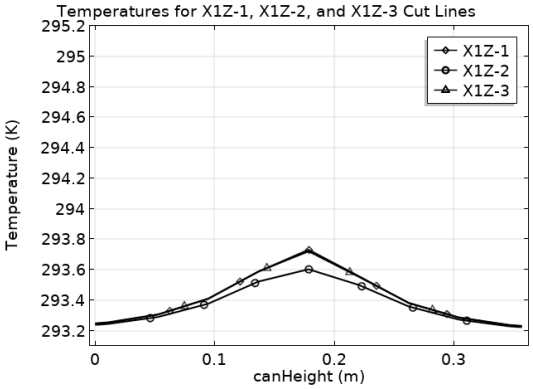

(d) X1Z cut lines

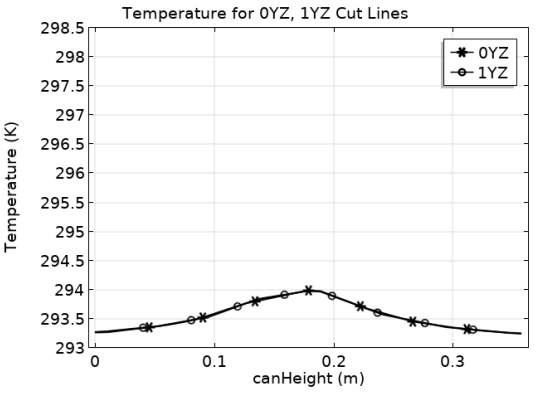

(e) 0YZ and 1YZ cut lines

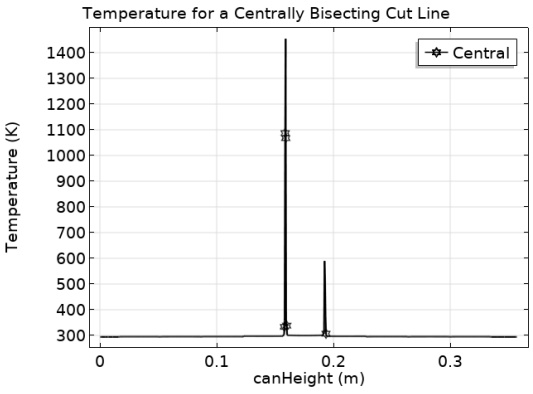

(f) A centrally bisecting cut line

Figure A.18: 40TPenultimate0-6W of a Set in Table 5.2

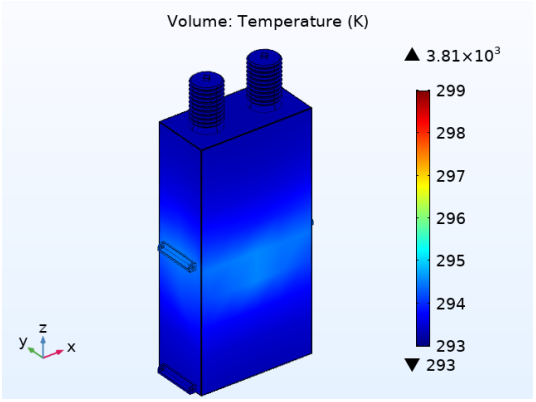

(a) Housing temperature

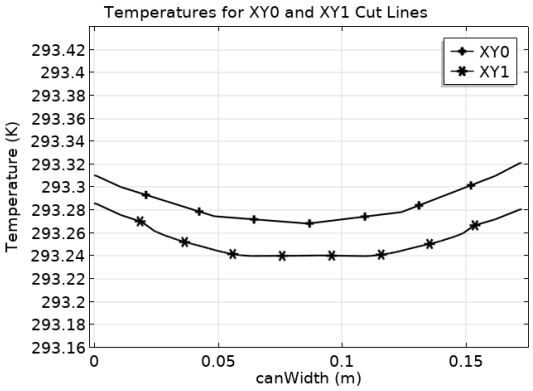

(b) XY0 and XY1 cut lines

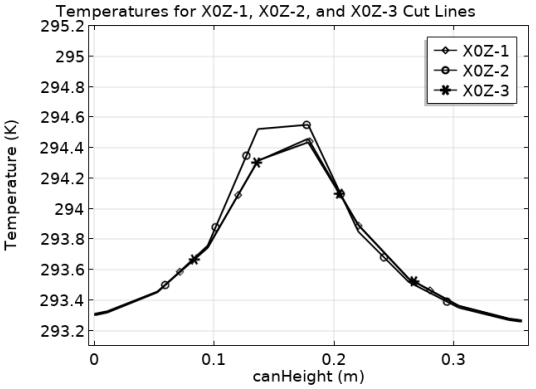

(c) X0Z cut lines

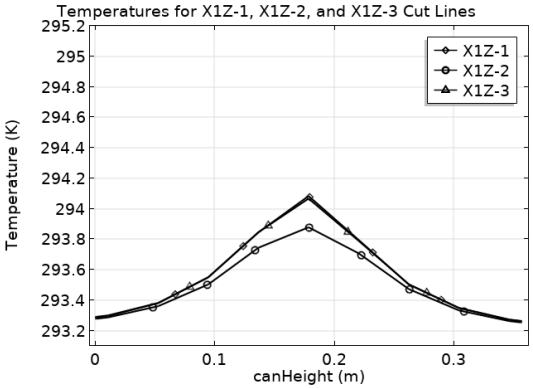

(d) X1Z cut lines

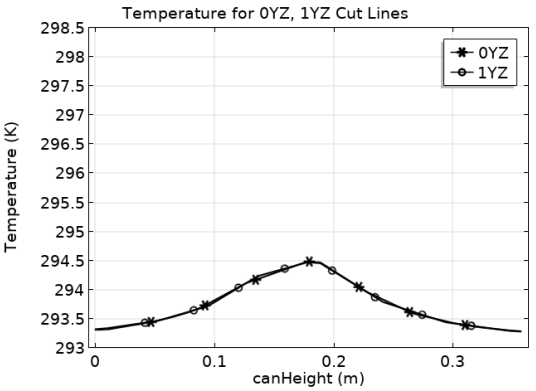

(e) 0YZ and 1YZ cut lines

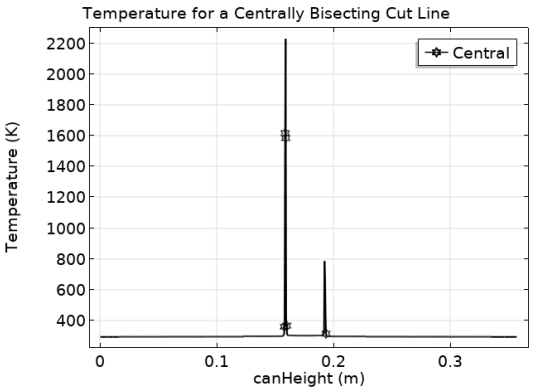

(f) A centrally bisecting cut line

Figure A.19: 40TPenultimate1W of a Set in Table 5.2

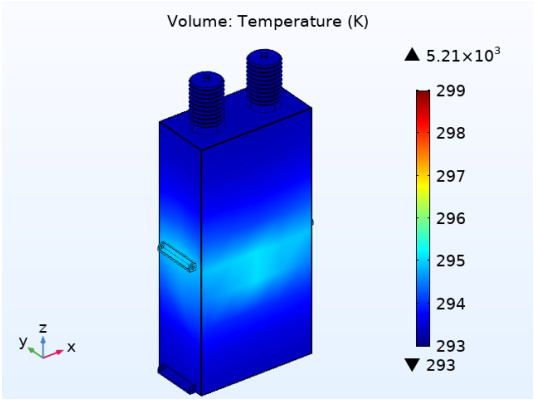

(a) Housing temperature

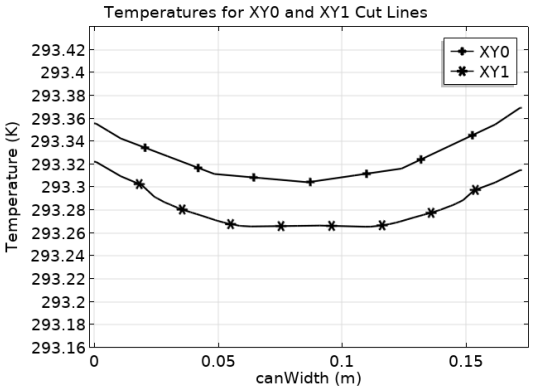

(b) XY0 and XY1 cut lines

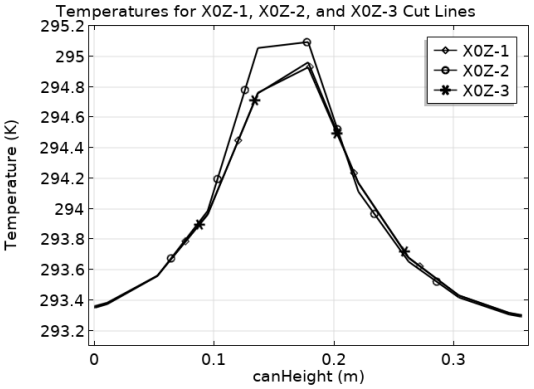

(c) X0Z cut lines

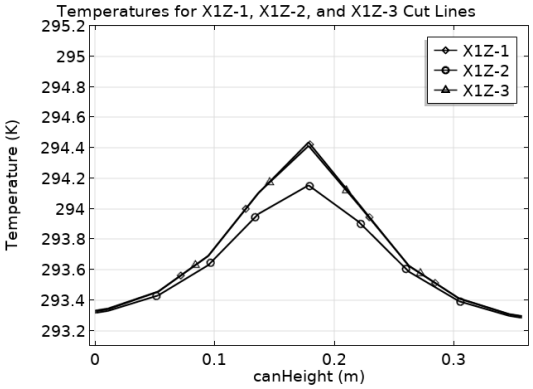

(d) X1Z cut lines

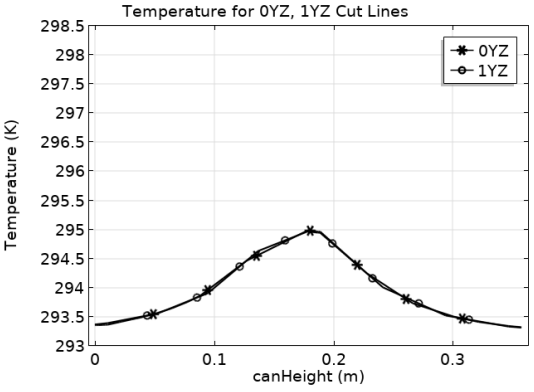

(e) 0YZ and 1YZ cut lines

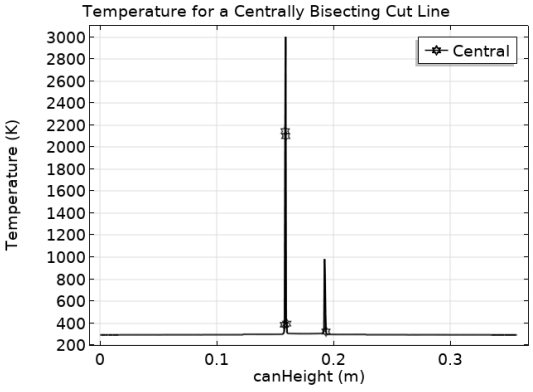

(f) A centrally bisecting cut line

Figure A.20: 40TPenultimate1-4W of a Set in Table 5.2

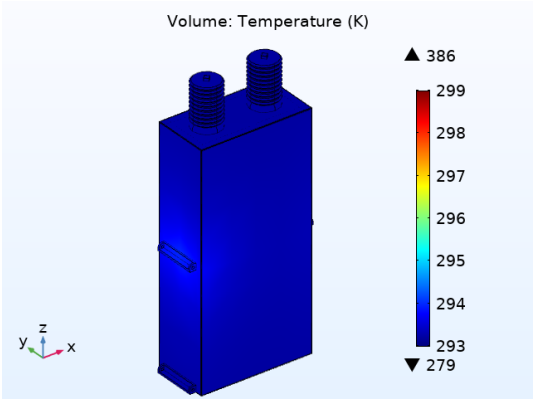

(a) Housing temperature

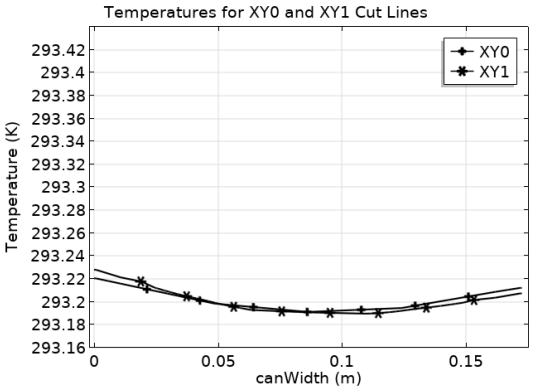

(b) XY0 and XY1 cut lines

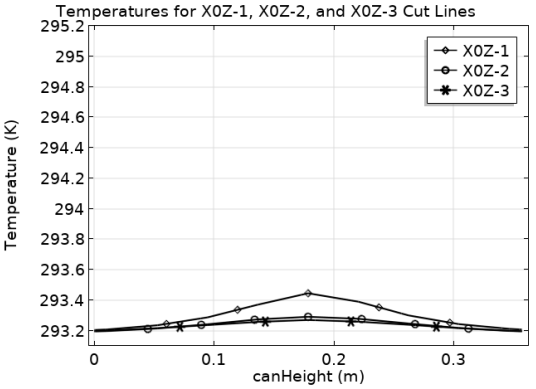

(c) X0Z cut lines

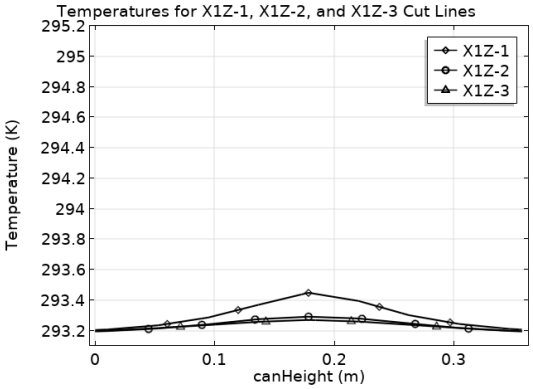

(d) X1Z cut lines

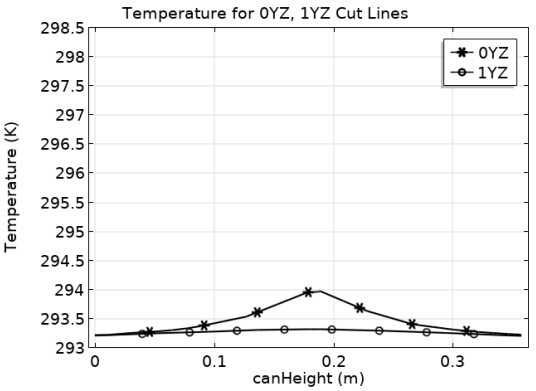

(e) 0YZ and 1YZ cut lines

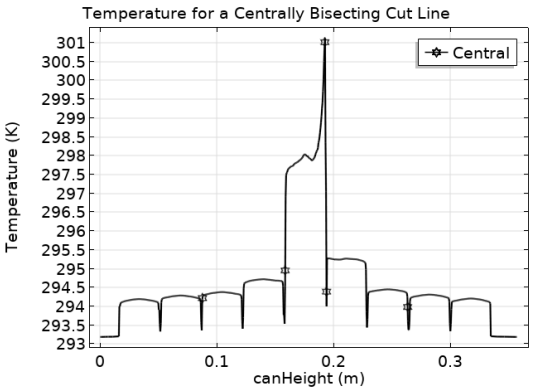

(f) A centrally bisecting cut line

Figure A.21: 40TOuter0-2W of a Set in Table 5.2

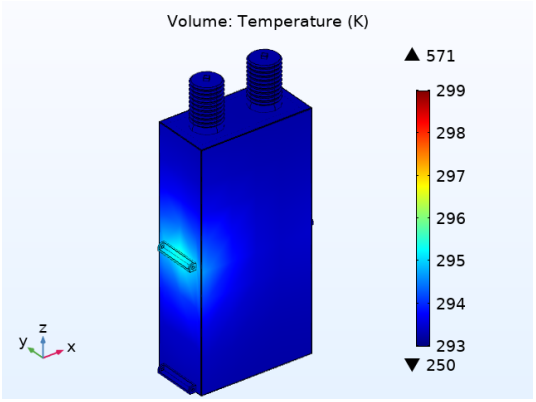

(a) Housing temperature

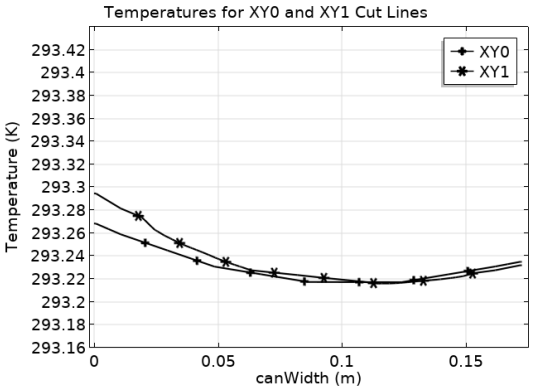

(b) XY0 and XY1 cut lines

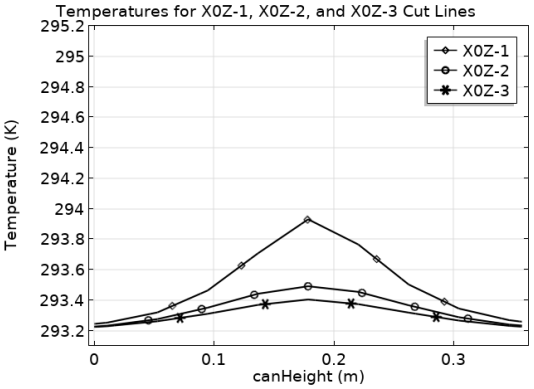

(c) X0Z cut lines

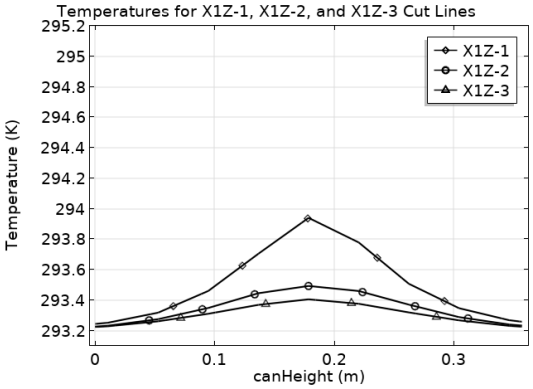

(d) X1Z cut lines

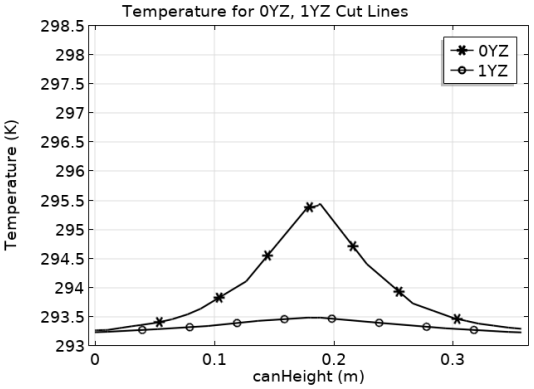

(e) 0YZ and 1YZ cut lines

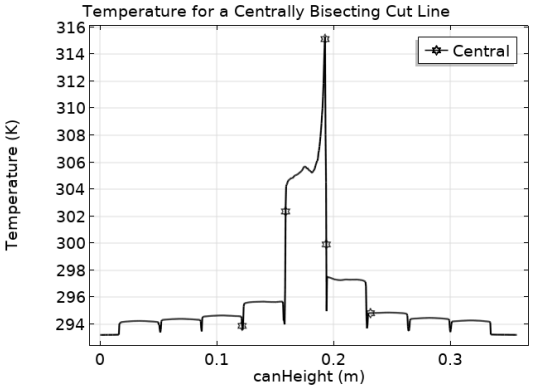

(f) A centrally bisecting cut line

Figure A.22: 40TOuter0-6W of a Set in Table 5.2

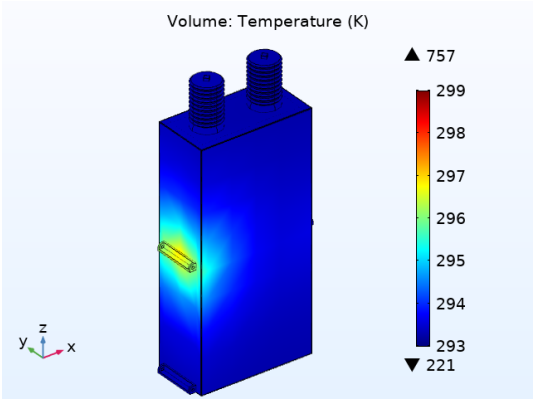

(a) Housing temperature

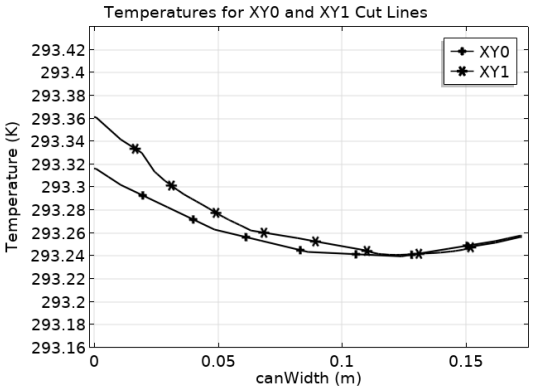

(b) XY0 and XY1 cut lines

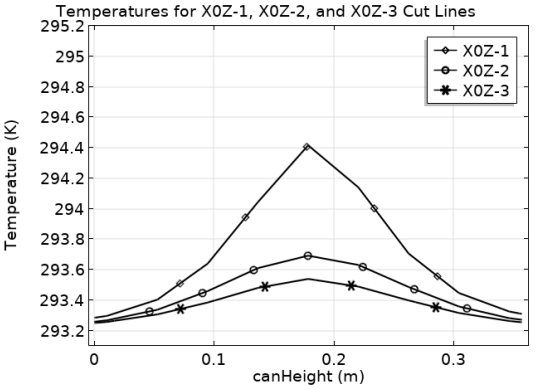

(c) X0Z cut lines

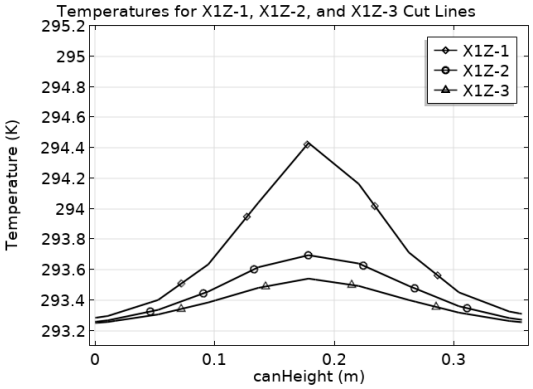

(d) X1Z cut lines

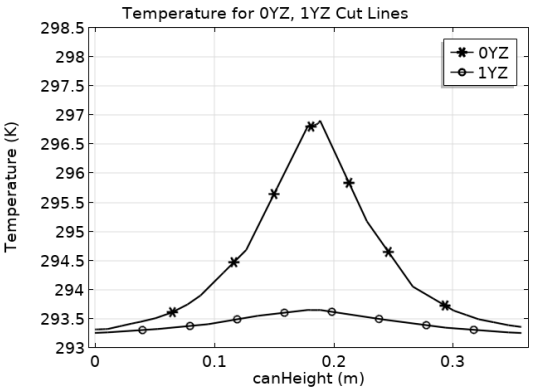

(e) 0YZ and 1YZ cut lines

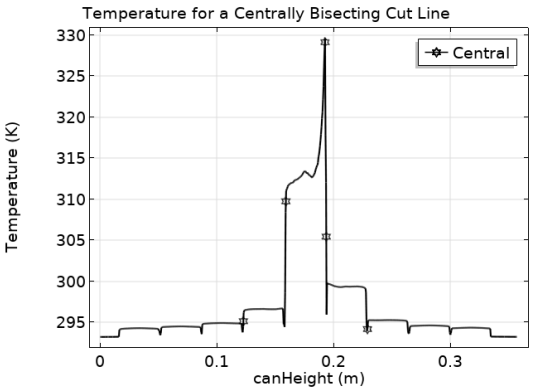

(f) A centrally bisecting cut line

Figure A.23: 40TOuter1W of a Set in Table 5.2

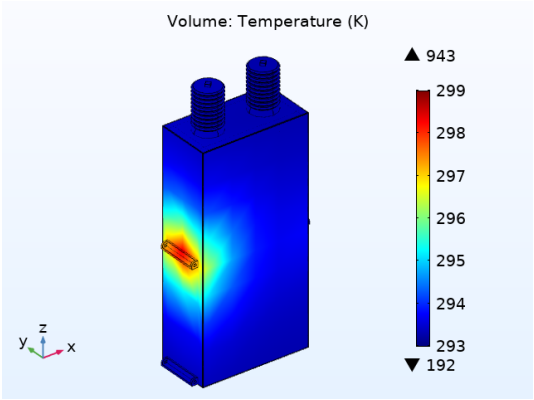

(a) Housing temperature

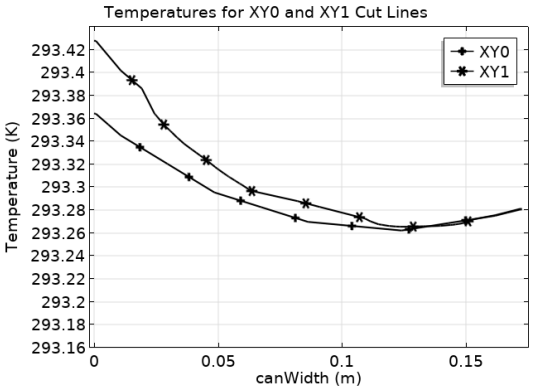

(b) XY0 and XY1 cut lines

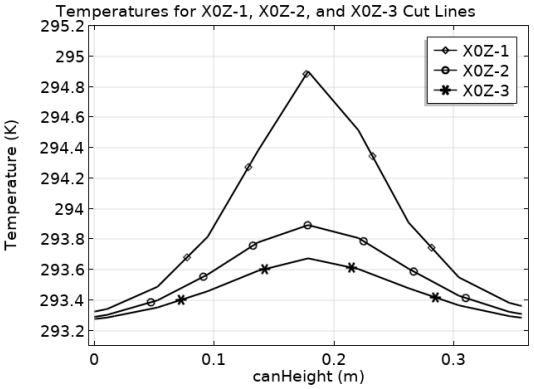

(c) X0Z cut lines

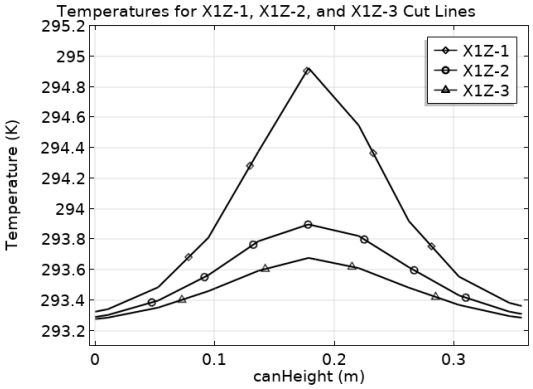

(d) X1Z cut lines

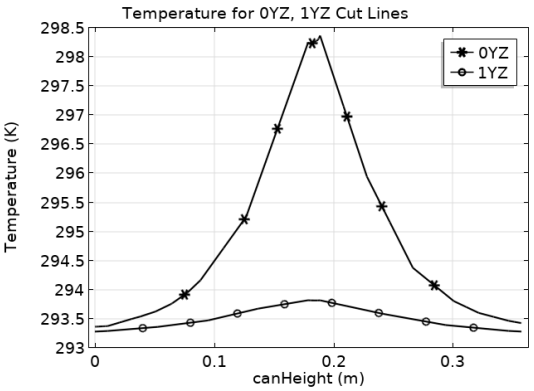

(e) 0YZ and 1YZ cut lines

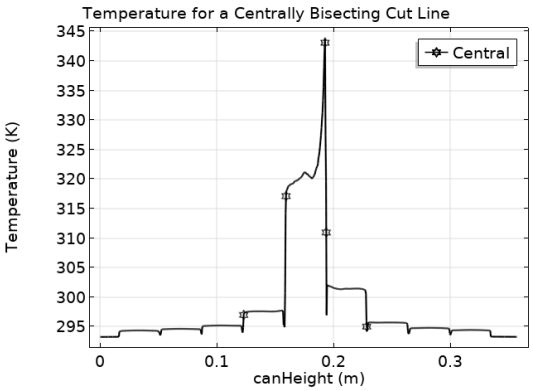

(f) A centrally bisecting cut line

Figure A.24: 40TOuter1-4W of a Set in Table 5.2

Bibliography

[1] M. Ellis, D. J. Meisner, and M. Thakur, "Innovative protection schemes for h configuration fuseless grounded shunt capacitor banks," in *Protective Relay Engineers, 2012 65th Annual Conference for*, pp. 449–458, IEEE, 2012.

[2] U. Nations, *The Sustainable Development Goals Report*. 2021.

[3] Ofgem, "Our strategy and priorities."

[4] O. Wyman, "World energy trilemma: Priority actions on climate change and how to balance the trilemma," *London, UK: World Energy Council*, 2015.

[5] J. McDonald, "Opening plenary session," All Energy Conference, May 2017. Plenary Session of All Energy, Glasgow, Scotland 10th May 2017.

[6] M. Weiser, "The computer for the 21st century," *Scientific American*, September 1991.

[7] ENTSO-E, "Completing the map: Power system needs in 2030 and 2040," tech. rep., ENTSO-E, 2020.

[8] D. Jovcic and K. Ahmed, *High voltage direct current transmission: converters, systems and DC grids*. John Wiley & Sons, 2015.

[9] A. Grid, "Hvdc-connecting to the future," *Alstom Grid*, pp. 226–228, 2010.

[10] B. M. Weedy, B. J. Cory, N. Jenkins, J. B. Ekanayake, and G. Strbac, *Electric power systems*. John Wiley & Sons, 2012.

[11] A. Lesnicar and R. Marquardt, "An innovative modular multilevel converter topology suitable for a wide power range," in *2003 IEEE Bologna Power Tech Conference Proceedings,*, vol. 3, pp. 6 pp. Vol.3–, 2003.

[12] R. Li and L. Xu, "Review of dc fault protection for hvdc grids," *Wiley Interdisciplinary Reviews: Energy and Environment*, vol. 7, no. 2, 2018.

[13] W. McDermid, A. Glodjo, and J. Bromley, "Analysis of winding failures in hvdc converter transformers," in *Electrical Insulation Conference and Electrical Manufacturing & Coil Winding Conference, 1999. Proceedings*, pp. 653–657, IEEE, 1999.

[14] W. McDermid, D. Grant, A. Glodjo, and J. Bromley, "Analysis of converter transformer failures and application of periodic on-line partial discharge measurements," in *Electrical Insulation Conference and Electrical Manufacturing & Coil Winding Conference, 2001. Proceedings*, pp. 577–582, IEEE, 2001.

[15] G. Bhuvaneswari and B. C. Mahanta, "Analysis of converter transformer failure in hvdc systems and possible solutions," 2009.

[16] F. A. M. Rizk and G. N. Trinh, *High Voltage Engineering*. CRC Press, Taylor and Francis Group, 2014.

[17] M. A. El-Hadidy and D. H. Helmi, "Failure analysis of medium voltage capacitor banks: The egyptian experience," in *Electricity Distribution-Part 1 2009. CIRED 2009. 20th International Conference and Exhibition On. 2009. IET*, IET, 2009.

[18] T. G. J. Ambrose, "Fire shuts one of uk's most important power cables in midst of supply crunch." https://www.theguardian.com/business/2021/sep/15/fire-shuts-one-of-uk-most-important-power-cables-in-midst-of-supply-crunch.

[19] T. G. Jillian Ambrose, "Firms agree scotland to england renewable energy 'superhighway'." https://www.theguardian.com/business/2020/nov/16/firms-agree-scotland-to-england-renewable-energy-superhighway.

[20] H. Jouybari-Moghaddam, T. S. Sidhu, M. R. D. Zadeh, and P. P. Parikh, "Shunt capacitor banks online monitoring using a superimposed reactance method," *IEEE Transactions on Smart Grid*, 2017.

[21] O. Macchione, "Western link hvdc." http://site.ieee.org/pes-hvdcfacts/files/2019/02/2018WesternLink.pdf, 2018.

[22] X. Wang, *Coordination of power flow control by using FACTS device and HVD transmission system*. PhD thesis, Cardiff University, 2016.

[23] P. Mohapatra, "Iet presentation." Talk at IET Glasgow, March 2018.

[24] J. A. Severin Carrell, "Scottish nuclear power station to shut down early after reactor problems." https://www.theguardian.com/uk-news/2020/aug/27/hunterston-scottish-nuclear-power-station-to-shut-down-early-after-reactor-problems, 2020.

[25] D. Ariely, "What makes us feel good about our work," *TED Talk, October*, 2012.

[26] F. MacDougall, G. McKee, J. Ennis, R. Cooper, R. Ness, and W. Reass, "High voltage capacitors designed to avoid catastrophic failure modes," in *Digest of Technical Papers. 12th IEEE International Pulsed Power Conference.(Cat. No. 99CH36358)*, vol. 1, pp. 97–101, IEEE, 1999.

[27] L. Xu, "Vsc transmission system using flying capacitor multilevel converters and hybrid pwm control," *Power Delivery, IEEE Transactions on*, vol. 22, no. 1, pp. 693–702, 2007.

[28] J. Matias, "Reactive power compensation," June 2013.

[29] W. J. Sarjeant, J. Zirnheld, and F. W. MacDougall, "Capacitors," *IEEE Transactions on Plasma Science*, vol. 26, October 1998.

[30] A. F. Bastos and S. Santoso, "Condition monitoring of circuit switchers for shunt capacitor banks through power quality data," *IEEE Transactions on Power Delivery*, vol. 34, no. 4, pp. 1499–1507, 2019.

[31] V. ESTA, "High voltage power capacitors," tech. rep., March 2018.

[32] "High voltage ac power capacitors," datasheet, Vishay ESTA, May 2018.

[33] L. B. Gordon and L. Cartelli, "A complete electrical hazard classification system and its application," in *2009 IEEE IAS Electrical Safety Workshop*, pp. 1–12, 2009.

[34] A. T. Carvalho, H. P. Amorim, C. F. Cunha, L. S. Ferreira, M. C. Cagido, and J. B. Borges, "Virtual instrumentation for high voltage power capacitors assessment through capacitance monitoring and acoustic partial discharge detection," in *IEEE Electrical Insulation Conference (EIC)*, pp. 511–514, IEEE, 2016.

[35] "Analysis techniques for system reliability: Procedure for failure mode and effects analysis (fmea)," Tech. Rep. IEC 60812:2006, Bureau Central de la Commission Electrotechnique Internationale, Genève, Suisse, 1985.

[36] "M4100: High voltage apparatus tester."

[37] "M4100 4^{th} gen high voltage apparatus tester," tech. rep., Doble Engineering Company, 2019.

[38] S. Samineni, C. Labuschagne, and J. Pope, "Principles of shunt capacitor bank application and protection," in *Protective Relay Engineers, 2010 63rd Annual Conference for*, pp. 1–14, IEEE, 2010.

[39] H. Jouybari-Moghaddam, T. Sidhu, I. Voloh, and M. Zadeh, "New method of capacitors failure detection and location in shunt capacitor banks," in *2018 71st Annual Conference for Protective Relay Engineers (CPRE)*, pp. 1–10, IEEE, 2018.

[40] P. S. Pavan and S. Das, "Novel method for location of internal faults in ungrounded double wye shunt capacitor banks," *IEEE Transactions on Power Delivery*, vol. 36, pp. 899–908, April 2021.

[41] "Iec 61165, application of markov techniques," Tech. Rep. IEC 61165, International Electrotechnical Commission, 2006.

[42] B. EN, "61025: 2007 fault tree analysis (fta)," *London: British Standards Institute*, 2007.

[43] C. Reed and S. Cichanowskil, "The fundamentals of aging in hv polymer-film capacitors," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 1, no. 5, pp. 904–922, 1994.

[44] J. Ennis, F. MacDougall, R. Cooper, and J. Bates, "Repetitive pulse application of self-healing high voltage capacitors," in *Conference Record of the Twenty-Fifth International Power Modulator Symposium, 2002 and 2002 High-Voltage Workshop.*, pp. 634–638, IEEE, 2002.

[45] I. S. Association, "C37:99-2012 ieee guide for the protection of shunt capacitor banks," Standard C37:99-2012, IEEE Power and Energy Society, 2013.

[46] N. Clark and S. Farnham, "Connection and protection of capacitor banks," *Electrical Engineering*, vol. 69, no. 3, pp. 230–230, 1950.

[47] J. Schaefer, S. Samineni, C. Labuschagne, S. Chase, and D. J. Hawaz, "Minimizing capacitor bank outage time through fault location," in *2014 67th Annual Conference for Protective Relay Engineers*, pp. 72–83, IEEE, 2014.

[48] T. M. Blooming and D. J. Carnovale, "Capacitor application issues," *Transactions on Industry Applications*, vol. 44, pp. 178–190, August 2008.

[49] H. Jouybari-Moghaddam, T. S. Sidhu, M. R. D. Zadeh, and P. P. Parikh, "Enhanced fault-location scheme for double wye shunt capacitor banks," *IEEE Transactions on Power Delivery*, vol. 32, no. 4, pp. 1872–1880, 2017.

[50] Z. Li, H. Li, X. Huang, H. Li, W. Wang, B. Wang, F. Lin, and Q. Zhang, "Temperature rise of metallized film capacitors in repetitive pulse applications," *IEEE Transactions on Plasma Science*, vol. 43, no. 6, pp. 2038–2045, 2015.

[51] S. Qin, J. Ho, M. Rabuffi, G. Borelli, and T. Jow, "Implications of the anisotropic thermal conductivity of capacitor windings," *IEEE Electrical Insulation Magazine*, vol. 27, no. 1, pp. 7–13, 2011.

[52] A. Hjert, "Multiscale modelling of a metallized film capacitor for hvdc applications," Master's thesis, 2017.

[53] E. M. Vaciloto Filho, A. L. O. e. de Silva, R. C. Campos, L. E. B. da Silva, L. S. e. Silva, and R. Rossi, "Methodology for mitigation of interferences in the protection system of capacitor banks in medium and high voltage," in *2018 IEEE PES Transmission Distribution Conference and Exhibition - Latin America (T D-LA)*, pp. 1–5, Sep. 2018.

[54] R. Arora and W. Mosch, *High voltage and electrical insulation engineering*, vol. 69. John Wiley & Sons, 2011.

[55] G. Jancke, N. Fahlen, and O. Nerf, "Series capacitors in power systems," *IEEE Transactions on Power Apparatus and Systems*, vol. 94, no. 3, pp. 915–925, 1975.

[56] Siemens, "Power capacitors and capacitor banks: Selected projects." Online.

[57] Meher, "Medium and high voltage capacitors, capacitor banks and systems," 2010.

[58] ABB, "Abb power capacitors and harmonic filters buyer's guide." Online, 2013.

[59] ABB, "Drydcap: Enhancing eco-efficiency." Online, 2014.

[60] K. A. O'Connor and R. D. Curry, "Recent results in the development of composites for high energy density capacitors," in *2014 IEEE International Power Modulator and High Voltage Conference (IPMHVC)*, pp. 496–499, IEEE, 2014.

[61] R. Curry, S. Dickerson, A. Howard, B. Lamb, and S. Mounter, "Scaled dc lifetime, test and evaluation of advanced nanocomposite materials for compact high voltage capacitors," in *2017 IEEE 21st International Conference on Pulsed Power (PPC)*, pp. 1–5, IEEE, 2017.

[62] Q. Li and S. Cheng, "Polymer nanocomposites for high-energy-density capacitor dielectrics: Fundamentals and recent progress," *IEEE Electrical Insulation Magazine*, vol. 36, pp. 7–28, March 2020.

[63] C. Zhang, S. Cao, and X. Wang, "Research on power capacitor insulation state detection based on polarization depolarization current method," in *2019 IEEE Innovative Smart Grid Technologies-Asia (ISGT Asia)*, pp. 244–249, IEEE, 2019.

[64] C. Wu, "Thermal modelling of high voltage capacitors," Master's thesis, University of Strathclyde, 2018.

[65] T. G. Aakre, A. V. Torbjørn, and Ø. L. Hestad, "Conductivity and permittivity of midel 7131: effect of temperature, moisture content, hydrostatic pressure and electric field," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 23, no. 5, pp. 2957–2964, 2016.

[66] C. Rappe, S. Marklund, L.-O. Kjeller, P.-A. Bergqvist, and M. Hansson, "Strategies and techniques for sample collection and analysis: experience from the swedish pcb accidents.," *Environmental health perspectives*, vol. 60, pp. 279–292, 1985.

[67] Y. Liang, D. Wang, and D. Liu, "Distribution and influence factors of electric field in high-voltage power capacitor," *IET Science, Measurement & Technology*, vol. 9, no. 6, pp. 702–708, 2015.

[68] B. S. Institution, "62001-4 high-voltage direct current (hvdc) systems - guidance to the specification and design evaluation of ac filters: Equipment," 2016.

[69] W. J. Lee, K. Narayanan, T. Maffetone, and P. Didsayabutra, "The design of a capacitor bank early warning system," *IEEE Transactions on Industry Applications*, vol. 39, no. 2, pp. 306–312, 2003.

[70] R. Moxley, J. Pope, and J. Allen, "Capacitor bank protection for simple and complex configurations," in *2012 65th Annual Conference for Protective Relay Engineers*, pp. 436–441, IEEE, 2012.

[71] R. C. de Lima Furnas and P. Lindberg, "Operational experience of 800 kv series capacitors," in *2005 IEEE Power Engineering Society Inaugural Conference and Exposition in Africa*, pp. 399–402, IEEE, 2005.

[72] P. H. Thiel, J. E. Harder, and G. E. Taylor, "Fuseless capacitor banks," *IEEE transactions on power delivery*, vol. 7, no. 2, pp. 1009–1015, 1992.

[73] J. Wang, M. Ibrahim, Z. Gajić, and M. Saha, "Internal failure detection and protection on capacitor banks," 2016.

[74] S. P. Polisetty and S. Das, "Novel method for location of internal faults in ungrounded double wye shunt capacitor banks," *IEEE Transactions on Power Delivery*, 2020.

[75] L. Fendrick, T. Day, K. Fender, J. McCall, and A. Chaudhary, "Complete relay protection of multi-string fuseless capacitor banks," in *Conference Record of the 2002 Annual Pulp and Paper Industry Technical Conference (Cat. No.02CH37352)*, pp. 194–198, June 2002.

[76] H. Jouybari-Moghaddam, T. Sidhu, P. Parikh, and I. Voloh, "Enhanced fault location method for shunt capacitor banks," in *2017 70th Annual Conference for Protective Relay Engineers (CPRE)*, pp. 1–11, IEEE, 2017.

[77] H. Jouybari-Moghaddam and T. S. Sidhu, "A study of capacitor element failures in high voltage shunt capacitor banks," in *2017 IEEE 30th Canadian Conference on Electrical and Computer Engineering (CCECE)*, pp. 1–4, IEEE, 2017.

[78] S. R. McCormick, K. Hur, S. Santoso, A. Maitra, and A. Sundaram, "Capacitor bank predictive maintenance and problem identification using conventional power quality monitoring systems," in *IEEE Power Engineering Society General Meeting, 2004.*, pp. 1846–1850 Vol.2, June 2004.

[79] A. F. Bastos and S. Santoso, "Identifying switched capacitor relative locations and energizing operations," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*, pp. 1–5, IEEE, 2016.

[80] "How to set up and run a simulation with comsol multiphysics," 2015.

[81] A. Abdelmalik, A. Nysveen, and L. Lundgaard, "Partial discharges in liquid embedded power electronics: Effects of pressure and liquid nature under negative pulse voltage stress," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 23, no. 2, pp. 1119–1125, 2016.

[82] G. G. Solutions, "Ge high voltage capacitor and equipment product application and selection guide," tech. rep., 2018.

[83] COMSOL Multiphysics, *AC/DC Module User's Guide*, 2019.

[84] COMSOL Multiphysics, *Heat Transfer Module User Guide*, 2018.

[85] C. J. Mackinnon and B. G. Stewart, "Regions of electrical stress in high voltage capacitor units," in *COMSOL Europe*, COMSOL, September 2019.

[86] M. G. Kong and Y.-P. Lee, "Electrically induced heat dissipation in metallized film capacitors," *IEEE transactions on dielectrics and electrical insulation*, vol. 11, no. 6, pp. 1007–1013, 2004.

[87] W. Sarjeant, F. MacDougall, D. Larson, and I. Kohlberg, "Energy storage capacitors: aging, and diagnostic approaches for life validation," *IEEE Transactions on Magnetics*, vol. 33, no. 1, pp. 501–506, 1997.

[88] S.-H. Lee, B.-Y. Lee, H.-K. Kim, and H.-G. Kim, "Local heat source approximation technique for predicting temperature rise in power capacitors," *IEEE transactions on magnetics*, vol. 45, no. 3, pp. 1250–1253, 2009.

[89] "Ds1920 temperature ibutton," tech. rep., Maxim Integrated Products, Inc., 2009.

[90] L. A. Morocho, L. B. Peralta, L. F. Guerrero-Vásquez, J. O. Ordoñez-Ordoñez, J. P. Bermeo, and P. A. Chasi-Pesantez, "Dielectric permittivity measurement of kraft paper by the resonant ring method," in *Advances in Emerging Trends and Technologies*, Advances in Intelligent Systems and Computing, pp. 271–279, Cham: Springer International Publishing, 2019.

[91] J. Ambrose, "Energy firms fined after subsea cable delay pauses scottish windfarms." https://www.theguardian.com/business/2020/nov/16/firms-agree-scotland-to-england-renewable-energy-superhighway, November 2021.

[92] J. Ambrose, "Energy firms face inquiry over £1.3bn green power cable." https://www.theguardian.com/business/2020/jan/28/energy-firms-power-cable-national-grid-scottish-power.

[93] COMSOL Multiphysics, *COMSOL Multiphysics Reference Manual*, 2019.

[94] C. Mackinnon, "Email correspondence," 2020.

[95] H. Burns, "What causes wellness, tedx glasgow." Online, accessed at https://www.youtube.com/watch?v=yEh3JG74C6s, 2014.

[96] S. A. Boggs, J. Ho, and T. R. Jow, "Overview of laminar dielectric capacitors," *IEEE Electrical Insulation Magazine*, vol. 26, no. 2, pp. 7–13, 2010.

[97] C. J. Mackinnon, "Hvdc condition monitoring first year review," tech. rep., University of Strathclyde, 2017.

[98] S. Blair, "Email correspondence," 2020.

Index