Characterisation of Arresters for Harmonic Overvoltage Studies

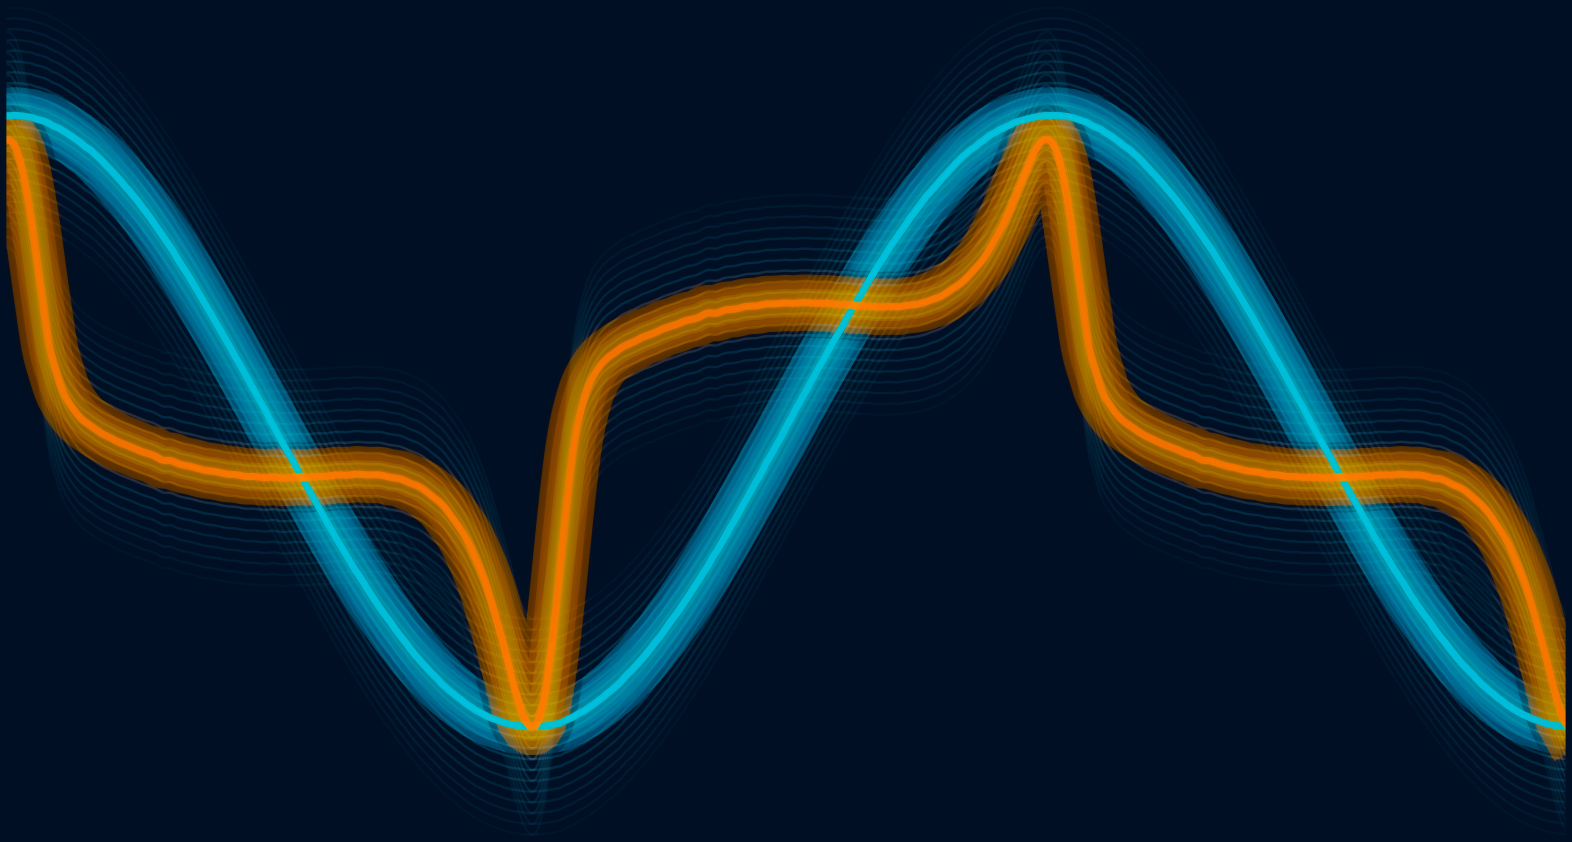

Pavan Dhulipala

TU Delft

Characterisation of Arresters for Harmonic Over-voltage Studies

Evaluating Surge Arrester TOV Withstand Characteristics in Transients

Master Thesis

by

Pavan Dhulipala

TU Delft

University:	Delft University of Technology, Delft
Study Programme:	MSc Sustainable Energy Technology
Thesis Committee:	Prof. Dr. Ir. Marjan Popov
	Dr. Ir. Fabio Muñoz Muñoz
	Mr. Christiaan S. Engelbrecht
	Dr. Ir. Dennis van der Born
Date:	27 August 2025

Prof. Dr. H. K. Høidalen of NTNU-Norway and Ir. K. Velitsikakis of TenneT TSO provided invaluable discussions and insights that significantly contributed to this thesis.

Printed by:	Ridderprint, The Netherlands
Cover:	Non-linear behaviour of a surge arrester MOV as it hovers around its conduction region — by Pavan
ISBN/EAN:	978-94-6384-822-0

An electronic copy of this thesis is available at
https://repository.tudelft.nl/

न जायते म्रियते वा कदाचिन्
नायं भूत्वा भविता वा न भूयः।
अजो नित्यः शाश्वतोऽयं पुराणो
न हन्यते हन्यमाने शरीरे॥२.२०॥

”It is never born nor dies, nor having once existed,
does it ever cease to be. It is unborn, eternal,
permanent, and primeval. It is not destroyed
when the body is destroyed.”

— *Bhagavad Gītā*, 2 · 20

Abstract

Metal oxide surge arresters constitute the primary overvoltage protection in power systems, yet their behaviour under harmonic-rich conditions remains inadequately understood. Modern grids face unprecedented challenges from three converging trends: extensive cable networks that shift resonant frequencies, renewable energy integration with inverter-based generation, and massive data centre loads with power electronic interfaces. These developments create harmonic resonance conditions that stress surge arresters beyond traditional design assumptions.

This research develops comprehensive frequency-dependent analysis techniques to investigate surge arrester behaviour across the operational frequency spectrum. Through systematic characterization of gapless zinc oxide varistors, the study reveals previously overlooked loss mechanisms arising from distributed grain boundary effects within the polycrystalline microstructure. These findings demonstrate that conventional frequency-independent models significantly underestimate thermal stress during harmonic temporary overvoltages, explaining discrepancies between predicted and observed failure rates.

To address these limitations, a novel fractional-order circuit model is developed that captures both dielectric relaxation phenomena and voltage-dependent nonlinear conduction. The modelling framework employs phase-sensitive decomposition techniques to separate capacitive and resistive current components, enabling accurate representation of frequency-dependent behaviour. Validation against experimental data confirms that harmonic content fundamentally alters energy dissipation patterns in ways that existing models cannot predict.

The research establishes that surge arrester quality assessment must consider frequency-dependent effects as well as voltage magnitude. The developed characterization methodology and modelling tools provide essential capabilities for evaluating surge arrester performance in modern cable-intensive, renewable-integrated grids with significantly intermittent loads, contributing to more resilient protection systems for evolving power networks.

Keywords: Metal oxide surge arresters, harmonics, temporary overvoltages, frequency-dependent characterization, fractional-order modelling, high voltage testing, ATP-EMTP

Preface

This thesis was written in partial fulfilment of the requirements for the Master of Science degree in Sustainable Energy Technologies at Delft University of Technology. It proposes a novel methodology to measure, analyse and model metal oxide surge arresters operating in their leakage current region to better assess temporary overvoltage withstand characteristics for these devices, thereby filling a critical gap in the field of power systems overvoltage protection.

The report assumes that readers possess a basic understanding of circuit theory and electrical engineering principles related to power systems. Readers with a specific interest in the analysis of surge arresters across harmonics and development of automation tools are encouraged to read Chapter 2. Those who are curious about frequency-dependent modelling of surge arresters should refer to Chapter 3.

In the spirit of transparency, I acknowledge the assistance of Anthropic's Claude in optimizing this work's analysis, as well as in organization of the technical writing throughout this thesis.

This journey of learning–much like a surge arrester's behaviour–has been decidedly non-linear in nature.

I am forever indebted to my committee members, who have over the past year, been pivotal in my academic growth and for giving me the confidence to stand proud of my work.

To my parents, who built bridges over waters they would never cross, just so I could reach the other side. And to every younger version of myself who chose to try again.

Pavan Dhulipala
Delft, August 2025

Contents

List of Figures

List of Tables

All models are wrong,
but some are useful.

— George E.P. Box

1

Introduction

The lights flickered for just 90 ms on July 10, 2024, in Northern Virginia. In that instant—faster than a human blink—a single surge arrester failure nearly brought down 70 % of global internet traffic. Sixty data centres housing 1500 MW of computing power disconnected simultaneously, their protective systems responding to a voltage disturbance that traditional engineering wisdom said should not have occurred. This near-catastrophe exemplifies a profound transformation reshaping electrical power systems worldwide: the infrastructure designed for yesterday's predictable power flows struggles to protect tomorrow's dynamic, interconnected grids [1] [2].

Across the globe, transmission networks are undergoing their most fundamental restructuring since the advent of alternating current. The imperative to integrate renewable energy—particularly from remote offshore wind farms—collides with mounting public resistance to overhead transmission corridors and increasingly stringent environmental regulations. The solution appears deceptively simple: bury the cables underground or lay them beneath the sea. Yet this shift from aerial conductors to subterranean and submarine cables introduces electrical characteristics that alter how power systems behave, creating vulnerabilities in components once considered utterly reliable.

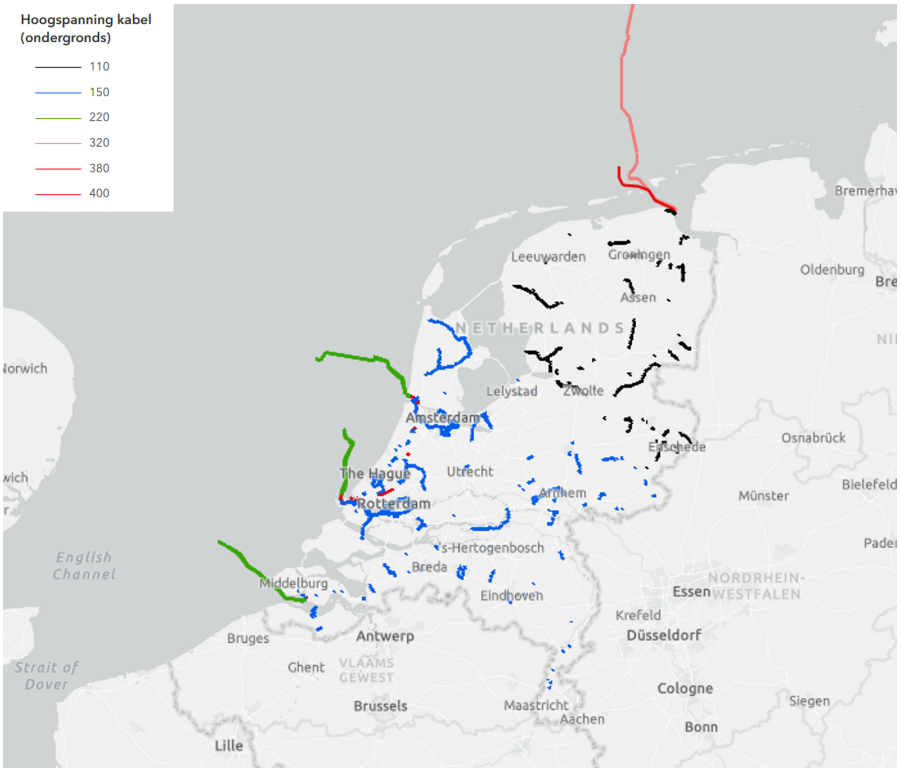

Figure 1.1: The expansive HVAC cable connections in the Dutch transmission grid

The Netherlands exemplifies this transformation. Figure 1.1 reveals an electrical landscape where submarine cables snake across the North Sea to harvest offshore wind, while underground cables thread through densely populated regions where overhead lines would face insurmountable opposition. For a nation where 26 % of the territory lies below sea level and pumping stations represent the thin electrical line between prosperity and catastrophe, grid reliability transcends technical specification—it becomes existential [3].

This cable-intensive future brings an unexpected consequence. Unlike overhead lines with their minimal capacitance, cables exhibit capacitance values 20–30 times higher per unit length. When these high-capacitance elements integrate into predominantly inductive transmission networks, they create resonant circuits at frequencies far lower than traditional grid planning anticipated. Where overhead systems might resonate at kilohertz frequencies safely above normal operations, cable-rich networks can exhibit resonances at 100 Hz–300 Hz—perilously close to the second and third harmonics of the fundamental power frequency [4] [5].

The technical implications also have economic consequences. In practice, surge arresters priced in the thousands of euros protect primary grid assets valued in the millions. A single arrester failure can lead to transformer damage, forced outages, and regulatory costs that exceed the arrester's purchase price by orders of magnitude.

Beyond economics lies human vulnerability. Modern society's complete dependence on electricity transforms power outages from inconvenience to crisis. During failures, 14.5 million households lose access to electricity-dependent medical equipment, with ventilator batteries lasting mere hours [6]. The 2003 Northeast Blackout demonstrated this starkly: mortality rates

spiked not just during the event but for days afterward as overwhelmed medical systems struggled to recover [7]. Low-income communities suffer disproportionately, experiencing restoration delays twice as long as affluent areas while lacking resources for backup power [8].

These cascading vulnerabilities converge at a critical component: the surge arrester. These compact devices protect transformers, cables, and system stability by exhibiting voltage-dependent conductivity. Under normal operating conditions, they maintain high impedance, conducting only microampere-level currents. During lightning or switching surges, their impedance drops dramatically, providing a low-resistance path to ground that clamps voltages within equipment withstand limits. This conduction persists only for the surge duration before the arrester returns to its high-impedance state [9].

Yet the transition to cable-rich networks fundamentally alters the stresses these protective devices must withstand. When routine operations like transformer energization coincide with cable-created resonances, the resulting overvoltages differ profoundly from the millisecond transients arresters were designed to handle. Instead of brief spikes, these resonance-induced Temporary Overvoltages (TOVs) can persist for hundreds of milliseconds with significant harmonic content, subjecting arresters to thermal stresses their designers never anticipated [10].

The physics underlying this challenge reveals itself through the grain structure of modern metal oxide arresters. Within each arrester, millions of zinc oxide (ZnO) crystals form a polycrystalline matrix where grain boundaries act as voltage-dependent switches; a simplified image of this can be seen in Figure 1.2a. Under normal voltage, these boundaries present high impedance, allowing only microamperes of current. As voltage rises, the barriers progressively weaken, permitting controlled conduction that limits system overvoltage.

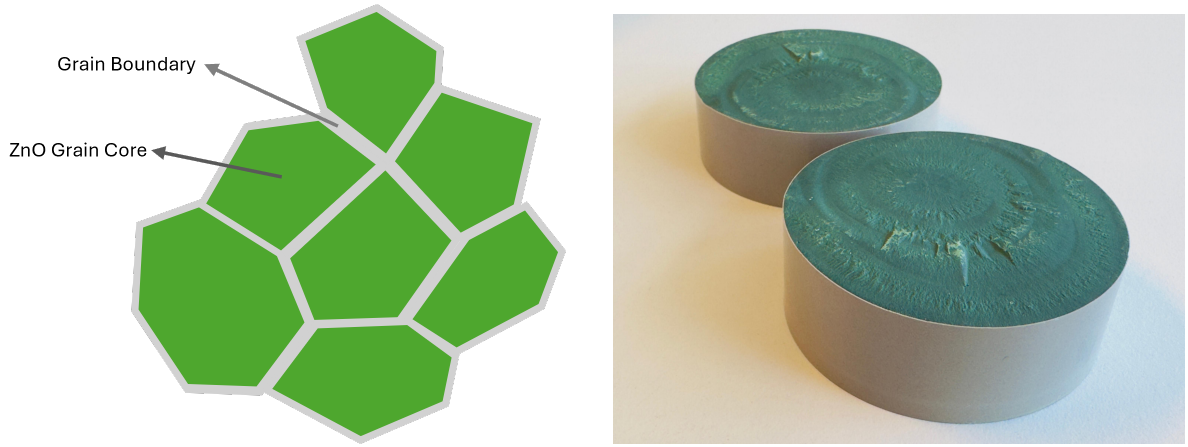

(a) Representation of intergranular ZnO crystals in the block

(b) Inside view of an MOV block

Figure 1.2: ZnO Block Structure

Extended harmonic-rich TOVs violate the assumption that arresters cool down after thermal accumulation. Even modest currents, when sustained for hundreds of milliseconds, generate heat faster than the ceramic structure can dissipate it [11]. Temperature rise reduces grain boundary resistance, increasing current flow in a potentially catastrophic positive feedback loop. The arrester that faithfully protected equipment for decades can fail within seconds when subjected to conditions outside its design envelope.

1.1. The Evolution of Power System Protection

The historical development of surge protection illuminates why current challenges prove so vexing. Early power systems employed spark gaps—crude but effective devices that physically separated conductors until voltage exceeded breakdown threshold. These gaps provided reliable isolation during normal operation and decisive action during overvoltages, but suffered from inconsistent operation and limited interrupting capability.

The introduction of silicon carbide arresters with series gaps in the mid-20th century improved coordination and energy handling, yet still required the arrester to interrupt follow current after operation. This limitation became increasingly problematic as system voltages and power levels rose, demanding more sophisticated protection.

The revolution came with ZnO varistors in the 1970s. These gapless arresters offered unprecedented advantages: consistent clamping voltages, instantaneous response, and no follow current to interrupt. A properly designed metal oxide arrester could conduct microamperes at normal voltage and kiloamperes during surges—a dynamic range spanning six orders of magnitude achieved through the non linear conduction mechanism dependent on the dielectric field strength in the granular layer of the MOV [12].

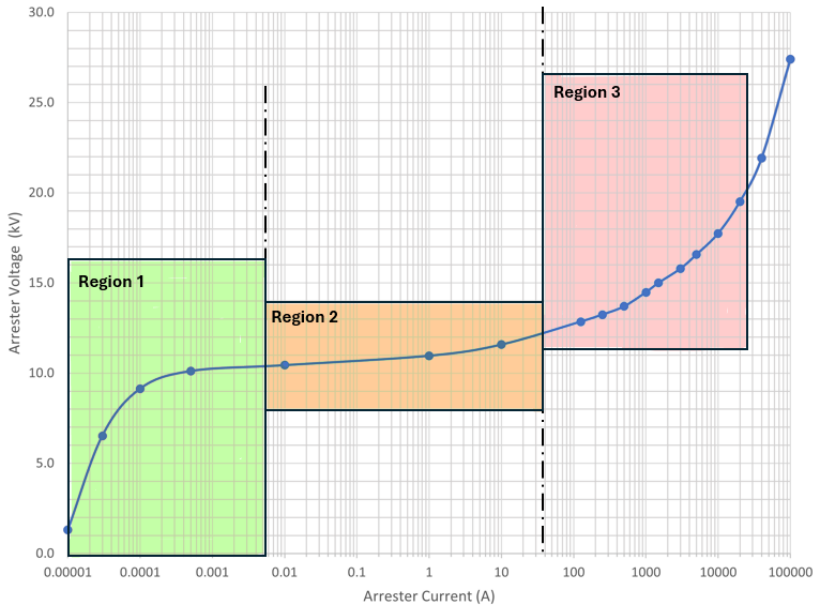

Figure 1.3: Typical V-I Characteristics of the MOV block used for this thesis study [13]

Figure 1.3 reveals this nonlinearity. Below rated voltage, the arrester appears nearly as an open circuit. As voltage rises through the transition region, current increases exponentially until the protection region where the device effectively clamps voltage despite enormous current variations. This characteristic enables a single device to protect against everything from switching surges to direct lightning strikes.

Success bred complacency. With metal oxide arresters proving remarkably reliable under traditional stresses, the industry standardized testing around well-understood threats: lightning impulses lasting microseconds and switching surges decaying within milliseconds and TOV testing. The IEC 60099-4 standard codified these tests, establishing TOV withstand require-

ments based on pure power-frequency overvoltages of defined amplitude and duration [14] [11].

Yet even as standards crystallized around traditional threats, prescient researchers noted emerging vulnerabilities where TOVs are no longer avoidable. CIGRE Working Group 33.06 warned three decades ago that harmonic resonances could subject arresters to “severe energy stresses through mechanisms not captured by standard testing” [12]. Their concerns, largely theoretical at the time, now manifest as operational failures in cable-rich networks worldwide. Detailed explanation of this is given in Section 1.4.

1.2. The Resonance Challenge

Understanding why cable systems create such profound challenges requires examining the main electrical differences between overhead lines and cables. An overhead transmission line, with conductors separated by air and suspended high above ground, exhibits relatively low capacitance—typically 10 nF km^{-1}–15 nF km^{-1}. This capacitance, combined with the system equivalent inductance, creates natural resonances at frequencies well above normal power system operations in weak grids [15] [16].

Cables tell a different story. With conductors surrounded by insulation and grounded shields, geometric constraints force electric fields into compact volumes, yielding capacitances of 150 nF km^{-1}–400 nF km^{-1}. When kilometres of such cables connect to a network designed around overhead line characteristics, resonant frequencies that once resided safely in the kilohertz range plummet toward power frequency harmonics [11] [4].

The physics are unforgiving. A resonant circuit’s frequency follows:

$$f_r = \frac{1}{2\pi\sqrt{LC}} \tag{1.1}$$

As cable capacitance C increases twenty-fold while system inductance L remains largely unchanged, resonant frequency f_r drops by a factor of 4.5.

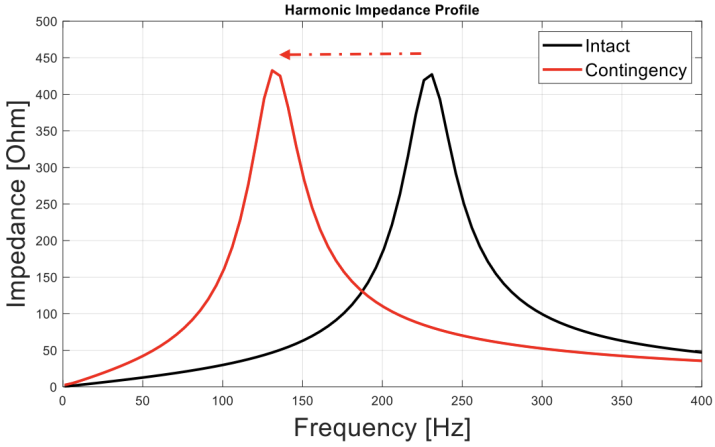

Figure 1.4: Illustrative shift of typical harmonic impedance profile [17]

Real networks are more complex. Multiple cables, transformers with nonlinear magnetizing characteristics, and varying load conditions create a rich spectrum of potential resonances. When routine operations inject harmonic currents—transformer energization producing second and third harmonics, capacitor switching generating high-frequency transients—these currents find resonant circuits ready to amplify them into sustained overvoltages [17].

The Dutch transmission system provides a compelling case study. Massive offshore wind farms connect through submarine cables to onshore substations, where underground cables continue the path to load centres. Each transition point between cable and overhead sections creates impedance discontinuities that can trap and amplify harmonic energy. System studies reveal resonances clustering around 100 Hz–300 Hz, precisely where transformer inrush and other switching operations concentrate their harmonic content [10] [18].

1.3. Surge Arrester Vulnerability

The convergence of cable-induced resonances and surge arrester characteristics creates conditions for vulnerability. Traditional overvoltage events—lightning and switching surges—subject arresters to high energy for brief periods. The ceramic structure readily absorbs and dissipates this energy, with cooling occurring between events. Resonance-induced TOVs overturn these assumptions by applying moderate overvoltages for extended periods with rich harmonic content.

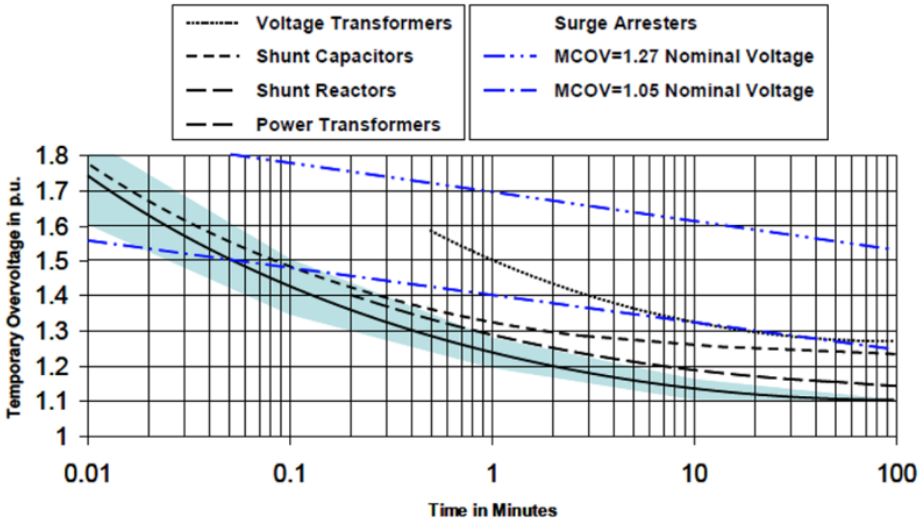

Figure 1.5: Typical power frequency TOV voltage-duration withstand curves [19]

Figure 1.5 illustrates the challenge. Among all major power system components, surge arresters (blue curve) exhibit the lowest tolerance to sustained overvoltages below 600 ms. While transformers and cables can withstand 1.5 p.u. voltage for several seconds, arresters may fail within 100 ms at the same level. This vulnerability, manageable when TOVs were rare and brief, becomes critical when resonances create frequent, sustained overvoltage conditions.

Each ZnO grain within the arrester has finite heat capacity and limited thermal conductivity to neighbouring grains and the external housing. During brief surges, temperature rise remains localized and quickly dissipates. Extended conduction, even at moderate currents, allows heat

to accumulate faster than it dissipates. The resulting temperature rise triggers a cascade of effects:

1. Elevated temperature reduces grain boundary resistance
2. Lower resistance increases current at given voltage
3. Higher current generates additional heat
4. Positive feedback accelerates toward thermal runaway

Harmonic content intensifies these thermal challenges through mechanisms poorly understood until recently. At power frequency, surge arresters exhibit well-characterised impedance dominated by grain boundary capacitance in the normal operating region. Harmonic frequencies alter this impedance as capacitive reactance decreases with frequency, potentially allowing more current flow within the arrester [20].

1.4. Standardized Testing Versus Emerging Realities

The disconnect between laboratory validation and field performance reveals itself in the standardized testing of surge arresters. The IEC 60099-4 standard, in its section 8.8, prescribes a procedure for verifying TOV withstand capability, a test sequence that embodies decades of industry experience with traditional power systems. Yet, the underlying assumptions may no longer hold true in cable-dominated networks.

The TOV type test unfolds as a stress sequence designed to represent worst-case field conditions. Six arrester samples undergo evaluation: four subjected to prior duty and two tested without prior duty. Prior duty involves injecting the arrester's rated thermal energy to heat the blocks to approximately 100 °C–120 °C, simulating the thermal stress from recent surge absorption. For station arresters, this consists of 1.0 to 1.1 times their rated thermal energy (W_{th}) through rectangular current impulses lasting 2 ms–4 ms; for distribution arresters, 1.0 to 1.1 times their rated thermal charge (Q_{th}) via paired 8/20 µs lightning impulses. This energy injection must complete within a three-minute window for station class or one minute for distribution class. Samples tested without prior duty begin at 60 °C. Both groups may undergo two-pulse preconditioning as a burn-in procedure before testing begins.

Following prior duty, the TOV is applied within 100 ms—less than two power-frequency cycles—ensuring no significant cooling. Each sample experiences different TOV durations to map the complete withstand characteristic: 0.1 s–1 s, 1.1 s–10 s, 10.1 s–100 s, and 101 s–3600 s. The test voltage at 10 s must equal or exceed the arrester's rated voltage (U_{r}), ensuring adequate margin for system conditions. Following TOV application, continuous operating voltage (U_{c}) stresses the samples for at least 30 min while monitoring systems track resistive current or power dissipation. The critical test for prior duty samples is whether they can cool down from their elevated starting temperature. Success demands meeting all three criteria: demonstrated thermal recovery through declining power dissipation, absence of physical damage, and residual voltage at nominal discharge current remaining within ± 5 % of pre-test values [14].

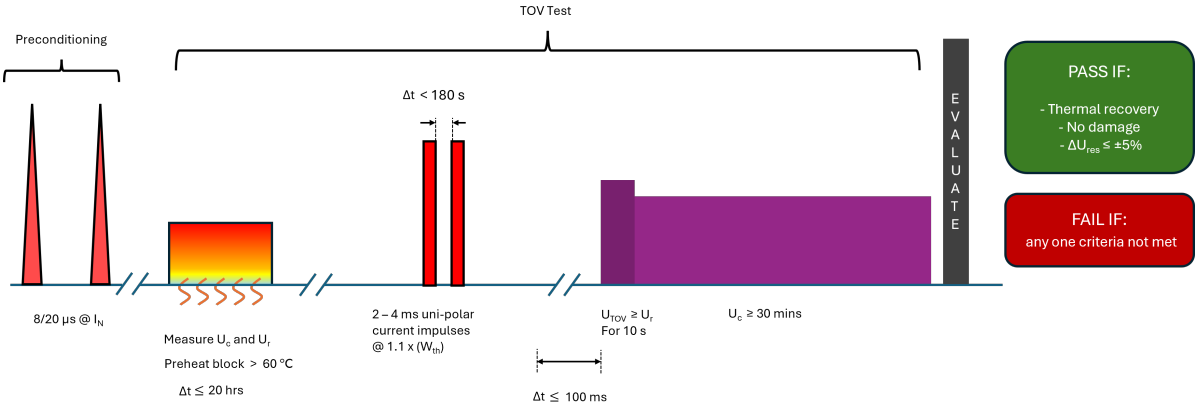

Figure 1.6: IEC 60099-4 TOV test sequence for station class arresters showing energy injection followed by overvoltage application and thermal recovery monitoring.

This standardized approach assumes TOVs manifest as elevated power-frequency voltages with well-defined amplitudes and durations—an assumption that held when overhead lines dominated, system behaviour remained predictable and extended TOVs were avoided by protection and control, operating restrictions [11]. Crucially, the test focuses primarily on resistive heating, neglecting frequency-dependent dielectric losses that become significant at harmonic frequencies. Even as early as 1990, researchers recognized these limitations. The CIGRE Working Group 33.10 acknowledged that real-world TOV waveforms rarely maintain constant amplitude [21]. Their solution—representing complex waveforms as envelopes and summing fractional energy contributions—offered a pragmatic approximation:

$$\frac{t_1}{T_1} + \frac{t_2}{T_2} + \cdots + \frac{t_n}{T_n} < 1 \tag{1.2}$$

where t_i represents the duration at voltage level i, and T_i the corresponding withstand time from the TOV curve. This envelope method assumed that harmonic content remained negligible and that thermal effects could be linearized—assumptions increasingly violated in modern systems.

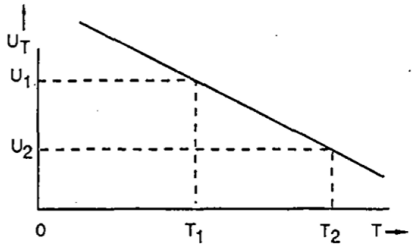

(a) Power frequency vs Time curve of a surge arrester MOV

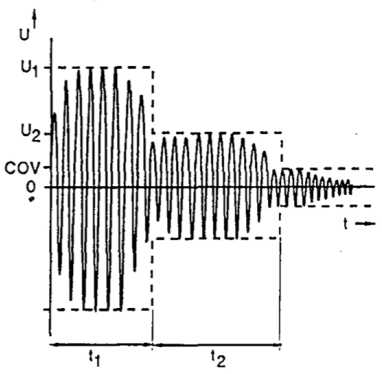

(b) TOV wave shape of varying amplitudes

Figure 1.7: CIGRE envelope method for complex TOV waveforms. The approach approximates time-varying overvoltages as discrete amplitude levels, enabling comparison against power-frequency TOV curves despite inherent limitations[21].

The replacement of overhead lines with cables changes system harmonic behaviour. Consider transformer energization, a routine operation that demonstrates this difference. When energized, transformers draw inrush currents that differ from load currents. These currents, initially reaching several times rated magnitude, contain second harmonic content due to core saturation and the nonlinearity of the magnetizing inductance. In traditional overhead systems, these harmonics dissipate rapidly through line resistance and radiation.

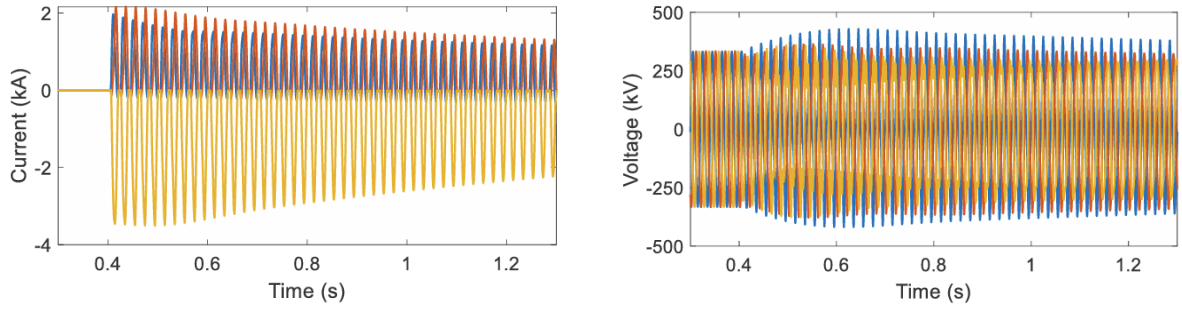

(a) Transformer inrush current

(b) Resulting TOV reaching 1.4 p.u.

Figure 1.8: Transformer energization in cable systems: Harmonic-rich inrush currents (left) excite network resonances, producing sustained overvoltages with matching harmonic content (right)[11].

Harmonic currents from transformer saturation encounter the resonant circuit formed by transformer inductance and cable capacitance. When these harmonics align with network resonance frequencies—often at the 2nd, 3rd, or 5th harmonic in cable-rich systems—amplification occurs. The resulting TOV no longer resembles the clean rectangular pulse assumed by standards but manifests as a sustained oscillation combining elevated fundamental voltage with significant harmonic content.

The implications for surge arresters extend beyond simple voltage magnitude. Each harmonic component induces its own current through the arrester, contributing not only resistive losses captured by traditional models but also frequency-dependent dielectric losses within the ZnO grains and their boundaries. As demonstrated in this thesis, these dielectric losses, negligible at power frequency, increase with frequency. An arrester experiencing a 1.4 p.u. TOV with 20 % third harmonic content absorbs substantially more energy than one facing 1.4 p.u. at pure power frequency—yet both scenarios map to the same point on conventional TOV curves.

Recent experimental evidence confirms that these theoretical concerns translate directly to reduced arrester life. Accelerated ageing tests under harmonic distortion revealed that even modest harmonic content—6.24 % third harmonic and 5.58 % fifth harmonic—increased the probability of premature failure by 58.93 % and reduced mean time to failure by 40.91 % [20]. The observed 92.51 % increase in resistive current under harmonic stress correlated directly with accelerated degradation, validating that frequency-dependent losses drive thermal runaway through mechanisms unaccounted for in standards. This experimental validation underscores the urgency of developing models that capture these frequency-dependent degradation mechanisms.

This divergence between standardized testing and operational reality frames the central challenge: how can engineers predict or evaluate surge arrester survivability when the stresses they face increasingly fall outside the bounds of conventional characterisation? The answer requires not merely refined testing but fundamental reconsideration of how we model arrester

behaviour under complex electrical stress, particularly the frequency-dependent loss mechanisms that transform harmonic content from nuisance to threat—a gap that existing approaches have yet to bridge.

1.5. The Modelling Gap

Confronted with field failures and recognizing the inadequacy of existing standards, the industry faces a critical knowledge gap. Current surge arrester models, developed for lightning and switching surge analysis, employ simplified representations adequate for microsecond phenomena but wholly inappropriate for sustained harmonic-rich stresses. These models typically represent the arrester as a nonlinear resistor, perhaps with parallel capacitance, ignoring the frequency-dependent behaviour crucial to understanding harmonic resonance effects.

More sophisticated models like that shown in Figure 1.9 attempt to capture the distributed nature of grain boundaries through multiple resistive and capacitive elements. Yet even these enhanced representations struggle to accurately predict behaviour under harmonic-rich TOVs. The fundamental challenge lies in characterising how ZnO grain boundaries respond to sustained harmonic currents—behaviour never fully explored because it was never previously necessary.

This modelling gap creates cascading uncertainties. System planners cannot accurately assess arrester stress during resonance events. Protection engineers lack tools to determine appropriate arrester ratings for cable-rich networks. Manufacturers design products using test standards that omit critical stress conditions. The result: arresters selected using traditional methods may fail prematurely especially when networks are most vulnerable.

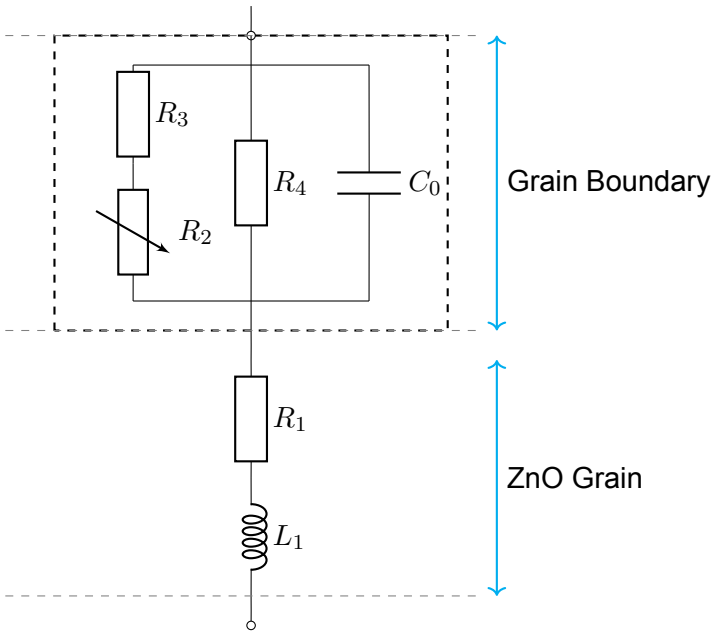

Figure 1.9: Equivalent circuit for an MO arrester block [22]

1.6. Research Question

The convergence of technological transformation and societal dependence creates an urgent research imperative. As renewable integration drives cable deployment and society deepens its reliance on uninterrupted power, surge arrester reliability evolves from technical detail to critical infrastructure concern. The €40 billion annual cost of grid congestion in the Netherlands alone demonstrates the economic stakes. The 14.5 million households dependent on medical equipment reveal the human dimension. The near-catastrophe in Northern Virginia warns of systemic vulnerabilities in our most advanced infrastructure.

This context frames the research question:

> *"How can we estimate electrical parameters and develop accurate modelling for Metal Oxide Surge Arresters operating under low-order harmonic resonance conditions to ensure their appropriate selection and application?"*

Answering this question demands new characterisation methods that capture frequency-dependent behaviour, enhanced models that predict energy accumulation under harmonic stress, and validation approaches that bridge laboratory testing to field performance. Most critically, it requires translating technical understanding into practical tools that protection engineers can apply to ensure grid resilience. The overarching research challenge naturally decomposes into specific technical questions, each addressing an aspect of surge arrester performance under harmonic resonance conditions:

1. **How can the behaviour of a surge arrester be captured and accurately classified for analysis?**

   Traditional characterisation methods focus on voltage-current relationships at power frequency or impulse response. Understanding harmonic resonance effects requires new approaches that reveal frequency-dependent impedance characteristics, thermal response to sustained conduction, and the interaction between capacitive and resistive current components across the frequency spectrum.
2. **How can we develop a frequency-dependent surge arrester model to better evaluate TOV duration curves and improve selection of surge arresters in their respective applications?**

   Existing models fail to capture the distributed nature of grain boundary impedances and their frequency-dependent behaviour. A new modelling framework must represent both the nonlinear conduction mechanisms and the complex impedance variations that determine energy absorption under harmonic-rich conditions.
3. **What methodology can be implemented by manufacturers and grid operators alike to correctly assess the rating of surge arresters during TOV scenarios to ensure grid resiliency?**

   Beyond theoretical understanding, practical application demands standardized approaches for evaluating arrester performance under realistic stress conditions. This includes test protocols that incorporate harmonic content, rating methodologies that account for thermal accumulation, and selection criteria adapted to cable-rich networks.

1.7. Thesis Structure and Contributions

This thesis addresses these questions through systematic investigation, progressing from fundamental characterisation through model development to practical application. Each chapter builds upon previous findings while maintaining focus on the ultimate goal: ensuring surge arrester survivability in evolving power systems.

Chapter 2 applies these methodological principles through laboratory testing. Multi-frequency characterisation reveals the complex impedance behaviour of surge arresters across the spectrum relevant to harmonic resonance, providing the empirical foundation for enhanced modelling.

Chapter 3 transforms characterisation insights into a practical frequency-dependent model implemented in ATP-EMTP. The chapter navigates the challenges of representing distributed grain boundary effects within transient simulation constraints, developing solutions that balance accuracy with computational efficiency.

Chapter 4 validates the modelling framework through extensive comparison with experimental data. Multiple test scenarios examine model performance under varied harmonic content and voltage levels, while application to different arrester types confirms the methodology's generalizability.

Chapter 5 synthesizes technical findings into practical guidance. The chapter examines implications for testing standards, selection procedures, and network planning, translating research outcomes into actionable recommendations for industry application. Socio-economic impact assessments quantify the benefits of enhanced surge arrester survivability.

Chapter 6 consolidates the research contributions and charts future directions. Beyond summarizing key findings, the chapter proposes future recommendations for these gapless metal oxide varistors.

Research Innovation

This thesis pioneers three advances in surge arrester technology:

- **Characterisation**: First comprehensive multi-frequency impedance analysis revealing previously unaccounted frequency-dependent loss mechanisms
- **Modelling**: Novel fractional-order circuit representation capturing distributed grain boundary effects
- **Application**: Practical methodology enabling accurate TOV assessment in surge arresters

Together, these contributions establish a basis for predictive engineering practice in surge-arrester characterisation and modelling.

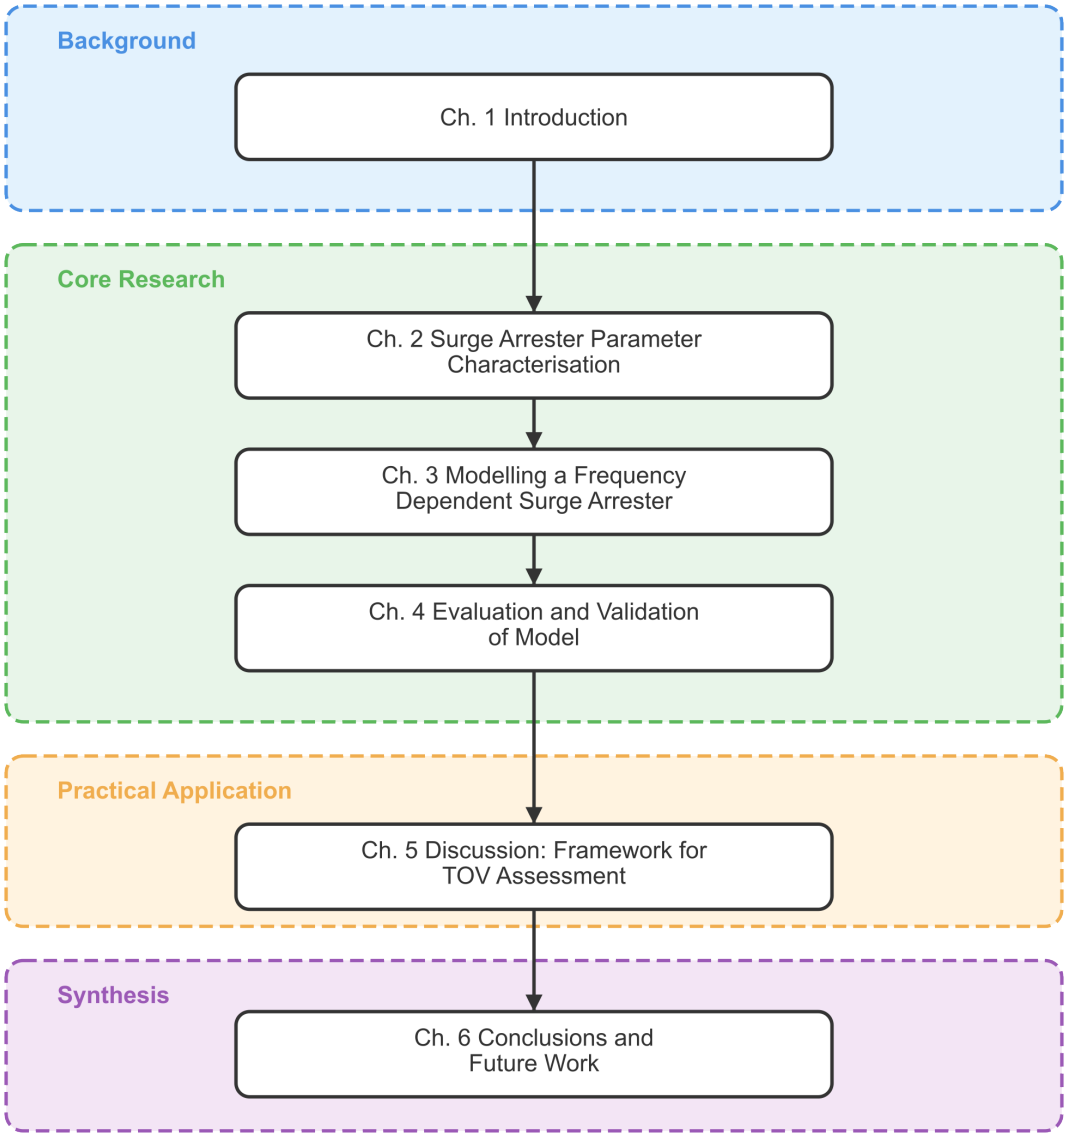

Figure 1.10: Thesis Outline

2

Surge Arrester Parameter Characterisation

How does an MOV truly respond when subjected to the complex voltage stresses of modern power systems? How can we extract meaningful parameters from a device that operates at the intersection of linear dielectric idle behaviour and highly non-linear conduction? These questions drive the experimental investigation at the heart of this chapter.

Surge arresters operate in a regime where currents span from microamperes to kiloamperes, where impedances range from megohms to ohms, and where the transition between these extremes can occur within single multiples of the rated voltage. Traditional characterisation methods, developed for either linear components or switching devices, do not capture the nuanced behaviour of these grain boundary-dominated systems. We need new approaches that can distinguish frequency-dependent losses from voltage-dependent phenomena.

Section 2.1 establishes the experimental foundation, describing the measurement setup used to extract the surge arrester behaviour across a frequency range spanning DC to the tenth harmonic. Section 2.2 then unveils the signal processing techniques that transform raw oscilloscope traces into quantified electrical parameters, including a decomposition method that separates capacitive, linear resistive, and non-linear resistive components. The results presented in Section 2.3 will guide the development of enhanced surge arrester models in subsequent chapters.

2.1. Measurement Setup

Accurate surge arrester characterisation demands a measurement system that can navigate between extremes—capturing both the subtle dielectric currents at low voltages and the emerging non-linear conduction as grain boundaries begin to break down. The experimental arrangement must maintain stability across five decades of impedance variation while providing the resolution necessary to distinguish frequency-dependent phenomena from voltage-dependent effects.

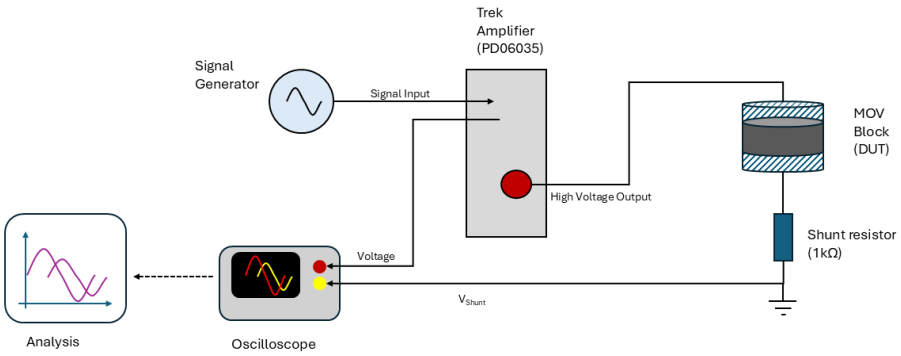

Figure 2.1: Single-line diagram of the measurement arrangement.

The laboratory setup, shown in Figure 2.1, begins with a signal generator that produces either pure sinusoidal waveforms or a steady DC level. Its high-impedance output drives a Trek 30/20A amplifier (Figure 2.2b) whose fixed gain of 3000 V/V and 47 Ω output impedance of the voltage monitor are small in comparison with the mega-ohm impedance of the test specimen; hence the amplifier behaves as an ideal voltage source over the test range. The amplifier's ±20 mA peak current limit provides an intrinsic safeguard against accidental overstress of the MOV under investigation.

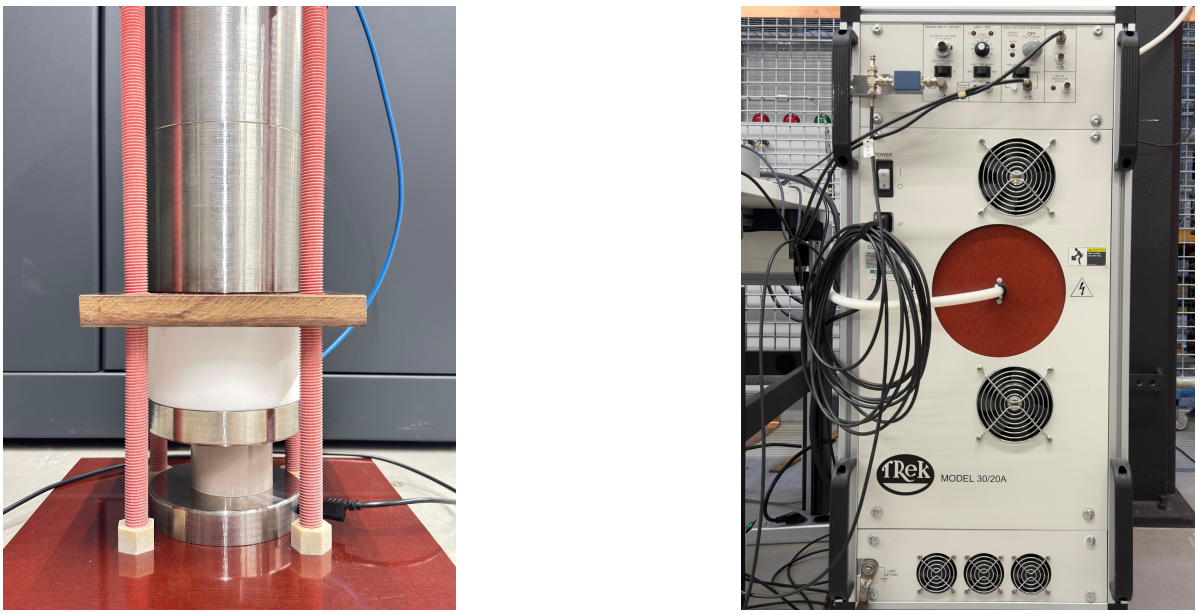

(a) Designed testing fixture

(b) Voltage Amplifier

Figure 2.2: Measurement arrangement in lab

The amplified waveform passes to a clamping rig in which two plane-parallel aluminium blocks apply a compressive force of approximately 20 kg. This load corresponds to a contact pressure of 60 kPa and guarantees low-resistance electrical contact with the thin aluminium electrodes

that coat the MOV surface.

The current-monitor output provided by the high-voltage amplifier includes not only the conduction current through the surge arrester but also the current produced by its own stray (parasitic) capacitances. Additionally, the amplifier's internal current monitoring introduces phase delays that compromise the phase-sensitive decomposition required for capacitive and resistive current separation. To avoid these errors, the test circuit employs a 1 kΩ shunt resistor connected between the arrester's bottom electrode and earth. This direct measurement ensures the voltage developed across the shunt is proportional only to the true arrester current, with phase relationships preserved across the frequency range including harmonics that would otherwise be distorted by parasitic effects.

The PicoScope 6424E oscilloscope serves as the data acquisition part of the system. Its specifications—8-12 bit vertical resolution and 5 GS/s maximum sampling rate—are carefully matched to our measurement needs. The actual sampling frequencies employed are frequency-dependent: 1 MS/s for 10 Hz and 150 Hz measurements, 400 kS/s for 17 Hz–100 Hz range, 2 MS/s for 300 Hz, and 4 MS/s for 500 Hz tests. The built-in low pass filter within the PicoScope has selectable cutoff frequencies. For sub-50 Hz measurements, the 100 Hz cutoff frequency provides 3 dB attenuation while for signals $\geq$50 Hz, the cutoff frequency was selected to be 5 kHz

The test is done maintaining stable conditions at $\approx$23 °C without any active control. This temperature control proves essential—preliminary tests revealed that even small temperature variations could shift leakage currents, potentially masking the frequency-dependent effects we seek to quantify and creating a biased measurement environment. The test was done on a station class arrester MOV block manufactured by TDK Electronics (type E64NR163E) [13]. It is shown in Figure 2.3 and its specifications are summarized in Table 2.1.

Table 2.1: Principal specifications of the test block (values from the manufacturer's data sheet [13]).

Parameter	Value
Diameter	64.5 mm
Thickness	44 mm
Maximum continuous operating voltage (rms)	4.97 kV
Reference current	5 mA
Reference voltage (rms)	6.98 kV
Nominal discharge current (8/20 µs)	20 kA
Residual voltage at 10 kA (8/20 µs)	$\approx$15.9 kV

Figure 2.3: MOV Block used for testing

The selection of test parameters reflects both theoretical requirements and practical constraints. Twenty-one voltage levels span from 3 kV_p to 10 kV_p, chosen to bracket the transition from linear dielectric behaviour to non-linear conduction. The lower bound ensures measurable currents rise above the noise floor at the lowest test frequency, while the upper bound is at the amplifier's current limit while still exceeding the peak continuous operating voltage. This voltage range, with its ≈350 V average increment, provides sufficient resolution to capture the exponential transition in conductivity that occurs near 0.7 times the rated voltage (U_r).

Table 2.2: Test parameters for dielectric characterisation.

Parameter	**Values**
Voltage (kV peak)	3.0, 3.5, 3.7, 4.0, 4.4, 4.7, 5.0, 5.3, 5.6, 5.9, 6.3, 6.8, 7.2, 7.7, 8.2, 8.6, 9.0, 9.2, 9.5, 9.7, 10.0
Frequency (Hz)	DC, 10, 17, 27, 50, 100, 150, 300, 500

Beyond the power frequency and its harmonics, sub-synchronous frequencies are measured to understand the frequency-dependent behaviour that emerges from the distribution of relaxation times in the polycrystalline structure.

The upper frequency limit of 500 Hz reflects distinctions in power system phenomena. In insulation-coordination practice, a TOV is any undamped or weakly-damped power-frequency overvoltage whose spectrum sits wholly below about 500 Hz ($\approx 10 \times f_0$); above that threshold the network no longer behaves as a lumped power-frequency circuit, the stresses shift from over-fluxing to impulse behaviour, and IEC therefore categorises such events as slow-front or fast-front transients rather than TOV. By constraining our investigation to this frequency range, we ensure that our characterisation remains relevant to the temporary overvoltage phenomena that represent the most thermally stressful conditions for surge arresters.[11]

Data acquisition employs frequency-dependent sampling rates: 400 kS/s for 17 Hz–100 Hz, 1 MS/s for 10 Hz and 150 Hz, 2 MS/s for 300 Hz, and 4 MS/s for 500 Hz measurements. Acquisition windows range from 50 ms to 1 s, capturing 10-50 cycles per measurement. The oscilloscope vertical range is set to ±10 V for voltage channels and ±2 V to ±20 V for current channels depending on test conditions. High sampling frequencies ensure accurate waveform digitization with sufficient resolution for zero-crossing detection and phase angle computation.

Each acquisition yields approximately 200k data points at 10-bit resolution. Signal processing algorithms subsequently extract capacitive and resistive current components.

Measurement discipline extends beyond equipment specifications to procedural details that ensure data quality. The nested measurement sequence—stepping voltage at each frequency rather than frequency at each voltage—minimizes electrical stress on the MOV block. After completing twenty-one voltage points at a given frequency, a fifteen-minute cooling period allows thermal equilibrium to be restored. This systematic approach ensures that each measurement begins from identical initial conditions, eliminating temperature as a variable in the analysis.

2.2. Data Analysis

Signal conditioning begins with hardware filtering inside the PicoScope to remove undesired high frequency noise before digitization. The filter cut-off frequency adapts to the measurement frequency: for tests below 50 Hz, the cut-off is set to 100 Hz, while measurements at 50 Hz and above employ cut-offs between 1 kHz and 5 kHz. This approach maintains a flat passband over the frequency range of interest while providing attenuation of high-frequency noise.

Following removal of the initial transient period, the steady-state data is truncated to the nearest power of two—for this analysis—2^{19} samples for 10 Hz tests and 2^{17} samples for higher frequencies.

The next filtering stage employs a 100 µs rectangular moving-average window:

$$\tilde{v}[n] = \frac{1}{M} \sum_{k=0}^{M-1} v[n-k], \qquad M = \lfloor 0.0001\, F_s \rfloor,$$

where $\tilde{v}[n]$ is the filtered voltage at sample n, $v[n-k]$ represents the raw voltage k samples earlier, M is the window length in samples, and F_s is the sampling frequency. This window removes spike noise while maintaining consistency with later resistive current smoothing. Within the passband, the window's frequency response introduces 3.92 dB attenuation at 5 kHz. The mean value is subtracted from each trace to centre the signals on zero, which simplifies zero-crossing detection, prevents spectral leakage into the DC bin and reduces the effect of any offset introduced by the oscilloscope.

Accurate determination of the applied test frequency is required because the phase-sensitive decomposition assumes a 90° relationship between the capacitive and resistive components. The frequency and phase are obtained from the voltage signal using a Hann-windowed FFT with three-bin interpolation; the derivation and corrections are given in Appendix C and [23, 24]. The magnitude spectra corresponding to this step are shown in Figure 2.4. The resulting applied frequency and phase are then used wherever a sinusoidal reference is required.

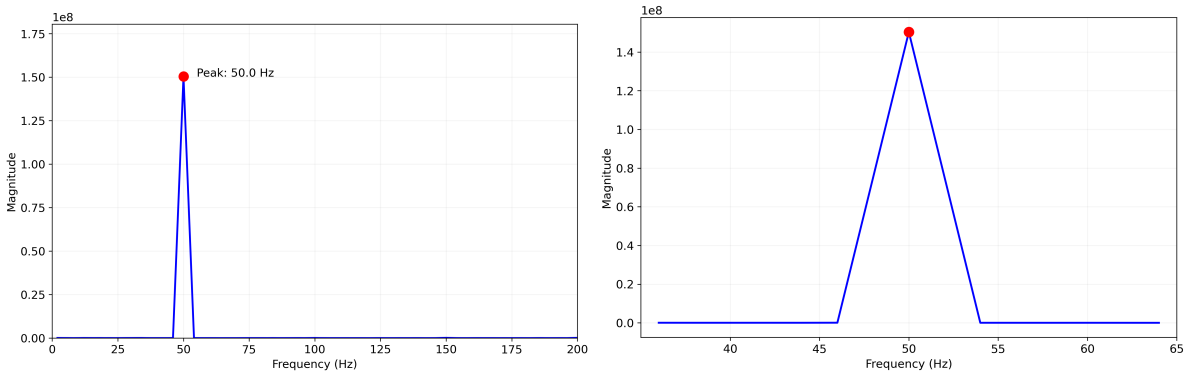

(a) Full spectrum **(b)** Zoomed fundamental

Figure 2.4: Magnitude spectrum of the Hann-windowed voltage waveform.

A time-domain check is performed on the same steady-state window by measuring the spacing of multiple voltage zero-crossings. Agreement between the spectral estimate and the median zero-crossing period is used as a quality measure; in these measurements the difference is typically within ± 0.5 Hz.

A Goertzel implementation is available to evaluate the discrete Fourier transform at selected frequencies [25]. In this work it is retained only as a diagnostic to check suspected harmonic contamination and is not used to construct the capacitive or resistive waveforms. Both FFT and Goertzel exhibit spectral leakage when the frequency is off-bin, hence the consistent use of a Hann window; details are provided in Appendix C.

The decomposition isolates the applied test frequency. Harmonics, when present, are treated as interference for parameter extraction and are not used in forming the capacitive or resistive components. Record length is determined by the acquisition settings in Section 2.1; the steady-state window timebase contains 10–50 cycles depending on test frequency.

Following frequency determination, the voltage signal is integrated to establish a phase reference:

$$v_{\mathrm{int}}[n] \;=\; \sum_{m=0}^{n-1} v_{\mathrm{AC}}[m]\,\Delta t,$$

where $v_{\mathrm{int}}[n]$ is the running integral at sample n and $\Delta t = 1/F_s$ is the sampling interval. Integration multiplies each Fourier component by $1/(\mathrm{i}\omega)$, reducing spectral magnitude as $1/\omega$ and shifting phase by $-90°$. The integrated waveform in Figure 2.5 shows reduced noise and stable zero-crossings with less than 20 µs jitter.

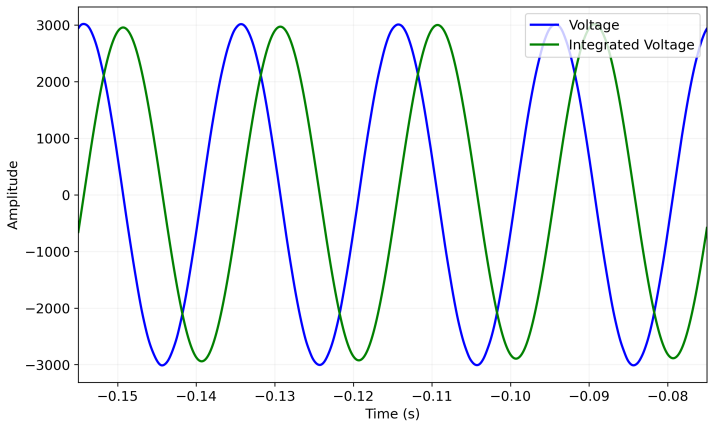

Figure 2.5: Integrated voltage signal used as phase reference for signal decomposition.

Capacitive and resistive current separation uses AC circuit phase relationships. In a capacitor, current leads voltage by 90°, reaching maximum at voltage zero-crossing. Sampling total current at voltage zero-crossings captures capacitive current at maximum while resistive current equals zero.

$$I_{\text{cap,peak}} \approx \text{mean}\left|i_{\text{AC}}(t_{\text{zc}})\right|, \qquad C_{\text{s}} = \frac{I_{\text{cap,peak}}}{\omega\, v_{\text{peak}}}, \qquad \omega = 2\pi f_{\text{applied}},$$

where $i_{\text{AC}}(t_{\text{zc}})$ is measured current at voltage zero-crossings, C_{s} is series capacitance, v_{peak} is peak voltage amplitude, and f_{applied} is the applied test frequency.

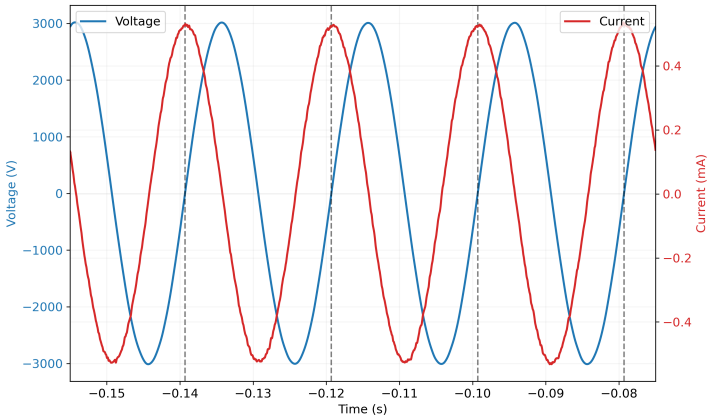

Figure 2.6: Voltage zero-crossings used to identify capacitive current peaks.

Complete capacitive waveform reconstruction from peak measurements maintains the 90° phase relationship (Figure 2.7). The sinusoidal capacitive current maintains phase lead throughout the cycle:

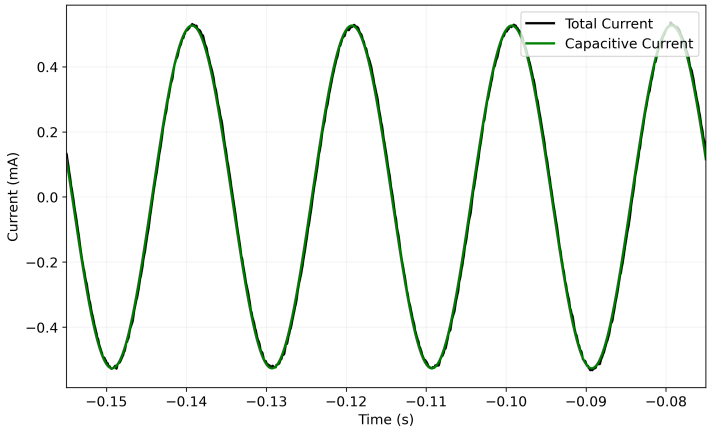

Figure 2.7: Reconstructed capacitive current with 90° phase lead.

Resistive current is obtained through subtraction:

$$i_{\text{res}}(t) = i_{\text{AC}}(t) - i_{\text{cap}}(t),$$

where $i_{\text{res}}(t)$ is extracted resistive current, $i_{\text{AC}}(t)$ is measured total AC current, and $i_{\text{cap}}(t)$ is reconstructed capacitive current. High-frequency components in resistive current are removed using the 100 µs moving-average window. Figure 2.8 shows the separated components.

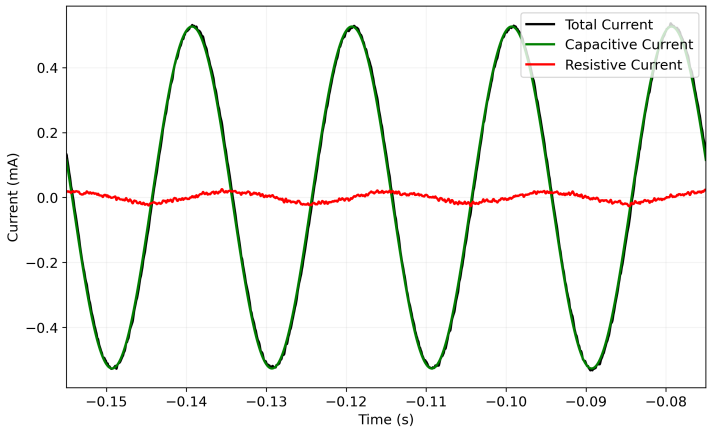

Figure 2.8: Current decomposition into capacitive and resistive components.

The 3 kV measurement as shown in Figure 2.9 provides the reference for linear behaviour. At this voltage below the conduction knee, nonlinear effects are minimal.

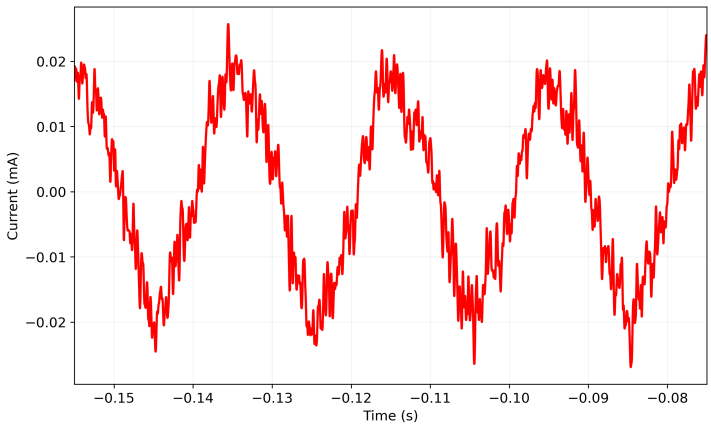

Figure 2.9: Reference resistive waveform at lowest applied voltage.

Reference and test waveform comparison identifies non-linearities arising in the current waveform through shape distortion. Both waveforms are normalized to one period and resampled to 1000 points using linear interpolation. This resampling enables efficient circular cross-correlation while maintaining adequate waveform resolution:

$$R_{xy}[m] = \sum_{n=0}^{N-1} x[n]\, y[n-m]$$

where $R_{xy}[m]$ is the cross-correlation at lag m, $x[n]$ is the reference waveform, $y[n-m]$ is the test waveform shifted by m samples, and N is the number of samples (1000). The lag m^* that maximizes this correlation is used to phase-align the waveforms, and the normalized correlation coefficient:

$$\rho = \frac{R_{xy}[m^*]}{\sqrt{R_{xx}[0]\; R_{yy}[0]}}$$

quantifies the quality of this alignment, where $R_{xx}[0]$ and $R_{yy}[0]$ are the zero-lag autocorrelations. When correlation is poor, indicating significant waveform distortion that prevents meaningful shape comparison, the algorithm falls back to simpler zero-crossing alignment.

The scaled reference waveform (Figure 2.10) represents the current that would flow if the MOV maintained linear behaviour at the higher voltage. The subtraction process isolates this purely nonlinear component:

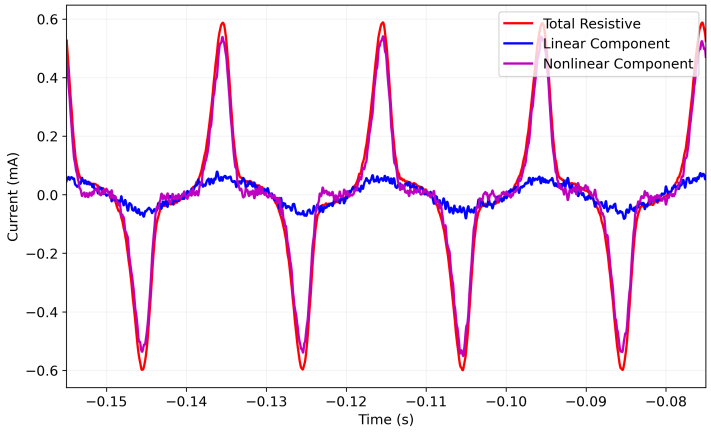

Figure 2.10: Isolated nonlinear resistive residual after reference subtraction.

The extraction of electrical parameters from these decomposed currents follows established circuit theory while respecting the measurement limitations. The total resistance derives from Ohm's law applied to peak values:

$$R_{\text{tot}} = \frac{v_{\text{peak}}}{i_{\text{res,peak}}},$$

where $i_{\text{res,peak}}$ is the peak value of the resistive current. Modelling the MOV as two parallel conduction paths—one linear R_{lin} and one nonlinear R_{nl}—allows the nonlinear resistance to be calculated from:

$$\frac{1}{R_{\text{tot}}} = \frac{1}{R_{\text{lin}}} + \frac{1}{R_{\text{nl}}}, \qquad R_{\text{nl}} = \frac{R_{\text{lin}}\, R_{\text{tot}}}{R_{\text{lin}} - R_{\text{tot}}},$$

This calculation is valid only when $R_{\text{tot}} < R_{\text{lin}}$.

For subsequent analysis and visualization, measurements where the nonlinear resistance exceeds 900 MΩ are excluded, as the corresponding nonlinear currents fall below the measurement floor of the system. These high impedances correspond to small resistive current magnitudes. Similarly, capacitive reactance values above this threshold indicate unreliable capacitance extraction. These limits ensure that only measurements with sufficient signal-to-noise ratio contribute to the characterisation results.

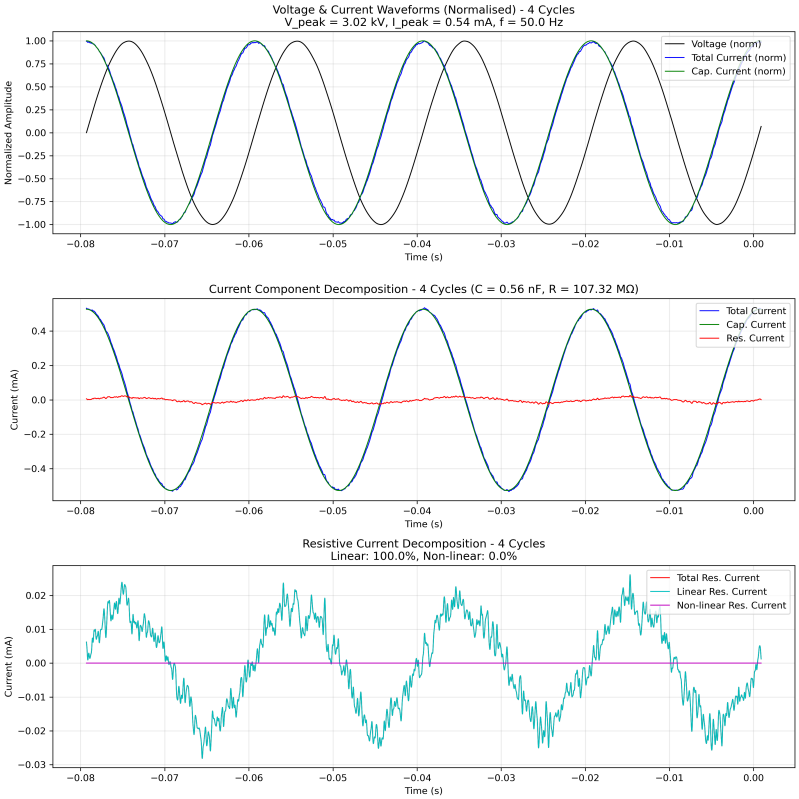

(a) Snapshot of an MOV block at 50 Hz voltage signal being decomposed with the automated analysis pipeline developed above

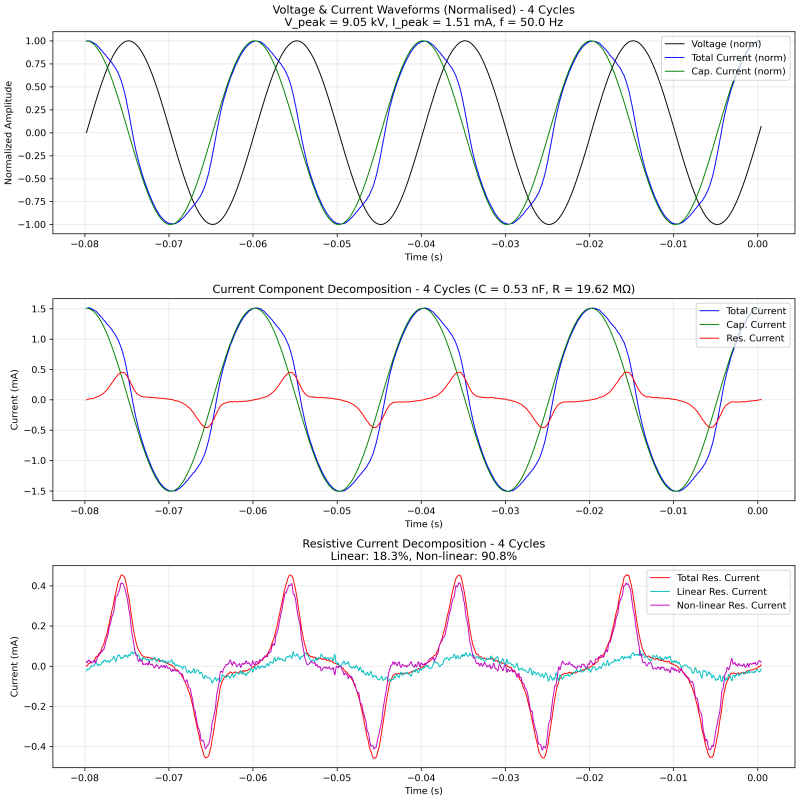

(b) The same MOV block at 50 Hz, but at a higher voltage point, where the conduction knee (in blue) is visible

Figure 2.12 presents the complete signal processing pipeline as a unified workflow, illustrating how raw oscilloscope measurements are systematically transformed into validated electrical parameters derived earlier.

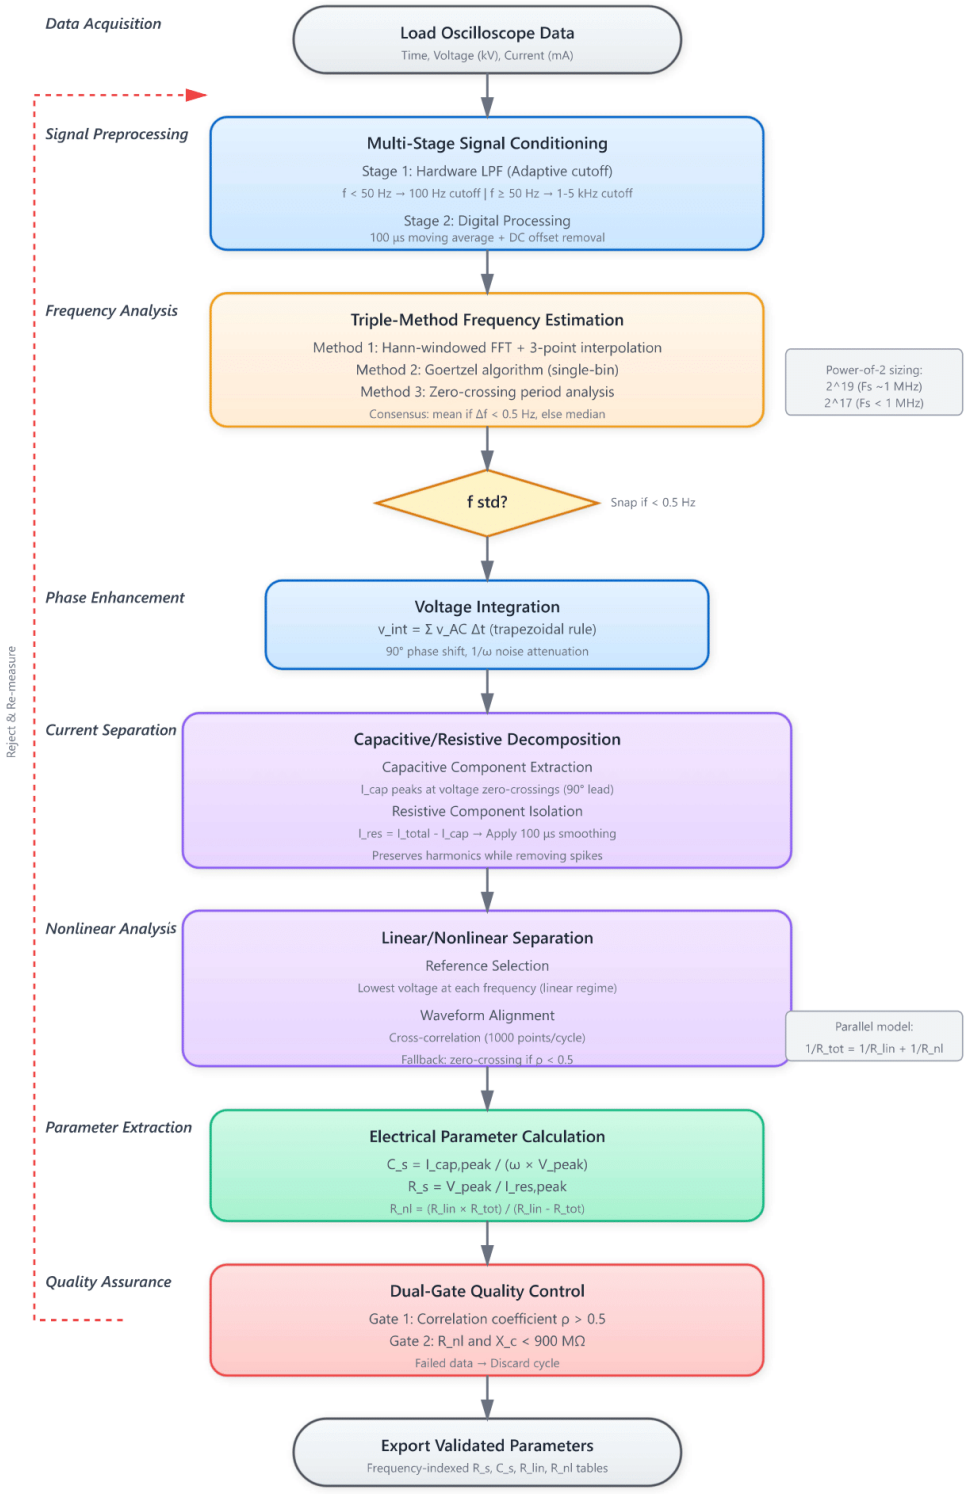

Figure 2.12: MOV analysis flowchart

2.3. Results

The experimental characterisation yields three key parameters across frequency and voltage domains. From the results discussed in this section, it can be observed that the capacitance measurements reflect grain boundary barrier geometry, linear residual resistance exhibits frequency-dependent dielectric relaxation, and non-linear resistance tracks voltage-dependent barrier modulation.

Capacitance Observations

MOV capacitance is frequency-independent but voltage-dependent near rated voltage (Figure 2.13). For TOV harmonic analysis, a single capacitance value suffices. Sub-harmonic frequencies require separate consideration.

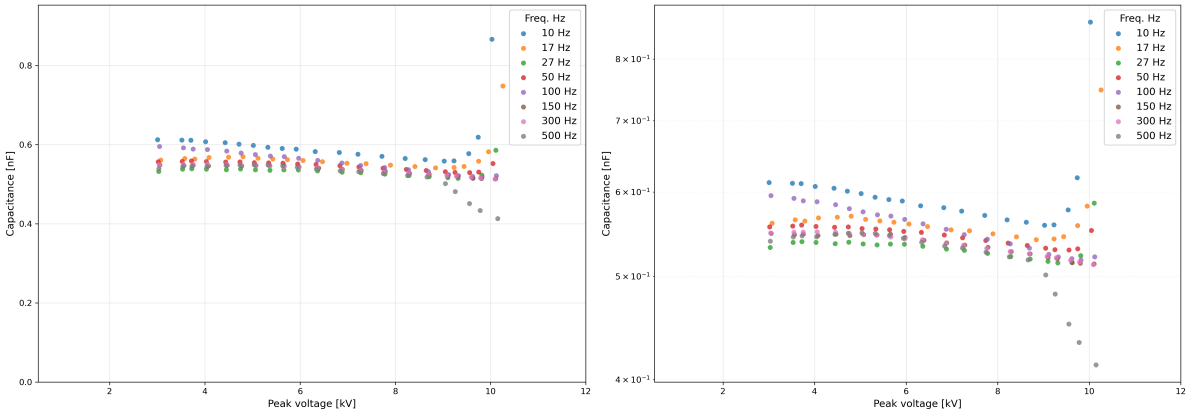

(a) Linear view **(b)** Logarithmic view

Figure 2.13: Extracted Capacitance C_s of the MOV as a function of peak applied voltage

Across the range from 10 Hz to 500 Hz, capacitance values cluster tightly between 0.52 nF and 0.6 nF—a 13 % variation. This stability indicates geometric capacitance, set primarily by the grain-boundary geometry with minor contributions from frequency-dependent polarisation [26].

The slight fall from 0.6 nF at 10 Hz to 0.52 nF at 500 Hz hints at residual relaxation [26], yet its magnitude is small. Either the relaxation lies well below our range, or modern varistors minimise it.

The capacitance shows a different behaviour with electric field. For peak voltages below 7 kV ($\approx 0.7\,U_\mathrm{r}$), C_s is essentially flat at $\sim$0.6 nF, indicating that the double-Schottky barriers (DSB) at the ZnO grain boundaries remain intact; the field is too weak to narrow the depletion width appreciably. Once the bias exceeds 8 kV the depletion region contracts and previously inaccessible interface states begin to store charge. Consequently C_s rises steeply, reaching 0.9 nF at 10 kV—a 50 % uplift that cannot be explained by geometry alone and therefore is evidence of barrier modulation [26].

With the increase in bias voltage, two effects combine to increase capacitance. First, the depletion regions at grain boundaries narrow under forward bias, effectively reducing the distance between the “plates” of each microscopic capacitor. The relationship $C = \varepsilon A/d$ dictates

that halving the separation doubles the capacitance. Second, and perhaps more importantly, interface states that were energetically inaccessible at low bias become available for charge storage as the barrier height decreases. These states, associated with defects and impurities at grain boundaries, can trap and release charges on the timescale of our measurements, contributing additional capacitance that manifests as the observed voltage dependence [26].

Resistance Observations

Resistance spans three decades in magnitude and almost two decades in frequency, highlighting how voltage and frequency jointly govern conduction.

Below 6 kV the logarithmic plots show resistance is strongly ordered by frequency. At any given voltage, resistance decreases monotonically with increasing frequency—10 Hz measurements yield resistances 3 to 5 times higher than their 500 Hz counterparts.

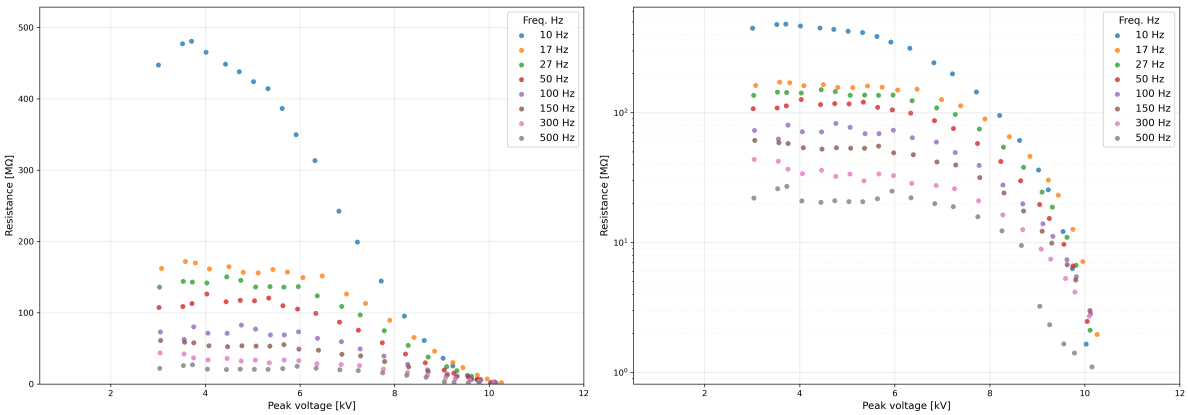

(a) Linear view **(b)** Logarithmic view

Figure 2.14: Extracted Resistance R_s of the MOV as a function of peak applied voltage

Classical dielectric-relaxation theory [27] gives

$$R_{\text{eff}} = \frac{\rho}{1 + \omega^2 \tau^2}, \tag{2.1}$$

so R is large when $\omega\tau \ll 1$ and falls as $\omega\tau$ approaches unity.

The anomalously high resistance values at 10 Hz approach the limits of the test set and show more scatter than higher-frequency data. At 10 Hz the 100 ms half-cycle matches the longest relaxation times, leaving residual charge that opposes the next half-cycle and inflates the apparent resistance.

On the semi-log plot (Figure 2.14b) each frequency trace is almost flat below 6 kV, indicating high, field-independent resistance in the pre-conduction region [26]. Above that threshold the traces descend exponentially, and their vertical separation narrows until all frequencies converge at 10 kV.

The transition at 6 kV indicates a change in conduction mechanism. Above this voltage, resis-

tance plummets exponentially, following the field-enhanced emission relationship [26]:

$$R \propto \exp\left(-\beta V^{1/2}\right) \tag{2.2}$$

where $\beta = \left[\frac{e^3}{4\pi t \varepsilon_0 \varepsilon_r}\right]^{1/2}$ is the field lowering coefficient dependent on the barrier thickness t and dielectric constant ε_r.

This field-enhanced conduction arises from Schottky barrier lowering at grain boundaries. The applied field progressively reduces the effective height of the double Schottky barriers, increasing the thermionic emission probability as $\exp(\beta V^{1/2})$, confirming that grain boundary barriers dominate high-field current flow in this operating region of the MOV [26].

All frequency traces in Figure 2.14b converge as the voltage approaches 10 kV. At low voltage each frequency probes a different part of the relaxation spectrum, producing dispersion; once non-linear emission dominates, barrier modulation overwhelms relaxation and the curves merge.

In the low-field regime, thermionic emission over grain-boundary barriers is modulated mainly by temperature, while interface-state relaxation and Maxwell-Wagner polarisation—the charge accumulation at interfaces between conductive grains and resistive grain boundaries—add frequency-dependent loss. The relaxation times range from microseconds to seconds, explaining why dispersion persists up to the knee voltage. [26].

The knee voltage can be defined as the point where the current waveform first shows visible deviation from a pure sinusoid due to the emergence of resistive current peaks near voltage maxima. The centre panel of Figure 2.11a shows oscilloscope trace below knee (pure capacitive sinusoid) while in Figure 2.11b, above knee (pronounced resistive peak) trace can be seen. As the field rises, DSB lowering and then tunnelling through thinned barriers successively take over, giving the exponential drop in R and the high-field convergence [26].

Between 50 Hz and 300 Hz the linear resistance falls by about a factor of three, so harmonic currents dissipate disproportionately more heat. Models that assume frequency-independent resistance under-estimate thermal stress in harmonics.

The onset of nonlinear conduction at $\approx 0.6\,U_r$ marks the beginning of varistor action. Below this threshold, the arrester behaves as a stable, if lossy, insulator. The safe continuous operating margin extends up to U_c ($\approx 0.8\,U_r$). Above this point, positive feedback develops—current causes heating that reduces barrier heights, enabling more current flow and risking thermal runaway.

Separating R_s into linear and non-linear components clarifies these behaviours. As shown in Figure 2.15, the non-linear branch is frequency-independent, confirming its origin in instantaneous voltage-controlled barrier modulation.

The linear branch drops from about 450 MΩ at 10 Hz to 25 MΩ at 500 Hz, just as dielectric-loss theory for a broad relaxation-time spectrum predicts [26]. Its flat voltage dependence confirms it is an intrinsic material property.

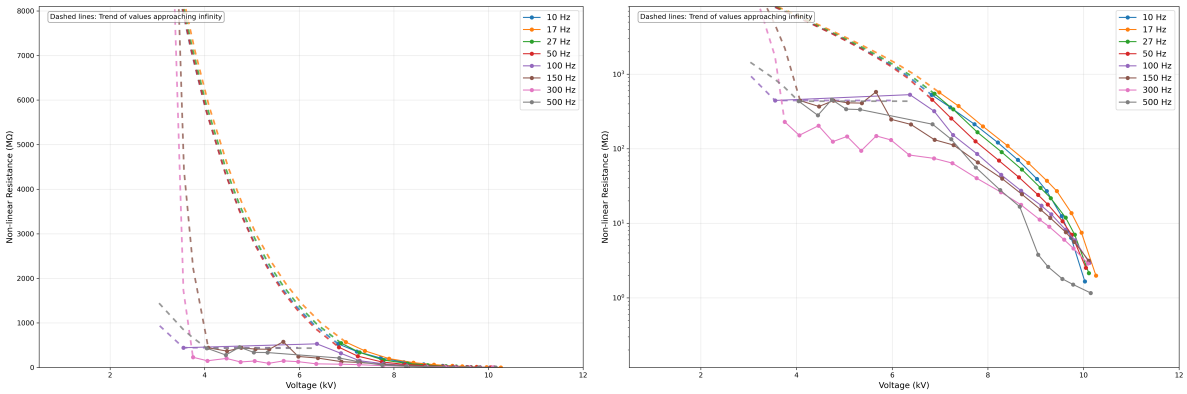

(a) Linear view **(b)** Logarithmic view

Figure 2.15: Non-linear resistance $R_{\text{non-linear}}$ of the MOV as a function of peak applied voltage

Below 5 kV the non-linear current sinks beneath the noise floor, so the curve is dashed; this threshold (0.5–$0.6\,U_r$) marks the onset of varistor action. Above it $R_{\text{non-lin}}$ collapses by three decades between 6 kV and 10 kV, demonstrating non-linearity and electronic, frequency-independent nature.

Subtracting the non-linear branch leaves the residual linear resistance (Figure 2.16); it shows strong frequency dependence but negligible voltage dependence—the mirror image of the non-linear branch.

A lossy dielectric has $\varepsilon^* = \varepsilon' - j\varepsilon''$; because $\sigma_{\text{eff}} = \sigma_0 + \omega\varepsilon''(\omega)$, $R_{\text{lin}} \propto 1/\sigma_{\text{eff}}$ and inevitably falls with frequency.

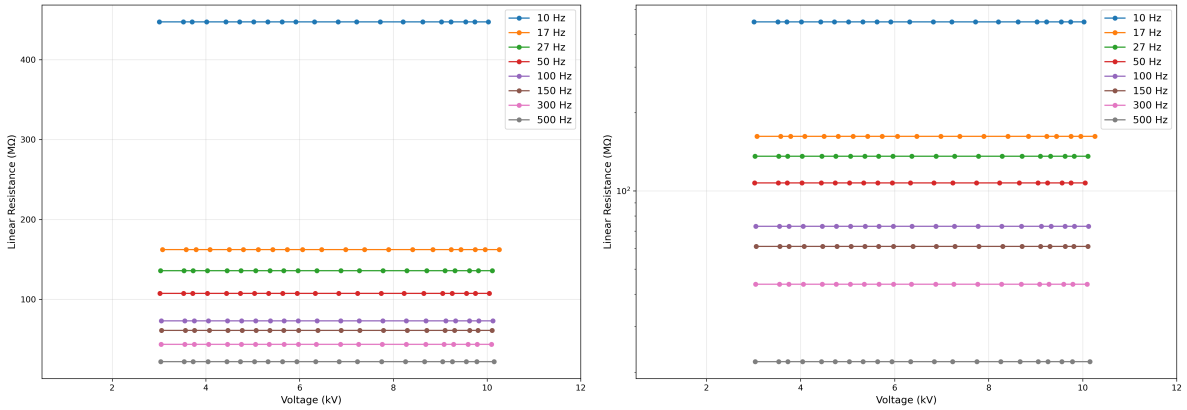

(a) Linear view **(b)** Logarithmic view

Figure 2.16: Residual (linear) resistance R_{linear} of the MOV as a function of peak applied voltage

The voltage-current characteristics reveal non-linear coefficients (α) in the range of 25-45, typical for modern surge arresters [28]. The slight frequency dependence of the knee voltage—lower frequencies showing marginally lower knee voltages—occurs because capacitive current magnitude scales with frequency ($Ic \propto \omega C$). As current increases and non-linear conduction dominates, these frequency effects vanish, confirming the successful separation of linear and non-linear mechanisms.

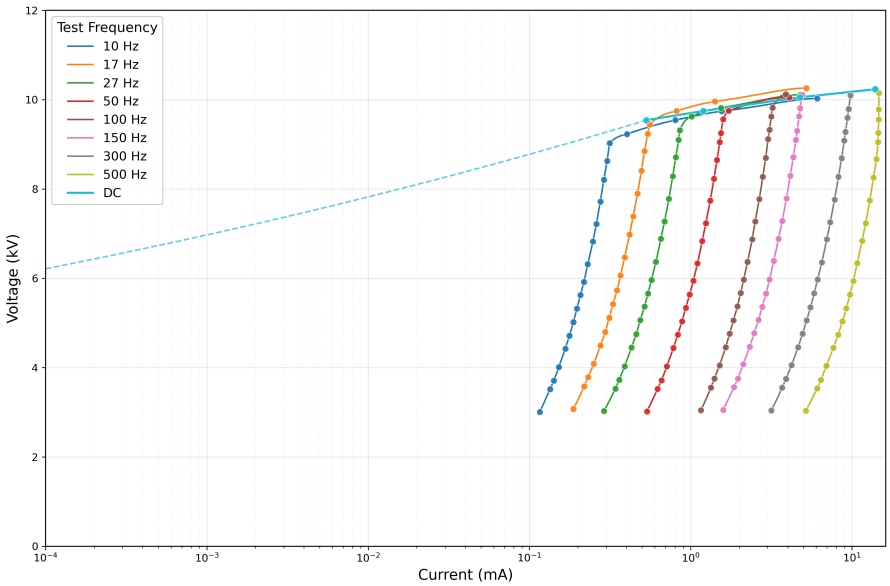

Figure 2.17: MOV V-I characteristics in the leakage current region.

2.4. Summary

This chapter's systematic characterisation of an MOV block across nine frequencies (DC-500 Hz) and twenty-one voltage levels (3 kV–10 kV peak) reveals four insights:

First, capacitance remains frequency-stable (13 % variation) but strongly voltage-dependent (50 % increase approaching protection level), validating power frequency testing while requiring multi-point voltage characterisation.

Second, resistance behaviour separates into distinct regimes: total resistance varies three orders of magnitude with voltage versus one with frequency. Decomposition reveals frequency-independent nonlinear resistance and strongly frequency-dependent linear resistance (twenty-fold decrease from 10 Hz–500 Hz).

Third, the linear resistance frequency dependence confirms grain boundaries act as lossy dielectrics with distributed relaxation times, challenging conventional frequency-independent models.

Fourth, the measurement methodology successfully extracts parameters across five decades of impedance through adaptive filtering, phase-sensitive decomposition, and correlation-based separation.

These results establish that surge arresters transition from lossy dielectric behaviour at low fields to electronic conduction at high fields. The frequency independence of nonlinear conduction confirms its electronic nature, while linear resistance frequency dependence reveals polycrystalline relaxation processes. This experimental foundation guides subsequent model development, particularly motivating fractional-order impedance models to represent the observed distributed relaxation phenomena.

3

Developing a Frequency Dependent Surge Arrester Model

This chapter addresses the following questions: How can the continuous spectrum of relaxation times inherent in heterogeneous materials be mathematically represented? This is discussed in Section 3.1 and subsequently in Section 3.2 where the Cole-Cole framework is presented as the mathematical approach, while Section 3.3 examines how fractional-order impedance behaviour can be realised using conventional circuit elements. Section 3.4 addresses practical implementation in ATP-EMTP, including details of component arrangement and nonlinear characteristic extraction. Through these sections, a circuit model is developed that represents the physics of an MOV block.

3.1. The Foundation: Understanding Frequency-Dependent Losses

Characterisation studies reveal that while surge arrester capacitance and nonlinear resistance remain stable across frequencies, dielectric losses vary significantly. These frequency-dependent losses arise from charge transport across grain boundaries and affect model accuracy.

In polycrystalline ZnO varistors, grain boundaries create potential barriers for charge carriers. Under alternating fields, charges accumulate during one half-cycle and must traverse or discharge during the opposite half-cycle. This finite transport time creates phase lag between voltage and current. With millions of boundaries having varying heights, widths, and conductivities, the result is a distributed spectrum of relaxation times manifesting as frequency-dependent impedance.

The experimental characterisation separates total current into capacitive and resistive components through phase decomposition. This separation reveals that the surge arrester behaves as a parallel combination of frequency-dependent resistance (representing dielectric losses) and capacitance (from grain boundary polarisation). For Cole-Cole analysis, the total impedance of this parallel combination is required, as shown in Equation 3.1.

For the parallel combination of frequency-dependent resistance R_4 (of Figure 1.9) and block capacitance C_0, the complex impedance Z yields the following real and imaginary components:

$$Z_{\text{real}} = \frac{R_4 \cdot X_c^2}{R_4^2 + X_c^2}, \quad Z_{\text{imag}} = -j\frac{R_4^2 \cdot X_c}{R_4^2 + X_c^2} \tag{3.1}$$

where R_4 is the frequency-dependent parallel resistance and $X_c = 1/(\omega C_0)$ is the capacitive reactance magnitude. The negative imaginary component indicates capacitive behaviour. Even with constant C_0, both the real and imaginary components exhibit complex frequency dependence through the interaction between R_4 and X_c.

This charge accumulation and relaxation behaviour distinguishes surge arresters from simple resistors. In their normal operating state—below the conduction threshold—surge arresters behave as lossy dielectrics rather than conductors. The applied electric field polarises the grain boundaries without allowing significant charge flow through the material. Energy is stored in the polarisation field and partially dissipated through relaxation processes, not through direct conduction. This dielectric behaviour dominates until the electric field becomes strong enough to lower the grain boundary barriers sufficiently for electronic conduction to begin.

The frequency dependence of this dielectric loss is observed through impedance measurements. For a pure resistor, impedance would remain constant with frequency. For an ideal capacitor, impedance would decrease proportionally to the inverse of the applied frequency, f^{-1}. Surge arresters show a fractional power law, requiring the fractional-order impedance models discussed in subsequent sections.

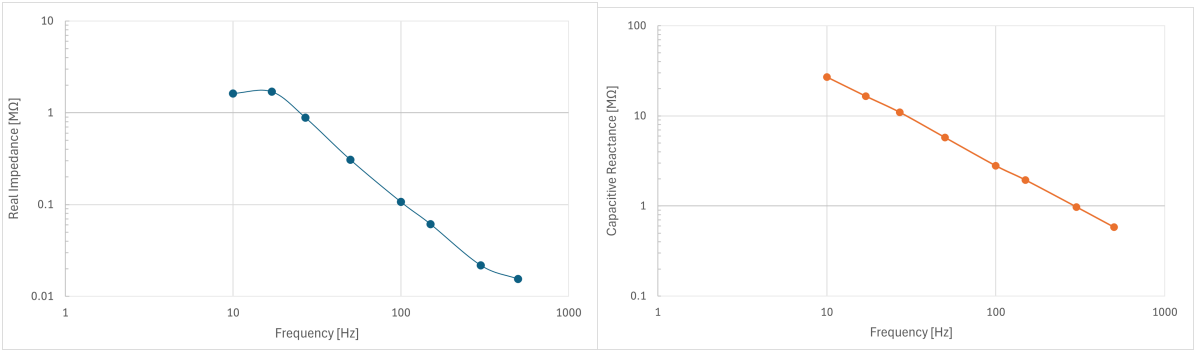

(a) Real impedance Z_{real} vs f

(b) Imaginary impedance $|Z_{\text{imag}}|$ vs f

Figure 3.1: Measured impedance components at 5.5 kV showing power-law frequency dependence characteristic of distributed relaxation processes.

The voltage level of 5.5 kV was selected for analysis as it represents the midpoint of the measurement range (3 kV to 10 kV) and marks the beginning of measurable nonlinear current components. The measured characteristics in Figure 3.1 show power-law relationships on logarithmic scales. The real impedance exhibits a plateau at low frequencies, decreases through the middle frequency range, and approaches saturation at the highest measured frequencies. The imaginary impedance magnitude (capacitive reactance) decreases monotonically with frequency, as expected for capacitive behaviour.

The real impedance behaviour reflects the transition from low-frequency response dominated by slow grain boundary relaxation processes to high-frequency response where these processes cannot follow the applied field. The capacitive component remains relatively constant, with the observed decrease in reactance magnitude arising from the $1/\omega$ relationship. The interaction between frequency-dependent resistance R_4 and capacitive reactance produces the complex impedance behaviour observed in the measurements.

The frequency-dependent loss tangent $\tan\delta = Z_{\text{real}}/|Z_{\text{imag}}|$ increases substantially at lower frequencies. Surge arresters therefore dissipate proportionally more energy under low-order harmonics. During resonant overvoltages in cable networks, where second and third harmonics reach significant magnitudes, fixed resistance models underestimate energy dissipation.

3.2. The Cole-Cole Framework: Capturing Distributed Relaxation

Real dielectrics produce depressed semicircular arcs in complex permittivity plots rather than the perfect semicircles predicted by Debye theory [29], indicating distributed rather than single relaxation times. In MOVs, the grain boundary network creates energy barriers with varying heights and widths, each with its own relaxation time.

The Cole-Cole model incorporates distributed relaxation through a fractional exponent:

$$\frac{\varepsilon^*(\omega) - \varepsilon_\infty}{\varepsilon_0 - \varepsilon_\infty} = \frac{1}{1 + (j\omega\tau_0)^{1-\alpha}} \tag{3.2}$$

The parameter α quantifies distribution breadth. When $\alpha = 0$, single Debye relaxation is recovered; as α approaches unity, the distribution broadens. MOVs typically exhibit α between 0.5 and 0.7 and the capacitance for dielectric capacitors is given by [29]:

$$C^*(\omega) = C_\infty + \frac{C_0 - C_\infty}{1 + (j\omega\tau_0)^{1-\alpha}} \tag{3.3}$$

High-frequency capacitance C_∞ represents instantaneous response from electronic and atomic polarisation [29]. Low-frequency capacitance C_0 includes interfacial polarisation at grain boundaries, manifesting when fields change slowly enough for charge accumulation. The characteristic time τ_0 centres the relaxation spectrum, with characteristic frequency $f_0 = 1/(2\pi\tau_0)$.

Complex impedance relates to complex capacitance as:

$$Z^*(\omega) = \frac{1}{j\omega C^*(\omega)} \tag{3.4}$$

Substituting Cole-Cole expressions and separating components yields frequency-dependent resistance and reactance following established patterns for fractional impedances [29]:

1. For $\omega \gg \frac{1}{\tau_0}$, the capacitance equals C_∞.
2. For $\omega \ll \frac{1}{\tau_0}$, the capacitance equals C_0.
3. Low-frequency polarisation mechanisms ensure that C_0 exceeds C_∞.
4. The characteristic angular frequency $\frac{1}{\tau_0}$ lies close to the minimum of the impedance-ratio $|X|/R$.
5. On a log–log scale, R varies with a slope of $-\alpha$ for $\omega \ll \frac{1}{\tau_0}$.
6. For $\omega \gg \frac{1}{\tau_0}$, the slope of R changes to $-(2-\alpha)$.
7. Increasing the distribution parameter α increases the phase angle (or X/R ratio) at $\omega = \frac{1}{\tau_0}$.
8. As C_0 nears C_∞, the overall resistance decreases.

These observations allow parameter extraction from measured impedance data. Initial estimates derive from asymptotic behaviour: C_∞ from the high-frequency capacitance plateau, C_0 from the low-frequency plateau, and α from the low-frequency resistance slope. The characteristic time τ_0 is estimated from the frequency where either the loss tangent peaks or the reactance-to-resistance ratio minimises.

Refinement uses nonlinear optimisation, minimising the weighted sum of squared errors:

$$\chi^2 = \sum_i \left[\frac{(R_{eq,model} - R_{eq,meas})^2}{\sigma_{R,i}^2} + \frac{(|X_{eq,model}| - |X_{eq,meas}|)^2}{\sigma_{X,i}^2} \right] \tag{3.5}$$

The optimisation enforces physical constraints: positive capacitances, $C_0 > C_\infty$, positive relaxation time, and $0 < \alpha < 1$. Differential Evolution [30] algorithm is used for its global search capability in complex parameter spaces.

The use of an evolutionary algorithm is required when examining the error surface for Cole-Cole parameter extraction. Figure 3.2 reveals multiple valleys and ridges in the objective function, particularly in the τ_0-α parameter space when gauging the eight observations made earlier. These local minima would trap gradient-based optimisers, yielding suboptimal fits that misrepresent the surge arrester's frequency-dependent behaviour. The population-based DE algorithm explores multiple valleys simultaneously, ensuring convergence to the global minimum.

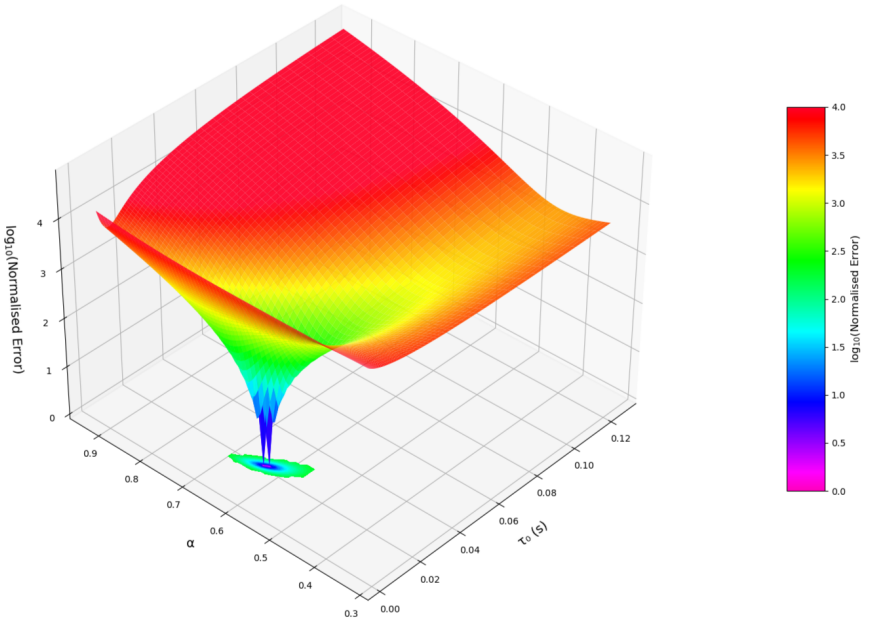

Figure 3.2: Error surface of Cole-Cole parameters of the MOV showing multiple valleys.

Application to surge arrester impedance data at 5.5 kV yields: $C_\infty = 0.523$ nF (instantaneous polarisation of ZnO grains), $C_0 = 0.692$ nF (includes 0.169 nF from grain boundary polarisation).

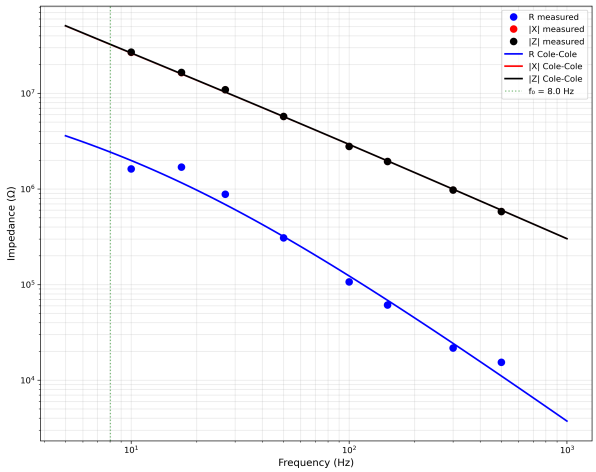

Figure 3.3: Cole-Cole model fit to measured surge arrester impedance data at 5.5 kV.

The characteristic time $\tau_0 = 19.8$ ms corresponds to $f_0 = 8.0$ Hz, positioning the loss peak below the measurement range. Measurements capture the transition region where slope changes most rapidly. The distribution parameter $\alpha = 0.626$ indicates a broad distribution of relaxation times.

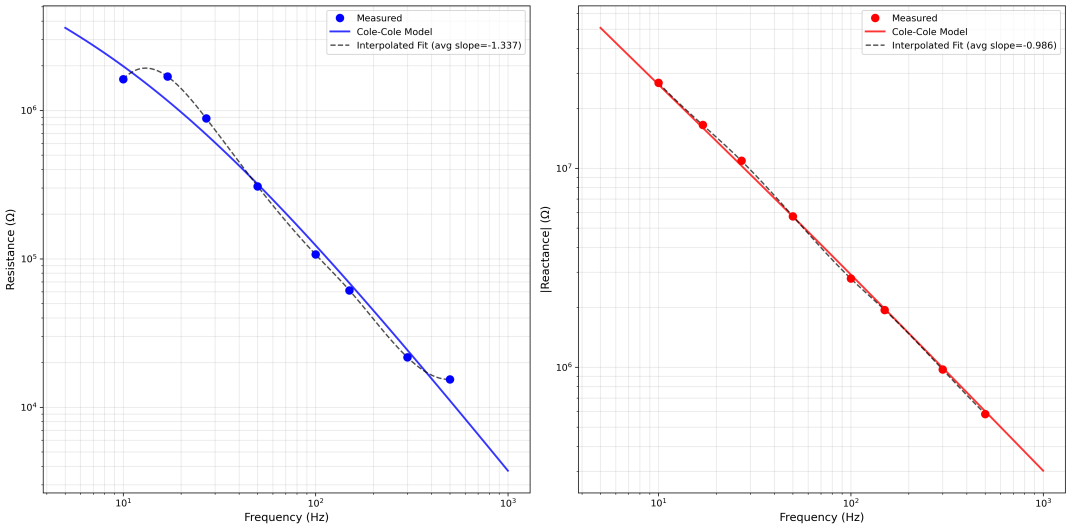

Figure 3.4: Separated resistance and reactance components with Cole-Cole model fit.

Figure 3.4 shows the separated components. Resistance follows power law across three frequency decades. Reactance maintains slope near -1 (theoretical capacitor value), with deviations from frequency-dependent losses.

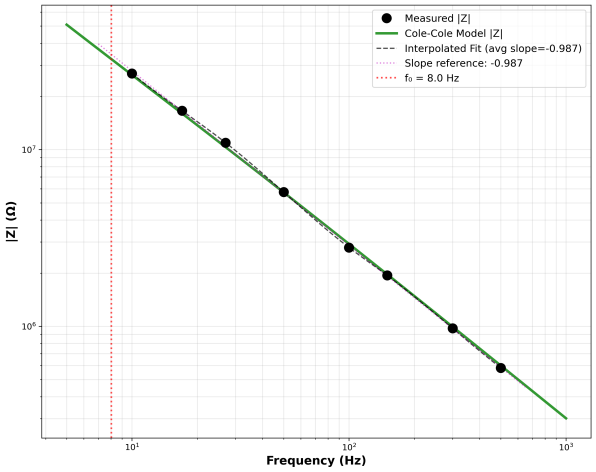

Figure 3.5: Impedance magnitude versus frequency on log-log scale showing measured data and Cole-Cole model.

The impedance magnitude (Figure 3.5) follows power law with average slope -0.987, near the ideal capacitive value -1. This near-capacitive behaviour reflects reactive dominance over resistive components, while deviation from -1 captures frequency-dependent loss influence. Figure 3.6 summarises the parameter extraction algorithm.

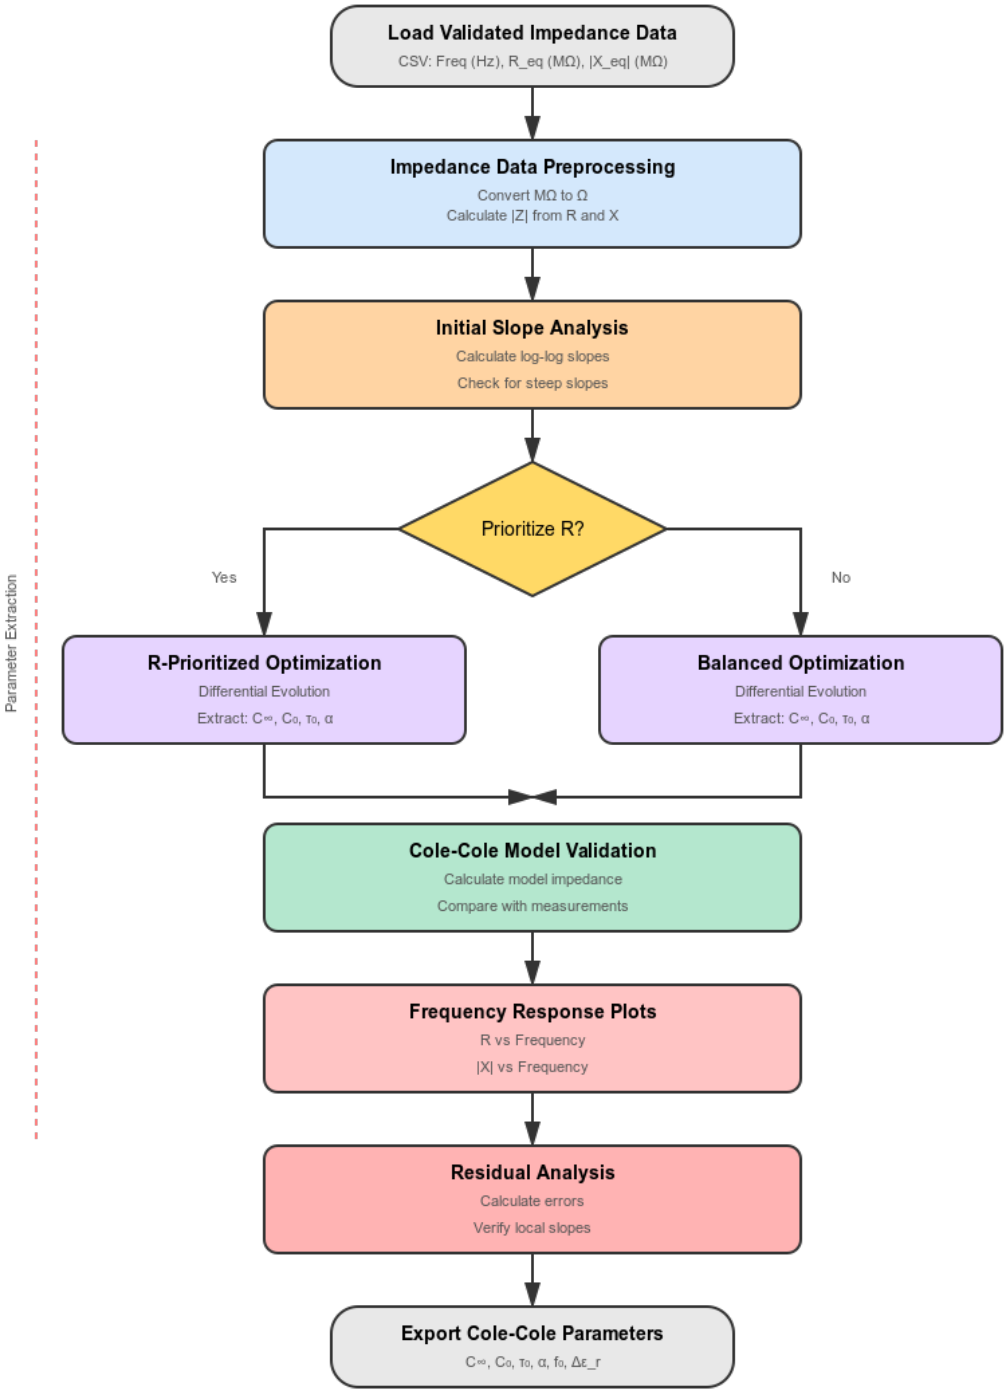

Figure 3.6: Flowchart for extracting Cole-Cole parameters

3.3. From Fractional Calculus to Realisable Circuits

The Cole-Cole model's fractional exponent $(j\omega\tau_0)^{1-\alpha}$ is not implemented with ideal resistors and capacitors—any RC network produces only rational functions with integer powers. One approach uses approximation: a fractional-order impedance can be approximated by a rational function with sufficient terms. On a logarithmic frequency scale, fractional impedance $Z_A \propto (j\omega)^{-\alpha}$ appears as a straight line with slope $-\alpha$. Since resistors produce horizontal lines (slope 0) and capacitors produce lines with slope -1, any intermediate slope can be approximated by alternating between resistive and capacitive behaviour.

This leads to the Oustaloup approximation method [31] as visualized in Figure 3.7. The method constructs a zigzag pattern following the ideal fractional impedance. Each corner in this zigzag corresponds to a pole frequency of one RC branch, where impedance transitions from resistive to capacitive behaviour.

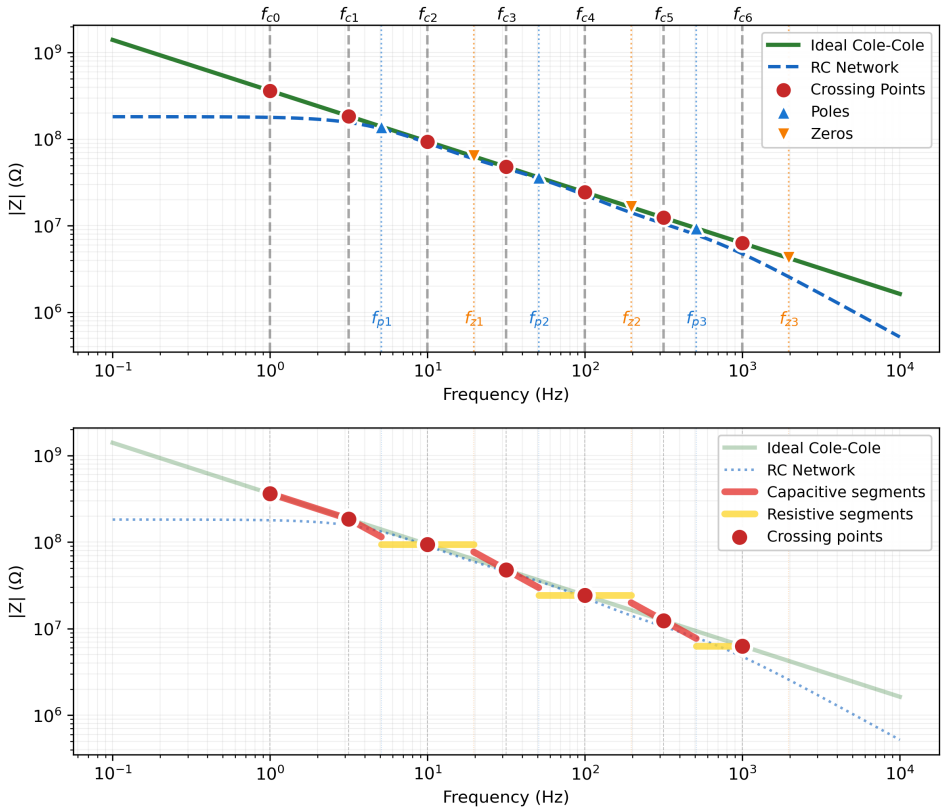

Figure 3.7: Piecewise linear approximation of fractional impedance Z_A on log-log scale. Green curve: ideal fractional impedance with slope $-\alpha$. Red segments: capacitive behaviour (slope -1). Yellow segments: resistive behaviour (slope 0). Blue dashed curve: synthesized RC network response. Crossing points f_{ci} mark where approximation intersects the ideal curve.

The crossing frequencies—where the zigzag intersects the ideal curve—are distributed logarithmically:

$$f_{ci} = f_0 \left(\frac{f_1}{f_0}\right)^{\frac{i}{2N}}, \quad i = 0, 1, 2, ..., 2N \tag{3.6}$$

where f_0 and f_1 define the frequency range of interest, and N is the number of RC branches. Logarithmic spacing ensures uniform approximation quality across all frequency decades.

To maintain average slope $-\alpha$, the horizontal distance (resistive portion) and diagonal distance (capacitive portion) must satisfy specific ratios. This geometric constraint yields:

$$f_{p,k} = f_0 \left(\frac{f_1}{f_0}\right)^{\frac{2k-1-\alpha}{2N}}, \quad f_{z,k} = f_0 \left(\frac{f_1}{f_0}\right)^{\frac{2k-1+\alpha}{2N}} \tag{3.7}$$

The ratio between each zero and preceding pole equals $(f_1/f_0)^{\alpha/N}$, maintaining consistent geometry throughout. The blue dashed curve in Figure 3.7 shows the frequency response of the synthesized RC network, calculated as

$$Z_{RC}(\omega) = \sum_{k=1}^{N} R_k/(1 + j\omega R_k C_k) \tag{3.8}$$

This response differs from the piecewise linear approximation due to the continuous nature of RC transfer functions. The RC network impedance matches the ideal fractional impedance at each crossing frequency f_{ci}, with deviations occurring between these points. The maximum error occurs midway between adjacent crossing frequencies, where the smooth RC response deviates from the zigzag path. For the frequency range f_0 to f_1, the approximation error decreases as the number of branches N increases. For this document, the approximation uses 3 branches, but a 5 branch model has also been discussed in [ADD REFERENCE].

The complete fractional element becomes:

$$Z'_A(s) = R_0 \frac{\prod_{k=1}^{N}(1 - s/\omega_{z,k})}{\prod_{k=1}^{N}(1 - s/\omega_{p,k})} \tag{3.9}$$

where R_0 normalises impedance at the geometric centre of the approximation range.

For circuit implementation, two approaches exist: Foster I and Foster II networks [29]. Foster I implements the impedance form through partial fraction expansion:

$$Z'_A(s) = R_\infty + \sum_{k=1}^{N} \frac{R_k}{1 + sR_kC_k} \tag{3.10}$$

where each term represents a parallel RC branch connected in series. Foster II implements the dual admittance form:

$$Y'_A(s) = G_\infty + \sum_{k=1}^{N} \frac{sC_k}{1 + sR_kC_k} \tag{3.11}$$

where each term represents a series RC branch connected in parallel. Foster II has the following characteristics for surge arrester modelling: The explicit parallel conductance G_∞ avoids numerical difficulties with very small values common in high-quality dielectrics and represents the DC resistance of the block. Each series RC branch can be interpreted as a conduction path through the grain boundary network. The parallel conductance accommodates temperature-dependent leakage current, relevant for thermal modelling.

For the extracted surge arrester parameters, the fractional element DC impedance is:

$$R_0 = \frac{\tau_0}{C_0 - C_\infty} = \frac{19.8 \times 10^{-3}}{0.170 \times 10^{-9}} = 116.5\,\text{M}\Omega \tag{3.12}$$

The admittance residues yield capacitance values through partial fraction expansion:

$$C_k = B_k = -\frac{1}{R_0} \cdot \frac{\prod_{m=1}^{N}(1 - \omega_{z,k}/\omega_{p,m})}{\omega_{z,k} \prod_{m \neq k}(1 - \omega_{z,k}/\omega_{z,m})} \tag{3.13}$$

Resistance values are determined by ensuring each branch resonates at its zero frequency:

$$R_k = \frac{1}{2\pi f_{z,k} C_k} \tag{3.14}$$

The high-frequency parallel conductance becomes:

$$G_\infty = \frac{1}{R_0} \prod_{k=1}^{N} \frac{\omega_{p,k}}{\omega_{z,k}} = 0.897\,\text{nS} \tag{3.15}$$

Table 3.1: Foster II Network Parameters for Surge Arrester Model

Branch	C_k (nF)	R_k (MΩ)	τ_k (ms)	$f_{z,k}$ (Hz)
1	0.027	77.6	2.131	74.7
2	0.019	23.7	0.459	346.7
3	0.020	5.0	0.099	1609.4
External Components:				
G_∞	0.897 nS (1114.6 MΩ)			
C_∞	0.522 nF			
C_s	0.170 nF			

The Foster network parameters result from the Oustaloup approximation of the Cole-Cole fractional impedance. Three branches were selected to approximate the fractional-order behaviour across the measurement frequency range of 10 Hz to 500 Hz, spanning approximately three decades. The Oustaloup method typically requires one RC branch per frequency decade for adequate approximation accuracy. The resulting logarithmically-spaced time constants

collectively approximate the continuous distribution of relaxation times in the grain boundary network. Branch 1 with time constant 2.131 ms contributes primarily to the low-frequency response below 100 Hz. Branch 2 with time constant 0.459 ms and zero frequency at 346.7 Hz covers the range relevant for power system harmonics. Branch 3 with time constant 0.099 ms and zero frequency at 1609.4 Hz extends the approximation beyond the measurement range to maintain accuracy at the upper frequency boundary. These branches work together to reproduce the fractional-order impedance behaviour arising from the distributed grain boundary relaxation processes captured by the Cole-Cole model.

The complete surge arrester model combines the fractional element approximation with the static capacitances from the Cole-Cole model:

$$Z_{total} = \left[\left(Z'_A \right)^{-1} + j\omega C_\infty \right]^{-1} + \frac{1}{j\omega (C_0 - C_\infty)} \tag{3.16}$$

At very low frequencies, the total capacitance approaches C_0, while at high frequencies it approaches C_∞. The frequency-dependent losses are represented throughout the transition region, which is relevant for TOV analysis.

The complete procedure developed to obtain the RC network parameters from the Cole-Cole parameters can be seen in Figure 3.8 in the next page.

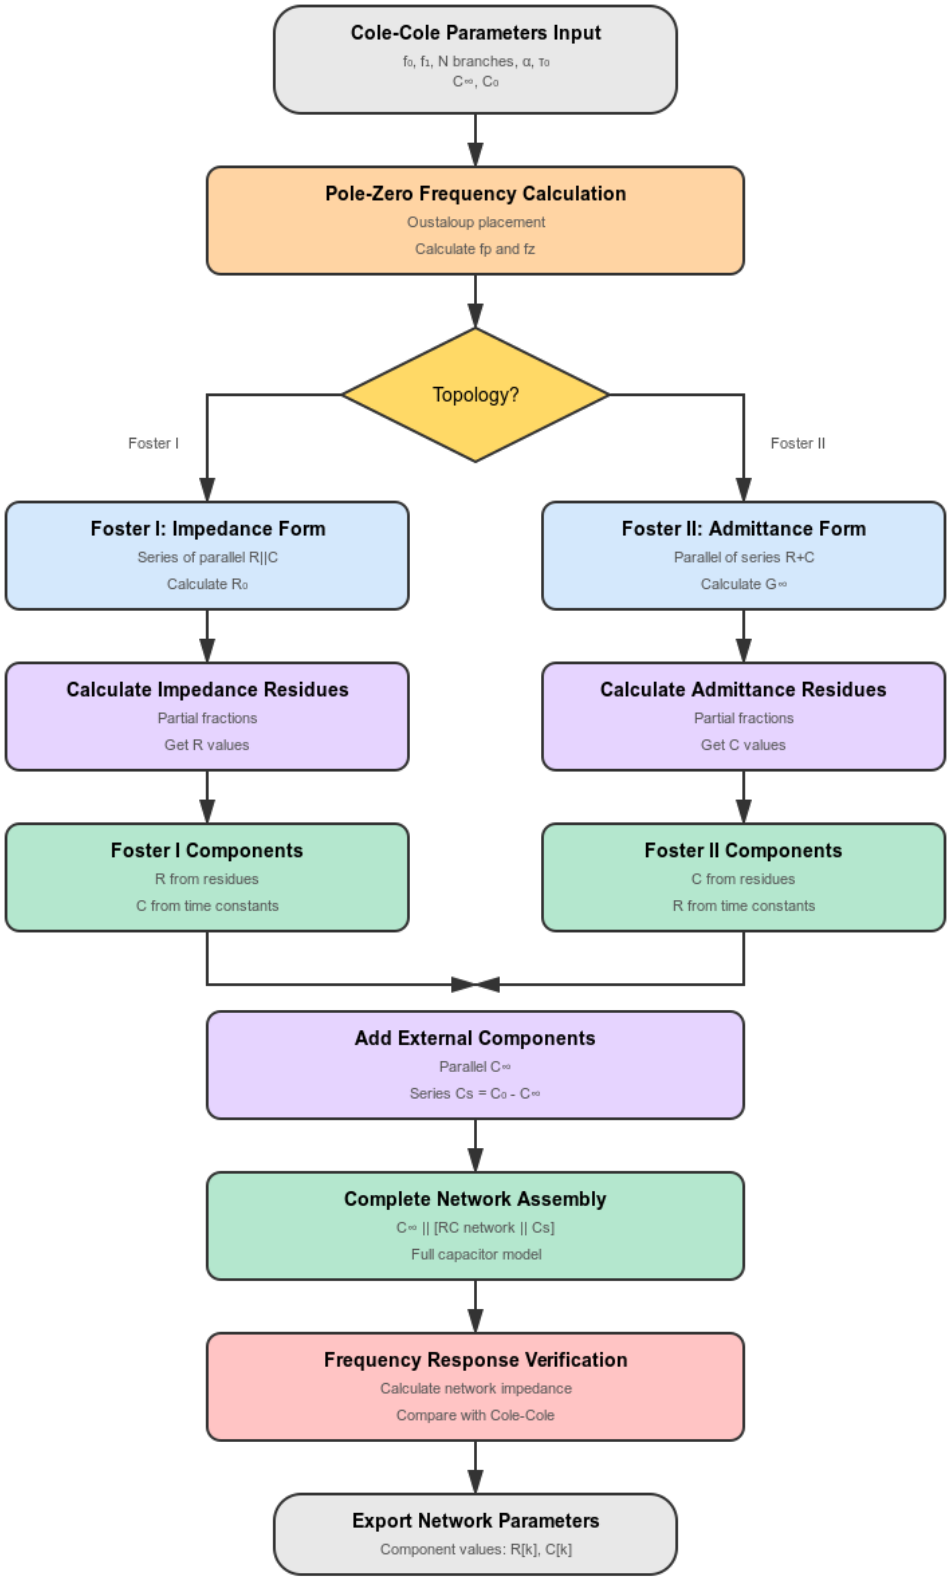

Figure 3.8: RC circuit parameter extraction flowchart

3.4. Implementation in ATP-EMTP

ATP-EMTP (Alternative Transients Program) is the industry-standard software for electromagnetic transient analysis in power systems, widely used for insulation coordination and protection studies. ATP was selected for implementing the frequency-dependent surge arrester model due to its capability to handle user-defined nonlinear characteristics and complex circuit topologies. The software enables time-domain simulation of the Foster network branches alongside the nonlinear V-I characteristic, essential for validating the model's response under both steady-state and transient conditions. ATP's built-in Type 92 non linear resistor element provides the framework for incorporating the frequency-dependent impedance network while maintaining computational efficiency required for system-level studies. The graphical preprocessor ATPDraw [32], developed by Dr. Hans Kristian Høidalen at NTNU, facilitates circuit construction and parameter adjustment during the iterative model refinement process.

3.4.1. Circuit Architecture and Component Arrangement

The frequency-dependent model implementation shown in Figure 3.9 requires consideration of component topology. The architecture represents a hierarchy of physical processes operating at different timescales.

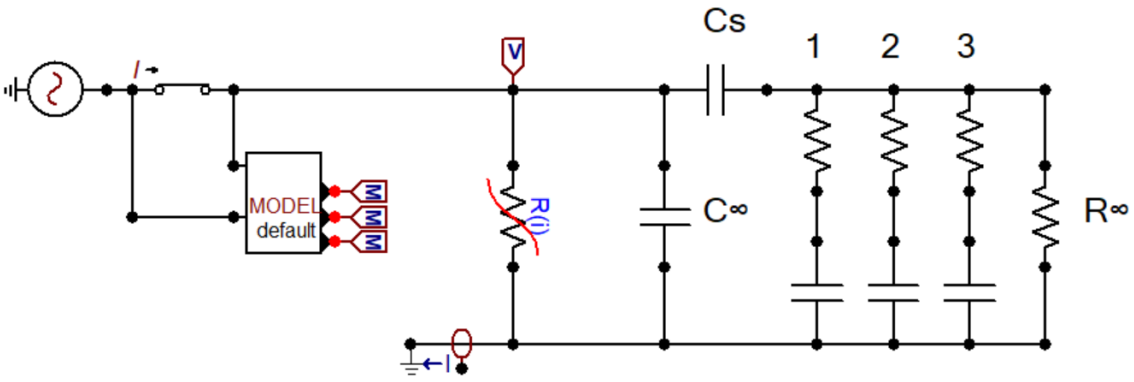

Figure 3.9: Frequency-dependent surge arrester model implemented in ATP-EMTP

At the input, the parallel capacitor $C_\infty = 0.522$ nF provides a path for high-frequency currents, representing the instantaneous polarisation mechanisms within ZnO grains. This capacitance responds without delay to voltage changes at any practical frequency. Following this, the series capacitor $C_s = 0.170$ nF introduces frequency-dependent impedance that increases at lower frequencies. Physically, this represents the additional polarisation capacity available when the electric field changes slowly enough for interfacial charges to accumulate at grain boundaries.

The frequency-dependent behaviour is implemented in the Foster II network: three series RC branches in parallel, supplemented by the high-frequency conductance G_∞. Each branch creates a frequency-dependent admittance:

$$Y_k = \frac{j\omega C_k}{1 + j\omega R_k C_k} \tag{3.17}$$

At frequencies well below $f_{z,k} = 1/(2\pi R_k C_k)$, the branch impedance is dominated by the resistance R_k, contributing minimal admittance. As frequency approaches $f_{z,k}$, the capacitive reactance becomes comparable to the resistance, and the branch begins conducting significantly. Above $f_{z,k}$, the branch admittance approaches $j\omega C_k$, contributing primarily reactive current.

The parallel conductance $G_\infty = 0.897$ nS represents a limiting case. Even when all capacitive paths are short-circuited at very high frequencies, this conductance provides a residual loss mechanism. Physically, it represents direct tunnelling or hopping conduction through grain boundaries that cannot be eliminated by bypassing capacitive effects.

The specific component ordering affects operation. C_∞ must be parallel to provide the high-frequency limiting capacitance. The series C_s must precede the Foster II network to ensure proper DC blocking; if placed after, the conductance G_∞ would provide a DC path that bypasses C_s. This ordering results in the following behaviour across all frequencies:

$$\lim_{\omega \to 0} C_{total} = C_0, \quad \lim_{\omega \to \infty} C_{total} = C_\infty \tag{3.18}$$

MOVs operating in their leakage region exhibit behaviour similar to dielectric absorption in conventional capacitors. Each grain boundary acts as a nanoscale capacitor with its own relaxation time determined by local barrier properties. The collective response of millions of such boundaries creates the distributed relaxation spectrum represented by the Cole-Cole model. The Foster II network provides a lumped-element approximation to this distributed behaviour, with each RC branch representing a subset of grain boundaries with similar relaxation times. The choice of three branches represents a balance between accuracy and complexity, with each branch covering one decade of frequency from 10 Hz to 10 kHz—covering power frequency, low-order harmonics, and switching transients.

3.4.2. Critical Implementation Details: Type 92 Element Selection

The choice between the MOV Type 92 block and the R(i) Type 92 element in ATP-EMTP requires consideration. The MOV Type 92 block implements a complete surge arrester model based on IEEE and IEC recommendations using exponential relationships:

$$i = p\left(\frac{v}{V_{ref}}\right)^q \tag{3.19}$$

where parameters p and q are derived from manufacturer data. This block includes automatic region switching to prevent numerical underflow, transitioning to linear approximation below a threshold. While beneficial for lightning and switching studies, this automatic behaviour creates complications for frequency-dependent modelling.

The MOV Type 92 block's internal impedance assumptions interfere with external RC networks. When parallel branches are added to create frequency-dependent behaviour, the built-in impedance characteristics dominate current distribution, masking the crafted frequency response. The automatic linear-to-exponential switching disrupts phase relationships required for harmonic analysis.

The R(i) Type 92 element, in contrast, implements a pure current-dependent resistance without hidden impedance effects or automatic switching. Users directly specify the voltage-current characteristic through data points, maintaining control over the nonlinear behaviour. The element includes an optional linear resistance R_{lin} in series, which must be set to zero for frequency-dependent applications to avoid interference with the external RC network. The high resistance values specified in the leakage region interact with the parallel RC branches to create the desired frequency-dependent behaviour.

3.4.3. Deriving the Nonlinear Characteristic

The voltage-current characteristic required for the R(i) Type 92 element derives directly from the characterisation performed in Chapter 2, details of which are provided in Appendix B. This measurement-based approach ensures the model represents the physical device across its entire operating range. The process begins with the nonlinear resistance values R_{nl} extracted through the linear-nonlinear decomposition technique. For each measurement voltage, the nonlinear current is calculated:

$$I_{nl} = \frac{V_{peak}}{R_{nl}} \tag{3.20}$$

Quality filtering removes data points where R_{nl} exceeds 900 MΩ, eliminating measurements dominated by noise rather than genuine nonlinear conduction and focuses the characterisation on the region where the arrester transitions from capacitive to resistive behaviour.

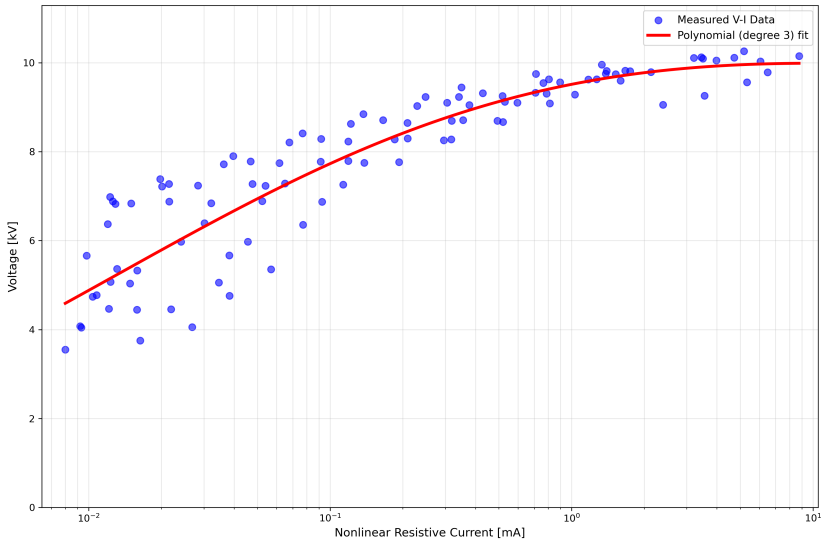

Figure 3.10: Polynomial fit of nonlinear current characteristic. Transition at 0.6-0.7 U_r marks change from capacitive to resistive behaviour. Red points: discretised Type 92 values.

The resulting voltage-current data exhibits the characteristic behaviour expected from MOVs. Below approximately 0.6 times the rated voltage, the nonlinear current remains negligible—the arrester operates in its capacitive region where dielectric behaviour dominates. The transition region between 0.6 and 0.7 U_r marks the onset of significant nonlinear conduction. To create a smooth characteristic suitable for numerical simulation, polynomial fitting is performed in logarithmic space:

$$\log(I) = a_3[\log(V)]^3 + a_2[\log(V)]^2 + a_1 \log(V) + a_0 \tag{3.21}$$

This transformation linearises the exponential relationship typical of MOV behaviour, allowing accurate fitting with low-order polynomials. The fitted curve is then transformed back to linear coordinates and discretised at voltage intervals suitable for the Type 92 element's interpolation algorithm. The scatter visible at low voltages has minimal impact on model accuracy since total arrester current is dominated by the capacitive component flowing through the Foster II network. The details of the extracted nonlinear characteristics can be found in Appendix B.

3.4.4. Validation Through Simulation

The frequency-dependent surge arrester model was implemented in ATP-EMTP using parameters derived from the 5.5 kV characterisation data. Single-phase simulations replicated the laboratory test conditions across the full frequency range from 10 Hz to 500 Hz. Simulation parameters balance accuracy with computational efficiency, using time step $\Delta t = 1 \times 10^{-6}$ seconds for adequate resolution up to several kilohertz. The simulation duration varies with frequency to capture steady-state behaviour.

The simulation results demonstrate frequency-dependent behaviour. As frequency decreases, the voltage-current ellipses broaden, indicating increasing losses. The phase angle between voltage and current changes systematically, transitioning from nearly 90° at high frequencies (capacitive) to smaller angles at low frequencies (increasingly resistive). This frequency-dependent phase shift is required for TOV energy calculations.

Deviations between model and measurement appear near the conduction knee, with the model initiating significant conduction at voltages 3 %–8 % higher than observed experimentally. These deviations stem from limitations in the ATP-EMTP surge arrester model, which uses simplified hysteresis implementation assuming fixed-width loops rather than the frequency-dependent loop shapes observed in real devices. Temperature effects represent another limitation. The model assumes constant temperature, neglecting the dynamic heating that occurs during conduction. In reality, even small resistive currents cause localised heating at grain boundaries, reducing barrier heights and increasing conductivity [26]. This positive feedback mechanism enhances conduction, particularly during the extended durations typical of TOVs.

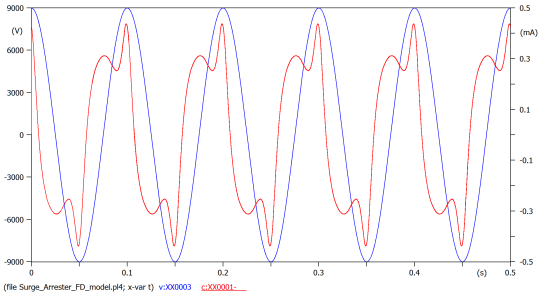

(a) 10 Hz

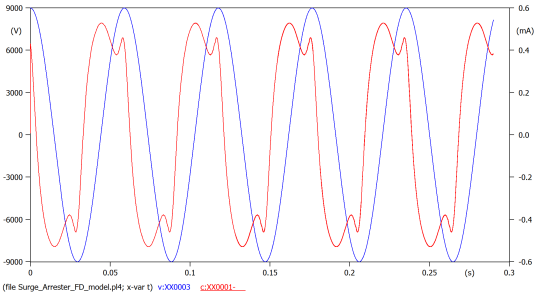

(b) 17 Hz

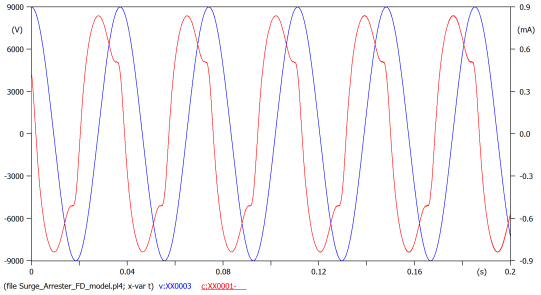

(c) 27 Hz

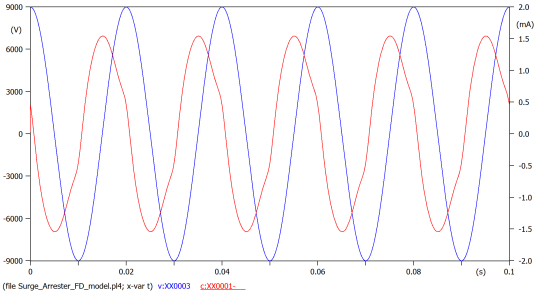

(d) 50 Hz

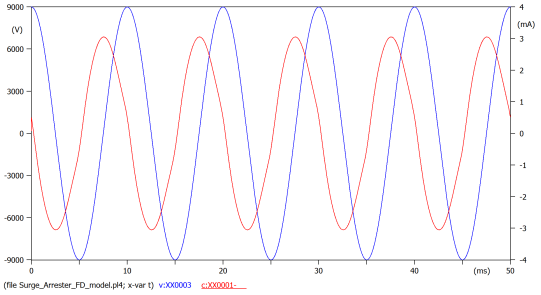

(e) 100 Hz

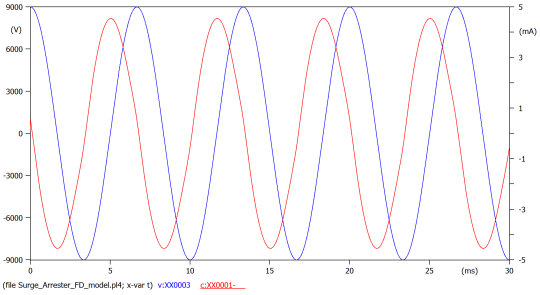

(f) 150 Hz

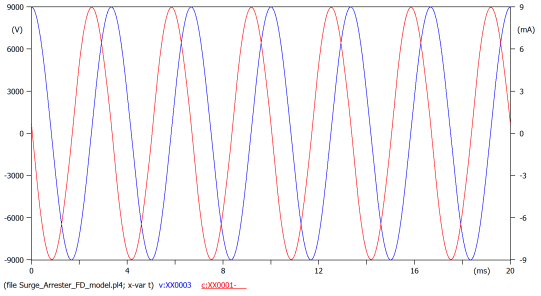

(g) 300 Hz

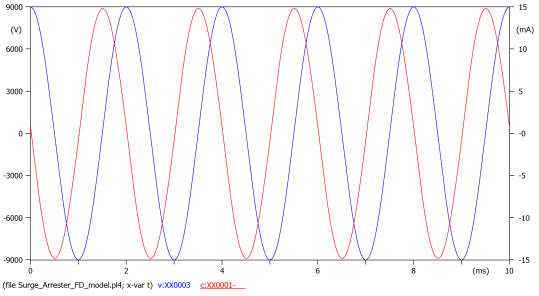

(h) 500 Hz

Figure 3.11: Frequency-dependent voltage-current characteristics of surge arrester model at 9 kV test voltage, showing the progressive transition from capacitive to resistive behaviour as frequency decreases.

Despite these limitations, the model provides representation of frequency-dependent behaviour in the leakage region where surge arresters spend most of their operational life. For system studies focused on insulation coordination and TOV withstand, the model provides the required accuracy. Applications requiring precise loss calculations in the conduction region should apply correction factors derived from the experimental data, particularly for harmonic-rich scenarios where frequency-dependent effects are most pronounced.

3.5. Summary: A Complete Frequency-Dependent Framework

This chapter has developed a frequency-dependent surge arrester model from physical understanding to practical implementation. The analysis began by noting that simple capacitive and resistive elements do not capture the complex impedance behaviour arising from distributed grain boundary relaxation. The Cole-Cole framework provided the mathematical approach to represent this distributed behaviour through physically meaningful parameters directly related to the arrester's microstructure.

The transformation from fractional-order impedance to realisable RC networks required approximation techniques. The Oustaloup method created a piecewise linear approximation on the logarithmic frequency scale, with each RC branch contributing to the overall fractional behaviour. The Foster II topology was selected for implementation, providing numerical stability and physical interpretation.

Implementation in ATP-EMTP required consideration of component arrangement and element selection. The frequency-dependent dielectric network combines with nonlinear resistance to create a model that captures both small-signal frequency response and large-signal nonlinear behaviour. Validation against laboratory measurements demonstrated the model's ability to represent the frequency-dependent phase shifts and loss variations.

The resulting model provides representation for surge arresters in modern power systems. As cable installations create conditions for low-order harmonic resonances, assessment of surge arrester thermal stress requires models that capture frequency-dependent losses. This model provides that capability while maintaining the computational efficiency required for large-scale system studies.

The next chapter will demonstrate the model's performance compared to measured MOV behaviour, revealing the accuracy of the frequency-dependent model.

4

Evaluation and Validation of Surge Arrester Model

Model validation requires comparison with experimental data to assess predictive capability. The frequency-dependent surge arrester model developed in the preceding chapters incorporates fractional-order impedances representing distributed relaxation, Foster networks approximating continuous spectra, and extracted nonlinear characteristics. Without rigorous validation against experimental data, these mathematical constructs remain theoretical exercises.

The validation challenge extends beyond point-by-point comparison. Surge arresters transition from linear dielectric response to highly nonlinear conduction across wide frequency and voltage ranges. A model that performs adequately at power frequency may fail at harmonics; one that captures low-voltage behaviour might diverge at higher stresses. Therefore, multiple validation approaches are employed, each examining different aspects of model performance.

Section 4.1 examines small-signal impedance analysis, evaluating whether the Foster network represents the distributed RC structure inherent in grain boundaries. Section 4.2 presents nonlinear validation through power and energy dissipation. Section 4.3 tests the methodology's applicability to other MOV blocks. Section 4.5 synthesises these findings—both the model's capabilities and its deviations emerge, indicating areas for refinement while demonstrating practical utility.

4.1. Impedance Response Validation

Before examining the model's performance under full operating conditions, verification of its small-signal behaviour is required. This validation approach isolates the frequency-dependent linear components from nonlinear effects, allowing assessment of the Foster network's ability to approximate the measured fractional-order behaviour.

The validation methodology employs AC analysis with a constant-amplitude sinusoidal current source, measuring the complex impedance $Z = R + jX$ across frequencies from 17 Hz to 500 Hz. Here, R represents the resistive component responsible for power dissipation, while

X denotes the capacitive reactance that governs energy storage and phase relationships. By maintaining operation in the linear region, this approach provides a test of the model's frequency-dependent structure without the complicating effects of voltage-dependent nonlinearity.

Figure 4.1 presents the comparison between model predictions and measured impedance components at 5 kV applied voltage. The logarithmic scales reveal power-law relationships characteristic of distributed RC networks, with both resistance and reactance magnitude decreasing monotonically with frequency. The resistance drops nearly two orders of magnitude from approximately 1.7 MΩ at 17 Hz to 15 kΩ at 500 Hz, while reactance magnitude spans from 16.5 MΩ to 0.58 MΩ over the same range.

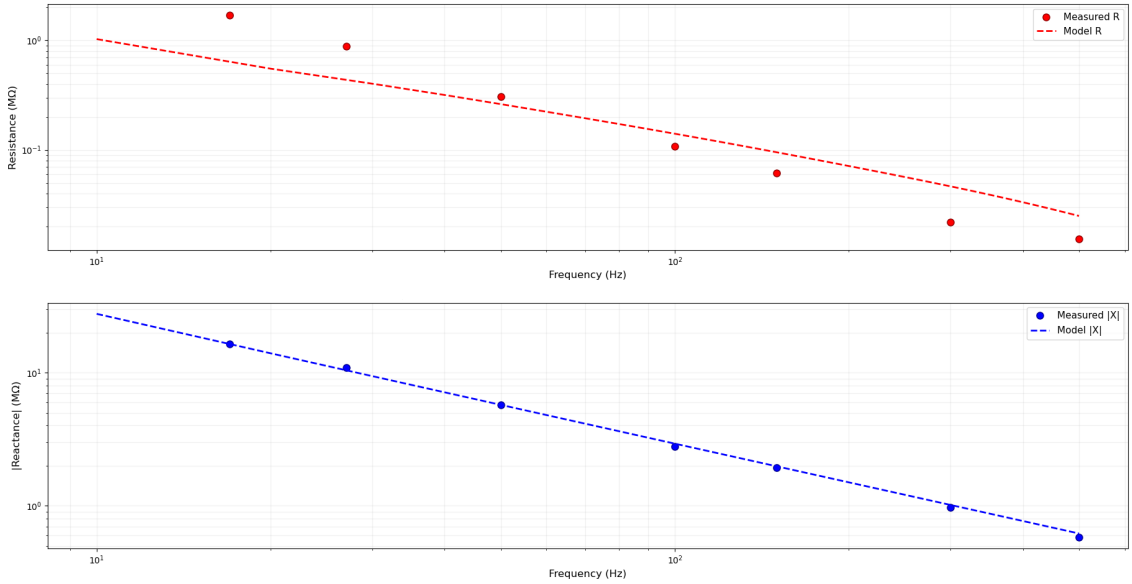

Figure 4.1: Validating impedance response of the model with the measurement values at 5.5 kV

Across the entire frequency range, the model tracks the measured capacitive reactance with errors remaining below 6 % for most frequencies. The maximum deviation of 12 % occurs at the lowest test frequency of 17 Hz, where measurement uncertainties are also greatest due to the extended measurement times required and higher impedance. The mean absolute error across all frequencies is 4.3 %, indicating that the selected capacitance values $C_\infty = 0.522$ nF and $C_s = 0.170$ nF, combined with the Foster network structure, represent the surge arrester's dielectric properties.

The resistance comparison reveals larger deviations than the reactance. While the model captures the general frequency-dependent trend of resistance—the power-law decrease with increasing frequency—systematic deviations appear across the spectrum. At low frequencies (17 Hz–50 Hz), the model underestimates resistance, with errors reaching −76.2 % at 17 Hz. This underestimation gradually diminishes with increasing frequency, crossing through near-perfect agreement at 100 Hz (0.2 % error) before transitioning to overestimation at higher frequencies. By 300 Hz, the model overshoots the measured resistance by 94.1 %, though the absolute magnitude of this error is smaller due to the lower resistance values at high frequencies.

Table 4.1: Impedance model validation results

Frequency (Hz)	R_{equ} (MΩ)	R_{model} (MΩ)	Error (%)	$\lvert X_{\text{equ}}\rvert$ (MΩ)	$\lvert X\rvert_{\text{model}}$ (MΩ)	Error (%)
17	1.690	0.402	-76.2	16.500	18.483	12.0
27	0.883	0.277	-68.6	10.900	10.951	0.5
50	0.308	0.189	-38.6	5.740	5.783	0.7
100	0.107	0.107	0.2	2.790	2.942	5.4
150	0.061	0.075	22.5	1.940	1.980	2.1
300	0.022	0.042	94.1	0.974	1.009	3.6
500	0.015	0.025	60.8	0.582	0.617	6.0

Mean Absolute Error - Resistance: 51.6 %
Mean Absolute Error - Reactance: 4.3 %

The pattern of resistance errors indicates the model's limitations. The transition from underestimation to overestimation suggests that while the three-branch Foster network captures the general frequency dependence, it cannot perfectly reproduce the continuous relaxation spectrum of the real device. The crossing point near 100 Hz corresponds to the geometric centre of the measurement range on a logarithmic scale (17 Hz to 500 Hz), where the optimization algorithm makes a compromise between low and high frequency errors.

These impedance validation results establish expectations for the full nonlinear validation that follows. The reactance agreement indicates that the model will predict phase relationships and reactive power flow accurately. The resistance errors, while significant in percentage terms, follow a predictable pattern. This validation demonstrates that the model structure—parallel combination of frequency-dependent and frequency-independent elements—represents the surge arrester's small-signal behaviour.

4.2. Validation of the Model

The transition from impedance characteristics to power and energy dissipation represents a comprehensive test of model fidelity. While impedance validation examined the model's linear response, practical surge arrester operation involves the complex interaction between frequency-dependent linear elements and voltage-dependent nonlinear conduction. This section validates the model's ability to predict actual power and energy losses—parameters that directly determine thermal stress and ultimately surge arrester survival.

The validation methodology employs a custom MODEL block implemented in ATP-EMTP that performs real-time calculation of electrical parameters during transient simulation. This approach ensures that all nonlinear interactions are captured, including the phase shifts arising from frequency-dependent components and the conductivity changes as voltage approaches the protection level. The implementation computes three interconnected parameters that fully characterise energy dissipation.

The equations implemented in both model and measurement analysis are:

$$p(t) = v(t) \times i(t) \tag{4.1}$$

$$E(t) = E(t - \Delta t) + \frac{\Delta t}{2}[p(t) + p(t - \Delta t)] \tag{4.2}$$

$$P_{avg} = \frac{E(T)}{T} \tag{4.3}$$

$$E_{cycle} = \frac{E(T)}{N_{cycles}} \tag{4.4}$$

Different measurement durations at various frequencies could introduce bias if not properly addressed. Therefore, all waveforms were truncated to contain equal numbers of complete cycles, eliminating edge effects and ensuring fair comparison across the frequency spectrum. The validation dataset comprises 55 matched measurement-model pairs spanning eight frequencies (10, 17, 27, 50, 100, 150, 300, and 500 Hz) with voltages ranging from approximately 3 kV to 9 kV—covering both the linear dielectric region and the onset of nonlinear conduction. The following figures present representative comparisons at three key frequencies that span the validation range:

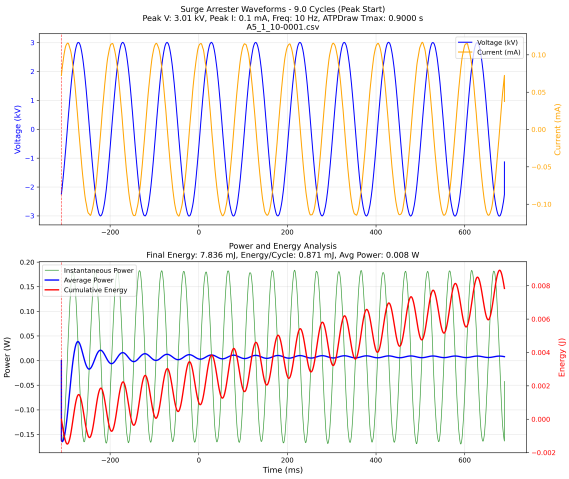

(a) Measurements

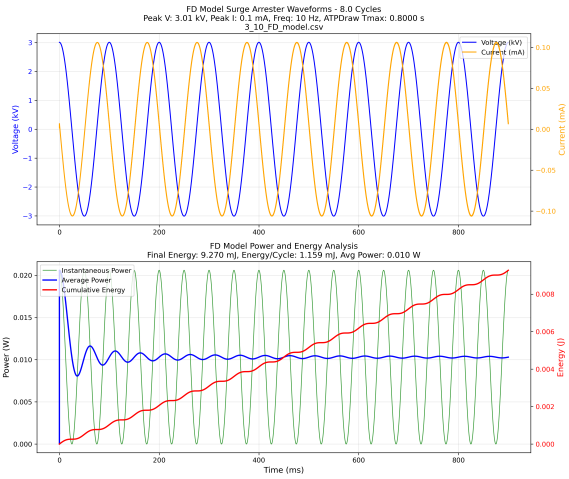

(b) Model

Figure 4.2: Energy and Power losses at 10 Hz for 3 kV

At 10 Hz and 3 kV (Figure 4.2), the surge arrester operates in the linear region where dielectric losses dominate. The measured cumulative energy reaches 7.836 mJ, while the model predicts 9.27 mJ—an 18 % overestimation. The average power shows similar behaviour with 0.008 W measured versus 0.01 W modelled. This overestimation at low frequencies appears to contradict the resistance underestimation observed in impedance validation. However, the dominant capacitive reactance, which is overestimated by 12 % at 17 Hz, results in an overall impedance magnitude overestimation that leads to current underestimation.

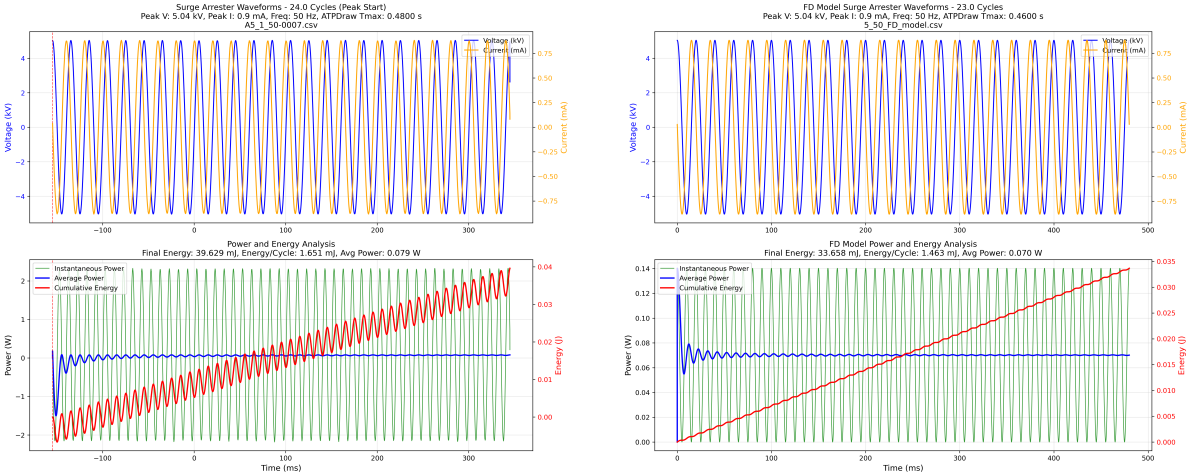

(a) Measurements **(b)** Model

Figure 4.3: Energy and Power losses at 50 Hz for 5 kV

The 50 Hz validation at 5 kV (Figure 4.3) represents typical power frequency operation. The model shows measured average power of 0.08 W versus modelled 0.07 W—a 12.5 % underestimation. The cumulative energy comparison shows 39.63 mJ measured against 33.66 mJ modelled, while energy per cycle yields 1.651 mJ measured versus 1.463 mJ modelled. This underestimation at 50 Hz is particularly relevant given that surge arresters spend most of their service life operating at this frequency.

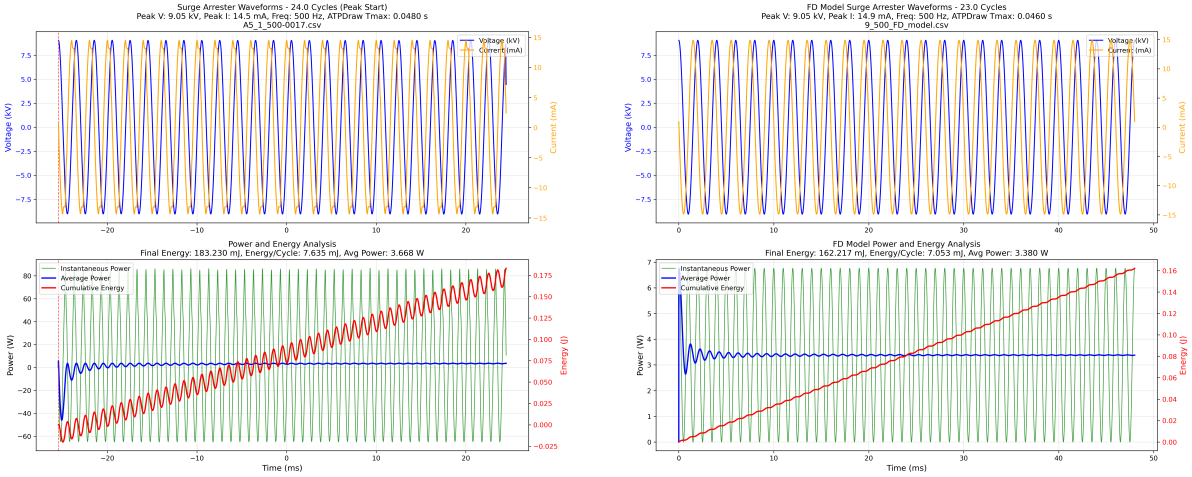

(a) Measurements **(b)** Model

Figure 4.4: Energy and Power losses at 500 Hz for 9 kV

At the upper frequency extreme of 500 Hz and 9 kV (Figure 4.4), the surge arrester approaches its nonlinear conduction region. The substantial power dissipation—3.668 W measured versus 3.38 W modelled—reflects the combined effects of increased conduction and higher frequency operation. The model shows consistent underestimation with cumulative energy of 183.23 mJ measured against 162.22 mJ modelled (11.5 % error). At 500 Hz, the model underestimates energy per cycle by 7.6 % (7.053 mJ modelled vs 7.635 mJ measured).

The comprehensive validation dataset enables statistical analysis of model performance across the full operating range:

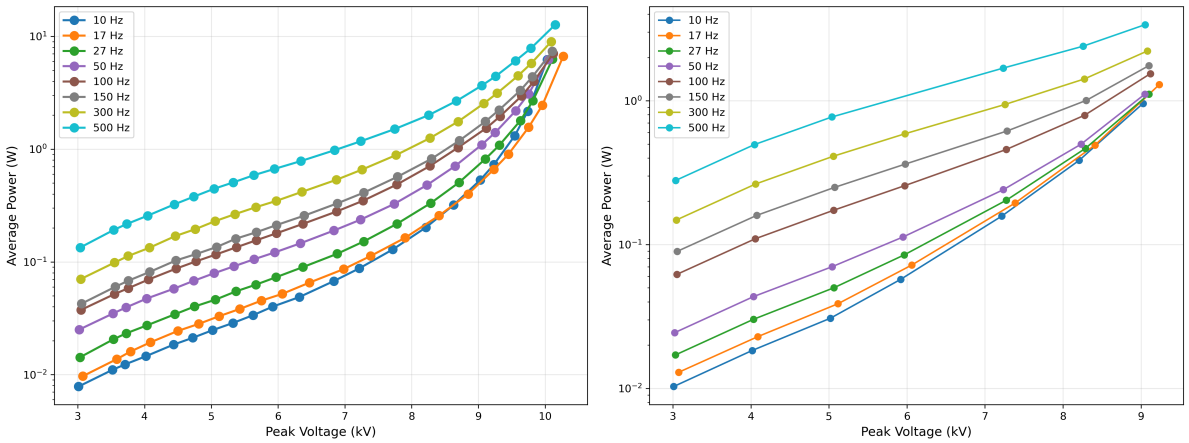

(a) Measurements **(b)** Model

Figure 4.5: Overview of average power losses

The average power comparison in Figure 4.5 reveals that the model captures the characteristic behaviour of surge arrester losses. The proper frequency ordering is maintained throughout—power consistently decreases from 10 Hz to 500 Hz at any given voltage. The exponential growth of power with voltage, characteristic of the emerging nonlinear conduction, appears in both model and measurements with similar slopes on the semi-logarithmic plots.

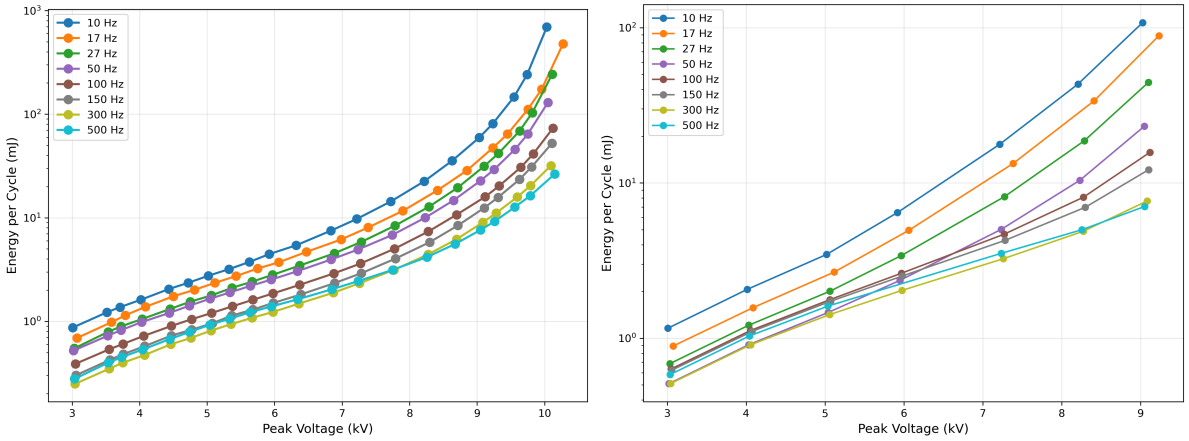

(a) Measurements **(b)** Model

Figure 4.6: Overview of energy losses per cycle

The energy per cycle analysis (Figure 4.6) reveals frequency-dependent behaviour by removing the time-scaling effects. The validation reveals a distinctive pattern: at 10 Hz, the model overestimates energy per cycle (1.159 mJ model vs 0.871 mJ measured), and this overestimation continues at most frequencies. However, near power frequency (50 Hz), the model achieves improved accuracy with slight underestimation (1.463 mJ vs 1.651 mJ). At higher frequencies, the model returns to overestimation, though the magnitude varies with voltage as nonlinear effects become significant.

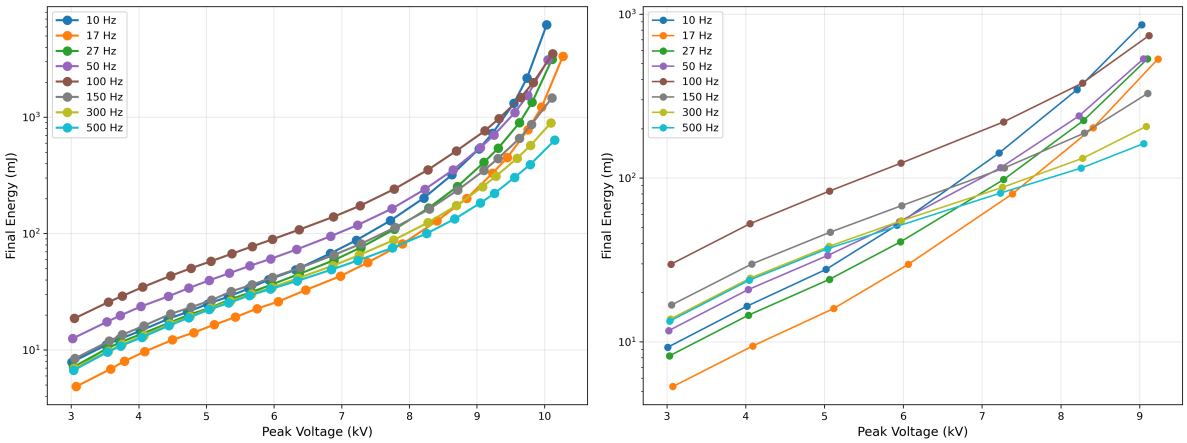

(a) Energy losses from measurements

(b) Energy losses from model

Figure 4.7: Comparing Cumulative energy losses in mJ

At lower frequencies, particularly subharmonics (10, 17, 27 Hz), energy losses (as shown in Figure 4.7) increase more steeply with voltage than at higher frequencies, eventually matching or exceeding the losses at 500 Hz despite lower current magnitudes. This behaviour could stem from time-dependent conduction mechanisms in the ZnO varistor structure. The longer voltage half-cycles at low frequencies allow complete development of nonlinear conduction processes—including thermionic emission, space charge accumulation, and grain boundary charging—resulting in more pronounced current peaks at lower frequencies for the same voltage.

The observation that cumulative energy peaks at 100 Hz—the 2nd harmonic—may reflect either material-specific relaxation phenomena or the influence of power system harmonic patterns on surge arrester design evolution. This correlation merits further investigation, particularly given the prevalence of 2nd harmonic content in transformer energisation and other system transients. At harmonics relevant to TOV, the dielectric losses are as significant as nonlinear resistive losses.

Direct comparison plots provide additional visualisation of model performance:

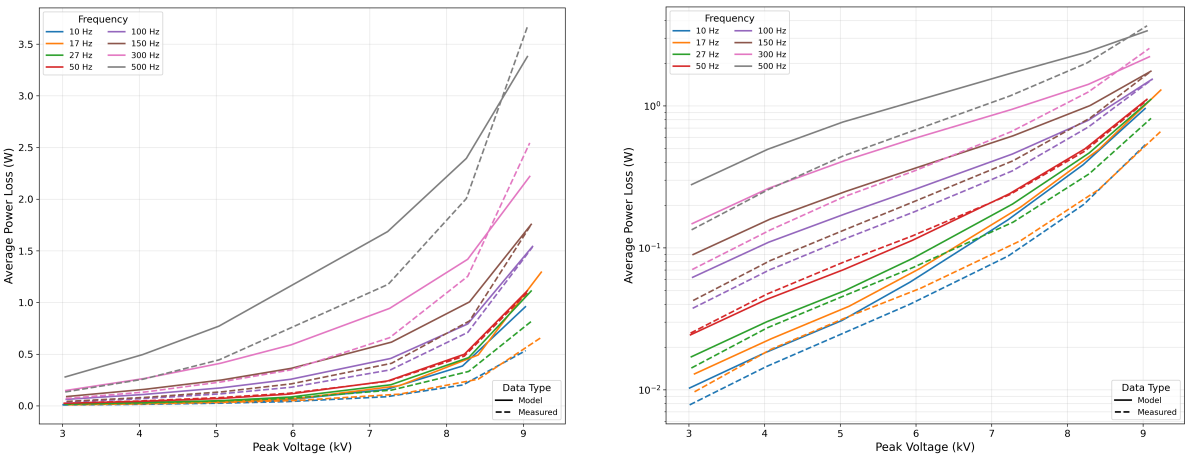

(a) Linear view

(b) Logarithmic view

Figure 4.8: Overview of average power losses

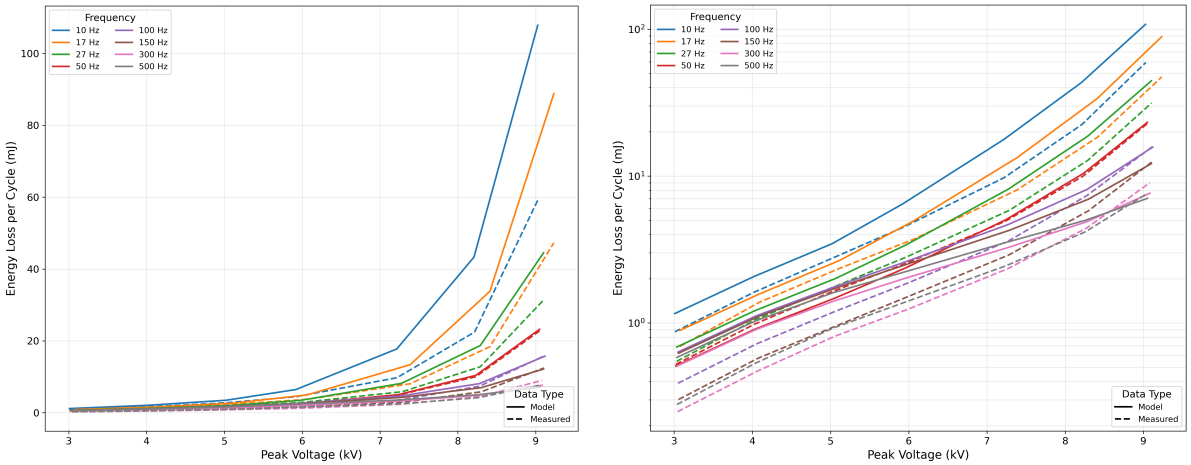

(a) Linear view **(b)** Logarithmic view

Figure 4.9: Overview of energy losses per cycle

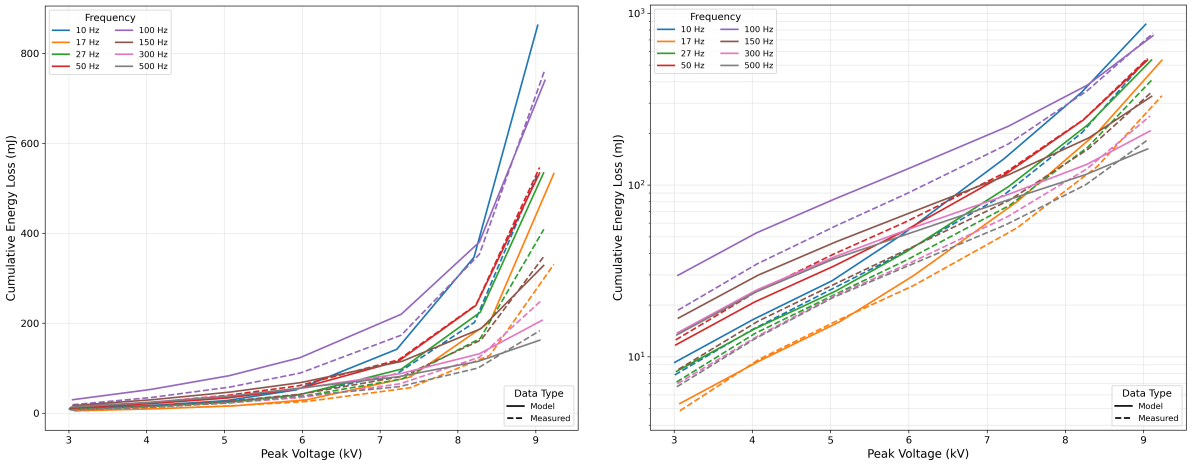

(a) Linear view **(b)** Logarithmic view

Figure 4.10: Overview of cumulative energy losses

Quantitative error analysis across the entire validation dataset reveals systematic patterns:

Table 4.2: Model validation error statistics by frequency

Frequency (Hz)	Points	Energy Error (%)	Power Error (%)
10	7	38.15 ± 25.91	53.49 ± 28.78
17	7	25.69 ± 27.44	52.53 ± 33.30
27	7	18.89 ± 12.53	23.57 ± 13.02
50	7	-6.99 ± 5.59	-3.18 ± 5.83
100	7	32.13 ± 23.07	36.90 ± 23.89
150	7	52.98 ± 38.01	62.10 ± 40.22
300	7	47.48 ± 42.20	57.12 ± 44.92
500	6	48.97 ± 43.04	54.91 ± 44.66
Overall		31.86 ± 33.43	41.95 ± 36.52

The error statistics in Table 4.2 reveal distinct frequency-dependent behaviour. The model achieves its best accuracy at 50 Hz with mean errors below 7 %—the only frequency showing underestimation. This power frequency performance is particularly relevant since surge arresters operate predominantly at 50 Hz throughout their service life. At all other frequencies, the model overestimates losses, with errors increasing from 18.89 % at 27 Hz to over 50 % at higher frequencies.

The large standard deviations indicate strong voltage dependence in model accuracy. At 500 Hz, for instance, the mean shows 48.97 % overestimation, yet the 9 kV case demonstrates 11.5 % underestimation. This reversal suggests that nonlinear conduction at higher voltages partially compensates for the frequency-dependent resistance errors observed in linear operation.

The error pattern—underestimation only at 50 Hz with overestimation elsewhere—reflects the limitation of approximating a continuous relaxation spectrum with a discrete Foster network. Physical mechanisms not captured by the model, including voltage-current hysteresis at low frequencies and temperature-dependent effects during extended operation, likely contribute to these deviations. Nevertheless, the model preserves all essential qualitative features: correct frequency ordering of losses, exponential voltage dependence, and the characteristic grain boundary relaxation behaviour.

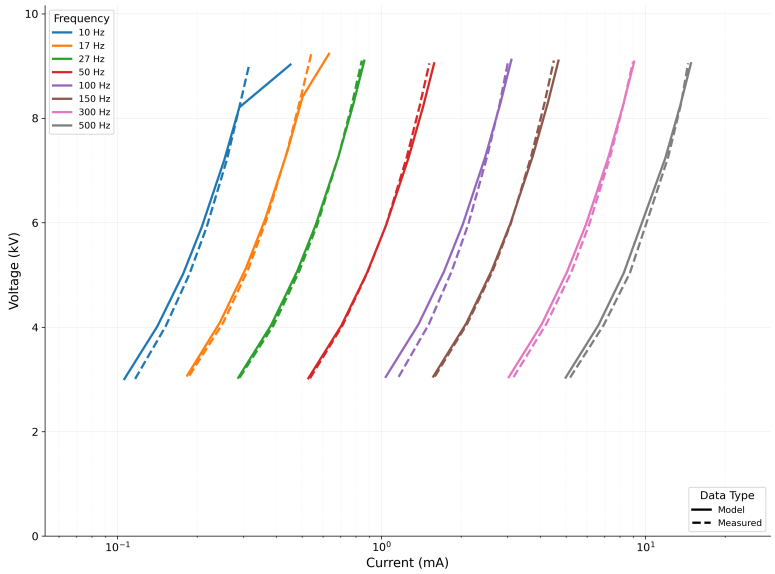

Figure 4.11: VI characteristics of the model and measurements compared

Despite these limitations, the validation indicates that the frequency-dependent model represents an advancement over traditional frequency-independent approaches. By explicitly incorporating the distributed relaxation processes inherent in polycrystalline metal oxide materials, the model provides capabilities that simpler models cannot achieve. The demonstrated ability to predict frequency-dependent losses within reasonable bounds enables more accurate assessment of surge arrester thermal stress in harmonic-rich environments—the primary motivation for this modelling effort.

4.3. Validation of the Methodology

The measure of a modelling methodology lies not in its ability to represent a single device, but in its generalisability to the broader class of components it describes. The central question remains: *Can this approach be applied to other MOV blocks? How accurately can the methodology model other MOV blocks with different voltage ratings, energy capabilities, or manufacturer-specific compositions?*

To address this question, the characterisation methodology was applied to a different MOV geometry—a station-class surge arrester block manufactured by Tridelta Meidensha (MZE100B23). This disc differs from the original test specimen: while maintaining comparable diameter (100 mm versus 64.5 mm), its height is reduced by nearly half (22.5 mm versus 44 mm). This geometric variation provides a test case for validating the measurement approach.

Table 4.3: Specifications of the station class surge arrester provided by Tridelta Meidensha[33]

Parameter	Value
Diameter	100 mm
Height	22.5 mm
Rated voltage (rms)	3.0 kV
Maximum continuous operating voltage (rms)	2.4 kV
Reference current	1 mA
Reference voltage (rms)	4.18 kV to 4.82 kV
Nominal discharge current (8/20 µs)	10 kA
Residual voltage ratio at 10 kA	1.57
Residual voltage at 10 kA (8/20 µs)	6.56 kV to 7.57 kV

Figure 4.12: MOV Block used for validation

The characterisation measurements for this validation block were conducted over a frequency range of 10 Hz–300 Hz, capturing harmonics up to the 6th order. This differs from the original test protocol, which extended to 500 Hz. The frequency limitation to 300 Hz was due to the power supply's maximum current capability. The higher capacitance of this MOV block (2.5 nF vs 0.5 nF) resulted in the current limit being reached at a lower frequency than the first block. This phenomenon is addressed in Appendix A. Despite this measurement limitation, a frequency-dependent model has been reconstructed.

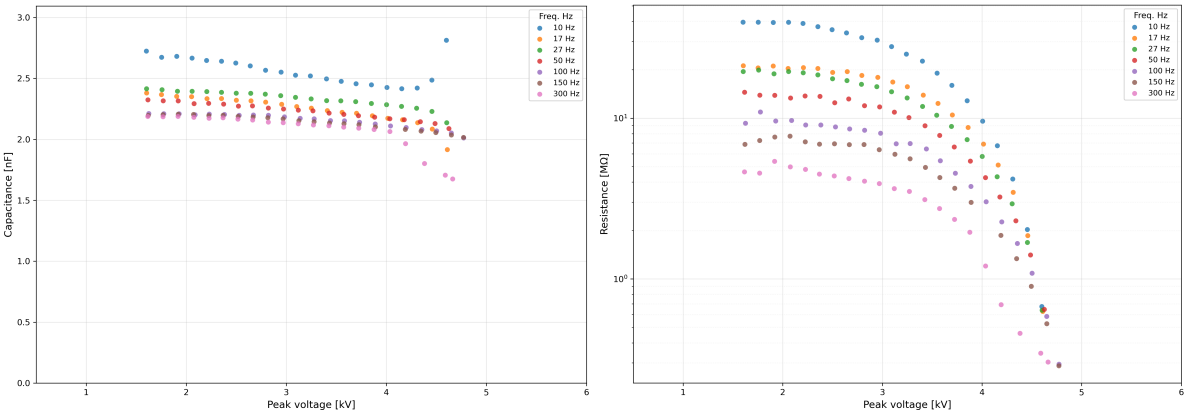

(a) Capacitance **(b)** Resistance

Figure 4.13: Extracted capacitance and resistance of the validation MOV as a function of peak applied voltage

Figure 4.13 reveals the consistency of MOV behaviour across different geometries. The capacitance maintains its characteristic stability in the pre-conduction region at approximately 2.5 nF—a value 4.5 times higher than the original block. This enhancement directly reflects the geometric differences, as capacitance depends on the ratio of area to thickness ($C \propto A/t$). The flatter disc combines a 2.4-fold increase in cross-sectional area with a twofold reduction in height, yielding the observed capacitance scaling.

The resistance exhibits the expected monotonic decrease with both voltage and frequency, with absolute values reduced by an order of magnitude compared to the thicker block. This reduction, combined with the preserved frequency dependence, demonstrates the conduction mechanism identified in Chapter 2. In the pre-conduction region, the MOV behaves as a parallel combination of high-value resistance and capacitance, yielding an apparent resistance that varies inversely with frequency ($R_{\text{app}} \approx 1/\omega C$). The fourfold increase in capacitance thus translates directly to a proportional decrease in apparent resistance, shifting the entire resistance-voltage-frequency surface downward while preserving the characteristic trends observed in the original measurements.

The decomposition into linear and nonlinear resistance components is shown in Figure 4.14. The linear resistance component maintains the same voltage-invariant, frequency-dependent decay observed in the original characterisation—a power-law relationship that reflects the conduction mechanisms within the ZnO grains and grain boundaries established in Section 2.3. The preservation of this trend, despite the order-of-magnitude change in absolute values, indicates that the measurement approach captures intrinsic material properties rather than sample-specific artifacts.

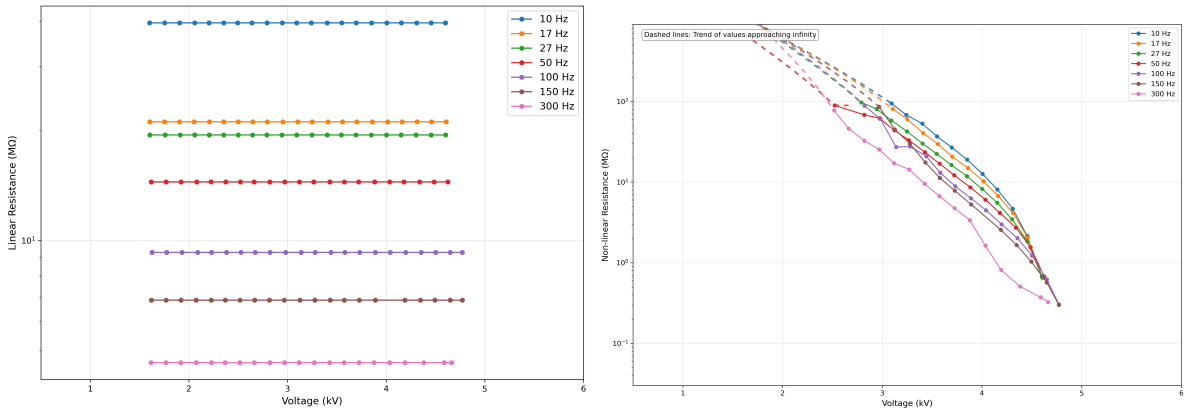

(a) Linear resistance **(b)** Non-linear resistance

Figure 4.14: Decomposed resistance components of the validation MOV as a function of peak applied voltage

The nonlinear resistance demonstrates agreement with the expected trends, exhibiting smooth voltage dependence with minimal scatter. This improved consistency likely stems from this block's lower degree of polarisation. Unlike the original test specimen, which had experienced repeated high-current impulses during qualification testing, this validation block retained more uniform grain boundary properties throughout its volume.

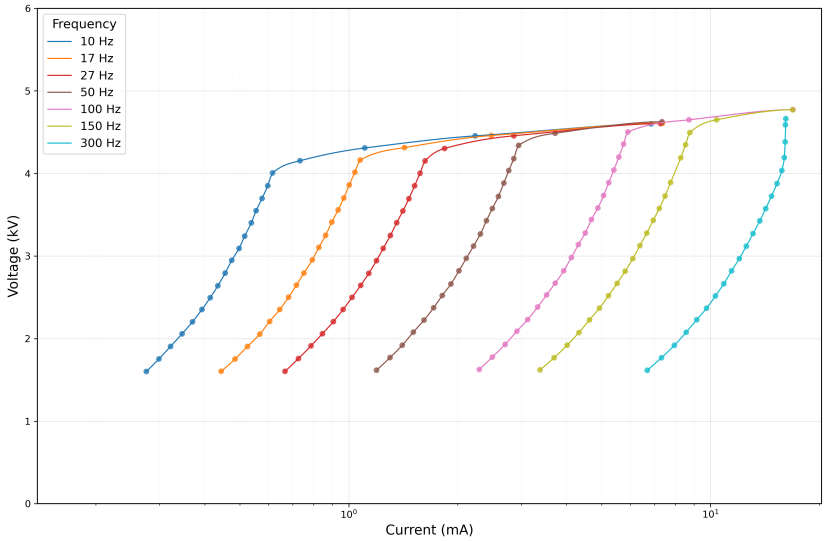

Figure 4.15: V-I characteristics of validation MOV block

The measured V-I characteristics (Figure 4.15) show the expected leakage-region behaviour compressed into a lower voltage range. The transition from leakage to conduction occurs at approximately 4.5 kV peak—about half the voltage observed in the original block, indicating the linear relationship between operating voltage and disc height. This geometric scaling arises from the electric field distribution: for a given field strength that triggers grain boundary conduction, the total voltage scales directly with the number of grain boundaries in series, which is proportional to the disc thickness.

This characterisation of a geometrically distinct MOV block validates the methodology's broader applicability and demonstrates that the approach captures the fundamental physics governing

MOV behaviour. The improved measurement quality achievable with less polarised samples further suggests that characterising and modelling should be performed on depolarised blocks in future surge arrester studies.

Following the characterisation of this block, the complete frequency-dependent modelling methodology developed in Chapter 3 was applied. The measured impedance spectra were first analysed using Cole-Cole theory to extract the dielectric relaxation parameters, as shown in Figure 4.16.

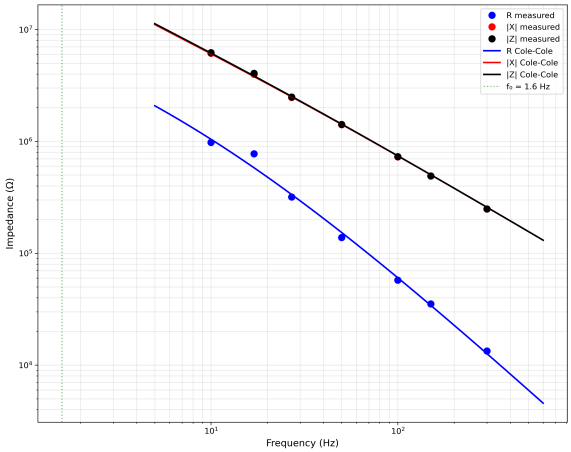

Figure 4.16: Impedance spectra of the validation MOV block with Cole-Cole fit.

Parameter	Value
C_∞	1.96 nF
C_0	4.51 nF
f_0	1.59 Hz
τ_0	0.10 s
α	0.55

Table 4.4: Cole–Cole fit parameters for the validation MOV.

The extracted Cole-Cole parameters reveal differences from the original specimen. The higher capacitance values ($C_\infty = 1.96$ nF versus 0.522 nF) directly reflect the geometric differences, while the fractional exponent $\alpha = 0.55$ indicates a slightly broader distribution of relaxation times compared to the original block's $\alpha = 0.64$. These parameters served as inputs to the Foster network synthesis algorithm, which employed the Oustaloup approximation to transform the fractional-order impedance into a realisable circuit with three RC branches, maintaining consistency with the original modelling approach.

Table 4.5: Foster network parameters for the validation MOV model.

Branch	C_k (nF)	R_k (MΩ)	τ_k (ms)	$f_{z,k}$ (Hz)
1	0.56	15.14	8.449	18.8
2	0.22	3.85	0.845	188.4
3	0.11	0.77	0.084	1883.6
External components:				
G_∞	0.831 nS (1203.9 MΩ)			
C_∞	2 nF			
C_s	2.5 nF			

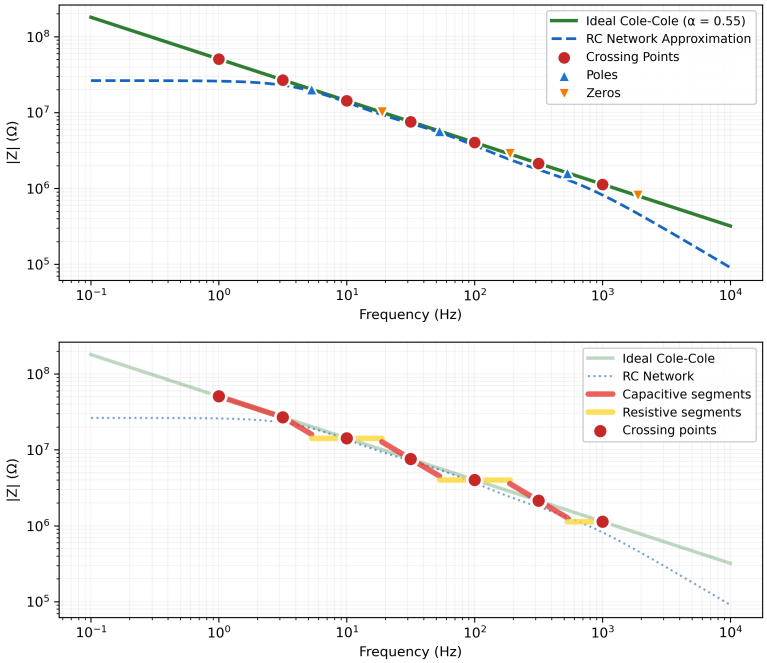

Figure 4.17: Oustaloup approximation of the validation MOV's fractional impedance.

The synthesised Foster network parameters (Table 4.5) exhibit the expected scaling with geometry: resistance values are reduced by approximately one order of magnitude while capacitances increase fourfold, consistent with the changed disc geometry. The decade spacing of the characteristic frequencies (18.8, 188.4, and 1883.6 Hz) provides uniform coverage of the frequency spectrum, ensuring accurate representation of the distributed relaxation processes. Before proceeding to full nonlinear validation, the small-signal behaviour was verified using the approach established in Section 4.1, with results shown in Figure 4.18.

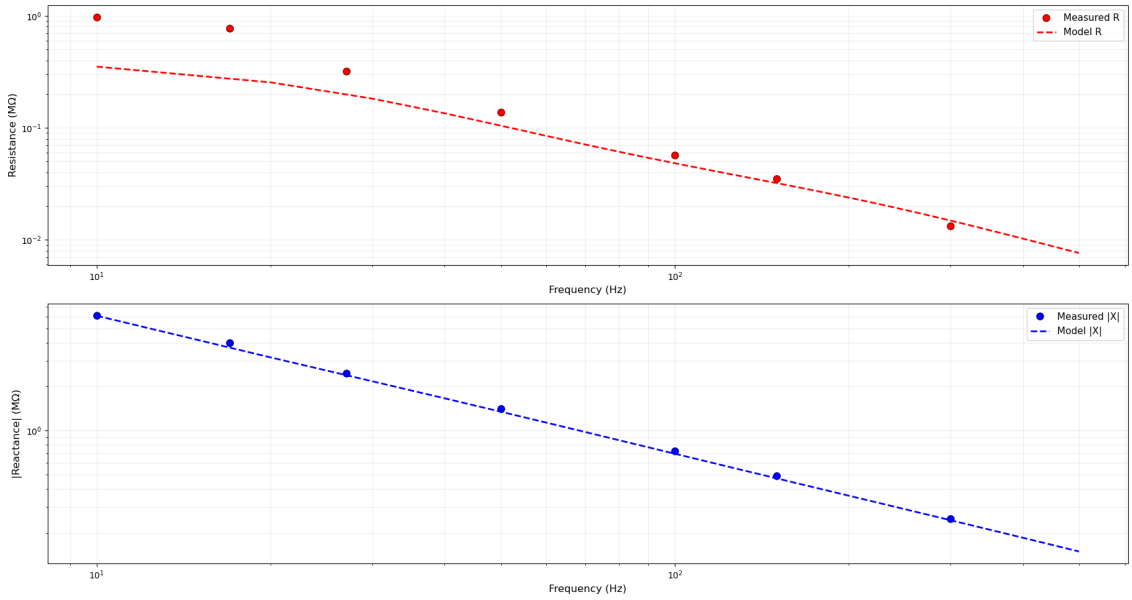

Figure 4.18: Validating impedance response of the model with the measurement values at 3 kV

Mean Absolute Error - Resistance: 31.8 %

Mean Absolute Error - Reactance: 2.5 %

The validation dataset comprises 49 matched measurement-model pairs across seven frequencies (10, 17, 27, 50, 100, 150, and 300 Hz), providing comprehensive coverage of the operational spectrum.

The power loss comparison (Figure 4.19) demonstrates the methodology's successful adaptation to this different device. The model maintains proper frequency ordering and captures the exponential voltage dependence across nearly three orders of magnitude in power dissipation. The visual agreement between model and measurement appears consistent with that achieved for the first MOV block.

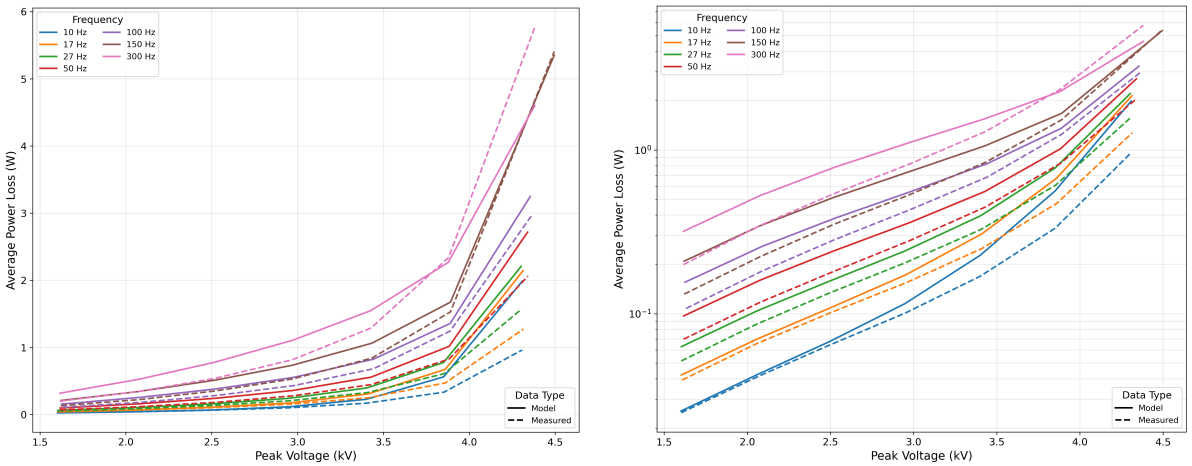

(a) Linear view **(b)** Logarithmic view

Figure 4.19: Average power loss comparison for the second MOV block

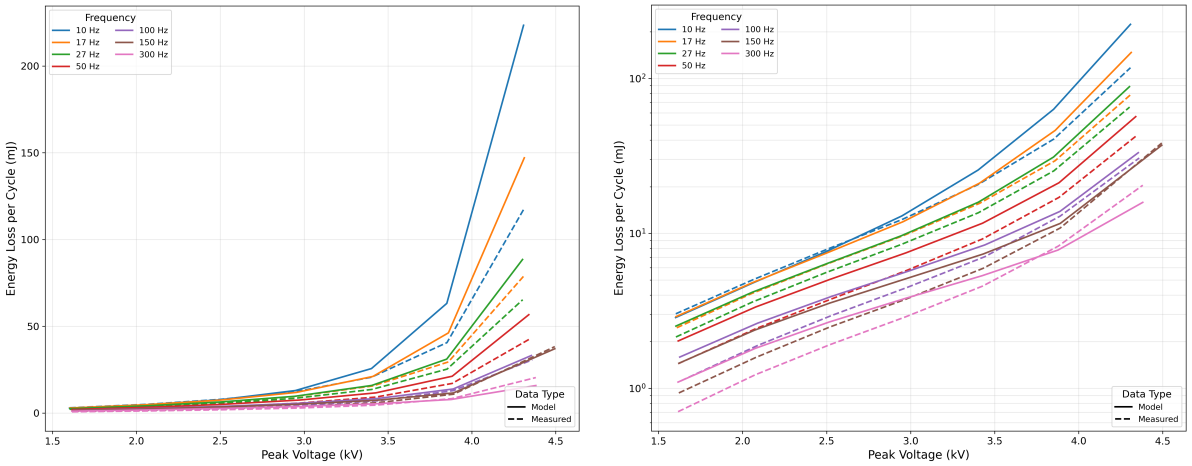

(a) Linear view **(b)** Logarithmic view

Figure 4.20: Energy per cycle comparison for the second MOV block

The energy per cycle analysis (Figure 4.20) indicates the transferability of the methodology. While the first MOV exhibited a clear pattern of overestimation at frequencies away from 50 Hz, this second device shows a more nuanced response. At 10 Hz, the model underestimates energy per cycle by 7.9 %—a complete reversal from the 55.4 % overestimation observed in the first specimen.

The quantitative comparison in Table 4.6 reveals differences in model performance between the two devices. The second MOV achieves overall mean absolute errors of 22 %—approximately half that of the first device.

Table 4.6: Comparative error statistics between MOV blocks

Frequency (Hz)	First MOV Block		Second MOV Block	
	Energy Error (%)	Power Error (%)	Energy Error (%)	Power Error (%)
10	55.4 ± 29.1	53.5 ± 28.8	-7.9 ± 50.4	-0.3 ± 54.6
17	46.6 ± 32.0	52.5 ± 33.3	7.1 ± 37.2	-2.6 ± 33.9
27	28.8 ± 13.6	23.6 ± 13.0	-3.5 ± 16.1	0.3 ± 16.8
50	-2.9 ± 5.8	-3.2 ± 5.8	10.9 ± 7.1	9.6 ± 7.0
100	34.9 ± 23.6	36.9 ± 23.9	10.2 ± 9.1	11.8 ± 9.3
150	58.6 ± 39.4	62.1 ± 40.2	18.6 ± 17.5	21.3 ± 17.9
300	52.9 ± 43.8	57.1 ± 44.9	18.7 ± 27.1	22.0 ± 27.8
500	55.4 ± 44.9	54.9 ± 44.7	—	—
Overall MAE	43.0	43.8	22.0	22.4

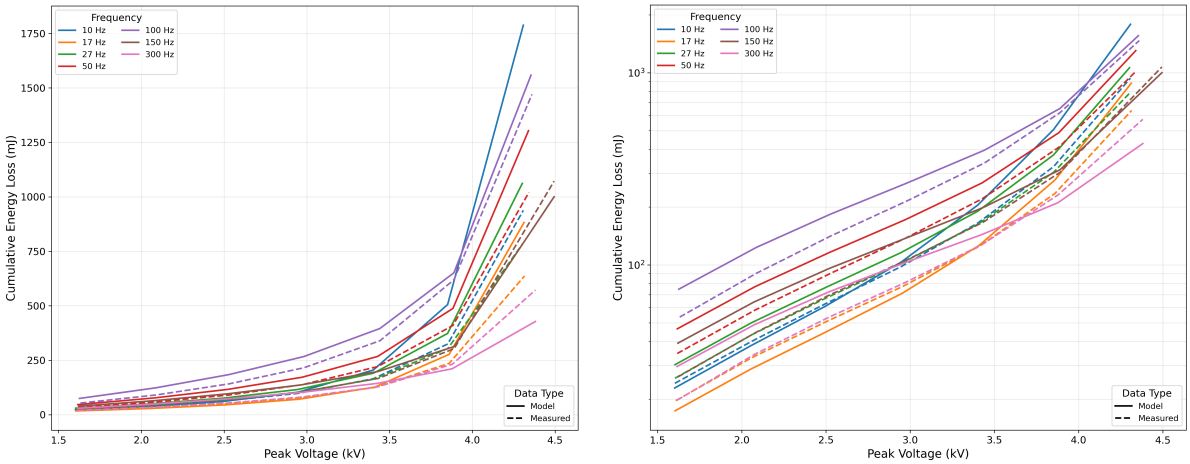

(a) Linear view **(b)** Logarithmic view

Figure 4.21: Cumulative energy loss comparison for the second MOV block

The cumulative energy comparison (Figure 4.21) demonstrates consistent tracking between model and measurement across the full operational range. The logarithmic view particularly emphasises the model's ability to maintain proper scaling relationships even as the operating point transitions from linear dielectric behaviour to the onset of nonlinear conduction. The consistent parallelism between model and measurement curves at each frequency indicates that the fundamental physics—frequency-dependent grain boundary relaxation coupled with voltage-dependent conduction—is properly captured regardless of the specific MOV block characteristics.

The voltage-current characteristics (Figure 4.22) provide evidence of successful methodology transfer. Despite operating at lower voltages and exhibiting different microscopic properties, the model accurately reproduces the transition from ohmic to nonlinear conduction. The frequency-dependent separation of curves, arising from the complex impedance of grain boundaries, is captured across the entire current range spanning three orders of magnitude.

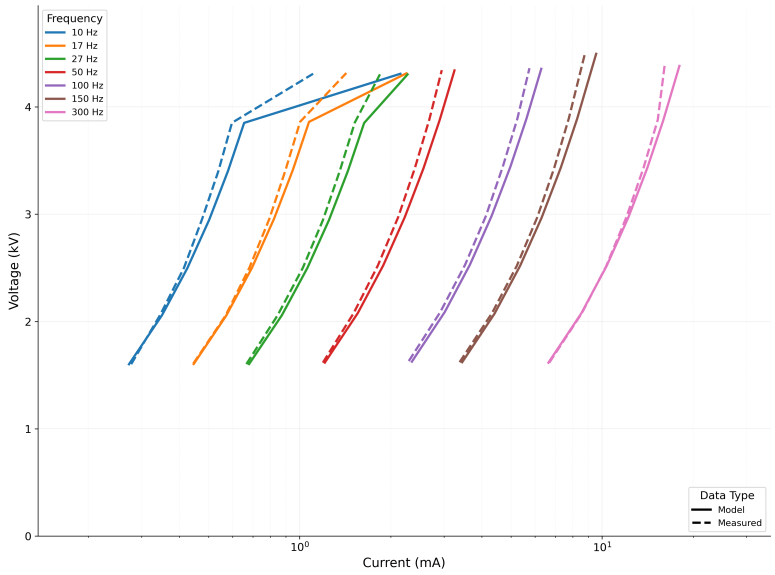

Figure 4.22: V-I characteristics comparison for the second MOV block

Analysis of voltage-dependent error patterns reveals another dimension of methodology robustness. The second MOV shows improved accuracy in the mid-voltage range (2.18 kV–3.34 kV) with errors of only 2 %–3 %, increasing to 27.8 % at the highest voltages. This voltage-dependent variation, also observed in the first device but with different characteristics, suggests that the nonlinear conductance model may require voltage-range-specific tuning for enhanced accuracy across the full operational spectrum.

The application to a fundamentally different MOV block validates several key aspects of the methodology:

1. Transferability: The impedance spectroscopy and Foster network synthesis procedures extract device-specific parameters without requiring *a priori* knowledge of material properties or manufacturing details.

2. Physical Fidelity: Different error patterns between devices indicate that the model captures real physical differences rather than imposing artificial behaviour. Each MOV's unique grain structure and dopant distribution manifests in its specific Foster network parameters.

3. Robustness: Despite operating at different voltage ranges and exhibiting different relaxation spectra, both models maintain proper qualitative behaviour—frequency ordering, voltage scaling, and transition to nonlinear conduction.

4. Practical Accuracy: With mean errors of 22 % for the second device compared to 43 % for the first, the methodology demonstrates that accuracy depends on device-specific characteristics rather than fundamental limitations of the approach.

This validation on a distinctly different MOV block indicates that the frequency-dependent modelling methodology generalises beyond single-device representation. By adapting to device-specific characteristics through impedance-based parameter extraction, the approach provides a framework for modelling the diverse range of MOV blocks encountered in surge pro-

tection applications. The methodology's ability to capture both common physics and device-specific variations positions it as a practical tool for surge arrester design and analysis in increasingly complex power system environments.

Extending the Foster network from three to five branches (covering 0.1 Hz–10 kHz versus 10 Hz–1000 Hz) improved resistance matching from 31.8 % to 16.7 %. However, energy prediction showed 29 % overestimation versus 22 % for the three-branch model. This pattern suggests that increased network complexity would benefit from multi-objective optimisation considering both impedance and energy metrics simultaneously. The complete five-branch analysis is presented in the next section.

4.4. Extended Foster Network Analysis: Five-Branch Implementation

The validation of the frequency-dependent surge arrester model on a second MOV block achieved 22 % mean absolute error with a three-branch Foster network. This section examines whether a five-branch Foster network improves frequency-dependent representation.

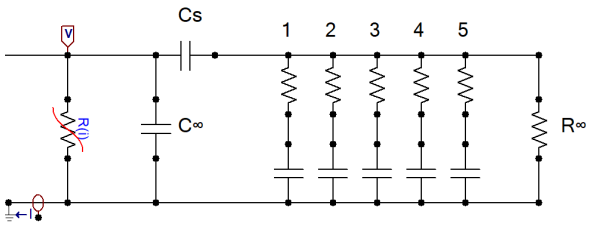

Figure 4.23: Extended five-branch frequency-dependent surge arrester model in ATP

The Oustaloup approximation extends to higher-order Foster networks. Analysis of the three-branch implementation showed systematic deviations at frequency extremes, where the approximation transitions to its asymptotic behaviour. Extending the approximation range from 10 Hz–1000 Hz (three decades) to 0.1 Hz–10 kHz (five decades) moves the edge effects beyond the operational frequency band of 10 Hz–300 Hz.

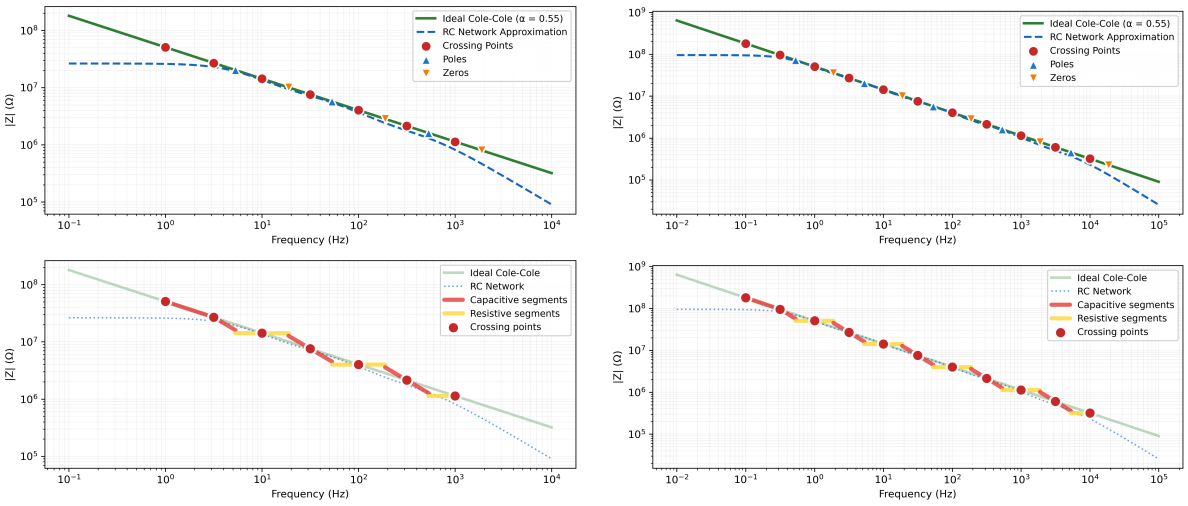

(a) 3-branch approximation (10 Hz–1000 Hz)

(b) 5-branch approximation (0.1 Hz–10 kHz)

Figure 4.24: Oustaloup approximations with extended frequency range

The extended approximation positions the zigzag transitions of the Oustaloup method further from the frequencies of interest. This allows scaling to different levels of modelling fidelity based on application requirements.

Applying the Cole-Cole parameter extraction and Foster synthesis procedures with five RC branches produces a distributed representation of the frequency-dependent behaviour. The resulting network parameters span five frequency decades with logarithmic spacing, as shown in Table 4.7.

Table 4.7: Five-branch Foster network parameters for the validation MOV

Branch	C_k (nF)	R_k (MΩ)	τ_k (ms)	$f_{z,k}$ (Hz)
1	1.57	53.80	84.493	1.9
2	0.60	14.04	8.449	18.8
3	0.22	3.92	0.845	188.4
4	0.08	1.08	0.084	1883.6
5	0.04	0.22	0.008	18836.5
External components:				
G_∞	0.019 nS (53 690 MΩ)			
C_∞	2.0 nF			
C_s	2.5 nF			

The distribution of relaxation times across five decades changes the high-frequency conductance G_∞ from 0.831 nS in the three-branch model to 0.019 nS in the five-branch implementation. This reduction occurs as the additional RC elements represent the high-frequency behaviour.

The five-branch implementation improves impedance matching. As shown in Figure 4.25, the model tracks both resistance and reactance across the measurement frequency range. The resistance matching shows mean absolute error reducing from 31.8 % to 16.7 %. At 50 Hz,

the error decreases from −24.1 % to −6.0 %.

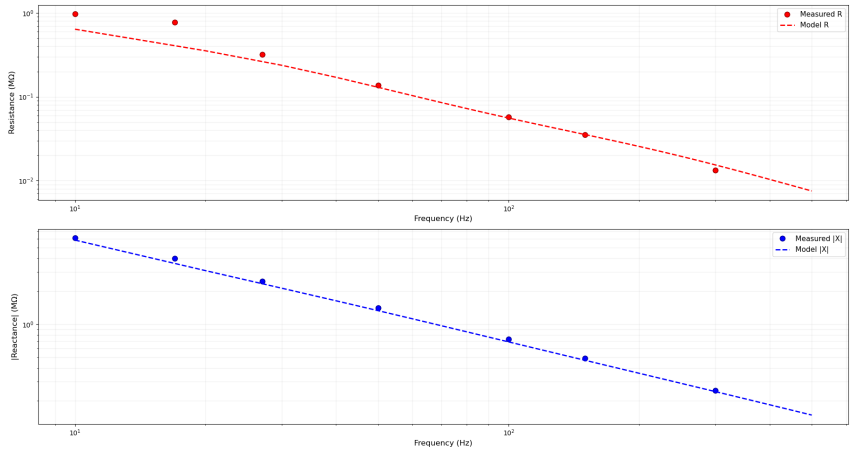

Figure 4.25: Impedance response comparison: five-branch model versus measurements at 3 kV

Table 4.8: Impedance validation error comparison between three-branch and five-branch models

Frequency (Hz)	Resistance Error (%)		Reactance Error (%)	
	3-Branch	5-Branch	3-Branch	5-Branch
10	-64.0	-34.7	-0.6	-4.7
17	-63.3	-43.3	1.7	-1.7
27	-36.0	-14.7	0.0	-2.3
50	-24.1	-6.0	-4.5	-5.7
100	-15.7	-2.8	-4.6	-5.2
150	-8.6	0.8	-3.8	-4.3
300	11.0	14.8	-2.1	-2.5
Mean Absolute Error				
Resistance	31.8 %	16.7 %	–	–
Reactance	–	–	2.5 %	3.8 %

The reactance error increases from 2.5 % to 3.8 %, remaining within acceptable bounds for practical applications. This degradation in capacitive representation results from the optimization algorithm distributing effort across more parameters.

Validation under operating conditions shows that while the five-branch model maintains correct frequency ordering and exponential voltage dependence, it exhibits consistent positive bias in power and energy predictions across all operating conditions.

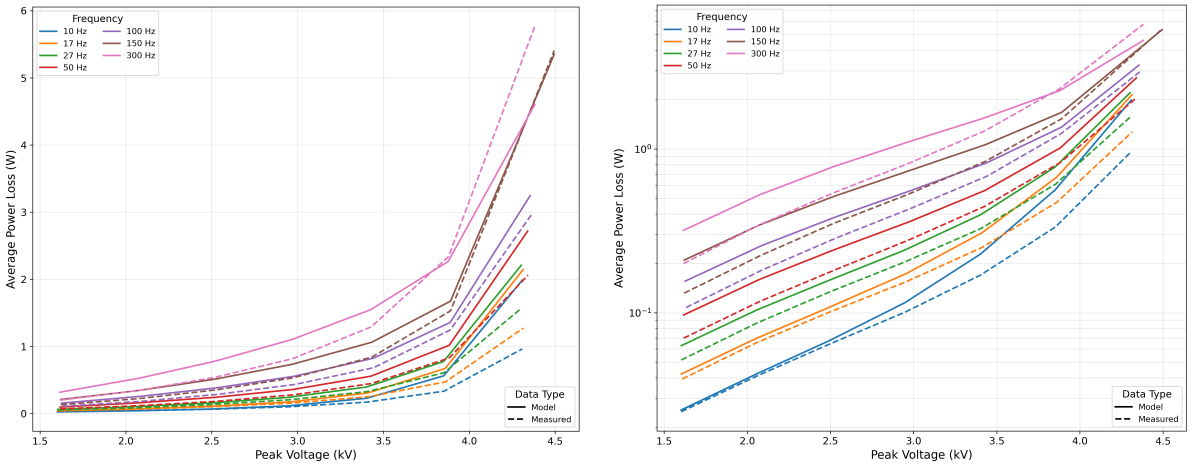

(a) Linear scale **(b)** Logarithmic scale

Figure 4.26: Average power loss comparison for the five-branch model

This bias contrasts with the three-branch model's variable error patterns. While the three-branch model exhibits both positive and negative deviations depending on frequency and voltage, the five-branch model consistently overestimates losses. This uniform behaviour indicates that the sequential optimization approach—first fitting impedance parameters, then extracting nonlinear characteristics—requires refinement when applied to higher-order networks.

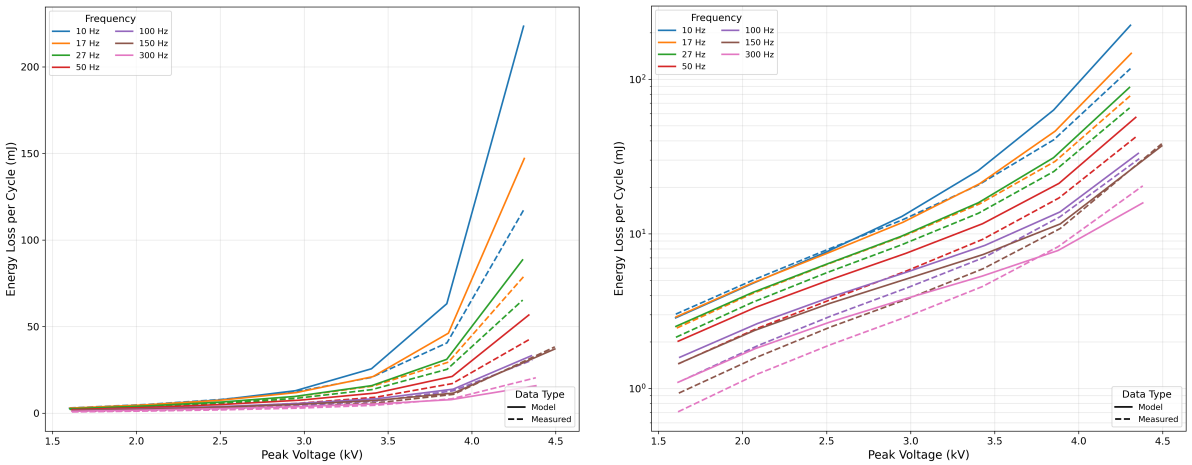

(a) Linear scale **(b)** Logarithmic scale

Figure 4.27: Energy per cycle comparison for the five-branch model

The energy per cycle comparison confirms this systematic behaviour. Across all frequencies and voltage levels, the model maintains correct relationships but with consistent overestimation. The preservation of frequency ordering and voltage scaling shows that the physics representation remains valid, while the uniform bias indicates a calibration issue rather than a modelling error.

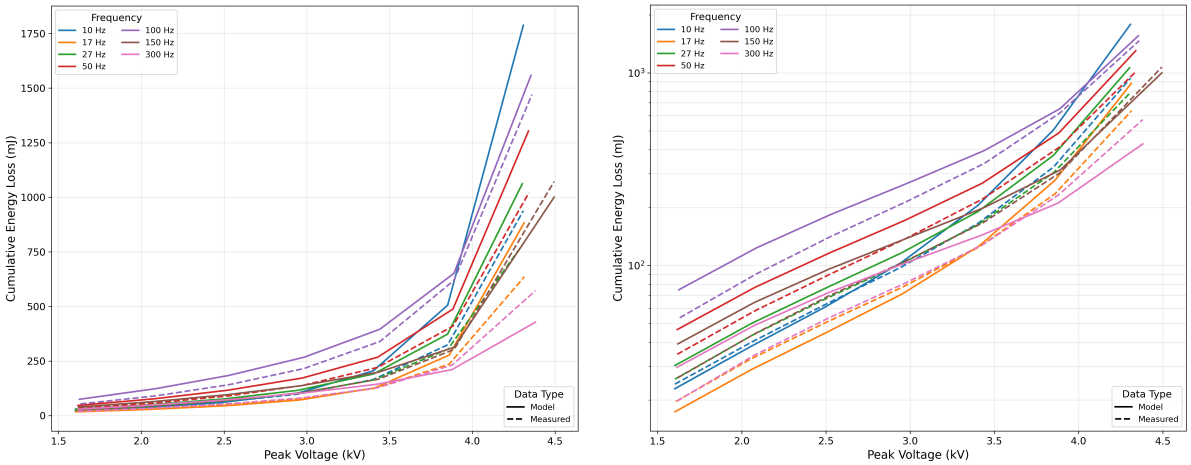

(a) Linear scale **(b)** Logarithmic scale

Figure 4.28: Cumulative energy loss comparison for the five-branch model

Quantitative analysis shows a trade-off between consistency and accuracy. The five-branch model exhibits reduced standard deviations, particularly at lower frequencies where the three-branch model shows greatest variability. At 10 Hz, the standard deviation decreases from $\pm$54.6 % to $\pm$40.2 % for power predictions. This improved consistency comes with higher mean absolute error: 29.5 % for power and 29.0 % for energy, compared to 22.4 % and 22.0 % for the three-branch model.

Table 4.9: Error statistics comparison between three-branch and five-branch models

Frequency (Hz)	3-Branch Model		5-Branch Model	
	Energy Error (%)	Power Error (%)	Energy Error (%)	Power Error (%)
10	-7.9 $\pm$ 50.4	-0.3 $\pm$ 54.6	23.2 $\pm$ 37.0	33.4 $\pm$ 40.2
17	7.1 $\pm$ 37.2	-2.6 $\pm$ 33.9	36.1 $\pm$ 25.8	23.8 $\pm$ 23.5
27	-3.5 $\pm$ 16.1	0.3 $\pm$ 16.8	19.4 $\pm$ 7.7	24.1 $\pm$ 8.0
50	10.9 $\pm$ 7.1	9.6 $\pm$ 7.0	32.5 $\pm$ 5.9	30.9 $\pm$ 5.8
100	10.2 $\pm$ 9.1	11.8 $\pm$ 9.3	25.2 $\pm$ 14.2	27.1 $\pm$ 14.4
150	18.6 $\pm$ 17.5	21.3 $\pm$ 17.9	30.1 $\pm$ 22.0	33.1 $\pm$ 22.5
300	18.7 $\pm$ 27.1	22.0 $\pm$ 27.8	23.4 $\pm$ 28.8	26.9 $\pm$ 29.6
Overall MAE	22.0	22.4	29.0	29.5

The analysis indicates that increasing Foster network order from three to five branches reduces impedance matching error from 31.8 % to 16.7 %, while energy prediction error increases from 22.0 % to 29.0 %. This trade-off results from the sequential parameter extraction methodology, which optimizes impedance fit before determining nonlinear characteristics. The consistent positive bias observed in the five-branch model contrasts with the variable errors of the three-branch implementation.

These results suggest that higher-order Foster networks require modified parameter extraction approaches. Multi-objective optimization that simultaneously considers impedance matching and energy prediction accuracy may better utilize the additional degrees of freedom provided by five-branch networks.

4.5. Summary

This chapter presented a comprehensive validation of the frequency-dependent surge arrester model through three complementary approaches. The impedance response validation examined the model's small-signal behaviour, comparing the Foster network's ability to reproduce frequency-dependent resistance and reactance characteristics against measured data. This analysis revealed the model's strengths in capacitive reactance representation while identifying systematic patterns in resistance deviations.

The full nonlinear validation extended the analysis to realistic operating conditions, examining power dissipation and energy absorption across 55 test points spanning multiple frequencies and voltages. This validation demonstrated how the interaction between frequency-dependent linear elements and voltage-dependent nonlinear conduction manifests in practical surge arrester operation. The analysis covered instantaneous power, cumulative energy, average power, and energy per cycle metrics to provide a complete characterisation of model performance.

The methodology validation tested the transferability of the modelling approach by applying it to a geometrically distinct MOV block with different electrical characteristics. This validation examined whether the impedance spectroscopy and Foster network synthesis procedures could adapt to different device configurations without manual intervention. The comparative analysis between the two MOV blocks revealed how device-specific characteristics influence model accuracy and error patterns, demonstrating that the methodology can be applied to other MOV blocks as well.

5

Discussion

The experimental characterisation and frequency-dependent modelling of metal oxide surge arresters presented in this thesis provides insights into surge arrester behaviour under harmonic-rich conditions. This chapter synthesizes these technical findings into practical guidance, interpreting results within the broader context of evolving power systems while addressing protection challenges facing cable-rich networks.

The comprehensive investigation spanning characterisation, modelling, and validation has shown three findings. First, MOV resistance exhibits frequency dependence, decreasing from 10 Hz to 500 Hz—differing from the assumption of frequency-independent parameters. Second, the decomposition of surge arrester impedance demonstrated that dielectric losses in the grain boundary network contribute to energy dissipation under harmonic-rich conditions. Third, the development of a fractional-order circuit model provides an adaptable framework for analysing surge arrester performance under temporary overvoltages.

Validation results show the model's applicability with mean absolute errors of 7 % at power frequency while maintaining proper frequency ordering across the spectrum. The application to geometrically distinct MOV blocks (22 % mean error for the second device versus 42 % for the first) indicates the methodology's transferability. The frequency dependence of the resistance shows the distributed nature of grain boundary relaxation in MOVs. At 300 Hz—the sixth harmonic—resistance drops to approximately 10 % of its power frequency value, suggesting harmonic currents experience lower impedance and possibly disproportionate heating effects.

The Cole-Cole framework captures this distributed nature through power-law frequency dependence. The frequency independence of non-linear resistance confirms its electronic origin, while the frequency dependence of an MOV reveals the dielectric nature of the pre-breakdown regime.

The validation methodology showed both agreement and systematic deviations. Accuracy at 50 Hz (−7 % mean error) validates the approach for traditional applications, providing conservative margins for the fundamental frequency. However, at harmonic frequencies, the model tends toward overestimation with errors reaching 50 %.

Investigation of a five-branch network demonstrated the methodology's scalability, achieving

impedance matching improvement from 32 % to 17 % error. The observed systematic behaviour in large-signal performance (29 % error with consistent overestimation) reveals that as Foster network complexity increases, the interaction between frequency-dependent linear elements and voltage-dependent nonlinear conduction becomes more pronounced. This consistent performance characteristic—contrasting with the three-branch model's variable error patterns—may be advantageous for practical implementations where systematic calibration is preferable. The model's ability to maintain steeper slopes near the knee point confirms its value in complementing existing surge arrester models, particularly in the leakage current region important for TOV assessment.

The investigation operated within several defined boundaries: exclusive use of pure sinusoidal waveforms (though real conditions contain capacitive harmonics that become problematic at higher frequencies where nonlinear current exhibits distortion, as detailed in Appendix A); absence of temperature-dependent effects on ZnO grain boundaries; and characterisation limited to two MOV blocks of different geometries. The thesis time frame prioritized establishing and validating the frequency-dependent methodology across multiple devices rather than implementing advanced parameter optimization techniques, providing a foundation for future improvements through computational methods as suggested in the subsequent Chapter 6.

Despite limitations, the research delivers contributions. The automated measurement pipeline enables rapid characterisation with consistent accuracy. The systematic capture of frequency-dependent behaviour provides a comprehensive dataset linking electrical response to physical mechanisms. Model transferability across different MOV geometries confirms adaptability. Identifying previously neglected dielectric loss mechanisms opens new paths for design optimization, while ATP-EMTP implementation ensures practical applicability.

Manufacturers should include harmonic characterisation beyond traditional power-frequency testing. While accelerated ageing tests at MCOV and standard TOV tests remain necessary, testing protocols should expand to include harmonic distortion (particularly up to the 6th harmonic frequency) that reflects actual field conditions. The frequency-dependent model developed in this thesis enables quantification of harmonic impacts on energy dissipation. Comparing losses under harmonic-rich conditions against power-frequency TOV curves reveals when standard ratings may underestimate electrical stress—important for product specifications and warranty considerations.

Additionally, as environmental factors increase leakage current over time, frequency-dependent dielectric behaviour becomes more significant. Modelling these characteristics provides understanding of surge arrester performance throughout service life, enabling better end-of-life prediction and maintenance recommendations. This characterisation approach ensures surge arresters can protect modern cable-rich networks where harmonic content is expected, not exceptional.

For Grid Operators, harmonic assessment should become part of TOV evaluation. Operators should map system resonant frequencies and evaluate surge arrester exposure at critical points. Where high harmonic content is unavoidable, derating factors from the frequency-dependent model guide appropriate selection. Modern digital fault recorders enable post-event analysis to identify arresters experiencing accelerated ageing from harmonic exposure.

For Protection Engineers, The frequency-dependent model provides the quantitative frame-

work for answering “Do we need additional MOV blocks?” Engineers should apply these models for cable-rich networks, high renewable penetration systems, or resonance-prone configurations. When applying TOV duration curves, engineers should calculate equivalent TOV stress for complex waveforms—a 1.4 p.u. overvoltage with harmonic content may produce stress equivalent to higher magnitude power frequency-only overvoltage.

IEC 60099-4 could benefit from updates to address frequency-dependent phenomena. Multi-frequency characterisation may warrant consideration for cable-rich or harmonic-prone applications. Section 8.8 could be expanded to incorporate harmonic content in test waveforms. Future methodologies might translate complex TOV waveforms to standardized ratings through frequency decomposition, application of loss factors, and energy integration.

The frequency-dependent framework addresses assessing surge arrester performance under complex TOVs. The traditional envelope method is inadequate when frequency-dependent losses vary by orders of magnitude. The translation methodology from complex TOV to equivalent thermal stress provides quantitative tools for evaluation. Engineers can calculate energy dissipation for specific configurations, compare against thermal capacity, and determine required safety margins.

Understanding that harmonic impedance can be up to 80 %–90 % lower than power-frequency values explains traditionally-selected arrester vulnerability. Selection procedures should consider frequency content alongside magnitude. The model enables coordination studies accounting for frequency-dependent effects throughout the protection scheme. Risk mitigation options include additional thermal margin, harmonic exposure monitoring, or system reconfiguration to shift resonant frequencies.

The understanding of surge arrester behaviour translates to improved grid reliability, particularly for cable-rich urban networks and renewable integration projects. Offshore wind farms, solar installations, and HVDC terminals all create conditions where frequency-dependent effects are significant. By preventing failures that could cascade into widespread outages, the research contributes to infrastructure resilience while removing technical barriers to renewable deployment.

Economic benefits arise from avoided failure costs and optimized protection investments. The frequency-dependent model enables specification based on actual conditions rather than conservative assumptions, potentially reducing protection costs while improving survivability. Environmental benefits include reduced equipment failures and extended asset life, supporting sustainable infrastructure operation and climate goals through renewable integration.

Continuing cable deployment trends increase the relevance of frequency-dependent modelling.The proposed framework provides tools for these evolving systems while enabling international knowledge sharing through standardized approaches. The shift from impulse-focused to harmonic-aware protection represents a change in approach. Traditional engineering emphasized brief, high-magnitude events; modern systems face sustained, harmonic-rich stresses. This research provides quantitative tools for this approach.

6

Conclusions

This thesis has addressed a vulnerability in modern power system overvoltage protection through the investigation of metal oxide surge arrester behaviour under harmonic-rich temporary overvoltage conditions. The research methodology progressed through three interconnected phases: multi-frequency characterisation revealing frequency-dependent loss mechanisms, development of fractional-order circuit models capturing distributed grain boundary physics, and validation demonstrating transferability across different surge arrester geometries. Through automated signal processing pipelines processing 168 individual test scenarios spanning frequencies from 10 Hz to 500 Hz and voltages from 3 kV to 10 kV peak, this work extracted frequency-dependent impedance characteristics that differ from conventional surge arrester models. The investigation revealed that dielectric losses in grain boundary networks also contribute to stress under harmonic resonance conditions.

The results extend beyond theoretical understanding to practical implementation. By developing Foster network synthesis methodologies compatible with ATP-EMTP constraints, this research bridges the gap between fractional-order impedance theory and application. The frequency-dependent models represent the interaction between linear capacitive behaviour and nonlinear resistive conduction, maintaining accuracy across three orders of magnitude in power dissipation. The validation framework, applied to MOV blocks with contrasting geometries—one cylindrical (64.5 mm diameter, 44 mm height) and one disc-shaped (100 mm diameter, 22.5 mm height)—demonstrates the methodology's adaptability while revealing specimen-specific variations in grain boundary characteristics. By quantifying how surge arresters behave under harmonic stress, this work provides tools for protecting cable-dominated grids where resonances are inevitable.

Research Questions Addressed

Research Question 1: *"How can the behaviour of a surge arrester be captured and accurately classified for analysis?"*

Chapter 2 addressed this question through the development of an automated signal processing framework that transforms raw oscilloscope measurements into validated electrical parameters. The methodology employs phase-sensitive decomposition utilizing voltage integration

as a stable phase reference, enabling separation of capacitive currents (leading voltage by 90°) from resistive components. The framework ensures parameter extraction across the entire 10 Hz–500 Hz measurement range. The classification system decomposed total current into three distinct components: capacitive current revealing frequency-stable but voltage-dependent behaviour, linear resistive current exhibiting frequency dependence but voltage independence, and nonlinear resistive current showing voltage dependence while remaining frequency-independent. This decomposition revealed that grain boundaries behave as lossy dielectrics with distributed relaxation times.

Research Question 2: *"How can we develop a frequency-dependent surge arrester model to better evaluate TOV duration curves and improve selection of surge arresters in their respective applications?"*

Chapter 3 delivered a solution through the synthesis of fractional-order equivalent circuits implemented within transient simulation constraints. The modelling framework employs Cole-Cole impedance representation to capture the distributed nature of grain boundary relaxation, then transforms this fractional-order behaviour into a Foster-II network topology realizable in ATP-EMTP. By incorporating frequency-dependent RC branches in parallel with voltage-dependent nonlinear resistance, the model reproduces both the capacitive-to-resistive transition and the power-law frequency dependence observed in measurements. The three-branch model demonstrated consistent performance from DC to the tenth harmonic. The five-branch model suggested possible opportunities to improve the methodology.

Research Question 3: *"What methodology can be implemented by manufacturers and grid operators alike to correctly assess the rating of surge arresters during TOV scenarios to ensure grid resiliency?"*

Chapter 4 and Chapter 5 established a framework for evaluating surge arrester adequacy under complex overvoltage conditions. The validation process, comparing model predictions against experimental measurements across 55 distinct test configurations, confirmed the methodology's reliability with mean absolute errors of 22 % for average power and 7 % at power frequency. The comparative analysis between two geometrically distinct MOV blocks demonstrated methodology transferability—the second device achieved 22 % mean error versus 42 % for the first, confirming adaptability to different surge arrester designs. The investigation revealed that harmonic currents at 300 Hz experience resistance values approximately 10 % of those at power frequency, leading to disproportionate heating effects that conventional models cannot predict. This framework enables translation of complex harmonic-rich TOV waveforms into equivalent thermal stress assessments, providing quantitative tools for surge arrester selection that account for the frequency-dependent losses in modern cable-rich networks.

Main Research Question: *"How can we estimate electrical parameters and develop accurate modelling for Metal Oxide Surge Arresters operating under low-order harmonic resonance conditions to ensure their appropriate selection and application?"*

This thesis demonstrates that surge arrester assessment under harmonic resonance conditions requires consideration of frequency-dependent impedance characteristics arising from distributed grain boundary relaxation processes. Through investigation progressing from characterisation to validated modelling and practical application, the research establishes that a

frequency-dependent model can represent surge arrester behaviour when harmonic content is present. The developed methodology captures frequency-dependent effects through automated parameter extraction revealing a resistance variation across the 10 Hz–500 Hz spectrum, represents distributed relaxation via fractional-order circuit models compatible with ATP-EMTP, and validates predictions through comparison confirming model accuracy while identifying systematic deviations that inform practical application. The frequency dependence of resistance reveals the dielectric nature of grain boundary networks.

Recommendations for Future Research

Based on the work presented in this thesis, the following recommendations can guide future research in this field. While the proposed methodology operates within its framework, it represents a first step toward practical application. The objective remains in developing a surge arrester model that can improve TOV duration curves.

- Develop measurement setups with higher current capability to characterise surge arresters above their rated voltage, capturing the transition to full conduction. Improved voltage and current transducers would also enhance resolution in the high-impedance regime where measurement uncertainty currently limits accuracy.

- Advance the analysis beyond pure sinusoidal excitation by implementing frequency-domain decomposition techniques. Since operational surge arresters experience harmonic content, the methodology should evolve to preserve these harmonic components rather than assuming purely sinusoidal capacitive current.

- Integrate temperature-dependent grain boundary characteristics and hysteresis effects into the frequency-dependent model. This multi-physics approach would capture the coupled electrical, thermal, and mechanical behaviour that governs surge arrester performance during extended TOV events.

- Implement optimization frameworks for automated parameter extraction, potentially employing machine learning techniques to transform impedance measurements into Foster network parameters across the complete operational envelope.

- Extend validation to encompass a broader range of surge arrester types, manufacturers, and voltage ratings to confirm methodology transferability.

- Investigate how ageing mechanisms impact frequency-dependent behaviour by tracking impedance changes across the frequency spectrum over time. Since harmonic currents experience lower resistance paths, long-term studies should determine whether grain boundary degradation occurs differently under harmonic stress compared to power frequency operation, enabling frequency-specific ageing models.

- Correlate model predictions with field performance data from cable-rich networks experiencing harmonic resonances. Documented surge arrester failures should be analysed to validate thermal stress calculations and establish practical safety margins for arrester selection.

- Establish standardized depolarization protocols for MOV blocks through investigation of DC bias voltage levels, duration, and recovery periods. While preliminary results using $0.003 \times V_{\text{ref}}$ DC bias achieved approximately 70 % depolarization improvement, studies should determine parameters across different MOV formulations and voltage classes. This would enable consistent initial conditions for frequency-dependent characterisation and improve measurement reproducibility across laboratories.

References

[1] Data Center Dynamics. *Virginia Narrowly Avoided Power Cuts When 60 Data Centers Dropped Off the Grid at Once*. July 2024. URL: https://www.datacenterdynamics.com/en/news/virginia-narrowly-avoided-power-cuts-when-60-data-centers-dropped-off-the-grid-at-once/.

[2] NERC. *Incident Review: Considering Simultaneous Voltage-Sensitive Load Reductions*. North American Electric Reliability Corporation. Technical Report. 2025. URL: https://www.nerc.com/pa/rrm/ea/Documents/Incident_Review_Large_Load_Loss.pdf.

[3] About the Netherlands. *How Much of the Netherlands is Actually Living Below Sea Level?* 2024. URL: https://aboutnl.com/how-does-the-nederlands-live-below-sea-level/.

[4] CIGRE Working Group C4.502. *Power System Technical Performance Issues Related to the Application of Long HVAC Cables*. Technical Brochure 556. CIGRE. 2013.

[5] Christiaan Engelbrecht et al. In: *Challenges and Mitigations for the Energization of Large Offshore Grids in the Netherlands*. June 2019.

[6] NPR. *Power Outages Can Be Risky for Those Relying on Medical Devices*. 2024. URL: https://www.npr.org/sections/shots-health-news/2024/07/11/nx-s1-5035514/power-outage-safety-medical-devices-equipment.

[7] C. Dominianni et al. “Health Impacts of Citywide and Localized Power Outages in New York City”. In: *Environmental Health Perspectives* 126.6 (2018). URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6084843/.

[8] J.A. Casey et al. “Power Outages and Community Health: A Narrative Review”. In: *Current Environmental Health Reports* 7 (2020), pp. 371–383. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC7749027/.

[9] J. A. Martinez-Velasco. *Power System Transients: Parameter Determination*. 1st. Boca Raton, FL, USA: CRC Press, 2010. DOI: 10.1201/9781420065305.

[10] O. Lennerhag et al. “Towards an Impedance-Based Criterion for Efficient Analysis of Resonant Overvoltages in the Swedish Transmission System”. In: *2022 20th International Conference on Harmonics & Quality of Power (ICHQP)*. IEEE. 2022.

[11] CIGRE Working Group C4.46. *Evaluation of Temporary Overvoltages in Power Systems due to Low Order Harmonic Resonances*. Technical Brochure 913. 2023.

[12] CIGRE Working Group 33.06. *Metal Oxide Surge Arresters in AC Systems*. Technical Brochure 60. 1991.

[13] TDK Electronics AG. *Energy Varistors: Product Brief*. Tech. rep. Accessed 24 June 2025. EPCOS AG, 2018. URL: https://www.tdk-electronics.tdk.com/download/173738/2a40555e614a4af140f9480d2bc262ed/energy-varistors-pb.pdf.

[14] International Electrotechnical Commission (IEC). *Surge Arresters – Part 4: Metal-oxide Surge Arresters Without Gaps for A.C. Systems*. IEC 60099-4:2014, Geneva, Switzerland. 2014.

[15] R. Grünbaum, M. Halonen, and S. Rudin. “Power factor, ABB static VAr compensator stabilizes Namibian grid voltage”. In: (Jan. 2003), pp. 43–48.

[16] M. Halonen et al. “SVC for resonance control in NamPower electrical power system”. In: *2001 Power Engineering Society Summer Meeting. Conference Proceedings (Cat. No.01CH37262)*. Vol. 2. 2001, 860–865 vol.2. DOI: 10.1109/PESS.2001.970168.

[17] K. Velitsikakis and I. Tannemaat. “Surge Arrester Stresses due to Harmonic Resonance Temporary Overvoltages in Transmission Systems – A Case Study of the Dutch Grid”. In: *Proceedings of the International INMR Conference*. Berlin, Germany, 2022.

[18] Ajibola Akinrinde, Andrew Swanson, and Innocent Davidson. “Investigation and Mitigation of Temporary Overvoltage Caused by De-Energization on an Offshore Wind Farm”. In: *Energies* 13 (Aug. 2020), p. 4439. DOI: 10.3390/en13174439.

[19] CIGRE. *Temporary overvoltage withstand characteristics of extra high voltage equipment*. ELECTRA Paper 179. Paris. 1998.

[20] P. Bokoro and I. Jandrell. “Failure Analysis of Metal Oxide Arresters under Harmonic Distortion”. In: *SAIEE Africa Research Journal* 107.3 (Sept. 2016), pp. 146–152. DOI: 10.23919/SAIEE.2016.8532240.

[21] A. Schei and A. Ekström. *Temporary Overvoltages: Causes, Effects and Evaluation*. CIGRE Paper 33-210. Presented at the 1990 CIGRE Session. Paris, France: CIGRE Working Group 33.10 and IEEE Task Force on TOV, 1990.

[22] W. Schmidt et al. “Behaviour of MO-surge-arrester blocks to fast transients”. In: *IEEE Transactions on Power Delivery* 4.1 (1989), pp. 292–300. DOI: 10.1109/61.19216.

[23] Krzysztof Duda. “Interpolation Algorithms of DFT for Parameters Estimation of Sinusoidal and Damped Sinusoidal Signals”. In: *Fourier Transform – Signal Processing*. IntechOpen, 2011. DOI: 10.5772/16154. URL: https://www.intechopen.com/chapters/16154.

[24] Rado Lapuh. *Sampling with 3458A*. Left Right d.o.o., 2016. ISBN: 978-961-94476-0-4. URL: https://www.amtest-tm.com/en/product/r-lapuh-sampling-with-3458a-isbn-978-961-94476-0-4/.

[25] Gerald Goertzel. “An Algorithm for the Evaluation of Finite Trigonometric Series”. In: *The American Mathematical Monthly* 65.1 (1958), pp. 34–35. DOI: 10.1080/00029890.1958.11989143.

[26] He Jinliang. *Metal Oxide Varistors: From Microstructure to Macro-Characteristics*. Wiley, Apr. 2019. ISBN: 9783527333820. DOI: 10.1002/9783527684038.

[27] A. K. Jonscher. *Dielectric Relaxation in Solids*. Comprehensive theory of dielectric absorption (relaxation) in solid dielectrics; links slow polarization phenomena to energy dissipation (dielectric loss) across a broad frequency spectrum. London: Chelsea Dielectrics Press, 1983. ISBN: 0950871109.

[28] CIGRE WG A3.17. *Metal Oxide (MO) Surge Arresters - Stresses and Test Procedures*. Technical Brochure 544. Reference: 544. CIGRE, Oct. 2013. URL: https://www.e-cigre.org/publications/detail/544-metal-oxide-mo-surge-arresters-stresses-and-test-procedures.html.

[29] Kenneth S. Kundert. *Modeling Dielectric Absorption in Capacitors*. Version 2e. Originally written March 1982, revised October 2001. Aug. 2021. URL: https://designers-guide.org/modeling/da.pdf (visited on 05/21/2025).

[30] Rainer Storn and Kenneth Price. “Differential Evolution - A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces”. In: *Journal of Global Optimization* 11 (Jan. 1997), pp. 341–359. DOI: 10.1023/A:1008202821328.

[31] A. Oustaloup et al. “Frequency-band complex noninteger differentiator: characterization and synthesis”. In: *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* 47.1 (2000), pp. 25–39. DOI: 10.1109/81.817385.

[32] Hans Kristian Høidalen, Laszlo Prikler, and J Hall. “Atpdraw-graphical preprocessor to atp, windows version”. In: *Proceedings of International Conference on Power Systems Transients, IPST*. 1999, pp. 20–24.

[33] Meidensha. *Major Ratings of Zinc Oxide (ZnO) Blocks*. Tech. rep. Accessed 03 July 2025. Tridelta, 2018. URL: https://www.meidensha.com/products/energy/prod_01/prod_01_07/__icsFiles/afieldfile/2020/08/27/F20-011.pdf.

[34] Kenneth S Cole and Robert H Cole. “Dispersion and absorption in dielectrics I. Alternating current characteristics”. In: *The Journal of chemical physics* 9.4 (1941), pp. 341–351.

A

Appendix A

Measurement System Limitations at High Frequencies

During the characterisation of MOVs at frequencies above 300 Hz, measurement limitations were encountered due to the power supply's maximum current capability. This appendix documents these limitations to explain why certain frequency ranges could not be measured reliably.

The high-voltage amplifier used for MOV characterisation has a maximum current output capability of 20 mA. When measuring capacitive loads at high frequencies, the required current increases according to:

$$I = V \cdot \omega \cdot C = V \cdot 2\pi f \cdot C \tag{A.1}$$

As frequency increases, the capacitive current demand rises linearly. When this current demand exceeds the amplifier's 20 mA maximum capability, the supply automatically switches from voltage source mode to current source mode. This mode switching results in distorted voltage and current waveforms that do not represent the actual MOV characteristics.

Figure A.1 shows the waveform distortions observed when the power supply current limit was reached. The current waveform transforms from sinusoidal to an almost square wave, while the voltage across the MOV becomes triangular. These distortions are artifacts of the power supply mode switching, not MOV material behaviour.

The frequency at which current limitations occur depends on the MOV's capacitance. For the two blocks tested, the first block with approximately 0.5 nF capacitance reached the current limit at 500 Hz, while the second block with approximately 2.5 nF capacitance encountered limitations at 300 Hz. Table A.1 shows the measured currents at the conditions where distortion occurred.

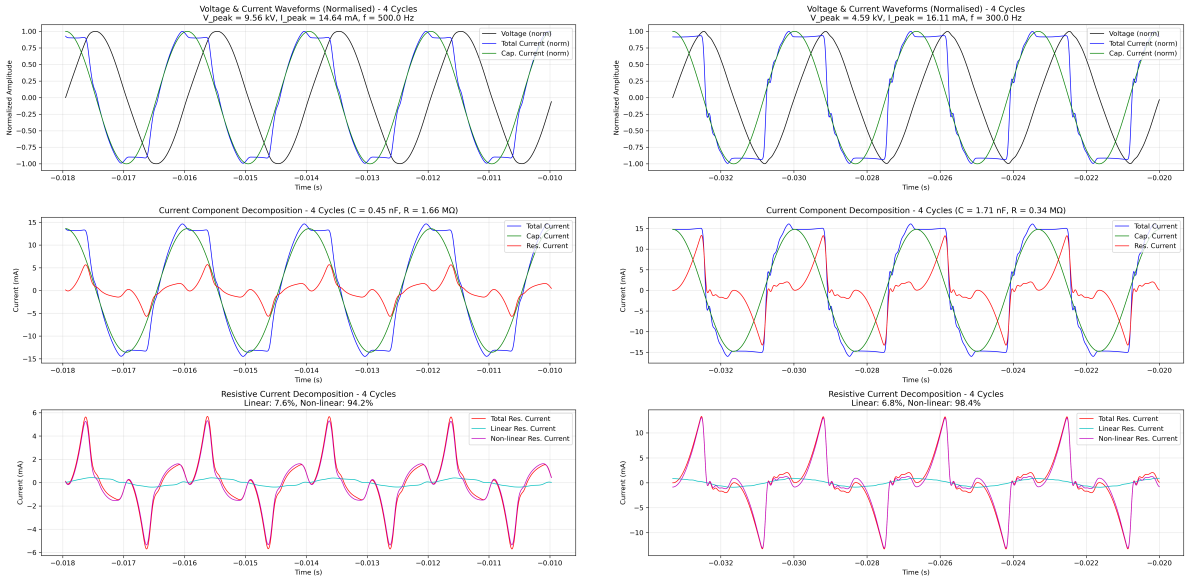

(a) First MOV block at 500 Hz

(b) Second MOV block at 300 Hz

Figure A.1: Waveform distortions caused by power supply current limitations

Table A.1: Measured currents at distortion conditions. *Oscilloscope measurements show lower values due to low-pass filtering, but actual raw currents exceed the 20 mA amplifier limit. **Current demand would far exceed amplifier capability.

MOV Block	Frequency	Voltage (peak)	Measured Current	Note
First block	500 Hz	9 kV	14.3 mA	Exceeds 20 mA (raw)*
Second block	300 Hz	4.2 kV	15.5 mA	Exceeds 20 mA (raw)*
Second block	500 Hz	—	—	Not measurable**

The measured values in the table appear below 20 mA due to the oscilloscope's low-pass filtering, which attenuates the high-frequency harmonics generated by the mode switching. The actual instantaneous currents, including all harmonic components, exceed the amplifier's 20 mA limit in all cases where distortion was observed.

These amplifier current limitations necessitated restricting measurements to frequencies where reliable voltage source operation could be maintained. Measurements were limited to frequencies where the total current, including harmonics, remained below the amplifier's 20 mA capability. The first MOV block could be characterised up to 500 Hz before reaching this limit, while the second MOV block was limited to 300 Hz due to its higher capacitance. All measurements presented in the main thesis were conducted within these limitations to ensure data reliability.

The observed waveform distortions serve as a clear indicator of when the amplifier has reached its current limit. This limitation is common when testing high-voltage capacitive devices at elevated frequencies and represents a fundamental constraint of the available test equipment rather than any material property of the MOV blocks.

B

Appendix B

Non-linear VI characteristics data

Table B.1: Non-linear V-I Characteristic Data of E64NR163E

I(A)	U(V)
8E-6	4590.11
1.15E-5	5066.15
1.66E-5	5544.14
2.39E-5	6018.00
3.49E-5	6499.32
5.02E-5	6946.60
7.23E-5	7372.70
0.000104125	7773.07
0.000152057	8157.49
0.000218962	8494.37
0.000315303	8796.40
0.000454035	9062.32
0.000663041	9299.78
0.000954775	9491.58
0.001374869	9648.35
0.001979802	9771.85
0.002891167	9867.36
0.004163261	9930.69
0.005995068	9969.08
0.008754786	9986.21

Table B.2: Non-linear V–I Characteristic Data of MZE100B23

I(A)	U(V)
2.59E-6	1911.04
4.08E-6	2067.46
6.43E-6	2229.94
1.01E-5	2397.74
1.62E-5	2576.62
2.56E-5	2752.25
4.03E-5	2929.95
6.35E-5	3108.32
0.000101733	3292.65
0.000160251	3467.58
0.00025243	3638.19
0.000397632	3802.63
0.000637398	3964.83
0.001004039	4110.75
0.001581577	4244.65
0.002491324	4364.64
0.003993558	4472.58
0.006290712	4558.75
0.009909222	4625.97
0.015884346	4674.32

Appendix C

Signal processing details for frequency and phase estimation

This section documents the signal processing methods used in Chapter 2.

Pre-processing and notation

Let $x[n]$ denote the voltage signal of length N, sampled at rate F_s with sampling interval $\Delta t = 1/F_s$. After removing the DC component (sample mean),

$$x_{\text{ac}}[n] \;=\; x[n] - \frac{1}{N}\sum_{m=0}^{N-1} x[m].$$

To reduce spectral leakage, a Hann window

$$w[n] \;=\; \tfrac{1}{2}\Big(1 - \cos\frac{2\pi n}{N-1}\Big), \qquad 0 \le n \le N-1,$$

is applied to obtain $x_w[n] = w[n]\,x_{\text{ac}}[n]$. The coherent gain of the Hann window is

$$\text{CG}_{\text{Hann}} \;=\; \frac{1}{N}\sum_{n=0}^{N-1} w[n] \;=\; \tfrac{1}{2}.$$

The N-point DFT of the windowed signal is

$$X[k] \;=\; \sum_{n=0}^{N-1} x_w[n]\, e^{-j2\pi kn/N}, \qquad 0 \le k \le N-1,$$

with bin frequency $f_k = kF_s/N$.

Hann-windowed three-bin interpolation (frequency)

The true signal frequency typically falls between DFT bins. Following [23, 24], let k be the index of the peak magnitude bin. The magnitudes of this bin and its neighbours are:

$$M_{-1} = |X[k-1]| \quad \text{(left neighbour)} \tag{C.1}$$
$$M_0 = |X[k]| \quad \text{(peak bin)} \tag{C.2}$$
$$M_{+1} = |X[k+1]| \quad \text{(right neighbour)} \tag{C.3}$$

The fractional bin offset $\delta \in (-\frac{1}{2}, \frac{1}{2})$ estimates how far the true frequency lies from bin k:

$$\delta \;=\; \frac{\frac{1}{2}\big(M_{+1} - M_{-1}\big)}{2M_0 - M_{-1} - M_{+1}}.$$

The refined frequency estimate combines the integer bin location with the fractional offset:

$$\hat{f} \;=\; \frac{(k+\delta)F_s}{N}.$$

Amplitude and phase at the refined frequency

After determining $\hat{f}$, the complex DFT is re-evaluated precisely at this frequency (not at a bin centre):

$$\hat{X} \;=\; \sum_{n=0}^{N-1} x_w[n]\, e^{-j2\pi \hat{f} n/F_s}.$$

The peak amplitude corrects for the window's attenuation:

$$\hat{V}_{\text{peak}} \;=\; \frac{2\,|\hat{X}|}{N\,\text{CG}_{\text{Hann}}} \;=\; \frac{4\,|\hat{X}|}{N},$$

where the factor of 2 accounts for single-sided spectrum and the division by CG_{Hann} compensates for window attenuation. The phase angle of the sinusoid is:

$$\hat{\phi} \;=\; \arg(\hat{X}).$$

Zero-crossing check (time domain)

As an independent verification, zero-crossings of $x[n]$ are located using linear interpolation between samples to obtain precise crossing times $\{t_i\}$. The median period between crossings:

$$\tilde{T} \;=\; \text{median}\,\{\, t_{i+1} - t_i \,\}$$

gives a time-domain frequency estimate $\tilde{f} = 1/\tilde{T}$. Agreement $|\hat{f} - \tilde{f}| < 0.5\,\text{Hz}$ validates the spectral estimate.

Goertzel evaluation (diagnostic)

The Goertzel algorithm evaluates the DFT at a single frequency f_0 without computing the full spectrum. With normalised radian frequency $\omega_0 = 2\pi f_0/F_s$, the recursive filter states are:

$$s[n] = x_w[n] + 2\cos(\omega_0)\, s[n-1] - s[n-2], \quad s[-1] = s[-2] = 0,$$

where $s[n]$ is the current state, $s[n-1]$ and $s[n-2]$ are the two previous states. After processing all samples, the complex DFT coefficient at f_0 is:

$$X_G(f_0) \;=\; s[N-1] - e^{-j\omega_0} s[N-2].$$

The amplitude and phase at frequency f_0 are:

$$V_{\text{peak}}(f_0) = \frac{2\,|X_G(f_0)|}{N\,\text{CG}_{\text{Hann}}}, \qquad \phi(f_0) = \arg\big(X_G(f_0)\big).$$

This method requires $O(N)$ operations per frequency versus $O(N \log N)$ for FFT [25]. Used here only for diagnostic harmonic checks.

Mathematical Foundations of Dielectric Relaxation and Parameter Extraction

Physical Origins of Dielectric Absorption in Surge Arresters

Dielectric absorption in metal oxide surge arresters manifests through voltage recovery phenomena. A long-charged capacitor, when briefly short-circuited, recovers to 10^{-4}–10^{-1} of its original voltage as shown in Figure C.1 [29]. A similar effect occurs in polycrystalline ZnO, but with a higher recovery voltage v_f compared to low-loss plastics. This effect dominates the losses generated in the surge arrester in the low frequency range relevant to its application, determining thermal equilibrium, energy capability, and eventually, service life.

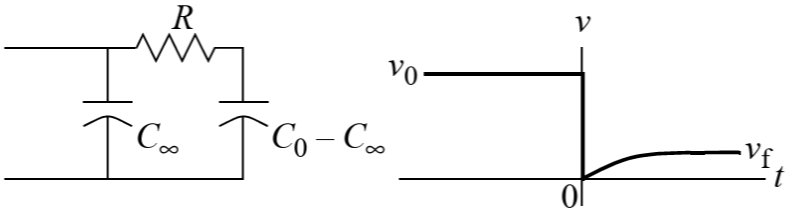

Figure C.1: Absorptive capacitor and its voltage recovery due to a negative current impulse [29]

This behaviour is usually expressed in terms of a complex permittivity:

$$\varepsilon = \varepsilon' - j\,\varepsilon'', \tag{C.4}$$

describes material response, where ε' represents stored energy and ε'' indicates dissipation per cycle.

Impedance measurements reveal frequency dependence clearly. Real impedance is proportionally inverse to the applied frequency, f^{-1} from mHz to kHz (absorption dominance), then

rises in proportion with $f^{1/2}$ above hundreds of kilohertz (electrode effects). The dissipation factor

$$\tan\delta = \frac{\varepsilon''}{\varepsilon'}$$

remains nearly constant, yielding

$$R = \frac{\tan\delta}{C\,\omega}, \tag{C.5}$$

making resistance the key diagnostic for dielectric loss.

Applied fields activate progressively slower polarisation mechanisms as frequency decreases: electronic polarisation above 10^{15} Hz, ionic polarisation in the 10^{12} Hz–10^{14} Hz range, dipolar relaxation from 10^{2} Hz to 10^{6} Hz, and—critically for surge arresters—interfacial polarisation below 10^{2} Hz. At power frequency (50 Hz–60 Hz), grain-boundary polarisation drives absorption and dissipation, occurring precisely where space-charge relaxation dominates (10^{-2} Hz to 10^{2} Hz). Multiple time constants span this behaviour: fast interface states produce half the voltage recovery within seconds (relaxation frequencies around 1 Hz–10 Hz), whereas deep traps require minutes (relaxation frequencies of 10^{-3} Hz to 10^{-2} Hz). This microsecond-to-minute continuum, corresponding to frequencies from megahertz down to millihertz, shapes the frequency-dependent impedance that determines harmonic performance [29].

From Debye to Cole-Cole: Mathematical Development

The Debye model assumes single relaxation [29]:

$$\varepsilon(\omega) = \varepsilon_\infty + \frac{\varepsilon_0 - \varepsilon_\infty}{1 + j\,\omega\tau}, \tag{C.6}$$

predicting perfect semicircles in ε''-ε' plots. ZnO measurements, as seen in this thesis however, show depressed arcs (Figure C.2) similar to the ε''-ε' plot in Figure C.2b, indicating distributed relaxation.

This can be represented by adding fractional exponents as introduced by Cole-Cole [34]:

$$\varepsilon(\omega) = \varepsilon_\infty + \frac{\varepsilon_0 - \varepsilon_\infty}{1 + \left(j\,\omega\tau_0\right)^{1-\alpha}}, \tag{C.7}$$

where α ($0 \leq \alpha \leq 1$) quantifies distribution breadth. Zero recovers Debye; larger values reflect microstructural variations.

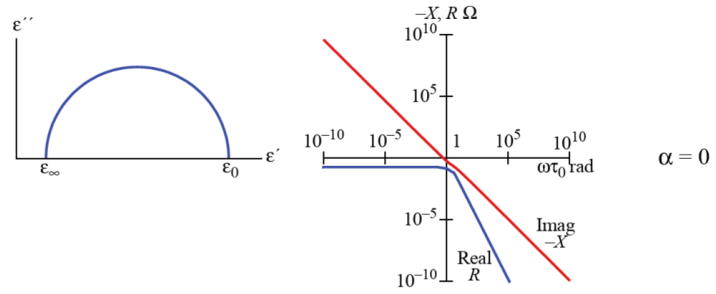

(a) Debye Model - Perfect semicircle

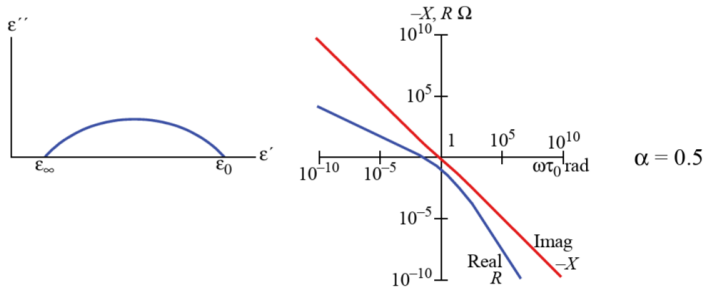

(b) Cole-Cole Model for α = 0.5 - Depressed semicircle

Figure C.2: Impedance and dielectric Cole-Cole plots showing the transition from single to distributed relaxation.

In the above figure, the Debye model produces perfect semicircles while Cole-Cole captures the depressed arcs observed in real surge arresters.

Differential Evolution Algorithm for Cole-Cole Parameter Extraction

Algorithm Structure and Implementation

Differential evolution (DE) employs population-based global optimization particularly suited for the multi-modal error surfaces encountered in Cole-Cole parameter extraction [30].

DE operates on a population of N_P parameter vectors $\mathbf{x}_{i,G}$ at each generation G, where $\boldsymbol{\theta} = [C_\infty, C_0, \tau_0, \alpha]^\mathsf{T}$. For each target vector, mutation creates a donor by adding a scaled difference of two randomly selected individuals to a third:

$$\mathbf{v}_{i,G+1} = \mathbf{x}_{r_1,G} + F(\mathbf{x}_{r_2,G} - \mathbf{x}_{r_3,G}), \tag{C.8}$$

where scale factor $F \in [0.4, 1.0]$ controls exploration magnitude. The expected value of this operation,

$$\mathbb{E}[\mathbf{v}_{i,G+1}] = \mathbb{E}[\mathbf{x}_{r_1,G}] + F\,\mathbb{E}[\mathbf{x}_{r_2,G} - \mathbf{x}_{r_3,G}],$$

reveals DE's adaptive behaviour: early generations explore broadly when $\mathbb{E}[\mathbf{x}_{r_2,G} - \mathbf{x}_{r_3,G}] \approx 0$, while later generations exploit promising directions as the population converges.

Physical constraints are enforced throughout: $C_\infty < C_0$, $10^{-6} < \tau_0 < 10^2$ seconds, and

$0 < \alpha < 1$. Solutions yielding negative resistance at any frequency are rejected.

Crossover combines donor and target vectors according to a binomial distribution:

$$P(k) = \binom{D}{k} CR^k (1 - CR)^{D-k}, \tag{C.9}$$

where $CR \in [0, 1]$ is the crossover probability, $D = 4$ is the problem dimensionality corresponding to the four Cole-Cole parameters $[C_\infty, C_0, \tau_0, \alpha]$, and $k \in \{0, 1, ..., D\}$ is the number of parameters inherited from the donor vector. This binomial scheme balances preservation of successful combinations against diversity.

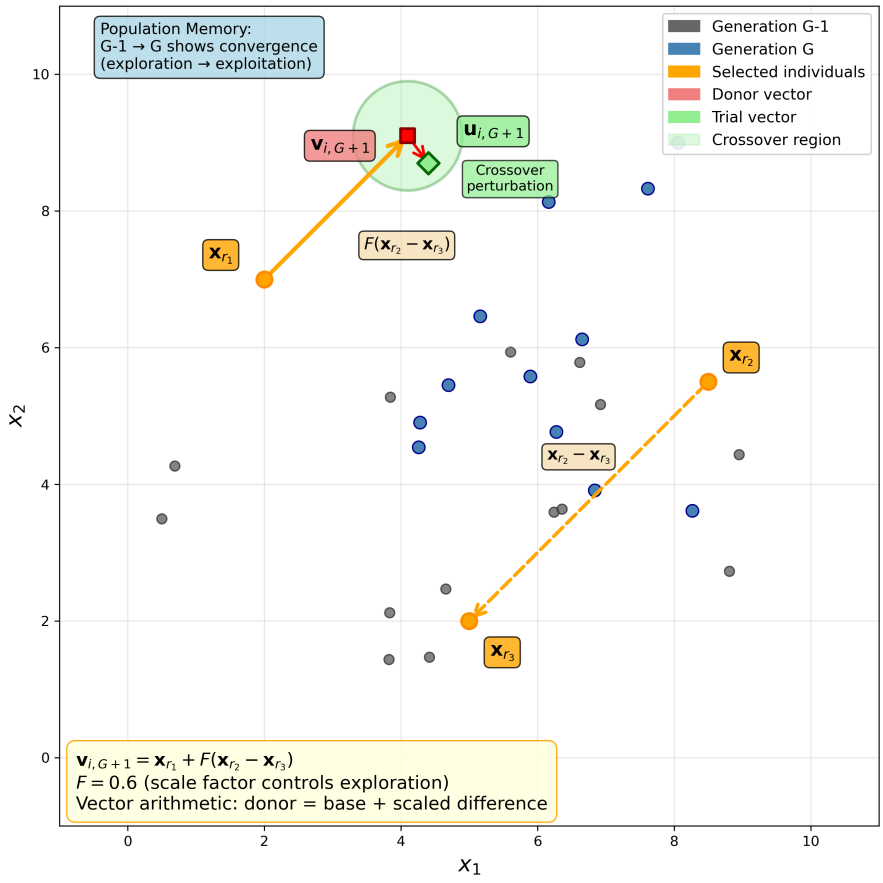

Figure C.3: Population mechanics in Differential Evolution.

The population progression illustrated in Figure C.3 demonstrates DE's inherent memory mechanism through generational inheritance. Each blue point (Generation G) derives from a corresponding grey point (Generation $G - 1$) through the mutation-crossover-selection cycle, with successful parameter combinations propagating forward while unsuccessful variations are discarded. This elitist selection ensures $f(\mathbf{x}_{i,G}) \leq f(\mathbf{x}_{i,G-1})$ for all population members, creating a ratchet effect toward optimal solutions. The visible clustering from grey to blue populations reveals the algorithm's transition from exploration to exploitation: early generations maintain diversity across the parameter space to avoid premature convergence, while selection pressure gradually concentrates the population around promising regions. This convergence pattern emerges naturally from the population dynamics—as individuals cluster, the difference vectors $(\mathbf{x}_{r_2,G} - \mathbf{x}_{r_3,G})$ in Equation C.8 become smaller, reducing step sizes and enabling

fine-tuning around discovered optima.

Figure C.3 illustrates one mutation-crossover-selection cycle. This approach handles the coupled nature of Cole-Cole parameters automatically. When τ_0 shifts the relaxation frequency, the population dynamics discover necessary adjustments in α. The ratio C_0/C_∞, controlling loss arc depth, emerges through collective search rather than manual tuning, making DE particularly suited for this challenging parameter extraction problem.

Initialization

The algorithm begins by randomly distributing NP candidate solutions across the parameter space:

$$C_{\infty,i} \sim \mathcal{U}(C_{\infty,\text{min}}, C_{\infty,\text{max}}) \tag{C.10}$$
$$C_{0,i} \sim \mathcal{U}(C_{0,\text{min}}, C_{0,\text{max}}) \tag{C.11}$$
$$\tau_{0,i} \sim \log\mathcal{U}(\tau_{\text{min}}, \tau_{\text{max}}) \tag{C.12}$$
$$\alpha_i \sim \mathcal{U}(0.1, 0.9) \tag{C.13}$$

where $i = 1, 2, ..., NP$ indexes the population members. Logarithmic sampling for τ_0 ensures uniform coverage across multiple time decades.

Mutation

For each generation g and each candidate $\mathbf{x}_i^g$, a mutant vector is created:

$$\mathbf{v}_i^{g+1} = \mathbf{x}_{r1}^g + F \cdot (\mathbf{x}_{r2}^g - \mathbf{x}_{r3}^g) \tag{C.14}$$

where $r1, r2, r3$ are distinct random indices different from i, and $F \in [0, 2]$ is the differential weight controlling mutation strength.

Crossover

The trial vector $\mathbf{u}_i^{g+1}$ combines elements from the mutant and current vectors:

$$u_{i,j}^{g+1} = \begin{cases} v_{i,j}^{g+1} & \text{if rand}(0,1) \leq CR \text{ or } j = j_{\text{rand}} \\ x_{i,j}^g & \text{otherwise} \end{cases} \tag{C.15}$$

where $CR \in [0, 1]$ is the crossover probability and j_{rand} ensures at least one parameter is inherited from the mutant.

Selection

The trial vector replaces the current vector only if it improves the objective function:

$$\mathbf{x}_i^{g+1} = \begin{cases} \mathbf{u}_i^{g+1} & \text{if } f(\mathbf{u}_i^{g+1}) \leq f(\mathbf{x}_i^g) \\ \mathbf{x}_i^g & \text{otherwise} \end{cases} \tag{C.16}$$

Objective Function Formulation

The objective function minimizes the weighted squared error between measured and modelled impedance:

$$f(\mathbf{x}) = \sum_{k=1}^{N_f} \left[w_R \left(\frac{R_{\text{model},k} - R_{\text{meas},k}}{R_{\text{meas},k}} \right)^2 + w_X \left(\frac{X_{\text{model},k} - X_{\text{meas},k}}{X_{\text{meas},k}} \right)^2 \right] \tag{C.17}$$

where N_f is the number of frequency points, and w_R, w_X are weighting factors (typically equal).

Constraint Handling

Physical constraints are enforced through penalty functions:

$$f_{\text{penalized}}(\mathbf{x}) = f(\mathbf{x}) + \lambda \cdot P(\mathbf{x}) \tag{C.18}$$

where the penalty term:

$$P(\mathbf{x}) = \max(0, C_\infty - C_0)^2 + \max(0, -\tau_0)^2 + \max(0, \alpha - 1)^2 + \max(0, -\alpha)^2 \tag{C.19}$$

Convergence Criteria

The algorithm terminates when any of the following conditions are met:

1. Maximum generations reached ($g_{\max} = 100$)
2. Population variance falls below threshold:

$$\text{var}(f(\mathbf{x}_i^g)) < \epsilon_{\text{var}} = 10^{-6} \tag{C.20}$$

3. Best fitness improvement over 10 generations:

$$\frac{|f_{\text{best}}^g - f_{\text{best}}^{g-10}|}{f_{\text{best}}^{g-10}} < \epsilon_{\text{improve}} = 10^{-4} \tag{C.21}$$

Parameter Settings for Surge Arrester Application

Based on extensive testing with surge arrester impedance data, the following parameters provide robust convergence:

Table C.1: Differential Evolution Parameters for Cole-Cole Extraction

Parameter	Value	Rationale
Standard Optimization		
Population size (NP)	50	Balance exploration/computation
Differential weight (F)	(0.5, 1.5)	Dithered for robustness
Crossover probability (CR)	0.9	High mixing for complex surface
Maximum iterations	5000	Ensure convergence
Strategy	Multiple*	Try different approaches
Refinement Stage (if needed)		
Population size (NP)	100	Larger for fine-tuning
Differential weight (F)	(0.3, 1.0)	Smaller for local refinement
Crossover probability (CR)	0.95	Very high mixing
Maximum iterations	3000	Focused refinement

*The implementation tests multiple DE strategies (”best1bin”, ”best2bin”, ”rand1bin”, ”current-tobest1bin”) and selects the one yielding the lowest error. This multi-strategy approach increases robustness against the complex error surface shown in Figure 3.2.

The algorithm employs a two-stage approach when resistance fitting is prioritized:

1. **Global search**: Wide parameter bounds with multiple strategies
2. **Refinement**: Tighter bounds around best solution with larger population

Parameter bounds adapt based on measured data characteristics.

Notes