

Physics-Based Assessment of EV Battery Degradation in V2G Services using a Coupled Electrical-Electrochemical-Thermal-Fluid Model

P.A. van Dijk



Physics-Based Assessment of EV Battery Degradation in V2G Services using a Coupled Electrical-Electrochemical- Thermal-Fluid Model

by

Patrick Armando van Dijk

in partial fulfillment of the requirements for the degree of

Master of Science

in Electrical Engineering

at the Delft University of Technology,

to be defended publicly on Thursday July 23, 2026 at 14:00 AM.

Student number:	5024021		
Project duration:	November 2025 – July 2026		
Thesis committee:	Dr. ir. G. R. Chandra Mouli	TU Delft	Supervisor
	Dr. ir. M. Rana	TU Delft	External committee member

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

All (good) things must come to an end; my time at the TU Delft is no exception, with this thesis marking the end of my Master's study in Electrical Engineering. It has been an incredibly challenging yet rewarding period upon which I will certainly look back with much glee. I am thankful for the many opportunities that have made it possible for me to end up where I am today, and I would not have had it any other way.

I would like to express my sincere gratitude to Dr. Gautham Ram Chandra Mouli for his guidance throughout this work. I am also thankful to Koen, both for developing the basis upon which this entire research stands, and for putting up with our long-winded conversations and discussions, which were of great help for navigating the, at times, vague and uncertain areas of this research. I wish you the best of luck with finishing your own research during your remaining years as a PhD candidate. Additionally, I would like to thank Dr. Moumita Rana for joining my graduation committee.

Finally, I would like to thank my family and friends. To my parents, thank you for your endless support and inspiration. By no means has my educational path been straightforward or easy, and I am glad that, in all these years, you have never attempted to steer me in a specific direction and have always encouraged me to follow my passion and motivation. I am also grateful for the many friends I have made along the way during my studies. Studying together made it possible to pass even the most difficult courses, and I will fondly remember the many moments and adventures we shared in between our academic obligations. Last, but absolutely not least, I would like to thank my dear friends at the Delft Solar Team. I truly appreciate and cherish our time together, both during our gap year, and also while finishing our studies.

*P.A. van Dijk
Delft, July 2026*

Abstract

Vehicle-to-grid (V2G) allows electric vehicles to provide grid ancillary services, but its impact on battery degradation remains unclear. Current V2G literature relies predominantly on computationally efficient, but limited, semi-empirical methods at constant temperatures, while high-fidelity models are too computationally intensive for long-term simulations. This research bridges this gap by developing a coupled physics-based thermal-electrochemical framework to accurately model battery degradation without the high computational cost of Computational Fluid Dynamics (CFD) models. A Single Particle Model with electrolyte dynamics (SPMe) is augmented with heat generation and side reaction kinetics. This is integrated with a Lumped Parameter Thermal Network (LPTN) and a 1D thermal-fluid Thermal Management System (TMS) model utilising empirical flow correlations. The single-cell framework is validated against experimental voltage and surface temperature data of cylindrical 21700 LG M50 cells under 0.1C, 0.5C, 1C and 1.5C constant current discharge experiments, while the pack-level LPTN and TMS are validated using CFD data from a module consisting of 28 cylindrical 21700 LG M50 cells with serpentine channel cooling under 1C and 5C constant current discharge experiments.

Validation confirms that the reduced-order multi-physics framework accurately replicates pack-level thermal-fluid behaviour ($R^2 > 0.94$) at a fraction of the computational cost required for full volumetric meshing. Utilising this framework, subsequent ablation studies demonstrate that, for relatively mild load profiles ($\leq 0.1C$), thermal and electrical coupling yields virtually no divergence in macroscopic degradation compared to simple isothermal models. Meaningful degradation spread only emerges at more aggressive currents ($\geq 0.75C$), where localised heat generation distinctly accelerates temperature-dependent solid electrolyte interphase (SEI) growth. Finally, long-term V2G simulations demonstrate that standard Frequency Containment Reserve (FCR) services impose a negligible thermal and degradative burden on the battery pack. Even under more extreme stress-test scenarios with tenfold power limits, heat generation remains minimal, failing to approach the thresholds required to trigger active cooling. Therefore, this research concludes that for standard V2G frequency regulation, the implementation of complex, anticipative thermal management strategies offers no tangible benefit, and standard reactive cooling logic is entirely sufficient to preserve pack health.

Contents

Preface	i
Abstract	ii
Nomenclature	v
1 Introduction	1
1.1 Background and Motivation	1
1.2 Research Scope and Objectives	2
1.3 Research Questions	3
1.4 Thesis Outline	3
2 Background	4
2.1 Electrochemical Fundamentals	4
2.2 Degradation Mechanisms	6
2.2.1 Degradation Effects	6
2.2.2 Solid Electrolyte Interphase (SEI) Growth	6
2.2.3 Lithium Plating	7
2.2.4 Particle Cracking	7
2.2.5 Degradation Feedback Loops	8
2.3 Electrochemical Modelling	8
2.3.1 Semi-Empirical	8
2.3.2 Physics-Based	10
2.3.3 Data-Driven	11
2.4 Thermal Dynamics and Coupled Effects	12
2.4.1 Heat Kinetics	12
2.4.2 Heat Transfer	13
2.4.3 Pack-Level Dynamics	14
2.5 Literature Review	15
2.6 Research Gaps and Objectives	16
3 Methodology	18
3.1 Framework Architecture	18
3.1.1 Model Selection	18
3.1.2 Delineation of Contributions	19
3.1.3 Data Flow and Domain Wiring	19
3.2 Electrochemical and Heat Kinetics	20
3.2.1 Electrochemical Principles	20
3.2.2 Heat Kinetics Augmentation	20
3.2.3 Degradation Mechanisms	21
3.3 Computational Paradigm: Acausal Networks and DAE Stiffness	22
3.3.1 Causal versus Acausal Modelling	22
3.3.2 Symbolic Compilation and Monolithic Failure	22
3.4 Pack-Level Integration and Multi-Physics Coupling	23
3.4.1 Electrical Domain: Monolithic Failure and Nodal Partitioning	23
3.4.2 Thermal Domain: Core-Shell Topology and Inter-Cell LPTN	24
3.4.3 Fluid Domain: 1D Curved-Pipe Modelling	26
3.5 Partitioned Multi-Rate Time-Stepping	28
3.5.1 Analytical and Numerical Partitioning	28
3.5.2 First-Order Hold (FOH) Co-Simulation	29
3.5.3 Dynamic Time-Stepping and Threshold Optimisation	29

3.6	Anticipative TMS Control Logic	30
3.6.1	Thermal Inertia and Control Limitations	30
3.6.2	Reactive Baseline: Hysteresis Control	30
3.6.3	Moving-Average Lookahead	30
3.6.4	Proactive Control Logic	31
3.7	Parametrisation and Validation Strategy	31
3.7.1	Electrochemical Validation	31
3.7.2	Thermal-Fluid and Module-Level Validation	32
4	Results and Discussion	33
4.1	Framework Parametrisation and Validation	33
4.1.1	Electrochemical Validation (Single-Cell)	33
4.2	Thermal-Fluid Calibration (Pack-Level)	35
4.3	Ablation Study: Impact of Coupling Domains	39
4.4	V2G Degradation Assessment	43
4.5	Sensitivity Analysis	46
5	Conclusion	49
5.1	Recommendations	50
A	Framework Parameterisation	55
A.1	Electrochemical Parameters (LG M50)	55
A.2	Thermal-Fluid and Pack-Level Parameters	56
A.2.1	Temperature-Dependent Specific Heat Capacity Polynomials	57

Nomenclature

Abbreviations

Abbreviation	Definition
BESS	Battery Energy Storage System
BMS	Battery Management System
CFD	Computational Fluid Dynamics
DAE	Differential-Algebraic Equation
DFN	Doyle-Fuller-Newman
ECM	Equivalent Circuit Model
EHC	Entropic Heat Coefficient
EV	Electric Vehicle
FCR	Frequency Containment Reserve
FOH	First-Order Hold
KCL	Kirchhoff's Current Law
KVL	Kirchhoff's Voltage Law
LAM	Loss of Active Material
LCO	Lithium Cobalt Oxide
LFP	Lithium Iron Phosphate
LLI	Loss of Lithium Inventory
LPTN	Lumped Parameter Thermal Network
MPC	Model Predictive Control
NCA	Lithium Nickel Cobalt Aluminium Oxide
NMC	Lithium Nickel Manganese Cobalt Oxide
OCV	Open Circuit Voltage
ODE	Ordinary Differential Equation
OEM	Original Equipment Manufacturer
P2D	Pseudo-Two-Dimensional
PID	Proportional-Integral-Derivative
PV	Photovoltaic
QNDF	Quasi-Constant Step-Size Backward Differentiation Formula
RES	Renewable Energy Source
RMSE	Root Mean Square Error
SEI	Solid Electrolyte Interphase
SoC	State of Charge
SoH	State of Health
SPMe	Single Particle Model with electrolyte dynamics
TMS	Thermal Management System
V2G	Vehicle-to-Grid
V2X	Vehicle-to-Everything
WLTP	Worldwide Harmonised Light Vehicles Test Proce- dure
ZOH	Zero-Order Hold

Symbols

Symbol	Definition	Unit
A	Cross-sectional or surface area	m^2
a	Hydraulic radius (fluid)	m
a_k	Active surface area density (electrodes)	m^{-1}
b	Bruggeman coefficient	—
C	Thermal capacitance	J K^{-1}
c	Lithium concentration	mol m^{-3}
c_p	Specific heat capacity	$\text{J kg}^{-1} \text{K}^{-1}$
D	Diffusion coefficient	$\text{m}^2 \text{s}^{-1}$
D_h	Hydraulic diameter	m
De	Dean number	—
f	Friction factor (Darcy/Fanning)	—
h	Convective heat transfer coefficient	$\text{W m}^{-2} \text{K}^{-1}$
I	Macroscopic electrical current	A
j	Local intercalation/reaction current density	A m^{-2}
K	Minor loss coefficient (form drag)	—
k	Thermal conductivity	$\text{W m}^{-1} \text{K}^{-1}$
k_{sei}	Side reaction rate constant	A m^{-2}
L	Length or domain thickness	m
M	Molar mass	kg mol^{-1}
m	Mass	kg
$\dot{m}$	Mass flow rate	kg s^{-1}
Nu	Nusselt number	—
P	Pressure	Pa
Pr	Prandtl number	—
Q	Heat generation / convective heat flux	W
R	Electrical or thermal resistance	Ω or K W^{-1}
R_c	Radius of curvature	m
R_k	Particle radius	m
Re	Reynolds number	—
T	Temperature	K or $^{\circ}\text{C}$
t	Time	s
t_+	Cation transference number	—
U	Open-circuit potential	V
V	Voltage	V
$\bar{V}$	Partial molar volume	$\text{m}^3 \text{mol}^{-1}$
v	Fluid velocity	m s^{-1}
z	Stoichiometry	—
α	Tuning parameter / transfer coefficient	—
ϵ	Porosity	—
η	Overpotential	V
μ	Dynamic viscosity	Pa s
ρ	Density	kg m^{-3}
σ	Electronic or ionic conductivity	S m^{-1}
ϕ	Electrical potential	V

Introduction

1.1. Background and Motivation

The global energy transition has accelerated significantly in an effort to meet strict carbon emission targets. A major consequence of this shift is the annual exponential growth of installed renewable energy sources (RESs), such as wind and photovoltaic (PV) systems, which are currently placing significant strain on existing grid infrastructure [1].

The fundamental challenge accompanying this shift is the inherent intermittency of weather-dependent generation, which introduces a severe timing mismatch between power supply and demand; localised generation typically spikes during peak daylight hours [2, 3], followed by a steep increase in demand at sundown [3]. To avoid curtailment and grid instability, operators are often forced to heavily oversize RESs to account for both daily and seasonal fluctuations [3]. Additionally, the rapid replacement of conventional synchronous fossil-fuel power plants with inverter-interfaced, voltage-based renewables significantly reduces the total rotational mass of the grid, which manifests in decreased grid stability margins and a more frequent triggering of grid congestion bottlenecks [4].

Upgrading traditional centralised transmission and distribution infrastructure to handle these decentralised and bidirectional power flows is a slow and capital-intensive process. While dedicated grid-scale lithium-ion battery energy storage systems (BESS) technically present a viable alternative to manage peak shaving and frequency control, committing to the massive capital investment required for such energy capacities is often financially infeasible. This is where vehicle-to-grid (V2G) technology offers a scalable and economical alternative.

Projections show that the majority of passenger cars in Europe will be electric by 2031 [5], which lines up with the introduction of new mandates such as the European Union's "Fit for 55" regulation that aims to reduce carbon emissions in new cars by 55% by 2031, and 100% by 2035 [1]. This mass electrification provides unique opportunities for energy storage: because passenger cars are generally stationary for approximately 95% of their operational lifespan [6], the collective of all electric vehicles (EVs) represents a highly distributed and underused reserve of energy capacity [7]. For example, Germany's current national energy storage of approximately 52 GWh could theoretically be expanded by bidirectional EV fleets to 170 GWh by 2030, and nearly 1300 GWh by 2045 [7].

By establishing bidirectional power pathways, EVs can act as virtual power plants that store excess renewable energy during times of overproduction, discharge power back into the network during peak usage, and even emulate the missing rotational inertia of legacy power plants [8]. Despite these clear macroeconomic benefits, the widespread deployment of V2G remains hindered by a social, rather than a technical, barrier: consumer acceptance and battery health anxiety [9]. Native bidirectional support and standardised operational guidelines from original equipment manufacturers (OEMs) are currently still severely lacking [10], and it is practically guaranteed that lithium-ion cells will experience irreversible degradation during active usage [8]. For consumers to confidently lease their EVs for grid ancillary services, the effect of V2G-induced degradation must be well understood; without it, formulating fair

financial compensation frameworks or resolving warranty implications becomes impossible.

Evaluating V2G degradation parameters remains a complex challenge. The vast majority of studies in current literature rely heavily on semi-empirical modelling techniques [8]. While computationally lightweight, these grey-box models are fundamentally limited by their lack of underlying physical principles, relying instead on parametrisation through curve-fitting against static experimental data. As a result, they fail to capture the highly dynamic pack-level interactions, such as thermally induced degradation feedback loops, that contribute to accelerated ageing in real-world conditions. Additionally, the wide variety of inconsistently simplified simulation scenarios leads to highly scattered and contradictory research [8].

V2G presents unique opportunities for degradation mitigation. Whereas load profiles resulting from standard driving behaviour are highly erratic and difficult to predict, V2G load profiles are typically well-organised and predetermined. This inherent predictability offers an unexploited opportunity to implement proactive, instead of reactive, thermal management systems (TMS). Traditional cooling systems generally only activate reactively once a surface temperature threshold is exceeded. As a result, they suffer from severe thermal lag due to the poor through-plane thermal conductivity of battery cells [11]. Heat generated deep within the core of the cell struggles to conduct outwards through alternating layers of active material and plastic separators before reaching the surface where the coolant can extract it.

By utilising load forecasting algorithms, an anticipative TMS can proactively adjust coolant flow rates to effectively pre-chill the system. This directly minimises transient internal temperature spikes and limits maximum thermal gradients, which theoretically translates to a significant reduction in long-term degradation. Accurately investigating such a mitigation strategy requires a shift away from grey-box approximations towards higher fidelity, computationally lightweight physics-based simulation frameworks.

1.2. Research Scope and Objectives

To address the identified limitations of grey-box estimations and the infeasibly high computational costs of three-dimensional continuum simulations, this research establishes a "middle-ground" multi-physics simulation framework designed specifically for long-term battery degradation analysis. The scope of this thesis is strictly bounded by the operational parameters of a module-level EV battery pack operating under constraints representative of typical V2G loads.

The primary aim of this research is achieved through the execution of the following four sequential technical objectives:

1. **Mathematical formulation and coupling:** augmentation of an electrochemical single-cell model with internal heat generation equations and side reaction kinetics, including solid electrolyte interphase (SEI) growth and lithium plating. This augmented model is then dynamically coupled to a thermal model to accurately capture internal temperature gradients and heat dissipation.
2. **Pack-level scaling:** integration of the single-cell models into a pack-level architecture, coupled with a fluid model that dynamically computes convective heat transfer along an active cooling channel structure.
3. **Multi-level module validation:** execution of a systematic validation process across decoupled physical domains. The isolated electrochemical-thermal cell model is validated against experimental constant-current data, while the pack-level thermal and fluid networks are benchmarked against high-fidelity reference data at varying coolant velocities and currents.
4. **Framework application and mitigation analysis:** application of the fully validated framework to execute comparative degradation scenarios. The framework assesses capacity and power fade over multi-cycle tests to evaluate the qualitative impact of thermal feedback loops. Leveraging the predictable nature of V2G load profiles, a proof-of-concept anticipative TMS is deployed to evaluate the theoretical potential of proactive cooling in dampening thermal transients and extending pack life.

1.3. Research Questions

To achieve the outlined objectives, this thesis addresses three main research questions, each of which is supported by sub-questions that are described in more detail in Section 2.6:

- **RQ1:** How can a computationally efficient multi-physics framework be formulated and validated to accurately capture coupled thermal-electrochemical degradation at the pack level?
- **RQ2:** To what extent do pack-level scaling and thermal feedback affect degradation predictions compared to standard isothermal and single-cell assumptions?
- **RQ3:** How effectively can an anticipative TMS, leveraging the inherent predictability of V2G load profiles, mitigate cyclic battery degradation?

1.4. Thesis Outline

To systematically address these research questions and construct the necessary simulation framework, this thesis is structured as follows:

- **Chapter 2 (Background):** establishes the theoretical foundation required to understand battery degradation and the modelling choices made in this research. While the chapter briefly introduces fundamental battery chemistry and mechanics to provide necessary context, its primary focus is on establishing the critical electrochemical kinetics and thermal dynamics that drive cyclic ageing. Furthermore, it details the spectrum of available modelling paradigms and outlines the fundamental fidelity-versus-computational-cost trade-off. The chapter concludes with a literature review that identifies unresolved gaps in V2G degradation assessments. These gaps directly justify the selected framework architecture and yield the specific research questions of this thesis.
- **Chapter 3 (Methodology):** details the mathematical formulation and architecture of the proposed multi-physics framework. Moving beyond the theoretical options presented in Chapter 2, this section explicitly justifies the selection of the Single Particle Model with electrolyte dynamics (SPMe), the core-shell Lumped Parameter Thermal Network (LPTN), and the 1D fluid dynamics. It details the governing equations, the wiring logic used to connect the domains, the multi-rate time-stepping strategy used to bypass mathematical stiffness, and the development of the anticipative cooling algorithm.
- **Chapter 4 (Results and Discussion):** presents the execution of the experiments and answers the research questions. It begins with the decoupled validation of the framework against experimental and reference CFD data. Following this, an ablation study evaluates the framework's various domains across multiple scales of physical complexity to isolate the impact of coupled thermal and electrical effects. Finally, it provides a qualitative assessment of the degradation mitigation achieved during comparative anticipative versus reactive V2G cooling tests.
- **Chapter 5 (Conclusion):** summarises the core findings of the research, reflects on the limitations and necessary assumptions of the established framework, and provides recommendations for future work.

2

Background

2.1. Electrochemical Fundamentals

Lithium-ion batteries, currently the most prominent type of secondary batteries [12], consist of four distinct components that are essential for understanding what enables them to repeatedly store and release energy. These components are the positive and negative electrodes, the liquid electrolyte, and the separator. The separator is a porous membrane placed between the anode and cathode that does allow ion transport, but strictly blocks electrons from passing through, which prevents an internal electrical short circuit from being created. Batteries can then be thought of as electrode sheets placed in an electrolyte, where the positive and negative electrodes are alternated with separator sheets in between, and the electrodes connect to the external terminals through highly conductive current collectors.

Backwards reasoning indicates that, for a battery to deliver energy, a positive current must flow from its positive terminal to the negative terminal through the external circuit, which implies that electrons must flow in the opposite direction. Therefore, a fully charged battery must be storing electrons at its anode. Discharging the battery causes electrons to flow from the anode to the cathode, so electrons will start to accumulate at the cathode. Nature likes stability, and batteries are no exception, so to balance the electrostatic charges the external transport of electrons must have an internal counterpart, which is the transport of lithium-ions, hence their name.




Figure 2.1: High-level schematic of a lithium-ion cell discharging. Current flows from the cathode to the anode through the external circuit, which means that electrons and lithium-ions flow in the opposite direction from the anode to the cathode, with electrons flowing through the external circuit, and lithium-ions passing through the internal liquid electrolyte.

Typically, the electrolyte of a lithium-ion battery contains a lithium salt, such as lithium hexafluorophosphate (LiPF_6), dissolved in organic solvents [13]. These solvents act as a liquid medium, allowing ions to move around freely. Whenever a non-zero current is applied, electrons will start to accumulate at one of the electrodes, which forces lithium-ions to be transported towards that same electrode through a combined process of both diffusion, caused by concentration gradients, and migration, caused by electrostatic forces.

In the electrodes, lithium-ions and electrons need to be able to be accepted and released repeatedly without the lithium-ions reacting to form permanent compounds, which requires specific structures like crystal lattices. Most commonly, anodes are made of graphite sheets, and cathodes are often made of transition metal oxides (TMOs), such as nickel-manganese-cobalt (NMC) [13]. At the anode, electrons will enter the conduction band of the graphite, causing it to become slightly negatively charged. This will attract the lithium-ions to react with the charged graphite to sit right in the middle of a hexagonal carbon ring of the graphite crystal lattice, making that region electrostatically neutral again. This process of lithium-ions entering electrode particles is also referred to as intercalation. Because the lithium-ions are loosely held by weak electrostatic interactions, they are able to safely leave the electrode particles again, which is what enables the secondary, bi-directional behaviour of these batteries.

Different materials are used at the electrodes to create the potential energy that is responsible for a battery's ability to store energy, with cathode structures typically forming stronger bonds with lithium-ions than the anode structures—this is the chemical equivalent of a boulder “storing” potential energy when it is unstably at rest on top of a high hill that neighbours a valley. The resulting combined properties of this chemical potential can be described using the open-circuit voltage (OCV), which simply is the difference between the cathode and anode potentials:

$$\text{OCV}(x, y) = U_{\text{cathode}}(x) - U_{\text{anode}}(y) \quad (2.1)$$

The electrode potentials themselves are a function of the stoichiometries, which are a measure of how much lithium is stored inside the local crystal structure. It is important to note that stoichiometry is generally not the same as the macroscopic quantity state of charge (SoC). Since the electrode potential curves typically do not completely overlap perfectly, one of the electrodes will generally be entirely filled or depleted sooner than the other. Additionally, manufacturers typically place operational voltage limits on batteries, which should prevent the electrodes from becoming entirely filled or depleted. The resulting range is what defines a battery's state of charge range.

2.2. Degradation Mechanisms

Among the many degradation mechanisms that can occur in a battery, there are three that are widely considered to be the most prominent, namely solid electrolyte interphase (SEI) film formation, lithium plating, and particle cracking. While other mechanisms might occur more slowly and independently, these three mechanisms form severe positive feedback loops and are triggered much more frequently [14].

2.2.1. Degradation Effects

The most important imperfection of batteries is their guaranteed obsolescence after a certain amount of time or number of cycles, which is typically caused by a combination of both capacity fade and power fade.

The previous section showed that energy can be stored through the process of lithium-ion intercalation, from which it can be deduced that energy capacity is directly proportional to the amount of useful lithium-ions that a battery holds. When lithium-ions become unusable through irreversible parasitic processes and are permanently trapped, it directly results in a loss of capacity. This specific physical phenomenon is defined as loss of lithium inventory (LLI) [15]. Similarly, if the active materials responsible for hosting these ions fracture and become disconnected from the main structure, they can no longer store energy. This corresponding effect is referred to as loss of active material (LAM) [15].

As byproducts from these parasitic reactions build up, they block conductive pathways, resulting in resistance increases across the battery. This leads to a power rate decrease, or power fade, which is mainly an issue for use cases such as EVs, where older batteries will hit their strict thermal limits sooner and limit the user's ability to extract full performance from the device. Rather unfortunately, degradation is a highly nonlinear process that accelerates over time due to degradation mechanisms being coupled, and degradation putting proportionally higher strain on the remaining "healthy" parts of the battery. This is one of the main reasons for why the end-of-life (EOL) boundary is often put at 80%, as past this point many batteries will typically experience a "knee point" tipping region beyond which the battery will degrade significantly faster in an exponential fashion [16].

2.2.2. Solid Electrolyte Interphase (SEI) Growth

SEI layer formation is the most famous of the three degradation mechanisms, and in part, it is actually also responsible for batteries being able to operate safely in the first place. As is implied by its name, the SEI layer is the solid byproduct of a parasitic reaction, also known as passivation, that occurs at the interphase between the solid phase electrode particles and the liquid electrolyte.

When a freshly manufactured battery is charged for the first time, electrons flow into the anode, causing the local potential to drop below the stability window of organic electrolyte solvents, such as ethylene carbonate. This creates the ideal environment for solvents and electrolyte salts to react at the anode's surface with the incoming electrons through a redox reaction that splits the molecular bonds of the solvent. The resulting solid compounds precipitate directly on the anode's surface into an inorganic layer facing the anode structure and an organic layer facing the electrolyte [17]. While the irreversible consumption of lithium-ions does directly cause LLI, the resulting inorganic layer is actually a key safety feature that ensures safe reversible lithium intercalation, as the dense molecules (such as LiF and LiO₂) are highly effective electrical insulators. Beyond a certain thickness of several nanometres, electrons will not conduct across the material, ensuring that the parasitic redox reactions will not occur anymore in the bulk electrolyte [18]. Lithium-ions, however, are still capable of reversible intercalation through microscopic defects in the inorganic layer.

Interestingly, following this line of reasoning, one could assume that passivation only occurs during the initial charge, as electrons are not able to travel through the interphase layer to reach the electrolyte for new reactions. However, experiments show that, even in the absence of other degradation mechanisms, the SEI layer continues to grow at a highly predictable rate that scales with the square root of time [14]. One of the ways in which the SEI layer can continue to grow works as follows: organic solvent molecules can slowly diffuse through the existing SEI layer towards the anode boundary until they are in close enough vicinity for electrons to tunnel towards the solvent molecule, triggering a reaction in the middle of the SEI layer [19]. These continuous forms of SEI layer formation do not require active current flow, and therefore it is also referred to as calendar aging, as it is the main driving force behind batteries degrading even when they are not actively used.

2.2.3. Lithium Plating

Similar to SEI layer formation, lithium plating also acts through chemical pathways with undesirable reactions, but instead of occurring continuously, plating occurs much more frequently under specific, aggressive conditions. These conditions are that the battery has to be charging at a high state of charge, and the degradation process is exacerbated by anything that might increase the resistance or the overpotential at the anode, such as low temperatures or high charge rates [20].

The potential of the anode is usually already quite low, with most graphite anodes sitting slightly higher than 0.2V when empty and reaching just below 0.1V when almost completely full. Charging causes a voltage drop across the battery and also lowers the anode potential, which can relatively easily reach negative values where it is thermodynamically more efficient for intercalating lithium-ions to grab an electron and form pure lithium metal at the anode surface. Intuitively, this can also be thought of as the anode lattice filling up to such a degree that it becomes increasingly difficult for incoming lithium-ions to find and navigate to the remaining vacant spots, because there are less of them and there will be more and more electrostatic repulsion from the already intercalated lithium-ions. As a result, lithium-ions will start piling up near the anode's surface, and it will take more energy to force the lithium-ions through the traffic jam into the particles, which manifests itself as an increased resistance. This increase in resistance means that more energy is required for intercalation reactions to occur, which is more formally expressed as an increase in activation overpotential, and at a certain point the reaction for the formation of lithium metal will require less energy.

This line of reasoning also provides an intuitive explanation for why lower temperatures and high charge rates speed up the plating reaction. At lower temperatures there is significantly less thermal energy present to help lithium-ions shed their transporting solvent shells upon arriving at the anode's surface when having travelled through the electrolyte, drastically slowing down intercalation. Lithium-ions will arrive faster than they can be accepted, and therefore they will start to bunch up at the anode's surface, leading to an increase in resistance. Again, the same line of reasoning can be used for high charge rates: high charge rates will speed up the mass transport of lithium-ions through the electrolyte, but the intercalation rate cannot match this speed as the anode fills up, which also ends up causing a severe build-up of lithium-ions near the anode's surface.

A rather nasty property of lithium plating that SEI layer formation does not exhibit, is that lithium metal will form needle-like dendrites through a self-reinforcing loop. Because electrical charges favour accumulation at sharp points and curved surfaces, lithium-ions will be directed towards existing plating sites, increasing the bunching up effect. As the resulting dendrites start to branch outwards, the electric field lines only get stronger and stronger, further accelerating this process. The main danger with these dendrites is that, once they grow tall enough, they will become capable of physically penetrating the porous separator, leading to a violent, internally induced thermal runaway event [20, 21].

2.2.4. Particle Cracking

While SEI layer formation and lithium plating are both chemically driven processes, particle cracking is a strictly mechanical failure mode. As one could derive from its name, it literally describes electrode particles fracturing and disconnecting from the main conductive structure.

Crystal lattices expand when the structure houses lithium-ions, and they contract again once the lithium-ions deintercalate. Deeply cycling the battery will cause the electrode structures to repeatedly expand and contract asymmetrically, which causes severe mechanical stresses that, over time, can lead to

fractures being created [22]. This is similar to how macroscopic objects can obtain structural defects as a result of repeated deformation—think about how the clip of a plastic pen will snap off more easily after it has been bent and straightened over and over again.

In the case of electrodes, intercalation occurs sooner and more frequently in particles near the electrodes' exterior surfaces, which means that the surface will experience compressive stress, while the un lithiated particles near the inner core experience tensile stress. This structural mismatch stress that follows from internal concentration gradients is what causes cracks to begin forming and propagate along weaknesses from structural defects and imperfections in the crystal lattice [23]. Particle cracking is directly linked to the degradation effect LAM, as fragments that get disconnected from the main structure typically become unsuitable for intercalation, as they have no way to obtain electrons from the current collectors.

2.2.5. Degradation Feedback Loops

Much more significant than the effects that particle cracking has on its own is the role that it has in the feedback loop with SEI layer formation and lithium plating. The true severity of battery aging lies in the fact that these distinct degradation mechanisms do not operate in isolation; they are highly coupled and form self-accelerating positive feedback loops.

The most critical coupling here can be found between particle cracking and SEI growth, as particle cracking exposes fresh, unpassivated new surfaces on which additional SEI film can grow [18, 24]. This instantly triggers a new wave of rapid SEI film formation, aggressively consuming more lithium-ions (LLI), which in its place leads to more non-uniform SEI growth, with certain areas getting “clogged” more easily, forcing the current to take alternative paths. This drives an increase in concentration gradients, which will further accelerate the mechanical stress leading to particle cracking.

Lithium plating is also similarly exacerbated by particle cracking, as the newly exposed highly irregular surfaces will create more favourable geometric conditions for dendrites to start forming. It is also exacerbated by increasingly thicker growing SEI layers, as the increased thickness and resistance will lead to more lithium-ions bunching up and the overpotential to increase, making the formation of pure lithium metal more likely [25]. Since plating does not appear uniformly, it also contributes to an increase in concentration gradients, which leads to a further increase in particle cracking. Finally, to make each of the mechanisms interconnected with one another, lithium plating also worsens SEI growth, as lithium metal is highly conductive, and once dendrites grow from freshly cracked sites with thin SEI layers, electrons can much more easily reach the boundary to contribute to further passivation.

Most of these coupled pathways, with the addition of individual and less significant mechanisms, more commonly take place when the battery is actively used under highly dynamic and transient thermal conditions, and therefore this is often also referred to as cyclic aging, as opposed to calendar aging, which primarily involves continuous SEI growth.

2.3. Electrochemical Modelling

2.3.1. Semi-Empirical

One could say that electrical engineering can at times be referred to as the field of justifiable neglect, as approximations are applied in many cases to simplify and speed up various types of calculations. Circuit modelling often involves making idealised assumptions where parasitic behaviour and properties are left out, as many non-ideal components act in ways that usually do not significantly affect the behaviour of the circuit on larger timescales compared to their ideal counterparts, while potentially adding unnecessary bloat. When batteries are not the primary focus of a circuit simulation, they are often simply modelled as constant voltage sources, or maybe as a capacitor if it useful to understand how the declining voltage during usage affects the rest of the system. This approach provides plenty of information in the majority of cases, but falls short once energy balance calculations come into play.

The first step up in fidelity that is required for such calculations leads to the introduction of the first of three categories in the field of battery modelling, namely semi-empirical models. As is implied by its name, it is not a pure black-box model based only on data, but it is partially rooted in physics, making it a so called grey-box model. At this level of detail, it is not of particularly great interest what happens

inside of a battery, so there is no need to think about intercalation and lithium-ion transport. Instead, look at the black-box external behaviour of a battery when it is interacted with in a circuit. What jumps out are the following key-characteristics:

- Batteries are not an infinite energy source; they can become depleted to the point where they cannot deliver anymore energy, and they can also only accept a limited amount incoming energy.
- Discharging a battery will lower its resting voltage, while charging a battery will do the opposite.
- Current can be pulled from and pushed into batteries at a finite rate due to resistive limitations.
- When a battery discharges, it experiences a sudden voltage drop proportional to the current. Charging leads to the opposite effect, causing a sudden increase instead.
- Once current drops from a non-zero value to zero, the battery voltage will slowly settle towards a resting value.

With the knowledge from the previous section about intercalation and lithium-ion transport, it is relatively straightforward to link each of these characteristics to their corresponding physical phenomena, but this is not a necessary step for semi-empirical modelling. The most common approach here is to use a so-called equivalent circuit model (ECM), which uses a variable voltage source, resistors and capacitors to model all of these key-characteristics.




Figure 2.2: Schematic of a two-stage equivalent circuit model, consisting of a SoC-dependent voltage source, a series resistor and two parallel RC networks connected in series. The output of the network is what represents the terminal voltage of the modelled battery.

Starting with the variable voltage source, this component can be used to model the voltage versus SoC relation, capturing the energy capacity, and power fade due to decreasing voltage at deeper charge states. It is common for a look-up-table or curve-fitted equation to be used here [26], which requires experimental data from either a full depth charge or discharge of the to be modelled cell. Basic ECMs might use only a single curve, but when doing SoC estimation for lithium-iron-phosphate (LFP) batteries, for example, using a single curve will lead to massive errors due to a combination of the flat OCV curve of LFP cells and the property of batteries that the resting voltage is also a function of its history, a phenomenon known as hysteresis. This comes as a result of many combined effects rooted in thermodynamic changes and ion transport that lead to the internal state never being exactly the same, even though the state of charge might be exactly equal. Direction-dependent look-up-tables can help to combat this, or alternatively a correction term for the hysteresis voltage state could be tracked along the way [26]. The SoC itself can be kept track of by Coulomb counting, which involves integrating the current to track the capacity.

After the voltage source, either a single or multiple resistor-capacitor (RC) pairs can be found. These

typically have different sizes, as they are used to capture the voltage drops and recovery over different timescales, and these pairs can even be expanded by adding a differently sized parallel resistor and diodes to account for hysteresis. When changing the current in the circuit, the voltage over the resistor cannot instantly change due to the presence of the parallel capacitor in a pair, which effectively models the voltage lag. Since these capacitors hold charge, they will also nicely mimic how a battery's voltage slowly settles at its resting voltage once the current changes from a non-zero value to zero. In the case of using a two-stage RC network, the first pair can be used for the small time-constant behaviour related to the voltage drops and lag, while the second pair can be sized with significantly larger values to model the much slower resting behaviour. Similarly to how a Taylor-series becomes an increasingly better approximation as more terms are added, an ECM can become more accurate through the addition of more RC pairs that each represents a different sliver of the time response. This is often avoided, however, due to the increase in computational cost and diminishing returns in accuracy.

While ECMs efficiently handle dynamic voltage responses, it is crucial to distinguish this from how these frameworks predict degradation. To perform degradation modelling, semi-empirical models typically separate battery aging into two distinct empirical functions, namely calendar aging and cyclic aging. Equations related to the former are curve-fitted to isolate the effects of time, temperature, SoC. Cyclic aging, on the other hand, is modelled by tracking the active usage of the cell, using stress factors such as capacity-throughput, current, and depth of discharge (DoD). By summing these curve-fitted equations, total capacity fade over time can be calculated without simulating internal chemistry.

The glaring problem with this approach, however, is that empirical degradation models are fundamentally extrapolative. Because they are parameterised using data from highly controlled, static climate-chamber tests, they struggle to accurately predict degradation under novel or highly dynamic conditions [27], where batteries might be subjected to erratic micro-cycles and severe localised thermal transients. These conditions could cause a battery to cross a physical threshold, such as the sudden onset of lithium plating at low temperatures, causing degradation to accelerate exponentially. Since semi-empirical models treat the battery itself as a black box, they are entirely blind to such phenomena; they can only show how a battery degraded in past experiments, offering no insight into why it degrades under new ones. Therefore, to actively optimise battery usage and mitigate degradation in a wider range of scenarios, other models might be more useful.

2.3.2. Physics-Based

Whereas semi-empirical models are grey-box models, physics-based models are true white-box models, being modelled from the ground up with roots in the fundamental physical principles and phenomena, giving significantly more accurate insights into what chains of action might lead to both desirable and undesirable outcomes. Superficially, physics-based modelling of batteries is actually quite elegant and relatively straightforward, as it leans on what can be described as three fundamental pillars:

- Conservation of mass
- Conservation of charge
- Electrochemical kinetics

These fundamental pillars are what make physics-based models so useful for modelling and understanding the root causes of degradation. Given that the exact physical pathways to degradation mechanisms are well enough understood, they can be mathematically simulated with great similarity. As opposed to the blind curve-fitting involved in the process of semi-empirical modelling, adapting the implementation of degradation mechanisms in physics-based models has the benefit of vastly more information being available to deduce what might be causing incorrect behaviour.

In the realm of physics-based modelling, the most popular approach is to assess the battery in a so-called pseudo-2D (P2D) reference frame, established by Doyle, Fuller and Newman in 1993 [28]—the terms DFN and P2D are often used interchangeably to refer to the same type of model. While it is perfectly possible to calculate the conservation of mass and conservation of charge in 3D, it is extremely computationally intensive and requires fine meshing to accurately solve the related partial differential equations (PDEs) everywhere locally. This is something that is usually only of interest to those who design the literal batteries themselves, as these simulations typically run significantly slower than real-time on consumer hardware.

The simplification that P2D makes is to assess the macro-scale interactions in 1D, and micro-scale interactions related to the individual particles in a separate radial reference frame, hence the name “pseudo 2D”. The micro-scale interactions related to lithium-ion diffusion stay of similar fidelity, but compared to the 3D simulation the main sacrifices are made on the macro-scale, as height and width are neglected. There are two core assumptions that can be made to show why this simplification does not lead to a significant loss in accuracy. Starting with the extreme geometric aspect ratio of battery layers, the through-plane thickness usually sits in the order of magnitude of hundreds of micrometres, while the in-plane height and width of a layer can be several centimetres up to several tens of centimetres, which constitutes a size difference ratio of roughly 1:1000 up to 1:10,000. This means that ions travel significantly faster across the thickness, making it the main operational bottleneck. Additionally, the in-plane solid layers, excluding the separator, exhibit both high electrical and thermal conductivity, meaning that under many conditions, in-plane heterogeneities almost immediately become uniform. The only exceptions to these assumptions are transient conditions at high currents, as those can quickly overwhelm the in-plane conductivity to cause in-plane gradients [29].

Even though P2D runs at least a single and often multiple orders of magnitude faster than 3D electrochemical models, for some applications, especially ones involving multiple cells, it is still too slow. The mathematical technique of model order reduction (MOR) that was used to obtain the simplified P2D model from the 3D model can be applied once more to obtain the single particle model with electrolyte dynamics, or SPMe. Whereas P2D simplified the macro-scale reference frame by moving from 3D to 1D, SPMe further simplifies the macro-scale by assuming that reactions occur uniformly across all particles in the electrodes. This assumption allows SPMe to model all particles using a single averaged particle, hence the name. When the single particle model was first introduced, early iterations often disregarded the electrolyte dynamics, but keeping the electrolyte dynamics does not significantly affect the computational cost while significantly improving the error margin at higher currents [29]. Similarly to how P2D becomes less accurate at extremely high currents, SPMe also becomes less accurate at higher currents, but significant deviations already start at a lower threshold due to the uniform reaction rate assumption.

When discussing the drawbacks of semi-empirical models in the previous section, it was mentioned that increasing the fidelity of those models could be achieved by adding more coefficients representing additional contributing phenomena. This brought with it the impracticality of requiring large amounts of experimental data to keep fitting the new additional coefficients. While physics-based models do not struggle with this in the exact same way, since experimental data might only be necessary for validation, the acquisition of data is still a necessity, and often poses realistic hurdles in obtaining truly accurate models. With these models being deeply rooted in the underlying principles of how batteries work physically, many fundamental incalculable parameters are necessary to perform the physics-based calculations. Many of these parameters are highly dependent on the exact formulation of the material chemistries, and as such they can vary widely across different cell types and models. Determining the parameters accurately requires extensive tear-down experiments—these involve battery cells being carefully deconstructed in gloveboxes, after which measurements can be done on individual components—which up until recently were done rather infrequently, and in many cases the obtained data would not be publicly available. One could be using the most advanced high fidelity model available, but without the right parameters, the output will be highly inaccurate—or as some researchers might say: garbage in is garbage out [30].

2.3.3. Data-Driven

The third and last category in the field of electrochemical modelling, data-driven modelling, involves feeding large amounts of data to machine learning (ML) algorithms that try to learn the non-linear behaviour and patterns in the data in order to accurately predict new situations. Critical readers might point out that this category is a sub-set of semi-empirical or purely empirical modelling, as it still involves curve-fitting to obtain a black-box model. The reason for it being an entirely separate category is the recent widespread popularisation and effective utilisation of these types of models in a wide variety of disciplines, industries and scientific fields. Electrochemical modelling is also a field with bright prospects for data-driven models, as the offloading of computational burden from online evaluation to offline training could revolutionise battery management systems (BMSs). Currently, BMSs have to resort to using extremely lightweight models for SoC and SoH predictions, if the latter is even modelled

at all. This often means that they are likely to be inaccurate in more extreme situations and they will most certainly drift over time, which many people might have already dealt with on multiple occasions when their phone shut down prematurely in cold weather, or when they have had to recalibrate their devices with a discharge test to be able to charge the batteries to higher SoCs again. Obtaining higher fidelity data-driven models does not significantly affect the online evaluation time, as it is mostly a matter of improving the training process. Evaluation with data-driven models typically has a similar level of computation complexity to using a look-up table.

Even though data-driven models can perform extraordinarily well when fitting non-linear behaviour compared to simpler semi-empirical and empirical methods, it is not uncommon for them to still be highly inaccurate. This is where the drawback of the black-box characteristic of data-driven models comes in, as the model itself does not provide any meaningful insights into how it actually produced the predicted value. Recently, solutions to this have been proposed in the form of hybrid models, where physics-based principles are added to data-driven models. This has the benefit of potentially increasing the accuracy of the models by giving the model “hints” about what it should be trying to learn, and it can push the models slightly more towards being grey-box instead of being completely black-box, since the model is at least partially rooted in physical reality [31]. A prominent approach is the physics-informed neural network (PINNs), which implements physical equations or relations, such as the OCV curve, in their loss function, reducing the required amount of training data and producing electrochemically feasible predictions [32]. It is also possible to combine physics-based models and data-driven models [31], with an example application being residual learning, where a lightweight model, such as an ECM, might be used to generate the initial predictions. The data-driven model then learns the difference between the output of the ECM and the experimental data in order to improve the ECMs predictions.

In theory, data-driven models could have the potential to outperform both physics-based and semi-empirical modelling in both accuracy and computational complexity [33]. The unfortunate reality is that it is currently either impossible or highly unlikely that such models can be trained, as these models require massive amounts of high-quality data [34]. In order to accurately perform SoC predictions, weeks of real-time data might be required to train the models under a variety of conditions. Degradation modelling is even worse due to the coupled effects and the exponentially greater amount of ways in which different cells can reach similar degradation states. This comes with the additional drawback of degradation testing usually taking literal months up to entire years, which would constitute a single set of conditions for a single combination of cells. With new batteries constantly being produced, and the recent innovations in novel cathode materials and solid state batteries, it is not unreasonable to assume that data-driven models will not widely replace semi-empirical and physics-based models anytime soon.

2.4. Thermal Dynamics and Coupled Effects

2.4.1. Heat Kinetics

The temperature of any object is a function of the heat that it generates internally and the heat that it dissipates to its environment, so modelling a battery’s thermal dynamics usually starts with modelling its heat generation. Typically, heat generation results from energy lost in the system, which in most electrical circuits is modelled using Joule’s first law of heating, or some variation of it. Naïve modelling of batteries could be performed in a similar fashion by simply multiplying the current by the difference of the OCV and the terminal voltage [35]. This, however, is a significant underestimation and also fails to account for the interesting property of batteries being able to actually cool down when charging at low SoC values.

Recall from the previous section about the implementation of electrochemical kinetics that there are more sources of inefficiencies than just Ohmic losses, which Joule heating exclusively models. This loss term, expressed through the ohmic overpotential, accounts for the resistive losses that electrons and lithium-ions experience during transport, but the total overpotential consists of at least two more major terms, namely the concentration overpotential and the activation overpotential. They can respectively be used to model the energy losses incurred from concentration gradients and lithium-ions crossing the interphase boundary.

These three overpotential terms are often grouped together and referred to as the irreversible loss term. In physics-based models, this total irreversible heat is formulated as the product of the applied current

and the total overpotential, effectively representing the energy required to force the reactions to occur. A more simplified macroscopic approximation might calculate this by multiplying the applied current with the difference of the OCV and the terminal voltage.

This irreversible heating is opposed to the so-called reversible, or entropic, loss term. As its name suggests, the reversible loss term constitutes processes that can generate, but also consume heat, due to changes in entropy—as a refresher, entropy represents the state of disorder of a system, with high entropy meaning that a system is highly disordered. The responsible changes in entropy are specifically related to the entropic states of the electrodes during charging and discharging, as their ordering and structure change when lithium-ions intercalate and deintercalate.

This is coincidentally also the process behind why lithium-ion batteries are able to cool down when charging from low SoC values. In general, it is the case that increasing a system's entropy is an endothermic process, meaning that heat is consumed, while the opposite is an exothermic process, meaning that heat is generated. Charging at low SoC values constitutes lithium-ions leaving the cathode lattice and entering the anode lattice, which means that the former will be experiencing an increase in entropy due to the lattice becoming disordered as a result of lithium-ions deintercalating, and the latter experiencing the opposite. While these result in opposing effects, in this specific case, the cathode's endothermic reaction dominates, potentially causing batteries to cool down if the irreversible heating term is low enough, which happens at low currents.

The irreversible heating term usually scales with the square of the current, while the reversible heating term is proportional to the current, meaning that at higher currents, the irreversible heating term starts to dominate, which is one of the reasons for why reversible heating could be neglected in some cases. Similarly to the process of modelling the electrochemical behaviour of a battery, accurately modelling heat kinetics is a more straightforward process in physics-based models compared to semi-empirical ones. As previously mentioned, a naïve approach of modelling heat kinetics could entail exclusively modelling the Ohmic losses. Interestingly, modelling the reversible heating is actually done in almost exactly the same way in both semi-empirical and physics based models. It is typically defined as the product of the current, temperature and the derivative with respect to temperature of the entropic heat coefficient (EHC) curve [36]. This derivative indicates by how much the OCV changes as a function of temperature at every SoC. The relation between the OCV and the temperature is directly linked to the net endothermic or exothermic effect, which is why it can be used to determine the reversible heating term.

Modelling the remaining terms included in irreversible heat generation is relatively straightforward in any physics-based model, as simply multiplying the local current density with the overpotential terms gives the related heat generation at every point in the battery [36]. Semi-empirical modelling will once again struggle here, as accurately modelling additional heat terms will involve blind curve-fitting against experimental data.

2.4.2. Heat Transfer

Once the fundamental principles of internal heat generation have been established, the next step in thermal modelling is understanding how the generated thermal energy actually moves through a system. In physics, heat transfer is governed by three fundamental mechanisms, namely conduction, convection and radiation [37]. Conduction is the transfer of kinetic energy between neighbouring atoms, which, in a battery, dictates how heat moves through solid layers, such as the current collectors, separators and outer casing. Convection, on the other hand, dictates how heat is carried away by moving fluids or gases. In the automotive industry, this principle is the primary mechanism behind cooling systems that use forced air or liquid cooling channels to remove heat from the external surfaces of batteries. Finally, radiation is related to the emission of electromagnetic waves, which can commonly be experienced when feeling the heat from a nearby campfire, for example. Its contribution is generally negligible at the relatively low operating temperatures of batteries, and only really becomes a dominant factor during extreme events like thermal runaway, where cell temperatures exceed several hundred degrees.

Just as electrochemical models range from simple empirical circuits to highly detailed physical simulations, heat transfer modelling offers a similar spectrum of fidelity and computational cost, where each model addresses the fundamental mechanisms of heat transfer in different ways. At the highest

end of the spectrum are the so-called continuum mechanics approaches, such as finite element analysis (FEA) and computational fluid dynamics (CFD), which are true white-box physics-based thermal models that get their name from the continuum hypothesis: a fundamental assumption that ignores the atomic or molecular structure of matter, and instead treats materials and fluids as a continuous, unbroken medium. The underlying physical principles, similar to those of electrochemical modelling, are again physical laws of conservation related to physical quantities such as mass, momentum and energy. Since the corresponding governing equations are often expressed as partial differential equations (PDEs), exact analytical solutions are often impossible to solve directly for complex real-world geometries.

FEA and CFD use numerical methods such as finite elements or finite volumes to break the continuous domains into more manageable discrete pieces. To assess conduction in the solid structures of cells, FEA often discretises these areas in millions of elements and solves Fourier's law of heat conduction—which is basically the thermal analogue to Fick's first law of diffusion, as it indicates that the flow of thermal energy points in the direction of steepest thermal descent, and is affected by thermal conductivity instead of a diffusion coefficient—across every boundary to capture the exact temperature gradients. CFD handles the convective side, solving the notoriously complex Navier-Stokes equations to simulate the flow and heat transfer of viscous fluids, such as the liquid electrolyte or the coolant in channels of a cooling plate.

These methods are indispensable to those who design the exact geometries of cooling systems or specific components inside battery cells, but this level of microscopic fidelity comes with a significant sacrifice of steep computational costs. Because these methods rely on the continuum hypothesis, incredibly tiny time-steps are required to maintain numerical stability in the form of continuity across the boundaries of the mesh. An hour of simulation time could easily take days or even weeks on consumer hardware, making these methods impractical for long-term degradation simulations.

Similar to how model-order-reduction could be applied to 3D electrochemical models to obtain P2D, the complexity of 3D FEA and CFD can be reduced by using volume averaging, which involves homogenising structures and internal components. The resulting homogenised blocks are treated as anisotropic, meaning that they share the exact same temperature everywhere, which implies that related models intrinsically recognise that conduction happens significantly faster in-plane than through-plane. This drastically reduces the geometric complexity and number of elements to be solved, but still relies on numerical mesh-based solvers, making it unsuitable for long simulations.

For applications where rapid calculations are of major importance, and microscopic accuracy is not of particular interest as long as macroscopic accuracy stays similar, lumped parameter thermal networks (LPTNs) are often the model of choice [37]. These discretise the spatial domain into just a handful of distinct macroscopic nodes, entirely abstracting the complex physical geometry needed for finite meshing away. A single battery cell might be simplified into just two or three nodes, maybe even just a single one. When multiple nodes are used, the core node often represents the place of internal heat generation, while the surface node represents that place of heat dissipation. The fundamental mechanisms of heat transfer are modelled by linking these nodes together using resistors and capacitors, respectively representing thermal masses and heat capacities. Mathematically, this makes an LPTN the exact thermal analogue to the ECM used in semi-empirical electrochemical modelling, and they can be analytically solved using the same circuit analysis principles, such as Kirchhoff's current and voltage laws, and Ohm's law. Current becomes heat flux, and voltage becomes temperature.

2.4.3. Pack-Level Dynamics

While an LPTN sacrifices its ability to predict micro-scale spatial temperature gradients within a single cell, it offers computational speeds several orders of magnitude faster than CFD and FEA. When transitioning from single-cell modelling to pack-level simulations, computational efficiency, apart from being a convenience, becomes a vital necessity. When modelling a single cell in isolation, it is common to assume idealised, uniform boundary conditions where the surrounding environment acts as an infinite heat sink. This assumption breaks down, however, when tens, hundreds, or maybe even thousands of cells are tightly integrated together, such as in EV battery packs, where cells physically interact and share thermal pathways, meaning that the boundary conditions of one cell are directly dictated by the behaviour of its neighbours.

These interactions are significantly affected by a pack's thermal management system (TMS). Whether using liquid cooling plates, forced air, or other methods of cooling, a TMS relies on some form of coolant flowing through a designated path to strip heat away from cells. As the coolant absorbs thermal energy from the first cells it encounters, its temperature rises, and consequently, by the time it reaches the cells located further downstream, its cooling capacity will have significantly diminished. This process of heat accumulation creates a spatial thermal gradient across battery packs, where downstream cells will consistently operate at higher temperatures than their upstream counterparts [38]. Similarly, cells located in the middle of a module will typically experience higher temperatures due to being surrounded by other cells, which limits their heat dissipation capabilities, and additionally, neighbouring cells might also inject more heat into each other. Cells near the extremities of a module can more easily dissipate heat to ambient, and have less surrounding cells to deal with, so aside from a TMS's geometric design, modules will already inherently experience spatial thermal gradients as a result of differences in ambient heat dissipation and incoming heat flux.

Macro-scale thermal imbalances are a practically unavoidable phenomenon, but they have a profound effect on the degradation of battery packs as a whole due to the introduction of thermal feedback loops. Many battery degradation mechanisms, such as SEI growth, are governed by Arrhenius kinetics, meaning that their reaction rates are exponentially sensitive to temperature variations, causing hotter cells to degrade faster. Recall from the section about degradation mechanisms that degradation typically leads to resistance increases, for example due to a thickening of the SEI layer and the loss of active material, so more heavily degraded cells will experience increased internal heat generation as a result of their increased resistance, which in turn leads to even higher temperatures again [39]. This positive feedback loop, caused by pack-level dynamics, can drastically shorten the lifespan of affected cells, and therefore also modules or packs as a whole, as their capacity will be limited by the weakest link—BMSs typically shut down battery packs once a single cell's voltage or temperature reaches outside of the allowed operational range.

Thermal divergence and degradation can be further exacerbated by inhomogeneous current distributions resulting from variations in initial parameters [40], such as capacity and internal resistance, which are bound to vary slightly from cell to cell due to manufacturing tolerances. When cells are connected in parallel, these minor electrochemical deviations will dictate the current distribution, as current preferentially flows through the path of least resistance, meaning that slightly healthier or warmer cells will shoulder a greater share of the electrical load. This uneven current distribution leads to varying rates of heat generation and SoC drift, which superimposes an electrochemical imbalance on top of the previously discussed thermal gradients.

A critical reader might suggest that these current distributions could potentially equalise the degradation spread across a module or pack, since healthier cells will experience higher currents, which might be beneficial. While this might be partially true, the reality is that the increased strain on the healthy parts of a battery will lead to an overall acceleration in pack degradation, which does not compensate for any potential homogenisation of the degradation distribution.

2.5. Literature Review

With the fundamental electrochemical and thermal dynamics governing lithium-ion battery degradation having been established, the following section will review recent applications of modelling paradigms related to these domains. The focus is specifically narrowed down to advancements in pack-level electrochemical modelling, degradation modelling, and reduced-order thermal and fluid networks. By evaluating the capabilities and limitations of these methodologies, the unresolved research gaps related to long-term V2G degradation can be established.

Recent advancements in pack-level multi-physics modelling have heavily favoured coupling high-fidelity electrochemical models with 3D continuum approaches to capture spatial variances. For instance, Jordan et al. [41] introduced a novel multi-physics modelling framework to simulate coupled electrochemical-thermal-electrical phenomena in lithium-ion battery packs. It combines a volume-averaged thermal (VAT) model with a tank-in-series battery model using DFN and a 0D electrical network. With this framework, they were able to show that 0.5% cell-to-cell variations could realistically lead to significant voltage variations of over 0.2 V. While they do mention that the framework could be used to assess

degradation, it is reasonable to assume that the model must be run on a high-performance setup to make long-term degradation simulations feasible due to the relatively high-fidelity models used. Building upon this need for heterogeneous assessment, Xia et al. [42] proposed a similar multi-physics framework to simulate capacity fade in realistic heterogeneous conditions. They coupled a P2D model with 3D thermal and fluid continuum approaches, with the addition of also modelling SEI growth. This allowed them to research the impact of coupled effects and inconsistencies in batteries on degradation, showing how state of health (SoH) distributions varied during multi-month experiments for modules containing 15 cells. However, the setup used high-performance hardware, which is not conducive to wide-spread usage on lower end hardware and high volume research, mainly caused by the inclusion of continuum approaches.

Attempting to reduce this computational burden, He et al. [43] used a lumped SPM and 3D thermal and fluid continuum approaches to simulate SoH distributions in a 4s6p electrically connected 24 cell module with a serpentine cooling channel. Heat kinetics were accurately modelled by appending the SPM with additional physics-based heat equations, but ageing was modelled using semi-empirical Arrhenius relations. While assumably this framework should benefit from the lighter SPM model, it suffers greatly from reduced accuracy due to not taking electrolyte dynamics into account, and by semi-empirically modelling degradation. Additionally, the framework is still slowed down significantly by the 3D continuum approaches used to model fluid mechanics and heat transfer. To bypass 3D bottlenecks entirely, Pang et al. [44] created an even more lightweight framework by coupling an extended SPM with a core-shell LPTN, and included degradation modelling through physics-based SEI growth equations. With this model, they were able to show how heat generation evolved non-linearly over time as a result of aging. The model did, however, only take into account SEI growth, and simulated only a single cell, effectively neglecting the impact of coupled pack-level effects on both degradation and heat generation.

Shifting from pack-level scaling to specific V2G applications, researchers have increasingly investigated the accelerated degradation caused by dynamic, bidirectional load profiles. Linders et al. [8] performed a quantitative review of research assessing the impact of V2G on battery degradation, and found that the majority of literature used mainly semi-empirical models. Additionally, the effect of SoC on degradation was often not accounted for, and isothermal assumptions are often made. Across the field of V2G related degradation research, this leads to conflicting and inconsistent results. Addressing this fidelity issue, Lee et al. [45] used a P2D model with heat kinetics and SEI growth to assess V2G related degradation for a single cell across a wide range of charge and discharge patterns. The simulation time was artificially decreased by compressing vehicle operating power for an entire week into a single day. While they did use a physics-based model with SEI growth included, which yielded higher accuracy results compared to previous semi-empirical research, coupled pack-level effects were neglected, and the impact of lithium plating and particle cracking was not taken into account.

Finally, the lack of coupled active cooling in V2G degradation research presents a significant blind spot, particularly given the predictable nature of V2G profiles, which makes the use of anticipative Thermal Management Systems (TMS) an interesting mostly unexplored option. Recent studies have begun to explore proactive cooling to mitigate thermal transients. Tomanik et al. [46] successfully developed and experimentally validated a predictive dynamic programming (DP) based controller that utilises future route information to proactively precondition a battery pack's temperature prior to high-current fast-charging events. The framework successfully demonstrated that anticipatory cooling significantly reduces TMS energy consumption while avoiding dangerous thermal peaks. However, to maintain real-time computational viability, their framework relied on a simple lumped-parameter model that explicitly assumes uniform temperature distribution across the cell.

2.6. Research Gaps and Objectives

The literature review reveals a distinct methodological gap. Frameworks are either too computationally heavy due to 3D continuum approaches, or too electrochemically simplified through isothermal, single-cell, or semi-empirical assumptions. Furthermore, while the energetic benefits of anticipative thermal management have been demonstrated macroscopically, no quantitative nor qualitative assessments of its ability to mitigate internal physics-based degradation have been made as of yet.

To resolve these discrepancies, there is a clear necessity for a computationally efficient, decoupled

multi-physics framework capable of executing long-term V2G profiles while preserving critical intra-cell, inter-cell, and fluid dynamics. To achieve the primary objectives of this thesis, this research aims to answer the following supporting sub-questions:

RQ1: How can a computationally efficient, multi-physics framework be formulated and validated to accurately capture coupled thermal-electrochemical degradation at the pack level?

Since most electrochemical and thermal-fluid models are already well-established individually in literature, the primary objective is successfully coupling the disparate physical domains without incurring severe computational penalties. Validation is performed by decoupling the modules and validating them against isolated experimental and numerical references. Specific sub-questions include:

- How accurately can the coupled electrochemical-thermal cell model predict voltage and temperature characteristics when validated against experimental cell data?
- How accurately can the fluid and thermal models replicate pressure drops and heat generation gradients of a battery module when benchmarked against high-fidelity computational fluid dynamics data?
- Does the fully coupled multi-physics framework execute at a computational speed viable for long-term degradation simulations?

RQ2: To what extent do pack-level scaling and thermal feedback affect degradation predictions compared to standard isothermal and single-cell assumptions?

Current V2G degradation studies often assume isolated, uniformly heated cells, completely neglecting thermal gradients and the effects of active cooling structures. By systematically increasing the scale and physical complexity of the simulations, this research evaluates the qualitative impact of inter-cell thermal interactions. Furthermore, to provide a controlled baseline, the framework's predictions are directly compared against standard empirical models. Specific sub-questions include:

- How do long-term degradation predictions differ across multiple scales of implementation: single-cell isothermal, single-cell thermally coupled, pack-level isothermal, and full pack-level with active cooling?
- How do variations in initial parameters and subsequent uneven current distributions exacerbate thermal feedback and accelerate pack-level cyclic aging?

RQ3: How effectively can an anticipative TMS, leveraging the inherent predictability of V2G load profile, mitigate cyclic battery degradation?

V2G grid services rely on scheduled, highly predictable load profiles, meaning that thermal spikes can theoretically be anticipated. While strict controller optimisation is beyond the scope of this research, a proof-of-concept anticipative cooling strategy is evaluated to demonstrate the potential of proactive thermal management—a concept that is potentially similarly applicable to navigation-assisted driving profiles. Specific sub-questions include:

- To what extent can a simple anticipative lookahead TMS reduce the maximum temperature difference across a module compared to standard reactive cooling?
- What is the corresponding long-term reduction in capacity fade and resistance increase achieved by proactively dampening thermal transients?

3

Methodology

3.1. Framework Architecture

As established in the literature review in Section 2.5, simulating battery degradation under dynamic multi-cycle experiments presents the significant challenge of requiring the assessment of multiple coupled physical domains. Accurately investigating how load profiles accelerate ageing necessitates that the framework must be capable of solving the temperature gradients that drive Arrhenius-dependent degradation, while simultaneously maintaining a computational execution speed viable for long multi-cycle simulations.

Addressing this specific research gap (RQ1) requires constructing a framework that strategically navigates the inherent trade-offs between physical fidelity and computational cost. This thesis proposes a highly coupled, multi-physics framework constructed entirely from reduced-order models, divided into four distinct domains: electrical, electrochemical, thermal, and fluid. The selection of the underlying models for each domain was driven by a careful evaluation of the modelling paradigms discussed in Chapter 2.

3.1.1. Model Selection

The foundational core of this framework requires an electrochemical model capable of tracking internal physical states, specifically local lithium concentrations and overpotentials, to accurately trigger and track degradation side-reactions. Data-driven models were immediately rejected due to their black-box nature and inability to provide root-cause analysis (Section 2.3). Semi-empirical ECMs offer exceptional computational speed, but because they rely on static curve-fitting, they are fundamentally blind to the physical triggers of degradation, such as localised potential drops at the anode boundary.

To achieve the necessary white-box fidelity, a physics-based approach is required. However, the gold-standard P2D model relies on solving heavily coupled partial differential equations across finite macroscopic and microscopic dimensions, making long-term simulation computationally prohibitive. Therefore, this framework utilises a SPMe augmented with side reactions, formulated by Planella et al. (2023) [47]. By mathematically homogenising the electrode particles into a single representative sphere, the SPMe achieves execution speeds approaching those of an ECM while preserving the highly informative internal physical states of a P2D model. While the SPMe relies on the assumption of uniform reaction rates, making it prone to accuracy deviations at higher C-rates, typically those that exceed 2C, this trade-off is highly justifiable for this specific research scope, as V2G grid services generally operate at lower, sustained C-rates.

Because degradation is sensitive to thermal accumulation across a module (Section 2.4.3), the electrochemical model must be coupled to a pack-level thermal and fluid environment. The highest fidelity approach would involve coupling the SPMe to a 3D CFD. However, solving the Navier-Stokes equations for thousands of volumetric fluid nodes alongside the electrochemistry of 28 separate cells would reduce the simulation speed to significantly slower than real-time.

To bypass this bottleneck, the thermal and fluid domains were discretised using reduced-order algebraic networks, leading to the choice of implementing a 0D LPTN for the representation of the thermal domain. A standard isothermal LPTN would contradict the physical reality of a cylindrical cell's poor radial conductivity [48], therefore, a core-shell topology was selected. By mathematically separating the cell into a discrete internal core node and an external shell node, the framework is able to accurately preserve some aspects of the internal radial temperature gradients, without the computational overhead of a 3D mesh. While this approach sacrifices the ability to detect highly localised and asymmetric hot-spots within the jelly roll, it provides a highly accurate macroscopic thermal gradient. Finally, convective cooling is handled by a 1D algebraic fluid network, which computes downstream heat extraction based on empirical flow correlations, entirely bypassing the need for finite element fluid solvers, while still allowing for cooling strategies to be researched.

3.1.2. Delineation of Contributions

The base implementation of the foundational isothermal SPMe [47] in the Julia programming language was provided as a starting point for this research. The primary contribution of this thesis involves taking this isolated model and architecting it into the fully coupled, thermally active pack-level framework required to answer the research questions. This was achieved through two major extensions:

- **Heat kinetics augmentation:** the base SPMe was augmented by systematically integrating the theoretical heat kinetic equations derived from an earlier study by Planella et al. (2021) [36]. This required extending the SPMe code to actively calculate both irreversible and reversible heating, accommodating for spatial variations by integrating the current density across the internal domains.
- **Pack-level domain coupling:** the augmented single-cell models were dynamically wired into a module-level electrical circuit to allow for heterogeneous current distribution, discretised into the intra-cell core-shell and inter-cell LPTN, and coupled to the 1D fluid network.

3.1.3. Data Flow and Domain Wiring

To ensure that the disparate reduced-order models behave as a cohesive physical system, they must continuously exchange active variables during execution. Their connections and corresponding data flows can be summarised as follows:




Figure 3.1: High-level block diagram of the four physical domains, their interconnections and the corresponding summarised data flows. The electrochemical domain does not actually pass cell resistances but a simplified proxy, which is discussed in more detail in Section 3.4.1.

- **Electrical domain (0D nodal analysis):** performing pack-level simulations means that multiple cells will be connected in specific series-parallel configurations that result in different current distributions depending on their equivalent series resistances. Calculating these heterogeneous current distributions across the parallel branches of the battery pack requires solving Kirchoff's laws. This yields the applied current and its rate of change (dI/dt) for each cell, which can be passed to the electrochemical domain. Since the current distribution is dependent on the internal resistances of the cells, the electrical domain receives the terminal voltage, which is simply a state variable of the electrochemical domain. This is used to approximate the instantaneous impedance, instead of using the actual internal resistance, to simplify and speed up the calculations. The specifics of this proxy approach are discussed in more detail in Section 3.4.1.
- **Electrochemical domain (SPMe):** operating at the single-cell level, the augmented SPMe receives the applied current and its trajectory from the electrical domain. Using its reduced-order solid phase and electrolyte equations, it calculates the internal lithium concentrations and the resulting overpotentials. These overpotentials are then used to calculate the total volumetric heat generation and its resulting rate of change (dQ/dt), which are subsequently passed to the

thermal domain. To accurately capture Arrhenius-driven degradation, this domain continuously receives the internal core temperature back from the thermal domain to dynamically update its temperature-dependent kinetic parameters.

- **Thermal domain (LPTN):** the thermal domain receives the volumetric heat generation and its slope (dQ/dt) from the electrochemical domain and directly injects it into the core node of the core-shell LPTN. The heat then conducts outward through a calibrated thermal resistance to the shell node. The core temperature is fed back to the electrochemical domain, while the shell temperature acts as the interaction boundary for the fluid domain. Additionally, the shell nodes of adjacent cells are linked via conductive thermal resistors to simulate inter-cell heat transfer.
- **Fluid domain (algebraic 1D network):** to simulate active pack-level cooling, this domain reads the shell temperatures of the cells and computes the convective heat extraction based on the coolant's flow rate and temperature. As the coolant flows downstream, it absorbs heat, directly influencing the boundary conditions of subsequent downstream cells in the thermal domain.

By decoupling the complex 3D physical reality into these four communicating reduced-order domains, the framework successfully retains the physical feedback loops while remaining computationally agile enough to perform long multi-cycle degradation assessments.

3.2. Electrochemical and Heat Kinetics

With the macroscopic architecture established, the internal physical states of the individual cells must be solved. As discussed in Section 2.3, the fundamental premise of a physics-based framework is that accurate degradation prediction requires explicitly tracking internal variables, specifically lithium concentrations, local potentials, and heat generation, which are impossible to observe using empirical grey-box models.

3.2.1. Electrochemical Principles

Recall from the working principles of batteries (Section 2.1) that the entire process of energy storage and exchange is made possible by the transport of lithium-ions. The foundational core of the SPMe [47] mathematically models this movement, which is driven by two separate physical forces: diffusion and migration.

The calculation of the diffusion force relies on variations of Fick's laws of diffusion. Conceptually, Fick's laws state that particles naturally move from areas of high concentration to areas of low concentration—similar to how ink disperses in water. In the SPMe, these equations calculate the exact concentration gradients of lithium-ions, determining how fast they can move from the core of a spherical particle to its surface, and subsequently through the liquid electrolyte across the macroscopic thickness of the cell.

Ions are also subjected to migration, which is driven by electrostatic forces. Nature favours creating a neutrally charged environment; when electrons accumulate at a current collector, they create an electric field that physically pulls the positively charged lithium-ions toward them. This is modelled using the physical law of conservation of charge. Together, these principles dictate the electrical potentials in both the solid phase and the liquid electrolyte.

Whenever lithium-ions are forced to move or intercalate against these natural gradients, it requires energy. This required energy is expressed through the overpotential. In the SPMe, these resistive barriers are grouped into three primary terms: the activation overpotential (the energy required to trigger the intercalation reaction at the particle boundary), the concentration overpotential (the energy lost overcoming diffusion limits in the electrolyte), and the ohmic overpotential (the pure physical resistance of the solid phase and electrolyte). By accurately calculating these three overpotentials across the internal geometry of the cell, the SPMe provides the exact variables needed to calculate the resulting heat generation.

3.2.2. Heat Kinetics Augmentation

As established in Section 2.4.1, battery heat generation is broadly separated into irreversible heat and reversible entropic heat. To implement these terms dynamically, the base SPMe was extended using the exact theoretical heat kinetic equations derived by Planella et al. (2021) [36]. Because the SPMe

preserves the macroscopic 1D thickness of the cell, the local current density and overpotentials vary as a function of space. Therefore, calculating the total lumped heat to pass to the 0D thermal node requires explicitly integrating these spatially dependent variables across the respective thicknesses of the internal layers.

The total volumetric heat generation is calculated as the sum of four distinct spatial integrals. The formulation of these integrals mirrors the equations used to calculate their corresponding overpotentials [47], with the addition of the spatially dependent current factor inside the integral.

This specific approach splits the irreversible heating into three distinct terms:

1. **Solid phase heating:** the heat generated by electrons struggling to move through the electrodes.
2. **Electrolyte heating:** the heat generated by ions travelling through the liquid electrolyte.
3. **Intercalation (activation) heating:** the heat generated or absorbed by the physical reaction of ions entering and leaving the crystal lattices.

These are mathematically implemented by integrating the local spatial currents ($i(x)$ or $j(x)$) multiplied by their respective potential gradients or activation overpotentials (η) across the thickness of the negative (n) and positive (p) electrodes:

$$Q_s = \sum_{k \in \{n,p\}} \int_0^{L_k} i_{s,k}(x) \cdot \nabla \phi_{s,k}(x) dx \quad (3.1)$$

$$Q_e = \sum_{k \in \{n,p\}} \int_0^{L_k} i_{e,k}(x) \cdot \nabla \phi_{e,k}(x) dx \quad (3.2)$$

$$Q_{\text{int}} = \sum_{k \in \{n,p\}} \int_0^{L_k} j_k(x) \cdot \eta_k(x) dx \quad (3.3)$$

The fourth term calculates the reversible entropic heating. This tracks the entropy changes of the crystal lattice during intercalation (as previously detailed in Section 2.4.1). It is calculated by integrating the spatial current multiplied by the absolute core temperature (T_{core}) and the EHC curve:

$$Q_{\text{rev}} = \sum_{k \in \{n,p\}} \int_0^{L_k} j_k(x) \cdot T_{\text{core}} \cdot \frac{\partial U_k}{\partial T} dx \quad (3.4)$$

The widely used Chen 2020 dataset was selected to provide the baseline electrochemical geometries and OCV parameters for the LG M50 21700 cell. However, this dataset fundamentally lacks the required thermal parameters, such as thermal mass, specific heat capacities, and the non-linear EHC curve. To bridge this gap, the Chen 2020 dataset was appended with the O'Regan 2022 dataset [49].

3.2.3. Degradation Mechanisms

With the internal lithium concentrations and core temperatures successfully implemented, the final stage of constructing the electrochemical domain involves introducing the side-reaction kinetics. Rather than relying on empirical stress-factor curves, this framework implements physical degradation mechanisms that dynamically respond to the transient internal states of the cell.

Following the methodologies outlined by O'Kane et al. (2022) [50] and Planella et al. (2023) [47], the chemically driven mechanisms, solid electrolyte interphase (SEI) growth and lithium plating, are modelled as parasitic side-reactions. In the mathematical formulation, these parasitic reactions introduce competing current densities that compete with the primary intercalation current at the anode boundary. The sum of these fluxes directly dictates the total current density:

$$j_{\text{total}} = j_{\text{intercalation}} + j_{\text{sei}} + j_{\text{plating}} \quad (3.5)$$

As electrons are consumed by j_{sei} and j_{plating} , active lithium-ions are permanently trapped outside of the crystal lattice. This physically reduces the energy capacity of the cell and directly manifests as a loss of lithium inventory (LLI). Furthermore, the growth of the SEI layer adds a continuously thickening physical barrier at the electrode interphase, which is mathematically implemented as a continuously updating film resistance that steadily increases the total internal resistance of the cell.

While the theoretical review in Chapter 2 establishes that cyclic degradation is driven by a highly coupled triad of SEI growth, lithium plating, and particle cracking, the current mathematical implementation of this framework is restricted to the chemically driven mechanisms SEI growth and lithium plating. The mechanical failure mode of particle cracking, which leads to LAM, was ultimately excluded from the final computational architecture due to time constraints. However, because particle cracking actively exposes fresh electrode surfaces that accelerate secondary SEI growth, its integration remains the ideal and necessary next step for future iterations of this framework to fully capture the balance between high-temperature chemical ageing and low-temperature mechanical fatigue.

3.3. Computational Paradigm: Acausal Networks and DAE Stiffness

As previously mentioned at the start of the preceding section, traditional multi-physics frameworks for battery simulations present significant computational bottlenecks. Even though this newly formulated framework should in theory be relatively lightweight due to the usage of reduced-order models, it still has to deal with similar mathematical challenges inherent to solving multi-physics systems.

In this specific case, connecting disparate physical domains generates massive, tightly coupled systems of differential-algebraic equations (DAEs). The system contains differential states, such as the lithium concentration gradients and core temperatures, that change as a function of time through numerical integration. Simultaneously, these states are tightly constrained by algebraic physical laws that must hold true at every instantaneous moment, which usually are applied conservation laws, such as Kirchhoff's circuit laws and heat balance equations.

These DAE systems often exhibit extreme mathematical stiffness, meaning that there is a high degree of disparity between the time constants at which the coupled domains operate. Domains at opposite sides of the spectrum hinder each other in these systems, as the slower systems can be destabilised by transients from faster systems, and faster systems force the slower systems to be run at the same small timesteps, which significantly hinders computational efficiency.

3.3.1. Causal versus Acausal Modelling

Traditional sequential scripting, which is the commonly used approach in standard Python-based environments, generally requires users to manually derive the explicit causal input-output relations of the system. This typically implies the presence of one-way relations, such as that current explicitly dictates the voltage, which makes this approach extraordinarily inflexible. Taking that same example, if the simulation suddenly requires a voltage-controlled constant-voltage charging phase, then the underlying physics equations must be manually rearranged to accommodate for this reference shift.

While this is perfectly functional for simple, isolated single-cell models, it is highly impractical for larger systems such as the one presented in this thesis, especially since the bidirectional coupling often results in algebraic loops that are both difficult to manually define and handle.

These fundamental limitations can fortunately be avoided by using an acausal component-based modelling system, such as the one present in the ModelingToolkit.jl (MTK) ecosystem of Julia. The key characteristic of an acausal paradigm is that, instead of users having to define the computational flow and order of operations, the framework automatically establishes the fundamental conservation laws. Electrical, thermal, and fluid models can be built as independent objects equipped with physical "pins" and "ports", corresponding to electrical terminals or thermal interfaces, for example, and defining the system simply consists of connecting components through these ports.

3.3.2. Symbolic Compilation and Monolithic Failure

The automatic formulation of these conservation laws unfortunately brings with it a unique mathematical hazard in the form of a singular Jacobian matrix. Numerical solvers rely on calculating the Jacobian matrix, which is a matrix of all first-order partial derivatives, to solve the non-linear equations. The

inherent danger lies in the fact that a Jacobian matrix becomes singular, meaning that it either has no solution or infinite solutions, when its determinant reaches zero.

This frequently occurs in the presence of topological loops, such as when ideal voltage sources are connected in parallel with zero interconnect resistance. This leads to an infinite number of possible current pathways that can satisfy the equal nodal voltages that are enforced by Kirchhoff's Voltage Law (KVL). Whenever the determinant reaches zero, it becomes impossible to perform matrix inversion, as that requires division by the determinant, which becomes undefined, preventing solvers from stepping forward in time or losing convergence and crashing due to uncontrolled oscillations.

Typically, MTK is able to bypass many topological issues during its pre-compilation phase, where it acts as a symbolic compiler to construct a bipartite graph of the entire coupled framework, mapping every variable to its corresponding equation and performing symbolic index reduction and tearing to create smaller, more manageable subsets in the system of equations. This can flag cyclic algebraic dependencies, underdetermined systems, and automatically substitute and rearrange equations to eliminate redundant variables. Additionally, further optimisations are applied during compilation by creating a highly optimised system of ordinary differential equations (ODEs) that can be rapidly evaluated during solving.

Despite this powerful symbolic compilation, it was discovered that attempting to couple the stiff non-linear electrochemistry, cyclic degradation kinetics, and bidirectional thermal feedback into a single, heavily constrained monolithic Jacobian matrix caused solvers to immediately hang once any non-zero current was applied.

Initialisation also became a significant bottleneck when coupling these domains, due to the increasing amount of guesses that had to be provided to the system in order to resolve cyclic dependencies, requiring pre-initialisation approaches with hundreds of manual mappings. In other words, finding a mathematically valid starting point where SEI layer thickness, core temperatures, and current distributions all satisfied the algebraic constraints proved to be either impossible for the fully connected DAE architecture, or led to an immediately diverging unstable system.

Following these observations, an architectural pivot towards a hybrid approach involving both analytical and numerical decoupling was implemented, the details of which are discussed in the following section.

3.4. Pack-Level Integration and Multi-Physics Coupling

3.4.1. Electrical Domain: Monolithic Failure and Nodal Partitioning

The example of homogeneous parallel current distributions leading to solver failures (as discussed in Section 3.3.2) was not mentioned coincidentally; it was a direct observation that followed from early attempts to wire multiple SPMe instances in parallel.

Initial iterations of this framework sought to remedy the singular Jacobian within the monolithic MTK architecture by adding deterministic symmetry breaking. This was implemented by adding “virtual” weld resistors with a deterministic offset to each parallel branch. In theory, these resistors should provide the necessary impedance for the solver to analytically solve the parallel current distribution, satisfying Kirchhoff's Voltage Law (KVL) without significantly skewing the macroscopic voltage characteristics of the module. While this topological band-aid was successful for the ideal isothermal case, severe mathematical issues arose once more stiff interactions, such as cyclic degradation kinetics and transient thermal feedback, were introduced into the framework.

Successful circumvention of these initialisation and stiffness issues was ultimately achieved through a structural pivot to analytical partitioning. This entails explicitly instantiating a single, electrically and thermally isolated ODE solver for each SPMe, rather than compiling a single massive 28-cell matrix. To dynamically couple these isolated cells, an external analytical nodal analysis is executed at the boundary of every simulation step.

This specific implementation calculates a “virtual resistance” for each cell. Accurately extracting the true equivalent series resistance via internal impedance calculations at every micro-timestep is far too computationally expensive for long-term simulations. Instead, the virtual resistance is evaluated as a highly dynamic approximation of the instantaneous impedance. It is calculated based on the difference

between the cell's OCV and its terminal voltage, divided by the applied current:

$$R = \frac{U_{\text{OCV}} - V_t}{I_{\text{app}}} \quad (3.6)$$

The resulting resistance value is then used to analytically apply Kirchhoff's current law to divide the total applied pack current across the parallel branches, before passing the newly calculated specific cell currents down to each isolated SPMe.

Because the virtual resistance acts as a mathematical proxy rather than a true physical impedance, it introduces transient errors during rapid, high-frequency current changes. However, this simplification is justified: once the current stabilises during a charge or discharge pulse, the dynamic variations in terminal voltage inherently capture the true physical distribution of the pack.

3.4.2. Thermal Domain: Core-Shell Topology and Inter-Cell LPTN

While solving for the electrical distribution ensures accurate electrochemical operation at the individual cell level, both cyclic and calendar ageing are significantly affected by temperature. As established in Section 2.4.2, evaluating a cell's thermal state requires accurately balancing internal heat generation with external heat transfer. To achieve this, the framework constructs a comprehensive 0D LPTN composed of both an intra-cell and an inter-cell topology, parametrised using the experimental dataset provided by O'Regan et al. (2022) [49].

Many computationally lightweight thermal models make the assumption of an isothermal cell, implying that internal conduction is so rapid that internal thermal gradients are negligible. However, cylindrical cells exhibit highly anisotropic thermal conductivity. Heat generated near the core struggles to travel radially outwards through the alternating layers of thermally resistive plastic separator. Isothermal assumptions artificially suppress the internal jelly roll temperature, thereby fundamentally underestimating Arrhenius-dependent degradation kinetics.

The intra-cell LPTN assigns specific thermal properties to a discrete core node, representing the jelly roll, and a discrete casing node. The thermal capacity (C) of the core is fundamentally dynamic; rather than using a static scalar, the specific heat capacity ($c_{p,\text{core}}$) is calculated as a third-order polynomial. This polynomial continuously adjusts as a function of temperature and is derived from the exact mass fractions of the jelly roll's constituent materials, namely the positive and negative electrodes, separator, metallic foils, and electrolyte.

$$C_{\text{core}}(T) = m_{\text{jelly roll}} \cdot c_{p,\text{core}}(T) \quad (3.7)$$

The core node receives the total volumetric heat generation calculated by the SPMe. The heat must then dissipate to the external shell. Unlike simplified models that use a single lumped internal resistor, this framework explicitly separates the internal conduction into parallel radial and axial pathways. This distinction is critical due to the physical winding of the jelly roll. Radial conduction must repeatedly cross the highly resistive plastic separators, whereas axial conduction travels uninterrupted along the highly conductive copper and aluminium current collectors.

However, as highlighted by O'Regan et al. (2022) [49], axial heat transfer in cylindrical cells is severely bottlenecked by the presence of dense plastic insulator caps situated at the top and bottom of the cell casing. Because proprietary dimensions and specific material compositions for these insulators are rarely disclosed by manufacturers, a back-of-the-envelope calculation was performed to mathematically account for this bottleneck. The volume-averaged internal axial resistance of the jelly roll was appended with an estimated boundary resistance for the insulator caps.

By explicitly separating these anisotropic nodes and accounting for physical bottlenecks, the framework accurately preserves some aspects of the internal temperature gradients without requiring a dense volumetric mesh.




Figure 3.2: Schematic representation of the intra-cell LPTN connections, consisting of a current source that models the heat generation in parallel with the cell's heat capacity, connected in series to two parallel thermal masses that represent the axial and radial pathways, and ending at the shell node that connects to the shell's thermal capacitance and to other components in the inter-cell LPTN.

In this framework, the shell node of each cell is connected to its adjacent neighbouring cells, the ambient environment, and the liquid cooling channel structure via discrete external thermal resistors. These thermal resistance values can be fitted against reference CFD data, or they can be analytically derived for known geometries and materials using the fundamental equation for thermal resistance:

$$R_{th} = \frac{L}{k \cdot A} \quad (3.8)$$

Here, L is the conduction distance between the specific nodes, k is the thermal conductivity of the interstitial material, which could be the thermal interface material or an air gap, and A is the effective cross-sectional area for heat transfer.




Figure 3.3: Schematic representation of the inter-cell module-level LPTN connections, consisting of a radial connection to adjacent cells, an axial connection to ambient, and a contact resistance with the coolant channel structure.

3.4.3. Fluid Domain: 1D Curved-Pipe Modelling

The mechanisms of fluid dynamics can be understood relatively easily through analogies with circuit theory, since the fundamental conservation laws share many similarities. In this 1D algebraic fluid network, the coolant mass flow rate ($\dot{m}$) acts as the electrical current, the pressure drop (ΔP) represents the potential difference or voltage drop, and the physical constraints of the pipe create a non-linear hydraulic resistance. To accurately quantify this hydraulic resistance, fluid mechanics relies on dimensionless numbers to characterise the state of the flow. The foundational metric is the Reynolds number (Re), which represents the ratio of inertial forces to viscous forces. It dictates whether the fluid flows in orderly, parallel streamlines (laminar) or in a chaotic, mixing state (turbulent). It is calculated as follows:

$$Re = \frac{\rho \cdot v \cdot D_h}{\mu} \quad (3.9)$$

Here, ρ is the fluid density, v is the velocity, μ is the dynamic viscosity, and D_h is the hydraulic diameter. Because battery cooling channels are rarely perfectly circular, the hydraulic diameter allows rectangular ducts to be mathematically treated as equivalent circular pipes, defined as $D_h = 4A_{\text{flow}}/P_{\text{wet}}$, where A_{flow} is the cross-sectional area and P_{wet} is the wetted perimeter.

Standard straight-pipe equations, such as the Darcy-Weisbach equation equipped with a baseline laminar friction factor of $f = 64/Re$, assume that fluid streamlines remain perfectly attached to the pipe walls. However, in a tight serpentine coolant channel structure wrapping around cylindrical battery cells, this is never the physical reality. The fluid is constantly forced through sharp 180° bends, which induces centrifugal secondary flows, also known as Dean vortices, and causes flow separation at the inner corners. To capture this complex 3D behaviour in a 1D model without resorting to arbitrary curve-fitting, this framework implements a hybrid approach leveraging continuous helical coil correlations.

The magnitude of these secondary centrifugal flows is quantified using the Dean number (De), referenced from Kakaç et al. (Chapter 5, Section 5.2, Page 5-4, Eq. 5.1) [51], and implemented in the framework as:

$$De = Re \sqrt{\frac{a_{\text{hyd}}}{R_c}} \quad (3.10)$$

Here, a_{hyd} is the channel's hydraulic radius ($D_h/2$), and R_c is the radius of curvature. Because strong centrifugal forces physically stabilise fluid streamlines, they suppress turbulent fluctuations, which means that the transition from laminar to turbulent flow is severely delayed. The framework dynamically calculates this curvature-delayed critical Reynolds number (Re_{crit}), proving the battery coolant remains largely laminar at standard operating speeds, according to Kakaç et al. (Chapter 5, Section 5.3, Page 5-9, Eq. 5.21 / 5.22) [51]:

$$Re_{\text{crit}} = 2100 \left[1 + 12 \left(\frac{R_c}{a_{\text{hyd}}} \right)^{-0.5} \right] \quad (3.11)$$

Depending on whether the dynamic flow remains in this extended laminar regime or breaches into turbulence, the modification of the viscous wall shear is predicted using helical correlations. For the extended laminar regime, the framework applies Srinivasan's curved friction ratio, referencing Kakaç et al. (Chapter 5, Section 5.3, Page 5-7, Eq. 5.16) [51]. If extreme high-speed pump conditions push the flow past Re_{crit} , Ito's turbulent curved friction factor is used to ensure mathematical stability and physical accuracy, referencing Kakaç et al. (Chapter 5, Section 5.4, Page 5-22, Eq. 5.48) [51]:

$$\text{Laminar (Srinivasan): } \frac{f_c}{f_s} = 0.1125 De^{0.5} \quad (3.12)$$

$$\text{Turbulent (Ito): } f_c \left(\frac{R_c}{a_{\text{hyd}}} \right)^{0.5} = 0.00725 + 0.076 \left[Re \left(\frac{R_c}{a_{\text{hyd}}} \right)^{-2} \right]^{-0.25} \quad (3.13)$$

While these equations originally derive from continuous helical coils, their application to serpentine structures is highly justified: Dean vortices develop almost immediately upon a fluid entering a curve. However, a serpentine channel is not a perfect helix, as it features sharp, discrete turns that cause the fluid to temporarily detach and tumble. To compensate for the form drag that the helical equations naturally miss, an empirical minor loss coefficient (K) is superimposed into the fundamental Darcy-Weisbach pressure drop equation:

$$\Delta P = \left(f_c \frac{L}{D_h} + K \right) \frac{\rho v^2}{2} \quad (3.14)$$

This hybridisation of analytical helical wall shear with empirical form drag ensures the total pressure drop exactly matches the exponential scaling observed in reference 3D CFD data.

Once the local velocity and pressure drop are established, the framework must couple the hydrodynamic behaviour back to the thermal network. To do so, a second fluid-specific property must be evaluated: the Prandtl number (Pr). The Prandtl number defines the ratio of momentum diffusivity (kinematic viscosity) to thermal diffusivity, effectively describing whether a fluid is better at transporting velocity gradients or temperature gradients. It is defined as:

$$Pr = \frac{c_p \cdot \mu}{k} \quad (3.15)$$

Using both the Reynolds and Prandtl numbers, the system calculates the Nusselt number (Nu). Conceptually, the Nusselt number represents the ratio of convective heat transfer across the fluid boundary to the purely conductive heat transfer that would occur if the fluid were entirely stagnant. Serpentine coolant channels typically exhibit cross-sections with extreme aspect ratios due to their thin and tall shape. Under a constant axial heat flux and uniform peripheral temperature boundary condition—the “H1” condition, which the highly conductive coolant channel walls emulate—the Nusselt number asymptotically approaches 8.235 as the aspect ratio approaches zero. Bypassing complex geometric correction factors and adopting this asymptotic limit provides a mathematically sound and computationally lightweight simplification for laminar flow, referenced from Kakaç et al. (Chapter 3, Section 3.4, Page 3-47, Eq. 3.161) [51].

If a severe thermal transient triggers the TMS to push the coolant past the curvature-delayed critical Reynolds threshold (Re_{crit}), the enhanced thermal mixing of the transitional and turbulent regimes must be captured. For this, the framework implements the industry-standard Gnielinski correlation, referenced from Kakaç et al. (Chapter 4, Section 4.2, Page 4-34, Table 4.4) [51]. It is paired with Petukhov's smooth-pipe turbulent friction factor (f_D) [51], pre-converted from the Fanning form to the Darcy form for direct implementation in the continuous LPTN solver:

$$f_D = (0.79 \ln(Re) - 1.64)^{-2} \quad (3.16)$$

$$Nu_{turb} = \frac{(f_D/8)(Re - 1000)Pr}{1 + 12.7(f_D/8)^{0.5}(Pr^{2/3} - 1)} \quad (3.17)$$

With the dimensionless Nusselt number established, it is converted back into the physical heat transfer coefficient (h), representing the thermal conductivity of the convective boundary layer in $W/(m^2K)$:

$$h = \frac{Nu \cdot k}{D_h} \quad (3.18)$$

This convective coefficient acts as the exact bridge between the fluid domain and the cell's solid shell node, dictating the convective heat flux (Q_{conv}) extracted from the metal pipe walls:

$$Q_{conv} = h \cdot A_{surface} \cdot (T_{solid} - T_{fluid}) \quad (3.19)$$

Finally, to tie the multi-physics network together, the resulting temperature of the coolant is governed by a 1D advection ordinary differential equation (ODE). This equation perfectly mirrors a capacitor being charged by multiple current sources: the fluid’s thermal mass ($m_{\text{fluid}} \cdot c_p$) is charged by the convective heat flux from the cell walls (Q_{conv}) and the advective heat flux carried downstream from the upstream fluid ($\dot{m} \cdot c_p \cdot \Delta T$). This establishes the nodal energy balance:

$$m_{\text{fluid}} \cdot c_p \frac{dT}{dt} = Q_{\text{conv}} + \dot{m} \cdot c_p \cdot (T_{\text{inlet}} - T_{\text{node}}) \quad (3.20)$$

Applying an arbitrary “friction multiplier” to standard straight-pipe equations forces 1D thermal models to match specific experimental or CFD data, but this approach lacks predictive power. By explicitly deriving the dimensionless flow parameters to separate linear viscous shear from quadratic form drag, and by dynamically calculating the resulting heat fluxes in the advection ODE, the formulated fluid domain achieves dimensional awareness. It predicts how centrifugal forces stabilise the flow, how transition to turbulence is delayed, and how thermal energy accumulates downstream. This ensures that the LPTN can simulate highly dynamic drive cycles, varying temperatures, and geometric iterations without requiring constant recalibration.

3.5. Partitioned Multi-Rate Time-Stepping

The previously discussed approach of applying analytical partitioning to the electrical connections of the SPMe successfully circumvented the singular Jacobian and topological stiffness. However, significant room for improvement regarding the computational efficiency of the framework remained due to the fundamental disparity in timescales between physical domains.

Thermal systems typically operate on a timescale several orders of magnitude greater than electrochemical systems due to the presence of large thermal masses and relatively slow convective fluid dynamics. If forced to assess these domains monolithically, standard implicit solvers are mathematically constrained to restrict their universal step size to match the fastest dynamic present in the system. This leads to the solver taking millions of microscopic time-steps to safely solve the stiff electrical and electrochemical transients, while the entire 0D and 1D thermal-fluid matrix is redundantly evaluated millions of times as well, without any meaningful physical changes occurring across those minute intervals. Logically, this unnecessarily inflates the computational cost, rendering long-term degradation research entirely infeasible on standard consumer hardware.

3.5.1. Analytical and Numerical Partitioning

To resolve this timescale disparity, the framework avoids evaluating the electrochemical and thermal-fluid domains as a single coupled system. Instead, it implements a hybrid approach utilising both analytical and numerical partitioning.

Analytical partitioning refers to explicitly severing the mathematical dependencies at the system level. The electrochemical and thermal domains are isolated and instantiated as distinctly separate Ordinary Differential Equation (ODE) problems within Julia’s DifferentialEquations.jl ecosystem. This practically results in 28 electrically and thermally isolated SPMe models being instantiated, alongside a singular thermal-fluid LPTN. Because the fluid domain consists exclusively of analytically solvable algebraic equations and correlations (as detailed in Section 3.4.3), it does not require its own differential solver and is instead solved analytically alongside the thermal ODEs.

Numerical partitioning follows as a direct result of this analytical severance. By isolating the ODE problems, the domains can be assigned highly specialised numerical solvers and discrete integrator objects, allowing them to traverse time at their respective, appropriate timescale-matching speeds. Both the electrochemical and thermal domains utilise QNDF (Quasi-Constant Step-Size Backward Differentiation Formula), which is exceptionally robust for stiff systems. Their error tolerances are aggressively tuned to their corresponding physical realities: the SPMe integrators operate with strict tolerances to precisely capture high-frequency internal lithium concentration gradients, while the slowly evolving thermal-fluid integrator operates with significantly looser tolerances.

3.5.2. First-Order Hold (FOH) Co-Simulation

What remains is handling the data exchange between these numerically partitioned domains. The traditional zero-order hold (ZOH) decoupling successfully bypasses timescale stiffness constraints but inherently introduces a mathematical lag error. Because an electrochemical integrator operating under ZOH uses a constant, “held” value over the thermal time-window, it continuously generates heat based on outdated boundary conditions. To eliminate this lag error without resorting to monolithic solver compilation, the framework implements a true multi-rate First-Order Hold (FOH) co-simulation with a dual-handshake architecture.

Instead of holding states constant, the FOH methodology calculates the first-order derivatives of the interacting variables—specifically the rate of change of the applied current (dI/dt) and the volumetric heat generation (dQ/dt)—and projects these linear slopes across the integration steps. At the boundary of every thermal macro-step, the master orchestrator executes a nested sequential algorithm:

1. **Thermal boundary targeting:** the orchestrator determines the next valid thermal macro-step (t_{target}) based on the current maximum temperature gradient across the pack.
2. **Electrical micro-stepping and projection:** before the thermal solver can advance, the electrical domain must catch up. The orchestrator divides the thermal macro-step into smaller, dynamic electrical micro-steps. For each micro-step, it projects the future total pack current, calculates the corresponding branch currents, and derives the rate of change (dI/dt).
3. **Electrochemical integration:** the isolated SPMe integrators are supplied with their specific branch currents, the dI/dt slopes, and the current core temperatures. They integrate forward to the next micro-step boundary, dynamically adapting to the changing current trajectory rather than a static block.
4. **Thermal integration:** once all electrochemical integrators successfully reach the thermal macro-step boundary (t_{target}), the orchestrator calculates the resulting rate of change in volumetric heat generation (dQ/dt) for every cell. These slopes are passed to the thermal-fluid integrator, which then steps forward to update the macroscopic temperatures and fluid dynamics.

3.5.3. Dynamic Time-Stepping and Threshold Optimisation

By using the FOH methodology, the integrators continuously anticipate the trajectory of the interacting domains. This linearisation of the inter-step gradients practically eliminates the severe lag errors that restrict ZOH methods, granting the implicit QNDF solvers the stability to take significantly larger time-steps without compromising physical validity.

To fully exploit this stability and maximise computational efficiency, the time-steps for both the electrical and thermal integrators are controlled dynamically using gradient-based heuristics. For example, the thermal macro-step ($\Delta t_{\text{thermal}}$) is calculated dynamically based on the maximum rate of temperature change across the pack:

$$\Delta t_{\text{proposed}} = \frac{\alpha}{\left| \frac{dT_{\text{max}}}{dt} \right| + \epsilon} \quad (3.21)$$

Here, α is a tuning parameter dictating the allowable state change per step, strictly clamped to an upper bound (dt_{max}), and ϵ (1×10^{-4}) prevents division by zero.

To find the optimal hyper-parameters (α and dt_{max}) across different V2G and driving regimes, a threshold-based grid search was implemented. Because the FOH architecture maintains a highly linear error response, previous optimisation strategies, such as identifying maximum curvature on a Pareto front, were rendered obsolete. Instead, the framework executes a high-speed parameter sweep against a nearly continuous “ground truth” reference simulation ($\Delta t = 0.1$ s).

The orchestrator evaluates the computational speed-up of each configuration while strictly mapping them against predefined error thresholds. For the electrochemical domain, these constraints include a target Root Mean Square Error (RMSE) and a minimum coefficient of determination ($R^2 \geq 0.95$). For the thermal domain, configurations are bounded by a maximum allowable peak temperature deviation and a strict thermal dose error percentage of less than 1%. The algorithm programmatically selects the

parameter combination that yields the highest execution speed-up while definitively satisfying these hard accuracy constraints, ensuring rigorous physical fidelity at a fraction of the computational cost.

3.6. Anticipative TMS Control Logic

3.6.1. Thermal Inertia and Control Limitations

Controlling the temperature of a high-density battery pack presents a fundamental physical challenge due to thermal inertia: by the time an external surface sensor detects a rise in the cell casing temperature, the internal jelly roll has already been operating at elevated temperatures for several minutes due to the poor radial thermal conductivity of the cell.

Traditionally, TMS control logic relies on proportional-integral-derivative (PID) controllers. However, a PID controller is fundamentally reactive; it calculates its control output based solely on the instantaneous error between the current temperature and a target set-point. While the derivative term can anticipate the slope of a temperature rise, a PID controller is entirely blind to the future.

A frequently proposed approach for similar problems involves using model predictive control (MPC). MPC is theoretically ideal for accurate predictions because it utilises a receding horizon to “look” into the scheduled future load and calculate an optimal cooling trajectory. However, MPC requires embedding a plant model, which in this case should be a mathematical replica of the battery, inside the controller, and it must solve a complex mathematical optimisation problem at every single control step to find the most efficient cooling output. Attempting to wrap an iterative MPC optimisation solver around a stiff 28-cell multi-physics DAE system, which already requires careful numerical partitioning, as detailed in Section 3.5, would result in an astronomical computational burden, rendering long degradation simulations entirely impossible.

3.6.2. Reactive Baseline: Hysteresis Control

Assessing the benefit of anticipative cooling requires a baseline that reflects current industry standards. The reactive control strategy implemented in this framework is modelled upon the thermal logic used in modern commercial EVs, drawing inspiration for operational setpoints from the high-voltage battery management systems detailed by Volkswagen [52].

This reactive strategy operates as a strict state machine governed by upper and lower temperature thresholds, creating a classic hysteresis loop. The logic monitors the maximum core temperature across the 28-cell pack:

- **Passive state:** while the pack remains below the upper threshold, the TMS circulates coolant at a baseline passive velocity at the ambient environmental temperature. This ensures uniform thermal distribution without expending significant compressor energy.
- **Active state:** if the maximum core temperature breaches the upper safety threshold, the TMS instantly switches to active cooling. The coolant velocity is increased and the inlet temperature is actively chilled via a refrigeration cycle to 10°C below ambient.
- **Hysteresis:** to prevent mechanical short-cycling, the system does immediately revert back to its passive state once it dips below the upper safety threshold. Instead, the active state is locked until the maximum core temperature is successfully driven down below the lower threshold, at which point it reverts to the passive state.

3.6.3. Moving-Average Lookahead

To circumvent the limitations of thermal inertia without the prohibitive cost of MPC optimisation, the framework implements a deterministic anticipative TMS strategy. Because V2G load requests are scheduled in advance by a grid aggregator, the vehicle’s internal BMS theoretically should have access to the future load profile array.

The anticipative strategy leverages this foreknowledge by calculating a moving-average window of the scheduled current demand. Rather than relying on a continuous integral, which introduces unnecessary computational overhead and can be unstable during rapid oscillations, the lookahead is formulated as a discrete mathematical summation. The orchestrator script samples the scheduled absolute pack current at N evenly spaced intervals across a predefined lookahead window.

The average future current per parallel cell branch ($I_{\text{future_avg}}$) is calculated as:

$$I_{\text{future_avg}} = \frac{1}{N \cdot \text{cols}_{\text{parallel}}} \sum_{i=1}^N \left| I_{\text{pack}} \left(t + i \cdot \frac{t_{\text{window}}}{N} \right) \right| \quad (3.22)$$

3.6.4. Proactive Control Logic

With the future load mathematically quantified, the anticipative strategy supervises the reactive baseline using a dual-condition logic gate. The state machine evaluates both the instantaneous maximum temperature (T_{max}) and the impending future load ($I_{\text{future_avg}}$) against their respective thresholds.

The control logic dictates that the TMS will trigger the active state if either of the following conditions is met:

1. $T_{\text{max}} > T_{\text{high}}$ (35°C) (reactive fail-safe)
2. $I_{\text{future_avg}} > I_{\text{threshold}}$ (anticipative override)

The hysteresis control logic is modified in two crucial ways: the system will only revert to the passive state if the pack is physically cool ($T_{\text{max}} < 32^{\circ}\text{C}$) and the future scheduled load has subsided ($I_{\text{future_avg}} \leq I_{\text{threshold}}$).

By implementing this deterministic control logic, the framework achieves true physical pre-chilling without iterative optimisation. If a significant V2G grid discharge is scheduled to begin within the time-frame of the averaged-window, the $I_{\text{future_avg}}$ variable instantly spikes above the threshold. The logic gate triggers the active cooling phase immediately, extracting heat from the cell casings before the current is applied.

As a result, when the heavy discharge actually begins, the rapidly generated internal Ohmic heat is absorbed by the pre-cooled thermal capacitance of the jelly roll and casing. This fundamentally dampens the peak core temperature transient, preventing the cell from lingering in the high-temperature regimes that exponentially accelerate SEI growth and lithium inventory depletion.

3.7. Parametrisation and Validation Strategy

A coupled, multi-physics framework relies heavily on the accuracy of its foundational parameters; assuming arbitrary physical constants compromises the validity of the simulations. Furthermore, validating the fully coupled 28-cell module monolithically presents a diagnostic challenge, as discrepancies can originate from miscalculations or incorrect coefficients in any interconnected domain. Systematically eliminating this ambiguity requires a decoupled hierarchical validation strategy. Execution of the validation strategy is divided into two discrete phases.

3.7.1. Electrochemical Validation

The cell chemistry selected for this framework is the commercial LG M50, a relatively modern high-density 21700 format cylindrical cell with a nominal capacity of 5.0 Ah. It was specifically chosen due to the availability of extensive, open-source tear-down datasets. The baseline electrochemical parameters, including geometric dimensions and maximum solid lithium concentrations, are referenced from the dataset compiled by Chen et al. [53]. To ensure accurate modelling of both the heat kinetics and heat transfer, these base parameters are appended with thermal mass, heat capacity and entropic heat coefficient data derived by O'Regan et al. [49]. (Note: The exhaustive tables of all baseline electrochemical and thermal parameters used in this framework are provided in Appendix B).

First, a singular SPMe, dynamically coupled only to its own intra-cell core-shell node, is simulated under standard constant-current load profiles. The simulated voltage and temperature outputs are then benchmarked against experimental data provided in the Chen dataset. This phase ensures that the fundamental lithium transport and heat generation mechanisms are mathematically sound before pack-level integration introduces heterogeneous current distributions.

3.7.2. Thermal-Fluid and Module-Level Validation

At the pack level, the 4s7p module topology and the inter-cell thermal resistance values are mapped directly to a 3D CFD study by Nema et al. [54]. By utilising an identical physical architecture and structure, the framework's thermal and fluid mechanics can be reliably validated against the data from this study.

Second, the thermal and fluid domains are completely severed from the electrochemical domain by omitting all SPMe components and instead feeding the 0D shell nodes static, pre-calculated heat generation profiles, which are obtained from high-fidelity DFN simulations from Nema et al. The 1D fluid model is applied across a sweep of coolant velocities, using which the empirical adjustment coefficients, such as the minor loss form factor, are calibrated. Rather than relying on arbitrary manual tuning, this calibration was governed by an objective function designed to minimise the Root Mean Square Error (RMSE) and maximise the coefficient of determination (R^2) against the 3D CFD reference data.

It is important to acknowledge a valid critique of this approach: relying on empirical calibration against a specific CFD dataset resembles the curve-fitting limitations that are often attributed to semi-empirical methods. Generating new CFD models for varying geometries is a time-consuming process that is outside the scope of this research. Therefore, the current 1D fluid calibration is strictly bound to this specific 28-cell geometry. It is important to note that, since calibration has been performed within a bounded velocity range of 0.1 m/s to 0.5 m/s, extrapolation beyond these bounds should be avoided, as the validity of the 1D algebraic approximation cannot be guaranteed due to friction mechanics being highly dependent on turbulent behaviour. Within these bounds, however, the calibration ensures that the pressure drop properly exhibits exponential scaling, forcing the 1D nodal flow to accurately replicate the convective heat extraction of the 3D CFD reference. Compared to semi-empirical approaches, this framework remains strictly physics-based, allowing for deep visibility into the internal degradation states, and providing a foundation for future geometric scaling.

4

Results and Discussion

This chapter presents the outcomes of the developed coupled multi-physics framework. Sections 4.1 and 4.2 detail the hierarchical validation of the framework, systematically validating the electrochemical and thermal-fluid domains against reference experimental and 3D CFD data. Section 4.3 presents an ablation study, isolating the impact of thermal feedback and electrical coupling on battery degradation under dynamic load profiles.

4.1. Framework Parametrisation and Validation

As was previously mentioned in the methodology, the framework's accuracy was assessed through a two-phase decoupled validation strategy. This approach isolates the electrochemical heat generation of the SPMe from the heat extraction of the coupled 0D LPTN and 1D fluid network, which ensures that inaccuracies and errors can be more confidently attributed to specific characteristics of these separated domains.

4.1.1. Electrochemical Validation (Single-Cell)

The foundational SPMe was validated against experimental climate chamber data provided by Chen et al. [53] for a single LG M50 cell. To validate the electrochemical heat generation and voltage response, the simulation subjected the cell to constant current discharges of 0.1C, 0.5C, 1C, and 1.5C using the exact current and ambient temperature data from the experimental dataset.




Figure 4.1: Electrochemical validation of the SPMe against experimental LG M50 climate chamber data (Chen et al. [53]) under 0.1C, 0.5C, 1C, and 1.5C constant-current discharges. The stacked plots display terminal voltage (top), external shell temperature and ambient temperature (middle), and applied current (bottom).

The voltage overshoots at high C-rates and the lack of a prolonged relaxation curve, that asymptotically rises to a resting voltage, highlights the mass transport limitations inherent to the SPMe's uniform reaction assumption. The simulated terminal voltage tracks the experimental data with high accuracy across the entire active discharge phase in all cases, except for 0.1C, where it slightly detaches by overestimating the voltage. The Root Mean Square Error (RMSE) ranges from 64.19 mV ($R^2 = 0.9808$) for the 0.1C discharge, up to 144.22 mV ($R^2 = 0.9392$) for the 1.5C discharge. The decreasing accuracy at higher C-rates is likely mainly exacerbated by the overshoots in the low regime, which are penalised heavily by the RMSE. At high C-rates (1C and 1.5C), the SPMe struggles to accurately capture the steep voltage cliff below 2.7 V, overestimating the voltage drop and even dipping below 2.0 V during the 1.5C discharge. As mentioned previously, this deviation is likely a direct consequence of mass transport limitations within the porous electrodes, which the SPMe's uniform reaction assumption cannot perfectly account for under aggressive concentration gradients. Furthermore, during the post-discharge rest period, the SPMe voltage immediately flatlines, failing to capture the slow, asymptotic voltage recovery driven by lingering lithium diffusion gradients. Capturing these would require upgrading the SPMe to

DFN, which typically is able to accurately model the relaxation phase. For the purposes of this research, these inaccuracies are permissible, but they should be carefully taken into account when assessing degradation comparisons.

Thermally, the core-shell LPTN tracks the experimental shell temperature remarkably well. The thermal RMSE remains tightly bounded between 0.33°C and 1.57°C across all evaluated discharge rates. While the 0.1C profile yields a low thermal R^2 (0.1641), which is likely a statistical artefact caused by the negligible absolute temperature change at such a low current, where minuscule ambient convection deviations disproportionately penalise the R^2 metric, the higher C-rates (0.5C, 1C, and 1.5C) demonstrate highly accurate thermal scaling with R^2 values consistently exceeding 0.97. A comprehensive summary of all electrochemical and thermal error metrics is provided in Table 4.1.

Table 4.1: Electrochemical and thermal validation error metrics.

Discharge Rate	Voltage RMSE (mV)	Voltage R^2	Temp. RMSE ($^{\circ}\text{C}$)	Temp. R^2
0.1C	64.19	0.9808	0.33	0.1641
0.5C	67.69	0.9850	0.36	0.9703
1.0C	101.85	0.9692	0.65	0.9857
1.5C	144.22	0.9392	1.57	0.9744

4.2. Thermal-Fluid Calibration (Pack-Level)

With the SPMe validated, the 1D fluid dynamics and inter-cell LPTN were calibrated against 3D CFD reference data from Nema et al. [54]. The SPMe was replaced with static 1C and 5C constant current discharge volumetric heat generation profiles to isolate the thermal-fluid mechanics. The first validation step evaluated the thermal accumulation of the 28-cell module under passive ambient conditions.




Figure 4.2: Maximum module temperature validation under passive, ambient cooling conditions for 1C and 5C constant-current discharges, compared against 3D CFD reference data (Nema et al., 2026).

The integration of a physical module casing with thermal mass and heat capacity prevents the 1D LPTN from overestimating the transient temperature spikes. Initial active simulations revealed that the raw LPTN systematically overestimated transient temperature spikes. To rectify this, the framework's physical geometry was appended to include the thermal mass of the module casing, physically grounded by approximating a 10 mm thick bounding box enclosing the cells. With this physical thermal buffer integrated, the passive cooling scenario yielded an RMSE of 0.94°C ($R^2 = 0.963$) at 1C, and 4.05°C ($R^2 = 0.972$) at 5C. An interesting caveat of this band-aid fix is that the 1C and 5C scenarios show opposite behaviour during the later parts of the ambient cooling simulations, with the difference between the two increasing as more and more mass was added to the module casing. Future work should focus on creating a wider variety of CFD or experimental setups to validate what physical aspect the

reduced-order models specifically struggle with fitting, as that could not be clarified using the available data. For the scope of this research, which is mostly focussed on proof-of-concepts and qualitative research, this band-aid fix should be permissible since it allowed for better error metrics, improving the accuracy of the model within the bounds of this specific setup.

The introduction of active cooling fluid was then evaluated to test the convective heat extraction boundaries.




Figure 4.3: Maximum module temperature validation under active liquid cooling (0.3 m/s) for 1C and 5C constant-current discharges.

Under active cooling (0.3 m/s), the framework achieved an exceptional RMSE of 0.15°C ($R^2 = 0.949$) at 1C, and 2.16°C ($R^2 = 0.957$) at 5C. The 5C active plot reveals the limits of the 1D approximation; the true CFD data exhibits a slight concave-to-convex transition as the cooling fluid saturates, which the 1D LPTN exaggerates more, leading to a significant underestimation of the temperature in the later parts of the simulation. Initially, during development, the degree of exaggeration here was much higher, which led to the introduction of the previously discussed module casing. While further increasing the mass of the module casing improved the results for these two simulation scenarios, it worsened the fits for the ambient cooling scenarios. Since this was already a band-aid fix, and further increases did not seem to mirror real-world implementations, it was decided that the resulting parameter set was an acceptable trade-off point. Nevertheless, the high R^2 values confirm that the model remains highly reliable for both steady-state and transient scenarios.

Beyond maximum pack temperature, the ability of the fluid network to accurately transport heat downstream was assessed by evaluating the maximum temperature difference across the pack.




Figure 4.4: Validation of the maximum downstream temperature gradient (ΔT) across the 28-cell module under a 5C active cooling load.

The downstream thermal gradient validation yielded an impressive RMSE of 0.08°C ($R^2 = 0.981$). The simulation tightly hugged the 3D CFD ΔT data, proving that the 1D fluid network, augmented by the module casing thermal mass, accurately replicates the spatial downstream accumulation of heat without requiring computationally expensive 3D volumetric meshing. It is worth noting, however, that similar to the previous simulation scenarios, the limits of the reduced-order models become clear here due to the shapes of the curve not aligning exactly, which could be signs of the reduced-order models simply not having the underlying physical principles to model certain intricacies of temperature development in a module. In the case of the temperature difference, this could also be a result of the cells' cores being modelled as isothermal masses, while in reality, these areas are highly non-uniform thermally, with hot spots at the centres of the anode. In any case, the high R^2 do again show that the model remains highly reliable in these setups.

To prove that the model can fit and perform well in a wider variety of simulation scenarios, the hydrodynamic correlations were evaluated across a broad velocity sweep, which should allow for accurate modelling of both passive and active modes of a TMS.




Figure 4.5: Hydrodynamic and convective thermal validation across a coolant velocity sweep (0.1 m/s to 0.5 m/s) under a 5C discharge load.

The velocity sweep (0.1 to 0.5 m/s) proves the robustness of the implemented Dean number and Gnielinski correlations. The 1D fluid model successfully scaled across all tested pump velocities, maintaining R^2 values consistently between 0.945 and 0.95. This confirms that the model dynamically adjusts the heat transfer coefficient accurately, even when the flow transitions at higher coolant velocity rates. However, validating the total hydrodynamic pressure drop against the Nema reference data uncovered significant physical inconsistencies within the published CFD study itself.




Figure 4.6: Pressure drop and pump power validation for 40 mm and 50 mm serpentine channel heights.

The pressure and pump power validation required addressing a clear post-processing error in the reference paper. Nema et al. report pump power in milliwatts (e.g., 19.6 mW at 0.5 m/s), yet fundamental hydraulic physics ($W = Q \cdot \Delta P$) using their own reported pressure drops dictates a power output exactly ten times higher (~ 196 mW). After correcting the validation data to match physical reality, the 40 mm channel achieved a near-perfect pump power fit ($R^2 = 0.999$).

However, the 50 mm channel fit is extremely poor ($R^2 = -0.043$). The Nema study claims a 40% pressure drop reduction when the channel height increases from 40 mm to 50 mm. Under a constant velocity boundary condition, an extreme 50 mm $\times$ 2 mm aspect ratio is dominated by the narrow 2 mm gap; altering the height should barely impact the viscous wall shear. Looking at the implemented fundamental physics, this is exactly what the 1D framework predicts, because the friction factor is dictated by the hydraulic diameter, which is dominated by the narrowness of the channel rather than its height. This apparent physical contradiction highlights a probable boundary error or mesh-saturation issue in the reference 3D CFD.




Figure 4.7: Thermal validation across a serpentine channel height sweep (40 mm to 70 mm) under a 5C discharge.

Despite this hydrodynamic discrepancy in the pressure drop comparison, the convective temperature height sweep (40 mm to 70 mm) performed remarkably well. The RMSE values hovered between 2.63°C and 3.25°C, maintaining R^2 values above 0.91.

The reason the thermal accuracy remains unaffected by the potential pressure drop discrepancy lies in the acausal architecture of the framework. Because the 1D fluid network is driven by an ideal mass-flow boundary condition, the calculated hydraulic pressure drop does not recursively throttle the flow rate. The convective heat transfer is strictly governed by the local fluid velocity and the curvature-delayed transition regimes (Re_{crit}). Consequently, the pressure drop calculation acts purely as a decoupled diagnostic metric, used exclusively to calculate the parasitic pump power required by the TMS.

The slight decrease in thermal accuracy at the extreme heights (40 mm and 70 mm) simply suggests the bounds of the 1D Nusselt approximations, confirming that while the convective thermal model is highly robust, geometric extrapolation significantly outside the 50 mm validated baseline should be approached with caution.

4.3. Ablation Study: Impact of Coupling Domains

With the framework being functional and having been validated, it is now possible to study the effect of simulating load profiles with different combinations of physical domains. This should give insights into whether it is actually necessary to use the framework in its entirety, or whether simpler models, such as an isothermal SPMe will suffice. Since the context of this research is V2G-related degradation, the ablation study has been performed using simplified load profiles that are representative of the average current that EV batteries might experience during grid services.




Figure 4.8: Pack voltage, current and maximum cell temperature of the 0.1C limited ablation study load profile.

Using the load profile depicted in Figure 4.9, the following scenarios were evaluated:

- S1: Single SPMe, isothermal
- S2: Single SPMe, thermal
- S3: Pack, isothermal, electrically coupled
- S4: Pack, thermal, electrically decoupled
- S5: Pack, thermal, electrically coupled
- S6: Same as S5, but with doubled standard deviation

The isothermal simulations have their temperature locked at ambient, which in this case is 25 degrees Celsius. Electrical coupling means that heterogeneous current distributions are made possible through Ohmic current division, while decoupling is achieved by forcing an equal amount of current through all cells. Lastly, all multi-cell simulations, except for the sixth scenario, are subjected to a normally distributed initial capacity variation with $\sigma = 0.005$, meaning that 99.7% of cells will fall within 3σ , or a $\pm 1.5\%$ variation in initial capacity.




Figure 4.9: Average SoH across six different scenarios during a 180 day degradation experiment, showing practically no meaningful spread for a 0.1C limited load profile.

Table 4.2: Ablation study 0.1C limited load profile SoH spread.

Scenario	Final SoH (%)
S1 (iso, dec)	99.13
S2 (therm, dec)	99.13
S3 (iso, coup)	99.13
S4 (therm, dec)	99.13
S5 (therm, coup)	99.13
S6 (double var)	99.13

With the 0.1C limited load profile, all six scenarios are practically equally degraded after 180 days. Initially, this seems quite strange, since it would be expected that the inclusion of multiple coupled domains should show more degradation due to feedback effects and increased pack temperature. However, as the load profile shows, pack temperature does not increase significantly. Apparently, the rise in temperature is so insignificant, that the temperature-dependent degradation kinetics do not behave differently enough across scenarios to lead to a visible spread. Since this load profile has been limited to 0.1C, which is a current that is representative of many relatively calm to moderate V2G grid services, a preliminary conclusion might be that V2G-related degradation can simply be assessed using only a single SPM. Ideally however, it would also be interesting to be able to accurately assess degradation that occurs during WLTP (Worldwide harmonised Light vehicle Test Procedure) load profiles, which are representative of typical commutes through urbanised areas. These typically feature slightly more aggressive currents that could potentially generate more heat. To find out whether the higher temperatures generated by more aggressive load profiles could lead to a more significant degradation spread, the same load profile was used, but the 0.1C limit was increased to 0.75C instead.




Figure 4.10: Pack voltage, current and maximum cell temperature of the 0.75C limited ablation study load profile.

A rough Joule loss estimate predicts that increasing the current by a factor of 7.5 should lead to an increase in heat generation of more than 50 times. This, of course, is not entirely accurate, since reversible losses contribute significantly to heat generation at currents below 1.0C, and those scale proportionally to the current, but it can indeed be noted that the maximum temperature now peaks near 35 degrees Celsius, which is roughly 10 degrees above ambient.




Figure 4.11: Average SoH across six different scenarios during a 180 day degradation experiment, showing a slight degradation spread for a 0.75C limited load profile.

Now, with the more aggressive load profile, a clear divergence can be seen. After 180 days, there is a 6.7% relative difference between the least and most degraded scenarios. As expected, thermal coupling leads to more degradation as a result of higher temperatures being reached in pack, where

Table 4.3: Ablation study 0.75C limited load profile SoH spread.

Scenario	Final SoH (%)
S1 (iso, dec)	98.85
S2 (therm, dec)	98.85
S3 (iso, coup)	98.88
S4 (therm, dec)	98.80
S5 (therm, coup)	98.80
S6 (double var)	98.80

more cells generate heat and the average heat dissipation per volume of cell is lower than in single cell scenarios. Thermal feedback effects are practically negligible, as the thermal gradient is less than a tenth of a degree, even during this more aggressive load profile. This does line up with the observations that were made during the validation procedure, where it could be seen that a 5C discharge led to a thermal gradient of just a few degrees, so realistically, any currents below 1.0C will probably not generate any meaningful thermal gradients for this 28-cell pack, given these exact thermal conditions. It might be the case that a differently fitted LPTN will show worse thermal gradients at similar currents, but this will be discussed in more detail during the sensitivity analysis section near the end of this chapter. Aside from the effect of temperature, the electrical coupling also seemed to have had an effect. The isothermal electrically coupled case experienced the least amount of degradation, even when compared to the single cell scenarios. Especially when comparing S1 and S3, they should practically be equivalent, aside from the amount of cells that are simulated and the electrical coupling. This implies that the electrical coupling somehow reduces degradation, likely by forcing more current to pass through healthier cells and less current through more degraded cells. With a nearly symmetrical distribution of initial capacity variation, one could expect the degradation to then average out to the same value as for the single cell scenarios, but degradation is a non-linear process, which could help explain the offset. Interestingly, the thermally coupled pack-level scenarios do not show any meaningful differences in their final average SoH values. As of now, it is difficult to speculate what the exact reason is for the lack of degradation spread among these scenarios. These fundamentals are not the main focus of the research, and as such, it is mainly significant to conclude that initial capacity variations and electrical coupling do not seem to have a meaningful effect on degradation when thermal coupling is present. For higher current load profiles, pack-level simulation might be desirable for its improved accuracy in predicting the higher temperatures in a pack compared to single cell simulations. Thermal feedback effects do not play a role here, since the thermal gradients are practically negligible, but they might become relevant in scenarios with more cells and worse ambient cooling. Within the bounds of this research, a careful preliminary conclusion could be that, for relatively mild load profiles, just a single SPMe should suffice for accurate physics-based degradation assessments, since the first ablation study showed no meaningful degradation spread.

4.4. V2G Degradation Assessment

To make a realistic assessment of how much degradation is incurred through grid services, a load profile was constructed that attempts to mimic a generic usage pattern for EVs owned by the average working civilian, with the addition of FCR (Frequency Containment Reserve) as grid service. A single 24-hour cycle looks as follows:

- 30 minutes WLTP
- Charge to 65% SoC
- Rest for 8 hours
- 30 minutes WLTP
- FCR overnight
- Charge to 65% SoC half an hour before the end of the cycle

For the sake of simplicity, this 24-hour cycle is then repeated every day for an entire year. The reference simulation uses the same load profile, but instead of performing FCR overnight, it simply rests. Both

test scenarios are preceded by a pre-aging period of 5 years consisting of 24 hour cycles where the battery rests for 80% of the time, and the other 20% of the time, it charges and discharges around half capacity at a rate of 0.1C. The WLTP load profile was constructed using data from a Volkswagen ID.4 [55], and it was scaled down relative to the energy capacity of the 28-cell pack and the ID.4 battery pack. Similarly, the FCR load profile was constructed using grid frequency data and a droop of 0.4% without deadband [56]. This means that the battery pack will respond with a linear power throughput for deviations from the nominal frequency, with no throughput at the nominal frequency, and increasing up to the maximum throughput at deviations of ± 0.2 Hz. Here, the power was also scaled using a representative ID.4 battery pack, given that the maximum power throughput would be limited to 11 kW for FCR. This results in the following pack voltage, current and temperature profiles:




Figure 4.12: Pack voltage, current and temperature during the 24 hour load profile for the V2G-related degradation assessment, with and without FCR.

Relatively little additional heat is generated during the FCR periods, and even the WLTP profile does not generate a concerning amount of heat; the temperature does not come close to the 35 degrees Celsius threshold for activating the TMS. Given that 1.0C corresponds to a current of 35 A, the power scaling does seem reasonable.




Figure 4.13: SoH during 5 years of pre-aging, followed by 1 year of repeated 24 hour cycles with either 13.5 hour FCR or rest periods. Visually, there does not seem to be a significant difference between the two.

A visual comparison of the SoH traces during the profiles with and without FCR, there does not seem to be a significant difference. The data shows that there is actually a small difference, with the FCR profile actually having experienced less ageing by having a final SoH that is 0.1% in absolute percentages higher than the rest profile. It would be expected that doing anything else than resting would incur more degradation, but this behaviour can be explained perfectly fine with a caveat of how electrical coupling has been implemented in this framework. Heterogeneous current distributions are made possible through Ohmic current division, but this does mean that a non-zero pack current has to flow in order to generate heterogeneous current distributions. In the real world, parallel connected cells with different voltages will generate internal current loops, as their nodes should have equal voltages. This process of passive balancing is not present in this framework. Indeed, when comparing the voltage spread during the 1 year period, it can be found that the FCR profile ends with a maximum voltage delta of about 250 mV, while the rest profile ends approximately at 285 mV. Balancing can only happen actively in this framework, and since the FCR profile has relatively gentle C-rates during the FCR periods, it is likely that some slight balancing occurs here. Therefore, the resting profile has a slightly greater estimation of degradation compared to the FCR profile. Without any additional research where passive balancing is present, it is impossible to exactly quantify how much less the rest profile would degrade in reality, but with the SoH values being so close to each other after simulating for a year, and the ablation studies showing that balancing has little to no effect when thermal coupling is present, it can likely be concluded that, under these specific conditions, FCR does not significantly degrade batteries when compared to simply resting the battery.

With heat generation being so minimal, it seems like anticipative cooling does not have a valid business case, at least for FCR. However, there are more aggressive grid services, and following the conclusions from these results, it might be possible that, in the future, more aggressive grid services can be confidently performed more frequently due to the minimal effect on degradation. More aggressive grid services generate more heat, and therefore, anticipative cooling might still be relevant. To assess the validity of this argument, a stress test was constructed at an increased ambient temperature of 30 degrees Celsius, and the power limit for FCR was increased tenfold.




Figure 4.14: Pack voltage, current and temperature comparison of a 24 hour load profile cycle with FCR at a standard power limit and a 10x power limit. Even when the power is increased, and the ambient temperature sits closer to the 35 degrees Celsius TMS threshold, temperatures do not reach high enough to trigger the TMS.

With the power limit having been increased by a factor of ten, the voltage swings more significantly, and as such, the SoC now swings between 40% and 90% instead of hovering near 60%. The current regularly peaks around 0.5C, but the heat generation is still not great enough to achieve a temperature difference of more than 5 degrees. Since the reactive TMS does not even get triggered during these simulations, it is unlikely that anticipative cooling has a valid business case for city driving and calm to moderate V2G profiles. Anticipative cooling probably only becomes interesting once load profiles start to include large and sudden temperature spikes, which are much more likely to be found in load profiles containing highway driving or fast-charging. However, a significant caveat that should be considered, is that these conclusions do depend on the thermal characteristics of this system. While it is true that heat generation is relatively low for the simulated profiles, heat dissipation also plays a big role in deciding the temperature outcomes. The current framework is limited to a 28-cell module, and it has been fit using data from a single CFD setup. The following section will look at results from a sensitivity analysis to further assess the rigour of these conclusions.

4.5. Sensitivity Analysis

To study the sensitivity of the thermal characteristics of this framework to variations in parameters that might be highly specific to how they have been fitted uniquely for these conditions, a sensitivity analysis was performed on the following parameters:

- `cell_pitch`: shortest distance separating adjacent cells

- h_{conv} : ambient convection coefficient
- casing_th : thickness of the module casing surrounding the cells
- $R_{\text{contact_val}}$: contact resistance between cells and the walls of the coolant pipes
- $R_{\text{radial_val}}$: resistance between adjacent cells
- $R_{\text{axial_val}}$: resistance between cells and the casing (which is connected to ambient)

All parameters, except for the cell pitch, were included in this test, due to the fact that they had to be fitted manually in order to match the thermal characteristics of the CFD data. Typically, altering the cell pitch would also affect the amount of material that is placed between cells, therefore affecting the thermal characteristics. Since the conductive paths have been defined using effective thermal resistances, cell pitch does not affect the framework in the same way. Either increasing or decreasing the casing thickness will reduce the overall thermal dose, indicating that the baseline configuration sits at a local maximum. This highlights a non-linear trade-off between thermal inertia and thermal dissipation. Increasing the thickness adds thermal capacitance, acting as a buffer that successfully suppresses peak transient temperatures. Conversely, decreasing the thickness reduces thermal inertia; while this causes a slightly higher peak temperature, it allows the active TMS and ambient environment to extract heat much more rapidly during rest periods. The baseline thickness acts as a "thermal trap", lacking sufficient mass to suppress peak transients, yet possessing too much mass to allow rapid thermal shedding.




Figure 4.15: Sensitivity analysis of thermal network using a $\pm 10\%$ variation.

Table 4.4: Sensitivity analysis of thermal network using a $\pm 10\%$ variation.

Parameter	$T_{\max}$ Low	$T_{\max}$ High	ΔT Low	ΔT High	Dose Low	Dose High
R_axial_val	0.00%	-0.00%	-7.62%	7.35%	-0.07%	0.08%
h_conv	0.37%	-0.75%	1.65%	-5.13%	3.15%	-4.60%
casing_th	0.65%	-1.08%	4.68%	-7.27%	-1.69%	-3.07%
R_radial_val	0.00%	-0.00%	-2.01%	1.83%	-0.00%	0.00%
R_contact_val	0.00%	-0.00%	-0.44%	0.33%	-0.03%	0.02%
cell_pitch	0.35%	-0.35%	-1.00%	0.47%	0.82%	-0.72%

The peak temperature does not seem to be sensitive to changes in the parameters, but the thermal gradient and thermal dose do show a significant sensitivity. The thermal gradient is affected most significantly by the parameters related to the heat dissipation of the module, namely the ambient convection coefficient, which determines how much power per unit area can be dissipated, the axial thermal resistance, which determines how quickly heat can dissipate from the cells towards the casing and then to ambient, and the casing thickness, which affects the thermal mass and thermal capacitance of the casing. Interestingly, the radial resistance, which is literally the thermal resistance between adjacent cells, has a lesser effect. It might be the case that this relation will change when more cells are simulated. Any parameter that affects the dissipation capabilities of the thermal system will likely have a great effect on the thermal characteristics for such a small pack, since it has a relatively small volume to area ratio compared to larger packs. In the context of this research, which is centred around V2G-related degradation, this sensitivity is no cause of concern, since the thermal gradient is less than a tenth of a degree for currents up to 0.75C, and the ablation studies showed no significant differences at more gentle currents. Thermal dose is probably the most interesting metric of the sensitivity analysis, as the thermal dose is a measure of how long the cells spend at certain temperatures, which directly affects degradation through the temperature-dependent degradation kinetics. Here, it can be seen that mainly the casing thickness and the ambient convection coefficient affect the thermal dose. Both increasing and decreasing the casing thickness will reduce the thermal dose, which implies that the current configuration sits at a local maximum. This means that there is a non-linear trade-off between thermal mass and thermal capacitance.

Fortunately, since the heat generation is relatively low in the simulated load profiles, and the ablation study showed a relative difference of 6.7% in SoH during a 180 day simulation with a over 10 degrees difference in thermal peak, it is highly likely that the conclusions from the V2G comparison are still valid. However, since future research should focus on fitting this framework on more configurations anyway, it would be highly recommended to attempt to recreate and validate the exact heat dissipation conditions that an EV battery experiences inside of a car. Ideally, different fits should be used for WLTP, resting and V2G profiles, since they happen in different environments, which changes the ambient convection coefficient.

5

Conclusion

The primary objective of this research was to develop a computationally efficient, multi-physics framework capable of assessing the degradation of electric vehicle battery packs subjected to dynamic V2G load profiles, and to evaluate the necessity of anticipative thermal management strategies. Based on the validation data, ablation studies, and long-term V2G simulations, the following core conclusions can be drawn:

First, the developed framework provides a computational middle-ground between resource-intensive 3D CFD models and simplified single-cell ECMs. By coupling multiple SPMe to a 0D LPTN utilising empirical hydrodynamic correlations, the framework accurately replicates pack-level thermal-fluid behaviour. The validation against 3D CFD reference data confirmed that the 1D nodal fluid network accurately predicts macroscopic temperature peaks ($R^2 > 0.94$), spatial downstream thermal gradients, and pressure drops, all at a fraction of the computational cost required for full volumetric meshing.

Second, the ablation study showed that no significant degradation spread could be observed between simpler single-cell SPMe and variations of the pack-level framework at lower currents (limited to 0.1C). Meaningful spread only emerges at higher currents ($\geq 0.75C$), where the increased heat generation begins to accelerate temperature-dependent SEI growth. Electrical coupling did not seem to significantly affect degradation, but the current model only allows heterogeneous current distributions to exist when current actually flows. Passive balancing is not present in the model, and future research should look into the effect of this characteristic on the results of similar ablation studies. The future inclusion of particle cracking could also alter these results, but a careful preliminary conclusion could be that at lower currents, such as the ones seen under normal urban driving and mild V2G conditions, simpler single-cell SPMe can be sufficiently accurate to perform physics-based degradation assessments. More aggressive load profiles do benefit from using the entire framework, as the degradation spread between models diverges significantly over the course of multiple years.

Finally, the long-term V2G assessment concludes that standard FCR services impose a negligible thermal and degradative burden on the battery pack. Over a simulated 1-year FCR profile, heat generation was minimal, with pack temperatures never approaching the 35°C threshold required to trigger active cooling. Even under a more extreme stress-test scenario, where the FCR power limit was increased tenfold, the battery temperature remained safely within passive cooling limits. Therefore, for standard V2G frequency regulation, the implementation of complex, anticipative TMS offers no tangible benefit. Standard reactive cooling logic is entirely sufficient to maintain pack health under these specific grid-service load profiles.

5.1. Recommendations

While the developed framework provides a robust foundation for pack-level V2G simulation, several physical limitations were identified that present clear opportunities for future research:

- **Implementation of mechanical degradation (LAM):** The current framework restricts capacity fade to SEI layer growth and lithium plating. As a result, it penalises higher temperatures while ignoring the mechanical fatigue, in the form of particle cracking coupled with LAM, that occurs at lower temperatures due to sluggish lithium diffusion. Future iterations should implement the coupled mechanical degradation mechanism and mode to accurately capture the balance between high-temperature chemical ageing and low-temperature mechanical degradation.
- **Implementation of passive balancing:** The long-term V2G simulations highlighted a minor discrepancy where resting batteries seemingly degraded slightly more than active FCR batteries. This was identified as an artefact of the framework lacking passive balancing, as heterogeneous cells currently require active pack current to equalise. This discrepancy could be eliminated by implementing a physically accurate balancing method, where current loops will be generated if voltages differ among parallel connections.
- **Dynamic ambient convection parameters:** The sensitivity analysis proved that the thermal dose is highly sensitive to the ambient convection coefficient. Currently, the framework uses a static convection coefficient. Future studies could dynamically alter this coefficient to reflect the physical environment of the vehicle, for example, by applying a high coefficient during WLTP driving to simulate aerodynamic airflow, and a heavily reduced coefficient during parked V2G or resting states to mimic stagnant garage air. Aside from dynamic altering, the preliminary conclusions from this research need further verification with thermal parameters that accurately represent the conditions inside of parked EVs, with great emphasis on the ambient convection coefficient.
- **Advanced CFD validation geometries:** The 1D thermal-fluid network was validated against a specific 28-cell serpentine module. Because empirical form-drag coefficients do not universally scale, future work must validate the framework against a wider array of CFD architectures, such as bottom-plate cooling and larger 100+ cell modules. This will form a robust library of validated 1D hydrodynamic parameters that can be scaled reliably without requiring unique CFD validations for every new configuration.

Aside from these limitations, expanding the variety of simulated load profiles could be interesting for future ablation and V2G degradation studies. While this research focused on WLTP and FCR, investigating highway driving and fast-charging profiles seem to be logical next steps, as these higher-stress use cases are becoming increasingly dominant. Furthermore, since this research demonstrated that FCR induces negligible additional ageing, future studies should use this physics-based framework to assess more aggressive V2G ancillary services.


Bibliography

- [1] D. Anggraini, M. Amelin, and L. Söder, "Business models for electric vehicle charging considering grid limitations: An extended review from the european electricity market perspective," *Applied Energy*, vol. 399, 12 2025.
- [2] A. Jain, P. Das, S. Yamujala, R. Bhakar, and J. Mathur, "Resource potential and variability assessment of solar and wind energy in india," *Energy*, vol. 211, 11 2020.
- [3] Y. Chen, J. Liu, and D. Yu, "Wind meets solar: crafting the perfect duck curve for future power systems," *Journal of Atmospheric and Solar-Terrestrial Physics*, vol. 277, 12 2025.
- [4] H. Morão, "The 2025 iberian peninsula blackout: Lessons for modern power systems and policy implications," 8 2026.
- [5] F. Jung, M. Schröder, and M. Timme, "Exponential adoption of battery electric cars," Dresden University of Technology, Tech. Rep., 6 2023. [Online]. Available: <http://arxiv.org/abs/2306.16152>
- [6] N. S. Pearre, W. Kempton, R. L. Guensler, and V. V. Elango, "Electric vehicles: How much range is required for a day's driving?" *Transportation Research Part C: Emerging Technologies*, vol. 19, pp. 1171–1184, 2011.
- [7] S. Hasselwander, M. Senzeybek, and G. Möring-Martínez, "Predictive modeling of vehicle-to-grid flexibility: A bottom-up approach demonstrated for a case-study in germany," *International Journal of Electrical Power and Energy Systems*, vol. 172, 11 2025.
- [8] K. Linders, S. Jenu, A. Hentunen, and G. R. C. Mouli, "The impact of v2x on battery degradation: A quantitative review," 12 2025.
- [9] R. Ghotge, K. P. Nijssen, J. A. Annema, and Z. Lukszo, "Use before you choose: What do ev drivers think about v2g after experiencing it?" *Energies*, vol. 15, 7 2022.
- [10] Z. Machardy, A. Khan, K. Obana, and S. Iwashina, "V2x access technologies: Regulation, research, and remaining challenges," *IEEE Communications Surveys and Tutorials*, vol. 20, pp. 1858–1877, 7 2018.
- [11] M. Choi, X. Du, S. Y. Choe, and T. R. Garrick, "Effect of temperature gradients in large-format pouch-type nmc/graphite lithium-ion battery on degradation," *Journal of Power Sources*, vol. 677, 6 2026.
- [12] S. A. Cetegen, T. Gundersen, and P. I. Barton, "Evaluating economic feasibility of lithium-ion battery energy storage systems in future us electricity markets," *Energy*, vol. 337, 11 2025.
- [13] A. K. Koech, G. Mwandila, F. Mulolani, and P. Mwaanga, "Lithium-ion battery fundamentals and exploration of cathode materials: A review," pp. 321–339, 10 2024.
- [14] J. S. Edge, S. O'Kane, R. Prosser, N. D. Kirkaldy, A. N. Patel, A. Hales, A. Ghosh, W. Ai, J. Chen, J. Yang, S. Li, M. C. Pang, L. B. Diaz, A. Tomaszewska, M. W. Marzook, K. N. Radhakrishnan, H. Wang, Y. Patel, B. Wu, and G. J. Offer, "Lithium ion battery degradation: what you need to know," pp. 8200–8221, 4 2021.
- [15] X. Han, L. Lu, Y. Zheng, X. Feng, Z. Li, J. Li, and M. Ouyang, "A review on the key issues of the lithium ion battery degradation among the whole life cycle," 8 2019.
- [16] R. Wittman, A. Fresquez, B. Chalamala, and Y. Preger, "Degradation of li-ion cells beyond 80

- [17] X. Jiang, X. Shi, Y. Yao, W. Han, X. Zhou, B. Jin, Y. Lei, B. Yang, and K. Zhang, "Solid electrolyte interphase (sei) for silicon anodes: From evolution to interface engineering and system-level optimization," *Energy Storage Materials*, vol. 88, 5 2026.
- [18] U. S. Meda, L. Lal, S. M, and P. Garg, "Solid electrolyte interphase (sei), a boon or a bane for lithium batteries: A review on the recent advances," 3 2022.
- [19] F. Single, A. Latz, and B. Horstmann, "Identifying the mechanism of continued growth of the solid-electrolyte interphase," *ChemSusChem*, 12 2018. [Online]. Available: <http://arxiv.org/abs/1812.03841><http://dx.doi.org/10.1002/cssc.201800077>
- [20] X. G. Yang, Y. Leng, G. Zhang, S. Ge, and C. Y. Wang, "Modeling of lithium plating induced aging of lithium-ion batteries: Transition from linear to nonlinear aging," *Journal of Power Sources*, vol. 360, pp. 28–40, 2017.
- [21] J. E, H. Xiao, S. Tian, and Y. Huang, "A comprehensive review on thermal runaway model of a lithium-ion battery: Mechanism, thermal, mechanical, propagation, gas venting and combustion," 8 2024.
- [22] K. Takahashi and V. Srinivasan, "Examination of graphite particle cracking as a failure mode in lithium-ion batteries: A model-experimental study," *Journal of The Electrochemical Society*, vol. 162, pp. A635–A645, 2015.
- [23] Y. Zhuang, J. Zhu, M. Sun, T. Wang, M. Yang, Q. Zhang, Y. Yao, X. Zhu, L. Gu, and H. Xia, "Dynamic lattice mismatch induced by local jahn-teller distortion in layered mn-rich cathode," *Journal of Materials Science and Technology*, vol. 277, pp. 85–92, 1 2027.
- [24] A. Karger, S. E. J. O'Kane, M. Rogge, C. Kirst, J. P. Singer, M. Marinescu, G. J. Offer, and A. Jossen, "Modeling particle versus sei cracking in lithium-ion battery degradation: Why calendar and cycle aging cannot simply be added," *Journal of The Electrochemical Society*, vol. 171, p. 090512, 9 2024.
- [25] C. J. Ko, C. H. Chen, and K. C. Chen, "Solid electrolyte interphase (sei) transport, evolution, and fracture: Governing factors for lithium dendrite formation," *Journal of Power Sources*, 6 2026.
- [26] M. K. Tran, A. Dacosta, A. Mevawalla, S. Panchal, and M. Fowler, "Comparative study of equivalent circuit models performance in four common lithium-ion batteries: Lfp, nmc, lmo, nca," *Batteries*, vol. 7, 9 2021.
- [27] W. Vermeer, G. R. C. Mouli, and P. Bauer, "A comprehensive review on the characteristics and modeling of lithium-ion battery aging," *IEEE Transactions on Transportation Electrification*, vol. 8, pp. 2205–2232, 6 2022.
- [28] M. Doyle, T. Fuller, and J. Newman, "Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell," University of California, Tech. Rep., 1993.
- [29] F. B. Planella, W. Ai, A. M. Boyce, A. Ghosh, I. Korotkin, S. Sahu, V. Sulzer, R. Timms, T. G. Tranter, M. Zyskin, S. J. Cooper, J. S. Edge, J. M. Foster, M. Marinescu, B. Wu, and G. Richardson, "A continuum of physics-based lithium-ion battery models reviewed," *Progress in Energy*, vol. 4, 10 2022.
- [30] K. B. O'Regan, "Developing parametrisation methods for physics-based models of lithium and sodium-ion batteries," University of Birmingham, Tech. Rep., 2024.
- [31] P. L. Ruiz, N. Damianakis, and G. R. C. Mouli, "Physics-based and data-driven modeling of degradation mechanisms for lithium-ion batteries - a review," pp. 21 164–21 189, 2025.
- [32] K. Gao, Q. Li, L. Hu, J. Huang, H. Li, and Y. Wu, "Physical informed neural network for soh estimation of lithium-ion battery with electrochemical mechanism," *Energy*, vol. 342, 1 2026.
- [33] M. Mama, E. Solai, T. Capurso, A. Danlos, and S. Khelladi, "Comprehensive review of multi-scale lithium-ion batteries modeling: From electro-chemical dynamics up to heat transfer in battery thermal management system," 2 2025.

- [34] V. Lucaferri, M. Quercio, A. Laudani, and F. R. Fulginei, "A review on battery model-based and data-driven methods for battery management systems," 12 2023.
- [35] H. E. Perez, J. B. Siegel, X. Lin, A. G. Stefanopoulou, Y. Ding, and M. P. Castanier, "Parameterization and validation of an integrated electro-thermal lfp battery model," University of Michigan, Tech. Rep., 2012.
- [36] F. B. Planella, M. Sheikh, and W. D. Widanage, "Systematic derivation and validation of a reduced thermal-electrochemical model for lithium-ion batteries using asymptotic methods," *Electrochimica Acta*, vol. 388, 8 2021.
- [37] L. Mattia, H. Beiranvand, W. Zamboni, and M. Liserre, "Lithium-ion battery thermal modelling and characterisation: a comprehensive review," 10 2024. [Online]. Available: <https://www.techrxiv.org/doi/full/10.36227/techrxiv.172927224.42356065/v1>
- [38] T. Iriyama, M. Carter, G. M. Cavaleiro, P. Poudel, G. J. Nelson, and G. Zhang, "Effects of non-uniform temperature distribution on the degradation of liquid-cooled parallel-connected lithium-ion cells," *Batteries*, vol. 10, 8 2024.
- [39] S. Li, C. Zhang, Y. Zhao, G. J. Offer, and M. Marinescu, "Effect of thermal gradients on inhomogeneous degradation in lithium-ion batteries," *Communications Engineering*, vol. 2, 12 2023.
- [40] M. N. Marlow, J. Chen, and B. Wu, "Degradation in parallel-connected lithium-ion battery packs under thermal gradients," *Communications Engineering*, vol. 3, 12 2024.
- [41] S. M. Jordan, C. O. Schreiber, M. Parhizi, and K. Shah, "A new multiphysics modeling framework to simulate coupled electrochemical-thermal-electrical phenomena in li-ion battery packs," *Applied Energy*, vol. 360, 4 2024.
- [42] Q. Xia, D. Yang, Z. Wang, Y. Ren, B. Sun, Q. Feng, and C. Qian, "Multiphysical modeling for life analysis of lithium-ion battery pack in electric vehicles," *Renewable and Sustainable Energy Reviews*, vol. 131, 10 2020.
- [43] H. He and X. Chen, "Analysing unbalanced ageing in ev battery packs using the low-cost lumped single particle model (lspm): The impact of temperature gradients among parallel-connected cells," in *Transportation Research Procedia*, vol. 70. Elsevier B.V., 2023, pp. 406–413.
- [44] H. Pang, X. Yan, N. Jiang, X. Qu, B. Jing, A. F. Burke, and J. Zhao, "Electro-thermal-aging modeling of li-ion batteries with active-material loss," *Applied Energy*, vol. 410, 5 2026.
- [45] H. G. Lee, J. W. Choi, S. T. Ryu, and K. J. Lee, "Life degradation of lithium-ion batteries under vehicle-to-grid operations based on a multi-physics model," *International Journal of Automotive Technology*, vol. 26, pp. 315–326, 4 2025.
- [46] V. Tomanik, P. Bares, A. N. Aronis, and B. Pla, "Integrated predictive control of battery thermal management in electric vehicles during driving and fast charging," *Applied Thermal Engineering*, vol. 289, 3 2026.
- [47] F. B. Planella and W. D. Widanage, "A single particle model with electrolyte and side reactions for degradation of lithium-ion batteries," *Applied Mathematical Modelling*, vol. 121, pp. 586–610, 9 2023.
- [48] X. Lin, H. E. Perez, S. Mohan, J. B. Siegel, A. G. Stefanopoulou, Y. Ding, and M. P. Castanier, "A lumped-parameter electro-thermal model for cylindrical batteries," *Journal of Power Sources*, vol. 257, pp. 12–20, 2014.
- [49] K. O'Regan, F. B. Planella, W. D. Widanage, and E. Kendrick, "Thermal-electrochemical parameters of a high energy lithium-ion cylindrical battery," *Electrochimica Acta*, vol. 425, 9 2022.
- [50] S. E. O'Kane, W. Ai, G. Madabattula, D. Alonso-Alvarez, R. Timms, V. Sulzer, J. S. Edge, B. Wu, G. J. Offer, and M. Marinescu, "Lithium-ion battery degradation: How to model it," *Physical Chemistry Chemical Physics*, vol. 24, no. 13, 2022.

- [51] S. Kakac, *Handbook of Single-Phase Convective Heat Transfer*. Wiley-Interscience, 1987.
- [52] Volkswagen, "The high-voltage system in the id.4 tablet format," LLC Volkswagen Group of America, Tech. Rep., 2020.
- [53] C.-H. Chen, F. B. Planella, K. O'Regan, D. Gastol, W. D. Widanage, and E. Kendrick, "Development of experimental techniques for parameterization of multi-scale lithium-ion battery models," *Journal of The Electrochemical Society*, vol. 167, p. 080534, 2020. [Online]. Available: <https://doi.org/10.1149/1945-7111/ab9050>
- [54] P. K. Nema, P. Muthukumar, and R. Thangavel, "Electrochemical-thermal coupled modeling of a serpentine-shaped liquid cooling channel for lithium-ion battery packs with high discharge rates," *International Journal of Thermal Sciences*, vol. 220, p. 110299, 2026. [Online]. Available: <https://doi.org/10.1016/j.ijthermalsci.2025.110299>
- [55] A. Menendez-Agudin, G. F. Aznar, P. Bauer, and G. R. C. Mouli, "Quantifying energy transport by electric vehicles: A monte carlo and optimization framework for flexible energy communities," *Smart Energy*, vol. 21, p. 100230, Feb 2026.
- [56] "Rte services portal," accessed: 2026-06-29. [Online]. Available: <https://www.services-rte.com/en/view-data-published-by-rte.html>



Framework Parameterisation

This appendix provides the complete parameter datasets utilised to formulate the coupled multi-physics framework. The parameters are divided into electrochemical properties (Section A.1) and thermal-fluid properties (Section A.2).

A.1. Electrochemical Parameters (LG M50)

The foundational electrochemical parameters for the LG M50 21700 cell are derived from the experimental tear-down dataset provided by Chen et al. [53].

Table A.1: Geometric and solid-phase parameters for the negative and positive electrodes.

Parameter	Symbol	Negative Electrode	Positive Electrode
Electrode thickness	L_k	$85.2 \mu\text{m}$	$75.6 \mu\text{m}$
Particle radius	R_k	$5.86 \mu\text{m}$	$5.22 \mu\text{m}$
Active surface area density	a_k	$3.84 \times 10^5 \text{ m}^{-1}$	$3.82 \times 10^5 \text{ m}^{-1}$
Maximum lithium concentration	$c_{k,max}$	33133 mol m^{-3}	63104 mol m^{-3}
Initial lithium concentration (100% SoC)	$c_{k,0}$	29866 mol m^{-3}	17038 mol m^{-3}
Solid-state diffusivity	D_k	$3.30 \times 10^{-14} \text{ m}^2 \text{ s}^{-1}$	$4.00 \times 10^{-15} \text{ m}^2 \text{ s}^{-1}$
Electronic conductivity	σ_k	215 S m^{-1}	0.18 S m^{-1}
Reaction rate constant	m_k	6.48×10^{-7}	3.42×10^{-6}
Electrode porosity	ϵ_k	0.25	0.335
Bruggeman coefficient	b_k	1.5	1.5
Stoichiometry Limits			
Stoichiometry at 0% SoC	z_0	0.0279	0.9072
Stoichiometry at 100% SoC	z_{100}	0.9014	0.2567

Table A.2: Electrolyte parameters (LiPF₆ in organic solvents).

Parameter	Symbol	Value
Separator thickness	L_s	$12.0 \mu\text{m}$
Separator porosity	ϵ_s	0.47
Separator Bruggeman coefficient	b_s	1.5
Initial concentration	$c_{e,0}$	1000 mol m^{-3}
Cation transference number	t_+	0.2594
Diffusivity (c in mol m^{-3})	$D_e(c)$	$10^{-10} \times [0.879(c/1000)^2 - 3.97(c/1000) + 4.86]$
Conductivity (c in mol m^{-3})	$\sigma_e(c)$	$0.13(c/1000)^3 - 2.51(c/1000)^{1.5} + 3.33(c/1000)$

Table A.3: Solid Electrolyte Interphase (SEI) degradation kinetics parameters.

Parameter	Symbol	Value
Reaction rate constant	k_{sei}	$1.0 \times 10^{-10} \text{ A m}^{-2}$
Transfer coefficient	α_{sei}	0.5
Molar mass of SEI product	M_{sei}	$0.162 \text{ kg mol}^{-1}$
Density of SEI product	ρ_{sei}	1690 kg m^{-3}
SEI film conductivity	σ_{sei}	$5.0 \times 10^{-6} \text{ S m}^{-1}$
Open-circuit potential of side reaction	U_{sei}	0.4 V
Initial SEI thickness	$L_{f,0}$	$5.0 \times 10^{-9} \text{ m}$
Partial molar volume	$\bar{V}_{\text{sei}}$	$9.585 \times 10^{-5} \text{ m}^3 \text{ mol}^{-1}$
Solvent diffusivity in SEI	D_{sol}	$2.5 \times 10^{-22} \text{ m}^2 \text{ s}^{-1}$
Solvent concentration	c_{sol}	$2636.0 \text{ mol m}^{-3}$

A.2. Thermal-Fluid and Pack-Level Parameters

The pack-level thermal and fluid domains were parameterised using reference 3D CFD geometries from Nema et al. (2026) and thermal boundary properties derived from O'Regan et al. (2022) [49].

Table A.4: Thermophysical properties of the solid materials and coolant fluid.

Material	Density (ρ) (kg m^{-3})	Specific Heat (c_p) ($\text{J kg}^{-1} \text{ K}^{-1}$)	Conductivity (k) ($\text{W m}^{-1} \text{ K}^{-1}$)
Aluminium (channel walls/casing)	2700.0	900.0	238.0
Bergquist TGF 1500 (potting)	2700.0	1000.0	1.8
Coolant (Liquid)	Density (ρ)	Specific Heat (c_p)	Viscosity (μ)
Water (baseline CFD reference)	998.0	4180.0	0.002 Pa s

Table A.5: Core-shell Lumped Parameter Thermal Network (LPTN) characteristics and calibrated resistances.

Parameter	Symbol	Value
Jelly roll mass	m_{jelly}	0.05708 kg
Jelly roll volume	V_{jelly}	$2.13 \times 10^{-5} \text{ m}^3$
Shell thermal capacitance	C_{shell}	5.075 J K^{-1}
Internal radial thermal resistance	R_{rad}	0.535 K W^{-1}
Internal axial thermal resistance (inc. insulator cap)	R_{ax}	2.225 K W^{-1}

Table A.6: Calibrated pack-level inter-cell LPTN thermal resistances. These values dictate the macro-scale heat transfer pathways across the 28-cell module and were parametrised against reference 3D CFD data.

Parameter	Symbol	Fitted Value
Axial resistance (cell to casing boundary)	R_{axial}	0.25 K W^{-1}
Radial resistance (between adjacent cells)	R_{radial}	0.20 K W^{-1}
Contact resistance (cell to TMS cooling channel)	R_{contact}	4.00 K W^{-1}

Table A.7: Calibrated 1D fluid network and hydrodynamic parameters. The minor loss coefficient was empirically fitted to capture the form drag of the 180° serpentine bends.

Parameter	Symbol	Implementation / Value
Minor loss form factor (Serpentine bends)	K	$42 \cdot L \text{ (m}^{-1}\text{)}$
Baseline laminar friction factor (Darcy)	f	$1/Re$

A.2.1. Temperature-Dependent Specific Heat Capacity Polynomials

To accurately capture the thermal inertia of the cell, the specific heat capacity of the internal jelly roll core ($c_{p,\text{core}}$) is calculated dynamically as a function of temperature (T in Kelvin) using the exact mass fractions of its constituent layers. It is defined as the sum of the following third-order polynomials:

$$\begin{aligned}
 c_{p,\text{pos}} &= 0.4571 \times (-0.0008414T^3 + 0.7892T^2 - 241.3T + 25080) \\
 c_{p,\text{neg}} &= 0.2892 \times (0.0004932T^3 - 0.4910T^2 + 169.4T - 18970) \\
 c_{p,\text{sep}} &= 0.0343 \times (0.001494T^3 - 1.4440T^2 + 475.5T - 51300) \\
 c_{p,\text{al}} &= 0.0503 \times (4.503 \times 10^{-6}T^3 - 0.006256T^2 + 3.281T + 355.7) \\
 c_{p,\text{cu}} &= 0.1025 \times (1.445 \times 10^{-6}T^3 - 0.001946T^2 + 0.9633T + 236.0) \\
 c_{p,\text{elec}} &= 0.0666 \times 229.0
 \end{aligned}$$

Table A.8: Module geometry and TMS dimensions.

Parameter	Symbol	Value
Cell centre-to-centre pitch	D_{pitch}	25.0 mm
Cell diameter	D_{cell}	21.0 mm
Coolant channel width	W_{channel}	2.0 mm
Coolant channel height	H_{channel}	50.0 mm
Channel wall thickness	t_{wall}	1.0 mm
Module casing thickness	t_{casing}	10.0 mm
Ambient convection coefficient	h_{conv}	$5.0 \text{ W m}^{-2} \text{ K}^{-1}$