Imperial College London
Department of Mechanical Engineering

The Effect of Micro and Nano Scale Phenomena on System Level Design and Engineering of Hybrid PEMFC Systems

Oisín Sebastian Shaw

Submitted in part fulfil ment of the requirements
for the degree of Doctor of Philosophy at
Imperial College London, June 2024



I hereby declare that this thesis and the work reported herein was composed by and originated entirely from me. Information derived from the published and unpublished work of others has been acknowledged in the text and references are given in the list of sources.

Oisín Sebastian Shaw (2024)

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Licence (CC BY-NC-ND).

Under this licence, you may copy and redistribute the material in any medium or format on the condition that; you credit the author, do not use it for commercial purposes and do not distribute modified versions of the work.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Hybrid PEMFC vehicles are an important technological solution in the electrification of transport and the viability of sustainable systems. The degradation of Lithium Ion batteries used in this application is heavily dependent on the thermal management of the system, as well as the ratio between total storage capacity and peak power required. These are also inherently linked in implementation due to the non-isotropic thermal conductivity of the devices, and so de-coupling their impact on long term performance is impossible in situ. To better understand the separate contributions, as well as the interaction between these two parameters with respect to degradation, custom cells composed of a single layer of active material were commissioned. these were then mounted in a custom rig for the purpose of ensuring uniform temperature across all of the active material, and cycled at a variety of temperatures and discharge rates. The results show that above a threshold discharge rate the degradation is significantly accelerated, as well as the shift between degradation modes across a range of temperatures. In addition, the specific power of a PEMFC can be increased by use of current perturbation, a well known technique. However, the mechanism through which this acts is poorly understood, as well as the implementation being rudimentary. A custom PEMFC system was built to demonstrate that modern automotive style PEMFC designs do still benefit from this technique, and that the parameters of the perturbation can be optimised in operando to yield greater benefit. Finally, the voltage response during and after a perturbation was measured at significantly higher sample rate than any present in the literature, giving greater insight into both the mechanism through which the technique acts, as well as informing more elegant future implementations to minimise drawbacks, such as accelerated catalyst and GDL degradation due to frequent voltage cycling.

Acknowledgements

Firstly, to my West London family, without whom I would be lost. Georgia, who keeps me from being the idiot I was born to be, Ren and their bottomless pit of optimism, Ollie who must put up with me when everyone else gets to go home, Sarah for helping me embrace my inner goblin and Beki, whose chaos will never fade.

Special thanks to Dr. Greg Offer for the opportunity of this studentship, but also for his understanding, guidance and kindness along the way. I doubt you knew what you were inviting in that first day a wavering undergrad dropped by with a wild project proposal, but your encouragement has been invaluable and will never be forgotten.

Thanks also to Dr. Billy Wu for his guidance, mysterious luck at poker and the small matter of a lab in which to complete my work.

Thanks for all of the guidance and support to Dr. Huizhi Wang and Dr. Monica Marinescu.

To all my friends and students at the Makerspace, a true escape when it is most needed. Kate, Becca, Richard, Marta, Chloe, Monica, Kyrillos and Ela, as well as more students than I can name, here's to many more wacky projects to come.

To everyone at ICAH, without whom this project would have been done 3 years earlier, but with none of the style.

To all my friends in the ESE research group, past and present, for putting up with my singing in the mornings. Max, Emma, Teddy, Waseem, Charlotte, Ryan, Gavin, Yan, Ian, Tom, James, Sam, both Lauras, Karthik, Anna and many more.

To all of the DE gang, wherever they are now, and the many parties that will either never be forgotten, or never fully remembered. Sam, Larissa, Audrey, Billy, Shayan, Anouk, Leila and many more!

To my friends from undergrad, who have been so patient with my inability to schedule, Rob, Joe, Claudia and more, here's to many more celebrations!

To Richard and Giorgos, for their friendship and help with many a wildly impractical circuit!

To Kiah, without your unquestioning kindness none of this would be possible, and to all the girls from KES.

Y finalmente a Ruth y Juana, por todo vuestro apoyo y cariño. Las quiero mucho.

The path I choose through that maze makes me what I am

Daniel Keyes
Flowers for Algernon

Contents

Abstract iii

Acknowledgements iv

List of Tables ix

List of Figures x

Nomenclature xiv

1 Introduction 1
- 1.1 Project motivation . . . 1
- 1.2 Project objectives and research questions . . . 1
- 1.3 Thesis structure . . . 3

2 Literature Review 5
- 2.1 Introduction . . . 5
- 2.2 Fuel Cells . . . 7
  - 2.2.1 Types of Fuel Cell . . . 7
  - 2.2.2 Fuels . . . 10
  - 2.2.3 Hybridisation . . . 13
  - 2.2.4 Current Perturbation . . . 16
- 2.3 Lithium Ion Batteries . . . 18
- 2.4 Scientific Techniques and Procedures . . . 22
  - 2.4.1 EIS . . . 22
  - 2.4.2 GITT . . . 25
  - 2.4.3 ICA . . . 26
  - 2.4.4 DMA . . . 27
- 2.5 Conclusion . . . 30

3 Decoupling the Impact of Discharge Rate and Temperature on the Degradation of Li-ion Cells - Experimental design and validation 32
- 3.1 Introduction . . . 32
- 3.2 Battery background theory . . . 33
  - 3.2.1 Mechanical structure . . . 33
  - 3.2.2 Materials . . . 36
  - 3.2.3 Performance Metrics . . . 37
- 3.3 Methods . . . 41
  - 3.3.1 Introduction . . . 41

3.3.2 Cells . . . 42
3.3.3 Mechanical fixture . . . 43
3.3.4 Active thermal management for pouch cell control experiments . . . 50
3.3.5 Cycling and Characterisation . . . 51
3.3.6 Practical considerations . . . 55
3.3.7 Selection of Experimental Conditions . . . 56
3.3.8 Terminology and other Miscellany . . . 57
3.4 Conclusion . . . 58

4 Decoupling the Impact of Discharge Rate and Temperature on the Degradation of Li-ion Cells - Results 60
4.1 Capacity analysis . . . 61
4.1.1 Beginning of life . . . 61
4.1.2 Round 1 Experiments . . . 62
4.1.3 Round 2 Experiments . . . 68
4.2 End of cycle voltages . . . 71
4.3 Incremental Capacity Analysis . . . 73
4.4 Degradation Mode Analysis . . . 87
4.4.1 Round 2 . . . 96
4.5 Pouch Control Experiments . . . 106
4.5.1 Beginning of life capacity . . . 106
4.5.2 Capacity Throughout Testing . . . 108
4.5.3 End of Cycle Voltages . . . 113
4.5.4 Incremental Capacity Analysis . . . 114
4.5.5 Degradation Mode Analysis . . . 116
4.6 Discussion . . . 122
4.7 Conclusion . . . 132

5 Current Perturbations for PEMFC Systems - Background theory and experimental design 135
5.1 Introduction . . . 135
5.2 Proton Exchange Membrane Fuel Cells - Principles of Operation and Background Theory . . . 137
5.2.1 Principles of operation . . . 137
5.2.2 Mechanical design . . . 138
5.2.3 Electrochemical operation . . . 143
5.2.4 Current perturbations . . . 149
5.2.5 Discussion . . . 154
5.3 Focus of work . . . 155
5.3.1 Research questions . . . 155
5.3.2 Experiment parameters . . . 158
5.3.3 Optimisation . . . 159
5.3.4 Analytical Solution . . . 162
5.4 Methods . . . 168
5.4.1 Selection of materials . . . 168
5.4.2 Design and Fabrication . . . 170

6 Current Perturbations for PEMFC Systems - Results and discussion 180
6.1 Introduction . . . 180
6.2 Results . . . 181
6.2.1 CVM Data Validity . . . 181
6.2.2 Voltage Response During a Perturbation . . . 183
6.2.3 Voltage Response Immediately After a Perturbation . . . 186
6.2.4 Voltage Response Between Perturbations . . . 188
6.2.5 Optimisation . . . 192
6.3 Discussion . . . 199
6.3.1 Optimisation . . . 201
6.3.2 Future Implementations . . . 202
6.3.3 Mechanism theory . . . 206
6.4 Conclusion . . . 208

7 Conclusion 211
7.1 Summary of Thesis Achievements . . . 212
7.2 Future Work . . . 215

A Diagrams of Single Layer Rig 239

A Permissions for reproduced works 242

List of Tables

2.1 A table showing the volumetric and gravimetric energy densities of popular automotive energy sources . . . 10
2.2 A table showing the classification of pressure vessels, as well as some brief notes about their use in PEMFC EVs . . . 12
2.3 A table listing the in- and through-plane thermal conductivity bounds for a selection of LIB chemistries, adapted from Wei *et al*[1] . . . 21

3.1 A table listing the properties of PE intercalation compounds, adapted from[2] . . 37
3.2 A table comparing various NE materials, adapted from[3] . . . 38
3.3 A table showing the manufacturer specification of the Belleville spring washers used 47
3.4 A table showing the experimental conditions applied to single layer cells, in which round of testing they were performed, and the colour used to represent them in figures . . . 58

5.1 A table showing the details of the terms of the Nernst Equation . . . 143
5.2 A table showing the details of the terms of the Butler-Volmer Equation . . . 145
5.3 Ohmic Resistances (Rm/mΩ) before and after 10s short circuit. Reproduced with permission from Silva *et al*[4] . . . 152

List of Figures

2.1 A diagram detailing the impact of degree of hybridisation on the performance of a vehicle . . . 15
2.2 Typical EIS results for LIBs and PEMFCs . . . 23
2.3 A hypothetical example of the problems with under- and over-fitting curves to real data for calculation of the differential curve. Reproduced with permission from Feng *et al*[5] . . . 28

3.1 A diagram showing the basic structure of a LIB Included with permission from[6] 35
3.2 CAD images of the single layer rig design . . . 44
3.3 Rig thermal performance validation results . . . 49

4.1 A histogram showing the BOL capacities of single layer cells tested in Round 1 of experiments . . . 61
4.2 A histogram showing the BOL capacities of single layer cells tested in Round 2 of experiments . . . 62
4.3 A histogram showing the BOL capacities of single layer cells tested in Rounds 1 and 2 of experiments . . . 63
4.4 A plot showing the Normalised Capacity against Equivalent Cycles of cells tested at 30°C and discharged at 1C . . . 64
4.5 A plot showing the standard deviation of capacity measurements against equivalent number of cycles for each round 1 experiment . . . 65
4.6 A plot showing the Normalised Capacity against Equivalent Cycles of cells tested at 30°C and discharged at a variety of C-rates . . . 66
4.7 A plot showing the Normalised Capacity against Equivalent Cycles of cells discharged at 1C at different temperatures . . . 67
4.8 A plot showing the Normalised Capacity against Equivalent Cycles of all round 1 experiments . . . 68
4.9 A plot showing the Normalised Capacity against Equivalent Cycles of selected round 2 experiments . . . 69
4.10 A plot showing the Normalised Capacity against Equivalent Cycles of round 2 experiments discharged at 4C . . . 70
4.11 A plot showing the standard deviation of capacity measurements against equivalent number of cycles for each round 2 experiment . . . 71
4.12 A plot showing the lower voltage reached during cycling against equivalent number of cycles for each round 2 cell . . . 72
4.13 A plot showing the lower voltage reached during cycling against equivalent number of cycles for each round 1 cell . . . 74
4.14 A plot showing the measured capacity against equivalent cycles of cells tested at 30°C and discharged at 4C . . . 74

4.15 A map of the major degradation modes impact on the ICA curve of an NMC 111/GIC cell[7] . . . 75
4.16 A plot of the ICA results of a cell cycled at 30°C, 1C discharge . . . 78
4.17 A plot of the ICA results of a cell cycled at 30°C, 4C discharge . . . 79
4.18 A plot of the ICA results of a cell cycled at 45°C, 1C discharge . . . 80
4.19 A plot of the ICA results of a cell cycled at 10°C, 1C discharge . . . 81
4.20 A plot of the ICA results at beginning of testing of all round 2 cells, overlaid with an average of the round 1 results . . . 82
4.21 A plot of the ICA results of a cell cycled at 20°C, 1C discharge . . . 83
4.22 A plot of the ICA results of a cell cycled at 20°C, 4C discharge . . . 84
4.23 A plot of the ICA results of a cell cycled at 55°C, 1C discharge . . . 85
4.24 A plot of the ICA results of a cell cycled at 20°C, 4C discharge . . . 85
4.25 A plot of the ICA results of a cell cycled at 10°C, 4C discharge . . . 86
4.26 A plot of the ICA results of a cell cycled at 45°C, 4C discharge . . . 87
4.27 A plot of the lithiation fraction bounds against charge throughput of cells tested at 30°C and discharged at a 1C, calculated using the DMA technique . . . 89
4.28 A plot of the electrode capacities against charge throughput of cells tested at 30°C and discharged at a 1C, calculated using the DMA technique . . . 89
4.29 A plot of the DMA results against charge throughput of cells tested at 30°C and discharged at 1C . . . 91
4.30 A plot of the lithiation fraction bounds against charge throughput of cells tested at 30°C and discharged at various rates . . . 93
4.31 A plot of the DMA results against charge throughput of cells tested at 30°C and discharged at various rates . . . 93
4.32 A plot of the lithiation fraction bounds against charge throughput of round 1 cells discharged at 1C, various temperatures . . . 94
4.33 A plot of the DMA results against charge throughput of round 1 cells discharged at 1C, various temperatures . . . 95
4.34 A plot of the DMA results against charge throughput of round 1 cells discharged at 1C, various temperatures . . . 96
4.35 A plot of the lithiation fraction bounds against charge throughput of cells discharged at 1C, various temperatures . . . 98
4.36 A plot of the electrode capacities against charge throughput of cells discharged at 1C, various temperatures . . . 99
4.37 A plot of the DMA results against charge throughput of cells discharged at 1C, various temperatures . . . 100
4.38 A plot of the lithiation fraction bounds against charge throughput of cells discharged at 4C, various temperatures . . . 101
4.39 A plot of the electrode capacities against charge throughput of cells discharged at 4C, various temperatures . . . 102
4.40 A plot of the DMA results against charge throughput of cells discharged at 4C, various temperatures . . . 103
4.41 A plot of the lithiation fraction bounds against charge throughput of cells discharged at 4C, various temperatures . . . 104
4.42 A plot of the electrode capacities against charge throughput of cells discharged at 4C, various temperatures . . . 105
4.43 A plot of the DMA results against charge throughput of cells discharged at 4C, various temperatures . . . 105

4.44 A histogram showing the BOL capacities of the pouch control cells 107
4.45 A plot showing the Measured Capacity against Equivalent Cycles of pouch control cells tested at 30∘C and discharged at 1C 109
4.46 A plot showing the standard deviation of capacity measurements against equivalent number of cycles for each pouch control cell experiment 110
4.47 A plot showing the Normalised Capacity against Equivalent Cycles of cells tested at 30∘C and discharged at a variety of C-rates 111
4.48 A plot showing the Normalised Capacity against Equivalent Cycles of cells discharged at 1C, various temperatures 113
4.49 A plot showing the lower voltage reached during cycling against equivalent number of cycles for each pouch control cell 114
4.50 A plot of the ICA results of a pouch control cell cycled at 30°C, 1C discharge . . 115
4.51 A plot of the ICA results of a pouch control cell cycled at 30°C, 4C discharge . . 116
4.52 A plot of the lithiation fraction bounds against charge throughput of pouch control cells tested at 30°C and discharged at various rates 118
4.53 A plot of the electrode capacities against charge throughput of pouch control cells tested at 30°C and discharged at various rates 119
4.54 A plot of the DMA results from pouch control cells tested at 30°C and discharged at various rates 120
4.55 A plot of the lithiation fraction bounds against charge throughput of pouch control cells discharged at 1C, various temperatures 121
4.56 A plot of the electrode capacities against charge throughput of pouch control cells discharged at 1C, various temperatures 122
4.57 A plot of the DMA results from pouch control cells discharged at 1C, various temperatures 123
4.58 A diagram to show the mechanisms through which various degradation modes occur. Primary mechanisms are in green, with secondary mechanisms in maroon. Reproduced with permission from[8] 124
4.59 A diagram to show how different degradation mechanisms interact with each other, as well as their contribution to degradation modes and observable repercussions. Reproduced with permission from[8] 126
4.60 A plot of the DMA results against charge throughput of round 2 cells at temperature points with both 1C and 4C discharges 129

5.1 Diagram of the internal structure of a PEMFC from Ji *et al*[9] 139
5.2 Data as presented by Silva *et al*[4] 151
5.3 EIS data presented by Silva *et al* in[4] 152
5.4 Graphical representations of the expected voltage and delivered power response of a PEMFC while using current perturbations. Times are not to scale, but have been expanded for clarity of communication. 162
5.5 Fig 2. from Gupta *et al* showing the voltage response of a 2 cell stack to short circuits of various lengths. Reproduced with permisison. 165
5.6 A diagram of the PEMFC balance of plant designed and built for this project . . 168

6.1 A plot of the voltage response of the stack from constant current equilibrium to a perturbation of 100ms. Operating conditions: 50A (estimated to be approximately 0.5 Acm^{-1}), P_{anode} = 0.8 bar, $P_{cathode}$ = 0.7 bar, $T_{coolant} \approx 60°C$. 182
6.2 A plot of the voltage response of the stack during a perturbation of 100ms . . . 184

6.3 A plot demonstrating the fit of the proposed model to the voltage response during a perturbation for values of n=2,3 186
6.4 A plot of the voltage response of the stack immediately after a perturbation of 100ms 188
6.5 A plot of the voltage relaxation of the stack after a perturbation 189
6.6 A plot showing the system optimising perturbation duration 194
6.7 A plot showing the system optimising both τ and d, using the adaptive τ method and P&O respectively 197
6.8 A plot showing the values of τ calculated by the adaptive τ method at 50A load 198
6.9 A plot showing the relaxation of the stack voltage following an optimisation using the adaptive τ method 199
6.10 A diagram to show the impact of various operating conditions on PEMFC degradation, including the relationship between mechanisms and modes as well as the observable consequences. 204

Nomenclature

d	Duration (of a current perturbation)
Q	Electrical charge
V	Potential difference/voltage
Z	Electrical Impedance

Greek Symbols

ϕ	Golden Ratio
τ	Period (of a current perturbation)

List of Abbreviations

ADC	Analogue to Digital Converter
BEV	Battery Electric Vehicle
BOL	Beginning of Life
BOP	Balance of Plant
CC	Constant Current
CFRP	Carbon Fibre Reinforced Polymer
CNC	Computer Numerical Control
CV	Constant Voltage
CVM	Cell Voltage Monitor
DMA	Degradation Mode Analysis
DOD	Depth of Discharge
DOH	Degree of Hybridisation
DUT	Device Under Test

ECM Equivalent Circuit Model

EIS Electrochemical Impedance Spectroscopy

EOL End of Life

ESS Energy Storage System

GDL Gas Diffusion Layer

GFRP Glass Fibre Reinforced Polymer

GIC Graphite Intercalation Compound

GITT Galvanostatic Intermittent Titration Technique

HFCEV Hybrid Fuel Cell Electric Vehicle

HOR Hydrogen Oxygen Reaction

HPCC High Precision Coulomb Counting

ICA Incremental Capacity Analysis

ICE Internal Combustion Engine

IPU Current Perturbation Unit

LAM Loss of Active Material

LIB Lithium Ion Battery

LLI Loss of Lithium Inventory

LTO Lithium Titanate (i.e. Lithium Titanium Oxide)

MCU Microcontroller Unit

MEA Membrane Electrode Assembly

MGU-H Motor Generator Unit - Heat

NE Negative Electrode

NiMH Nickel Metal Hydride

OCV Open Circuit Voltage

ORR Oxygen Reduction Reaction

P&O Perturb and Observe

PEMFC Proton Exchange Membrane Fuel Cell/Polymer Electrolyte Membrane Fuel Cell

PEM Proton Exchange Membrane

PE Positive Electrode

PID Proportional Integral Derivative

RTC Real Time Clock

SEI Solid Electrolyte Interface

SLPM Standard Litres per Minute

SOC State of Charge

SOFC Solid Oxide Fuel Cell

SPI Serial Peripheral Interface

Chapter 1

Introduction

1.1 Project motivation

Humanity faces an increasingly dire climate emergency, with global temperatures rising rapidly due to the emission of greenhouse gases. In order to combat this emergency, alternative technologies that can eliminate the sources of these emission are required. While this challenge extends to every sector of technology, transport and power generation/heating are key, as they contribute the most to harmful emissions.[10] Different data sources provide a range of values regarding the exact contributions of these sectors, but they are consistently the highest contributors, regardless of data source. Electrification of these sectors is a key step in reducing their emissions, but will only do so provided the sources of electrical power are also de-carbonised.

1.2 Project objectives and research questions

In order to contribute to the body of work required to achieve the de-carbonisation goals, this project is focussed on the use of electrochemical devices for discrete power systems (i.e. those that might be found in road vehicles, as compared with those connected to an electrical grid for generation or balancing). In particular, the impact of phenomena that occur on the scale of the active materials used in these devices is investigated, with a focus on those that have significant impact at the system level. These phenomena can have significant impacts on the performance

of the system, either by mitigating degradation, and therefore extending system lifetime, or by utilising them to improve system performance, allowing smaller, more performant systems to be used for a given application. These both have the potential to allow for significant reductions in emissions, either by reducing the number of systems that need to be produced, reducing the emissions produced in the production of a given system, or by improving the efficiency of the system itself (and therefore any emissions associated with its operation.

The specific devices under investigation are Lithium Ion Batteries (LIBs) and Proton Exchange Membrane Fuel Cells (PEMFCs), as these are among the most promising candidates for electrification of discrete power systems, particularly in the automotive sector. In LIBs, the key phenomena investigated is the impact of thermal gradients within the cells. In order to better understand the impact of these gradients, as well as provide an experimental method through which models can be more accurately parametrised, the focus of this section of work was the development of a scientific fixture that allowed the impacts of temperature and C-rate (the rate at which a device is charged/discharged, divided by its nominal capacity to allow easy comparison between devices of different scales) to be decoupled. This is not possible with typical commercially available cells, as they are formed of multiple layers of active material. These layers are thermally non-isotropic, exhibiting significantly higher thermal resistances perpendicular to the layers (i.e. through multiple layers). Combined with the internal heat generation of the devices during operation, which is dependent on the C-rate, this produces thermal gradients within the device. By commissioning the creation of custom cells with identical chemistry and materials as commercial cells (even being from the same material production batch as the control samples used), but with only a single layer of active material, the impacts of temperature and C-rate on the degradation of the device can be de-coupled. An initial study utilising this technique was performed, and is presented here.

In the field of PEMFC's, the phenomena under investigation is the use of current perturbations (or periodic "short circuiting", as it is sometimes referred to) to increase the specific power and efficiency of a given system. While this technique is fairly well known anecdotally, the available literature on it is incredibly sparse, and the behaviour and mechanism through which

the technique acts is poorly understood at best. The focus of this project in this field was therefore selected to be a better understanding of the mechanism through which the technique acts, as well as improving the method of applying the technique. Specifically, a demonstration of the technique on a modern automotive PEMFC, with *in operando* optimisation, with a high level of instrumentation. This allows the evaluation of the following research questions:

- Can current perturbation yield beneficial result when applied to modern automotive style PEMFCs?
- Is it possible to dynamically optimise current perturbations for PEMFC applications, and what are the quantified benefits of the technique for modern automotive style PEMFCs?
- Does the transient behaviour of a PEMFC under current perturbation lend credence to a particular theory as to the mechanism through which the technique gives improved performance, and can this then be used to find an analytical solution for the optimal time between perturbations?

1.3 Thesis structure

This thesis is structured as follows: First a broad literature review to evaluate the state of the field to which this work is intended to contribute. This will focus on the identification of the above research topics and questions from the rather broad initial project brief, as well as a brief overview of some of the specific scientific techniques used. From there, there are 4 chapters which represent the primary output of this project.

- Chapter 3: Decoupling the Impact of Discharge Rate and Temperature on the Degradation of Li-ion Cells - Experimental design and validation
  - This chapter outlines the specific theory of operation of LIBs, as well as the methods used to evaluate the research questions identified. This includes the design and validation of the custom fixture, as well as the definition of the experiments to be performed.

- Chapter 4: Decoupling the Impact of Discharge Rate and Temperature on the Degradation of Li-ion Cells - Results and discussion
  - This chapter presents the results of work carried out according to the methods described in the previous chapter, including analysis using a variety of techniques. While these results do have some direct practical use, they are also intended as a demonstration of how the methods proposed can be applied, as the data presented here is chemistry specific.
- Chapter 5: Current Perturbations for PEMFC systems - Experimental design and validation
  - This chapter outlines the specific theory of operation of PEMFCs, as well as the implementation of current perturbations. The creation of a custom, highly instrumented PEMFC balance of plant is described, including the development of electronic measurement apparatus to allow the investigation of the device behaviour during a current perturbation, data which is not previously available in any published literature. In addition, a theory for the optimisation of this technique is presented based on the available mechanistic theories present in the literature, including a suite of proposed optimisation strategies to be evaluated experimentally.
- Chapter 6: Current Perturbations for PEMFC systems - Results and discussion
  - This chapter evaluates the results gathered using the scientific apparatus described previously, including an evaluation of the quantifiable benefit of the various optimisation strategies proposed. In addition, the electrical behaviour of the PEMFC during a perturbation is presented, which gives significant insight into the mechanism through which the technique of current perturbations acts. Finally, a proposal for more mature implementations of this technique is presented, as the methods here are intended for scientific inquiry, as opposed to practical engineering application.

Chapter 2

Literature Review

2.1 Introduction

Energy systems at all scales are currently undergoing a renaissance due to the ongoing climate crisis. Electrification of transport, decarbonisation of the grid, decarbonisation of heating systems, these all require the development of their relevant technologies. Lithium Ion batteries have been a key technology in this process, including the electrification of vehicles. However, they are not a one-size-fits all solution, and the best results can be achieved using a mix of technologies.[11–14] Hybrid Fuel Cell electric vehicles (HFCEV's) solve many of the problems currently plaguing the Battery electric vehicle (BEV) market, such as range anxiety, refill/recharge times, and infrastructure concerns. By combining two complimentary technologies, the advantages of both can be utilised, while the disadvantages are reduced in impact. This does come at the cost of increased system complexity, however.[15] In order to aid design of systems better designed for their intended use case, the impact of phenomena that occur on the micro- and nano-scales within these electrochemical devices on the behaviour of the device and system level must be understood.

The primary difference between the two devices from a system perspective is that a FC allows the energy storage and the power production to be independently designed, giving a much greater range of freedom in the design space. LIBs, conversely, both store and convert the energy within the same device, limiting the ratio between energy stored and power capability as

well as the speed at which a system can be refuelled. LIBs have enjoyed a surge in research focus in the last 20 years, partly because their scale allows R&D for smaller products (such as handheld electronic devices), the economic and technological impetus of which can then be used to scale upwards, to applications such as vehicle propulsion and grid storage. Fuel cells, on the other hand, require a balance of plant to operate, and so have a minimum scale at which they are viable solutions. In addition, they significantly increase system complexity, as there is a separate HV system, the associated plumbing and flow control required for the fuel cell system (typically individual fuel, oxidiser and coolant flows) as well as the management and control system that controls the flow of power between fuel cell, energy storage system and traction system. This means there is a higher initial barrier to entry.[16] In addition, there are infrastructure and geopolitical drivers, such as the generation, storage, and distribution of Hydrogen as a fuel. Analysing the development of this infrastructure is beyond the scope of this work, which focuses on the device and vehicular system level.

The literature review presented here is intended only to narrow the wide scope of this projects initial brief into something both achievable with the available resources, but which maximises the positive impact. This is primarily done based on the engineering merits of the relative technologies, such as their theoretical performance, their technological maturity and the identification of key areas for novel research. The influence of economic or geo-political factors (such as the development of hydrogen fuelling infrastructure or the sourcing of minerals critical for the manufacturing of LiBs) are not strongly considered, as the focus of this work is on the engineering of an automotive system, not the production and sale of consumer automotive products. It has been intentionally kept as brief as possible, so that the bulk of this report might focus on the outcomes of the experimental work done. It is of course impossible to capture the wealth of knowledge available on even a single technology in such a document, never mind the multiple bridged by this topic (Batteries, fuel cells, system design and integration, vehicle design and electrical power systems and control, to name a few). This chapter is instead intended as to set the technological scene in which this work was performed, as well as serve as an initial primer for readers, allowing the presentation of the projects output to be streamlined somewhat.

The details of the operation of these devices will be examined in their respective chapters, alongside the challenges in probing the research questions identified.

2.2 Fuel Cells

There are a wide range of technologies which fall under the description of "Fuel Cell", comprising of different charge carriers, electrolytes, fuels and physical designs. There are even closely related devices (such as flow batteries) which share many of the technological developments and implementations. For the purposes of this project, it was necessary to narrow the scope to the technologies most applicable to the automotive sector. This is determined not only by device characteristics and performance, but by more qualitative factors such as ease of packaging, safety and supply infrastructure (of materials for both constructing and fuelling the devices).

2.2.1 Types of Fuel Cell

A shortlist of the fuel cell technologies that are sufficiently mature and applicable to the automotive sector is as follows:[17, 18]

- Polymer Electrolyte/Proton Exchange Membrane Fuel Cells (PEMFC's)
  - Traditional
  - Aqueous (e.g. Direct alcohol solutions)
- Solid Oxide Fuel Cells (SOFC's)
- Alkaline fuel cell
- Phosphoric acid fuel cell

We can reduce the field of competing technologies by examining each devices requirements. This will be done discompassionately, for while many of the problems limiting the application of these technologies can be overcome, doing so inevitably comes at a cost, either in the form of

limitations of the system produced (such as increased complexity, weight or size), or in the form of significant engineering effort to produce more exotic materials or complex designs. Phosphoric acid FCs utilise a liquid electrolyte of either pure or highly concentrated phosphoric acid. While this produces dangers associated with the manufacture, maintenance and use of the devices, these can be overcome relatively easily, and are not unique to this device. For example, lead acid batteries utilise concentrated sulphuric acid, but are present in the vast majority of internal combustion engine (ICE) based vehicles ever produced. Instead, the issue is with the physical properties of the electrolyte, which is held in a ceramic matrix, typically SiC. It freezes at 42°C, which means it would undergo frequent freeze-thaw cycles in a typical automotive application. This would mechanically stress the components and/or require significant complexity in the design and implementation.[19, 20]

Alkaline fuel cells also have a very corrosive electrolyte, albeit in the opposite form (hence the name), usually KOH. Their main limitation is their sensitivity to contaminants, however, particularly CO_2, requiring pure H_2 and O_2 reactants. This makes them both too fragile for realistic use, as well as the dangers and difficulties associated with refilling and on-board storage of both fuel and oxidiser. They also suffer from low power densities.[19, 21]

Aqueous PEMFC's are of interest as the fuels can be readily synthesised from a range of low carbon sources and can be stored in their pure forms for injection into a re-circulating aqueous flow, increasing the energy density of the system without need for high pressure vessels or plumbing. Unfortunately, they exhibit sluggish reaction kinetics (limiting power densities significantly), as well as accelerated degradation from potential cycling, as experienced in start-up/shut-down cycles.[19, 22] There is work being done to overcome these limitations, primarily in the form of organic catalysts, but this is not sufficiently mature for application.

This leaves PEMFC's and SOFC's. SOFC's have many benefits, such as their flexibility in fuels and resistance to contaminants.[18] Unfortunately, due to their high temperatures (typically 600-1000°C[17]) and solid ceramic electrolyte they tend to be bulky, exhibiting low

power densities (0.1-1 kWl^{-1}[23, 24]). While this has been significantly reduced in the more mature designs, such as Ceres Power SteelCell,[25, 26] the high thermal cycles experienced in an automotive application still causes accelerated degradation. In addition, they struggle to achieve their potential efficiencies on the scale required for automotive.[23] This combination of characteristics makes them of great interest for larger applications, such as de-carbonisation of shipping, where SOFC's can displace the cathedral ICE's without losing the excellent fuel flexibility the operators have come to benefit from.

PEMFC's, conversely, operate at much lower temperatures, reducing the degradation experienced by components due to thermal cycles, as well as requiring either significantly less thermal insulation than SOFC's, or none at all. This is beneficial for automotive applications due to the inherent bulk and difficulty packaging thermal insulation. PEMFC's can be designed to operate at 60-80°C, usually designated "low temperature" (although not typically specified, as this is the vastly more mature technology) or 110-180°C, referred to as "high temperature" PEM. While PEMFC's can be sensitive to contaminants (primarily CO and H_2S), the impact of this on device performance and lifetime has been drastically reduced thanks to developments in catalysts, both in the materials and the synthesis of the nanoparticles. Indeed, the PEMFC used for the experimental work in Section 5 is capable of running on reformate fuel, which typically contains significant concentrations of these species. They also require some care with respect to their thermal management, as well as humidity of reactant flows, the details for which are discussed further in Section 5.2, but this is significantly less difficult to manage than their SOFC counterparts. High temperature PEMFC's tend to be more resistant to poisoning by contaminants, as well as improved reaction kinetics yielding higher efficiency and specific power.[18] Unfortunately, the higher temperatures makes humidification management difficult, as well as increasing start-up time and accelerating degradation.[27, 28] For an automotive application, these last two issues are the primary difficulty, as consumers have become accustomed to (and therefore demand in future products) the extended lifetime and rapid start-up that ICE's have been extensively developed to provide.

Table 2.1: A table showing the volumetric and gravimetric energy densities of popular automotive energy sources

Fuel	Energy Density	
	Gravimetric (MJ/kg)	Volumetric (MJ/L)
Hydrogen (690 bar, 15°C)	5.323[29]	141.9[29]
Petrol	46.5[32]	34.4[32]
Diesel	38.6[30]	45.6[30]
Natural gas (250 bar)	9[32]	53.6[32]
Ethanol	24[30]	30[30]
Automotive battery (Tesla 85 kWh)	0.56[33]	15.8[33]

2.2.2 Fuels

One of the drivers for research into high temperature PEMFCs was their potential to be run on reformed fuels. With advances in catalysts and electrolyte polymers more resistant to poisoning, this is now feasible with low temperature PEMFC's, although still reduces lifetime significantly, limiting its implementation. The reason it is attractive, however, is to do with the difficulties storing the fuel for PEMFC systems. Gaseous hydrogen is incredibly gravimetrically energy dense, thanks to its low molecular mass, but is much poorer volumetrically. For this reason, it is typically stored in pressure vessels at 700bar for automotive use, even then only achieving 13% of the volumetric energy density of liquid petrol.[29, 30] Added to this is the complexity and cost of infrastructure required for distribution of this high pressure gas (average refuelling station upfront capital cost of $2.2 million[31]). A comparison of popular fuels energy densities can be seen in Table 2.1

There are a variety of potential solutions to this problem. One is storing hydrogen in a solid state, through either chemisorption or physisorption. Chemisorption technologies (such as hydrides) can have excellent volumetric performance, but typically lower gravimetric density. In addition, they have poor performance with respect to stability, cyclability and controlling the release of the gaseous fuel. Physisorption solutions boast better cyclability, but require cryogenic temperatures to reach acceptable weight percentages of fuel to system.[34] This would require a

cryogenic plant installed in a vehicle, which is not feasible for most applications due to the high cost, complexity and maintenance required. For this same reason use of liquid hydrogen in vehicles is also not a viable solution, as it requires temperatures $\leq$20 K to liquefy. This is useful in the infrastructure and transport of hydrogen as a fuel, as it can be integrated into larger storage systems or used for short distance transport (e.g. in tanker trucks, as might supply a fuel station).

Finally, onboard reformation of liquid fuel would allow the benefits that fuels volumetric density, while benefiting from a PEMFC's efficiency, clean output and low noise. Steam reformation of natural gas has been the dominant method of producing hydrogen for the majority of the 20th century, but greener hydrogen sources (such as electrolysers, bio-reactors or fermentation of organic waste) are available to decarbonise the supply chain. Light alcohols, such as ethanol and methanol are attractive options for on board reformation, as they boast excellent energy densities, are readily reformed, and are already produced and standardised as fuels. Ethanol is preferable, as it is more energy dense, and methanol is highly toxic. Performing this reformation on-board is possible, but adds system size, complexity and cost.

Hydrogen storage in vehicles is predominantly done as a compressed gas in Type 4 pressure vessels, which are formed from a composite material such as cabron fibre reinforced polymer (CFRP) or glass fibre reinforced polymer (GFRP) and lined with a high density polymer liner. A description of the type classifications of pressure vessels, as well as a comparison of their suitability for long term storage of automotive hydrogen can be seen in Table 2.2. Because they are required to withstand high pressure, and to maximise storage capacity while minimising structural weight, these tanks tend to be cylindrical. Due to the cost associated with making, testing and installing the tank (with all the appropriate safety equipment), the incentive is to have as few large tanks as possible. This unfortunately becomes difficult to package, especially compared to liquid fuels such as petrol or diesel, where storage can fill any irregular shaped void in the vehicle by use of blow moulded plastic tanks.

The benefit of gaseous hydrogen as a fuel is that it has a wide variety of green sources,

Table 2.2: A table showing the classification of pressure vessels, as well as some brief notes about their use in PEMFC EVs

Pressure vessel type	**Materials**	**Description**	**Notes regarding use for storing compressed gaseous hydrogen for automotive purposes**
1	Metal (typically steel)	All metal, easiest and cheapest to produce. Also heaviest	Steel embrittlement occurs due to contact with high pressure hydrogen, limiting lifetime. Additionally, high weight significantly reduces system gravimetric energy density. Typical for applications where weight is not critical, and where regular inspections are viable (e.g. on-site storage for industrial uses)
2	Metal (steel, aluminium), GFRP	Steel or aluminium vessel, with GFRP reinforcement in hoop direction. Metal and composite share hoop stresses, reducing weight but increasing cost	While somewhat lighter than Type 1, still heavy compared to later types. Metal embrittlement still problematic, as the inside of the vessel is still unprotected and required to contribute to structural integrity
3	Metal, GFRP/CFRP	Metal liner, with composite structure. Composite carries all structural stresses, with the liner providing sealing and form for manufacture	Significantly lighter than Types 1 & 2. Metal embrittlement less problematic, as liner is not structural, but can have some impact on contents purity
4	GFRP/CFRP, high density polymer	Composite structure with high density polymer liner to ensure contents purity	Significantly lighter than types 1 & 2. Liner is typically HDPE.

allowing progressive steps towards full decarbonisation. In addition, the use of a chemical fuel allows an operator to benefit from rapid refuel times (≈3 minutes for a fossil fuel passenger car,[35] 12 for long haul trucks[36] and 3.1 for HFCEV[31]) while also reaping the benefits of electrification, such as increased efficiency (electrochemical devices are not limited by Carnot efficiency as heat engines are), quiet operation and zero harmful emissions. While this does confer some benefits to general consumers, it is of most benefit for high utilisation vehicles. These are typically not personal vehicles, but instead public transport (i.e. buses, taxis) fleet vehicles (e.g. used by a company required to send technicians to customer sites), and as such vehicle downtime is critical.[37] In addition, some of these use cases are less hindered by limited refuelling infrastructure, as they would be refuelled at their fleet depot regardless of fuel type.

2.2.3 Hybridisation

FCs are typically hybridised in an automotive application because of their poor load following capabilities.[38, 39] In addition, a degree of hybridisation (DOH) can greatly increase system efficiency,[40] both by allowing energy capture techniques such as regenerative braking and by giving the system some freedom to optimise the operation of the FC to maximise its efficiency.[38] Hybridisation allows systems to be highly tailored to specific uses, often done by separating the conversion of chemical fuel to usable energy and the delivery of that energy to the components that use it. This comes at increased system complexity, and often some increase in size, cost and/or weight. This added complexity can also significantly increase maintenance required, as well as the skills and knowledge required by maintenance personnel. Fig 2.1 demonstrates some of the impacts of the degree of hybridisation on a vehicle's operation. ICE vehicles prior to electrification operated only at the first stage of this hybridisation, where the energy storage system (ESS) comprised solely of a flywheel attached to the engines output. In high performance applications, where the other characteristics of the engine had been improved to produce better load following, the flywheel mass would often be reduced, whereas in applications where efficiency was prioritised, heavier flywheels would limit the impact of aggressive throttle signals on system efficiency by acting as a low pass filter for the engines speed. However, with electrical hybridisation, much more can be achieved. The first major benefit is the ability to

capture waste energy. Most commonly, this is in the form of regenerative braking capturing kinetic energy, but it does take other forms, such as the MGU-H used in Formula 1 to capture waste energy from the engines exhaust. This typically has a drastic impact on system efficiency, as though the increase in DOH typically comes with some added weight and complexity, this is outweighed by the contribution of the added ESS. This is particularly the case because a well designed system can use the ESS to respond to loads that would typically force the primary energy converter (i.e. ICE or FC) to operate in its lowest efficiency regions. The most ubiquitous example is a vehicle pulling away from stationary, as this requires a pulse of high power to overcome the rolling resistance of the vehicle, applied for a short period at zero speed. For an ICE this typically means building the engine speed (storing energy in the flywheel) before engaging the clutch, which is a fuel intensive procedure, whereas for an un-hybridised FC system it would involve a step change high current demand for a short period. A FC's poor load following capabilities make this inefficient, but the larger problem is the need to increase the power capabilities of the FC to ensure that it can provide the peak power required, regardless of the duration. Conversely, in a hybridised system, the sum of the energy converter and ESS outputs must meet this peak power, allowing each to be smaller than in a non-hybridised system for the same application. As the DOH increases, this tends towards the primary energy converter being sized to meet the expected mean power of the system, with the ESS being sized such that it can provide the difference between this mean power and the expected load profile (both in terms of power and energy stored).

As an example, a small vehicle designed for personal use will cover short distances the vast majority of the time, but will still be expected to achieve much longer distances if required.[41–44] A fully BEV in this scenario would have its efficiency decreased by having to carry a battery large enough to assuage the consumers range anxiety during its normal short range operation. However, a small DOH, such as the use of a chemical range extender, could significantly reduce the system weight by allowing the battery to be sized for general use, and the range extender sized such that it can provide sufficient power and energy to allow for the edge use cases. In this case the ratio of peak power to average power is relatively

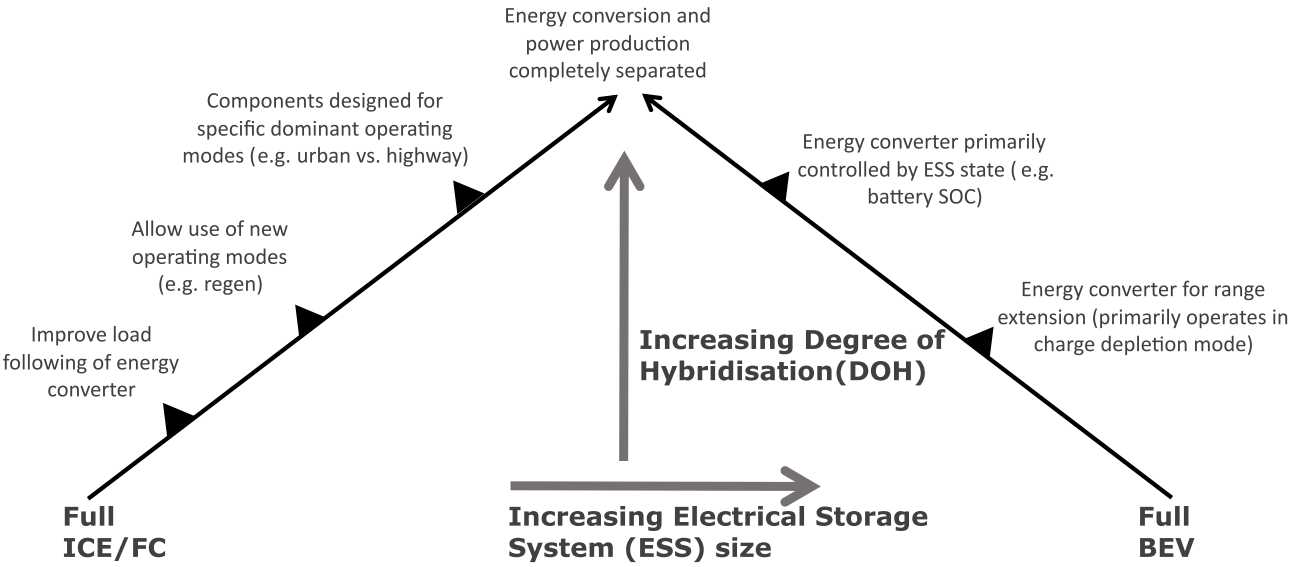


Figure 2.1: A diagram detailing the impact of degree of hybridisation on the performance of a vehicle

low, as a personal vehicle is not particularly heavy and can have tyres optimised for rolling resistance as opposed to traction. Conversely, the case of an inner city public transit bus has a surprisingly low average power, but a very high peak power, as it is required to constantly accelerate (and decelerate) a large mass, with tyres designed for safety and distribution of load more than for the reduction of rolling resistance. As such, a bus would require a more powerful but less energy dense ESS, (such as supercapacitors or Lithium Titanate (LTO) cells), with its primary energy converter sized to maintain the required state of charge (SOC) of the ESS.

To further detail the concepts expressed in Fig 2.1, the first step of hybridisation from a purely chemical energy source allows better load following of the energy converter, or for the load following to be appropriately filtered to maximise some other parameter, such as thermal management, system efficiency or operation in extreme conditions (e.g. fuel warmers for diesel truck in cold climates). The next significant threshold with increasing DOH is allowing the use of additional operating modes, such as the use of regenerative braking, or the shutdown of the energy converter in specific circumstances (e.g. high ESS SOC in stop-start traffic). Next, the sizing is components or subsystems for specific operating condition. An example of this is a "Plug-in hybrid" vehicle which operates solely of the ESS for urban driving, with the energy converter sized to facilitate longer distance highway driving. Finally, the extreme case of

hybridisation is where the primary energy converter is sized in order to meet the average power demand and the ESS is sized solely to meet the required load following. This is difficult to achieve, as it requires a very predictable and precisely quantified load profile to design accurately. As such, the systems typically fall one step either side of this fully hybridised system due to the unpredictable nature of real world applications.[13, 14, 45] Continuing onwards towards a full BEV, with DOH now reducing and the ESS being the primary component, the next step is systems where the ESS can perform the vast majority of short term load profiles themselves, with the energy converter focused on maintaining desired SOC and occasionally being required to contribute to meet peak power demands. The most common example of this would be the use of a Toyota Prius for inner city rideshare work, where the battery can provide sufficient power for the vast majority of routes, with the occasional need for the engine to provide tractive power on main roads (i.e. higher speed) or aggressive acceleration, but primarily operates to maintain the lower capacity batteries SOC. The final step before a fully BEV with no hybridisation is a BEV with a small range extender not capable of meeting the systems mean power demand. In this case the energy converter is used to extend the operating window, but cannot do so indefinitely, sometimes described as "charge depletion mode". The ESS SOC will drop over the expected load profile, but the ESS can be significantly lower capacity than a non-hybridised system for the same application due to the energy contribution from the range extender.

2.2.4 Current Perturbation

There is very little available literature on the subject of current perturbation, despite it being a well known technique by professionals in the field. Silva *et al*[4] investigate the impact of an extended short circuit on the operation and longevity of a PEMFC, but with the intention of designing a passive hybrid system with Ultracapacitors, as opposed to a control mechanism for the PEMFC itself. Because of this, they are investigating the impact of events such as acceleration of a vehicle powered by a passive PEMFC/UC powertrain, and so apply short circuits on the order of 10s, as opposed to the 10s-100s of ms typically used for improving stack performance. Similarly Cosse *et al*[46] attempt to model the behaviour of a PEMFC

during and immediately after a short circuit, but from the perspective of understanding its transient behaviour in order to aid system integration, for example ensuring that safety devices are triggered in the event of a failure, or sizing and impedance matching the input filters of DC-DC converters used to deliver power from a PEMFC to a system.

There are also a small selection of published works on the use of short circuits to aid humidification of PEMFC devices, typically for rapid start-up sequences. Kim *et al*[47] proposes its use for a small, open cathode stack powering unmanned aerial vehicles and Flach *et al*[48]propose a circuit allowing individual cells of a stack to be short circuited for this purpose. The theory that short circuiting has a significant impact on the humidification of a PEMFC is not believed to be a credible mechanism for its action, as will be discussed in Section 5.2.4.

There are a small number of papers investigating the use and optimisation of the technique for increasing the output power/efficiency. These are all on older, milled graphite stacks, with dated membrane electrode assemblies (MEA's), and in particular operating at lower current densities than those used in modern systems. Zhan *et al*[49] present a possible implementation of a short circuit unit in the context of a UPS system, although the data regarding the specific behaviour and benefit of the technique is poorly presented as screenshots of a LabView control panel, significantly limiting the resolution and therefore any insight that might be gleaned from it. Gupta *et al*[50] is the most well rounded work, presenting a systematic investigation of the two main parameters in a current perturbation (perturbation duration and period). This demonstrated that the increase in mean voltage (at constant current density) was convex with respect to both of these parameters. This work was repeated and built upon by Zhang *et al*[51] using a Horizon H-100 stack (graphite bipolar plate, dated MEA, open cathode, low maximum current density). This included some variation of the "short" impedance.

Finally, Trogadas *et al*[52] investigated the impact of repeated short circuiting on the longevity of a PEMFC. They used MEAs manufactured in house, which makes it difficult to determine their comparability to modern systems. In addition, they do not acknowledge some significant flaws in their methodology. For example, in order to accelerate the degradation of the devices,

they are operated at no load, open circuit voltage (OCV) conditions. This means that due to the application of the short circuit, one device is experiencing frequent voltage cycles across regions that are known to accelerate degradation mechanisms such as gas diffusion layer (GDL) corrosion and Pt nanoparticle agglomeration.[28] These voltage cycles would be significantly smaller in a realistic use case, where the stack is loaded while the short circuit protocol is applied. They conclude that the use of the technique significantly reduces device longevity by a factor of 2.5x, despite only performing testing for a total of 144h. While it is felt that the short circuit protocol applied in this publication is likely to impact the useable life of the device, this presentation is likely misleading, and significantly overestimates the impact. In addition, more informed implementations of the technique are expected to significantly reduce this impact, as will be examined in the discussion of this document.

2.3 Lithium Ion Batteries

Lithium Ion batteries have experienced a huge focus of attention in both the media and research over the last 30 years. This is primarily because of their wide range of applications, especially thanks to their excellent range of scales at which they can be applied. They are used in micro-cells, such as the coin cells used to power RTC's and ensure safe shutdown of computing components, all the way up to large grid balancing applications, with installations reaching 100's of MW. This range of scales also contributed to their accelerated development, as they had commercially viable applications beginning at small scales (such as use on handheld and portable electronics), which drove funding and research which allowed them to become viable solutions for the larger scale applications, while maintaining a lower barrier to entry than many of their larger scale alternatives. As a comparison with fuel cell technology (which will be discussed in Section 2.2), which requires a balance of plant, sealed pressurised fluid systems and often exchange of reactants with the surroundings (in the case of air breathing PEMFC's), Li-Ion cells are sealed units which can be manufactured at a range of scales with minimal impact on their performance characteristics. While it is far from impossible to power handheld electronics from a FC, it requires understanding of the system and careful integration. Conversely, a system

designer can easily use a Li-Ion cell without understanding it's internals significantly, requiring only that they be kept within certain electrical and thermal windows, an easy accommodation given that these are very similar to the requirements for many extant IC's, capacitors and other electrical components.

The primary difference between Li-Ion cells and other mature battery technologies is their use of intercalation reactions. Most battery chemistries rely on a redox reaction of some kind, usually stripping or depositing an oxide or metallic charge carrier. For example, Nickel Metal Hydride (NiMH) batteries convert their negative electrode (NE) alloy from its metallic form to its hydride form, while converting Nickel hydroxide to Nickel oxyhydroxide, and Lead-acid cells convert Lead dioxide positive electrode (PE) and metallic Lead (NE) to Lead(II) sulphate. In contrast, Li-Ion cells intercalate their charge carriers into the electrodes. This allows them to have higher specific energy and power densities, but requires them to be kept within acceptable voltage limits as the intercalation compound structures become unstable outside this window. One of the reasons for their higher energy and power densities is the increased cell voltage compared to other chemistries, with nominal voltages of 3.6V, compared to Nickel based cells of 1.2V and Lead-acid of 2.2V. This intercalation based chemistry has been expanded to other charge carriers and electrode materials, most notably Sodium chemistries, popular primarily because of Sodium's relative abundance compared with the rarer Lithium.

The passenger HFCEV's that have already been brought to market tend to fall to the far left of the diagram in Fig 2.1.[53] This means that their batteries are significantly smaller than their PEMFC. For example, the Toyota Mirai has a battery capacity of 1.6 kWh, but a FC power of 114 kW.[54] This is because the PEMFC is the primary power source, with the battery present only for load following and regenerative braking. Because of this, the battery must be significantly more rate capable than in a full BEV (due to its significantly smaller size), although the duration of the high rate event it is expected to handle is significantly shorter. I.e. the most extreme event it is likely to handle is hard (but still fully regenerative) braking from a high speed, which occurs over at most 10s of seconds, compared

to a full BEV accelerating and maintaining highway speeds for hours at a time. The other successful test cases for HFCEV's, buses and trucks, similarly have high peak power/mean power ratios in their expected drive cycles, requiring the components tasked with load following to provide high power for short periods. It is for this reason that Ultra-capacitors are often used instead, thanks to their high rate capability.[55] In addition, they have excellent cycle life, even under aggressive load profiles, and this is the key factor in selecting and sizing the ESS.[56, 57]

Because of these (and many more reasons that are beyond the scope of this work and have been excluded for brevity), the thermal management of the ESS in a HFCEV is key to ensuring its longevity,[8, 58–60] even more so than in a full BEV. There are a variety of thermal management solutions available to OEMs, but thermal gradients are an inescapable reality. That being said, Hunt *et al*[61] found that at high discharge rates, the direction of thermal gradients had a significant impact on the loss of useable capacity, with cells cooled via their tabs (and therefore having thermal gradients in the plane of the active material) showing a 1.2% loss over 1000 cycles, compared to the same pouch cells cooled via their surfaces losing 9.2%. Understanding the impact of these heterogeneities is an area of significant work.[62]

The primary reason for the findings in Hunt *et al* is the highly non-isotropic thermal conductivity of LIBs. Due to their layered structure, conductivity in the plane of the electrodes is significantly higher than in the perpendicular direction. This is mainly because of the thermal conductivity of the materials forming the electrodes, with the current collectors (aluminium and copper, with thermal conductivities of 398 $Wm^{-1}K^{-1}$ and 238 $Wm^{-1}K^{-1}$ respectively) being significantly more conductive than the active material (1.58 $Wm^{-1}K^{-1}$ for NMC and 1.04 $Wm^{-1}K^{-1}$ for Graphite Intercalation Compound - GIC).[63] In addition, heat can conduct through a single material in the plane of the electrodes, but must transfer between materials to propagate through the planes, with the contact resistance between the layers adding to the thermal resistance in this direction. The bounds of the values for in- and through-plane conductivities of different cell chemistries can be seen in Table 2.3. Because of this variation in properties in the through and in plane directions, the behaviour, magnitude and impact of these

Table 2.3: A table listing the in- and through-plane thermal conductivity bounds for a selection of LIB chemistries, adapted from Wei *et al*[1]

Chemistry	**In plane ($Wm^{-1}K^{-1}$)**	**Through plane ($Wm^{-1}K^{-1}$)**
LiFePO4 battery	10.09-45.06	0.66-1.04
LiCoO2 battery	10.3-45.65	0.75-1.47
LiMn2O4 battery	10.09-45.06	0.66-1.04
Li (NiCoMn)O2 battery	10.37-45.83	0.77-1.55

thermal gradients is strongly impacted by the design of both the cell and pack. The cells format is key, as this determines not only the orientation of the layers, but also the thermal pathways out of the cell through which heat is rejected. For example, a pouch cell with small tabs at one side is going to be less capable of rejecting heat through those tabs, and will therefore have to rely on surface cooling to reject enough heat to exhibit significant rate capability. It is therefore likely to experience significantly higher thermal gradients in the through plane direction than a cell with identical active material, but with large, flat tabs on opposing sides that have good thermal pathways to the current collectors within the cell. This latter case could potentially reject enough heat through the tabs under an aggressive load cycle, which would produce thermal gradients in the plane of the layers of active material, but have far less heterogeneity in the through plane direction. Thanks to the significantly higher thermal conductivity in the in plane direction, the thermal gradients induced will be lower in magnitude, and have less impact on the cells longevity according to the findings of Hunt *et al*.[61] Similarly, the physical properties of the materials used (which will differ across chemistries, manufacturers and cell designs) will have an impact, as well as the thickness of the layers they comprise. A more energy dense cell can be achieved by using thicker coatings of active material on the same thickness of current collector, therefore making a larger portion of the cells thickness and mass capable of storing charge. Conversely, a more rate capable cell can be achieved by using thinner slurries with smaller particle sizes, therefore increasing the fraction of the cell that the current collectors form, as well as the surface area/volume ration of the active material. Finally, the casing structure and material will play a role.

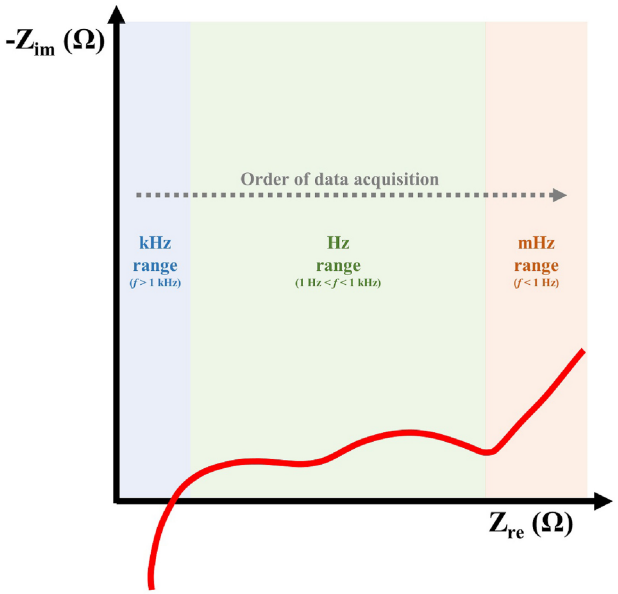
Zhao *et al*[64] produced a thermally coupled volume discretised model of a cell in order to

explain the findings of Hunt *et al*.[61] This required painstaking parameterisation and enumeration of the thermal pathways and their relative conductivities, and was able to demonstrate that (for the cell modelled, a 5 Ah Kokam pouch cell), while tab cooling was able to remove less heat for a given ΔT, and therefore the cell was at a higher mean temperature, the thermal gradients produced were significantly smaller. Unfortunately, in order to couple this model with a longer term degradation model with any accuracy, the impact of temperature and discharge rate would need to be decoupled. Their interdependence can be calibrated out when parameterising entire cells for modelling, but in order to produce a volume discretised model such as that presented by Zhao *et al* that also models the irreversible degradation of the device, the behaviour must be captured without the complex interplay of parameters that occurs at a cell level. This work, along with follow up work by Hales *et al* and Marzook *et al* defining the Cell Cooling Coefficient, does demonstrate the importance of thermal pathways in cell design. LIB's intended to be rate capable can only be used as such if their internal structures can be sufficiently cooled, which requires good thermal contact between both the cooling surface of the cell and it's thermal management system, as well as between the heat generating components inside the cell and whatever surface is used to cool it. It is for this reason that tab cooling is useful for high performance LIB's, as the need for low resistance current paths out of the cell can also be used to produce low resistance thermal pathways, given that thermal and electrical conductivity is often correlated in the materials available for use.

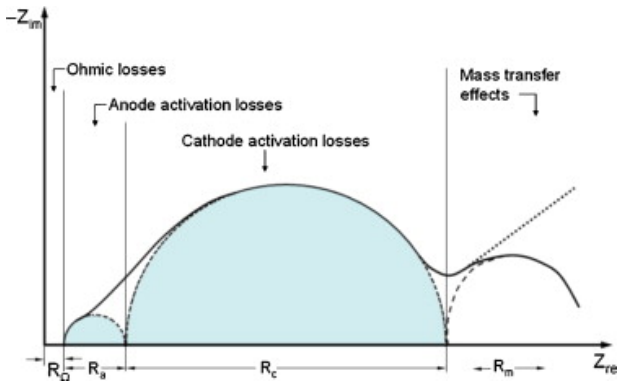
2.4 Scientific Techniques and Procedures

2.4.1 EIS

Electrochemical Impedance Spectroscopy (EIS) is a useful technique for the analysis of various electrochemical systems, including both batteries and fuel cells. It is performed by attaching the device under test to an instrument via 4 conductors. One pair carries the excitation signal, and the other measures the response. This excitation signal is typically a sinusoid, but can be in either the current or voltage domain (galvanostatic and potentiostatic respectively). The magnitude of this signal is fixed (hence the terms galvano*static* and potentio*static*), and the



(a) Hypothetical example of a typical LIB EIS measurement, presented on a Nyquist plot. Included with permission from Meddings *et al*[65]



(b) Hypothetical example of a typical PEMFC EIS measurement, presented on a Nyquist plot. Included with permission from Zhu *et al*[66]

Figure 2.2: Typical EIS results for LIBs and PEMFCs

frequency is swept across the desired range of test, typically from the tens of mHz into the kHz range. The measured impedance can then be plotted on a Nyquist plot, with a few key features. First, the point where the results cross the real axis represents the Ohmic (or instantaneous) resistance of the DUT. Next one or more arcs which represent the different time scales of processes occurring in the system, such as charge transfer resistances, mass transport losses etc. In the case of LIBs being the device under test (DUT), there is typically a long, low-frequency tail which represents the diffusion of the lithium ions in the solid phase electrode materials. Examples of the typical shapes of these results for both LIBs and PEMFCs can be seen in Fig 2.2.

The connection between the DUT and the instrument used to perform EIS is a key factor in ensuring the validity of the data. As has been mentioned, this requires separate conductor pairs for excitation and measurements, typically termed "4-wire impedance measurement" to ensure that the response is measured at the device terminals and does not include the resistance of the connection. These pairs are typically twisted in their low-current (i.e. measurement) and high-current (i.e. excitation) pairs for added rejection of noise from external EM sources, and must be properly shielded for the same reason. While this configuration does exclude the

contribution of the real impedance of the connection in the results, it cannot reject the presence of the imaginary component, or irregularities in the metrology equipment itself. Instead, this should be done with proper fixture correction/compensation, which involves performing the test on a series of calibration devices before connecting the sample to be evaluated. These calibration devices are typically:

- Open circuit: most commonly achieved using an air gap, as the voltages used in EIS are typically on the order of single digit volts. Z_{OC} in Eqn 2.1
- Short circuit: most commonly achieved using the largest piece of a high conductivity metal (aluminium, copper or brass) that will easily fit in the test fixture. Z_{SC} in Eqn 2.1
- Calibrated load: A component or component network designed to have an accurate and known impedance, typically measured and quoted by the manufacturer. Ideally this is in the same region as the samples to be tested for most accurate results.[67] Its nominal impedance is Z_{ref} in Eqn 2.1, and the measured value is Z_{Load} in Eqn 2.1

The impedance spectroscopy is performed on these three configurations separately, and then applied to the measurements taken from the samples using the relationship in Eqn 2.1.[68] $Z_{Measured}$ are the raw readings taken from the sample and Z_{DUT} is the impedance of the sample computed using this compensation technique. Data captured from EIS can be analysed in a variety of ways, although the most common is the fitting of Equivalent Circuit Models (ECMs).

$$Z_{DUT} = \frac{Z_{ref}(Z_{Measured} - Z_{SC})(Z_{OC} - Z_{Load})}{(Z_{OC} - Z_{Measured})(Z_{Load}) - Z_{SC}} \tag{2.1}$$

The magnitude of the excitation signal should be selected to be as large as possible while keeping the DUT in a region small enough to be consistent. For example, in LIB it should be chosen such that the SOC does not vary enough to produce significant changes in the impedance being measured, and also such that any heat generation from the current passed is minimal enough to have negligible effect on any measurements taken. This is difficult to define unilaterally, as it depends on any number of parameters such as cell form factor, chemistry, material loading of the electrodes, thermal management applied etc. In practice, it requires

experience and patience to yield the best results, especially as a single sweep can take many hours to perform, depending on how low a frequency of excitation is applied. While this technique can yield significant insight, it is very sensitive to errors or inconsistencies in the procedure, and correct fixture correction is often not applied in published data.

2.4.2 GITT

Galvanostatic Intermittent Titration Technique (GITT) is a technique for the characterisation of electrochemical materials and devices. Commonly used to parameterise models for devices such as LIBs, it is performed by applying a series of pulse loads, each followed by a period of no load, or a "rest". These pulse loads produce a measured change in the SOC of the device, and the voltage response of a cell during each pulse and rest can have curves fitted to them in order to produce the parameters of the desired model. It is especially useful because it can capture the desired parameters across the entire SOC range. For example, it has been used to calculate the diffusion coefficient for lithium across the entire SOC range by Li *et al*,[69] and is useful for comparing electrode materials[70–74] as they will exhibit different behaviour across their SOC range, making results only measured at a single SOC potentially unreliable or misleading. The major disadvantage of GITT is that it takes a long time, as each pulse must be followed by a rest long enough for the voltage to relax fully back to the OCV, or within measurement error. Depending on the rate capability and chemistry of the cell being tested, this can take anywhere from 2-4 hours. Even at a very coarse SOC step of 10%, this requires 10 rests, as well as the constant voltage (CV) hold prior to discharge (which would be expected to be significantly longer than the rests, or to have a rest immediately following it. This still takes between 1-2 days to perform, not including any setup time or thermal equilibriation. This is a barrier to its use in testing such as degradation studies, which are already a resource intensive process. GITT is very useful for the parameterisation of models, especially ECM's, as it can give excellent fidelity if much smaller pulses are used (e.g. 1% of nominal capacity), and is relatively easy to perform while maintaining an accurate temperature, provided the pulses are short enough that the heat generated does not significantly impact the behaviour during and immediately following the pulse. This allows accurate parameterisation across the entire operating window

of the device, especially useful when modelling inhomogeneities in packs, for example.

2.4.3 ICA

Incremental Capacity Analysis (sometimes referred to as dQdV) allows the diagnosis and prognosis of cells, and is the first of the techniques presented here that could reasonably be considered to be applicable *in situ*.[75] This is because it is very simple to apply, and requires only the acquisition of the OCV (or more frequently pseudo OCV) of the cell throughout the SOC range. This is typically achieved by discharging the cells slowly, as might be done when measuring the capacity, although can also be achieved from charge or a combination of charge/discharge data. The latter yields the best results, as the impact of overpotential can be corrected out where need be, and some chemistries exhibit features differently in the charge/discharge directions. While it takes longer to acquire, it does take less than twice the time as best results come after the cell has been properly conditioned. For example, a constnat current - constant voltage (CC-CV) charge should be used prior to acquiring data for the OCV, which can require several additional hours. However, OCV measured in the charge direction can be performed immediately after the discharge direction, with no additional preparation or conditioning. These OCV measurements are taken at rates as slow as reasonably possible, and best practice is usually C/25 or C/10.[7] It is preferable to have sufficient sampling rate that readings are taken at least every 2 mV. Ideally, a higher sample rate is applied, and then down sampled to remove duplicate readings, such that an array of results evenly spaced across the voltage window of the cell is obtained.

One of the advantages of this technique is the fact that the x-axis remains uniform across the life of a cell, allowing easy comparison between cells at different points in their life, such as might be done in a degradation study. It is an unfortunate reality that there will always be noise in any measurement process, so it is common for the data to be smoothed prior to differentiation. Care must be taken when doing this, as over-zealous smoothing can easily erase key features in the output curve. If the data has not been downsampled, and has multiple consecutive readings at the same voltage level (i.e. the sampling rate is higher than the voltage resolution), overfitting can also be a problem, where the smoothing captures the discrete levels of the measurement

instrument, as opposed to the response of the DUT. For this reason, the Savitzky-Golay method is popular as it is easily applied with modern tools, captures features in the differential well, and has excellent noise rejection.[5, 76] It is applied by fitting polynomial functions to a moving window of the data, optimised using the least squares method. This makes it good at capturing subtle details and features spread over a large number of data points, provided the signal is continuous. Discontinuities or step changes cause significant issues, as the tendency is to overfit the polynomials to capture the sharp edges, at the expense of validity in other regions. This is also true when the window on which it is applied is too small, as the polynomials begin to capture the noise, as opposed to filter it out. This is best explained by Fig 2.3, which is included with permission from Feng *et al*.[5] Feng *et al* propose a method, termed the LEAN method, which solves this problem, although it assumes data with a constant time-step and a sufficiently high sampling rate and noise rejection as to be primarily formed of multiple readings at each given voltage level. This is typical from data recorded from high quality scientific instruments using best practice, but unfortunately this is not the case for all data, such as that recorded by a vehicle system from a battery *in situ*. Therefore an understanding of a variety of methods, including their pitfalls and shortcomings, is important to ensure scientific rigour and validity of data.

The major drawback of ICA is the need for sample data to compare results to. This is termed as a degradation map, and consists of results from the same chemistry (and ideally identical cells to the DUT) which have been degraded in a variety of known ways. The ICA results from these cells can then be plotted and the movement of key features, such as peaks, plateaus and minima, noted. Results can then be compared to these curves of cells in known states, and the contribution of the various degradation modes dissected.

2.4.4 DMA

Degradation Mode Analysis (DMA) is an evolution of the ICA analysis technique. Instead of requiring measurement or simulation of the DUT in various states of degradation, it requires only the half cell OCV's of its constituent electrodes.[77] The OCV of the DUT is measured in the same way as for ICA, with all the same experimental best practices and caveats. Instead

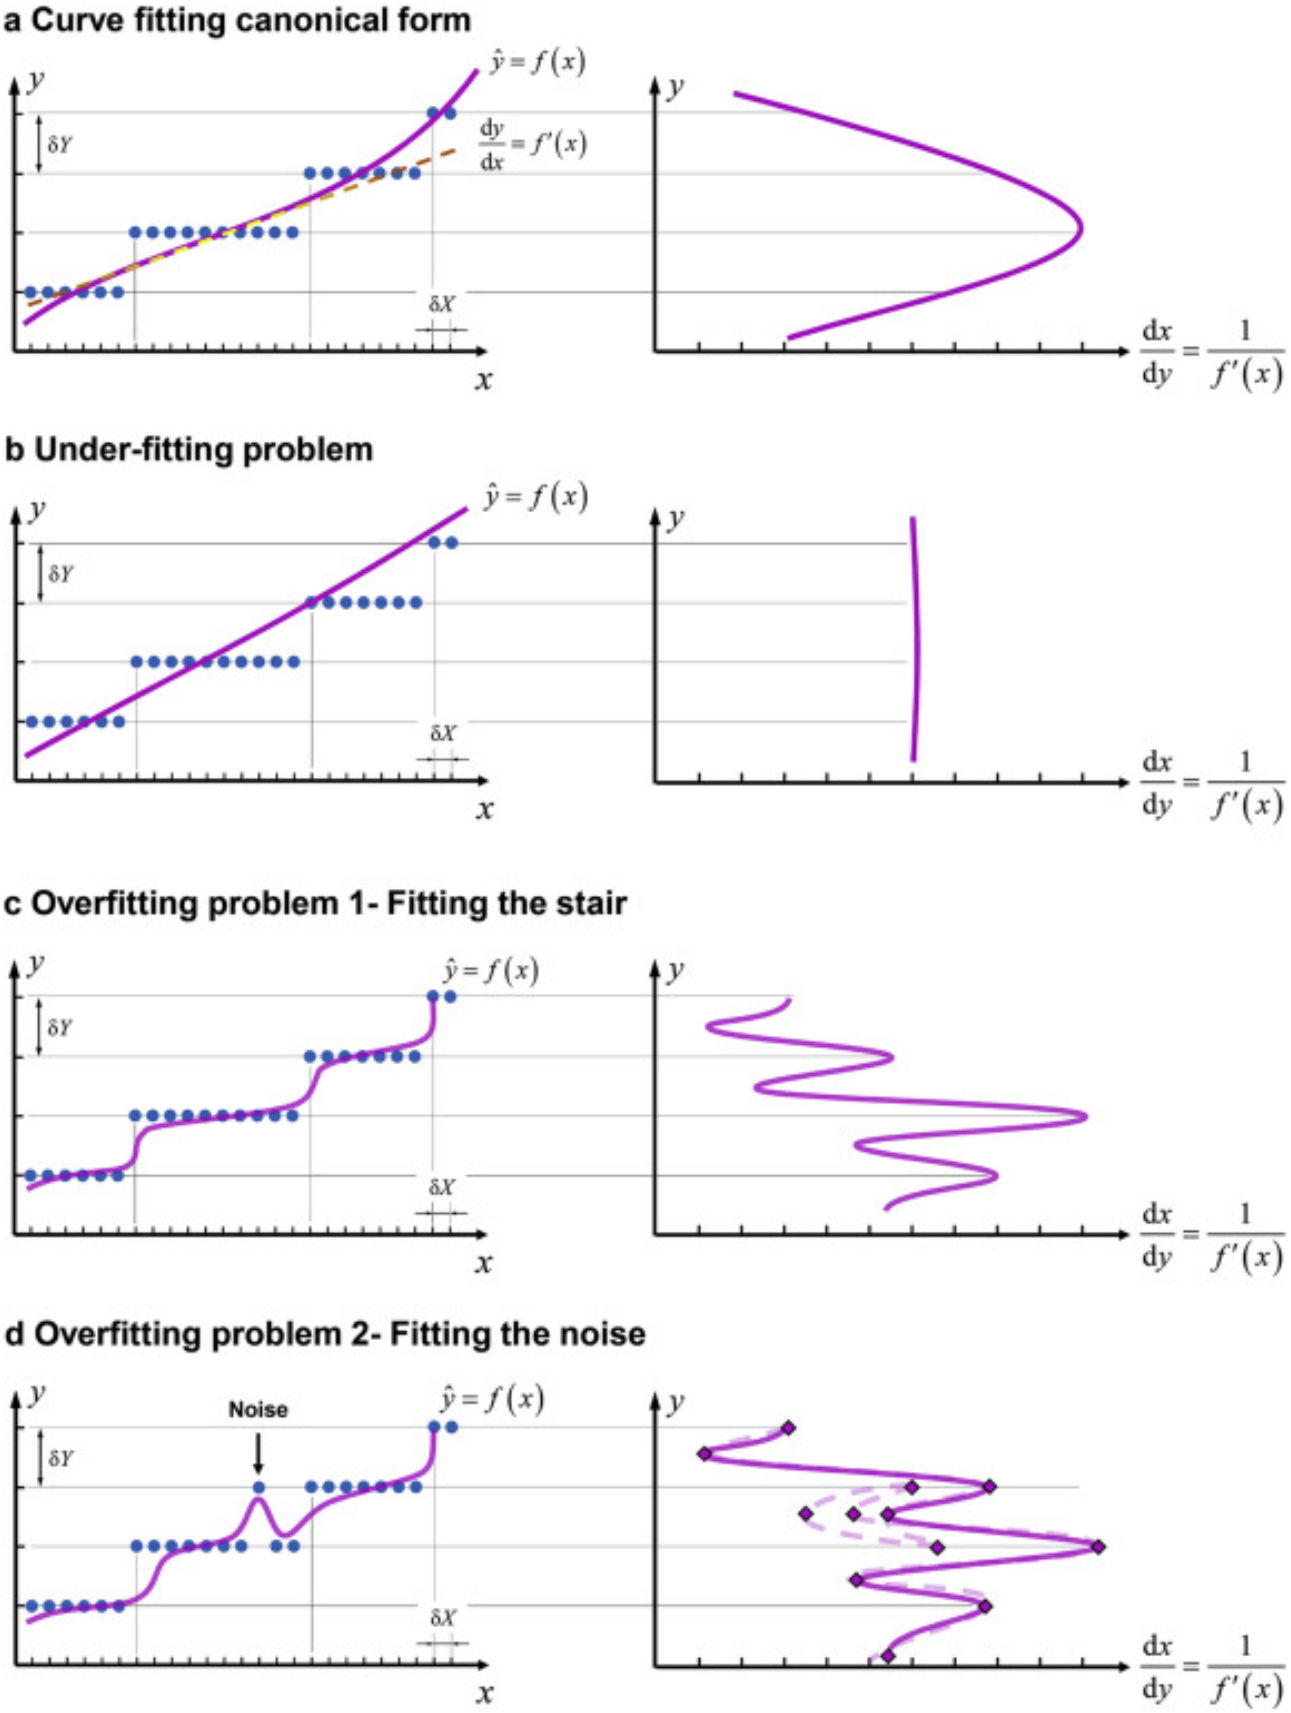


Figure 2.3: A hypothetical example of the problems with under- and over-fitting curves to real data for calculation of the differential curve. Reproduced with permission from Feng *et al*[5]

of analysing the differential of this curve, however, the half cell data is combined and fitted to it. The terminal voltage of any given cell is the sum of its two half cell voltages.[78] These half cells have their own complex relationship between SOC and voltage, determined by their chemistry and manufacture. Many of these are the same features that produce the features of interest in ICA, such as plateaus in the voltage where there are significant phase changes in the active material. By fitting the half cell data to the measured OCV, the relative capacities, lithiation fractions experienced, as well as the SOC offset between the two electrodes can be calculated. This allows quantitative measurement of the primary thermodynamic degradation modes (Loss of Lithium Inventory (LLI), Loss of Active Material (LAM) at both the PE and NE) from only the measurement of the OCV. While it does not matter in which direction this OCV was measured (i.e. charge or discharge), consistency is vital, and best results again come from the combination of the two in order to eliminate the impact of any overpotential.

The details of the process are mired in the minutiae of whatever coding language is used to apply it, but a high level approach is as follows. First, the half cell curves are converted into functions, which can then be fitted to the measured OCV data. This is done with a multi-variable optimisation, the fitting variables of which are the maximum and minimum lithiation fractions of each half cell experienced over the full SOC range of the cell. From these the apparent capacities of each electrode are calculated, as well as the offset between the electrodes SOCs. By applying this to a number of OCV measurements throughout the life of the cell, including a beginning of life (BOL) state, the LAM at each electrode is calculated (from the change in apparent capacity at each electrode), as well as the LLI (from the change in offset between the electrodes). This can produce remarkably accurate results, even from subpar input data, and yield insight that would be impossible to discern by eye from the ICA alone. It does struggle where the contribution of one or more of these degradation modes are poorly represented in the OCV, such as the case where the magnitude of one mode dwarfs others. In this case, the degradation modes are poorly represented in the OCV *because* their impact on cell capacity is minimal, so this is of less concern when it comes to cell diagnostics and prognostics, but problematic with regard to accurate parameterisation of a model to capture the cells behaviour.

2.5 Conclusion

In conclusion, the scope of the project has been narrowed to focus on PEMFC and LIB's, as these are the most promising technologies for electrifying the automotive sector based on a combination of theoretical performance and general maturity (i.e. actual performance of off the shelf units currently available, cost and scale of manufacturing). In the available literature for these devices, separate niches have been identified, and specific research questions posed which this project seeks to answer. Specifically, for PEMFCs the use of current perturbations or periodic short circuiting to improve both power output and efficiency of the device has been found to have both a high potential impact (>10% increase), as well as having very little available literature on its implementation, optimisation or mechanism. In particular, there is no available literature on the use of the technique on modern automotive style PEMFCs, which operate at significantly higher current densities (>2 Acm^{-2}), and are typically constructed from stamped metal plates with internal liquid cooling, compared to the older, bi-polar graphite plate construction with more limited current density (<800 $mAcm^{-2}$) which relies on its larger thermal mass and rejection of heat to the reactant flows. As such, the research questions posed address these gaps: Does the technique work on modern devices? What is the quantified benefits, and can they be improved via *in operando* optimisation? Does the behaviour of the device during the application of the technique give insight into the mechanism through which it acts? For LIBs, the opportunities for contribution identified are looking at longer time scales, specifically the degradation of the devices over their lifetime. Internal temperature and C-rate are both significant factors in the rate at which device performance is irreversibly lost, but are also intrinsically linked within typical commercially available devices. This is exacerbated by the non-isotropic thermal conductivity of the cells, making thermal management strategy a key area of engineering effort in the design of a mature, fit for purpose system. As such, the focus of this project in the area of LIBs is to produce a way to experimentally decouple the impact of these factors on the performance-through-life of a cell, allowing more accurate parameterisation of

degradation models for system design, specifically regarding the impact of thermal gradients within cells. This is of increased importance in Hybrid EVs as the ESS is typically of much lower capacity than a full BEV, but the power requirements are not similarly scaled down. As such, battery packs for hybrid systems are often required to be much more rate capable, and so are more dependent on good thermal management to maintain performance through life.

Chapter 3

Decoupling the Impact of Discharge Rate and Temperature on the Degradation of Li-ion Cells - Experimental design and validation

3.1 Introduction

The aim of the work this chapter represents was to produce a method by which the impacts of charge/discharge rate and temperature on the degradation of Lithium Ion cells might be decoupled, and to validate this method by way of a degradation study to produce indicative results for a selected chemistry. The chemistry selected here was NMC 532 due to manufacturer availability, and because it is a popular chemistry with a wide range of applications. Custom cells were commissioned from the manufacturer (Lifun), consisting of only a single layer of active material (compared to the many stacked layer typical in commercial cells). This allowed the bulk of the electrodes to be held at a uniform temperature, limiting any thermal gradients, facilitated by the use of a purpose designed fixture. The cells were then cycled at a range of discharge rate and temperatures, with characterisations taken periodically to allow comparison of their SoH. The degradation modes present are evaluated using a variety of techniques, and conclusions

regarding the impact of varying the discharge rate and temperature on device longevity are drawn. Cells assembled using material from the same manufacturing batch, but of a more conventional number of layers (and therefore capacity) were exposed to identical boundary conditions using a modified fixture, incorporating active thermal management using Peltier elements to produce the isothermal conditions required, and their degradation compared to the single layer experiments to validate the procedure.

3.2 Battery background theory

Lithium Ion batteries are used on a wide range of scales, from micro-devices to grid balancing. They have seen such an explosive growth in use primarily because of their excellent energy and power densities, as well as efficiency and scalability.[79] Fundamentally, a battery consists of two electrochemical reactions linked by two charge carrying systems. One system, the electrolyte, transports charge carriers (typically cations) from one reaction to the other within the device, and the other is typically an electrical circuit, transporting electrons through conductors.[80] By placing an electrical load in this circuit, the system can be made to drive it, discharging energy stored in the device, and by placing an electrical source in the circuit, this process is reversed, storing energy in the device.[81] These electrochemical reactions (typically termed "half cell" reactions) occur at a boundary between an electrode and an electrolyte.[82] In Lithium Ion batteries, these materials are designed to retain their structure across a wide range of SOC's by intercalating the charge carriers inside structured materials, instead of depositing/stripping them from a surface.[83]

3.2.1 Mechanical structure

Batteries are typically formed from layers of active material stacked and compressed together, then soaked with a liquid electrolyte to facilitate charge transfer. The most basic form of these layers (from negative to positive terminal) is as follows:[84]

- Negative current collector: Responsible for connecting the active materials electrically to the external circuit. Typically copper due to its material stability in this system.

- Negative electrode: Responsible for storing charge carriers in the "charged" state, and exchanging these carriers with the current collector and electrolyte during charge/discharge. See Table 3.2 for a comparison of materials

- Separator: Responsible for isolating the electrodes from each other electrically, forcing electrons through the external circuit, but allowing cations to be transferred through the liquid electrolyte.

- Positive electrode: Responsible for storing charge carriers in the "discharged" state, and exchanging these carriers with the current collector and electrolyte during charge/discharge. See Table 3.1 for a comparison of materials.

- Positive current collector: Responsible for connecting the active materials electrically to the external circuit. Typically aluminium due to its material stability in this system.

These layers are formed as individual electrodes, with active materials deposited onto their respective current collectors, then stacked. To improve performance, each electrode is typically double sided (with active material on both sides of the current collector sheet), so that when stacked the structure is mirrored.[85] There are a variety of form factors, the primary four being cylindrical, prismatic, pouch and coin cells.[86] In cylindrical and prismatic cells long rolls of the electrode materials and separator are wound around a mandrel and then encased in a rigid can.[86, 87] This allows for rapid manufacture, as the rolls can be produced in a continuous process, and therefore helps to reduce cost. Pouch cells are formed by cutting the electrodes into a uniform shape, typically a rectangle, and then stacking these layers.[86, 87] They are usually encased in a heat sealed flexible material, hence a "pouch". Finally, coin cells are formed by punching round sections from the electrode sheets and stacking them within a flat round can. Cylindrical and coin cells have their current collectors connected to different parts of their conductive cans, whereas prismatic and pouch cells have tabs or lugs protruding from their containers to allow electrical connection.

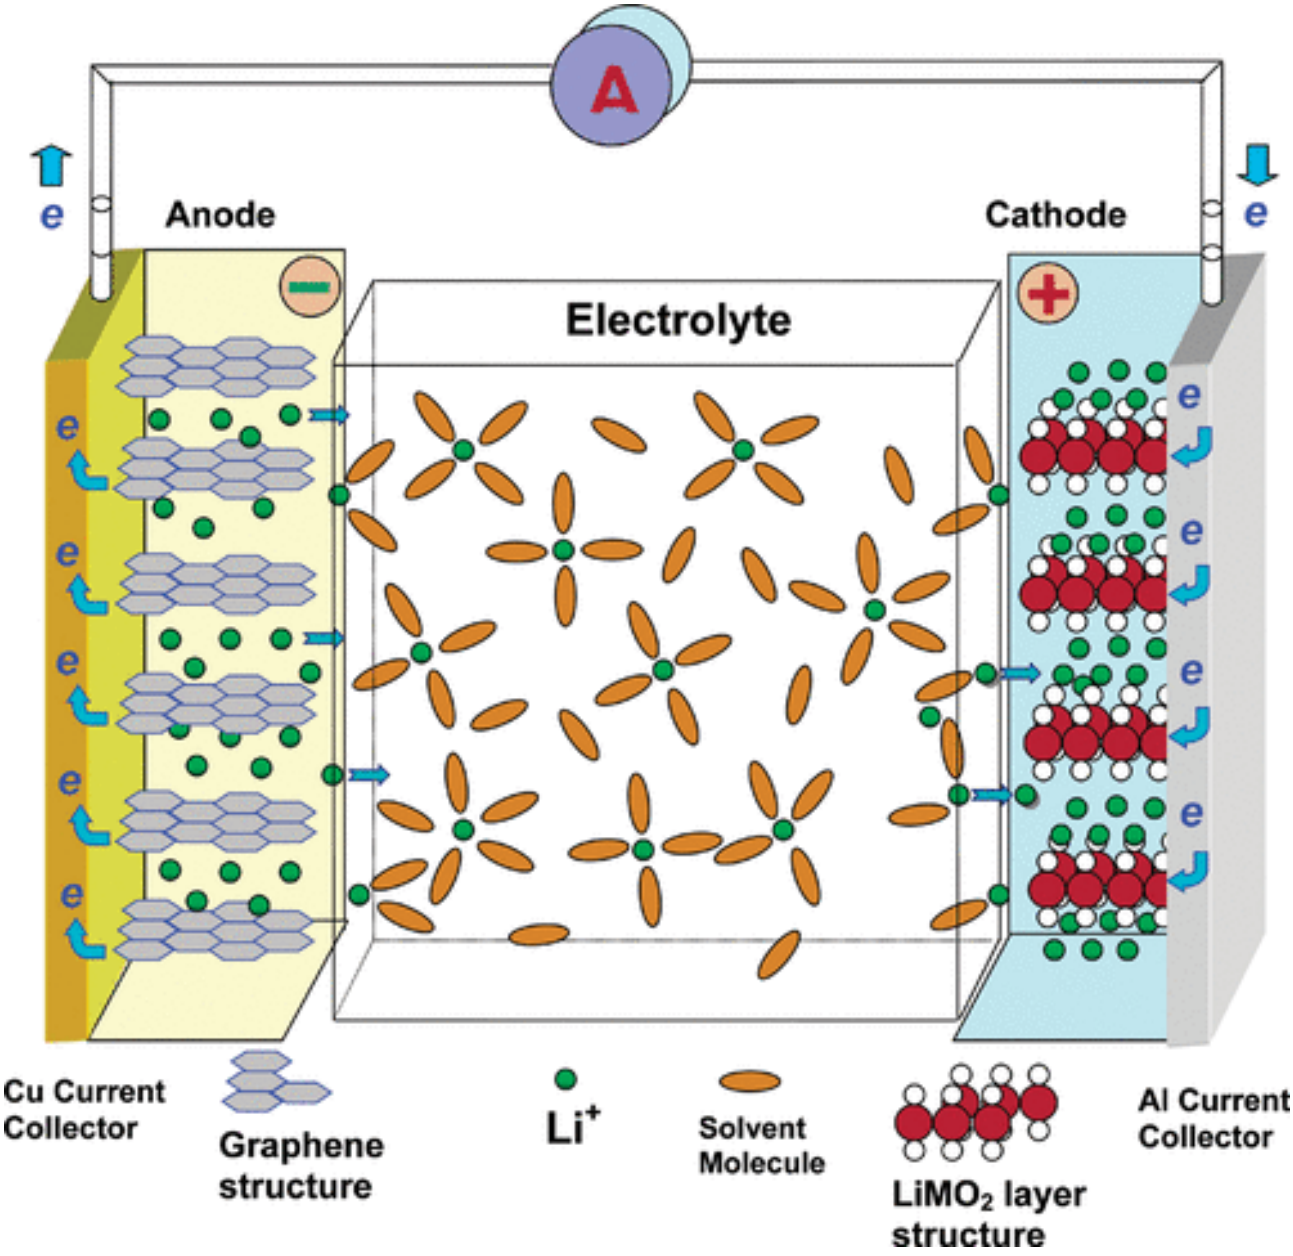


Figure 3.1: A diagram showing the basic structure of a LIB Included with permission from[6]

3.2.2 Materials

As mentioned previously, this form of battery intercalates the charge carriers within the electrodes, as opposed to intentionally depositing or removing them from its surface. This means that materials with a repeating structure that have pores or voids are required, as well as good mechanical stability to prevent degradation of the materials by repeated charge/discharge cycles.[88] Typical materials, including a brief comparison of their relevant properties can be seen in Tables 3.1[2] and 3.2.[3]

While there is interest in the use of a wide variety of materials for the NE, in practice only graphite, silicon-graphite mixes and Lithium Titanate ($LiTi_4O_5$) are available in commercial cells.[2, 3] Silicon is often used with graphite to increase the specific capacity, but exhibits a significantly larger volumetric change across its range of lithiation (≈300%),[89] and therefore suffers from reduced cycle life from the higher mechanical stress exerted during charge transfer. Lithium Titanate is extremely safe and low cost, as well as having excellent rate capabilities. Unfortunately, it suffers from significantly lower energy density thanks to the reduced cell voltage compared with graphitic anodes. This limits it's use somewhat, but it remains an excellent choice for applications where safety is vital, or where a high power to capacity ratio is needed, such as the ESS of a partially hybridised vehicle (e.g. PEMFC based public transit bus, which has a low mean power but frequent acceleration/deceleration). LTO cells fill a niche between super/ultra-capacitors and more traditional LIBs with regards to capacity and rate capability (power).

A much wider variety of PE materials are commercially available, with the differences being significantly more nuanced, and often driven by factors other than performance metrics, such as the availability/price of component minerals, recyclability or geo-political drivers. For example, the development of low cobalt NMC ($LiNi_AMn_BCo_CO_2$) compounds such as NMC811 over NMC111 due to the scarcity of Cobalt,[90] as well as the presence of various unethical practices in its mining. An examination of these materials, as well as the real world complexities involved in their manufacture, is beyond the scope of this project.

Table 3.1: A table listing the properties of PE intercalation compounds, adapted from[2]

Crystal structure	**Compound**	**Specific capacity (mAh g^{-1})** (theoretical/experimental/typical in commercial cells)	**Volumetric capacity (mAh cm^{-3})** (theoretical/typical in commercial cells)	**Average voltage (V)[91]**
Layered	$LiTiS_2$	225/210[92]	697	1.9
	$LiCoO_2$	274/148[93]/145	1363/550	3.8
	$LiNiO_2$	275/150[94]	1280	3.8
	$LiMnO_2$	285/140[95]	1148	3.3
	$LiNi_{0.33}Mn_{0.33}Co_{0.33}O_2$	280/160[96]/170	1333/600	3.7
	$LiNi_{0.8}Co_{0.15}Al_{0.05}O_2$	279/199[97]/200	1284/700	3.7
	Li_2MnO_3	458/180[98]	1708	3.8
Spinel	$LiMn_2O_4$	148/120[99]	596	4.1
	$LiCo_2O_4$	142/84[100]	704	4.0
Olivine	$LiFePO_4$	170/165[101]	589	3.4
	$LiMnPO_4$	171/168[102]	567	3.8
	$LiCoPO_4$	167/125[103]	510	4.2
Tavorite	$LiFeSO_4F$	151/120[104]	487	3.7
	$LiVPO_4F$	156/129[105]	484	4.2

3.2.3 Performance Metrics

Unlike many other forms of battery, Li-ion cells are remarkably shelf stable because they experience self-discharge significantly lower than their common counterparts, such as lead acid or NiMH cells.[127] In addition, they do not require complex cycling in order to achieve their maximum usable capacity.

There are two main metrics by which we can evaluate a cell's performance: capacity and rate capability.[128] Capacity is a measure of the total charge or energy a cell can store and release, and can be measured in a variety of ways. Rate capability is a measure of the rate at which a device can exchange this capacity. In order to allow more easy comparison between cells of different sizes, structures and form factors, this is usually non-dimensionalised by dividing the current in Amperes by the cell's total capacity in Ah. Typically, it's nominal capacity is used for this, defined by the manufacturer, although two cells with identical nominal capacities may deliver very different usable capacities. These non-dimensionalised rates are referred to as C-rates, where 1C is the current required to discharge the full nominal capacity of the cell in one hour. Values greater than 1C are written with "C" treated as units, such as 2.5C or 4C, whereas values less than unity are typically written as a fraction, with "C" as the numerator,

Table 3.2: A table comparing various NE materials, adapted from[3]

Materials	**Capacity (mA h g^{-1})**	**Advantages**	**Drawbacks**
Graphene[106]	770/1115	Safe for use	Greater irreversible capacity
Carbon nanotubes[107–110]	1115	Economical	Greater hysteresis voltage
Hard carbon[111–113]	200–600	Excellent working potential	Less coulombic efficiency
TiO_2[114]	320	Enhanced power capacity, extended life cycle	Reduced levels of energy densities
$LiTi_4O_5$[114]	176	Safer, low cost	Lower capacities
Tin[115]	992	Higher precise capacities, Increased energy densities, outstanding safety	Reduced irreversible capacity, Increased capacity fading, Shorter lifecycle.
Tin oxide[116]	793	Greater energy densities, excellent safety, greater specific capacities	More fading capacity, Lower irreversible capacity, Shorter lifecycle
Silicon oxide[117]	1562	Higher precise capacities, increased energy densities, outstanding safety	Lower irreversible capacity, More fading capacity, less life cycle
Antimony[118]	661	Greater energy densities, excellent protection, greater specific capacities	More fading capacity, Lower irreversible capacity, less life cycle
Silicon[119]	1562	Higher precise powers, Increased energy densities, Outstanding safety	Reduced irreversible capacity, Increased capacity fading, Shorter lifecycle.
Germanium[120, 121]	1623	Greater energy densities, excellent safety, greater specific capacities	More fading capacity, Lower irreversible capacity, less life cycle
Phosphides, nitrides, and sulphides of metals[122–124]	500–1800	Increased capabilities, reduced potential, diminished polarization	Limited rate capability, brief lifespan, elevated production cost
Transition metal oxides[116–126]	500–1200	High capacity, energy efficiency, low cost, eco-friendly	Reduced Coulombic efficiency, unstable solid electrolyte interface (SEI) formation, unfavourable lifecycle characteristics, and increased hysteresis potential.

for example $\frac{C}{10}$.

While these might seem like simple and straightforward definitions, they are, unfortunately, woefully inadequate. Firstly, there is a difference in measuring charge and energy as the voltage of the device shifts with SOC and the relationship between voltage and SOC can change for a wide range of reasons, both reversible and irreversible. The voltage delivered will also be impacted by the rate at which the measurements was taken as well as it's rate capability, due to the overpotential required to achieve the desired rate. Clearly, metrics for the performance of a cell must be carefully chosen and defined, and the method for acquiring them must be applied with rigour.

When stating a cells capacity, units are insufficient to communicate the details of the measurement, and the rate at which the measurement was taken must also be stated.[129] Temperature will also have a significant effect, and should be explicitly stated where possible, as well as the charge protocol preceding the measurement. In general, the appropriate procedure is to charge the cell to its maximum voltage at constant current (typically 4.2V for a standard Li-ion cell with a graphite NE), and then hold this voltage until the current decays below a given threshold (usually $< \frac{C}{10}$). This is defined as a CC-CV charge. The cell can then be discharged at a constant rate until it hits the lower voltage limit (between 2.5-3V, depending on cell chemistry and manufacturing tolerances), and the definite integral of this current with respect to time is a measure of the capacity of the cell.[130] This can be done at very low C-rates, to give a better idea of the capacity of the electrodes irrespective of rate capability, or at C-rates representative of those seen in use, to give an indicator of the cells "usable" capacity.[7] The usable capacity will of course be smaller, as it is limited by the overpotential required to achieve the desired C-rate, but will include the impact of the cell's rate capability and is a better metric for use in implementation. For clarity in scientific purposes, however, it is often necessary to decouple the capacity and rate capabilities as much as possible.

Similarly, rate capability can have multiple meanings. The Ohmic resistance is a single number, measuring the overpotential required for a given current across the majority of a cell's

SOC window, as in a Thevenin source. This is measured in Ohms, and is typically found by applying a step change current from a cell at rest and measuring the instantaneous voltage change exhibited by the cell. "Instantaneous" is important, because the overpotential will increase, asymptotically approaching a steady state as diffusion gradients are produced by the exchange of charge carriers between the electrodes and electrolyte. In order to measure this behaviour, a complex impedance is often used, as in electrical systems. However, this impedance is strongly influenced by the frequency of the excitation as well as the magnitude, and other factors such as temperature and SOC.[65] The two most common techniques for measuring this are Electrochemical Impedance Spectroscopy (EIS) and the Galvanostatic Intermittent Titration Technique (GITT).

The former consists of a series of excitations, with the response of the device measured in order to produce a plot of the impedance across a range of input parameters. For Li-Ion cells these excitations are typically a current, the magnitude of which follows a sine wave and the frequency of the wave is varied across an operating window, but the magnitude of the wave remains uniform. The voltage response of the cell is then measured, producing a Nyquist plot of the cell's complex impedance at a quasi- uniform SOC. The SOC will change somewhat during the test due to charge being exchanged, but should always fluctuate around the starting point, and the changes in impedance should be minimal provided the extreme edges of the SOC window are not reached.[65] GITT, on the other hand, is applied as a series of pulses, either in charge or discharge, each followed by a period at OCV (or "rest"). During each pulse the Ohmic resistance can be found from the instantaneous overpotential developed, while the more complex kinematic information (i.e. complex impedance) is found by fitting curves to the voltage changes exhibited both during a pulse and in the return to OCV during the rest period.[131, 132] This is more computationally complex and involved, but can yield more detailed results than just a single sweep of complex impedance over a range of frequencies. It is typically done over the majority of the SOC window, although if set up carefully the time required can be shortened by applying it at a fixed SOC repeatedly over the lifetime of a device. This is not the typical use, however, as one benefit of GITT compared to EIS is better resolution of changes in the

impedance details over the SOC window.

3.3 Methods

3.3.1 Introduction

The primary difficulty in investigating the decoupled impacts of temperature and C-rate on Li-ion cells are the thermal gradients produced by the internal heat generation of a cell. For a standard cell of any form factor there will be thermal gradients produced, as the heat is generated throughout the material, but can only be rejected through the boundaries, and even then, typically only a select area is thermally managed in implementation. While some heat will be generated from the electrical resistance of the current collectors this is negligible[133] compared to the amount generated by the reactions occurring in the active materials of the electrodes. This heat must be transferred out of the cell to maintain temperatures within the desired operating window, and this transport relies on a gradient to drive it, as there are no active cooling methods inside a typical cell. For cells operating at significant powers, external active cooling is usually required to be attached to one or more surfaces of the cell. For cylindrical cell, this can be either the base or sides of the can, prismatics can be cooled from any available side, and pouch cells are almost always cooled from their largest faces due to their aspect ratio typically being much higher than other form factors. Due to the layered structure of a cell, as well as the bulk properties of the materials used, the thermal conductivity is non-isotropic. The electrodes are much more thermally conductive in the plane of the current collectors, and much less perpendicular to this plane. This is in part due to the high thermal conductivity of the materials used (copper and aluminium having significantly higher thermal conductivities than the active materials they are coated with), as well as the interfacial resistance between layers. In order to travel in plane, heat can conduct along a uniform material with a low thermal resistance, but to travel across planes it must conduct through the active materials as well as the polymer separator, as well as passing from material to material, which will have its own added source of resistance based on surface features and applied pressure. As such, the thermal gradients are also non-isotropic. The intention behind this body of work was to make the first steps

in quantifying the decoupled impacts of C-rate and temperature on the degradation of Li-Ion cells.

Being able to characterise the de-coupled impact of C-rate and temperature would allow the parameterisation of more complex thermal and electrical coupled models, such as by volume discretisation of the device. This, in turn, allows much greater fidelity in evaluating the impact of design decisions regarding the thermal management of a system on the overall performance and lifetime of the system. This is especially true in systems with a higher peak power/energy stored ratio, such as in high performance applications or in HFCEV's, where the ESS is typically significantly smaller than in a full BEV, but is still required to meet nearly the full power demand of the powertrain, albeit for much shorter periods (i.e. to improve load following over a non-hybridised FC system). As it has not yet been possible to fully decouple these parameters, parameterising these types of models at a device level is difficult, often relying on material level data to model individual components, which can have inaccuracies due to the difficulty in modelling the interface between components or due to the complex interplay between components during operation.

3.3.2 Cells

De-coupling temperature and C-rate in a normal cell is therefore impossible, as the many layers of material will inherently be at different temperatures. As such, custom cells with only a single layer of active material were commissioned from cell manufacturer Li-Fun. These are formed of a single double sided negative electrode sandwiched between two positive electrodes. It was initially requested that these positive electrodes only be coated with active material on the sides facing the opposing electrode, but as this would require custom manufactured electrode material, this was not only prohibitively expensive, but it was also deemed that this might reduce the ease of comparison between the custom cells and the control cells. This does unfortunately mean that some of the longer time constant behaviours of these cells are atypical, as there is an imbalanced amount of active material, but this can be identified and accounted for in the design of the experiments. Because there are not multiple layers of active material, just a single symmetrical unit, the entire cells active material can be held at a much more uniform temperature regardless

of C-rate. In order to examine the difference this unique structure produces a set of control cells were also purchased. These are an extant size pouch cell sold by the manufacturer, which is the same footprint as the single layer versions, but with many more layers, as is typical. The material used for both the single layer and pouch cells was from the same manufacturing batch, in an attempt to further reduce any differences between them.

3.3.3 Mechanical fixture

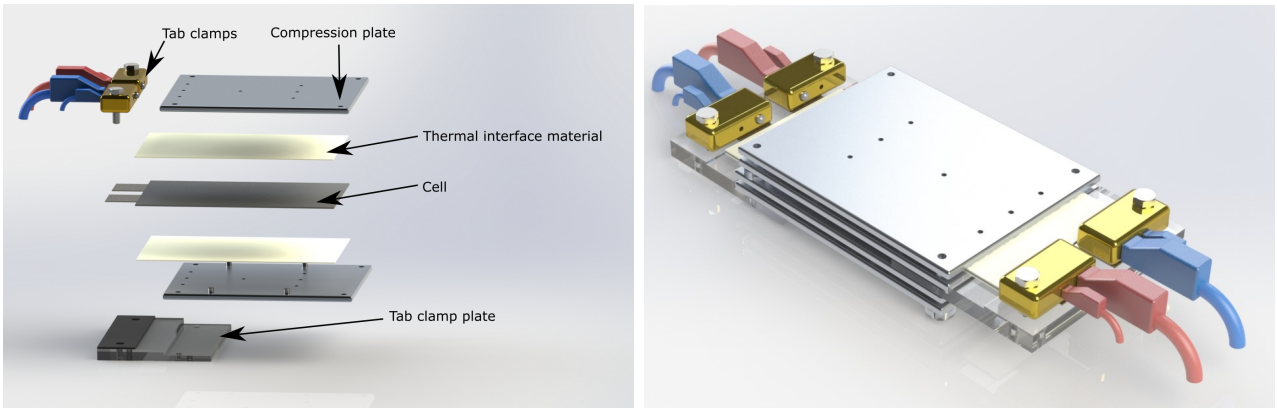
Requirements

The mechanical fixture required must apply tightly controlled isothermal boundary conditions on the surfaces of the cell, with sufficiently low resistance thermal pathways as to maintain the internal temperature of the cell as close to this boundary temperature as possible. It must also control the pressure applied to the cells, as this is not under investigation in these experiments, and so should be fixed and uniform across all devices under test. A pressure of 150 kPa was selected to be sufficient to ensure good thermal contact, while being low enough to fall into a window in which pouch cells are insensitive to changes in pressure.[134] The cells must also be given an electrical interface to allow them to be cycled and characterised. Finally, total fixture size must be minimised in order to maximise the number of cells than can be tested with the equipment and space available.

Design

In order to achieve these requirements, a fixture was designed which allows each cell to be mounted in its own unit, and these units can then be stacked together and compressed to achieve good thermal contact between the cell and its thermal management unit. A render of this design is shown in Fig 3.2. This allows all cells in a stack to experience uniform pressure of 150 kPa, as well as the instrumentation of that pressure, which would be prohibitively complex and expensive to achieve on an individual cell basis. Due to the high surface area/capacity ratio of the cells, it is not required to apply separate actively managed isothermal boundaries. Instead it is sufficient to confine the cells in comparatively large thermal masses in good thermal

contact, which in turn exchange heat with a coolant fluid at the desired temperature. For this purpose, 3 mm thick aluminium plates on both sides of each cell were used. These are coated with a thermal interface material, which both helps keep the pressure applied more uniform, as it can conform to any irregularities in the surfaces, as well as aid the heat transfer between the cell and the compression plates. These plates were located relative to each other within their cell unit using small dowel pins. The end of the plates next to the tabs of the cells are machined to mate with a plastic plate, which has a step machined into it. The thinner section of the plate formed by this step acts to space apart the cell units, and has through holes to keep it located by the tie rods holding the stack together. This is important to prevent any mechanical strain on the tabs, as well as give a sturdy interface for the electrical connections. This plastic plate also has threaded holes for the mounting of brass blocks, which clamp the tabs to the plastic surface using a bolt. These brass blocks hold the tabs in place and have holes in them to accept sprung banana plugs, providing the electrical connections. They are sized (and the mounting holes positioned) such that they can never touch each other, which is important for safe assembly and operation of the rig. Once clamped in place, they are also covered with a plastic cover, which has been 3D-printed with a compliant mechanism at one end, holding them in place on the blocks. This fully isolates the electrical connections, with the only surfaces exposed being the holes for connection to the instruments, which are covered by the housing of the banana plugs in operation. The total pitch between stacked cell units is 11mm per cell.



(a) Exploded diagram of a single cell holder (b) Render of a pair of cell holders stacked

Figure 3.2: CAD images of the single layer rig design

The stacks are held together using 4 silver steel tie rods, one in each corner. These tie rods are 3mm in diameter, as they need to transmit a relatively low force, and are threaded at each end to mount the compression hardware. Rods threaded over their entire length were not used to prevent friction between the tie rods and any cell units, ensuring they only constrain the units within their plane, and allow them to slide freely vertically. This ensures that the compressive forces are transmitted from unit to unit via the spacers only, producing uniform pressure distribution across all units in the stack, as well as across the surfaces of the cells within the units. Silver steel was selected due to its resistance to corrosion, keeping the surfaces smooth, as well as it's resistance to creep under mechanical load. This is important to maintain the compressive force applied to the stacks over a long period of time.

The compressive forces were applied using regular hex nuts. This applies a dimensional constraint on the stack height, but this does not directly translate to a uniform pressure. Due to the stiffness of the materials used, a small change in unloaded stack height would produce a very large change in pressure, and these variations in unloaded height are to be expected to occur from 3 main sources. Firstly, the thermal expansion or contraction of the fixture itself. Secondly, the active materials within a cell experience volumetric changes when cycled, as the charge carriers intercalate/de-intercalate with changing SOC and the positive and negative electrode active materials expand differently in response to this. Finally, the thermal interface material relaxing and flowing into surface irregularities over time. In order to ensure uniform pressure over time despite these sources of error, springs are placed between the hex nuts and the top compression plate. These mean that a change in stack height will have a much smaller impact on pressure, as they are much more compliant. However, for traditional helical compression springs to apply the required force they would have to be made from a relatively thick wire diameter, and would therefore also have a large bore diameter. While these springs are available, they would require a significant free length, greatly increasing the total height of a fixture stack, as well as storing sufficient energy to be both difficult to assemble, as well as potentially hazardous.

Instead, a type of spring called a Belleville washer was used. This is a conical washer,

stamped from sheet spring steel, and their manufacture is a very mature technology, allowing excellent uniformity and low cost. They are very low profile axially, and it is easy to find them with an appropriate bore diameter to fit the tie rods. Belleville washers do not have a constant relationship between force and deformation, as is found in helical springs, instead exhibiting an inflection point halfway between their free height and fully flat. This inflection point also occurs at half the force exhibited when the spring is completely flat, and the magnitude of the inflection depends on the ratio between the free height of the cone and the thickness of the washer.[135, 136] This can be further adjusted by stacking the washers to change the apparent spring constant. By stacking them all in the same orientation (e.g. all cones pointing in the same direction), their apparent spring constants are summed, as the same displacement must deform all of the springs the same amount. However, by stacking them in alternating orientations, for the same force their displacements sum. In this way, they act as a mechanical version of series and parallel circuits in electronics, allowing a combination of individual springs to deliver a wide range of desired force/displacement characteristics. For these tests they were used in pairs of alternating orientation at each end of every tie rod, reducing the apparent spring constant by a factor of 4 while taking up just a few millimetres of axial length. This allows accurate control of the pressure applied to the stack, as well as insensitivity to the order of magnitude of dimensional changes expected from the stack due to the sources listed above. One notable exception to this is the initial relaxation of the thermal interface material, which will flow under the applied pressure to fill any surface irregularities in the cell or compression plates. However, this occurs relatively quickly initially, as any inhomogeneities cause increased pressure in localised regions, pushing the thermal interface material away, rapidly equalising surface pressure. As such, the stacks can be compressed, left for a short period (20-60 minutes) to relax, and then adjusted to account for any changes, at which point they will maintain uniform pressure indefinitely. Validation of this will be discussed in Section 3.3.3. Washers were selected such that their load at flat was twice that of the desired load, and with a free-height/thickness ratio of approximately 1.4. This causes the spring constant to inflect to 0, with the spring acting as a constant force spring around the desired load, further decreasing the impact on changes in stack height on the pressure applied. Specifically, the chosen spring washers are Lee Spring

Part Number: 138-022-437, the specification of which can be seen in Table 3.3. The area of the cell between the two compression plates is approximately 90mm by 58mm (i.e. 5.2 x 10^{-3} m^2), therefore to apply 150 kPa a total force of 783 N is required, or 196 N per tie rod.

Table 3.3: A table showing the manufacturer specification of the Belleville spring washers used

Description	Specification
Spring Type:	Belleville Washers
Outside Diameter (mm):	11.1
Inside Diameter (mm):	3.51
Thickness (mm):	0.56
Calculated Load at Flat (N):	396.3
Free Height (mm):	0.81
Bolt Size:	M3.5
Material:	SS
Work Over Rod Diameter (mm):	3.51
Finish:	Passivate per ASTM A967

Instrumentation

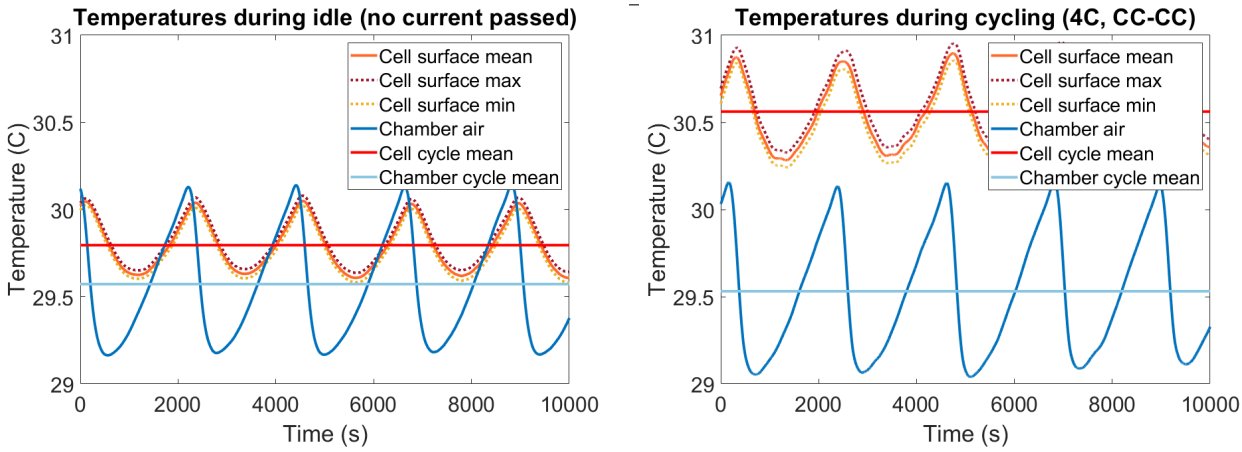
As well as the electrical instrumentation connected to the cell tabs, the fixtures have both thermal and compressive instrumentation. The former is provided by a pattern of small holes in the compression plates. These can accept either thermocouples or sufficiently small thermistors, and allow these measurement devices to measure the surface of the cell directly. The pattern is arranged to help identify any in plane thermal gradients, although this was primarily used for validation of the fixture, and in operation each cells temperature was only measured at a single point at the centre of the top surface. This measurement is only used to ensure that the cell is experiencing the correct boundary conditions, as opposed to any investigative or control techniques.

Compression was instrumented by the use of compressive load cells. One end of each stack had an extra pair of compression plates which sandwiched 4 load cells. These were placed as

close to each tie rod as possible, with their point of contact lying on the diagonal lines between opposite rods. This makes each load cell very sensitive to changes in length of the tie rod to which it is adjacent, while relatively insensitive to changes in the other rods. By careful and methodical cross tightening of the rods, the force applied by each one can be balanced, ensuring uniform pressure across the entire surface of the unit. The stack is then allowed to relax until there is no perceptible change in load over time, typically requiring about 20-60 minutes to do so. Any imbalances are then corrected, and provided they were within expected limits (<5% of load per rod), the load cells were disconnected from their measurement electronics. If a larger imbalance occurred, the stack would either be re-balanced and allowed to settle again or disassembled to identify the cause of the discrepancy. However, no stacks required disassembly, with only a couple requiring a small additional re-balancing before settling within expected limits.

Validation

In order to ensure that the fixtures meet the design requirements, as well as to characterise their performance, some simple validation steps were taken. Firstly, the load cells used to measure and balance the compressive loads were re-connected to their measurement electronics before any stack was disassembled at the end of testing, and the change in force measured to ensure it fell within measurement error of the force applied during assembly. Secondly, the rig was validated thermally to ensure that the cells experienced sufficiently stable and uniform boundary conditions, as well as ensuring that the thermal pathways between cell units were sufficiently small as to ensure that the behaviour of a cell would not impact adjacent cells in the stack. This was done by building a stack of 5 cell units, populated with cells degraded from initial planning experiments (i.e. highest resistance, and therefore greatest heat generation expected to be seen in any test performed). The cells were thermally instrumented, with extra measurement point placed on the very centre cell, as this would be expected to demonstrate the greatest impact of adjacent cells, as well as having the most resistance to airflow over the compression plates for heat transfer. This was allowed to acclimatise in a thermal chamber before cycling the cells at 4C continuously, with no constant voltage steps. This is far more aggressive cycling than used



(a) Temperatures at rest

(b) Temperatures during validation cycling

Figure 3.3: Rig thermal performance validation results

in any of the degradation profiles and was intentionally designed as such to produce the largest possible heat generation. The degraded cells were determined to have an Ohmic resistance of approximately 120 mΩ, and 4C is 640 mA, producing an expected heat generation of 49 mW each. The cells were allowed to cycle for a number of hours to ensure that they reached steady state, and to give an abundance of data to ensure confidence in the results. An excerpt of results can be seen in Fig 3.3.

The thermal chamber used (a Binder KB23) unfortunately suffers from an approximately 1°C peak to peak variation in air temperature under normal operation due to its control system, with a period of approximately 40 minutes. This period, as well as the appearance of this behaviour while the cells are at rest allows us to discount any contribution of cell heat generation to this fluctuation. This temperature swing is halved at the cell surface by the thermal mass of the compression plates acting to damp the response, as intended. Any perceptible difference in temperature between points on the cell surface, including between the top and bottom surfaces, is well within the measurement error, but the maximum, minimum and mean lines have been plotted in order to demonstrate how tightly grouped the results are. Under the aggressive cycling profile, the cell surface does see a temperature rise, but only of 0.81°C. While this is greater than initially expected, it is also well within acceptable limits. The variation between points on the cell surface does also appear to grow a small amount, but again is well within measurement

error, so any further conclusions drawn from this fact would be suspect. Interestingly, the variation in chamber air temperature appears to grow slightly also, perhaps due to the slightly increased load it must work against, but given how small a thermal load these cells produce it is as likely to be from a minor change in lab temperature as anything else. Again, the difference is within measurement error of the instruments, and so any further conclusions from this fact would be spurious at best.

3.3.4 Active thermal management for pouch cell control experiments

In order to demonstrate the validity of the experimental premise as a whole, namely that single layer cells will have minimal thermal gradients due to internal heat generation and will therefore act significantly differently in degradation compared to their conventional counterparts, a set of pouch cells of the same material batch and active area footprint was subjected to identical tests. These conventional pouch cells differed only in the number of layers (and therefore total capacity), but were otherwise identical. However, because of the larger heat generation produced, it was deemed necessary to actively manage the boundary conditions of these cells. In order to do this, the space between cell units was increased (also useful to minimise the impact of adjacent cells on each other), and Peltier elements were mounted on each surface. These rejected heat to a water block sandwiched between them and were actively managed by use of thin leaf thermistors placed on the cell surfaces. This allows the surface to be controlled incredibly accurately, with $<0.1°C$ variation, even in response to an aggressive step change in heat generation of the cell. This is primarily to do with the thermal resistance between the Peltier element surface and the cell surface being significantly lower than that of the cell itself (in the through plane direction, certainly), as well as the very rapid response times of the Peltier elements. Due to being a solid state device, they can react to perturbations much more rapidly than any other thermal management system, especially when implemented with a carefully tuned control loop and sufficiently accurate measurement device. Due to minor variations in device characteristics, as well as the unique and non-linear response characteristics of thermoelectric devices, each control loop was manually tuned for the specific experiment it was implemented in. It should be noted that the aim here is not to produce an average cell

internal temperature that is identical to that experienced by the single layer cells, but identical isothermal boundary conditions. To this end, as well as having active thermal management on their major surfaces, the rigs were also placed in thermal chambers at the desired surface temperature to ensure that all other surfaces experience the correct boundary conditions. The impact of this is likely minimal, as the heat flux due to free convection is negligible compared to that of the major surfaces, but it was done regardless from an abundance of rigour.

3.3.5 Cycling and Characterisation

Some initial characterisation of the cells was performed to determine their rate capabilities, cell to cell variation and rate of degradation in response to a variety of cycling types. Due to being long term experiments on complex electrochemical systems, degradation studies produce data that requires significant processing to yield reliable results, so it is vital to determine suitable cycling profiles to implement. This is especially true for these experiments, which utilise hand-made custom cells, which already differ from cells used in real world applications in their construction. Of particular note is the fact that there is a significant imbalance in capacity between the positive and negative electrodes active materials thanks to the double coating on the external layers having no opposing electrode. In addition, it is expected that these cells have a larger ratio of electrolyte to active material, due to the pouches being drawn less closely over their thin contents. In addition, each data point must be validated over a larger number of repeats than would be required from short term experiments on mass produced, commercially available cells.

Variables were eliminated primarily by focusing on what scientifically valid conclusions could be expected from such a study. As the direct comparison with any real world implementation would already be suspect, use of drive cycles or complex charging profiles (such as might be used to maximise charge current while avoiding specific degradation mechanisms such as lithium plating) was disregarded. Instead, it was determined that a CC-CV charge would likely yield the best clarity of results (while also ensuring a uniform SOC cycling window throughout the tests in a way that suitably reflects the implementation of these devices), and a CC discharge of uniform

capacity through coulomb counting. This is instead of other depth of discharge determinations such as energy or voltage, and was selected so that data from cells which exhibited different voltage characteristics (either because of variation in applied overpotential due to C-rate, or from a divergence of degradation mechanism pathways) would be more readily comparable. If a specific degradation mechanism was under investigation, perhaps which was sensitive to voltage window, a different determination may have been made. However, the intention here was to look at total degradation of the devices in way that can be related to their commercial counterparts, as well as an investigation of the impact of thermal gradients as a whole. To this end, a depth of discharge (DOD) based on initial capacity was selected, in this case 40%. This DOD was selected such that the most aggressively cycled data points will begin to fail to deliver it at the end of the expected experiment length, as they will hit the lower voltage limit.

The cells were cycled and characterised using the following protocol:

1. Full Characterisation - 30°C
   (a) Rest - 3 hours
   (b) CCCV Charge at C/4 to 4.2V, ending when I < 10mA
   (c) Rest - 10 minutes
   (d) CC Discharge at C/4, ending at 2.7V
   (e) Rest - 10 minutes
   (f) Repeat steps (b)-(e) for a total of 4 times
2. Full characterisation - Designated cycling temperature

   As above
3. Short characterisation
   (a) Rest - 30 minutes
   (b) CCCV Charge at C/2 to 4.2V, ending when I < 2mA

(c) Rest - 10 minutes

(d) CC Discharge at C/2, ending at 3.0V

(e) Rest - 10 minutes

(f) CCCV Charge at C/2 to 4.2V, ending when I < 2mA

(g) Rest - 10 minutes

(h) CC Discharge at 1C, ending at 3.0V

(i) Rest - 10 minutes

(j) CCCV Charge at C/4 to 3.8V, ending when I < 2mA

4. Cycling

   (a) CCCV Charge at C/2 to 4.2V, ending when I < 2mA

   (b) CC Discharge at designated rate, ending when 40% of nominal capacity has been discharged or at 2.7V, whichever occurs first

   (c) Repeat steps (a)-(b) for a total of 50 times

5. Repeat steps 3-4 for a total of 3 times

6. Repeat entire procedure

Characterisation Techniques

A wide variety of characterisation techniques were considered, and some were eliminated based on their reliability, relevance of output data and ease of use. For example, in early experiments, EIS was evaluated before being disregarded for the main body of experiments. This was due to its unreliable results in this scenario, producing results which differ by a significant margin on consecutive tests on the same cell. This variation was predominantly due to the unique manufacture of the cells, with the extreme overhang of the PE coupled by a long diffusion path, although part of the error will also be due to a lack of fixture correction, which is an inherent issue in the technique. In addition, the data is of limited use. While it can give some insight by either evaluating a Nyquist plot, or by fitting an ECN, these insights are often

tenuous, impacted significantly more by experimental minutiae than by the phenomena under investigation. Finally, the equipment required to perform EIS is limited, and would therefore have necessitated frequent transferal of equipment channels between experiments, which was logistically difficult, as well as a potential source of error in the data handling for such a large set of experiments. If the technique was expected to give useful insight, this would of course be an acceptable cost, but as its use was at best a shot in the dark, it was excluded from the final study.

GITT is another useful technique, and was considered. Full GITT analysis requires considerable time to perform, as pulses must be performed at a large number of steps over the device's SOC window, but with good analysis can grant significant insight into the device's performance over life, as well as indicate dominant degradation modes and facilitate model parameterisation. It was tested on these devices, but the results were once again unreliable. Significant work was done to increase the repeatability and accuracy of the data analysis techniques in an attempt to overcome this, but to little benefit. It is likely that the extreme overhang present in these cells due to the unorthodox electrode mismatch is once again the culprit, as the multi-variable curve fitting required is very sensitive to shifts in response time-constants. 1,2 and 3 RC pair ECMs were fitted, and again it was found that consecutive tests on a single cell often produced results that differed by a larger than acceptable margin. It was deemed that given the irregularity of the results achieved with proper GITT technique, as much insight could be gained by analysing the Ohmic resistances of the cells from the overpotential produced by any of the step changes in current present in the test protocol. As such, GITT was also disregarded from use in the full study.

It was learned that these cells appear to have a form of short term memory effect, where the same analysis would yield different results if it was preceded by cycling compared to being preceded by a characterisation technique. It is for this reason that the full characterisation protocol is repeated 4 times consecutively, as this yields significantly more reliable and illuminating data. In addition, consecutive results can be compared, and outliers discounted, or a mean taken. This was by far the most impactful alteration to the protocol with respect to reliable and trustworthy results. However, this increases the impact of slow techniques (such as EIS, GITT or HPCC) on

the time required to perform the experiments, as it is multiplied. As such, the above protocol was determined to yield the most useful results without increasing the experiments time and resource requirements beyond what was available. It comprises of capacity measurements taken at 1C, C/2 and C/4, performed at both a common temperature across all experiments (30°C), as well as at the designated cycling temperature of each experiment. Measurements at lower C-rates would have been desirable, but again produced high variability in output, so were not included. This protocol was deemed to give the most reliable and relevant measures of capacity and resistance over lifetime. In addition, the voltage response with respect to capacity can be investigated to give insight into specific degradation modes, either by using ICA and manually evaluating the plots in comparison with a degradation map, or through use of curve fitting for degradation mode analysis. Unfortunately, the latter technique was not considered until after the full set of experiments had begun, so beginning of life half cells were no longer an option. Instead, generic data of the electrode materials present was found in the literature and used to perform this analysis. Because of these limitations, as well as the difference between the DUT's and real world applications, any conclusions drawn from the data must be strictly comparative, i.e. "Cells tested at C-rate A exhibited behaviour X, while those tested at C-rate B exhibited behaviour Y".

3.3.6 Practical considerations

The final study required a total of 48 single layer cells and 24 pouch cells, and each experiment took >6 months to reach 1200 cycles. This required a large number of instrumentation channels for a long period of time, as well as frequent manual intervention (to change temperatures and load protocols). The first 6 experiments were begun at the beginning of 2020, with the intention of starting the remaining experiments as soon as equipment became available (either additional channels, or as the first round of experiments concluded). Unfortunately, this plan was impacted somewhat by a global pandemic, which caused both significant problems with maintaining the experiments already running, as well as beginning the follow up experiments. The first round of experiments were successfully performed, albeit with some interruptions due to power failure, miscommunications with other lab users having to perform some of the maintenance tasks and

over a slightly longer period than initially planned (due to limited access times for manual interventions). Due to circumstances outside the authors control, the remaining experiments were not completed until 2023. Unfortunately, the cells calendar aged significantly in this period. They were stored in fridges at 10°C to prevent this, however it is possible that this storage was interrupted (either by power failure, or by another lab user removing the cells from the fridges during lab maintenance tasks). This obviously has significant impact on the results, which will be evaluated in Section 4, but given the expense involved in the manufacture of these cells, the time and resources required to run this study, as well as the time available for the project, the only viable choice was to finish the experiments planned with the cells to hand, and evaluate the impact as best as possible, taking it into account for any conclusions drawn.

3.3.7 Selection of Experimental Conditions

As discussed above, the impact of only 2 variables is investigated in this study: Temperature, and C-rate in discharge only. Due to the reliability problems with common characterisation techniques, it was deemed necessary to have a minimum of 4 cells for any given set of conditions, the prudence of which will become clear when discussing the results. As such, the number of different conditions it was possible to investigate was limited. Using some kind of orthogonal array of conditions was considered, but it was feared that the complexity added by such an approach might obfuscate the results further. Instead, the initial plan consisted of a control experiment performed at 1C, 30°C, and other tests would vary either C-rate or temperature from this point (but not both). The first round of experiments followed this plan, with some cells remaining to either repeat problematic experiments or add additional conditions once the initial data had been analysed. The initial round consisted of 6 experiments, comprising C rates of 1,2,3, and 4 at 30°C, as well as 1C experiments at 10°C and 45°C. After this data had been analysed, it was deemed that a wider range of temperatures was necessary to draw meaningful conclusions, so experiments at 20°C and 45°C were added. It was also decided that an investigation of higher C-rates across a range of temperatures was warranted, so all temperatures selected were repeated with cells discharged at 4C. Table 3.4 shows which experimental conditions were used, and in which round of testing they were evaluated. Fewer pouch cells were available (30 in

total), so these experiments were performed only at 1C, all selected temperatures, and at 30°C, all selected C-rates, and only 3 cells per set of conditions.

3.3.8 Terminology and other Miscellany

Going forward, each set of experimental conditions shall be referred to as an "experiment" and each round of full characterisation shall be referred to as a "data point". Cells where given alphanumeric names based on their position in experimental stacks, for example, the bottom cell in the first stack was named "A1", the third from the bottom in the second stack would be "B3", and so on. Where possible, cells for a given experiment were spread over as many different positions across as many different stacks in order to help exclude any impact of variations in experimental conditions (while having the fewest possible stacks feasible). Some examples of these include variations in pressure across different stacks, differences in temperature across vertical positions in a stack due to differences in air speeds from fan dead-spots, or proximity to other cells which may impact thermal boundary conditions. This was done from an abundance of caution, and seemed like a prudent approach to take, but no variations or errors such as this were found. Cell names are used only for ease of reference and have no bearing on the conditions the cell was tested under. In addition, for ease of plotting, each experiment was given a colour based on each primary variable being changed having a single hue scale, and these scales were selected to be readable by those with deuteranopia, protanopia or tritanopia type colour-blindness. Table 3.4 shows the experimental conditions selected, as well as the round of testing in which each experiment was performed and the colour selected to represent that experiment in figures.

Table 3.4: A table showing the experimental conditions applied to single layer cells, in which round of testing they were performed, and the colour used to represent them in figures

C-rate/Temperature	10	20	30	45	55
1	1	2	1	1	2
2			1		
3			1		
4	2	2	1	2	2

3.4 Conclusion

In order to do empirical tests to quantify the decoupled contributions of C-rate and temperature to the degradation of a LIB, an experiment was designed that allows the practical elimination of thermal gradients within the device. This was achieved by the manufacture of custom cells consisting of only a single layer of active material, compared to the many typically found in a pouch cell. These cells were then mounted in an instrumented fixture, which held them in good thermal contact to large thermal masses (relative to the cell). These thermal masses were in the form of aluminium compression plates whose other surfaces were held at the desired temperature by forced convection (i.e. placed in a thermal chamber with a fan to aid circulation). These cell fixture units were then stacked, and the stack compressed using tied rods and Belleville washers, with the pressure instrumented using a set of load cells. In this way both temperature and pressure was controlled, although only temperature was a varied between experiments, with pressure applied uniformly across all tests at 150 kPa. The apparatus' ability to maintain the required temperature was validated with a "worst case scenario", consisting of a small stack filled with already degraded (and therefore higher resistance and heat generation) cells, which were cycled at 4C continuous current in both charge and discharged. The centre cell was thoroughly instrumented with thermocouples on both top and bottom surfaces, with measurements taken while the cells were electrically idle as well as under load. The average temperature rose by ¡0.81°C, less than the peak to peak variation of the thermal chamber. Additionally, the temperatures were well grouped, with the spread being well inside the error of the measurement

devices (k type thermocouples).

A degradation study was then designed to demonstrate the intended use of this rig, consisting of a number of cells cycled at a variety of temperatures and discharge rates. The cells were to be characterised periodically at a common temperature, allowing comparison of their performance loss throughout testing. Due to some slightly aberrant behaviour from the mismatched capacities of the electrodes, the characterisation could not include very low rate testing, as would be desirable for measuring their OCV for ICA. Additionally, for ease of comparison and to ensure that the loss of performance through life did not change the cycles experienced at the end of testing compared to the beginning, a fixed depth of discharge during cycling was selected (40%), which was applied using coulomb counting. The results of this study, as well as practical outcomes of the decisions made in designing the study, will be presented and evaluated in the following chapter.

Chapter 4

Decoupling the Impact of Discharge Rate and Temperature on the Degradation of Li-ion Cells - Results

This chapter outlines the data gathered using the experimental apparatus and methods described in the previous chapter. Succinctly, a fixture was created to mount custom manufactured Li-Ion cells consisting of a single layer of active material such that the entirety of the active material is held at a single temperature. This allows the impacts of temperature and C-rate on the degradation of the cells to be de-coupled, which is not possible with conventional cells and fixtures due to the thermal gradients generated by the internal heat generation of the device during operation. A degradation study was performed, with cells cycled at a variety of temperatures and discharge rates and characterisation performed throughout. The degradation was then analysed, initially from the capacity measurements, and then in further depth using ICA and DMA techniques.

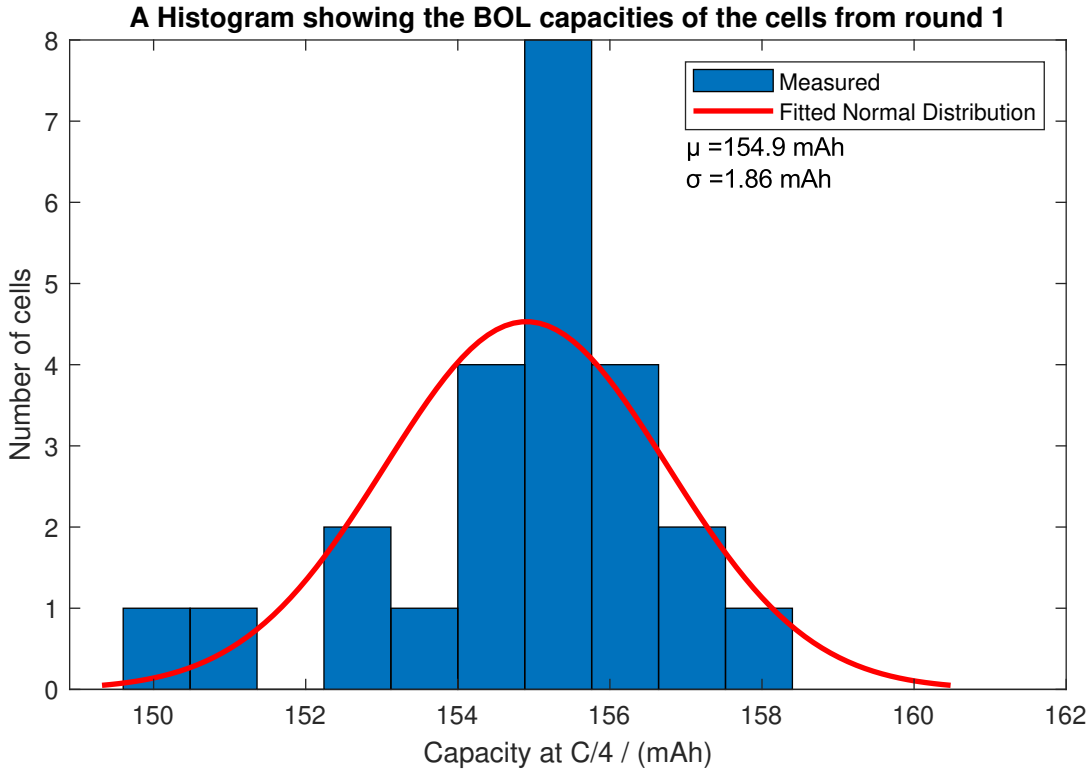


Figure 4.1: A histogram showing the BOL capacities of single layer cells tested in Round 1 of experiments

4.1 Capacity analysis

4.1.1 Beginning of life

The first thing to examine is the quality of the cells themselves. Their nominal capacity was 160 mAh, so the measured capacity (at C/4) was binned to evaluate cell to cell variability. A histogram of these results can be seen in Fig 4.1, including an overlay of a normal distribution fitted to the results. The cells have an average capacity of 154.9 mAh and a standard deviation of 1.86 mAh. This means that their standard deviation is just over 1.1% of their nominal capacity, and their mean is within 3.2% of the nominal value. This is better than expected, considering the hand made, small batch nature of the samples, and gives some confidence in the reliability of both the cells and the characterisation techniques applied.

Unfortunately, as mentioned in Section 3.3.6f, there was a significant delay between round 1 and round 2 of testing, and the cells calendar aged significantly in that time due to an unknown

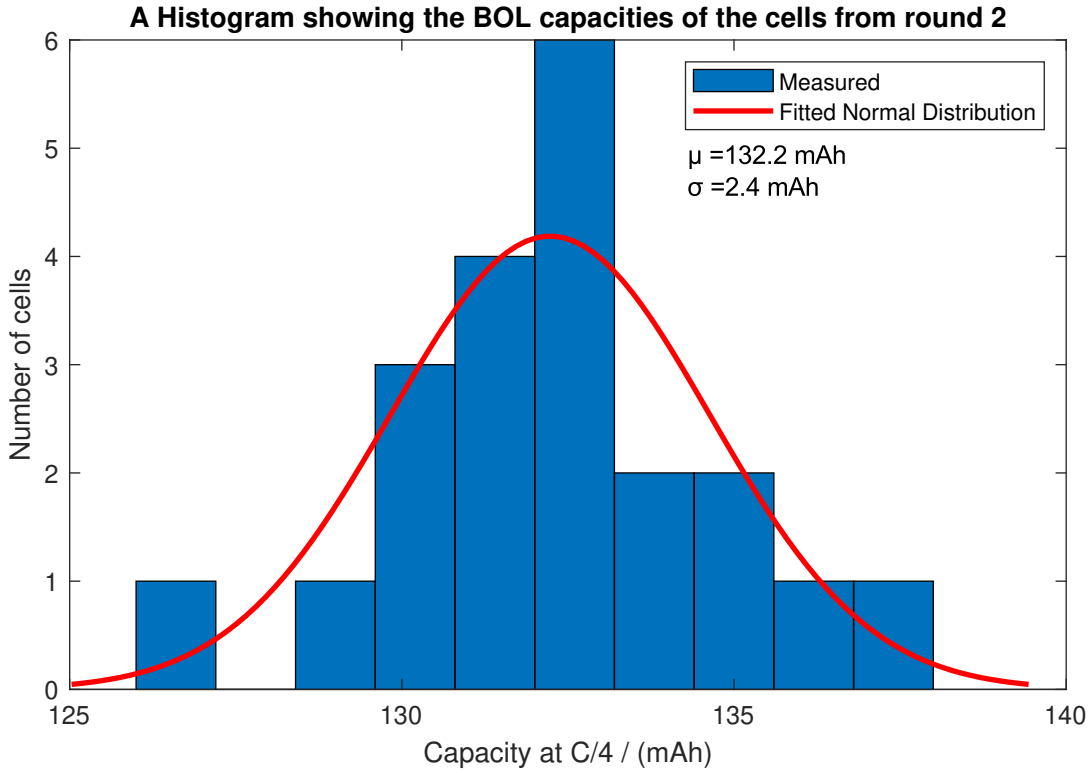


Figure 4.2: A histogram showing the BOL capacities of single layer cells tested in Round 2 of experiments

error in their storage. This can be seen by the capacities of the round 2 cells measured at the beginning of their testing in Fig 4.2. As is to be expected from calendar ageing, they have lost a significant amount of capacity (14% of their nominal compared with round 1, down to an average of 132.2 mAh), and have slightly more cell to cell variability, with a standard deviation of 2.4 mAh. This is still a respectably tight grouping however, and the manufacturer has clearly done well to produce such uniformity in custom made cells.

4.1.2 Round 1 Experiments

Now a starting baseline has been established, we can begin to evaluate the cells' performance throughout the degradation experiments. For ease of comparison, the capacity measurements are normalised for each cell by dividing them by that cell's capacity measured at the beginning of testing. The charge throughput each cell experienced since its delivery was calculated and divided by the nominal capacity of the cells to give a "Number of Equivalent Cycles" metric in

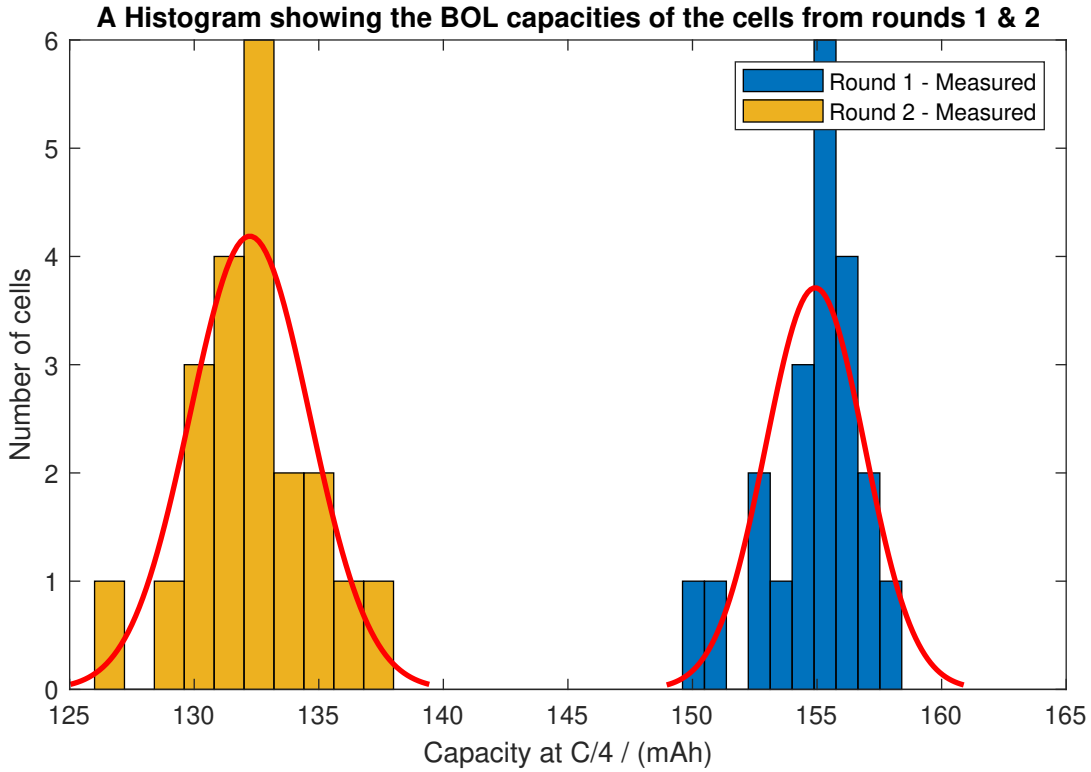


Figure 4.3: A histogram showing the BOL capacities of single layer cells tested in Rounds 1 and 2 of experiments

order to account for differences in testing due to cells tested at higher C rates hitting the lower voltage limits during cycling. This also accounts for the impact of any characterisation cycles done, including irregularities due to external factors. In addition, the capacity measurements are filtered for outliers by calculating a standard deviation for each data point, with measurements further than 3 standard deviations away from the mean value being excluded. These standard deviations (with outliers excluded) are included as error bars on all relevant figures.

To evaluate the reliability of the results prior to drawing any conclusions, we will begin by examining the capacity of all cells for a single experiment. As it is the centre of our matrix, and what might be considered "typical" conditions for a cell, we will look at cells cycled at 30°C and discharged at 1C. A plot of their Normalised Capacity over the testing can be seen in Fig 4.4. We can see that their grouping is reasonably tight, with a total standard deviation of 0.011 (representing 1.1% of nominal capacity). This trend is generally continued throughout the round 1 data, with some experiments having slightly more or less spread between the cells,

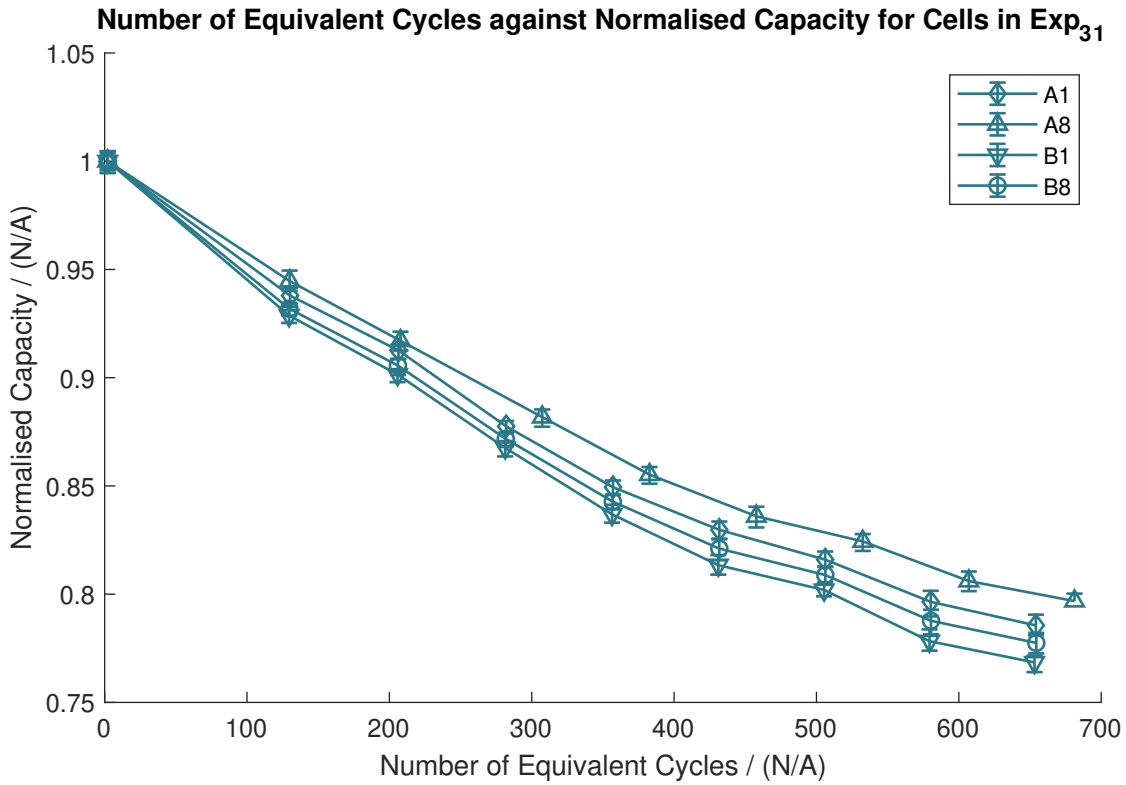


Figure 4.4: A plot showing the Normalised Capacity against Equivalent Cycles of cells tested at 30°C and discharged at 1C

but nothing significant. It can be seen that one cell varies more significantly than the others after 200 equivalent cycles, namely cell A8, including a significant shift on the x axis. This is due to a mistake in performing the testing, where some cycling was performed instead of the characterisation protocol. It is the only example of a mistake like this in the round 1 testing, hence its inclusion here.

In order to demonstrate that there are no significant cell to cell variations across the experiments performed in round 1 of testing Fig 4.5 presents the standard deviation of capacity measurements against the equivalent number of cycles for each experiment. Due to the way testing was setup, as well as the low cell to cell variability found in initial testing, it was not deemed necessary to bin the cells before assigning them to their experiments, so the standard deviations at the beginning of testing vary (as this was not controlled for). Even had it been logistically feasible, it would have been functionally difficult without putting significant cycles through the cells prior to the beginning of testing, and it is believed that any impacts due

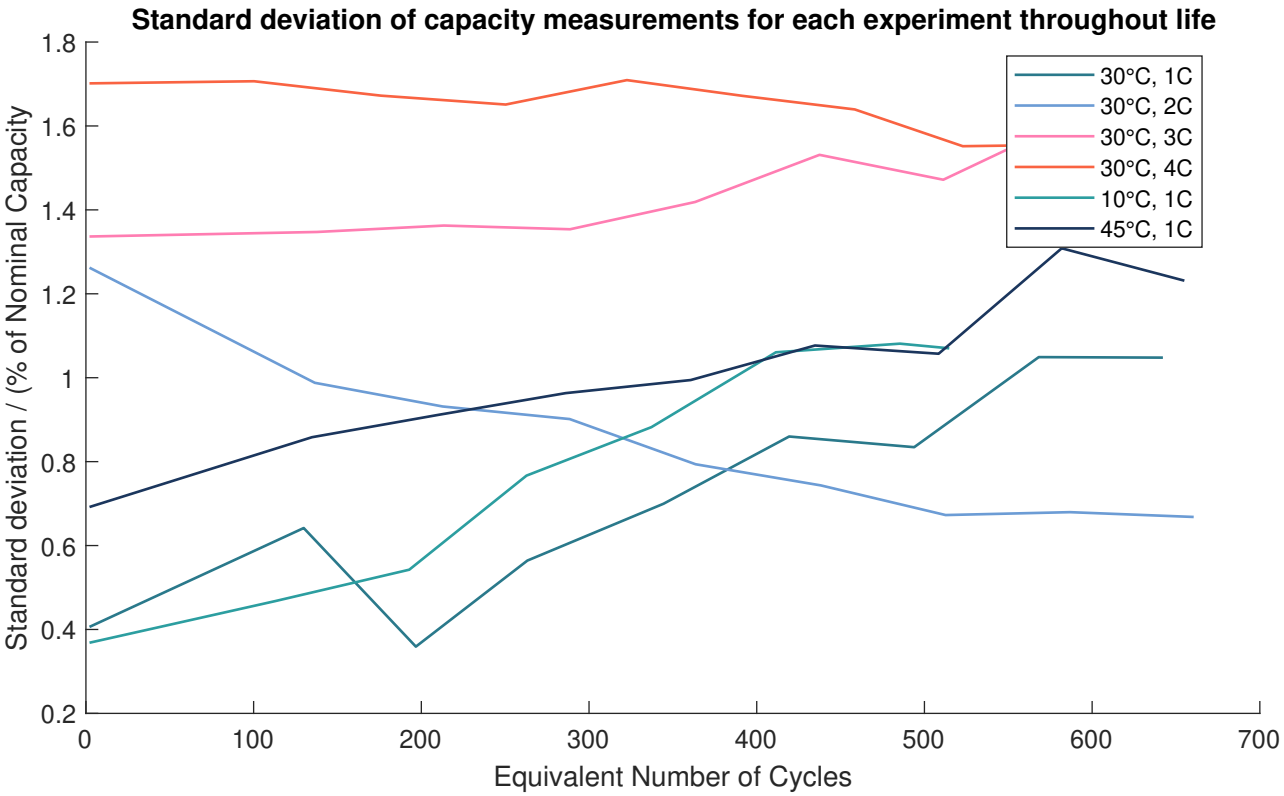


Figure 4.5: A plot showing the standard deviation of capacity measurements against equivalent number of cycles for each round 1 experiment

to the excessive handling, re-jigging and connection cycles required would likely have been as large as any potential benefit. The greatest variation between cells in an experiment at any point throughout testing is less than 1.8% of the nominal capacity, which is considered well within acceptable limits, and is present in an experiment where it changes very over the lifetime of the cells (i.e. those cells don't diverge so much as run parallel to each other, degrading at the same rates but with a minor constant offset from each other). Indeed, the largest increase in standard deviation over the course of testing (representing the greatest divergence) was exhibited by cells tested at 45°C, 1C, and is 0.7% of nominal capacity. This is again less than was expected, and well below what would raise doubts in the validity of the data.

Having examined the distributions of cells within their respective experiments, we can now begin to compare experiments with one another. Initially, we will examine the impact of changing discharge rate on capacity throughout life, as can be seen in Fig 4.6. Variations in the X direction between experiments are primarily caused by cells discharged at higher C-rates

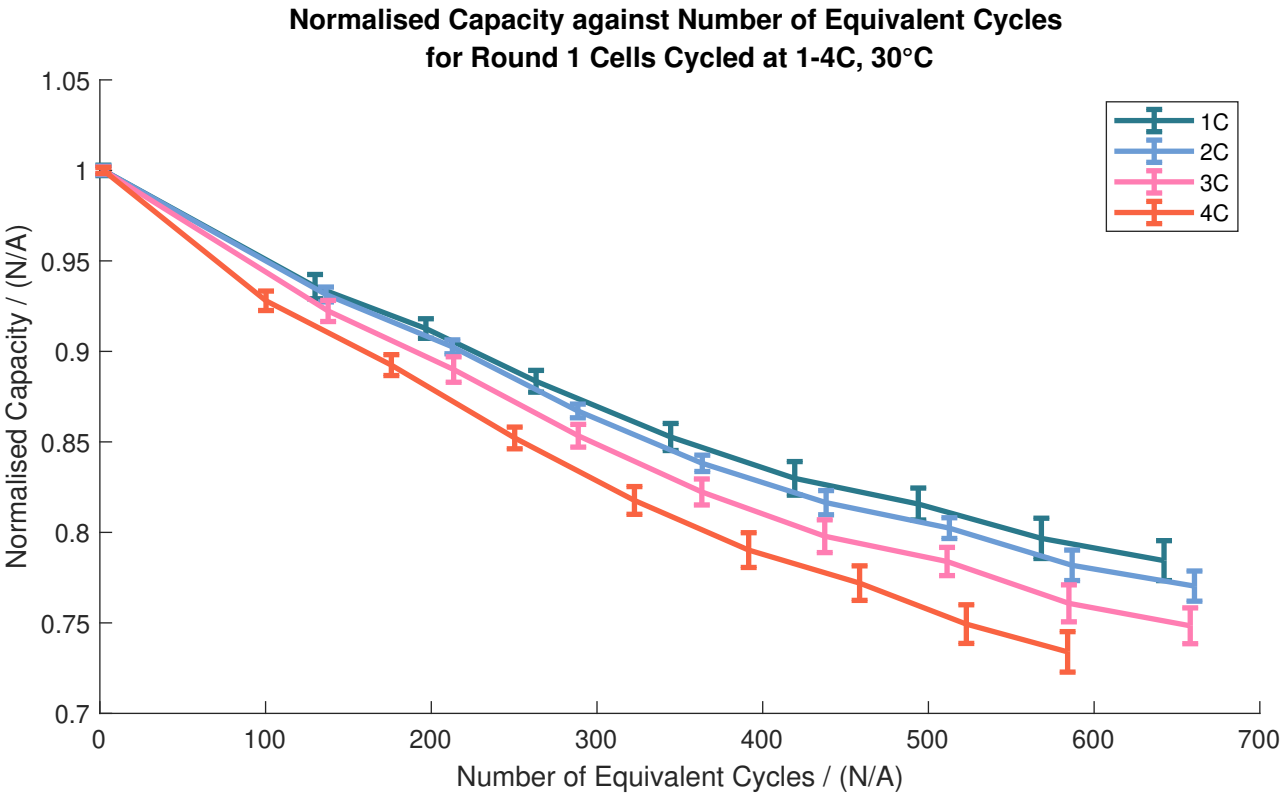


Figure 4.6: A plot showing the Normalised Capacity against Equivalent Cycles of cells tested at 30°C and discharged at a variety of C-rates

hitting the lower voltage limits during cycling, reducing their effective depth of discharge, and therefore charge throughput. What is interesting in this data is that the impact of discharge rate on degradation appears to be significantly non-linear, becoming significantly more damaging at higher rates. The impact is relatively small up to 2C, with approximately 1.2% of nominal capacity difference after 585 equivalent cycles between cells discharged at 1 and 2C, but is 2.1% between 2-3C and 2.7% from 3-4C. Greater resolution of C-rates would be interesting, and could be a topic for further work, but was beyond the resources of this project. This data does help demonstrate the impact of peaky load cycles, however, as the impact of those peaks is going to disproportionately impact degradation, even with a perfect thermal management system that eliminates all thermal gradients in active material. To further investigate this phenomenon, consider a load cycle that had 4C peaks for just 10% of its time, with the other 90% spent at 1C. This would mean that 40% of the capacity would be discharged at 4C, and so this data predicts that the cells would see as much degradation from that first 10% of time as the last 90%.

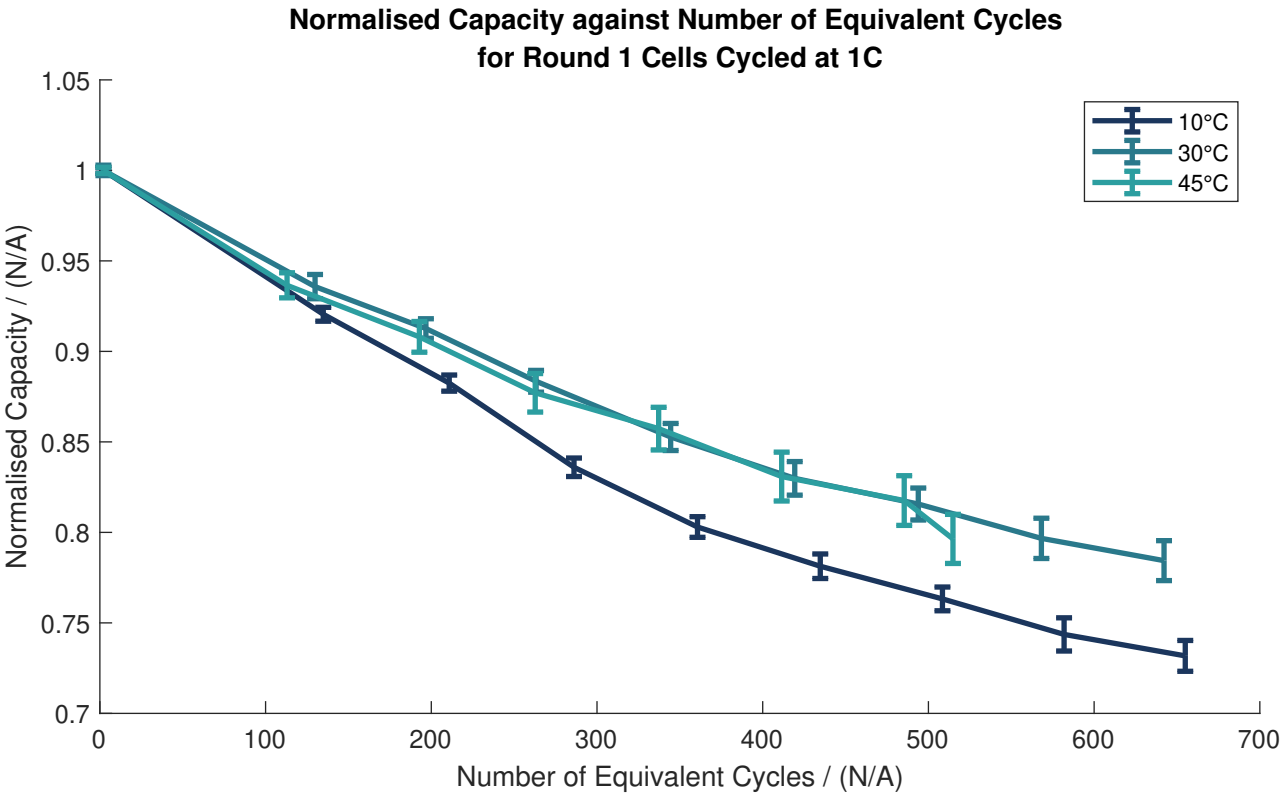


Figure 4.7: A plot showing the Normalised Capacity against Equivalent Cycles of cells discharged at 1C at different temperatures

The impact of temperature is more complex to examine, as the dominant degradation mode will shift more considerably across the temperatures tested than it does for C-rate. A more in depth analysis of this will be performed later, as it requires a variety of analysis techniques to yield any clear insights, but Fig 4.7 demonstrates how different the impact of temperature is on capacity loss than C-rate. It can be seen that there is no significant difference between cells cycled at 30°C and 45°C, indeed they are remarkably close together. However, cells cycled at 10°C saw much larger loss of capacity, with an increase of over 5.2% of the nominal capacity after 650 equivalent cycles. To put it in context of the total loss of capacity, the cells at 10°C saw 1.25 times the capacity loss of those at 30 and 45°C.

Indeed, if we wanted to compare the relative impacts of C-rate and temperature, we can plot all the round 1 experiments on the same graph, as in Fig 4.8.

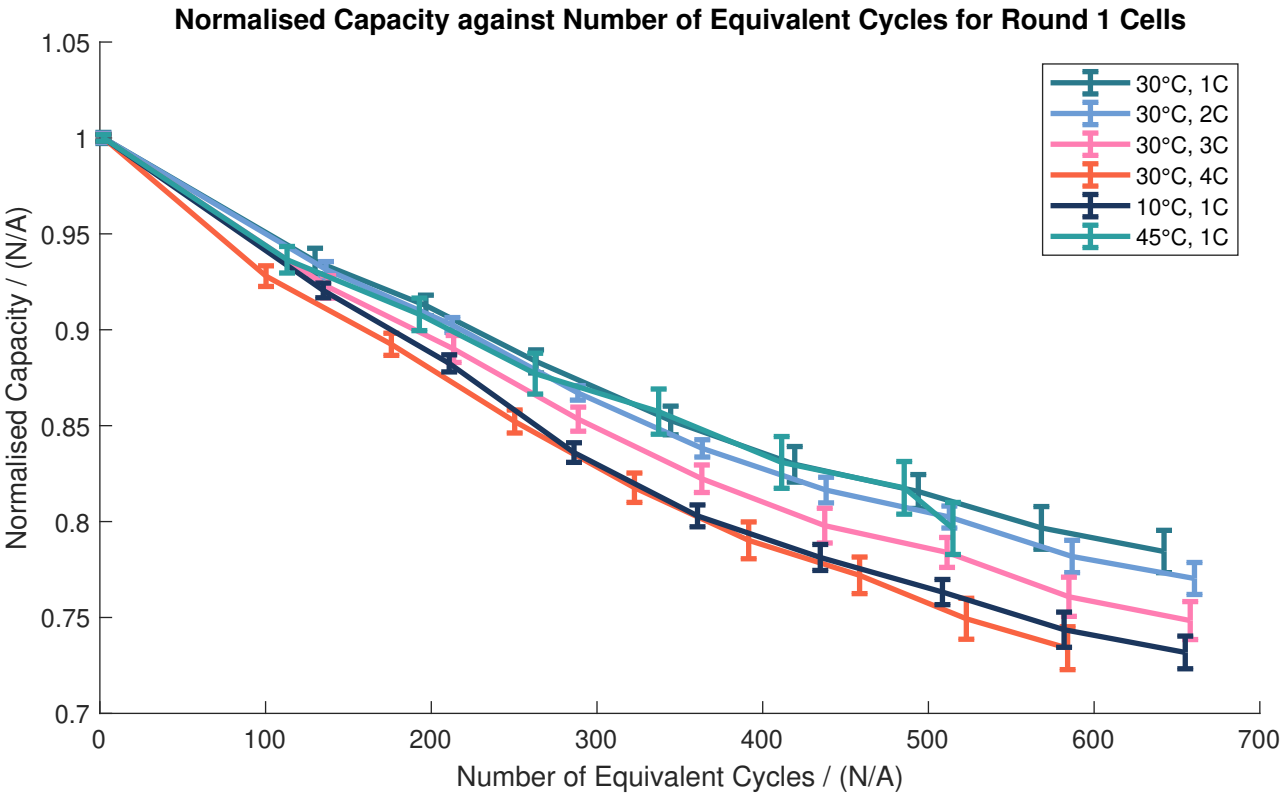


Figure 4.8: A plot showing the Normalised Capacity against Equivalent Cycles of all round 1 experiments

4.1.3 Round 2 Experiments

Initially, we shall examine these experiments only relative to each other, and a comparison to round 1 experiments shall be performed later. We only have 2 discharge rates represented here, 1C and 4C, so it is not possible to analyse how the impact of increasing C-rate changes, but we can look at how the difference between 1C and 4C changes with temperature. At both 20°C and 55°C, we have round 2 experiments at both 1C and 4C discharge rates, the results of which can be seen in Fig 4.9. The difference between capacity loss at 20°C, 1 and 4C discharge is large, similar to that seen at 30°C. Direct comparison is not presented here because the calendar ageing prior to these experiments causes difficulty, but there is certainly a marked difference between the impact of 1C and 4C discharge rates in both. However, the same discharge rates applied at 55°C have far less difference between them. Again, a greater number of experiments at a larger variety of discharge rates would probably be very illuminating here, but was not possible. It is expected that what would be shown is that as the temperature increases the C-rate at which significantly higher capacity loss is exhibited also increases. In this data we are seeing that the difference

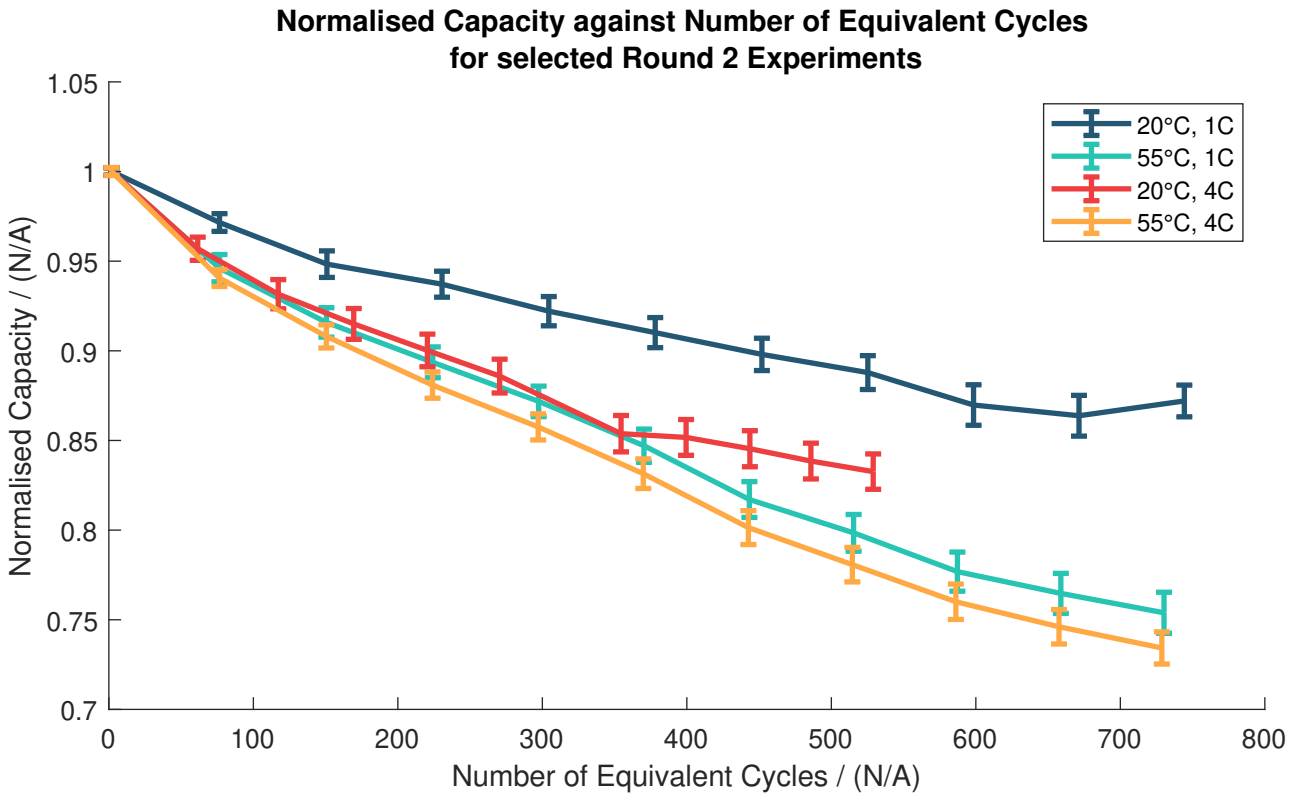


Figure 4.9: A plot showing the Normalised Capacity against Equivalent Cycles of selected round 2 experiments

between 4C and 1C at 55°C is comparable to the difference between 2C and 1C at 30°C. This is because the dominant degradation mode present here is likely particle cracking,[8] and at higher temperatures the diffusion gradients produced by increasing C-rates are significantly lessened, reducing the mechanical stress produced in the active material, and therefore the rate at which active material is lost. The total capacity loss is still greater, however, because while loss of active material might be reduced, the rate of SEI layer growth is increased, and this begins to dominate.

One of the purposes for adding new experiments in round 2 was to investigate the impact of temperature at higher discharge rates, hence the addition of experiments at 4C across all selected temperatures. The results from these can be seen in Fig 4.10, and it is here where the capacity results on their own become impossible to interpret with any confidence. Cells cycled at 55°C and 45°C behave somewhat as me might expect, with relatively smooth capacity loss over the duration of the experiments, with the higher temperatures degrading faster. The difference between the two lines is maybe a little greater than was initially expected, but not

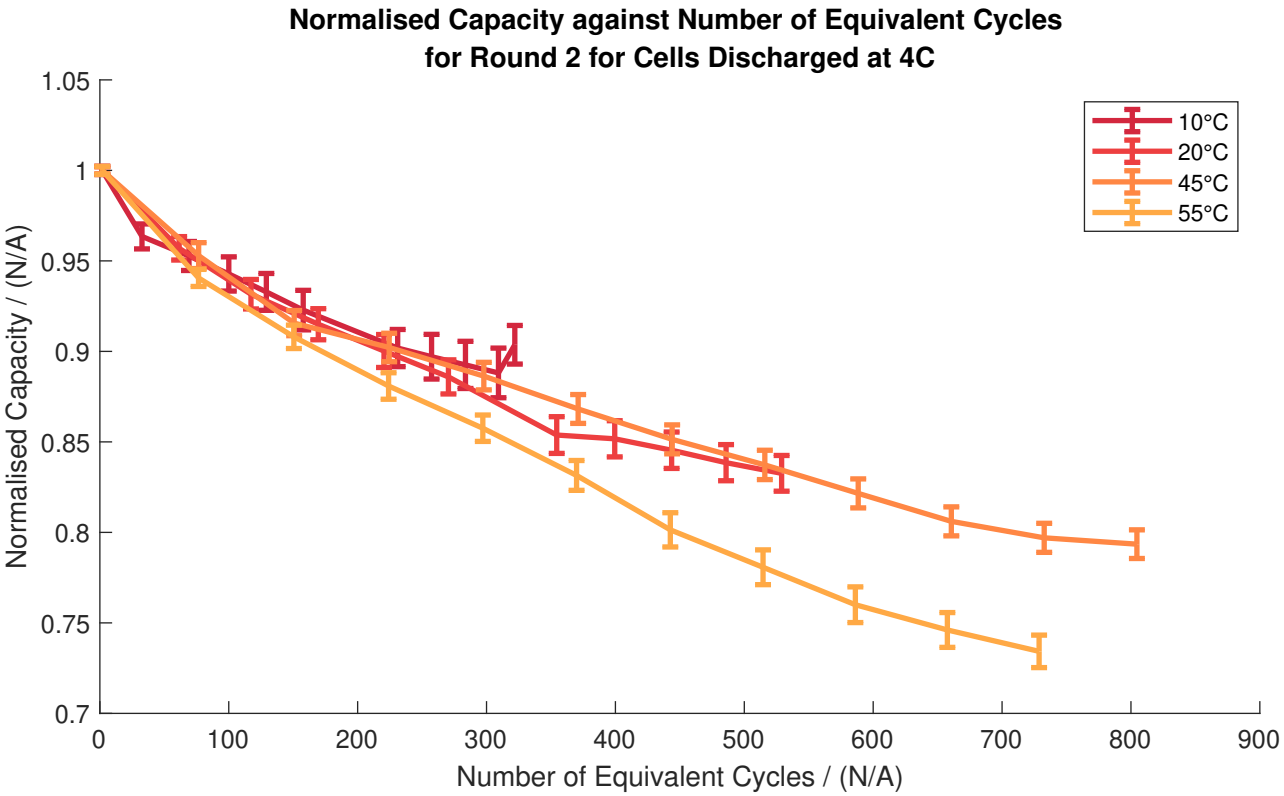


Figure 4.10: A plot showing the Normalised Capacity against Equivalent Cycles of round 2 experiments discharged at 4C

excessively. We might expect the cells at 20°C to exhibit less capacity loss than those at 45°C, but this does not appear to be the case.

If we look at the standard deviations for the round 2 experiments, we can see that they are a little more erratic in places, but all still well within acceptable thresholds. The 4C experiments appear to converge somewhat more than other experiments, perhaps due to them hitting the lower voltage cutoffs, limiting the throughput (and therefore degradation) of the worst performing cells. This is the drawback of using fixed capacity cycling for degradation experiments, especially as the initial depth of discharge in this case was chosen specifically knowing that the higher C-rates would hit the lower voltage limits at some point during cycling. This is expected to be the reason the lower temperature cells discharged at 4C converge. Unfortunately, Fig 4.11 does not shed much light on the peculiarities of Fig 4.10, but there are other forms of analysis that might.

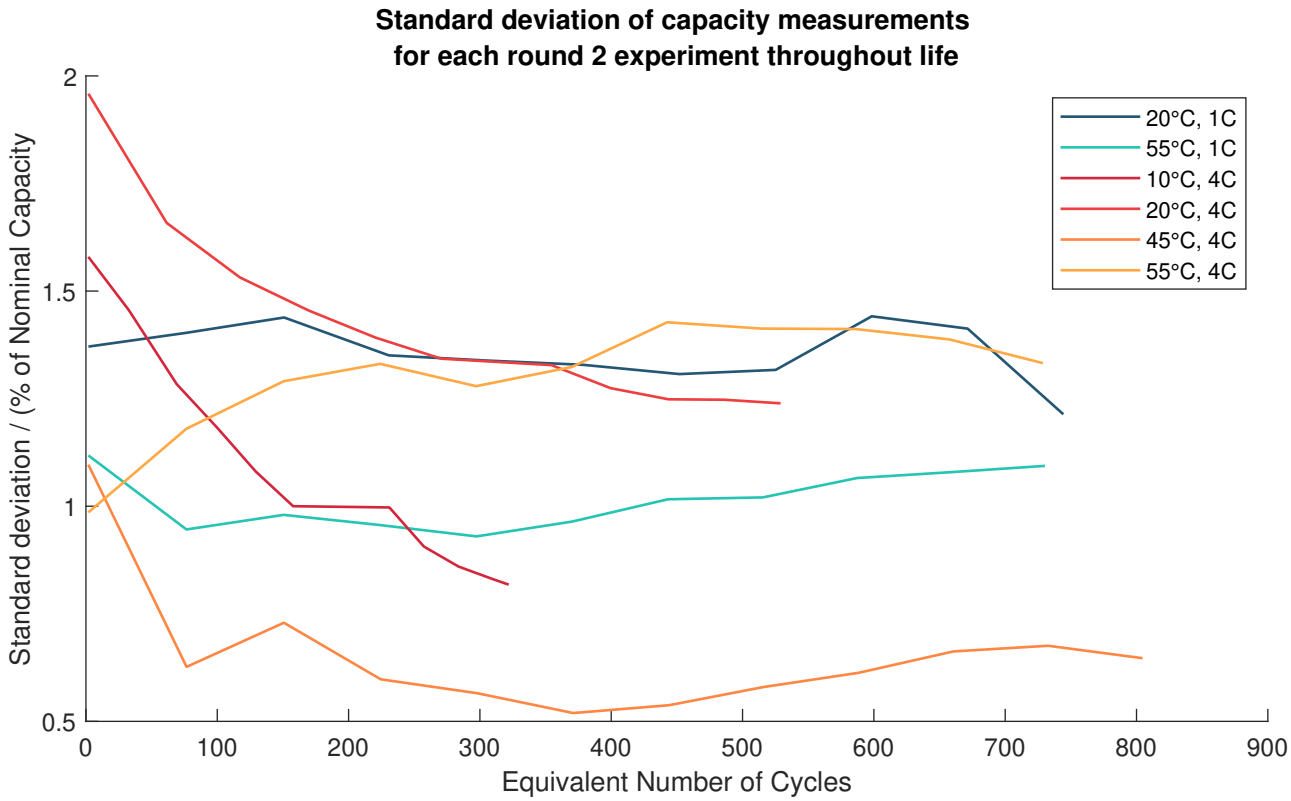


Figure 4.11: A plot showing the standard deviation of capacity measurements against equivalent number of cycles for each round 2 experiment

4.2 End of cycle voltages

As mentioned before, due to the nature of these cells many standard characterisation techniques do not yield reliable results. This is particularly true of those attempting to measure impedance, such as EIS and GITT. As such, no specific impedance values are presented in this report, as none were found with a degree of confidence that could be used for analysis. However, one value which is tied to the impedance is the voltage reached at the end of each cycle. This will be a function of both capacity loss and steady state resistance, but is of particular use in finding and explaining oddities in the results. For example, when examining the end of cycle voltages for round 2 (as presented in Fig 4.12, we can see that the lower temperature (10,20°C), higher C-rate cells did indeed hit the lower voltage limit during cycling, and in fact did so from the very beginning of the tests. We can also see that an error was made in testing the cells at 55°C, 4C, as they were clearly cycled at 1C instead of 4C for one period, hence the jump up to match the cells at 55°C, 1C. This doesn't appear to have had too significant an impact on the capacity loss, so it will be noted, but is not reason to discount the remaining data from that

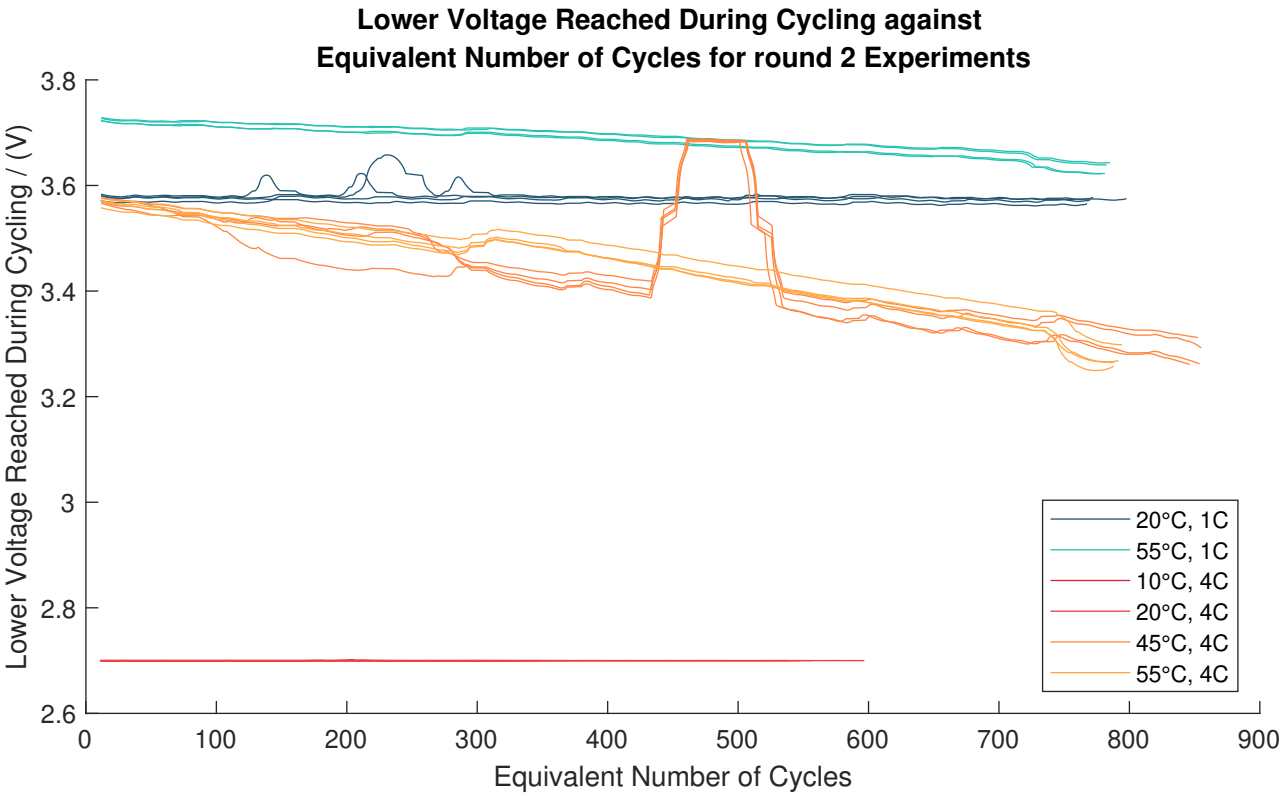


Figure 4.12: A plot showing the lower voltage reached during cycling against equivalent number of cycles for each round 2 cell

experiment. This is somewhat to be expected, as at the higher temperatures we see significantly less impact on degradation form increased discharge rate (due to the lower resistances produced by elevated temperature), and instead are expecting degradation to be dominated by SEI growth, which is primarily a function of charge throughput and temperature, not C-rate. What is interesting in this data is that cells cycled at 20°C, 1C show no significant change in end of cycle voltage, despite an approximately 10% of nominal capacity lost from the cell's total capacity.

Fig 4.13 shows the same data, but for round 1 experiments. Here we can see why 4C was selected as the maximum discharge rate, as only the cells discharged at 4C hit the lower voltage limits during cycling. Two of the cells do hit it much earlier than expected, which is of note, as from the capacity measurements (which can be seen in Fig 4.14) there is no significant discernible indicator that half the cells are performing significantly differently. For reference, the cells which hit the lower voltage limits significantly earlier are cells A4 and B4, which do exhibit lower capacity than the others, but not significantly so, and do not appear to exhibit significantly

higher capacity loss, but rather constant offset in capacity, including at beginning of life. As has been mentioned above, this may in part be due to them hitting that lower voltage limit and therefore seeing less degradation due to the cycles being shorter. Returning our attention to Fig 4.13, we can see that all the other cells respond as expected, with the highest end of cycle voltages being exhibited at 45°C, the gap between 10°C and 30°C being significantly larger than that of 30-45°C, and (with the exception of the 4C cells), C rate having an approximately linear impact on end of cycle voltage throughout life. 4C are the exception in this last statement because the damage due to particle cracking is so severe that it has a much more drastic impact on impedance, which also explains the significantly greater impact it has on capacity loss. We can begin to see a touch of this in the 3C data, which is significantly more varied than any at lower C rates, demonstrating that the hallmarks of different drivers of degradation modes are distinct, with temperature driven modes (such as SEI layer growth) exhibiting more uniform and predictable results, while discharge rate, which drives degradation modes such as loss of active material due to particle cracking, produces more erratic loss of performance. This is primarily due to the difference in scales at which these degradation modes occur, with the SEI layers physical features being orders of magnitude smaller than those of LAM due to particle cracking, and so producing less granular impact on the device as a whole.

4.3 Incremental Capacity Analysis

Incremental Capacity Analysis (ICA) can yield much more detail about the condition of a cell than a single measurement of capacity or resistance can. However, it is difficult to translate this into quantitative results based purely on a plot of dQ/dV against capacity. In order to do so with any level of rigour and certainty, one would require a comprehensive degradation map, consisting of plots of the cell at beginning of life, as well as at multiple steps along each possible degradation mode, ideally as isolated as possible and confirmed with post mortem analyses. This is obviously not available for these cells, but we can assemble one that is sufficient to allow a quantitative analysis of the curves from the literature. An example of one of these maps

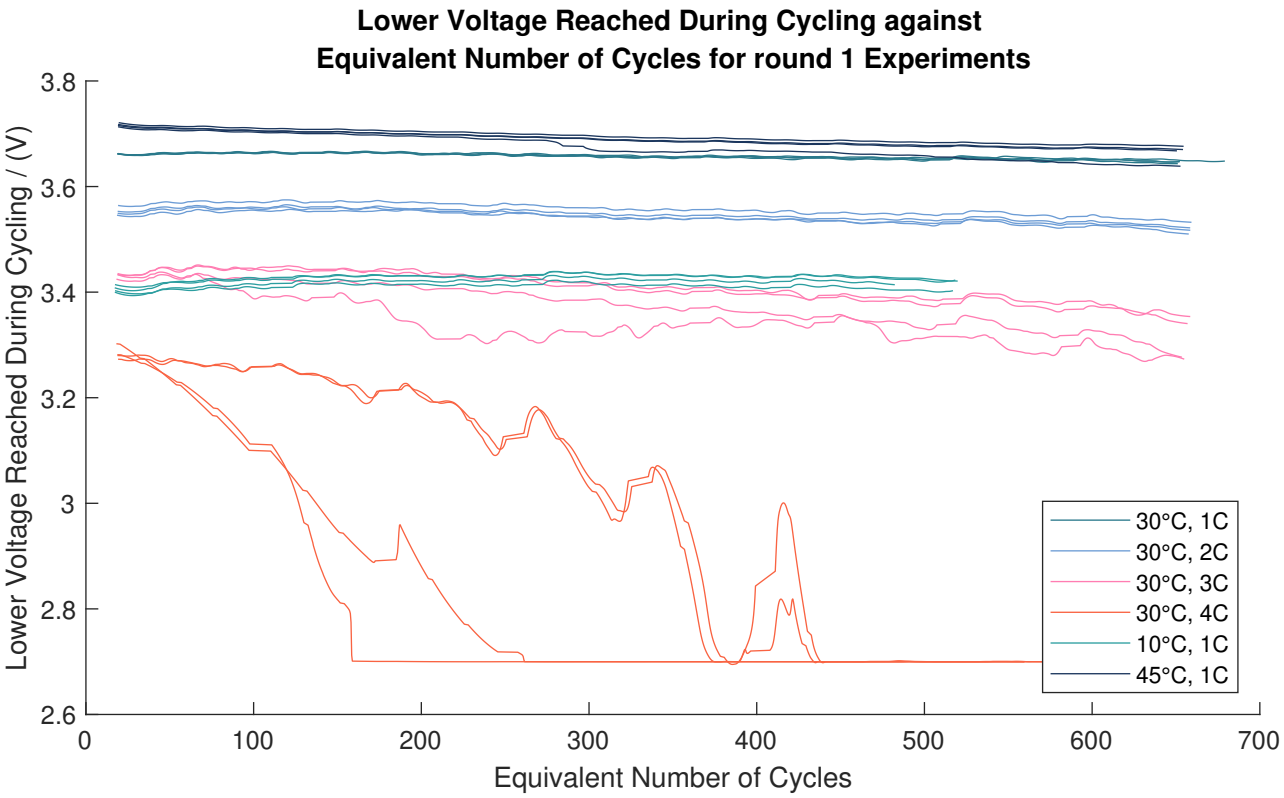


Figure 4.13: A plot showing the lower voltage reached during cycling against equivalent number of cycles for each round 1 cell

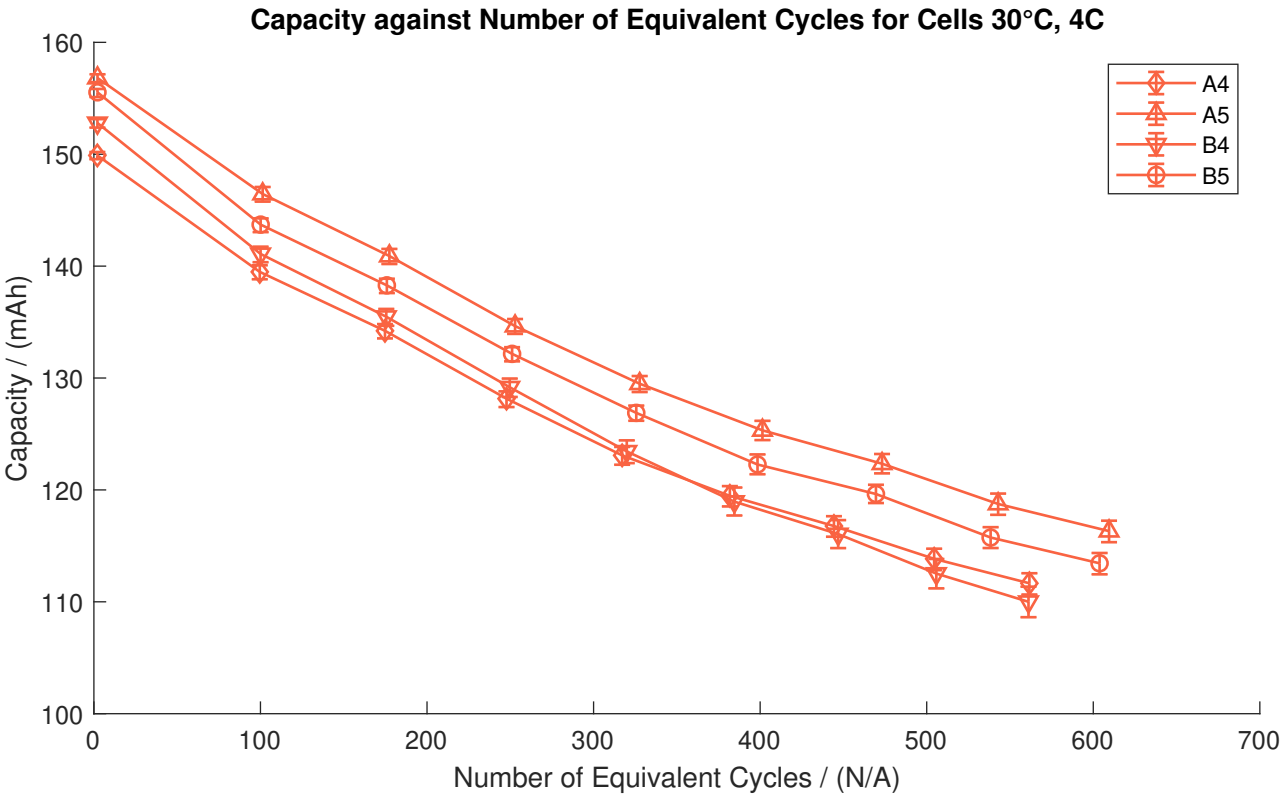


Figure 4.14: A plot showing the measured capacity against equivalent cycles of cells tested at 30°C and discharged at 4C

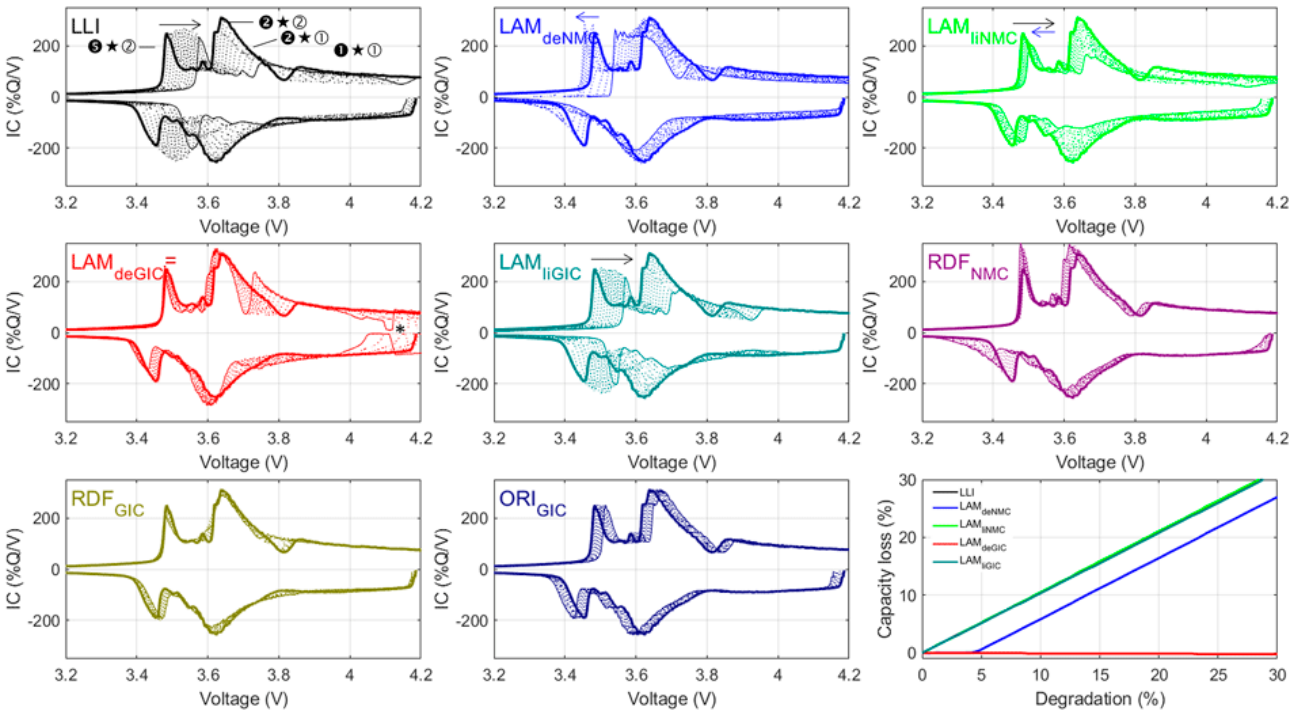


Figure 4.15: A map of the major degradation modes impact on the ICA curve of an NMC 111/GIC cell[7]

can be seen in Fig 4.15.[7] It is for NMC111, as opposed to the tested cells NMC 532, but the chemistries are very similar, and produce very similar results. NMC 811 varies much more, with a more distinct spread of peaks, but in practice the difference between cells marked as identical chemistries, but manufactured by different entities will have as much if not more difference between their ICA curves as will be seen between the data presented here and this degradation map.

An additional difficulty has already been mentioned multiple times, which is the unreliable results yielded by these cells at lower rates. This obviously presents a challenge for ICA, as it is typically performed at very slow C rates (typically <C/25). The slowest rate at which these cells were found to yield reliable results was C/4. While this will definitely impact the results, it has already been established that the data will not yield quantitative results with sufficient rigour, and well applied smoothing algorithms (such as a Savitzky–Golay filter) will produce results with discernible qualitative results. Dubarry *et al*[7, 129] find that while measurements taken at rates lower than C/10 (and ideally C/25) are best practice using this technique, higher rates can be used provided that the sampling rate is sufficiently high to capture the required features once

the data has been smoothed and differentiated (typically steps of 2 mV is more than sufficient). Data gathered at these rates will not have the accuracy required to pull quantifiable results from, but qualitative interpretation of the curves can still be illuminating, especially when combined with other techniques, as is presented here. The data gathered during these cycles is well over the threshold of 2000 points per regime suggested by Dubarry *et al*,[7, 129] and ICA is used here only to lend credence to theories produced by the more rigorous analysis techniques presented in this chapter.

We will begin with an example of a cell cycled at 30°C, 1C, as that is our reference point to compare other conditions to. In Fig 4.16, it can be seen that a reasonable output was obtained, although it lacks some of the fine detail and definition data produced on commercial cells with best practices can yield. The broad peak is typical of most NMC chemistries, with 2 subpeaks sometimes present in the charge direction (plotted on the positive Y axis of Fig 4.16). As is the long exponential tail that transitions this peak into the plateau present at the higher voltages, most identifiable at the beginning of life.[137, 138] The lower voltage of these two peaks is somewhat diminished, similar to the data presented by Noh *et al*. More importantly, these peaks shift continuously over the lifetime of the cell, and it is from these changes we can glean the most insight. Beginning in the charge direction, the major peak, centred around 3.7V, does not shift considerably in voltage over the testing, with a very minor shift to higher voltages discernible, more so in some cells than others. This peak does show the most dramatic change in height, however, and it is clear that this is where the bulk of the capacity is lost. This is most identifiable in the charge direction, because in the discharge the 2 peaks are more combined, masking the fall of one with the growth of the other. There is a general trend towards 0 in the discharge direction, as one would expect, the cell is, after all, loosing total capacity, therefore the area under this curve must shrink, but the tip of the peak falls from around 335 mAh V^{-1} to 215 mAh V^{-1} (i.e., within error of 2/3rds of its original value), compared to the charge directions peak falling from 385 mAh V^{-1} to 190 mAh V^{-1}, less than half of its original value. The difference can be seen in the second peak, which begins as a mere inflection point in the decay towards the lower voltages, around 210 mAh V^{-1} at 3.58V, shifting to the right and rising

to a clear maxima at 250 mAh V^{-1}, 3.62V approximately halfway through the testing. This peak retains its definition throughout the remainder of the testing, but does fall to 225 mAh V^{-1} by the end, which is still notably greater than its apparent size at the BOL.

Based on our degradation map, we can see that this behaviour can be explained with a combination of LLI, LAM of the negative, GIC electrode doing the bulk of the changes, and some initial degradation of the kinetics in the positive NMC electrode causing the initial rise and concentration of the peak at lower voltages. This is consistent with our expectations, as the dominant degradation mode expected here is to be particle cracking causing LAM, albeit not as strongly dominant as at higher discharge rates. It does tell us that much of the active material is lost in its lithiated state, therefore increasing capacity loss by compounding LLI with LAM at the negative electrode. It is difficult to discern the difference between this and loss of lithiated PE active material, but this is not believed to be the case due to the extreme overhang of this electrode thanks to the unique manufacture of these cells, as well as the fact that the higher voltage peak does shift slightly to the right, which is consistent with NE LAM, whereas in PE LAM it would be expected to shift slightly to the left. Of course some LAM will be occurring in the cell, and perhaps this has contributed to a muted voltage shift form this peak, but given there is twice as much PE active material as NE at the BOL, there would have to be extreme LAM to produce the capacity loss exhibited. This would likely manifest differently, producing a more sudden increase in capacity loss at some point in the testing, marking the point where the overhang had been consumed by this LAM and the total capacity of the PE began to fall below that of the NE. This is not present in the capacity measurements, or in the rate at which the ICA curves shift. These features and conclusions are consistent across all cells in this experiment, with no significant deviations identified. Only selected data is presented here for brevity and ease of communication.

Having established our reference point, and examined its likely implications for the degradation modes present in the testing, we will turn our attention to cells cycled at higher discharge rates, but the same temperature. Fig 4.17 is of a cell discharged at 4C, 30°C, which was also performed in round 1 of testing. We can see that the BOL shape is consistent with that seen in

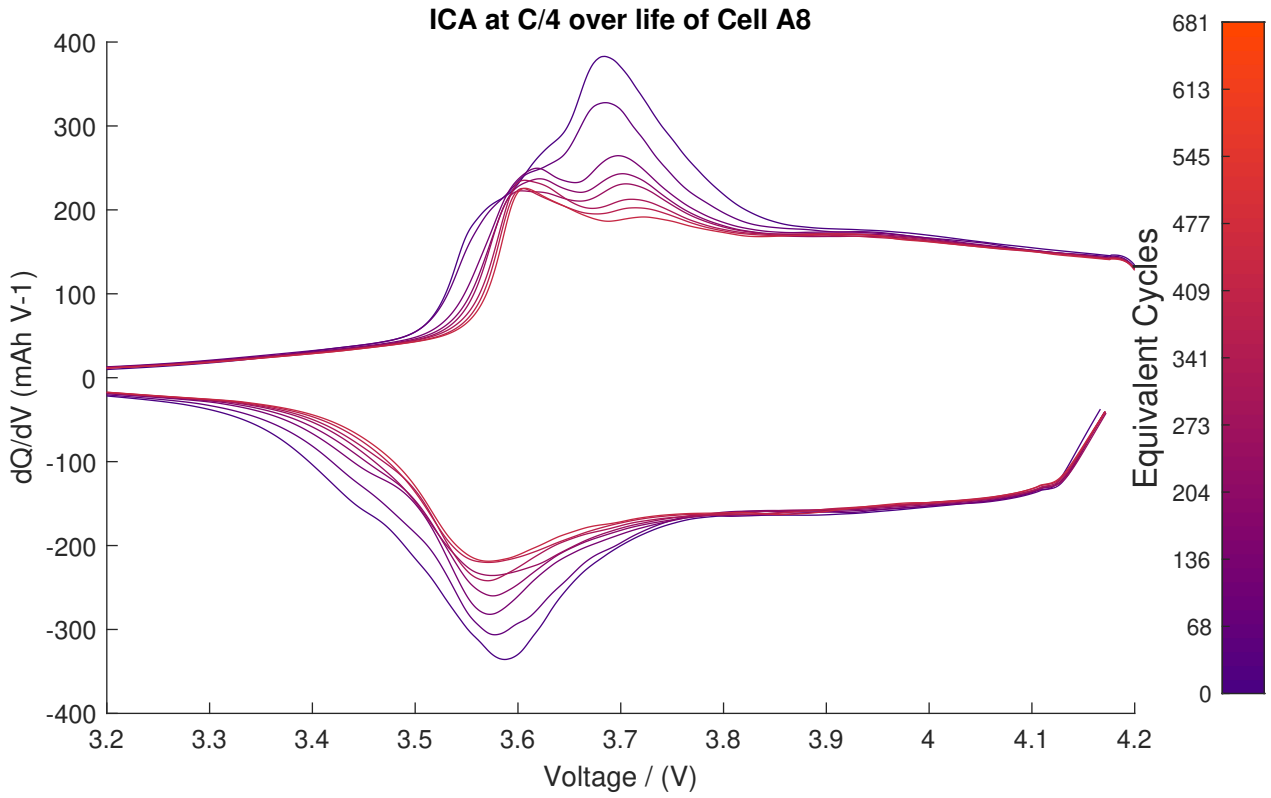


Figure 4.16: A plot of the ICA results of a cell cycled at 30°C, 1C discharge

Fig 4.16, giving us confidence in the reliability of these results, despite the unorthodox testing practices required. In general, the behaviour of the curve throughout lifetime is as expected, following the same trend as that of Fig 4.16, but with greater total loss over the course of testing. It should be noted here that the colour gradient is not identical from figure to figure, instead being fitted to the total range experienced by the individual cell being plotted, as the focus was readability rather than uniformity. As such, the cell in Fig 4.17 experienced nearly 100 fewer equivalent cycles over the course of testing than the cell in Fig 4.16, and its final trace should therefore be compared more with the penultimate trace of the cell discharged at 1C.

It can be difficult to discern without overlaying the plots (which has been avoided here as without interactivity it serves more to obfuscate than illuminate), but the initial rise of the peak at lower voltages in the charge direction is slightly more accentuated at the higher discharge rates. This lends further credence to the idea that this features shift is due to the degradation of the PE electrode kinetics, as this would be expected to respond to increased currents, regardless of the extreme overhang present, unlike LAM for the same electrode.

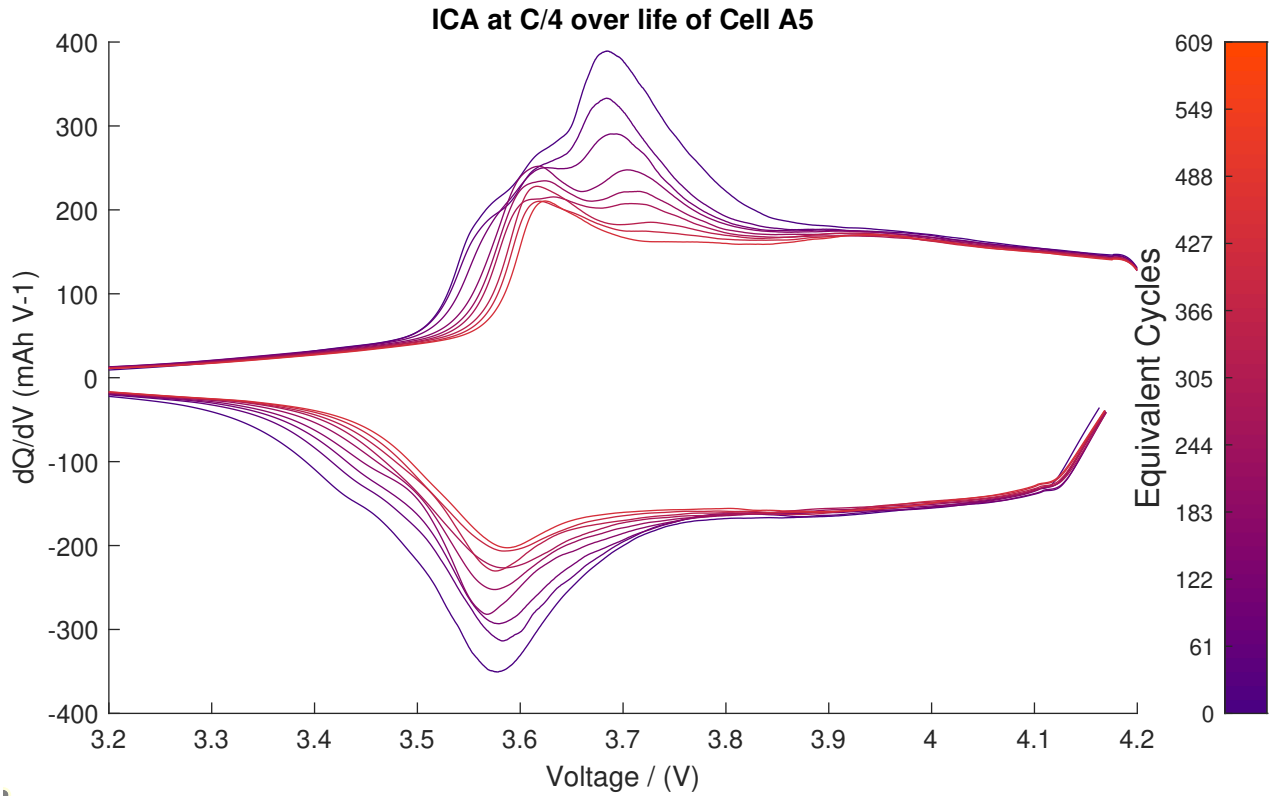


Figure 4.17: A plot of the ICA results of a cell cycled at 30°C, 4C discharge

These observations are carried out by all the cells tested at 30°C, with the trend following that of the capacity loss with respect to increasing discharge rate. Namely, there is little impact of increasing to 2C, but increasing difference between 2-3 and 3-4C.

We can next look at the different temperatures tested in round 1. Fig 4.18 shows the ICA results from a cell cycled at 45°C, 1C discharge. The results are largely similar to those at 30°C, as we would expect based on the capacity measurements presented above, but with one key difference. The sharpening of the lower voltage peak present in charge, which has here been attributed to changes in rate (i.e. power) capabilities, is much more accentuated at the higher temperature. It is difficult to draw much in the way of conclusions based on this observation, but it is of note, and combined with some of the other analysis techniques, may become illuminating.

Looking at cells cycled at 10°C, we would expect a more dramatic change in the ICA results. This is because lithium plating of the NE is expected to occur, which will be exhibited as a

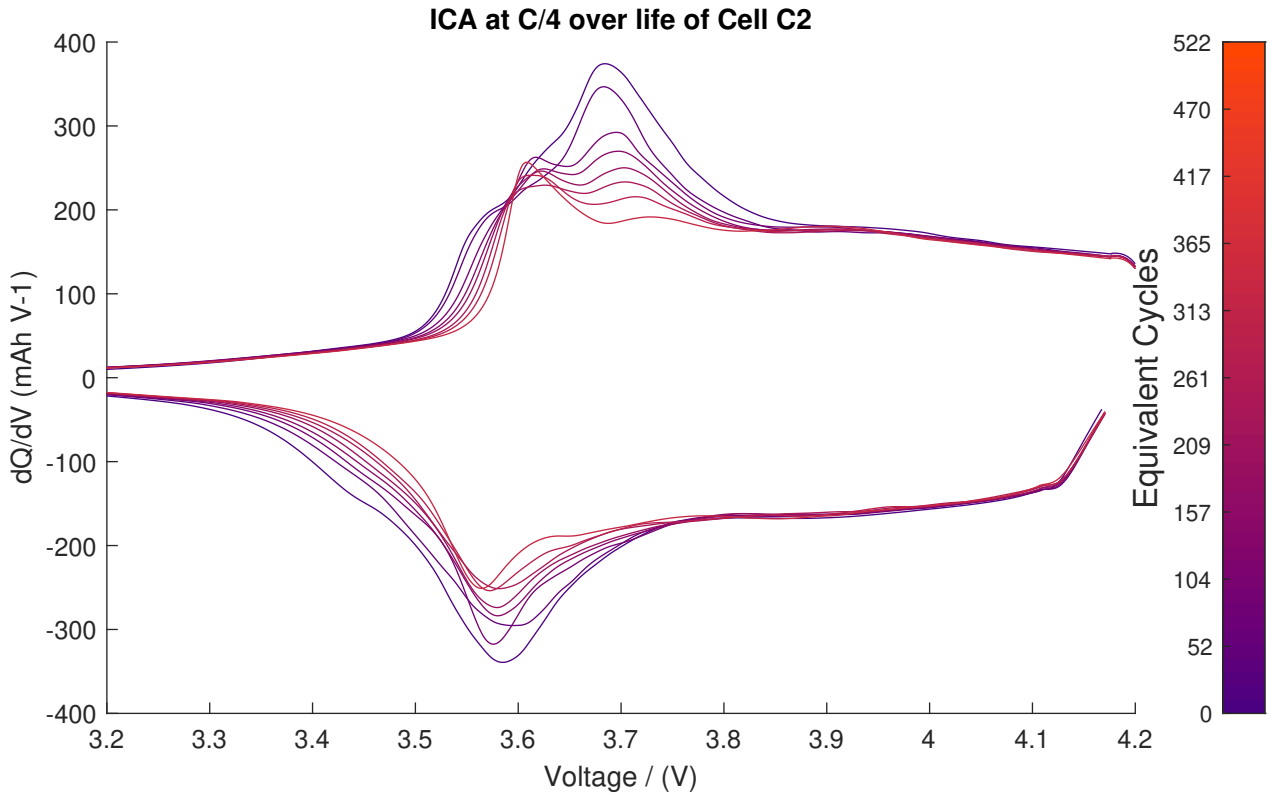


Figure 4.18: A plot of the ICA results of a cell cycled at 45°C, 1C discharge

large LLI, as well as the fact that this experiment (in concert with the 4C experiment) produced the highest capacity loss. Upon examining the data, an example of which is presented in Fig 4.19, it is fairly similar to other 1C experiments in trend, if not in severity. There is a much more dramatic loss of the major peak, so much so that the higher voltage peak in the charge direction is lost completely by the end of the testing, even becoming a local minima which rises gently to meet the plateau between 3.9 and 4.2V. This is the accelerated LLI we expected, and also shows its mark in the slightly increased voltage shift exhibited in the lower voltage peak in the charge direction. Once again, that peak also sharpens, however it is difficult to determine with any certainty if there is LAM in the NE also, as was found in the other round 1 experiments.

Now we must begin to examine the results from the second round of experiments. Before we look at the evolution of the cells over the course of testing, we must first evaluate the impact the unintended calendar ageing prior to testing has had. To this end, Fig 4.20 displays the ICA results of the first characterisation of each round 2 cell (coloured according to their experiments designated colour), overlaid with an average of the same results from all round 1 cells. We can

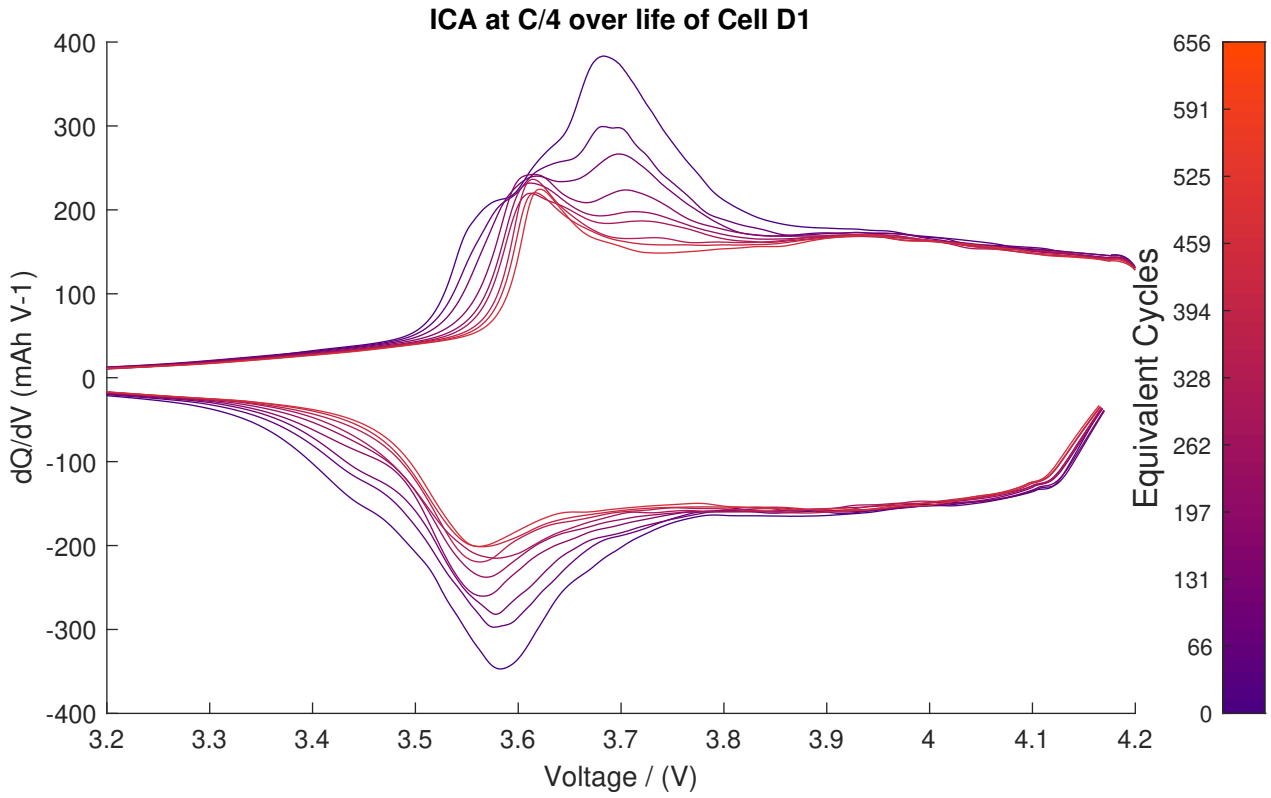


Figure 4.19: A plot of the ICA results of a cell cycled at 10°C, 1C discharge

see that significant degradation has occurred, primarily in the form of LLI, with an even more pronounced sharpening of the lower voltage peak in charge. This indicates that our assumption that this is due to calendar ageing alone appears to hold, as these are the hallmarks of SEI layer growth (which both consumes lithium, as well as inhibiting reaction kinetics). There does not appear to be any significant LAM, although LAM of the NE would be difficult to discern form this data. There is also clearly a slightly larger spread between the cells, as would be expected, as while they were stored together, there will obviously be some cells exposed to wider changes in temperature, as well as the process of calendar ageing exacerbating any cell to cell variability that was initially present. This data gives us enough information to put the round 2 cells into context, and while the analysis and conclusions drawn will not be nearly as clear cut as they would without this complication, the data is far from unusable.

Having established our unfortunate starting point, we can now look at the evolution of these curves throughout testing. We shall begin with the closest thing we have to a baseline for this round, which is 20°C, 1C discharge. We would not expect a significant amount of irreversible

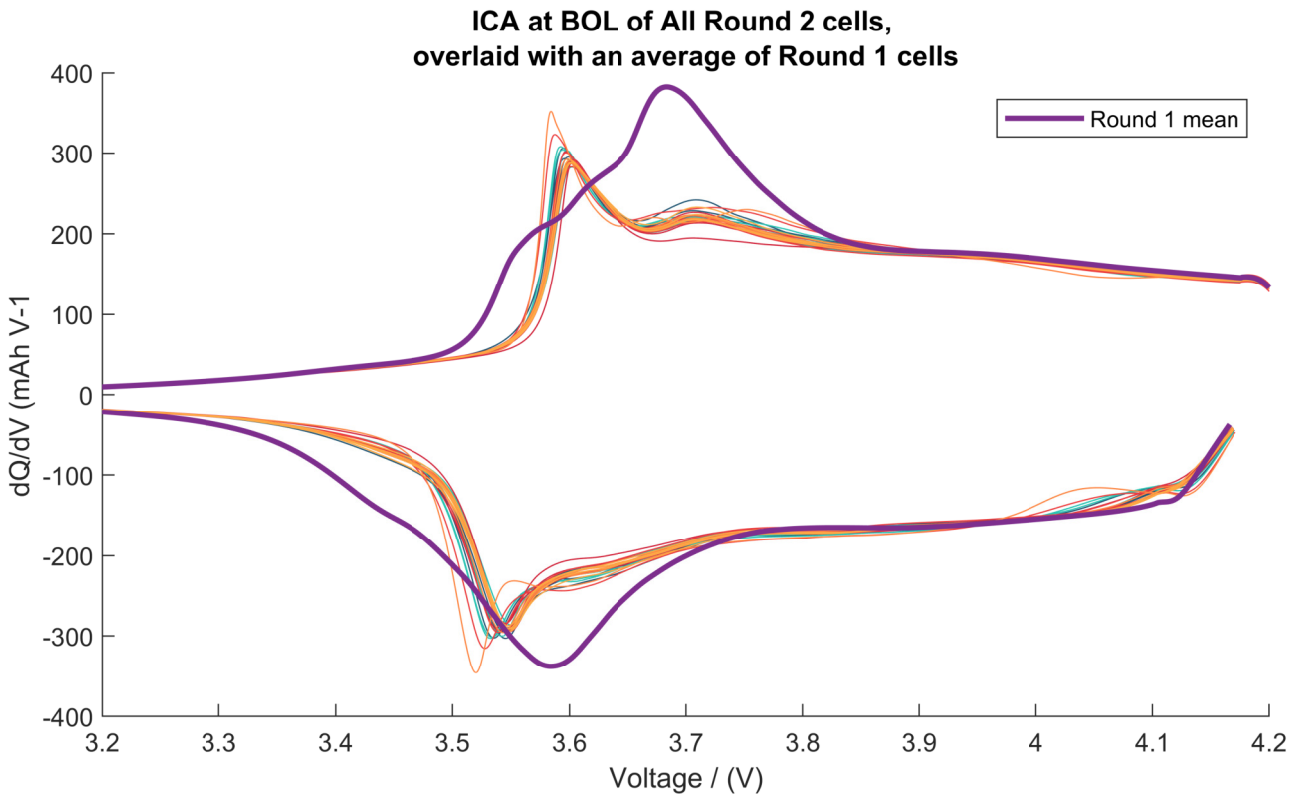


Figure 4.20: A plot of the ICA results at beginning of testing of all round 2 cells, overlaid with an average of the round 1 results

lithium plating to occur at this temperature, given that the cells are charged at C/2, therefore it is expected to behave similarly to cells tested at 30°C, 1C discharge, with perhaps slightly more aggressive LAM due to the lower temperatures increasing the magnitude of the diffusion gradients present during cycling, and therefore the mechanical stress which causes particle cracking. It is difficult to tell if this is true, given the incomparable starting points, but certainly the shift in the curves does appear to follow the same trend as seen in the 30 and 45°C data. There is less pronounced sharpening of the lower voltage peak in charge, but this is because it has already undergone the most noticeable shift prior to this testing. The higher voltage peak is once again lost completely, as it was in the round 1 experiments that exhibited the most severe degradation. In this case, of course, it is not because the experimental conditions are causing a faster degradation rate, but because of the compounded degradation pathway. This can be seen by the relatively small distance between adjacent traces, although this is also somewhat because the round 2 cells in general exhibited a slower rate of capacity loss. It is difficult to read a significant amount into this data without a better reference for comparison,

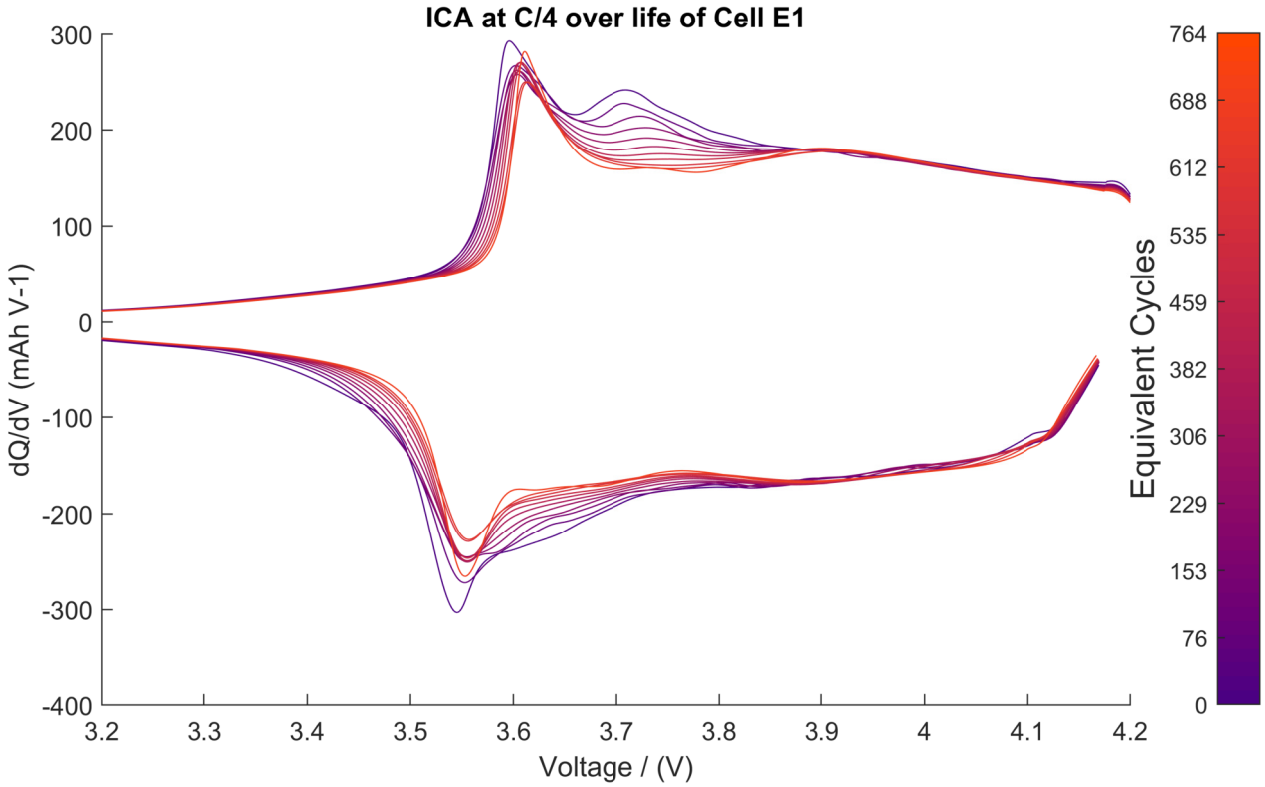


Figure 4.21: A plot of the ICA results of a cell cycled at 20°C, 1C discharge

so we must use this experiment as our stand-in reference point for the other round 2 experiments.

Beginning again with a comparison of data to investigate the impact of increasing discharge rate, Fig 4.22 shows the ICA development of a cell cycled at 20°C, 4C. It can be seen that the total capacity loss is higher as the depression of the curves is greater, once again completely removing the higher voltage peak in the charge direction and pushing it further into a local minima. The lower voltage peak is also significantly lower at the end of testing, and approximately 50-100mV further to the right compared to cells discharged at 1C. This indicates a greater LLI, but overwhelms the features that would help differentiate between LAM and any further changes in rate capability.

This trend is also present in the data from cells cycled at 55°C. Beginning with those discharged at 1C, an example of which is shown in Fig 4.23, it can be seen that the minima formed in place of the peak at higher voltage in the charge direction is pushed further down, but also widens significantly, pushing the maxima where it joins the plateau at higher voltages

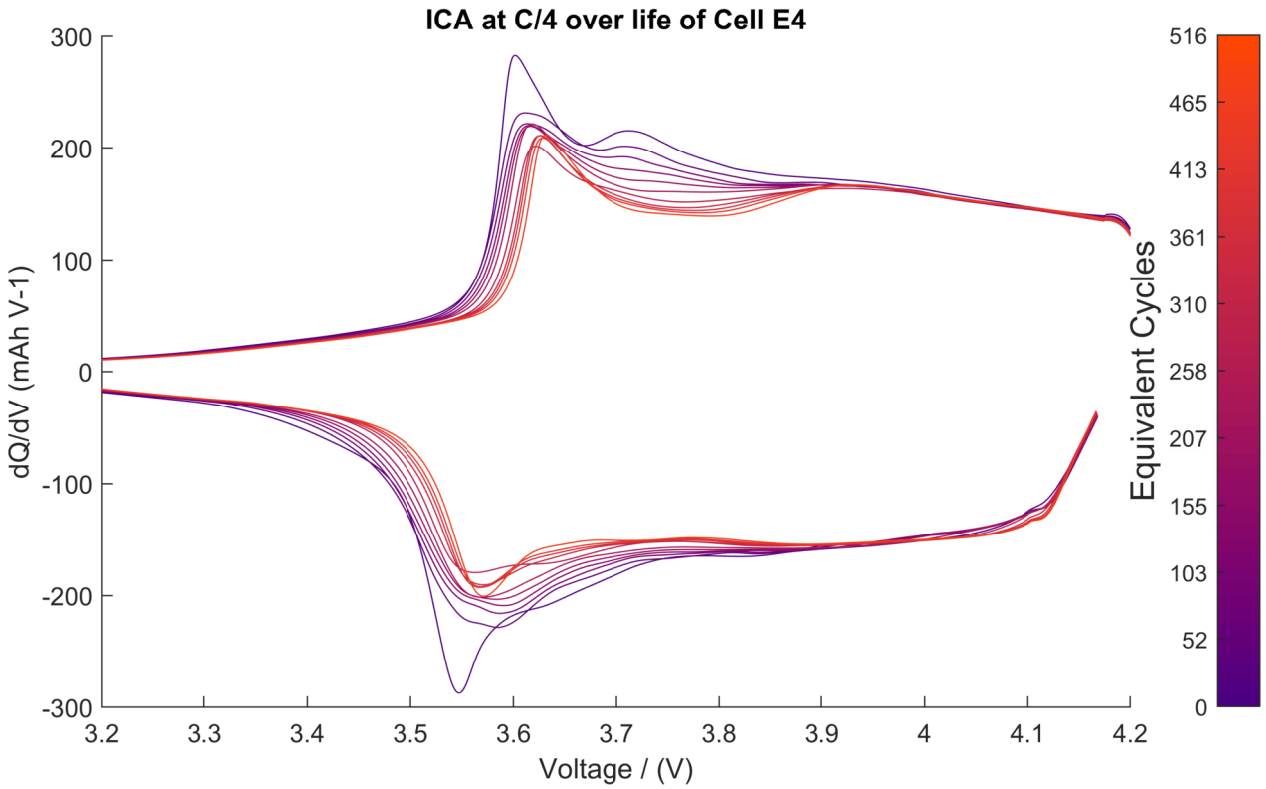


Figure 4.22: A plot of the ICA results of a cell cycled at 20°C, 4C discharge

further right, reaching approximately 4V by the end of the testing. The minima also broadens, forming a lower plateau between 3.7-3.9V. The shift to higher voltages is also accentuated, with the remaining peak in the charge direction moving by 150 mV between the beginning and end of testing. This shift indicates that perhaps LLI without LAM might be more prevalent at these higher temperatures, which we would expect, as once again the higher temperature serves to minimise particle cracking, but will accelerate side reactions that consume charge carriers, such as SEI layer growth.

Fig 4.24 shows the results from a cell cycled at 55°C, discharged at 4C. As with the capacity measurements, there is no discernible difference in the trend over testing between this and the 1C data at the same temperature, with a small increase in severity. This reinforces the assertion that the impact of discharge rate on balance of degradation modes is minimal at these higher temperatures, because the resistance to movement of charge carriers is reduced enough to similarly reduce the mechanical stress of diffusion gradients in the active material.

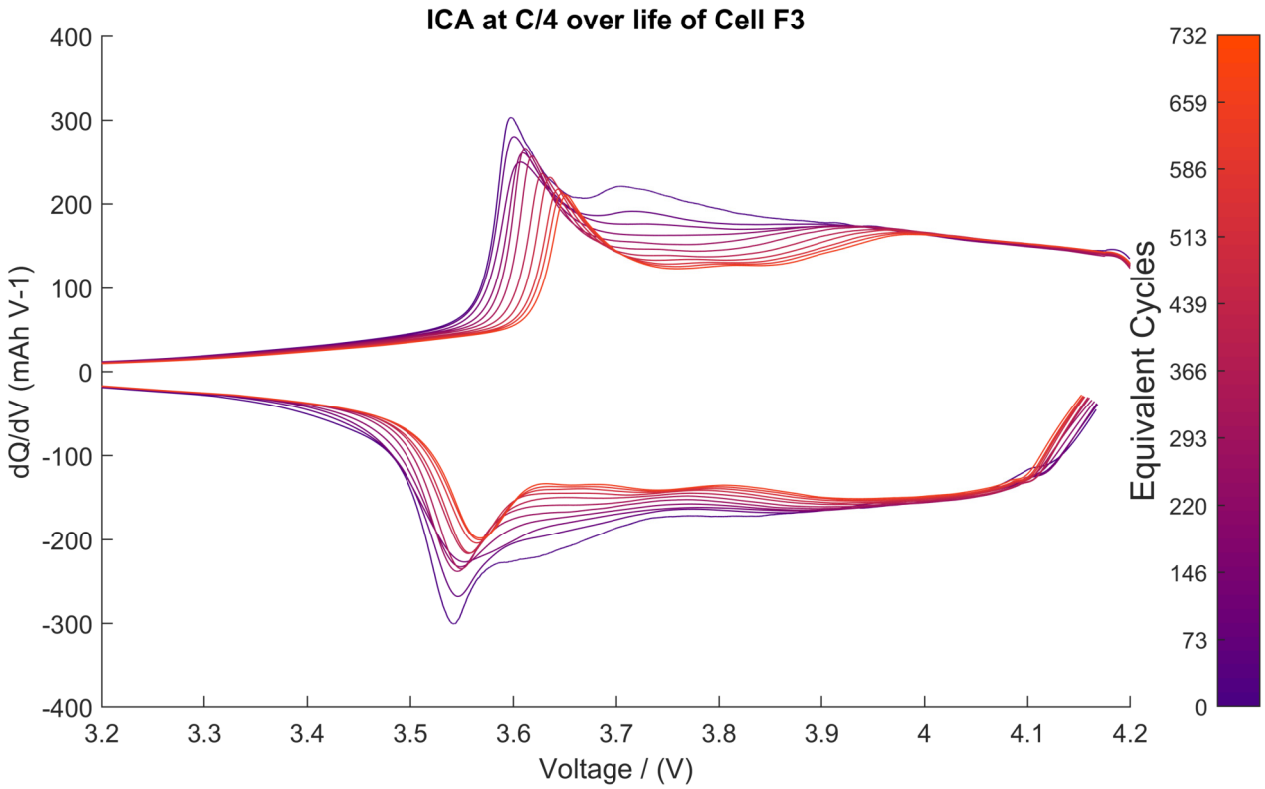


Figure 4.23: A plot of the ICA results of a cell cycled at 55°C, 1C discharge

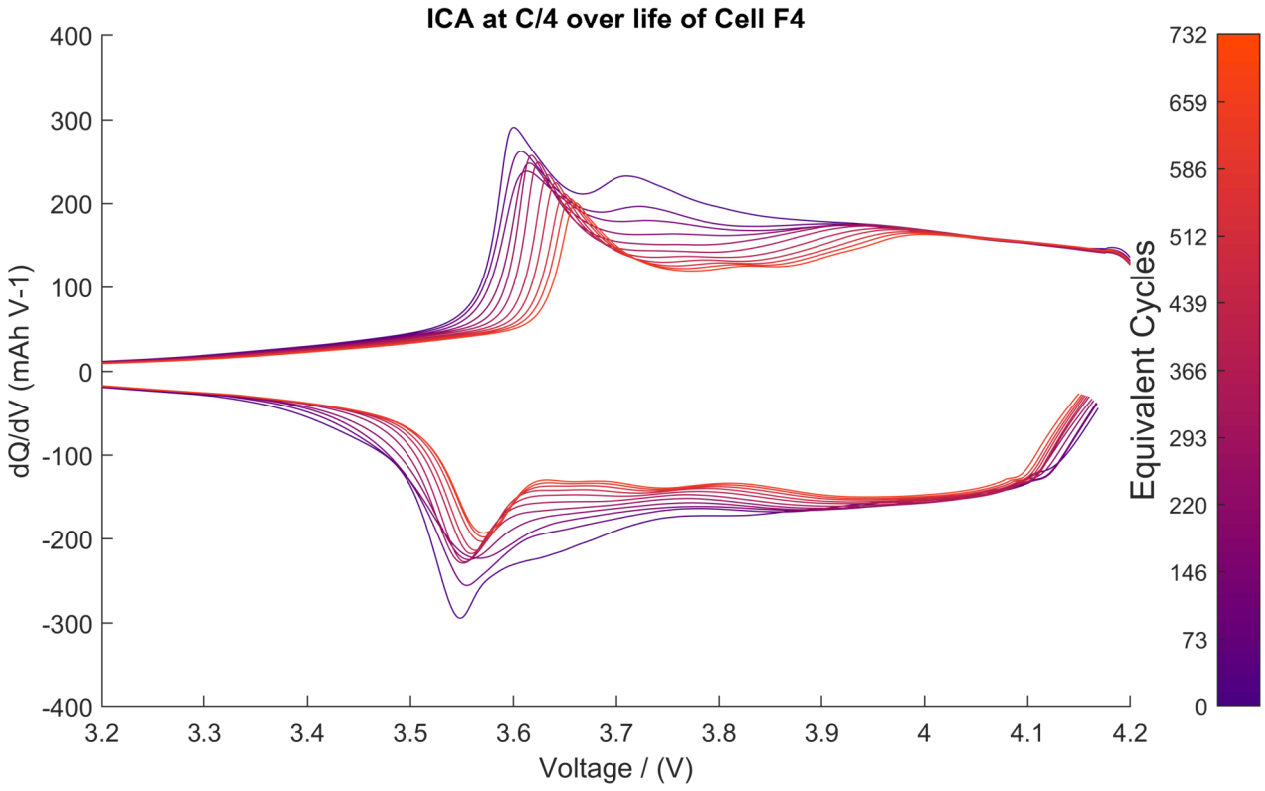


Figure 4.24: A plot of the ICA results of a cell cycled at 20°C, 4C discharge

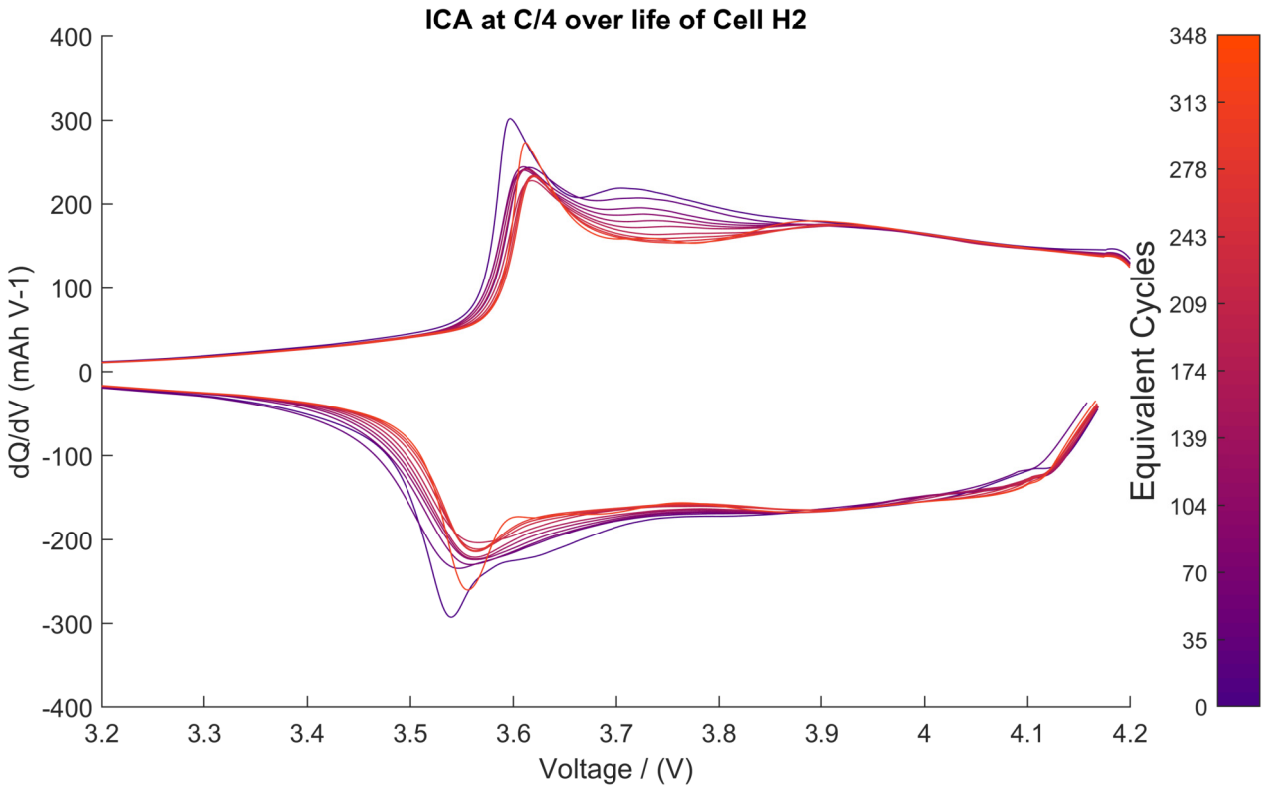


Figure 4.25: A plot of the ICA results of a cell cycled at 10°C, 4C discharge

Finally, comparing the ICA results from cells discharged at 4C, across 10,20,45 and 55°C, some clear trends are visible. While the capacity measurements show little difference between these experiments, the ICA demonstrates that there is a shift in degradation modes. At the lower temperatures, there is significantly more indicators of LAM, while the higher temperatures transition to LLI without significant LAM. This is demonstrated by the fact that at lower temperatures, there is a general depression of the curve, with the only major shift in features being the higher voltage peak in the charge direction beings transformed into a wide, flat local minima. There is no significant shift along the voltage axis, and the lower voltage peak in the charge direction is only significantly sharpened by the formation of the local minima directly to its right. In contrast, the higher temperatures see a much more marked shift towards higher voltages, as well as more significant sharpening of the same peak. This can perhaps be better seen in the discharge data, where the singular peak is left fairly broad and blunt at end of life (EOL) in the lower temperature experiments, but retains its full shape and point at the higher temperature. These features can be seen in Figs 4.25 and 4.26.

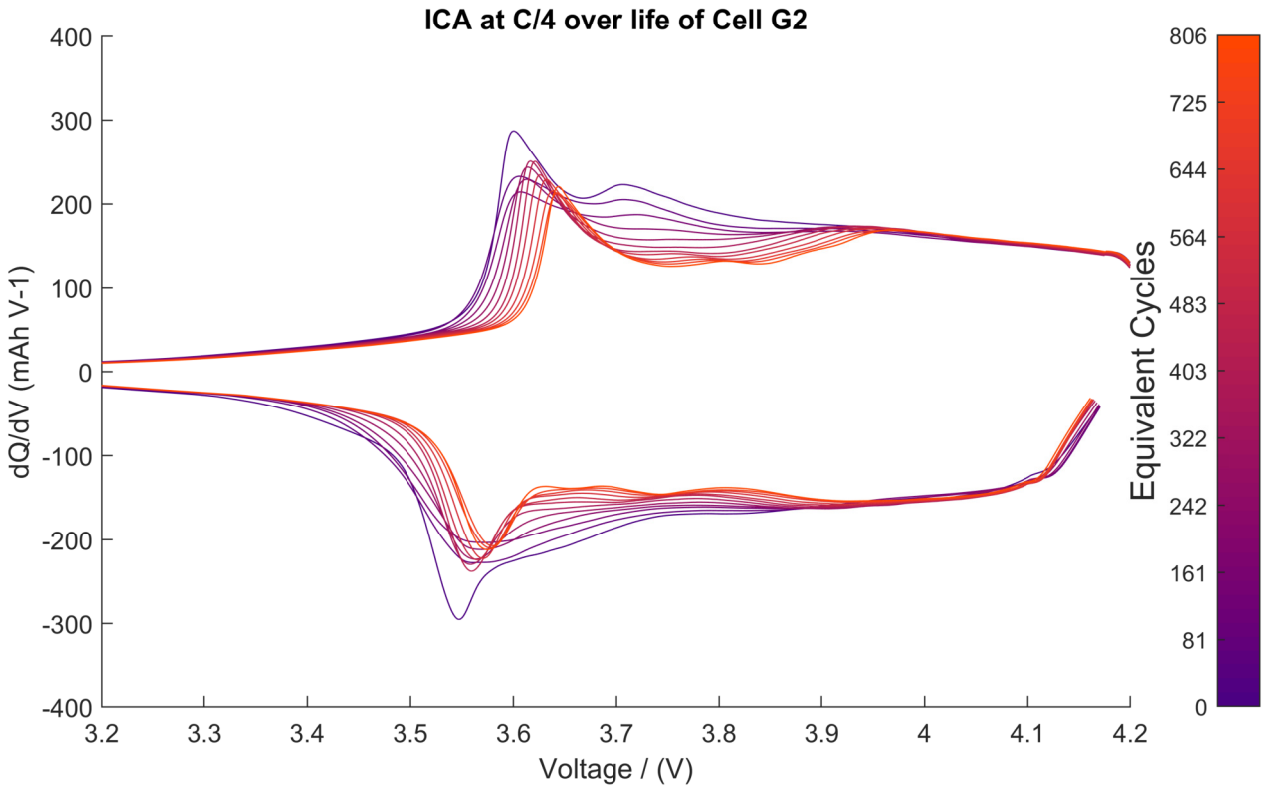


Figure 4.26: A plot of the ICA results of a cell cycled at 45°C, 4C discharge

4.4 Degradation Mode Analysis

As described in Section 2.4, DMA is a useful technique for quantifying the impact of the major thermodynamic degradation modes, but typically requires accurate half cell data to do so. The best practice for acquiring this half cell data is to disassemble a sample of the cells in question, make half cell coin cells from them, and then parameterise those. This is both time consuming, difficult, and has some additional sources of error (such as difficulty using identical electrolyte composition as a commercial cell). As there were insufficient cells to do comprehensive BOL half cell tests, and insufficient time to do full post mortem disassembly, fully parameterised DMA is not possible. However, we can see from comparison between the ICA results and literature that these cells do behave reasonably close to those of similar chemistry in the literature, and so it may be possible to utilise generic half cell data and still achieve a reasonable fit. Of greater concern was how the DMA fitting would behave on cells with such extreme overhang, but it proved to be a relative non-issue once fed the appropriate data. Code used was based on the technique used by Kirkaldy *et al*,[77] with minor modifications in order to interface it with the

data gathered in this work.

The DMA technique yields a variety of outputs. Initially, it calculates the lithiation fractions of each electrode by fitting the half cell curves to the full cell measurements. From this, electrode capacities and the offset between the electrodes are calculated. The degradation can then be separated into contributions from LAM at the PE, LAM at the NE and LLI in the system. Modes associated with changes in the kinetics of the cell, such as increase in Ohmic resistance, are not found as they are not found from the OCV of the system, but from its response to different C rates. A pseudo-OCV was found by combining the data from both charge and discharge at c/4 during full characterisation, DMA applied to each cell individually, and then these outputs were averaged with other cells in the same experiments. Fig 4.27 shows the upper and lower bounds of the lithiation fractions found, plotted against equivalent cycles, for cells cycled at 30°C and discharged at a 1C. It can be seen that the NE remains the limiting factor on capacity, as the fraction of lithiation it doesn't utilise is negligible. The bound that shows the most significant change over the course of testing is the PE lower bound, as the extreme overhang present in these cells means that the PE is never fully de-lithiated, and the depth of discharge it does experience is predominantly determined by the capacity of the NE.

Fig 4.28 shows the capacities calculated from this data, including the offset between the electrodes. It can be seen that the PE does experience capacity loss, but at less than half the rate of the NE (absolute as opposed to fraction of initial capacity). In addition, because it begins with superfluous capacity, this loss has negligible impact on the total capacity of the cell. Instead, it can be seen that cell capacity is nearly identical to that of the NE, albeit with a small amount of unused capacity. This is likely due to the overpotential present when acquiring this data, which would not be present if it had been possible to achieve reliable results at significantly lower C-rates. This again demonstrates that the capacity of the NE electrode is the limiting factor determining the capacity of the entire cell under these conditions.

If we look at the amounts of LAM and LLI calculated, as shown in Fig 4.29, we can see that

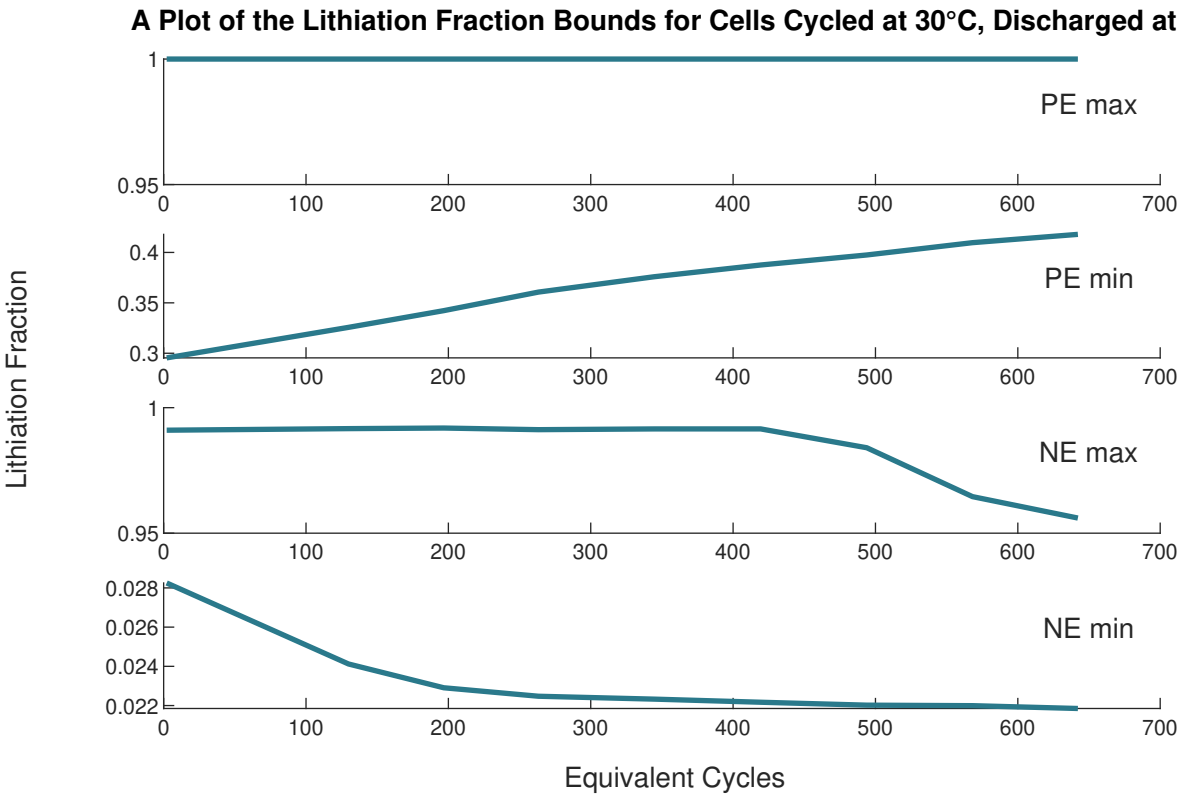


Figure 4.27: A plot of the lithiation fraction bounds against charge throughput of cells tested at 30°C and discharged at a 1C, calculated using the DMA technique

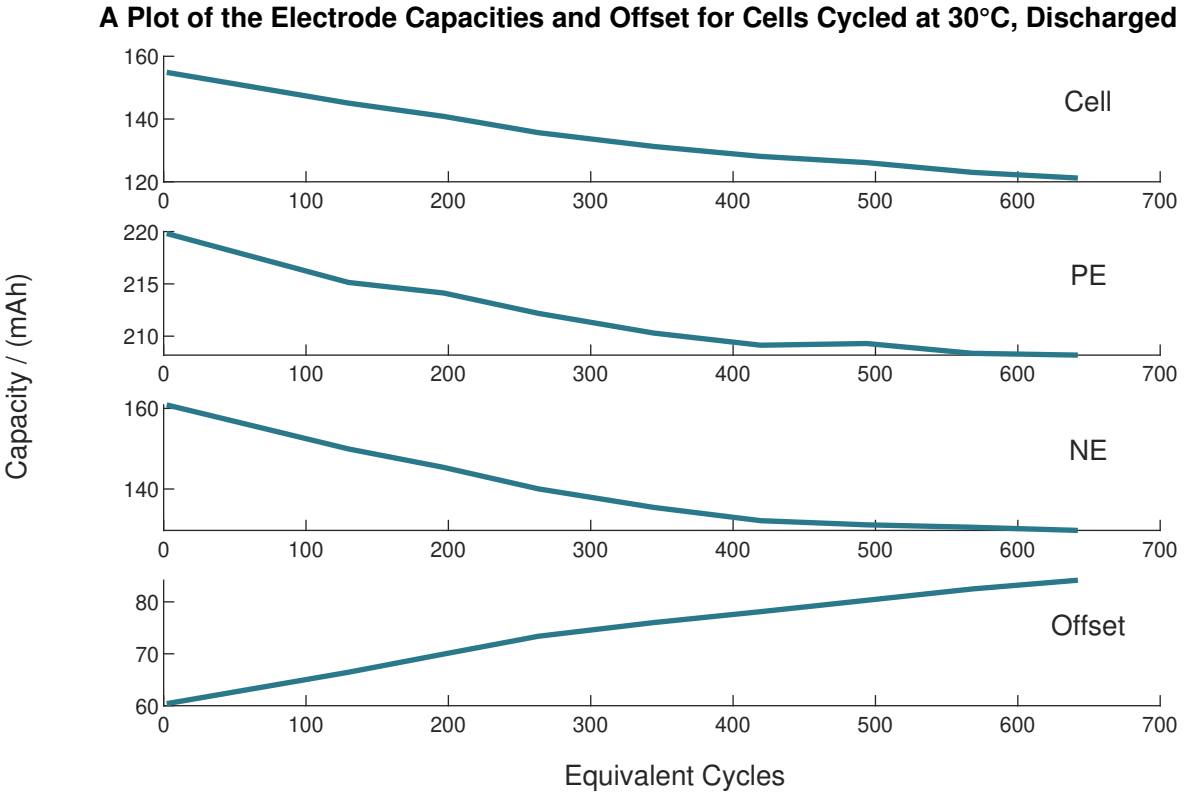


Figure 4.28: A plot of the electrode capacities against charge throughput of cells tested at 30°C and discharged at a 1C, calculated using the DMA technique

the NE loses significantly more active material than the PE (19.4% over the course of testing, compared with 5.3%). In addition, it can be seen that LLI follows the LAM NE closely up to 400 equivalent cycles, with a steadily growing positive gap. This implies that the LAM at the NE is occurring in the lithiated state, hence capturing charge carriers, with the difference primarily being due to the growth of the SEI layer. The PE experiences some LAM, but due to its significantly larger capacity the total capacity lost is smaller, as well as being a smaller fraction of the total PE active material. This is our final piece of evidence showing that the degradation of the single layer cells cycled at 30°C and discharged at 1C is dominated by LAM at the NE, likely from particle cracking, with a small measurable contribution from SEI layer growth. While the PE electrode does experience LAM, it is significantly less than at the NE, and does not impact the usable capacity of the device due to the extreme overhang present in these unorthodox cells.

After 400 equivalent cycles there is a shift in the dominant degradation mode. This can be observed in various parts of the data. Enumerating them systematically, in Fig 4.27 it can be seen that the upper bound of the NE lithiation fraction does begin to fall from this point on, albeit by a relatively small amount compared to the changes in the PE lower bound. This implies that from this point onwards, the full capacity of the NE is not being utilised. Next, in Fig 4.28, the NE capacity gradient shifts by a subtle amount at the relevant point, producing a sharp transition between gradients. This is a small enough change that on its own it may be disregarded as an artefact of the analysis process, but in combination with the other evidence does imply a change at this point. Finally, as has already been mentioned in Fig 4.29 the LAM at the NE stops following the LLI as closely, and the rate at which it is occurring (as a function of charge throughput) decreases. It is posited that the explanation that best fits this evidence is that prior to this point the dominant degradation mode is loss of lithiated active material at the NE, whereas afterwards SEI growth has a more significant impact on LLI without the loss of LAM tied to it. This is likely because the particle cracking experienced at the NE has exposed a significantly larger surface area for this SEI to occur. It could also indicate that a significant increase in the impedance of the cell is causing larger heat generation,

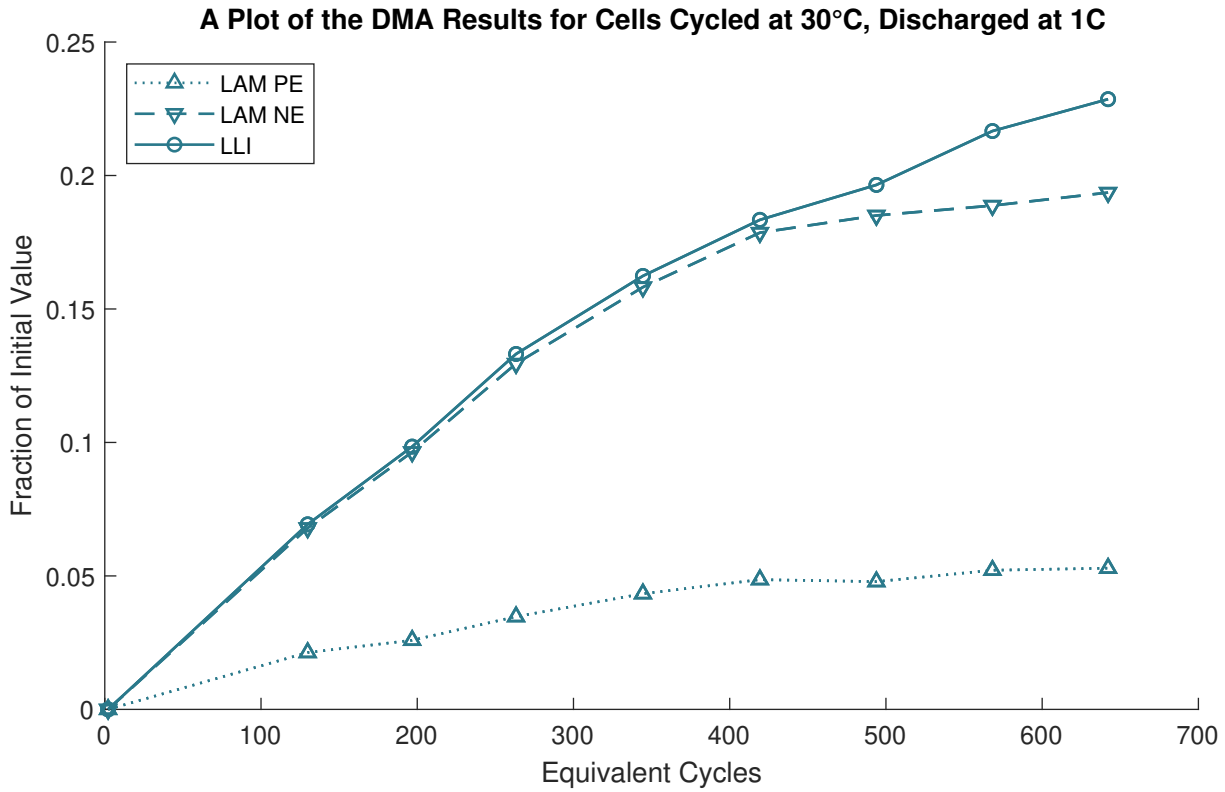


Figure 4.29: A plot of the DMA results against charge throughput of cells tested at 30°C and discharged at 1C

which is enough to raise the temperature of the active material and thus accelerate the SEI growth. However, based on the confidence in the thermal design of the test rig, supported by the validation data presented in Fig 3.3, as well as the fact that the end of cycle voltages do not show any significant shift (presented in Fig 4.13), it is felt that the exposure of new material surfaces for the SEI is much more likely to be the dominant factor in this shift in behaviour.

When examining the impact of changing conditions, these results allow insight into the shift in degradation modes, instead of just the magnitude of the capacity loss. Beginning with changing discharge rate, Fig 4.30 presents the lithiation fraction bounds calculated for all experiments conducted at 30°C. As expected, the PE upper bound and the NE lower bound stay at their limits (or near enough as to be a negligible difference in terms of cell performance). The PE lower bound tracks the total capacity loss of the cells, exhibiting the same trend as seen in Fig 4.6, namely, increasing discharge rate has little added impact on degradation up to 2C, and from then on is increasingly damaging. While this trend has been observed up until this

point, it has not been mechanically explained, but from this DMA analysis we can begin to see some possible mechanisms. The point at which the NE upper lithiation bound begins to deviate from its limit is nearly identical for cells discharged at 1-2C (certainly within experimental error of each other). However, this deviation occurs significantly earlier as discharge rate is increased from this point. When examining the 30°C, 1C discharge data, this deviation was posited to be indicative of a shift in the dominant degradation mode, as the particle cracking causing the LAM at the NE exposes enough surface area for LLI due to SEI growth to become dominant. This would explain the apparent threshold value that exist somewhere between 2-3C, above which significantly more capacity loss is measured, as the mechanical stresses (increasing with discharge rate due to the sharper diffusion gradients produced) begin to exceed the failure point of the material on a larger scales, as opposed to only where there is some stress concentrating geometry. This exposes more surface area earlier in the lifetime of the cell, and so has an exaggerated effect when accelerating the incidence of SEI growth becoming dominant over LAM. It is curious that this slows the rate of LAM at the NE, as opposed to accelerating the rate of LLI. One possible explanation is that the increased resistance to intercalation due to the thicker SEI begins to limit the magnitude of the diffusion gradients, therefore acting to limit mechanical stress, but there is little evidence for a significant shift in cell impedance.

Fig 4.31 shows that the trend with increasing discharge rate is carried over to the magnitude of the dominant degradation modes, with no discernible differences between the LAM at the PE, which of course does not contribute to any loss of usable capacity in the device. It can be seen that the shift in dominant degradation mode described above (based on the upper bound of the NE lithiation fraction) does occur, with the difference between the LAM at the NE and LLI following the trend in capacity, increasing significantly above 2C.

Turning our attention to the impact of temperature, we will begin with only the round 1 experiments. We are expecting little difference between 30-45°C, but higher LLI at 10°C due to lithium plating (as well as higher LAM due to increased particle stress caused by steeper diffusion gradients). Fig 4.32 shows the bounds of the lithiation fractions for these experiments. As

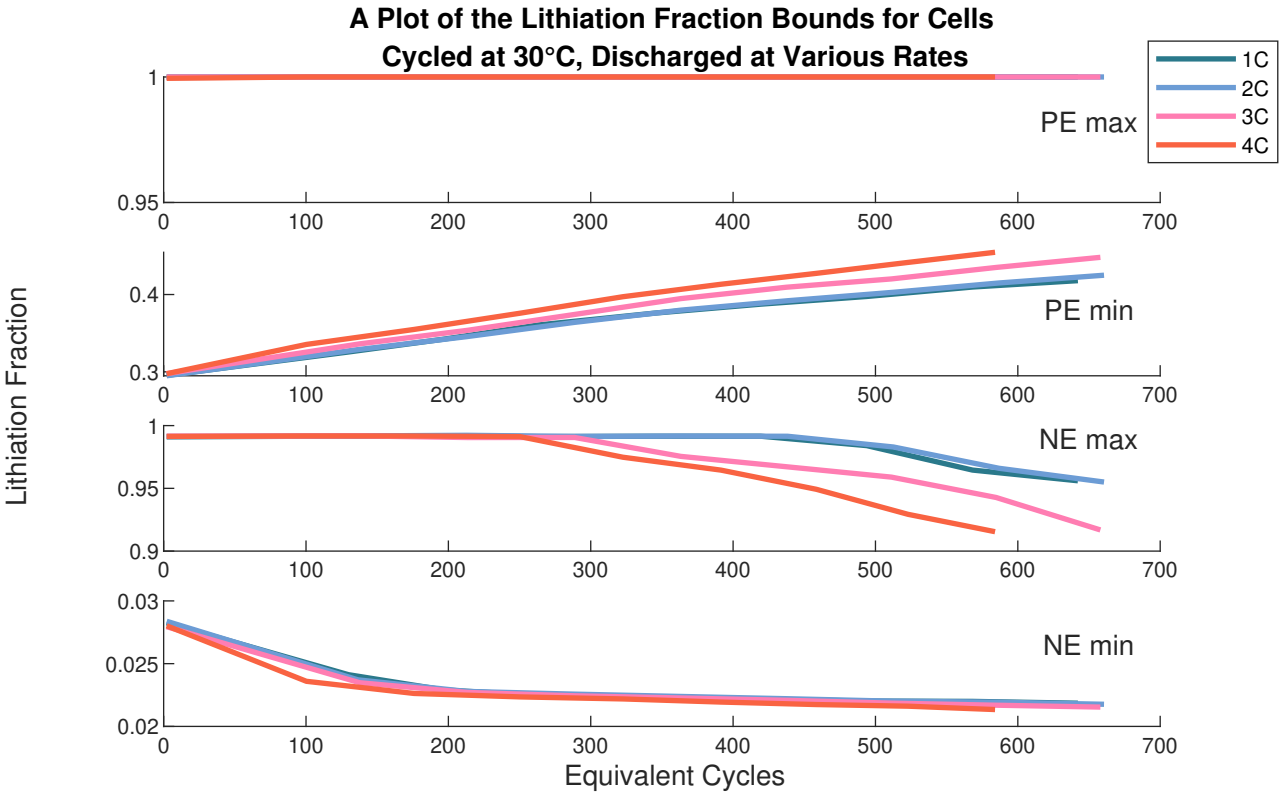


Figure 4.30: A plot of the lithiation fraction bounds against charge throughput of cells tested at 30°C and discharged at various rates

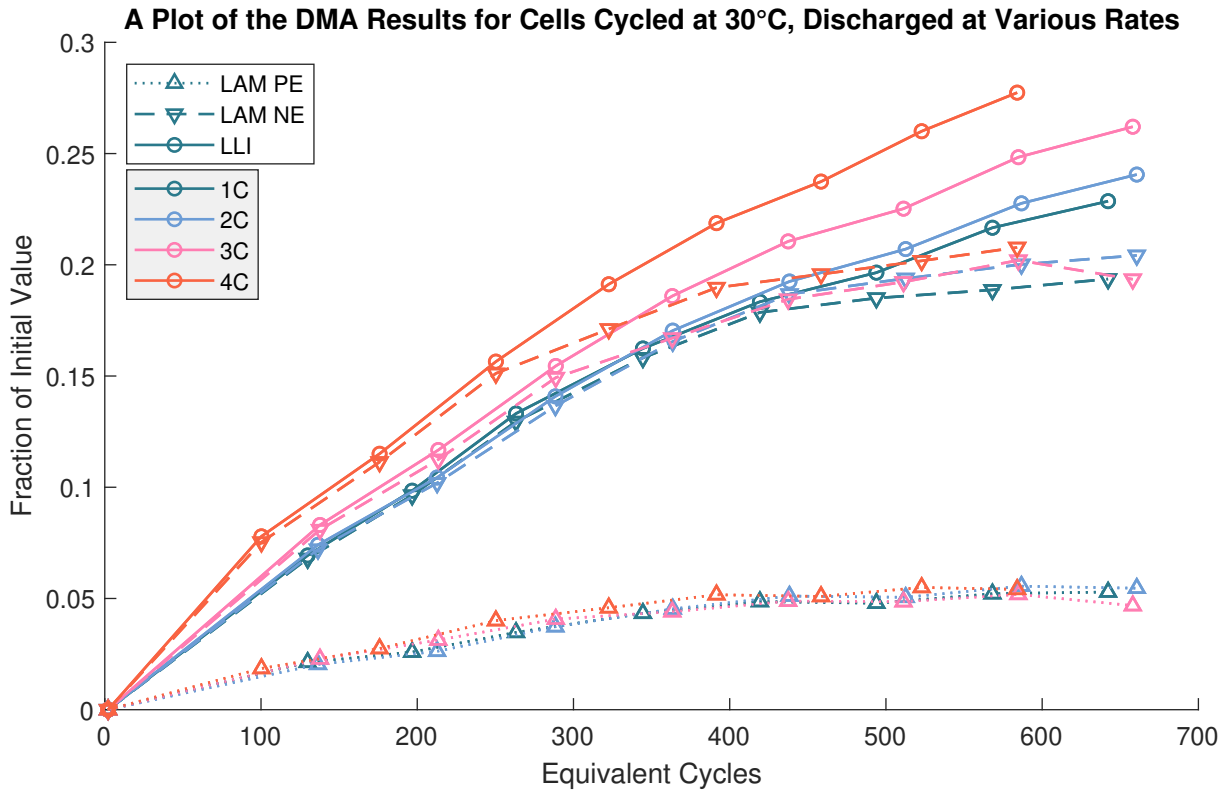


Figure 4.31: A plot of the DMA results against charge throughput of cells tested at 30°C and discharged at various rates

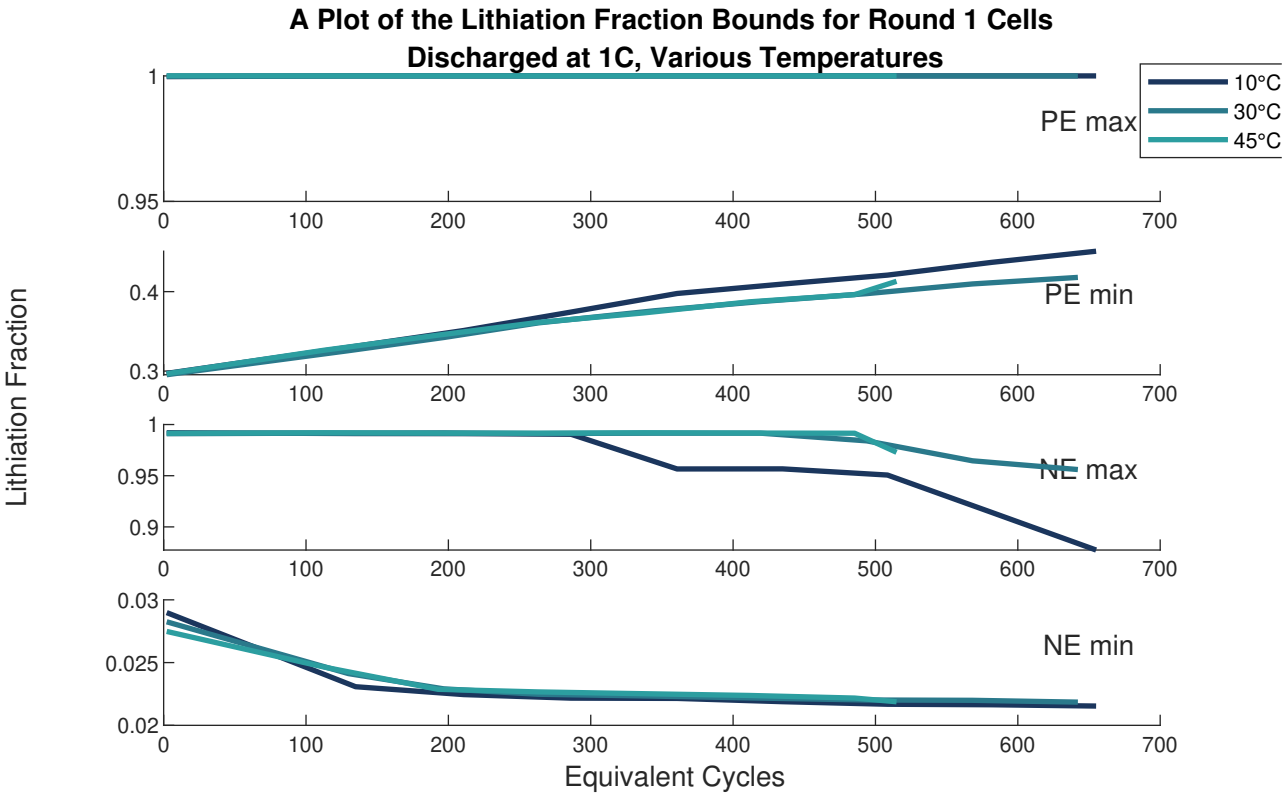


Figure 4.32: A plot of the lithiation fraction bounds against charge throughput of round 1 cells discharged at 1C, various temperatures

before, PE upper and NE lower lithiation fraction bounds stay at their limits, and the PE lower bound indicates the cells total capacity loss. The higher temperature (45°C) NE upper bound does deviate from its limit at the last data point, perhaps indicating the previously observed shift in degradation modes, although more data points would be necessary to say with any degree of certainty if this is so. This does occur later than at 30°C, but only slightly so. Because these cells were discharged at a relatively low rate (1C), the impact of increasing temperature is expected to be less than at higher rates, where the mechanical stress is significantly higher and so the impact of reducing it by increasing temperature is also greater.

At 10°C this deviation occurs much earlier, comparable to 30°C, 3-4C discharge experiments, which helps to explain some of the similarities between these experiments in the capacity data (Fig 4.8). It was predicted that these cells may experience some lithium plating, causing LLI to become dominant, but it could just as easily be explained by the more aggressive particle cracking leading to faster SEI growth. The transition from LAM at the NE to LLI being the

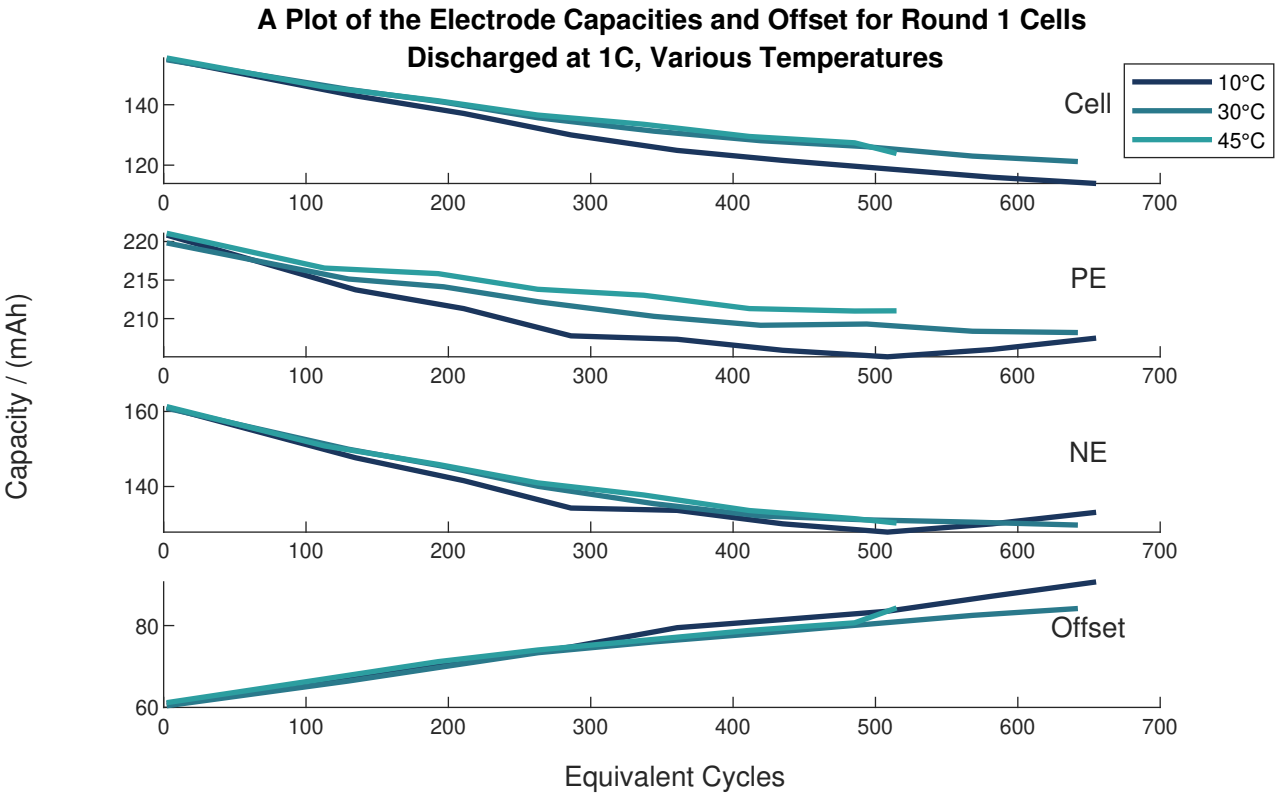


Figure 4.33: A plot of the DMA results against charge throughput of round 1 cells discharged at 1C, various temperatures

limiting factor in cell capacity can again be seen in the electrode capacities, shown in Fig 4.33, where the NE capacity stops following the trend of the total cell capacity at the relevant point in the experiments (just before 300 equivalent cycles). This change is again subtle, and given the variability in the calculated capacity of the NE after this point it would have been disregarded as noise in the measurements were it not for the other evidence presented.

The contribution of the degradation modes can be seen in Fig 4.34. The experiment at 10°C does exhibit significantly higher degradation across all modes, but the fact that it is mirrored so closely by the LAM at the NE casts some doubt on the assumption that lithium plating is dominant. It may be that at these lower temperatures and without through plane thermal gradients causing a greater mix of dominant degradation modes across the active material in the cells, lithium plating acts to replace SEI in consuming the charge carriers, but the fact it is largely undetectable in the gradient or shape of the LLI across the lifetime is disconcerting. In addition, it would be expected to see a larger difference between the LLI and LAM at the NE

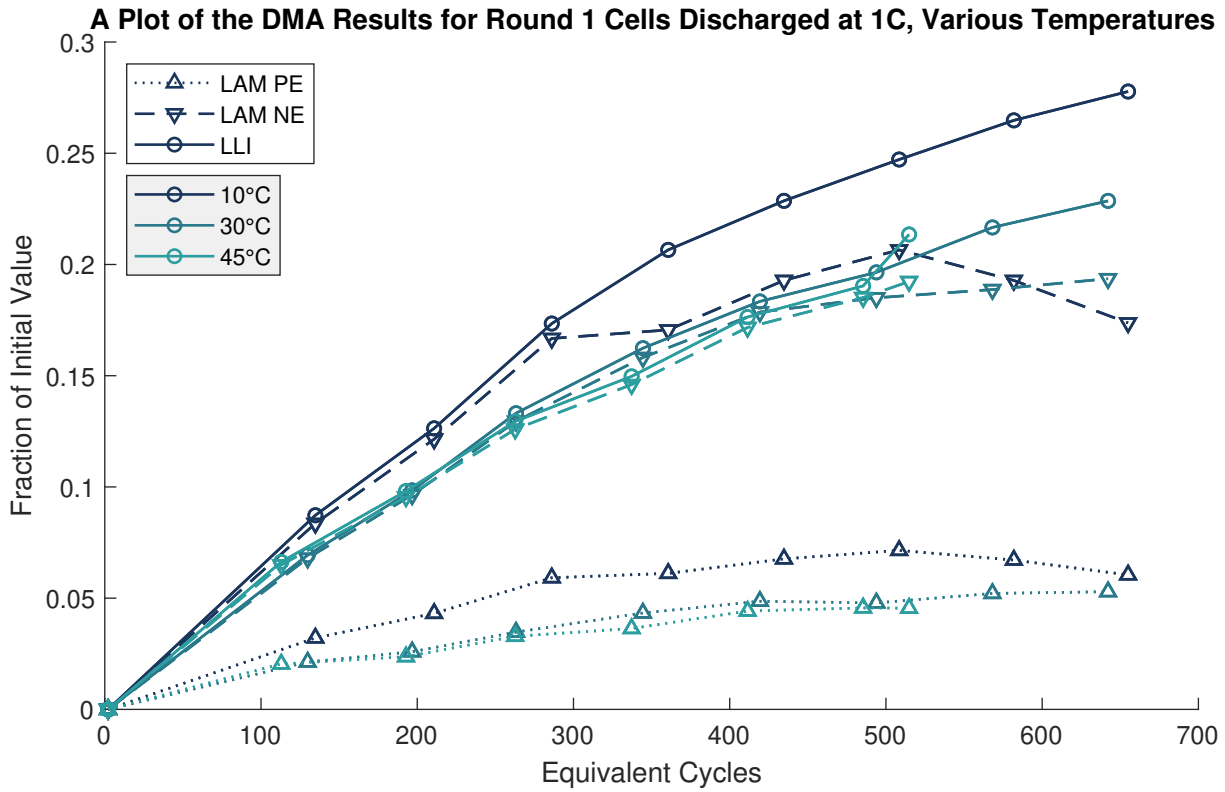


Figure 4.34: A plot of the DMA results against charge throughput of round 1 cells discharged at 1C, various temperatures

early on, as opposed to the sudden deviation exhibited part way through the testing. Indeed the same mechanism that is posited to explain this deviation in the other experiments (the increase in surface area due to particle cracking) might be expected to decrease the rate at which lithium plating consumes the charge carriers by reducing the overpotential at the surface, a key factor in lithium plating.

4.4.1 Round 2

Because DMA gives results normalised to the first reading, processing the round 2 results posed a conundrum. Given the already aged state when the first readings on these cells were taken, and the greater cell to cell variability (much more so in behaviour than in initial capacity), it was decided to input a false initial set of measurements to allow easier comparison to the round 1 data. This was done by taking an average of all the round 1 cells initial readings and using this as the initial readings for all round 2 cells. The initial readings for the round 2 cells were

also used, giving a quantified impact of the calendar ageing they experienced before cycling.

Continuing with an examination of the impact of temperature at 1C discharge, we can see the impact of the calendar ageing in the lithiation fractions in Fig 4.35. Round 2 experiments (performed at 20°C and 55°C) show a jump in the initial lower bounds of the lithiation fractions, demonstrating the significant loss of capacity in these cells prior to testing. Interestingly, their departure from the NE lithiation fraction upper bound is earlier than would be expected from the round 1 testing, with the 20°C experiment deviating 130 equivalent cycles before the 10°C (round 1), and the 55°C doing so after only 150 equivalent cycles. Based on the theory posited above, that this deviation is due to SEI growth becoming dominant thanks to the increased surface area available from particle cracking, it might be expected for the 55°C experiment to break the trend and exhibit this deviation earlier than the lower temperature experiment, as the increased temperature will accelerate the growth of SEI. While the increasing temperature will also act to reduce particle cracking, this will be balanced with the acceleration of parasitic side reactions, and the net outcome is very difficult to predict, hence the need for experimental data. Regardless, the early onset of this shift in degradation modes is undoubtedly due to the calendar ageing these cells experienced prior to testing, and obfuscates comparison between the 45°C and 55°C experiments for this purpose. The best we can do is compare the measured capacity of the cell as a whole at the point they begin to deviate from the NE lithiation fraction upper bound, which is 124 mAh and 121 mAh for the 45°C and 55°C experiments respectively. Alternatively, we could use the PE lithiation fraction lower bound as a guide, which yields 41.3% for both cases. This comparison tells us that they might have been expected to occur simultaneously, but the additional source of error from the different degradation pathways experienced makes this assertion untrustworthy. Only if there had been a large bias to one side or the other could this comparison be of further use.

To further complicate matters, it can be seen from Fig 4.36 that the DMA process has calculated that after 400 cycles, the cells tested at 55°C actually recovered capacity. If this had been a small amount, this might be considered to represent a real phenomenon, or be attributed

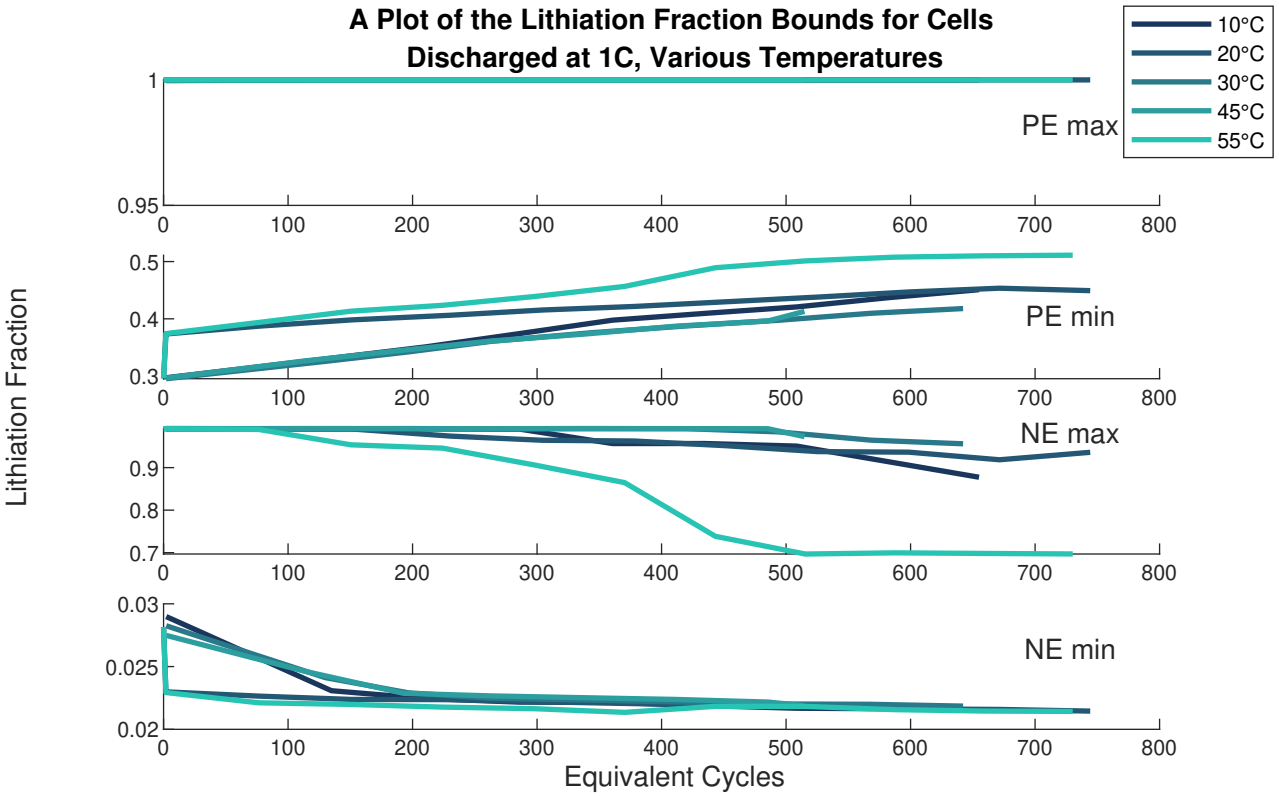


Figure 4.35: A plot of the lithiation fraction bounds against charge throughput of cells discharged at 1C, various temperatures

to experimental noise. However, as the capacity of both the PE and NE recover to greater values than their initial measurements, this is clearly an artefact due to the DMA process, and not representative of a real occurrence. Because LLI has become dominant at this point in the testing, these alleged increases in electrode capacities do not have an impact on the cells total capacity, and this is likely why the curve fitting has produced these erroneous results. The electrode capacities are no longer clearly observable in the input data, so their calculated values begin to vary unpredictably. This source of error can potentially also be seen in Fig 4.35, where the NE upper bound falls to 70% after 515 equivalent cycles, and then stays there consistently, despite the trend prior to this point being a rapid decay. What is likely occurring here is that the curve fitting is getting stuck at a local minima, or being limited by some other process which is entirely due to the application of the technique (i.e. the code), as opposed to a theoretical shortcoming of the technique. It is very possible that with better half cell data as an input (i.e. measured from disassembled samples of these cells, as opposed to generic half cell data of a similar chemistry), this would not occur. There are no clear ways to resolve this, as the limit is

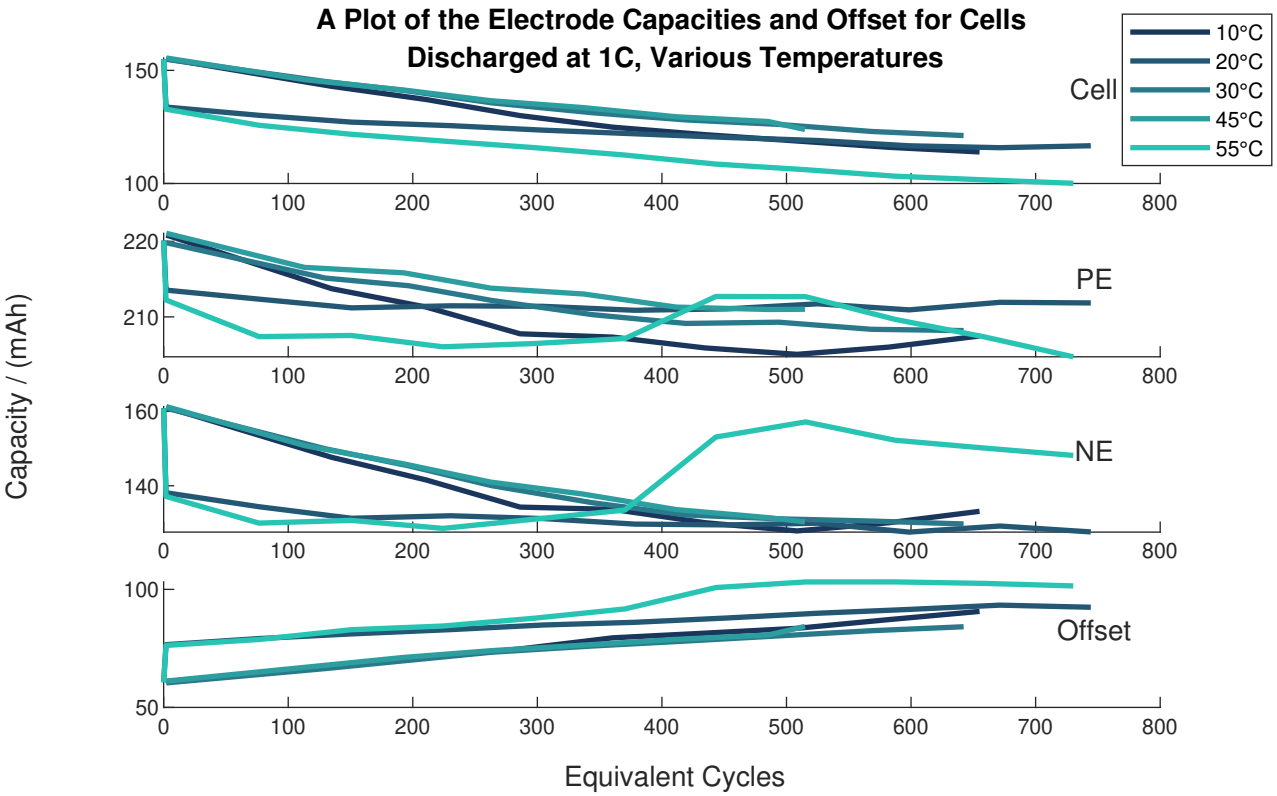


Figure 4.36: A plot of the electrode capacities against charge throughput of cells discharged at 1C, various temperatures

not something that is explicitly defined (indeed the actual limit explicitly defined for this curve fitting is 0), but rather the symptom of some more complex interaction that is not fully known or understood, and the project has insufficient time to overhaul the entire code base to find its source. It is sufficient for the purposes of this report to identify the error, determine with relative certainty that it is not representative of a real process (easily done in this case thanks to the magnitude of the fault), and ensure it is not used in any conclusions or predictions drawn from this data set. Essentially, we shall disregard the outputs of the DMA process as applied here for the PE and NE capacities of cells tested at 55°C, 1C discharge after 450 equivalent cycles, and be vigilant in identifying other examples of this fault which may be more subtle in their presentation.

We can see that this error continues in the calculation of the contribution of the degradation modes form Fig 4.37, with the LAM at both the NE and PE appearing to fall dramatically. Again, a small variation here could be indicative of a real phenomena, either from experimental error, or the re-connection of some amount of the active material that was previously

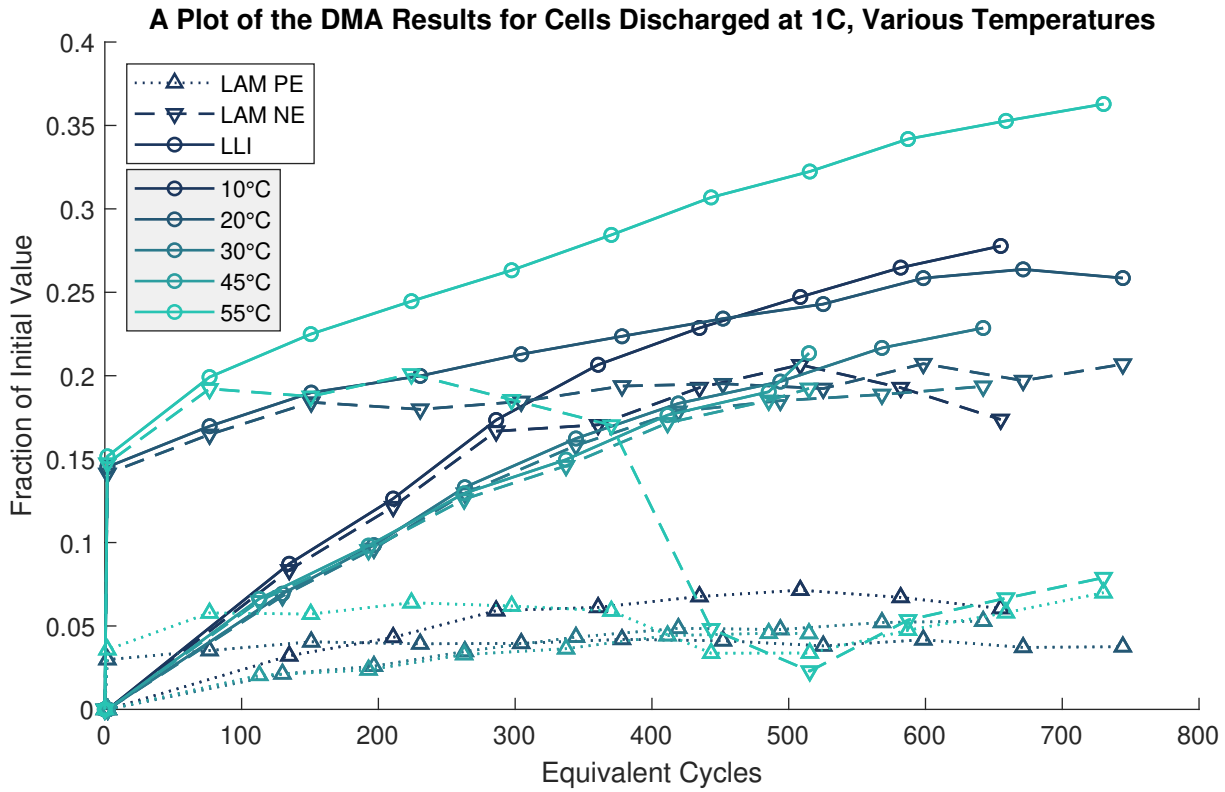


Figure 4.37: A plot of the DMA results against charge throughput of cells discharged at 1C, various temperatures

detached. However, once again the scale of the difference relative to the total electrode capacity, as well as the LAM prior to this point in both this experiment and others makes it clear that this is a fault in the processing, and not something we should attribute to a physical process.

Finally, the impact of temperature at higher discharge rates can be examined using the cells discharged at 4C. Because of the higher rate, these cells behaved more erratically, and so for clarity of presentation they have been split into two groups, 10-30°C and 30-55°C (both including the 30°C experiments as a point of mutual reference). It can be seen in Fig 4.38 that as soon as cycling begins, the cells cycled at 10°C did not utilise the full NE capacity, as the upper lithiation fraction bound falls, with the 20°C experiment following closely after at 60 equivalent cycles. Cells tested at 30°C showed this departure from the limit at 250 equivalent cycles, but were of course round 1 experiments, without the calendar ageing prior to testing the other data points experienced. The early onset of LLI as the limiting factor for cell capacity at the lower temperatures is most likely linked to lithium plating, which is made significantly

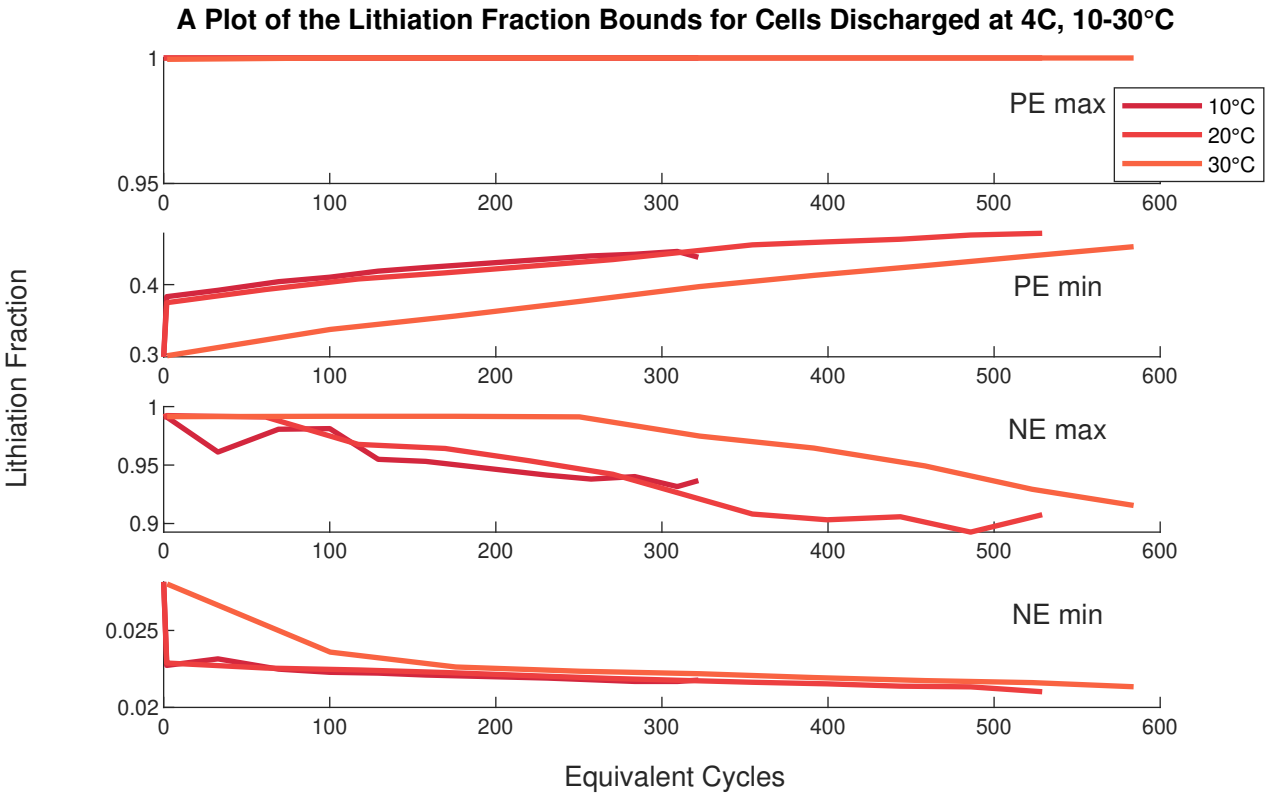


Figure 4.38: A plot of the lithiation fraction bounds against charge throughput of cells discharged at 4C, various temperatures

more aggressive by the increased discharge rate. While LAM will be occurring, it is difficult to ascertain its magnitude with any certainty as its impact on the cells capacity (as well as the measured OCV, used as the input for the DMA process) is minimal.

Given that we do not see the erratic behaviour, including the increase of the calculated NE capacity to values larger than feasible, that was discussed in the 1C, 55°C experiment, it is reasonable to place some cautious trust in the data presented in Fig 4.39. This implies small amount of LAM is occurring at the NE, with negligible amounts at the PE. The offset between the electrodes mimics the cells capacity most closely (albeit flipped on the y axis), in contrast to the lower discharge rates where NE capacity was more closely linked until the shift in dominant degradation mode later in the cells lifetime.

We can see the outcome of this difference in the calculated impact of the degradation modes in Fig 4.40. The LAM at the NE no longer shadows the LLI, with the exception of the first

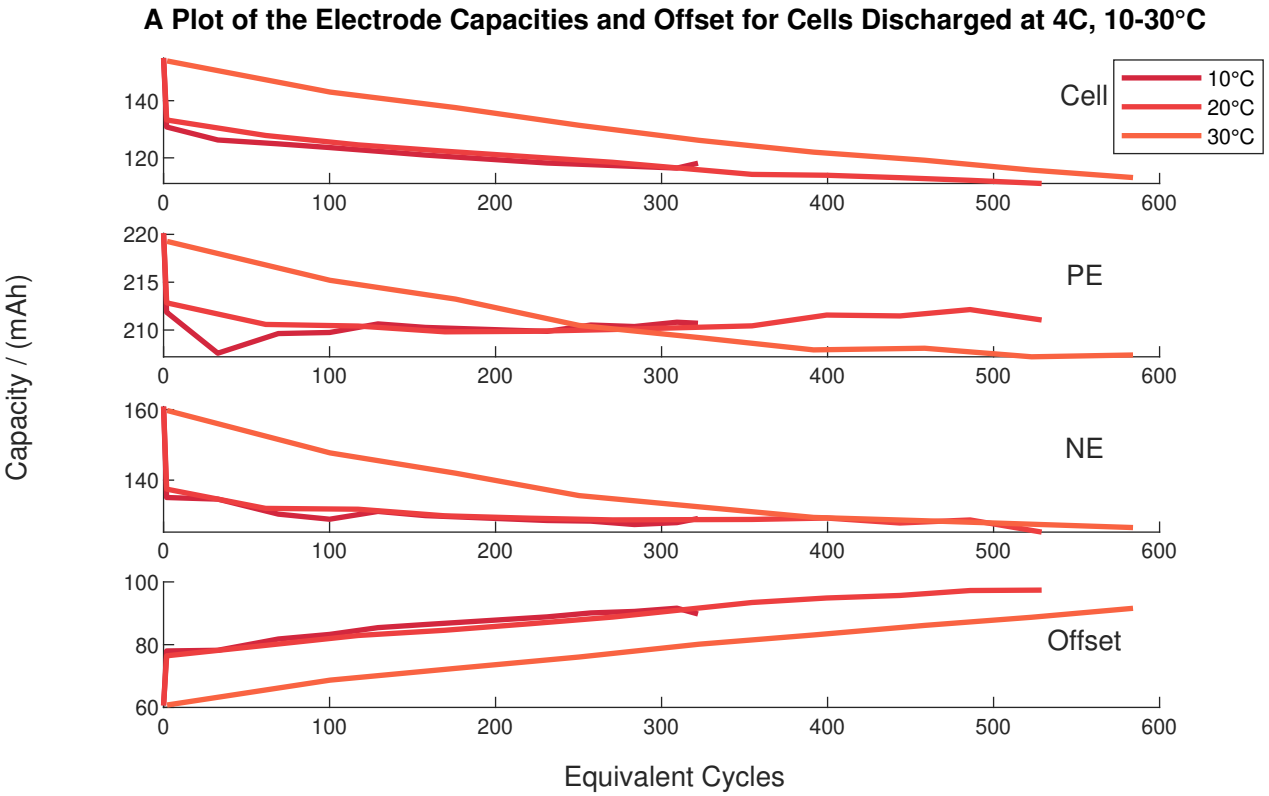


Figure 4.39: A plot of the electrode capacities against charge throughput of cells discharged at 4C, various temperatures

couple of measurements in the 20°C, instead showing a similar trend to the PE, with a small increase over the remaining testing. This is clearly not enough to produce the loss of cell capacity measured, for which LLI is responsible. This LLI could be from either SEI growth or lithium plating, and in reality will be some combination of the two (see Sec 4.6 for a discussion regarding how they are linked).

At higher temperatures, however, the DMA calculations are less reliable. For the 55°C cells this is largely due to the artificial limit found by the NE lithiation fraction upper bound at 0.7, which has been discussed above for the 55°C, 1C discharge experiment, where it is also found. The 45°C experiment does not find this limit (with the possible exception of the very last measurements taken), but does behave erratically after 600 equivalent cycles. This is once again attributed to a combination of limitations in the code and poor input data, the reasons for which have already been discussed. Specifically, attempting to fit the half cell curves to the available data can produce spurious results when the degradation of one of the electrodes has

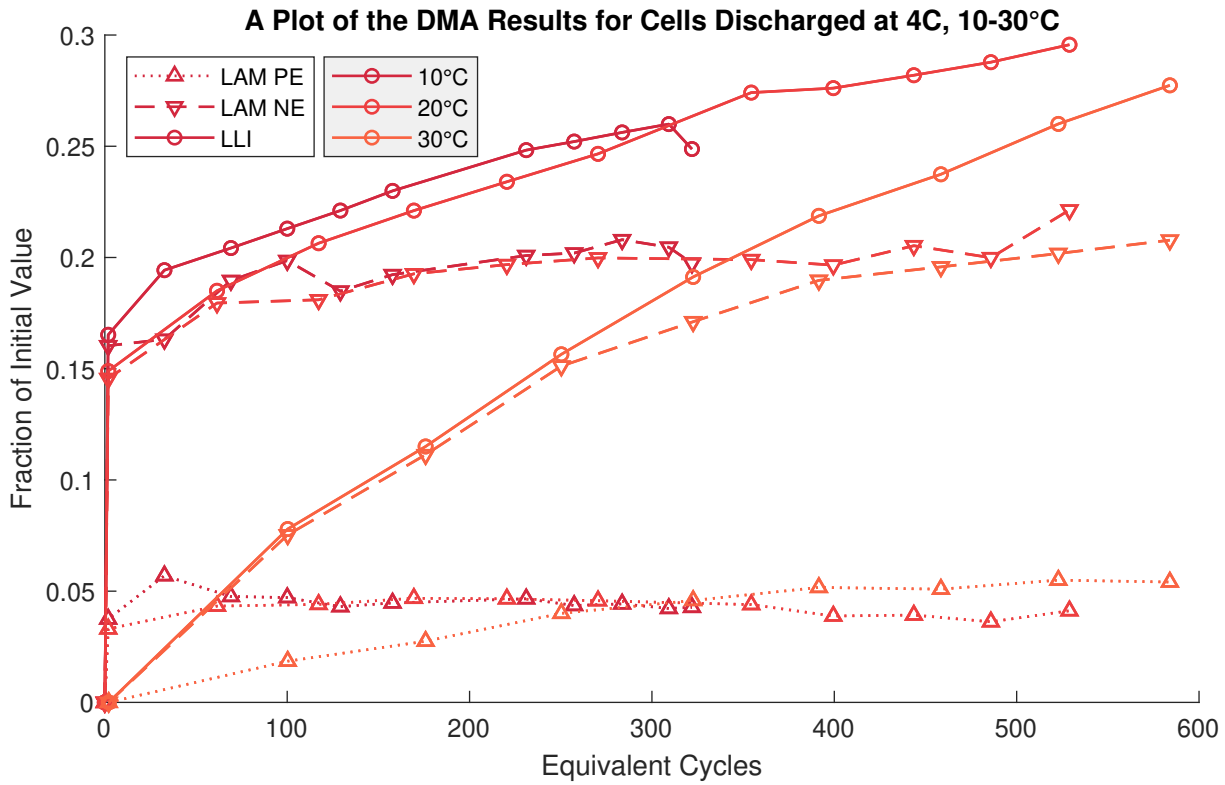


Figure 4.40: A plot of the DMA results against charge throughput of cells discharged at 4C, various temperatures

minimal impact on the behaviour of the device as a whole, especially when using generic half cell data, as opposed to half cells made from the specific cells under test (i.e. same product form the same manufacturer, ideally from the same manufacturing batch). It should be stressed that this is an extreme application for this technique, thanks to the cells unique manufacture and the lack of sufficient samples to produce better half cell OCV's.

This behaviour is more evident in the electrode capacities, again due to their impact on the cells OCV being minimal, and therefore the DMA process having little in the way of identifiable features from which to calculate these values. Because of this, it is felt that any conclusions drawn from the electrode capacities of the 55°C experiment after 350 equivalent cycles and the 45°C after 600 equivalent cycles would be spurious at best, and misleading at worst. This includes parameters calculated from these values, such as the LAM. Again, we see the offset between the electrodes correlating most closely with cell capacity, in contrast to the lower discharge rates.

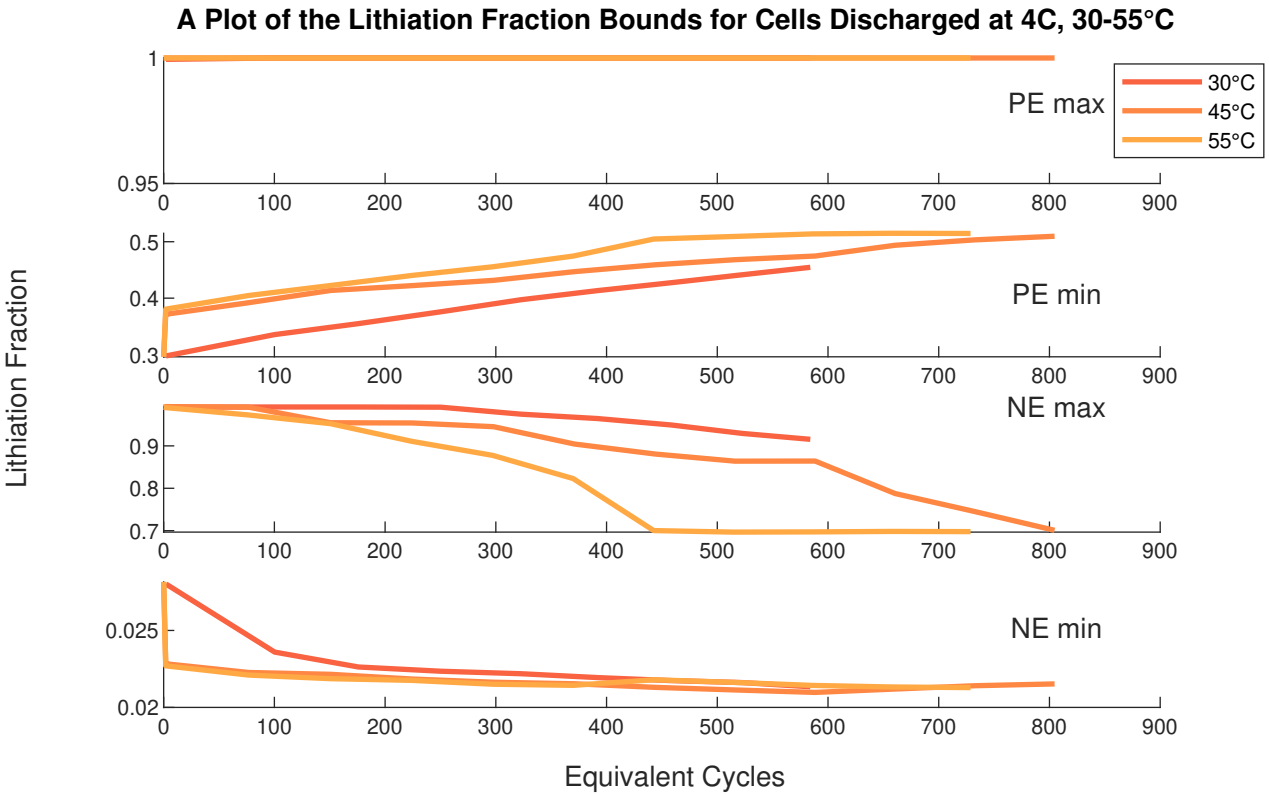


Figure 4.41: A plot of the lithiation fraction bounds against charge throughput of cells discharged at 4C, various temperatures

Despite the assertion that the data identified above is untrustworthy, it has been included in Fig 4.43 to illustrate the point. The LLI, however, by virtue of being so strongly dominant, yields much more stable and trustworthy values. The shape between the two higher temperature round 2 experiments is very similar, with the difference being the magnitude and gradient of the curve. This implies that the underlying cause of this LLI is accelerated by increasing temperature, and therefore most likely due to SEI growth, as opposed to LAM in the lithiated state or lithium plating, both of which are expected to have the opposite correlation based on both the available literature and the other data presented in this chapter.

Having established the features of interest in the results from these single layer cells, a more detailed discussion of their import, as well theories regarding what physical processes they represent will be presented in Section 4.6. In addition, the import of these conclusions with respect to real world engineering decisions will be examined, as well as suggestions for further

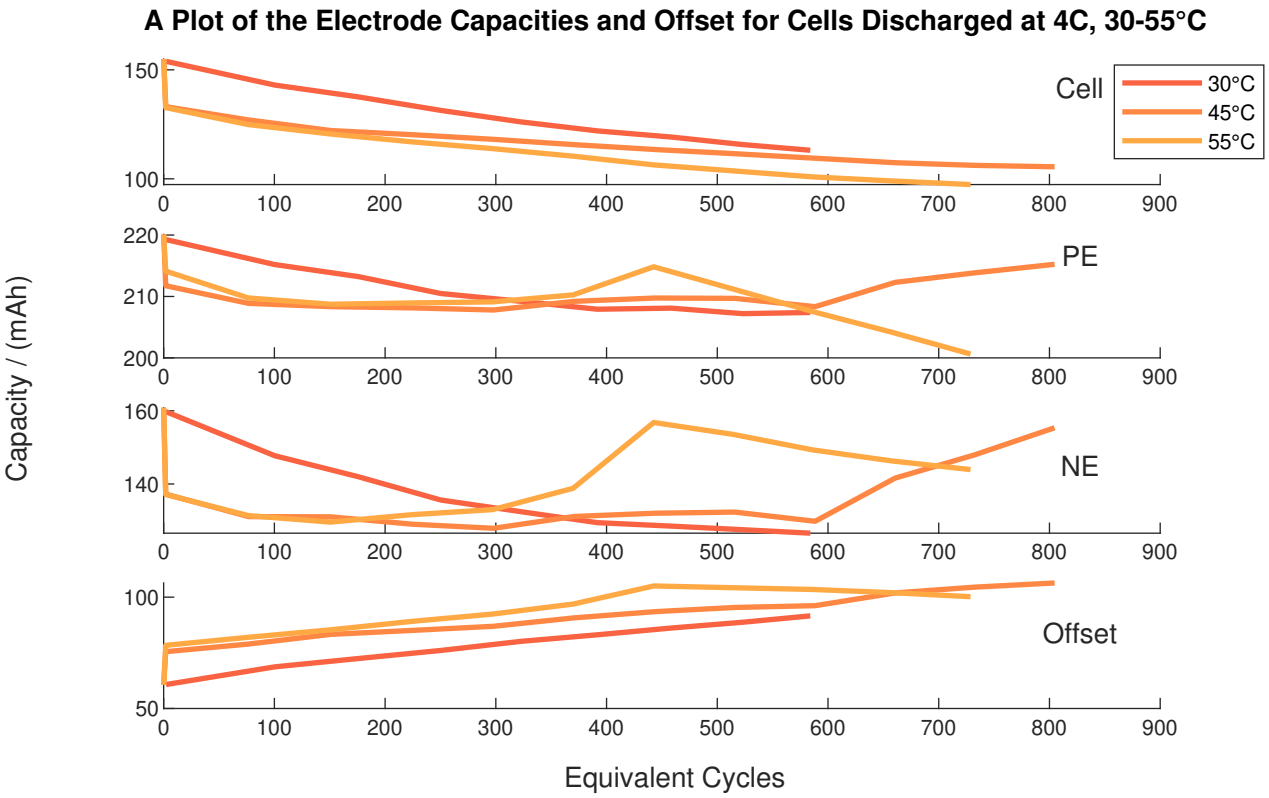


Figure 4.42: A plot of the electrode capacities against charge throughput of cells discharged at 4C, various temperatures

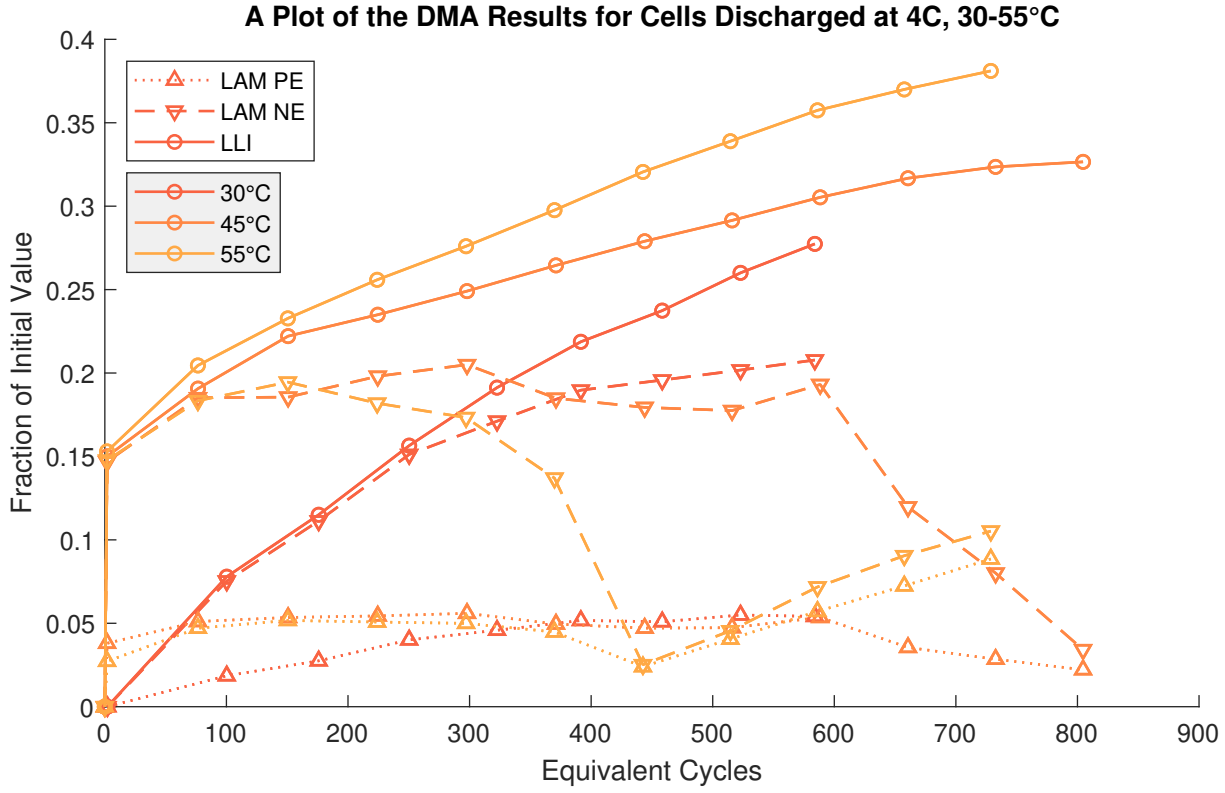


Figure 4.43: A plot of the DMA results against charge throughput of cells discharged at 4C, various temperatures

work based on the successes and failures of this initial use of single layer pouch cells for the purposes of de-coupling the impacts of C-rate and temperature on LIB degradation.

4.5 Pouch Control Experiments

As mentioned in Section 3.3.4, a set of pouch cells made from the same batch of material were purchased in order to compare their degradation with the single layer experiments. These were subjected to identical conditions as far as possible, with a modified version of the rig utilising active temperature control to impose isothermal boundary conditions on the major faces. The rigs were then placed in thermal chambers at the required temperature in order to maintain the correct thermal boundary conditions at the remainder of the surfaces. The intention was to demonstrate the impact of internal thermal gradients on the degradation of the cells, thus validating the use of single layer cells for further study. These control cells had the same footprint as the single layers used in this work, but were claimed to have a nominal capacity of 3 Ah thanks to their more typical number of layers. The presentation of the data from these cells will follow the approach set out for the single layer cell data presented above, with the exception of not being split into separate rounds of experiments.

4.5.1 Beginning of life capacity

After initial capacity measurements, they were found to have a significantly higher capacity than this, as can be seen in Fig 4.44. Initial measurements were taken in the same fashion as with the single layer cells (enumerated in Section 3.3.5), with 1C being 3A based on the data provided when purchasing. However, after examining the initial results, this was modified to 3.84 A moving forward, as this is the nearest whole multiple of the single layer cells nominal capacity above the mean measured capacity of the pouch cells, with the measurement mean being 96.9% of this adjusted nominal, compared to the first round of single layer cells value of 96.8%. Given the remarkable similarity of these numbers, it was felt that this was sufficient

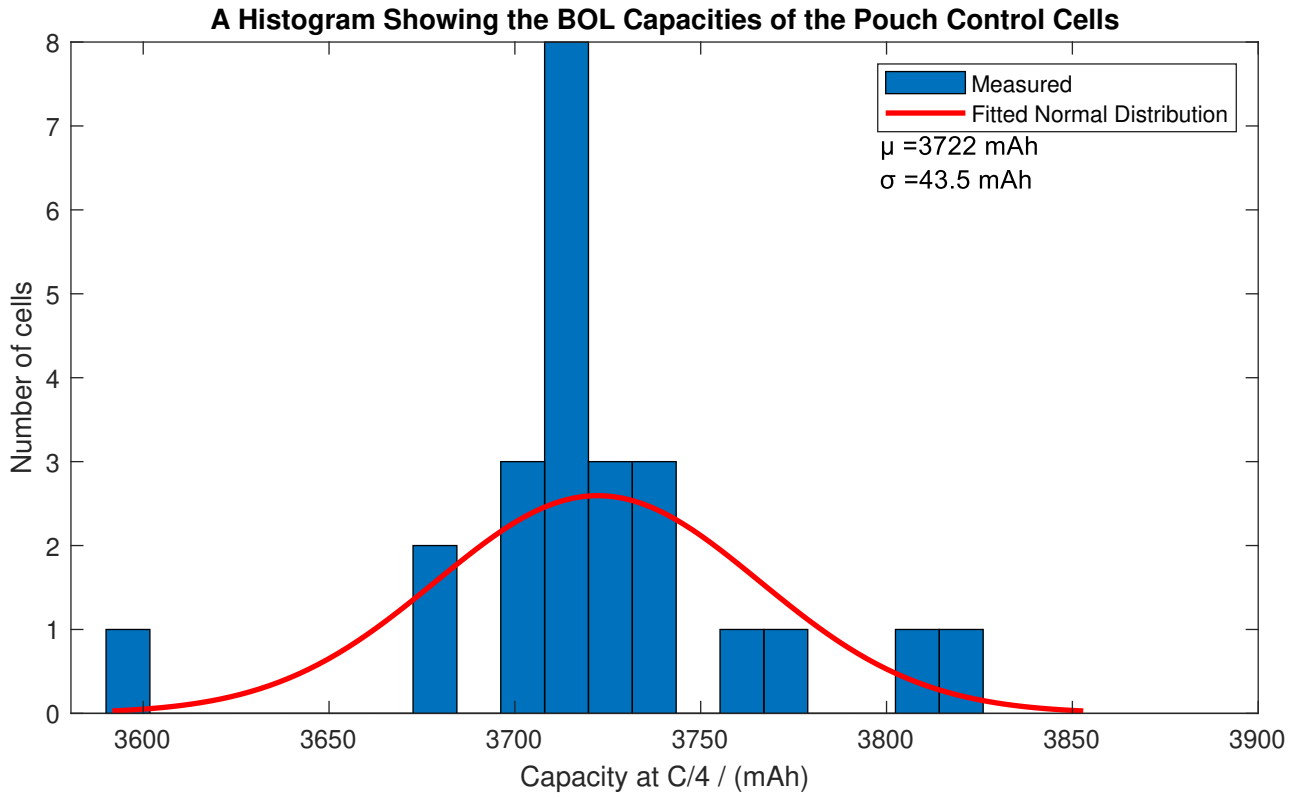


Figure 4.44: A histogram showing the BOL capacities of the pouch control cells

evidence to continue with this adjusted value, and that a more detailed analysis (such as a CT scan or other non-destructive imaging to count the layers of active material) was unwarranted. This work was performed concurrent with the second round of the single layer tests, but before the issues regarding calendar ageing were known. Had these issues been known, a more detailed evaluation would have been performed, but at the point they were discovered, this testing had already begun. It was felt that any adjustment part way through the experiments would be more disruptive and obfuscating than continuing with the selected values and taking the potential source of error into account during the analysis. Because all tests were done scaled to a nominal capacity of 3840 mAh, this will be treated as the nominal capacity for all further analysis and discussion from this point, unless explicitly stated otherwise.

While it might appear from first glance at the data in Fig 4.44 that these pouch cells had a larger cell to cell variability than the single layer cells, this is an illusion created by the difference in vertical scale of the fitted distribution. In fact the standard deviation of these results represents 1.13% of the nominal capacity, compared to 1.16% and 1.50% for the same

results from rounds 1 and 2 of the single layer tests, respectively. It is reasonable that the pouch control cells would have less cell to cell variability, as their manufacture is much more standardised. While this is true, the difference is less than might initially be expected. Whether they have also calendar aged or not would determine if a comparison to the round 1 or 2 single layers would be more apt, but even then the difference is small. The degree of calendar ageing these pouch control cells have experienced prior to testing will be evaluated later, based on the combination of the analysis techniques used. Because data from fresh cells is not readily available, as it was for the single layers (i.e. round 1 data could be used to evaluate the ageing of round 2 cells), this will be more challenging.

4.5.2 Capacity Throughout Testing

Beginning with an examination of the differences between cells exposed to identical conditions, our reference point is once again 30°C, 1C discharge. The data from these cells can be seen in Fig 4.45. The first feature of note is the increased inconsistency of the measurements, in particular the set taken around 460 equivalent cycles. It is unclear what the source of this is, but given that it occurs across all pouch cells tested at 30°C, it is likely to do with the environment they were tested in. It is not likely to be due to parameters such as pressure, as all of the cells presented in Fig 4.45 were in separate stacks. All three stacks of pouch control cells tested at 30°C were housed in a single chamber to maintain thermal boundary conditions on the surfaces not directly controlled via Peltier elements, and it is expected that there was an unexpected fluctuation in temperature during the relevant measurements. This was not found in any of the temperature measurements recorded, as they were all on the larger controlled surfaces of the cells, and so did not fluctuate outside of the measurement error throughout testing, but a separate measurement of the air temperature inside the enclosure was not recorded. It was never noted to be outside the expected range (30± 1°C), but as recordings (both digital and manual) were not taken, this is unfortunately anecdotal. The mistake is unfortunate, but does not invalidate the bulk of the data and can easily be accounted for when drawing any insights from these experiments.

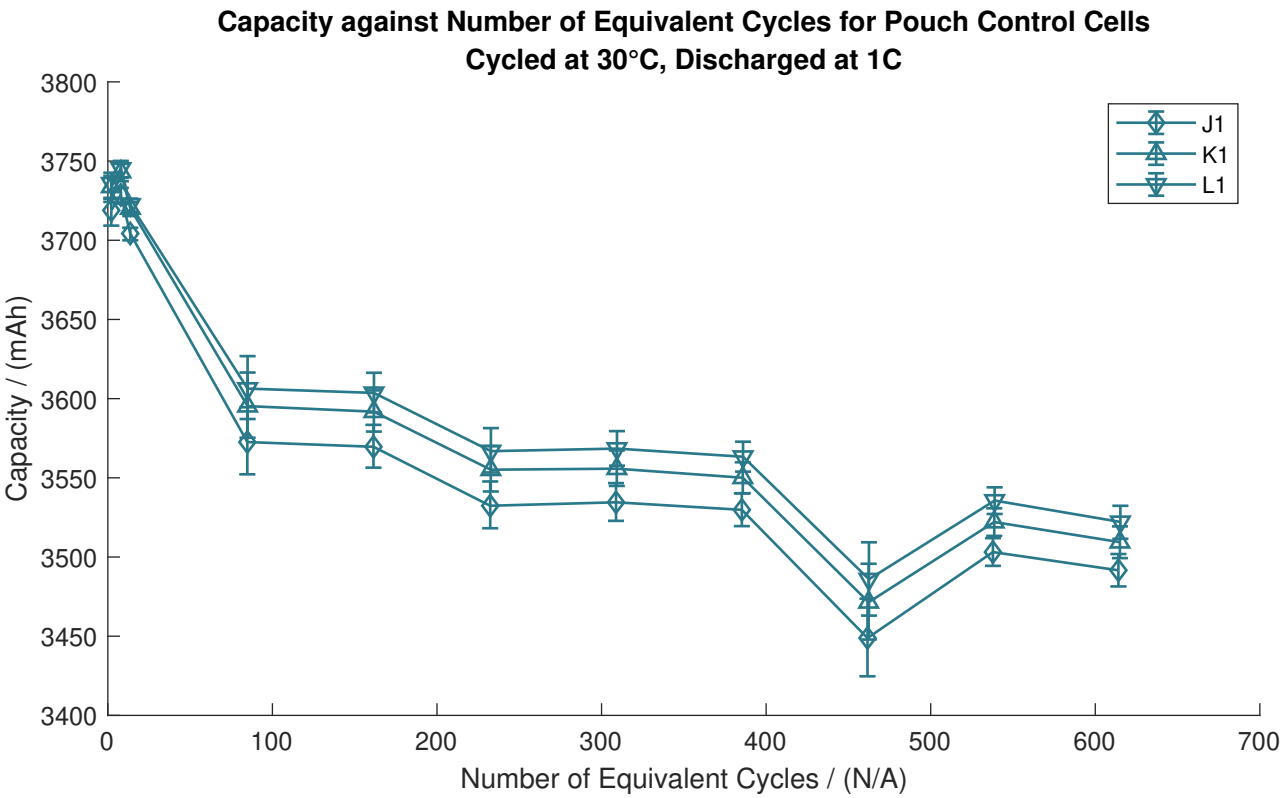


Figure 4.45: A plot showing the Measured Capacity against Equivalent Cycles of pouch control cells tested at 30∘C and discharged at 1C

The spread between cells over life is comparable to that seen in the single layer cells for the most part, with the majority of experiments falling in the 0.2-0.6% of the cell's nominal capacity, as can be seen in Fig 4.46. The exceptions to this are the lowest temperature (10°C) and the highest discharge rate (4C), which would be expected to be both the most erratic and the largest spread. These are larger, but this is also mirrored in the single layer cells, and has been discussed in Sec 4.1.

Next we can evaluate the impact of the variable parameters on the capacity of the cells throughout testing. Beginning with discharge rate, it can be seen in Fig 4.47 that these cells do not follow the same behaviour as the single layer counterparts. Instead, there is a greater difference between lower discharge rates (1-2C), with a much closer similarity in the middling rates (2-3C), and then the extreme impact of discharge rate appearing at 4C. It should also be noted that these cells saw significantly less capacity loss than their single layer counterparts,

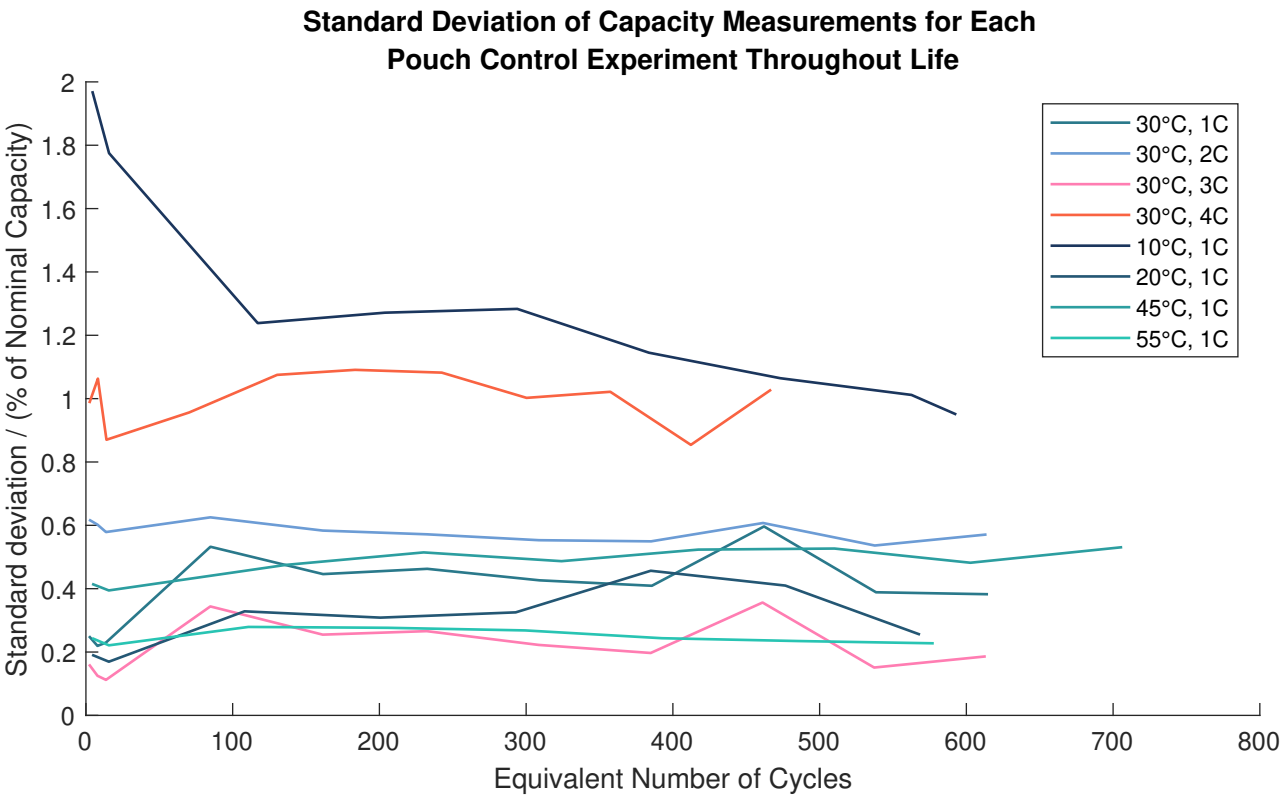


Figure 4.46: A plot showing the standard deviation of capacity measurements against equivalent number of cycles for each pouch control cell experiment

which was unexpected. It was predicted that they would see significantly more, as the thermal gradients would induce a greater variation in conditions and potentially cause a higher effective charge throughput of the active material due to the different regions or layers balancing as they return to more homogeneous conditions during a rest period following a load. It is unclear why there is such a marked difference between these cells and the single layer ones in this regard, and why it is in the direction observed. It can be seen from the distance between the measurements in the x direction that only the highest discharge rate hit the lower voltage limit during cycling, as intended, although this did occur from the beginning of testing, which was not. The intention was for this limit to be reached in the later part of the testing due to the loss of capacity and rate capability of the cells, but as these limits were set primarily to achieve this in the single layer experiments, it is unsurprising there is some difference in their behaviour here.

Looking at the data from changing temperature, the picture is significantly less clear. As can be seen in Fig 4.48, some experiments (in particular the higher temperatures) were very

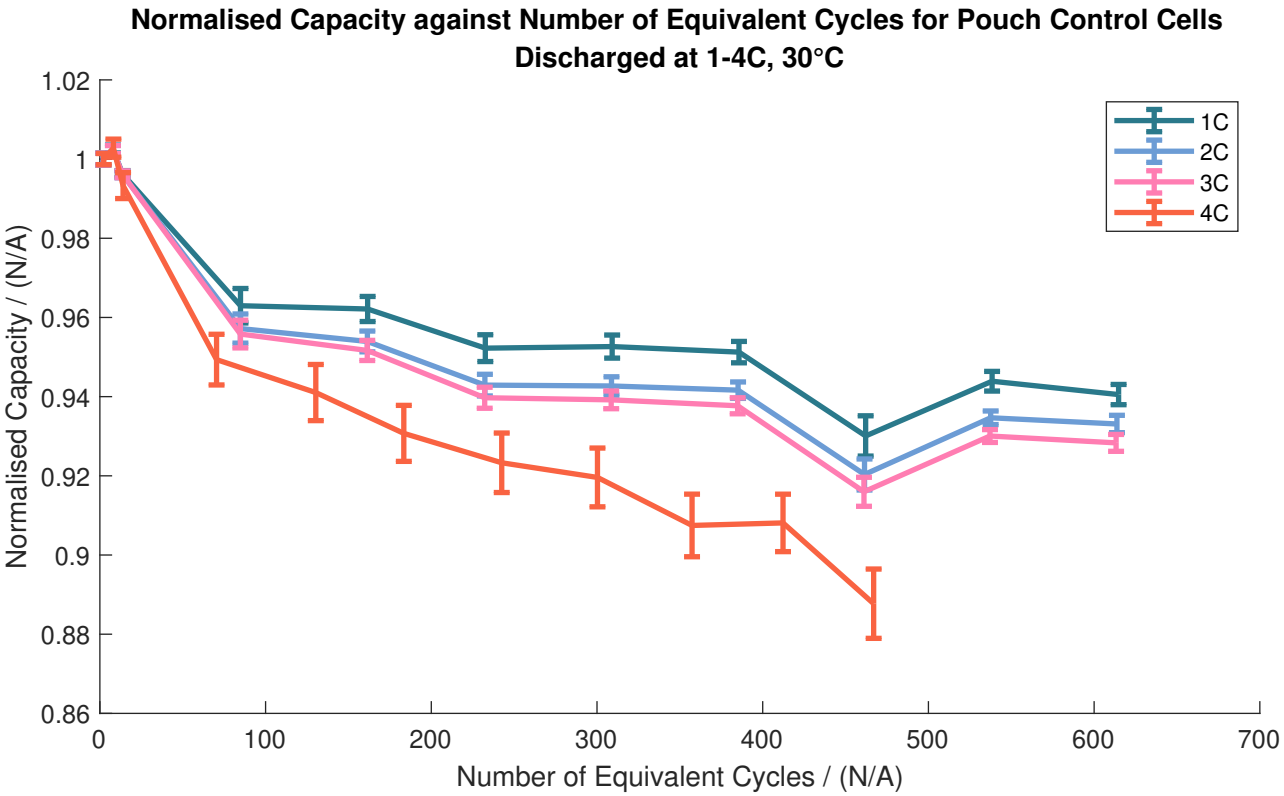


Figure 4.47: A plot showing the Normalised Capacity against Equivalent Cycles of cells tested at 30∘C and discharged at a variety of C-rates

well behaved, giving tightly grouped results with a smooth, consistent trend. The next two temperatures are somewhat more erratic, with significantly larger differences between consecutive measurements (shown by the error bars representing standard deviation across multiple measurements on multiple cells), as well as from characterisation to characterisation, with the lines following a smooth general trend less closely. The source of this variation is unknown, and has been briefly discussed for the 30°C experiments. The 20°C were in a more tightly controlled thermal chamber, however, as they shared space with the round 2 single layer experiments. It is possible that either due to irregularities in airflow around the stacks, or because of temperature fluctuations in the coolant fed into the chamber as a heat sink/source for the Peltier devices the bulk temperature of the rig was somehow unstable, but none of the recorded data shows any sign of this. All temperatures of both the pouch control and single layer rigs in this thermal enclosure were within measurement error of the set-points at all times, with the exception of the time spent acclimatising between characterisation and cycle temperature shifts. However, all measured temperatures were well within the set-points for

plenty of time before any cycling or characterisation measurements were taken, so it is unlikely the problem was due to the thermal mass of the rig lagging behind their measurements. The 20°C pouch cells were all in a single stack, as opposed to split across 3 stacks as in the 30°C experiments, so it is possible that there were some variations in pressure, but no significant difference was found when measuring the compression during disassembly. In addition, the 45 and 55°C experiments were also performed in a single stack each, and did not exhibit this issue.

Finally, perhaps the most striking set of measurements were those taken from pouch control cells cycled at 10°C, 1C discharge. While the trend of the line is relatively smooth, with the exception of the final set of measurements, the error at each data point is significantly larger than the other pouch control experiments. This is primarily from a large spread between the cells, although the error bars on each cell are considerably larger than other pouch control experiments also. They have been left in as there is no clear single outlier in this group, and because the absolute magnitude of the error on this normalised data is on the same order of magnitude as the single layer results, with the main difference being that the total capacity loss is significantly smaller, so the size of the error bars relative to the change being read is vastly different. Because the error is significantly larger than the difference in signal change between characterisations, these measurements will not be used to draw significant conclusions on the impact of temperature on the degradation of these cells. Instead, it is felt that the lower temperature experiments demonstrate the purpose of using the single layer cells for this study. While they do not provide a direct analogue in terms of rate of degradation, it is felt that for the purposes of more accurately detecting different degradation modes sensitivity to changes in various operating conditions (in this case discharge rate and temperature, but future work could look at any number of changes, including applied pressure, voltage windows, depth of discharge etc) they are a useful tool. In addition, it is believed they could help yield insight into the path dependency of different degradation modes, which will be discussed further in Sec 4.6.

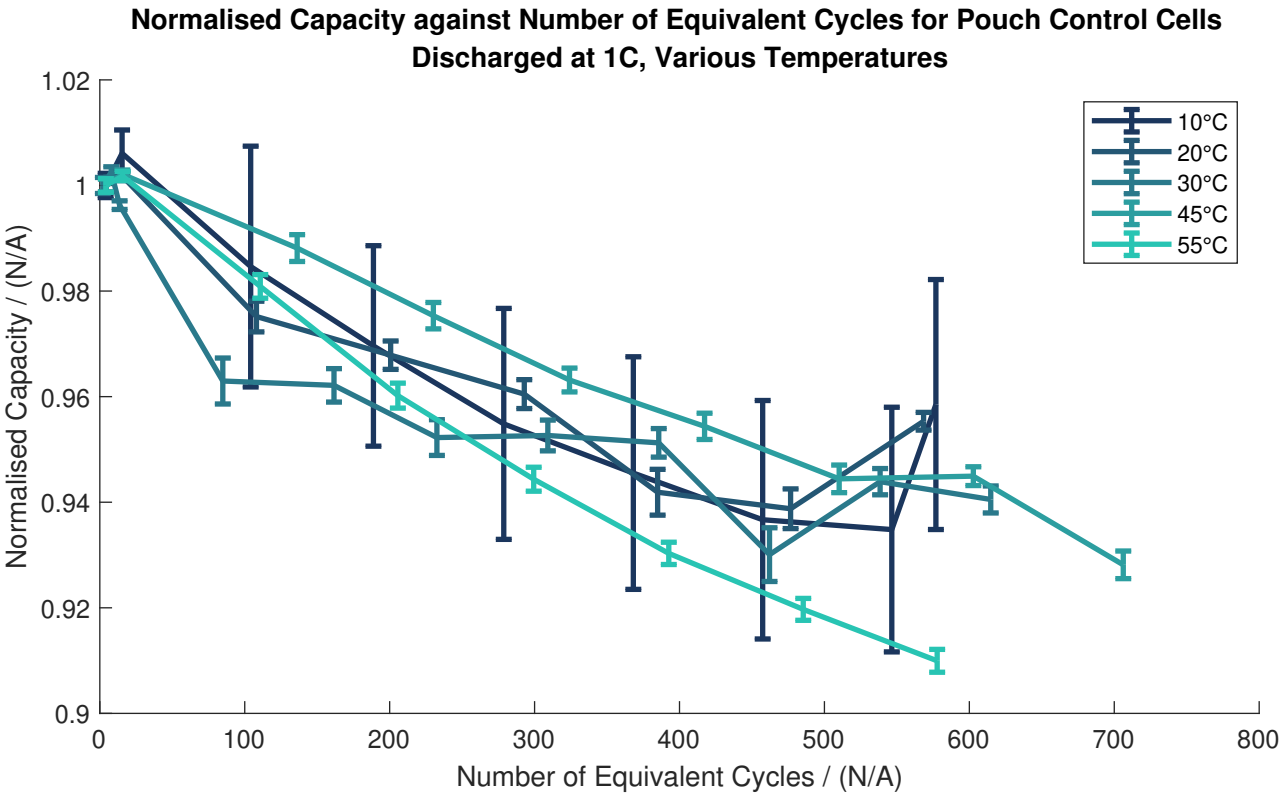


Figure 4.48: A plot showing the Normalised Capacity against Equivalent Cycles of cells discharged at 1C, various temperatures

4.5.3 End of Cycle Voltages

Next, we can evaluate the end of cycle voltages experienced by the pouch cells. Here we can see at least one source of the error in the 10°C experiments, which is that two of the cells were cycled at the wrong current for the first set of cycles. It is not clear how this happened, as it is obviously operator error, but the current at which they were cycled was not a selected cycling current. They were charged at the correct current, but discharged at 1.5A. While this will have contributed to the spread between the cells, the author is not convinced it is the primary source of the error. This is especially true because it would be expected that the lower discharge rate would have reduced the degradation these cells experienced during that first set of cycles, but they performed consistently worse than the correctly cycled cells, in both end of cycle voltage and in capacity.

On average, the pouch control cells exhibited end of cycle voltages approximately 100mV lower than their round 1 single layer counterparts, with the exception of those discharged at 4C,

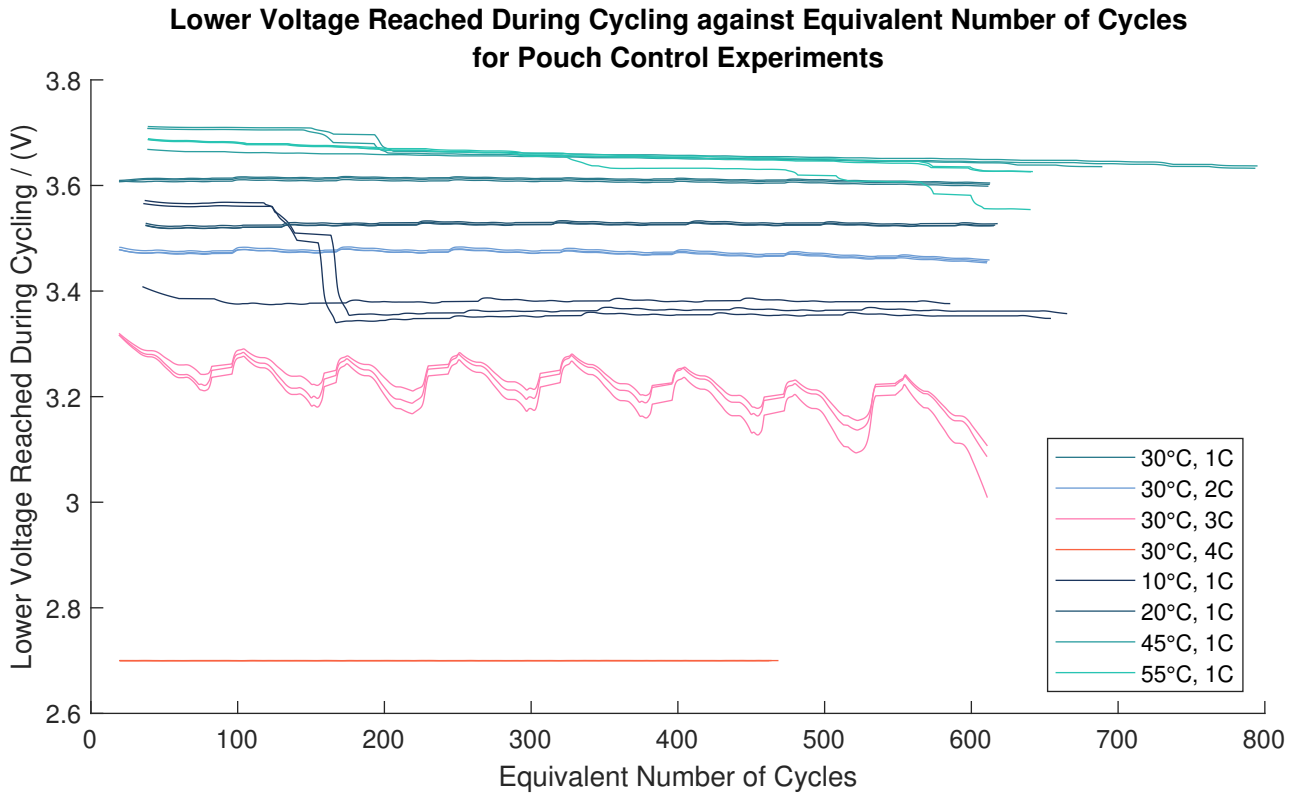


Figure 4.49: A plot showing the lower voltage reached during cycling against equivalent number of cycles for each pouch control cell

which immediately hit the lower voltage limit of 2.7V. Compared to the round 2 single layer cells, the pouch cells were still lower, but the difference was closer to 50mV. This demonstrates the impact of the unintended calendar ageing, but also that either the capacity with respect to voltage is different, or that the nominal capacity used for these cells was an over-estimate.

4.5.4 Incremental Capacity Analysis

The ICA results from these cells retains the features expected from an NMC 532 vs GIC cell, but do look different from the single layer cells. As can be seen by comparing Fig 4.50 to the ICA degradation map presented in Sec 4.3, the pouch cells produce ICA curves much closer to those found in literature. In particular, there are two clearly prominent peaks, even in the discharge direction, with an inflection point signalling a very small peak buried between the two, the magnitude of which is not sufficient to make itself fully distinct in this data. This is likely due to a combination of the higher than optimal rate at which this data was gathered, combined with

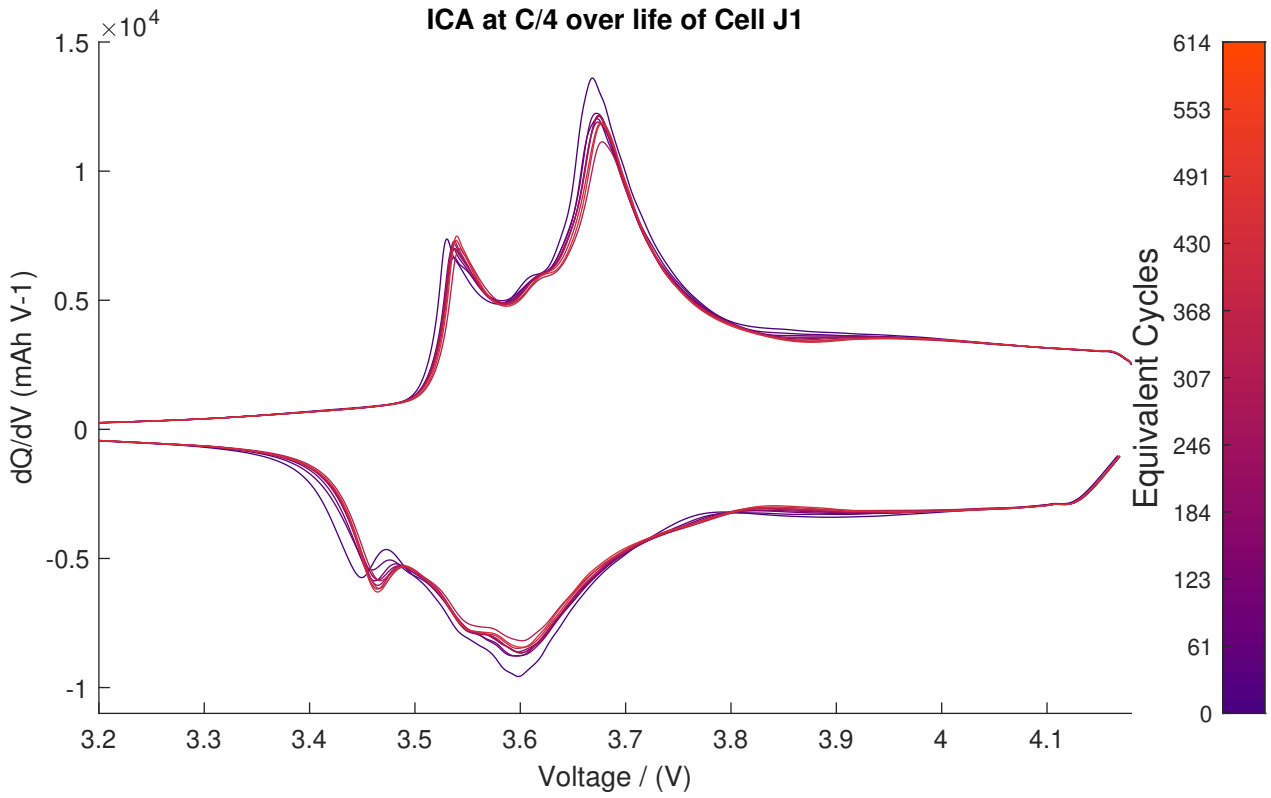


Figure 4.50: A plot of the ICA results of a pouch control cell cycled at 30°C, 1C discharge

the smoothing applied losing some of the very fine details. While it would have been possible to take capacity measurements from these cells at slower rates, and so yield more detailed ICA results, it was felt that it would be better to keep the test procedure identical between the single layer and pouch control experiments. Unfortunately, the magnitude of the shift in the ICA curves here is too small to make confident assertions about the degradation modes occurring.

As can be seen from Fig 4.51, the change in the ICA curve for cells which experienced more overall degradation during testing, such as those discharged at 4C, is greater, and some insight could potentially be gleaned from it. However, any conclusions would be somewhat tenuous, and the same insights can be found using the DMA process, in this case with a significantly greater degree of both accuracy and certainty. As such, the remainder of the ICA results will not be examined here.

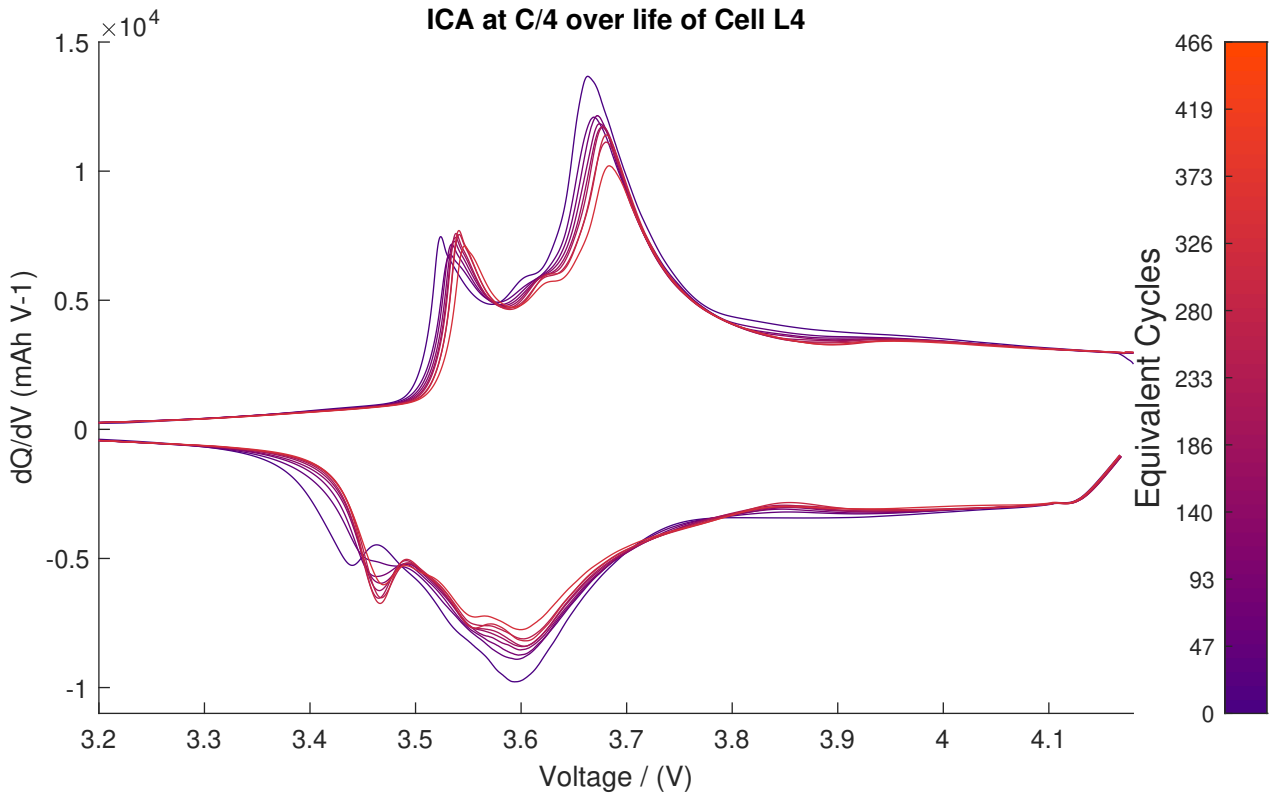


Figure 4.51: A plot of the ICA results of a pouch control cell cycled at 30°C, 4C discharge

4.5.5 Degradation Mode Analysis

As plenty of DMA results have already been presented in Section 4.4 and the results presented in this section are sufficiently similar, an examination of a single cell would be superfluous. Instead, the results from all pouch control experiments conducted at 30°C will be examined, with the intent of examining the impact of changing discharge rate. A plot of the lithiation fraction bounds for these experiments can be seen in Fig 4.52. There are a few features of note in this plot. Firstly, the PE upper bound and NE lower bound stay at their limits, as usual, with the exception of the cells discharged at 4C. These do drop by approximately 1-1.6% across all data points except the first. This is not thought to be significant, and is likely just an artefact of the DMA process. It could indicate that the CC-CV charge was insufficient to fully charge the cells, but it is felt that it is more likely to be a spurious reading in the very first result that is skewing the normalisation for this bound in subsequent readings.

Secondly, we can see that the NE upper bound does not leave its limit. This implies that

the NE capacity is always the limiting factor in determining the cell capacity, compared to the single layer cells where LLI became dominant partway through the testing. This is a key difference in these results, however it may be that these experiments would produce similar results if the cells were degraded further. Finally, because the PE electrode lithiation fraction lower bound follows the inverse trend of the cell capacity, as before, but the offset is smaller. The single layer cells began between 0.3-0.4 and rose to values between 0.4-0.5 over the course of the tests, the pouch control cells begin around 0.25. This is interesting, because we would expect it to be significantly closer to 0. Remember that the single layer cells had an uneven number of electrodes, which would not be expected in these cells. In addition, many commercial cells have overhang in the opposite direction, with slightly more NE capacity than PE. This helps prevent overcharge of the NE, of particular importance in reducing the impact of lithium plating.[8, 139] Therefore, either these cells were not manufactured with this NE overhang, or the calendar ageing they experienced was of both a sufficient magnitude and biased sufficiently towards the NE to remove it and offset the electrode capacities in the opposite direction.

Examining the electrode capacities calculated by the DMA process (presented in Fig 4.53), we can see the presence of this overhang. It can be seen from both the relative magnitudes of the PE and NE capacities, as well as from the offset. This does assume that the process calculates reasonably accurate capacities for the PE, despite it not being the limiting factor in determining cell capacity, but this is felt to be a reasonable assumption given that it is the PE that contributes many of the features exhibited in the OCV, as the magnitude of the changes in voltage across its range of lithiation is significantly larger than those present in graphite. In addition, the offset between the electrodes demonstrates that the loss of capacity in each electrode is relatively even across the testing, as it can be seen that each electrode loses 300-350mAh, with the offset not varying more than 25mAh from its initial value. This represents 6-7% of the PE initial value and 8-9% of the NE, with the maximum change in offset being 0.6% of the cell's nominal capacity. In comparison, many of the single layer experiments experienced a very similar change in the offset, approximately 25mAh, but they are of course much smaller in total capacity, and so this represents over 15% of their nominal capacity, greater

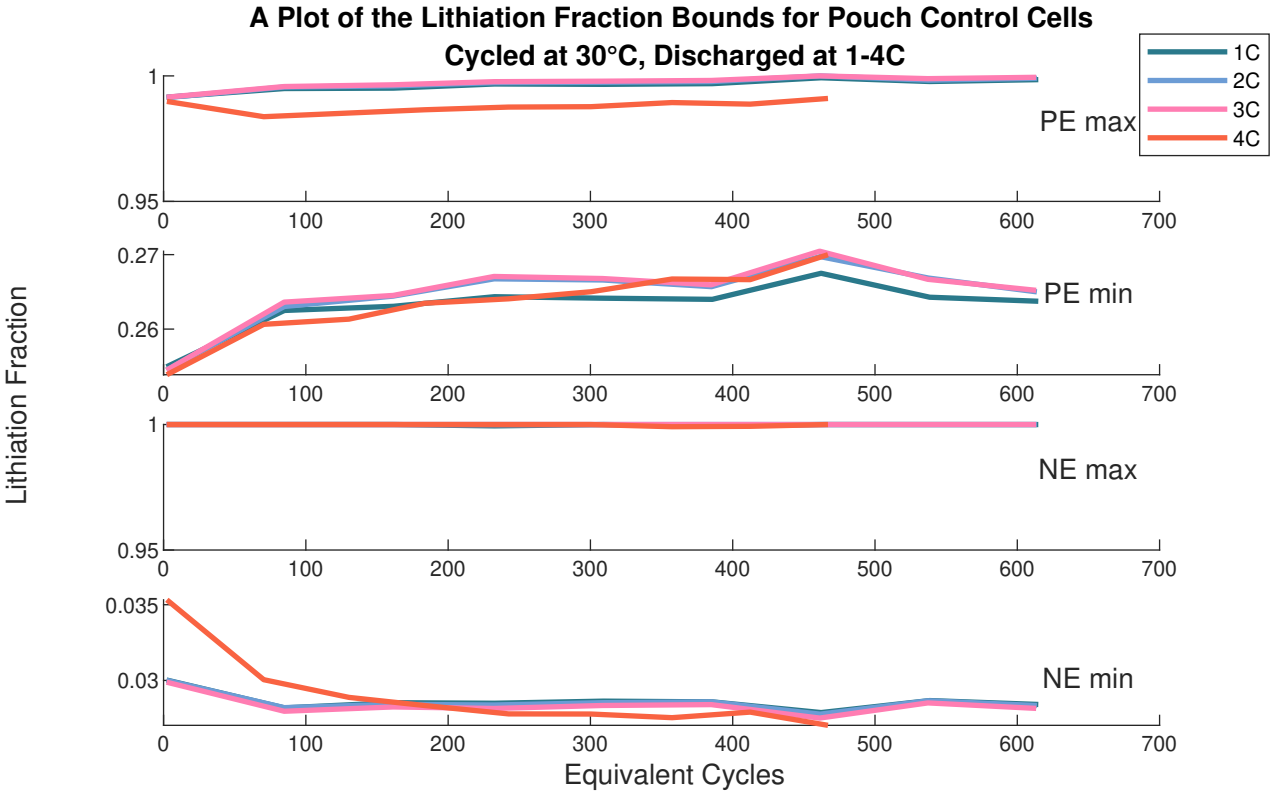


Figure 4.52: A plot of the lithiation fraction bounds against charge throughput of pouch control cells tested at 30°C and discharged at various rates

than the pouch control cells by a factor of over 20.

Examining what this means for the degradation modes, it can be seen that LAM at the NE and PE do correlate very well, with LAM at the PE appearing lower in Fig 4.54 because of its larger initial value. LLI is slightly higher than both, but not by a large enough factor to limit the cells capacity. Instead, the dominant degradation modes are the LAM at both electrodes, with the offset between them accounting for very little total loss (as can be seen from the fact that the NE electrode experiences a full 0-97% lithiation in Fig 4.52). This is perfectly normal, as fully de-lithiating graphite yields very little additional capacity and the voltage changes very rapidly, making it easy to over-discharge the electrode with a very small change in capacity. This is another reason the NE is typically manufactured with slightly more capacity than the PE, as it otherwise has a relatively flat voltage profile with respect to lithiation. It can also be seen that the increasing discharge rate has a fairly uniform impact on all three modes until 4C, where it has a disproportionately higher impact on LAM at the NE than at the PE, or on LLI.

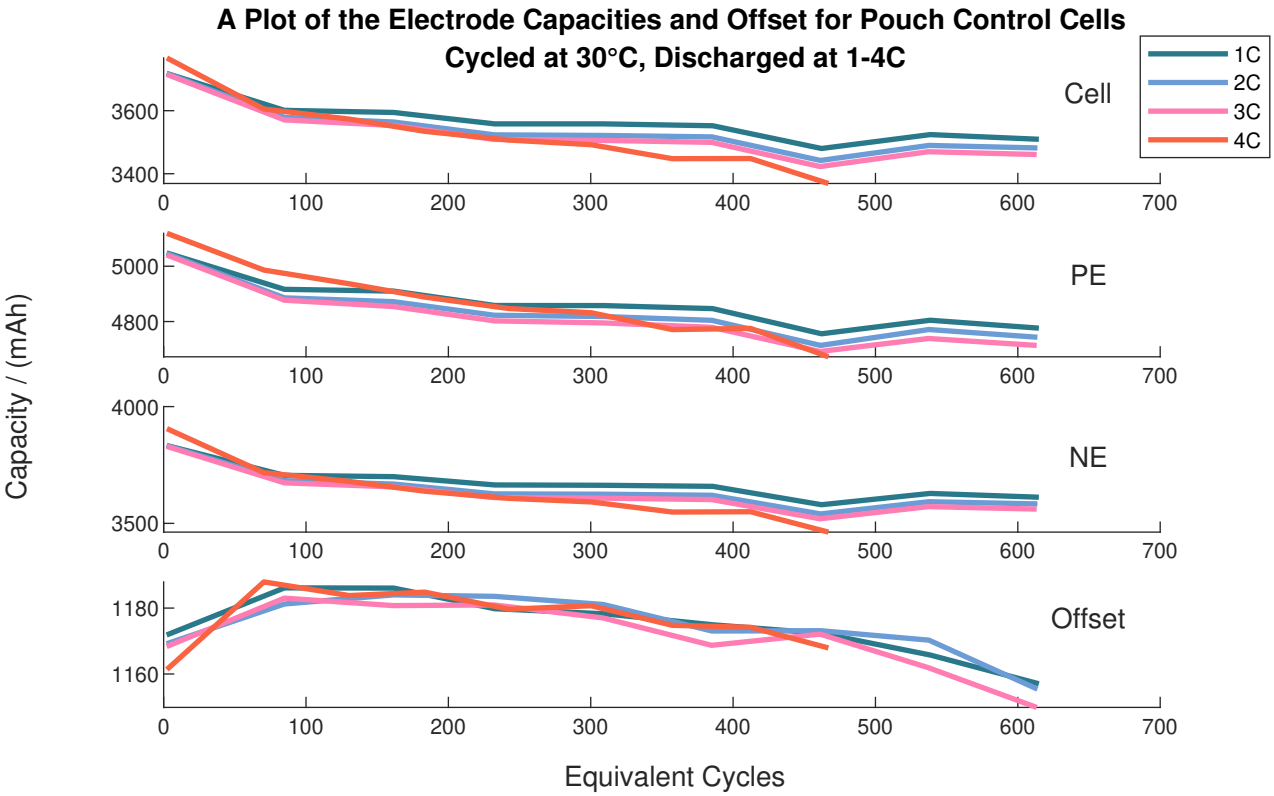


Figure 4.53: A plot of the electrode capacities against charge throughput of pouch control cells tested at 30°C and discharged at various rates

The last set of data to examine is the pouch cell control experiments performed at a range of temperatures. As mentioned in Section 3.3.7 there were not enough cells to perform these at both 1C and 4C discharge rates, so only 1C experiments were performed. The lithiation fraction bounds are shown in Fig 4.55, and it can be seen that the upper bound for the PE is no longer adhered to as strictly by the fitting process. The magnitude of this difference does not necessarily indicate that the PE is limiting the cells capacity, and it is unclear what might have caused this discrepancy, but it indicates that the capacity measurements did not fully discharge the cell, with 1-2% capacity left in the PE at the end of the measurements. This does not indicate a problem in the test procedure, as the appropriate voltage limits were reached, but instead that these cells potentially had a larger impedance, with the overpotential present when the voltage limits were reached responsible for the last scraps of the electrodes not being fully lithiated. The lower bound of the PE lithiation fraction no longer follows the cell capacity as closely, especially for the higher temperatures, which indicates more significant capacity loss

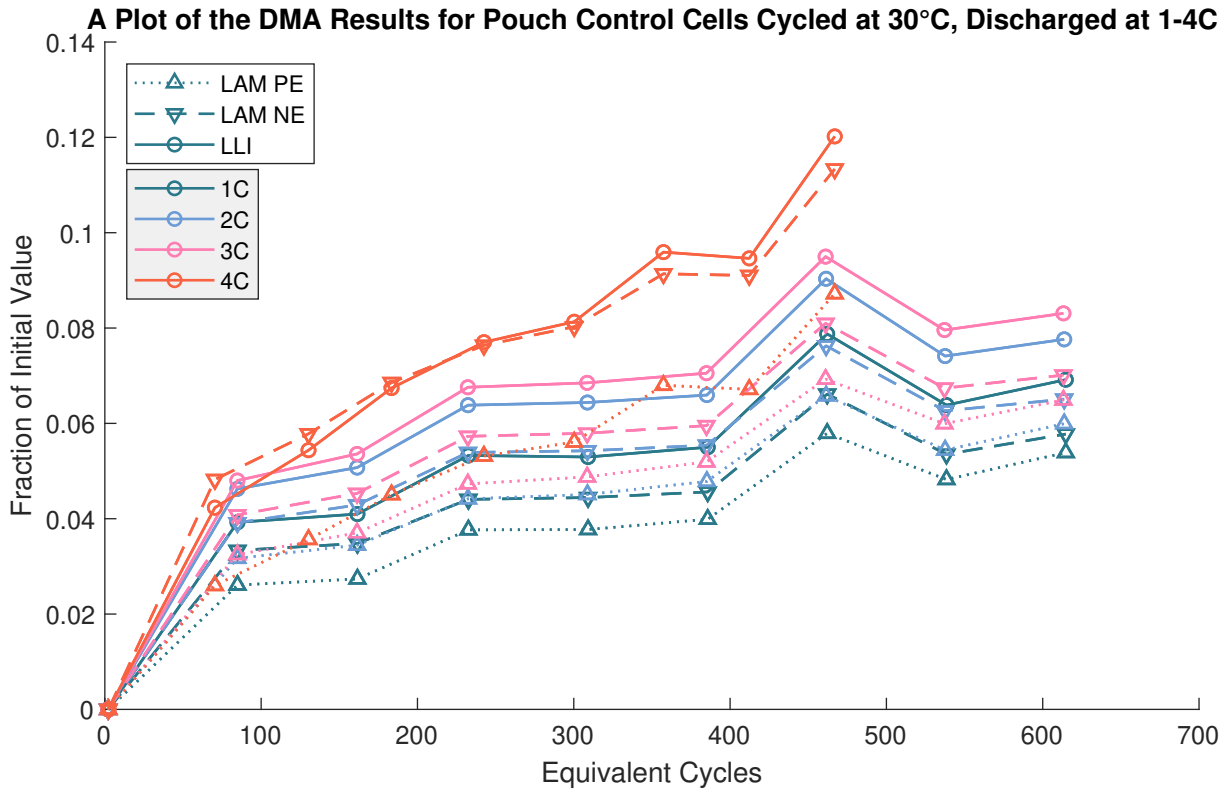


Figure 4.54: A plot of the DMA results from pouch control cells tested at 30°C and discharged at various rates

in the PE than other experiments. The NE upper bound does stay at its limit throughout testing for all these cells, again indicating that LLI does not limit the cells capacity, as there are sufficient charge carriers to completely lithiate the electrode when fully charged.

From the electrode capacities, presented in Fig 4.56, we can see that these cells follow the results from the 30°C tests. There is not significant difference between the different temperatures to read any meaningful conclusions from them based on electrode capacities alone, but the offset between electrodes does show a drastic difference. It can be seen that the offset correlates inversely with temperature, and that the higher temperatures do show a clear evolution of the offset. Especially interesting is that this evolution is in the opposite direction to that observable in the single layer experiments, with it decreasing over the course of the pouch cell testing. This again demonstrates a significantly higher loss of capacity in the PE, outstripping the magnitude of capacity loss in the NE significantly. Specifically, the PE for the 45°C and 55°C experiments see a capacity loss of 430 mAh and 485 mAh respectively, compared to their NE

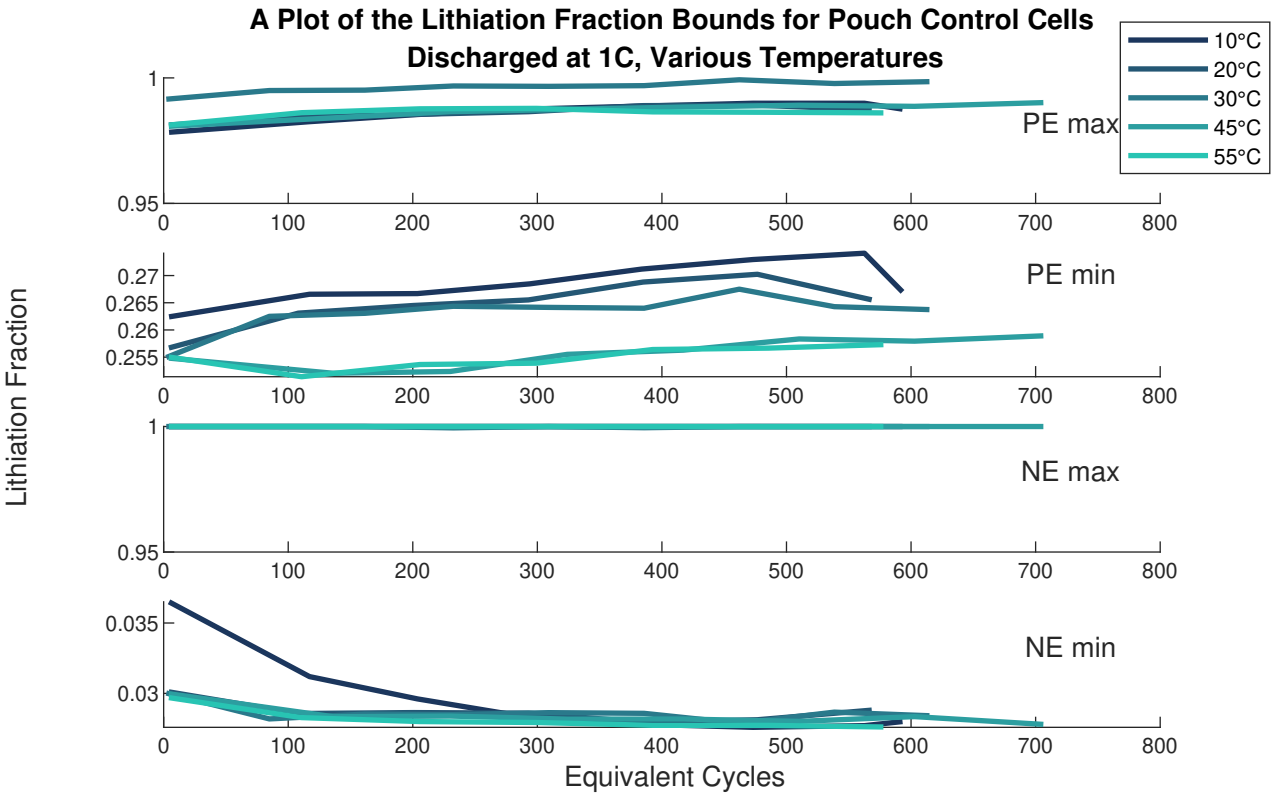


Figure 4.55: A plot of the lithiation fraction bounds against charge throughput of pouch control cells discharged at 1C, various temperatures

capacity loss of 300 mAh and 360 mAh. These are losses of 8.5-9.5% of the original value for the PE and 8.0-9.4% for the NE. This demonstrates a significant difference in the dominant degradation modes between the single layer and pouch control cell experiments, although this is admittedly mostly because of their construction, as opposed to from different reactions to the experimental conditions. The lower temperatures saw similar relative losses in the electrode capacities, but this did not translate into a shift in the offset between the electrodes, indicating that the effect was counteracted by something else. Specifically, it appears that the LLI at the lower temperatures matched the loss of capacity closely enough so as to cancel each other out with respect to the electrode offset.

This is confirmed by a plot of the calculated degradation modes, seen in Fig 4.57. The 45°C and 55°C are the only experiments where the LAM at the PE is larger than at the NE. This occurs throughout the testing, with no observable shift in the relative magnitudes of the degradation modes. However, the other temperatures are close, with relatively minor differences

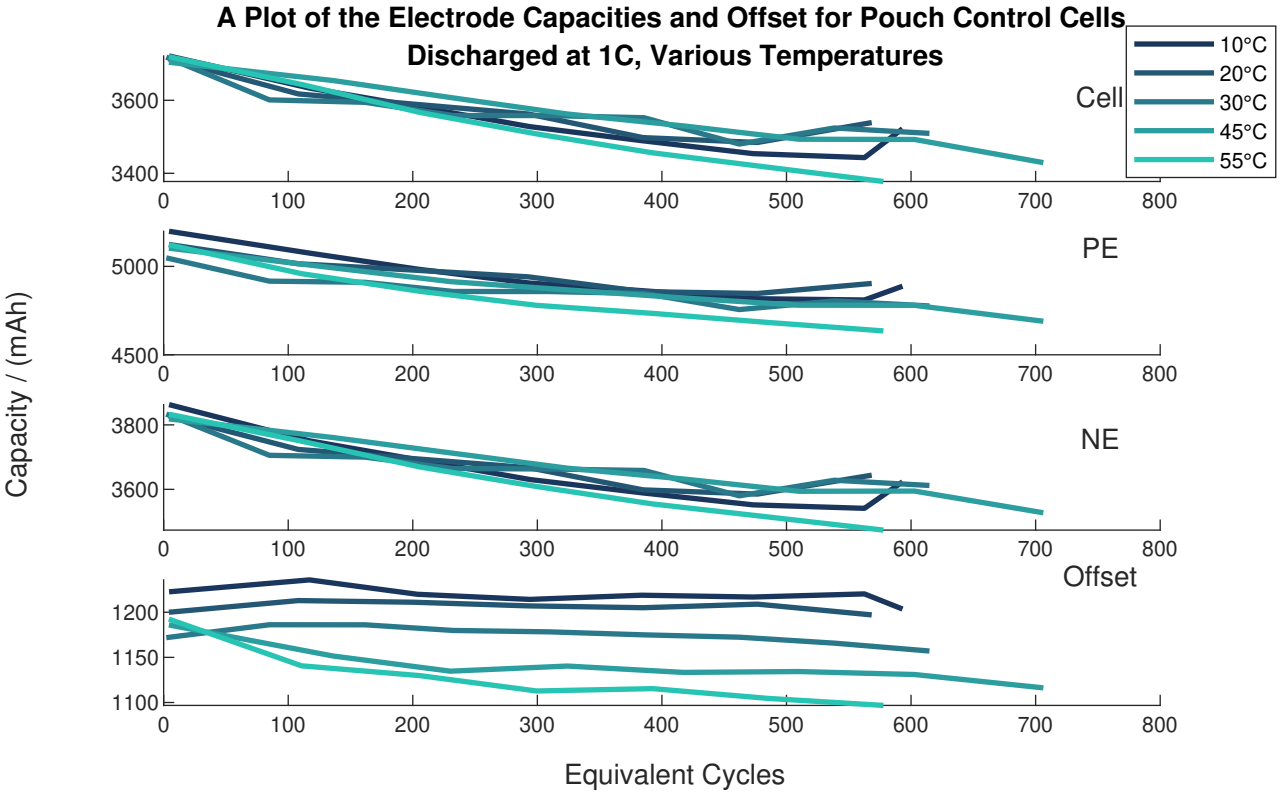


Figure 4.56: A plot of the electrode capacities against charge throughput of pouch control cells discharged at 1C, various temperatures

between the LAM at each electrode. The major difference between the temperatures is the difference between LAM and LLI, with the gap between them being much more pronounced at the lowest temperatures. This is especially true at the very lowest temperature, 10°C, which is unsurprisingly as it is expected that lithium plating would be strongly present at these temperatures, even given the moderate charge rate.

4.6 Discussion

Having presented the data gathered through the course of these experiments, as well as applied the relevant analysis techniques, we can begin a more thorough comparison of the impact of discharge rate and temperature on the degradation of these cells. To do so, let us first examine the expected degradation modes, and their relative impacts on cell performance. Fig 4.58 shows the most common and impactful degradation modes present in LIBs, and is reproduced with

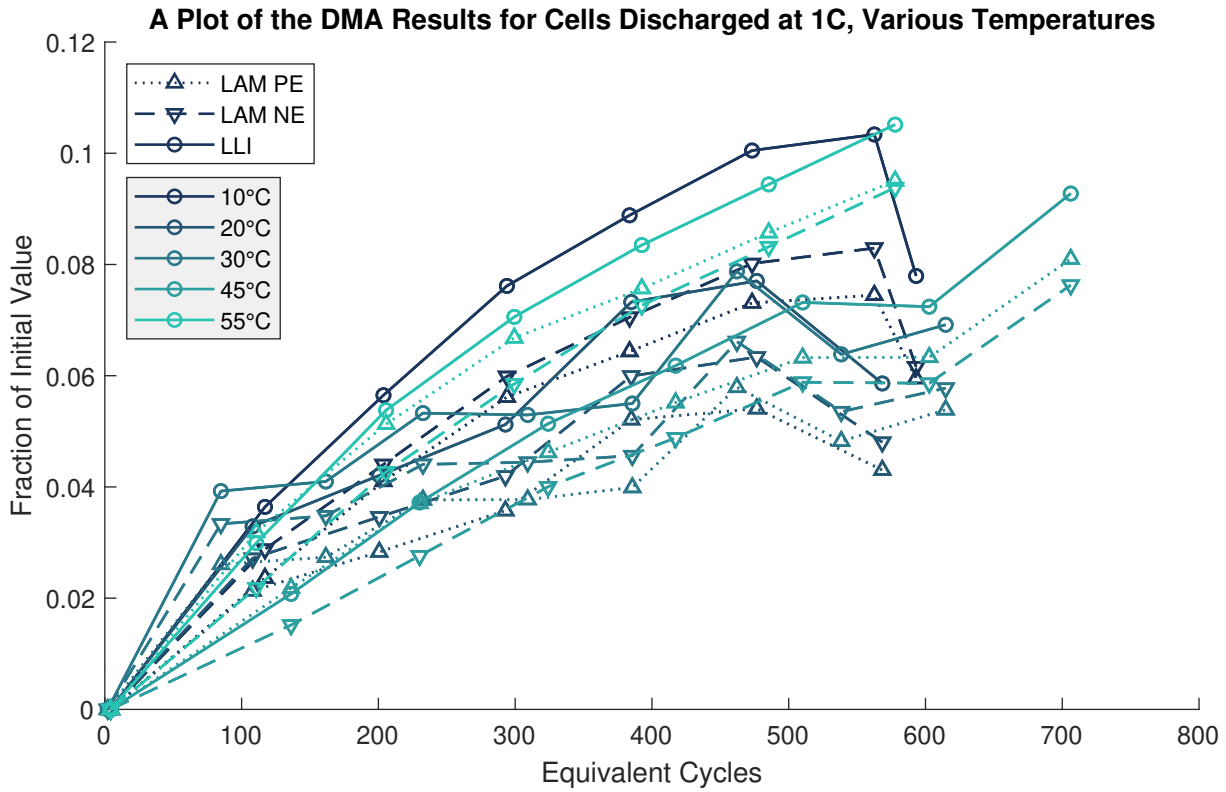


Figure 4.57: A plot of the DMA results from pouch control cells discharged at 1C, various temperatures

permission from.[8] Not all of these mechanisms will directly impact the cells performance in these experiments, but may significantly accelerate or retard other mechanisms that do. For example, particle cracking alone has relatively minimal impact on cell capacity, and primarily acts to change the impedance of the cell. This is because it disrupts the electrical contact between the active material and the current collectors. However, it can also lead to "island formation", where the fractures completely disconnect sections of the active material from the current collectors. When this occurs, the isolated volume of active material is no longer contributing to the cells function, any charge carriers intercalated in it will be removed from the accessible inventory, but it can still act to consume electrolyte and charge carriers through SEI layer growth and other parasitic side reactions. This can occur in large chunks, or can be the result of large groups of small cracks, which combine to pulverise sections of the active material. While particles are often represented in diagrams or models are spherical, this is for clarity and simplicity purposes, and does not represent the material produced by the majority of manufacturing methods. In practice, the particles are more irregular, often flatter and wider

due to both graphite's layered structure and the compacting and compression that is applied to ensure tight packing and good electrical contact between particles. This allows island formation to occur much more easily, as there is a smaller cross section required to sever in order to isolate volumes of the material.

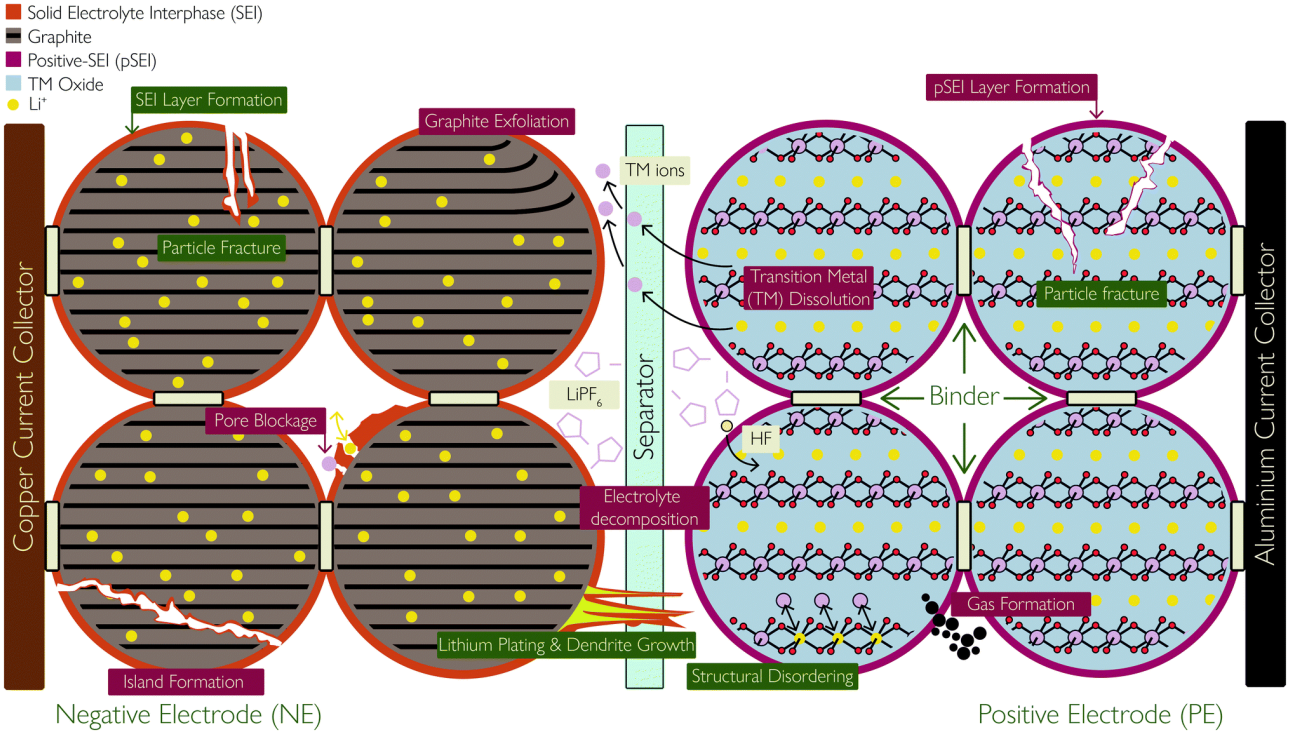


Figure 4.58: A diagram to show the mechanisms through which various degradation modes occur. Primary mechanisms are in green, with secondary mechanisms in maroon. Reproduced with permission from[8]

.

In addition, the surface area exposed by the fractures will form SEI, again acting to consume lithium inventory, reducing the capacity of the cell in extreme cases (as seen in the later stages of most of the single layer experiments performed), as well as reducing the conductivity of the electrolyte, therefore increasing the impedance of the cell. It is possible for the SEI layer itself to crack and separate due to the differences in material properties between the SEI and electrode materials. This acts to capture charge carriers, but then re-exposes the surface, allowing more SEI to grow at an accelerated rate, increasing the rate at which LLI occurs. The reduced capacity then acts to increase the effective rate seen by the remaining active material, as a lower capacity is available to fulfil the current demand, and so the current density in the material is

increased. This demonstrates how the various degradation mechanisms can act in a feedback loop, accelerating overall degradation. Particle cracking is exacerbated by any condition which acts to increase mechanical stress, which is primarily produced by the volumetric expansion intercalation compounds exhibit across their range of lithiation states, but can also be produced by improper mechanical mounting or containment of the cells (especially in form factors without a rigid can, such as pouch cells) as well as thermal gradients producing thermal stresses. This means that diffusion gradients in the active material will produce mechanical stress and therefore particle cracking, and the magnitude of these gradients will be dependent on the rate. Additives can produce stress because they exhibit a different rate or magnitude of volumetric change than the bulk of the active material. The most relevant and extreme example of this is the use of silicon in the NE. It can be used to increase the specific capacity of the NE, but exhibits a much more extreme change in volume between its lithiated and de-lithiated states (between 300-320%[89] compared to graphite's 12%[140]).

This increased growth of SEI can also reduce capacity by blocking the pores through which charge carriers migrate in the electrolyte. This prevents them from reaching surfaces through which they would otherwise intercalate, which can lead to increased charge transfer resistance (and therefore cell impedance), or act to prevent sections of the active material from participating in the intercalation reactions entirely even if they are in electrical contact, as there is no diffusion path between those sections and the opposing electrode through the electrolyte. This can again also lead to LLI, from lithium intercalated in the active material when it is lost, and also the charge carriers lost in the trapped electrolyte. A diagram showing the interactions between these various degradation mechanisms and modes can be seen in Fig 4.59, included with permission from.[8]

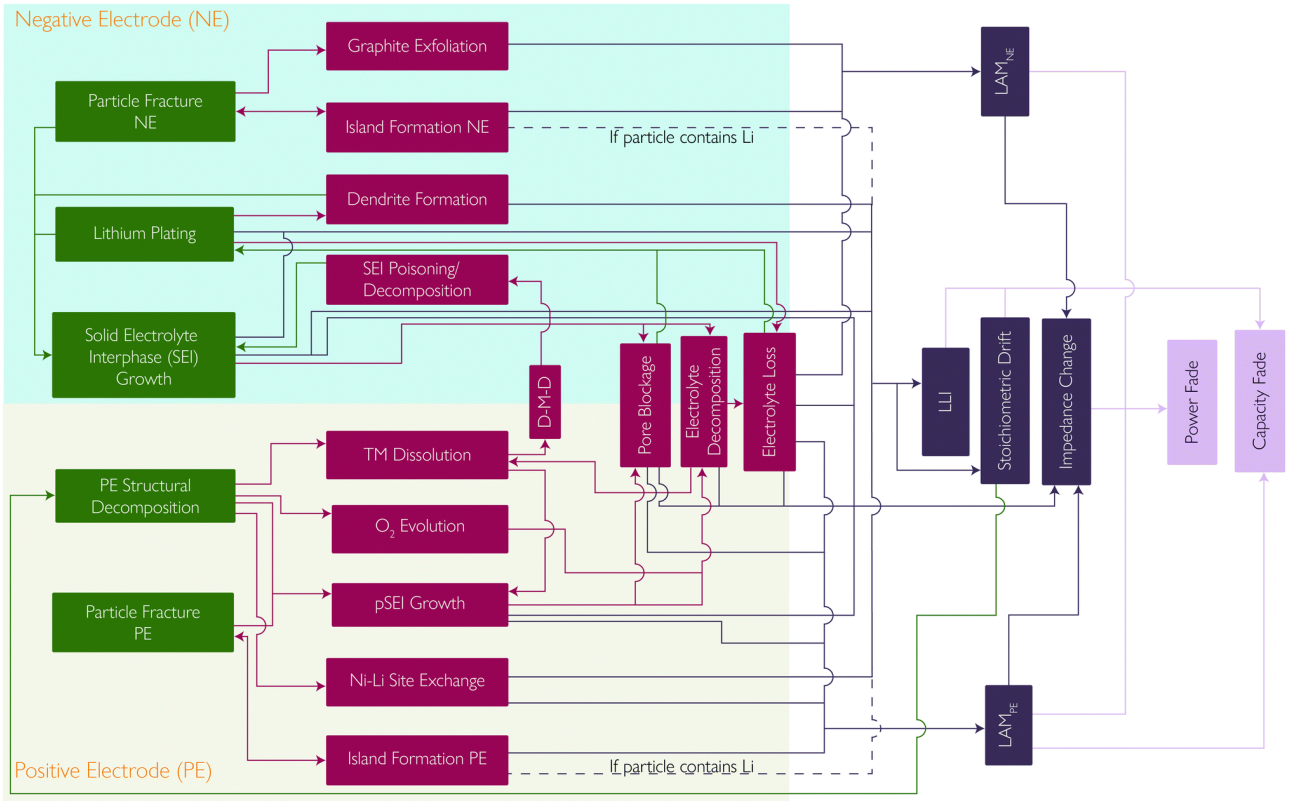


Figure 4.59: A diagram to show how different degradation mechanisms interact with each other, as well as their contribution to degradation modes and observable repercussions. Reproduced with permission from[8]

.

For single layer cells cycled at 30°C, increasing discharge rate had a minimal effect up to 2C, and from then on an increasingly dramatic increase in capacity loss was exhibited. The dominant degradation modes present did not change in the range of discharge rates tested here, with all cells losing capacity primarily to LAM at the NE from the beginning of testing, with a transition partway through testing to LLI being the limiting factor. This transition occurred significantly earlier at higher discharge rates, but is not solely responsible for the higher capacity loss. All of the cells appear to follow the same pathway, and the difference in the rate at which they lose capacity is fairly uniform across the testing. Due to the unorthodox balance of PE and NE active material in these cells, degradation at the PE is negligible, and does not impact the cells capacity at any point in the testing performed. LAM at the NE and LLI are identical from the beginning of testing until the aforementioned point where the dominant mode transitions. This implies a link between the two modes, as they are responding to changes in discharge rate identically. As can be seen in Fig 4.59, there are several ways in which this can occur. One

possibility is that the active material could be lost due to island formation from particle cracking, with these islands being in a lithiated state and therefore trapping lithium. The growth of SEI is present at these cracks, but the magnitude of the exposed surface is not large enough to allow LLI to overtake LAM. Alternatively, the LLI could be primarily from SEI growth, which in turn acts to block pores, severing regions of active material from exchanging charge carriers with the bulk of the electrolyte.

Due to the relationship between increasing discharge rate and the rate of capacity loss, as well as the fact that higher discharge rates cause LLI to become dominant earlier, it is proposed that the dominant degradation pathway in these experiments is the former. Remembering again that the particle shape and size in reality is not of uniform spheroids, but of flakes compacted together, tending to be thinner in one dimension than the other two, and of a range of sizes. At the BoL, particle cracking (caused by the mechanical stress of volumetric changes which occur when the graphite is lithiated) causes the extremes of these particles to break off, losing active material more rapidly at first as the more fragile or thinner edges of these flakes are isolated. This slows, eventually reaching a plateau, which in these cells is when approximately 20% of the initial capacity is lost. This can be seen from the DMA data presented in Section 4.4, Fig 4.31, where all of the LAM at the NE end between 19.4-20.8% of the initial value and does not follow any observable trend with increasing discharge rate. Pore blockages could also account for a portion of this LAM, but with the data gathered here it is not possible to distinguish its impact from island formation. Meanwhile the LLI continues, accounting for the differences between the experiments from that point onwards. It is for this reason we see the deviation of LAM NE from LLI earlier in the higher discharge rates, because the total loss of capacity is higher and so the LAM hits this ≈20% threshold much faster. Once the fragile extremities of the particles have been isolated by this particle cracking, the cracks continue to propagate, but no longer form islands as readily as the remaining particles have fewer readily isolated extremities. Instead, networks of cracks in the material form without forming islands, which exposes large amounts of new surface area for SEI to grow on. This SEI growth consumes the charge carriers, and is responsible for the continued capacity loss. It is posited that had testing

continued, there would be a sudden increase in the rate of capacity loss and LAM (termed a "knee point"), as is typically exhibited in LIB performance over life, and that this would be caused by the crack networks reaching the point at which they are fully pulverising the remaining material, with the fragments being isolated by the excessive growth of SEI. This theory allows relatively uniform rate of crack propagation across life, with some increase due to the loss of capacity requiring increasing current densities in the remaining active material to meet the total current demanded from the cell. This uniform rate of propagation does not produce a uniform rate of degradation because of the morphology of the active material particles.

The rate of crack propagation increases with increasing discharge rate, as the mechanical stresses increase due to more extreme diffusion gradients in the material. This relationship is non-linear, as below some threshold amount the mechanical stresses produced can only lead to fracture in the presence of stress concentrating geometries, of which there are many, but the rate is driven by their density more than by the discharge rate. However, above a threshold rate (which exists somewhere between 2-3C discharge in these cells) the mechanical stresses produced by the aggressive diffusion gradients are sufficient to overcome the material strength with minimal aid from stress concentrating geometries, and so increasing discharge rate suddenly has an exaggerated impact on cell capacity loss. This being the case, we would expect conditions which act to reduce the magnitude of the diffusion gradients to have minimal impact at discharge rates below the given threshold, but some initial benefit above these discharge rates. Increasing temperature is one such condition, and it can be seen from Fig 4.34 that this is the case, with cells tested at 45°C and discharged at 1C exhibiting almost identical degradation to those tested at 30°C, identical discharge rate. Unfortunately, because of the problems which occurred during the course of testing (enumerated in Section 3.3.6, it is difficult to determine if this holds true at 55°C. We can examine the difference between cells discharged at 1C and 4C at 20°C and 55°C for round 2 however, the DMA results of which are presented in Fig 4.60. Once again, it can be seen that LAM does not exceed 20% of the initial value, despite the different degradation pathway these cells experienced due to their calendar ageing prior to testing. It can also be seen that the impact of the increased discharge rate is significantly higher at 20°C,

whereas at 55°C there is a much more muted difference between the experiments. This holds with the proposed mechanisms of degradation present in these experiments, as the rate of crack propagation is minimally impacted by increasing discharge rate at the higher temperatures due to the reduced diffusion gradients in the material, and so the difference in exposed surface area from the fractures is small, and so is the LLI due to SEI growth at these exposed surfaces. At 22°C, however, the diffusion gradients are steeper, leading to more concentrated mechanical stresses and so the increased discharge rate causes significantly higher rate of crack propagation, and therefore the cells discharged at the higher rate reach the ≈20% plateau of LAM sooner, and the LLI due to SEI growth is higher throughout life, as there is significantly more surface area exposed by the fractures.

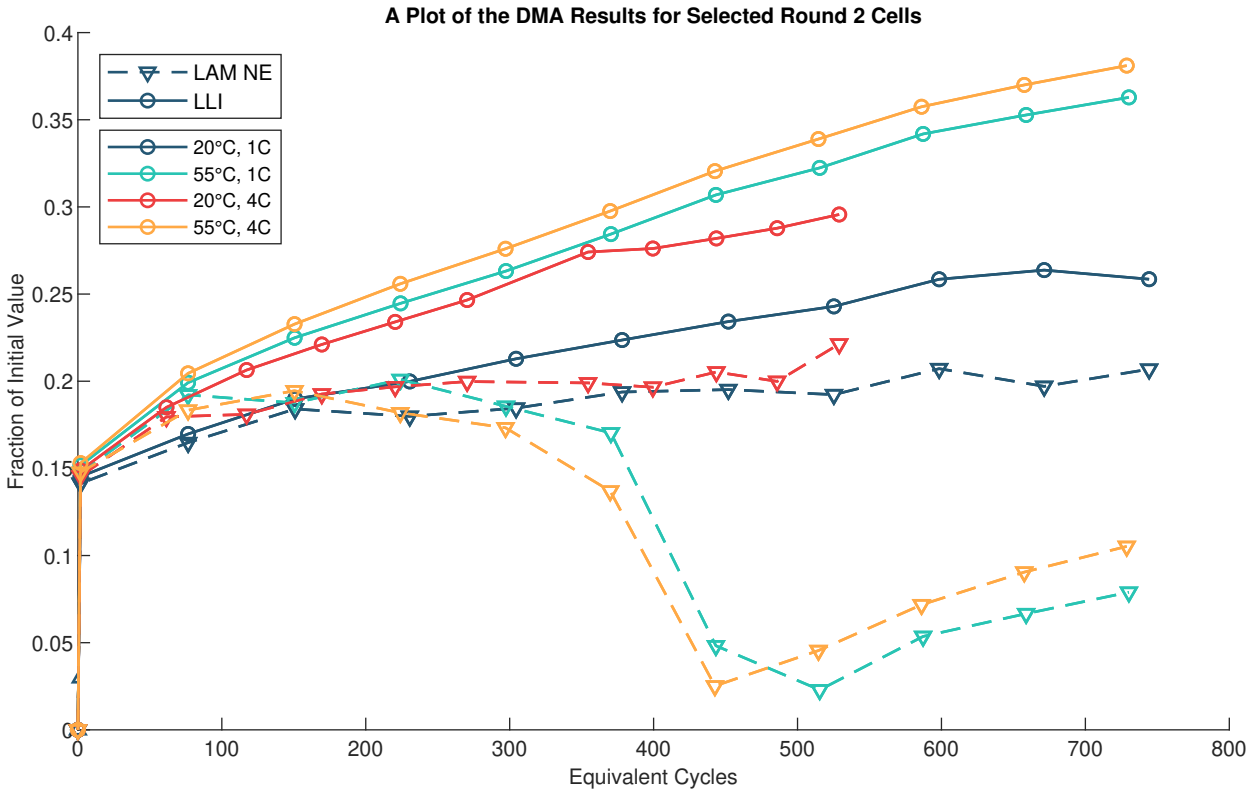


Figure 4.60: A plot of the DMA results against charge throughput of round 2 cells at temperature points with both 1C and 4C discharges

.

The only experiments which it is believed a different degradation pathway could become dominant are those performed at 10°C. Because of the low temperature, it is expected that

some lithium plating is occurring. This is not significant enough to cause extreme failure, such as from internal shorts brought about by dendrite growth puncturing the separator, but the deposition of metallic lithium at the active material surface can interact with the growth of SEI. In addition, if lithium is depositing on the surface, it means the diffusion gradients have hit a physical limit, limiting the magnitude of the mechanical stresses they produce in the material. However, LLI is still greater than experiments done at 20°C (as can be seen from Fig 4.40). this is because not all lithium plating is reversible, as SEI will grow on the metallic surface also, and the presence of the metallic lithium will act to accelerate SEI growth in its immediate vicinity. This combines to cause more lithium inventory to be captured by the SEI itself, but also the potential for islands of metallic lithium to be sealed by this layer, capturing further inventory at a rate not possible from SEI growth alone. One of the reasons discharge rate was varied, as opposed to charge rate (or both) was to limit the impact of this, as lithium plating is exacerbated by increases in the charge rate but not the discharge rate, as it is in the charge direction that the NE is being lithiated. This was largely successful, as the 10°C do generally follow the trend set by the other experiments, and there is minimal observable impact of lithium plating. In summary, it is expected that some lithium plating did occur at the lowest temperature point tested, but due to the reasonably modest charge rate (C/2) as well as the use of an extended CV hold to allow some of the reversible lithium plating to diffuse into the active material, the impact of this added degradation mechanism was minimal in these experiments, and acted very similarly to the mechanisms it might have replaced.

Finally, the pouch control cells did generally mirror the trends seen in the single layer cells, albeit at much lower rates of degradation. In addition, because of the more balanced initial capacities of the PE and NE they saw capacity loss approximately equal in magnitude in both electrodes, with the notable exception of the highest discharge rate (4C), where the LAM at the NE began to outstrip that at the PE, and initially LLI also. The significantly higher signal to noise ratio in these tests (especially at the lower temperature points) made it difficult to do a more detailed analysis, however, and due to a combination of experimental difficulties and the significantly lower rates of degradation in these cells compared to the single layers it is felt that

direct comparisons here would be spurious at best.

This does bring this section to its final conclusion, which is that the use of single layer cells for the purposes of investigating the interaction and decoupled impact of operating conditions on the degradation of LIBs is a useful and enlightening technique. It is hoped that the lessons learned from this project, as well as the test rig designed for this purpose will be useful for future work. The test rig design was given to the Multi-Scale Modelling consortium of the Faraday Institute, and at time of writing is in use for the investigation of rate, temperature and pressure in the degradation of various chemistries, as well as the comparison of these chemistries (varying PE, NE and electrolyte). In addition, it has been retro-fitted to aid the investigation of simulated multi-layer cells with controlled thermal gradients by connection of multiple single layer cells in parallel, with active thermal control (as was used in the pouch control experiments of this project) allowing a huge range of thermal gradients to be tested. The rigs worked as intended here, although some changes would be made based on the insights. Although pressure was intentionally selected to be in a range where LIBs are relatively insensitive to minor changes, for more complete collection of data it is suggested that future uses of this rig employ long term data loggers connected to the load cells, allowing the pressures to be measured during testing to ensure that this is not an additional source of error, as well as ensuring that the selected springs are configured appropriately to ensure that the volumetric changes in the cells do not act to significantly change the pressures during cycling. In addition, if higher pressures were required some minor modifications would be necessary, as the aluminium compression plates are not sufficiently rigid to ensure even pressure distribution above 300 kPa (at the current cell footprint). These alterations would also be necessary for cell footprints that were significantly longer than those tested here, as this obviously gives a larger distance between the points at which the compressive forces are applied (the tie rods in the corners) and so requires more rigidity to prevent a pressure gradient across the cell. While this could be achieved by making the plates significantly thicker, this somewhat defeats some of the design intents of the rig, which are to be as low cost as possible and to fit as many cell units as possible in a short height (i.e. small unit pitch in the stacks). A better solution would be to add some small spacers

located between units, with care to ensure that they are of a sufficiently thermally insulating material as to not thermally couple adjacent cells, as well as being either small enough or shaped correctly to ensure sufficient airflow between the units to maintain the desired temperature. This could easily be achieved with a small number of laser cut discs, identical to those used at the tie rods in the extant design, or of thin plastic ridges, depending on the configuration of the airflow over the stack. For the rig to be more self contained, it might also be prudent to mount fans directly to the stack to ensure this airflow. This was not necessary in these experiments, as the stacks were tested in thermal chambers with a large central fan at the rear (similar to a domestic fan oven), and could be oriented such that this airflow was forced to pass between the cell units, but were this not the case this fan could easily be mounted to the stack itself.

4.7 Conclusion

In commercial LIBs it is impossible to decouple the impact of C-rate and temperature due to the presence of thermal gradients produced by the internal heat generation of the cells. This is exacerbated by the non-isotropic thermal conductivities of the bulk material, in particular the significantly lower resistance in the plane of the layers comprising the cell compared to perpendicular to this plane. In order to allow an investigation of the degradation of these devices where the C-rate and temperatures are decoupled, custom devices were commissioned from a manufacturer comprising of a single symmetrical layer of active material, and a test rig designed to control the boundary conditions imposed on them. In particular, pressure was instrumented and controlled via the use of load cells, as well as the compression of a number of cell units together to eliminate the measurement error of the load cells and amplifiers as much as possible. In addition, the active material of the cells was held at a uniform temperature through the use of thermal masses applied in good thermal contact with the cells surface, which in turn were cooled by placing the rig in a thermally controlled chamber with good airflow over the remaining surfaces of these masses. This was validated, and in an extreme scenario significantly more demanding than the intended use was found that the maximum difference in

mean temperature between no-load and 4C continuous was less than 0.85°C, which is less than the peak to peak variation produced by the thermal chambers control system. In addition, the rigs were designed to minimise both cost and volume required, allowing a large number of cells to be tested simultaneously for minimal resources.

In addition to these tests, a set of control cells was also purchased, with active material from the same manufacturing batch as used in the single layer cells. These were tested under identical conditions, facilitated by a version of the test rig utilising active thermal management to impose the isothermal boundary. This active control was applied with the use of Peltier elements rejecting heat to water blocks, and managed via individually tuned Proportional + Integral + Differential (PID) loops. These maintained constant temperature at the cell surface within measurement error (measurements varied by a maximum of 0.01°C in use), regardless of the changing heat load from cycling the cells. These pouch cells had an identical footprint, but due to their more traditional number of layers had approximately 24 times the capacity of their single layer counterparts. Various difficulties occurred during testing, including the unintentional calendar ageing of approximately half the single layer cells and all of the control cells prior to testing, in major part due to the Covid-19 pandemic.

Despite these difficulties, it was demonstrated that increasing discharge rate was found to have minimal impact on the rate at which the cells degraded up to a threshold value, which for these cells was between 2-3C. Above this threshold, increasing discharge rate had a much more dramatic influence on the rate of degradation throughout life, but this influence was significantly diminished as temperature increased. At the highest temperature applied (55°C), there was only a small difference between the capacity loss of cells discharged at 1C and 4C, whereas at lower temperatures the effect became more pronounced. In addition, it was found that initially the LLI and LAM at the NE were essentially identical up to a value of 20% of the initial capacity lost, at which point the LAM NE plateaued, and the LLI continued to increase, becoming the dominant factor in limiting the cells capacity. This threshold amount remained constant regardless of the temperature or discharge rate applied, including the cells that experienced significant calendar

ageing prior to being tested. These relationships seemed to hold generally true for the pouch control cells, although the rate of degradation was significantly lower in these cells, likely due to the manual construction of the single layers. The data from the pouch control cells is also significantly less clear, both due to a higher variability in cells during testing (but not at BoL) and from the lower signal to noise ratio thanks to their limited capacity loss. This highlights the benefits of using single layer cells for this study, as well as the validity of the experimental procedure applied.

This work represents the first step in this new technique, and the rig design has been made available to various other researchers in the field so that it might be used to further investigate both this specific case (i.e. decoupling C-rate and temperature) as well as for other uses, such as the comparison of different chemistries (PE, NE and electrolyte), different material loadings, the influence of applied pressure, as well as the parameterisation and modelling of LIB's. This is considered by the author to be as important an output of the project as the specific data and conclusions presented, especially considering the limitations of said data due to circumstances beyond the projects control.

Chapter 5

Current Perturbations for PEMFC Systems - Background theory and experimental design

5.1 Introduction

This chapter will detail the work done by this project investigating the use of current perturbations (often referred to as "short-circuiting") to increase the power and efficiency of PEMFC systems. This technique is reasonably well known, and is the subject of a patent from Ballard,[141] as well as being included in the systems sold by Horizon.[142] However, the implementation has always been based somewhat on guesswork. The benefit of such a technique is that it has the potential to significantly improve the power density and efficiency of the PEMFC. This would allow PEMFC systems (such as HFCEVs) to be smaller for a given application, as well as requiring less fuel for a given application. This therefore also reduces the size of the fuel storage, compounding the reduction in total system size and weight, both significant factors in making a technology commercially viable. It is impossible to accurately quantify this potential benefit, however, as there are competing theories as to the mechanism through which the technique acts, and it has never been publicly demonstrated on a modern automotive style PEMFC stack. The aims of the project in this area of study were to determine if the technique functions on

such a device, quantify its benefit, provide evidence regarding the mechanism through which it functions, and finally demonstrate that it can be optimised live, *in operando*, with some determination of the quantified benefit of doing so.

This was accomplished by manufacturing a custom test rig to run a Powercell S2 stack with 24 cells, which differs from systems already in a variety of applications, including automotive, only in the number of cells present. This should have no significant impact on the use of current perturbations, and was done because the extra cells would yield no benefit while having a higher cost, danger and complexity. A full custom balance of plant (BOP) was produced, including a high level of instrumentation, as well as a cell voltage monitor (CVM) capable of measuring all 24 stacked voltages at up to 7 kSs^{-1} and a current perturbation controller. The purpose of creating the CVM was to allow measurement of the PEMFC's behaviour during and immediately after a perturbation to investigate the mechanism of the technique, as well as allowing additional validation of the length of each perturbation. The current perturbation controller controlled the switching of the perturbation circuit, allowing it to be controlled to sub μs accuracy, as well as being capable of running various control algorithms to optimise both the duration and the period of the perturbations. This was initially done using the "perturb and observe" (P&O) method, before verification of an analytical solution produced for the optimum value of the period, τ. The impact of these implementations is quantified, although a full analysis of the convergence is beyond both the resources and the scope of this project. Instead, the work presented here is intended to serve as a first step in an area of research the author feels has been neglected, to demonstrate that the technique work and provide some evidence as to how. With this in mind, some discussion of future implementations is presented based on the learnings of this project.

5.2 Proton Exchange Membrane Fuel Cells - Principles of Operation and Background Theory

5.2.1 Principles of operation

PEMFC's are a group of electrochemical devices which react a fuel and oxidiser electrochemically to produce electrical power. While there are some variations, such as High Temperature and aqueous PEM designs, by far the most promising and well developed is the traditional gas fed "Low Temperature" design.[18, 19] In such a device a fuel is supplied to the anode side of a cell, where it is ionised, aided by the presence of a catalyst. This cation then passes through a membrane, which is constructed such that it only allows positively charged particles below a given size to permeate it. Once through this membrane, the cations are in the cathode, which is supplied with an oxidiser, also gaseous. The anode and cathode are also linked by an electrical circuit, which accepts electrons from the anode and supplies electrons to the cathode reactions. The cations, oxidiser and electrons react, typically on a catalyst surface.[17, 143] The energy produced by this reaction is what provides the electrical potential, driving current through whatever load is attached to the device. The fuel used is predominantly hydrogen, although some other fuels have been tested with limited success, including reformate products.[144] Catalysts tend to be Platinum nanoparticles (or some alloy of Platinum)[145, 146] although a significant amount of work has been done to find catalysts with better lifetime performance, or catalysts which utilise less precious components,.[146–148] The membrane is most commonly a designer polymer utilising Perfluorosulfonic Acid reinforced with a structural polymer such as polytetrafluoroethylene (PTFE),[143] although finding support polymers with better mechanical properties, such as Polybenzimidazole (PBI) has also been studied.[149, 150] The most well known of these designer polymers is Nafion, but a wide variety exist which are constructed along the same principles.[151]

5.2.2 Mechanical design

In order for a FC to operate, the gaseous reactants must be delivered to a catalyst surface which is in good electrical contact with a current collector. This must be in close proximity to the exchange membrane, with another layer of electrically connected catalyst on the cathode side.[9, 17, 143] The cathode must not only supply the oxidiser but also remove the products. As such, PEMFC's are flow devices, as opposed to sealed units (e.g. Li-ion batteries). They are typically constructed from sandwiched layers compressed together, with the gaseous reactants and products conveyed to and from the reaction sites in a flow plate of some kind. This flow plate also acts as the current collector, and is typically bi-polar, meaning that it has the anode flow field of one cell on one side and the cathode flow field of an adjacent cell on the other. In older designs (pre-2010's) these were typically made from compressed graphite plates with machined flow channels in each surface.[17] Graphite was chosen because of its electrical conductivity as well as its resistance to corrosion. Due to the formation of free protons at the anode (i.e. the cationisation of hydrogen atoms), the materials inside the device experience an incredibly acidic environment, as well as humidity from the reactant products. As such, they must be able to resist this environment chemically. In more modern designs these bipolar plates are constructed from metal sheets, due to a range of benefits such as ease of mass manufacture, cost, mechanical robustness and reduced cell pitch.[152] A simplified diagram of the structure of a single cell can be seen in Fig 5.1

Metal plates are easier to mass manufacture for a number of reasons. First of all, the technology to stamp or punch sheet metal is incredibly mature, and has been used on a large industrial scale for well over a century. It is also much faster per operation, as it forms the desired part from stock material in just a few rapid operations.[153] In comparison, machined graphite plates must be carefully cut using a rotary cutter. This is typically a milling bit in a CNC mill, which must cut each desired flow channel one at a time, the speed of which is limited by the mechanical properties of the material being cut.[143, 153] Graphite is particularly difficult to process in this manner due to its hard but brittle nature, combined with its tendency to produce very fine but very conductive power when cut, requiring extra levels of protection for

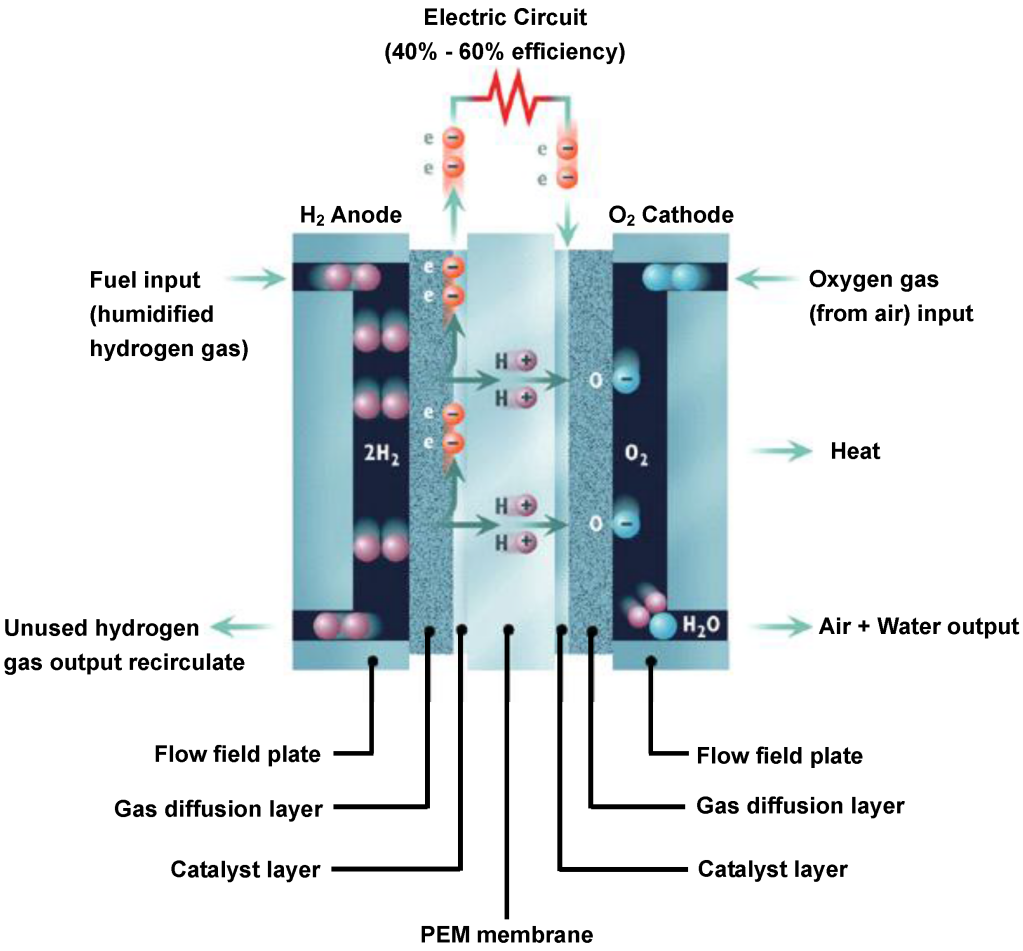


Figure 5.1: Diagram of the internal structure of a PEMFC from Ji *et al*[9]

the electronics controlling the process. In addition, due to the bi-polar nature of the parts they must be flipped to allow both sides to be machined, whereas punching a metal plate can form features on both sides in a single operation. In addition, stamped metal plates can be combined to form a flow field between cell, allowing the device to be liquid cooled. This, combined with better MEA's, has been vital in increasing the specific power of PEMFC's to increase viability for commercial use.[154]

There are some difficulties with the use of metal plates for this purpose. Firstly, the corrosion resistance required. This is typically overcome by treating the surface to passivate it in some way, and/or the use of highly corrosion resistant alloys, such as Titanium or stainless steel alloys.[152] Additionally, stamping has a lower limit on the feature size possible, and this can be limiting in the design of an optimal flow field. This lower limit is primarily linked to the thickness of the material stamped, which in turn cannot be reduced without sacrificing mechanical strength required for stable and reliable operation. One solution to this problem is to punch the required features out of multiple thinner layers of metal, and laminate these to produce the required flow fields and manifolds, also known as Roll-to-roll manufacturing or R2R.[155] This has the ability to reduce cell pitch (and therefore specific power) even further, as well as reaping the benefits of roll to roll production at scale (as opposed to stacking discrete parts).

Separating these plates is the Membrane Electrode Assembly.[9, 17, 143] This is formed from a number of layers of materials, each with specific roles. It is roughly symmetrical around the central membrane, although the catalysts used and properties of components can differ slightly between anode and cathode. Working in layers from the outside to this central plane of symmetry is the Gas Diffusion Layer (GDL), catalyst layer and finally proton exchange membrane. The GDL is formed of a porous mat of electrically conductive material, typically carbon nano-fibres.[156] It is responsible for allowing the relevant gaseous species to diffuse from the channels in the flow field to reach the entire surface of the catalyst and PEM.[157] The catalyst is made from nano-particles to increase surface area and therefore number of reaction sites, and is impregnated onto the GDL to allow electrons to flow to and from the

reaction sites via the current collecting bipolar plates.[145, 147] Finally, the membrane is at the centre, and the structure is then repeated on the other side, although details such as tortuosity or density of the carbon mat, size, loading and composition of the catalyst particles can be varied between anode and cathode side for optimum performance. For example, the catalyst loading on the cathode side is typically higher due to the Oxygen reduction reaction (ORR) being a slower reaction with a higher activation energy than that at the anode.[158, 159] This entire assembly operates best when it is well humidified, as this aids the transfer of charge carriers in the electrolyte.[150] However, it is important that water not block the diffusion of gases through the GDL, or the flow channels that supply the GDL. This occurrence is referred to as "flooding", and can cause significant degradation to occur if not properly managed.

In practice, the MEA is laminated from its constituent parts[160] and sealed between 2 flow plates, providing electrical isolation between them. This arrangement is repeated to create a "stack". The active area of a cell is proportional to the current it can provide, while the number of cells stacked in a single device determines the output voltage. Supply of the appropriate fluid flows is achieved by having the flow fields of each plate connect to input and output manifolds, which run through the entire stack.[153]

One of the limiting factors in the power a given mechanical design can produce is the mass transport of reactants to the catalyst surfaces. This is impacted by the design of the relevant flow fields, which are often formed of multiple parallel serpentine pathways that attempt to provide as uniform pressure at the boundary with the GDL as possible over the entire active area.[153] In addition, the density, structure and materials used in the GDL will play a role, as the gases must diffuse through the material in order to reach the catalysts. Too thick a GDL will cause significant diffusion losses, limiting current density within acceptable voltage windows, while too thin will prevent a uniform delivery of reactants over the entire active area, causing areas of "starvation", which rapidly accelerates degradation through a range of mechanisms.[17, 28]

The designs of the cathode flow field fall into 2 main categories. These are typically referred

to as "open" and "closed" cathode designs. In a closed cathode design, the cathode flows are treated similarly to the anode flows, in that they are supplied from an internal manifold running the length of the stack, and exhausted similarly. These are linked by a series of channels to provide a flow field, typically multiple parallel serpentine channels, and allow all the appropriate exchanges with the entire active area of the GDL.[153] This allows the cathode to operate at increased pressure, as well as allowing more control over other fluid properties such as temperature and humidity (provided an appropriate balance of plant). In contrast an open cathode flow field typically consists of many more straight channels which pass from one edge of the flow field to the other, allowing air to be blown directly through.[161] In this case the manifolds to provide uniform fluid properties to the entire stack are either external to the bipolar plates (i.e. an airbox similar to those used for ICE's), or may be forgone entirely, relying on fans closely mounted to the surface to provide adequate flow through the cathode.[142] This has the benefit of being cheaper and simpler to manufacture, but has the drawbacks of reduced control of the cathode fluid properties.[161] In addition, it is more likely to suffer from inhomogeneities in the cathode flow field due to differences in pressure from axial fans, eddies in the fluid flows, or blockages caused by external debris or particulates.[162] More fluid can be passed through the cathode, aiding in rejection of waste heat and water, but this heat rejection can no longer be controlled independently of humidity and stoichiometry. The importance of these fluid properties to the operation of the device is enumerated in following sections. In simplified terms, open cathode systems can often be lighter than closed cathode by forgoing much of the balance of plant required to condition the cathode flow, but cannot achieve as high a specific power within the stack due to the lack of control of the fluid properties, as well as being unable to operate in as extreme conditions (in particular cold starts), due to the lack of independent temperature, humidity and stoichiometric control.[161, 162] Because of this, open cathode systems rarely utilise an added coolant flow as they have added heat rejection via the cathode flow, lower current densities, and are typically used where system weight is a primary consideration. A popular example is use in low altitude aerospace applications (e.g. drones).[163]

Table 5.1: A table showing the details of the terms of the Nernst Equation

	Term	Definition	Value	Units
H	E	Reversible potential		V vs Reference
	E^0	Standard Potential		V vs Reference
	R	Universal Gas constant	8.31446261815324	J $K^{-1}mol^{-1}$
	T	Absolute Temperature		K
	F	Faraday constant	96485.3321	s A mol^{-1}
	n	Number of electrons transferred		mol
	$a_{Products}$	Activity of products		
	$a_{Reactants}$	Activity of reactants		
	j,k	Stoichiometric coefficients		

5.2.3 Electrochemical operation

The overall reaction taking place in a PEMFC is given in Eqn 5.1. Functionally, this does not occur as a single step reaction in a single reaction site, but rather as 2 half cell reactions linked by the charge transfer of an electrical circuit and electrolyte. These half cell reactions are the Hydrogen Oxidation Reaction (HOR) and Oxygen Reduction Reaction (ORR), and will be discussed further towards the end of this subsection.[19]

$$H_2 + \frac{1}{2}O_2 \rightleftharpoons H2O \tag{5.1}$$

The electrical potential produced by this reaction is given by the Nernst equation. A generalised form of this is given in Eqn 5.2 and the details of the terms in this equation are given in Table 5.1. Stoichiometric coefficients are the number of those species required for a balanced reaction. I.e. for the PEMFC reaction in Eqn 5.1, they are 1, 0.5 and 1 for hydrogen, oxygen and water respectively.[164]

$$\mathrm{E} = \mathrm{E}^0 - \frac{R \cdot T}{n \cdot F} \cdot \ln \frac{\prod \left(a_{\mathrm{Products}}\right)^j}{\prod \left(a_{Reactants}\right)^k} \tag{5.2}$$

For a PEMFC, the reactants species are gaseous, and so their activity is their partial pressures in atmospheres (atm). In the case of a pure hydrogen supply to the anode and air supplied to the cathode (with an oxygen content of 21%), this would be P_{Anode} and ($\mathrm{P}_{Cathode}$ x 0.21). PEMFC's operate below 100°C (or are designated "High Temperature" PEM, and are beyond the scope of this work). As such, the water produced is liquid (albeit dispersed in the fluid flows

as water vapour, but this is not thermodynamically equivalent to steam), and so its activity is set to 1. E^0 is equal to the sum of the standard state voltages of the half reactions of the system. These are defined as the standard state (i.e. temperatures and pressures) reversible voltages of any given half reaction (i.e. cathode or anode) against the potential required to reduce hydrogen under standard conditions. The anode of this device is of course reducing hydrogen, so has half cell potential of 0, and so the sum is just the standard state voltage of the Oxygen Reduction Reaction (ORR) at the cathode, +1.229V. Finally, the number of electrons transferred per reaction (n) is 2. With this, the generalised Nernst equation becomes Eqn 5.3.[164]

$$E = E^0 - \frac{R \cdot T}{2 \cdot F} \cdot \ln \frac{1}{P_{Anode} \cdot (P_{Cathode} \cdot 0.21)^{\frac{1}{2}}} \tag{5.3}$$

This gives the reversible potential of a PEMFC under a particular set of operating conditions. This reversible potential cannot be harnessed, however, as it represents the voltage at which both the forward and backwards rates of the reaction are balanced (i.e. no net charge transfer occurs, and therefore no current). Practically, this is defined as the "Open Circuit Voltage" of a device, although the value calculated theoretically tends to be an unattainable goal due to real world complexities such as parasitic side reactions, heat losses, leakage currents etc. However, it is reasonable to expect a modern PEMFC to achieve >1.1V OCV under standard conditions.[165]

While the net current density is null, this does not mean that no reactions are occurring. Rather, forward and backward reactions are constantly occurring, the relative rates of which are determined by the activation energy their species must overcome. Initially, the forward reaction would have a much lower activation energy, as the total reaction is thermodynamically favourable in that direction, hence releasing energy. However, without a pathway to transfer the charge carriers from one electrode to another, their charge will build up, producing a potential difference between the electrodes. As the species involved in the reactions are charged but must be in proximity with the electrode surface for the reactions to occur, the activation energies (and therefore reaction rates) are impacted by this charge. The OCV is simply the potential difference at which the forward and backward reactions equalise to produce no net charge transfer. The rates occurring in this state are defined as the "exchange current density", i_0.[166]

Table 5.2: A table showing the details of the terms of the Butler-Volmer Equation

Term	**Definition**	**Units**
I	Net charge transfer rate (current)	A
A	Area of electrode surface	mˆ2
i_0	Exchange current density	A mˆ-2
c_R^*, c_P^*	Reactant/Product surface concentration	mol mˆ-2
c_R^{0*}, c_P^{0*}	Reactant/Product surface concentration i_0 was measured under	mol mˆ-2
$\alpha_{a,c}$	Anodic/cathodic charge transfer coefficient	
n	Number of electrons transferred	
E	Electrode potential	V
E_{eq}	Equilibrium potential	V

For this voltage to be useful for driving an electrical load, the device must deliver a net charge transfer from one electrode to the other through the load. To achieve this, the forward reaction's activation energy must be biased by changing the voltage. As such, the voltage delivered to the load will always be smaller, often significantly so, than the OCV. The change in reaction rate relative to OCV is described by the Butler-Volmer equation[167] (Eqn 5.4), the terms of which are detailed in Table 5.2. Technically, the Butler-Volmer equation is derived for events that transfer only a single electron, but in practice is a good descriptor of most electrochemical reactions that occur in a single step, and even for more complex reactions where a single rate limiting step dominates the total reaction rate. Plotting the relationship between I and E given by the Butler-Volmer equation gives a good approximation of an ideal FC, with no Ohmic or mass transport losses.

$$I = A \cdot i_0 \cdot \left(\frac{c_R^*}{c_R^{0*}} e^{\alpha_a n F (E - E_{eq})/(RT)} - \frac{c_P^*}{c_P^{0*}} e^{-\alpha_c n F (E - E_{eq})/(RT)} \right) \tag{5.4}$$

If we want to improve the performance of the device, there are limited ways to do so (in this simplistic view). Voltage cannot be significantly increased, but the voltage loss required for a given current density can be decreased, or conversely the current density at a given cell voltage can be increased. This can be done by increasing the reactant concentration, temperature or number of reaction sites, or by decreasing the activation barrier.[168] Reactant concentration can be increased by increasing the pressure, but as most PEMFC's are designed to be "air breathing", they suffer from a nearly 5 times lessened oxygen concentration at the cathode.

This tends to make the cathode the major factor in determining the total reaction rate (and therefore delivered current of a device), and is one of the reasons that the catalyst loading tends to be higher in the cathode side of the MEA.[158, 159] Temperature increase is limited by the materials used,[169] as well as the need for liquid water to aid charge transfer in the electrolyte. It is also limited by system design, as pre-heating the fluid flows requires significant power (typically called a "parasitic load"), and may not increase the performance of the FC by the same amount.[170] This can be overcome by more complex systems that make use of heat exchangers and/or turbo machinery to capture energy (both heat and pressure) from the exhaust (or coolant) flows and transfer it to the intake flows, but this also comes with increased system size, weight, cost and complexity, so is not feasible for all use cases. The number of reaction sites is determined by the surface area (and to a smaller extent shape) of the catalyst particles, and loss of this area over time is a significant cause of irreversible degradation in PEMFC (from particles either becoming detached, or agglomerating to produce larger particles with a lower total surface area).[171] Finally, the activation barrier of the reaction is primarily impacted by the catalyst used, and is the area of much study.[146–148]

Comparisons of catalyst materials, structure and performance is beyond the scope of this project, but suffice it to say platinum group metals tend to produce the lowest activation energies because of their intermediate strength bonds with the species reacting in a PEMFC. They provide sufficient bond strength to adsorb the reactant species onto their surface and facilitate reaction, but not so strong a bond strength that the products are then captured, clogging active sites.[145]

It is not only the lower partial pressure/concentration of oxygen in the cathode flow that makes cathode kinetics dominant however. Even with pure oxygen supplied to the cathode, the ORR is significantly slower than the HOR.[172] This is primarily because it is much more complex, requiring multiple steps and considerable re-organisation of the constituents, whereas the HOR is as simple as adsorption, disassociation and desorption. Conversely, there are multiple different pathways that the ORR can take, not all of whose intermediates will interact with

each other in the forward direction,[173] and which are dependent on a variety of variables such as operating conditions, over-potential and the location of the reaction site on the catalyst surface (i.e. which lattice position in the metal structure and proximity to surface features such as edges and vertices).[174] All of this culminates in exchange current densities that are orders of magnitude lower for the ORR than the HOR. For example, the i_0 for the HOR under standard conditions with an acidic electrolyte (i.e. H^+ charge carrier, as opposed to OH^- in an alkaline cell) and Pt surface is on the order of 10^{-3} Acm^{-2}, whereas the i_0 for the ORR under identical conditions is on the order of 10^{-9} Acm^{-2}.[175] The three most dominant pathways in the literature are presented below:

Reaction Pathway presented by Gupta *et al*[50]

$$2Pt + O_2 \rightleftharpoons Pt_2O_2 \qquad \text{Adsorption}$$

$$Pt_2O_2 \rightleftharpoons 2PtO \qquad \text{Dissociation}$$

$$PtO + H^+ + e^- \rightleftharpoons PtOH \qquad \text{Reduction 1}$$

$$PtOH + H^+ + e^- \rightleftharpoons PtH_2O \qquad \text{Reduction 2}$$

$$PtH_2O \rightleftharpoons Pt + H_2O \qquad \text{Desorption}$$

Associative reaction pathway[176]

$$O_2 + \star \rightleftharpoons O_2\star \qquad \text{Adsorption}$$

$$O_2\star + H^+ + e^- \rightleftharpoons HO_2\star \qquad (5.5)$$

$$HO_2\star + H^+ + e^- \rightleftharpoons H_2O + O\star \tag{5.6}$$

$$O\star + H^+ + e^- \rightleftharpoons OH\star \tag{5.7}$$

$$OH\star + H^+ + e^- \rightleftharpoons H_2O + \star \tag{5.8}$$

Dissociative reaction pathway[176]

$$\frac{1}{2}O_2 + \star \rightleftharpoons O_2\star \qquad \text{Adsorption}$$

$$O\star + H^+ + e^- \rightleftharpoons OH\star \tag{5.9}$$

$$OH\star + H^+ + e^- \rightleftharpoons H_2O + \star \tag{5.10}$$

Finally, for these reactions to continue to occur in the forward direction, charge carriers must be transported between the two electrodes. The electrons are transported via the electrical load, and the resistance of this movement is dominated by the load's characteristics. There will also be some resistance involved in transferring them from the active sites of the electrodes to the current collectors (bipolar flow field plates). However, the charge transfer resistance tends to be dominated by the electrolyte, as this requires ionic conduction.[177] In a PEMFC, the electrolyte is formed from PFSA functional groups attached to a structural polymer. These form pores when polymerising, and these pores are lined with the functional groups.[178, 179] When permeated with water, the functional groups lose protons to the liquid to form hydronium complexes (H_3O^+). Due to the electrical field present, these are drawn from the anode to the cathode, hence transferring the required charge characters. However, this requires large ionic complexes to diffuse through nano-pores of a solid electrolyte filled with a liquid, which is significantly more difficult than electrons moving through metallic bonds, as in electronic charge transfer. It is for this reason that PEMFC's are so sensitive to humidity, requiring as much

liquid water present as possible to facilitate proton migration through the electrolyte, as this is the dominant charge transfer loss. This is aggravated by the fact that the charge carriers moving from the anode to the cathode move water molecules with them (referred to as osmotic drag), as well as the fact that water is produced at the cathode. The device therefore relies on water diffusion back through the electrolyte from cathode to anode, driven by the concentration gradient.[179] As such, the anode is significantly more humid than the anode under normal operation. The hydrophobic nature of PTFE is believed to be one of the reasons it works well as a structural polymer for PEM electrolytes, as it facilitates both proton conductivity in the forward direction as well as back diffusion of water.

5.2.4 Current perturbations

In order to increase power produced by a FC, it is possible to use periodic current perturbations (frequently achieved by short circuiting the stack). This produces an increase in power of approximately 10%,[49, 50] which in turn would allow the system to be 10% smaller for a given application. This is obviously a significant increase in power density, and it is surprising that more work has not been done in this area. There is not a consensus on the mechanism through which these perturbations increase power, and there is little available data on the impact on the longevity of the FC. Functionally, the phenomena has been found to appear as follows:

A PEMFC device under a given set of operating conditions, (i.e. reactant flow rates, device and fluid temperature and humidities) at a constant current (and therefore fuel consumption) will operate at a given voltage, here referred to as "system equilibrium voltage". A current perturbation (i.e. a short pulse of much higher current through the device) will cause the output voltage immediately after the perturbation to increase significantly. The length of these perturbations is typically on the order of 10s of milliseconds.[?, 50] The output voltage will then decay exponentially back to the system equilibrium voltage, with a typical time constant being on the order of 10s of seconds. By periodically applying these perturbations, the voltage will never be allowed to decay fully back to the system equilibrium voltage, and therefore the average voltage delivered by the device will be higher. This produces a higher average power output for

a given fuel consumption, therefore increasing specific power and efficiency of the system.[50]

These current perturbations can be easily achieved with a simple solid state switch, or at most 2 if the load cannot tolerate the fluctuations in voltage and must be disconnected during a perturbation. The practical challenges in achieving this will be discussed further in Section 6.3, but here it is sufficient to say that it does not require significant increase in system complexity or cost.

While this technique is reasonably well known, it is poorly understood. There is no academic consensus for the mechanism through which it acts, although 2 main theories are present in the literature.[50] The first suggests that the increase in charge throughput due to the perturbations, and therefore increase in the reaction product (water) increases the humidity of the stack, reducing charge transfer resistances. The second suggests that the perturbations strip oxide species from the catalyst surfaces, increasing the number of active sites available for reactions to occur.

Humidification theory

As described above, PEMFC's require high humidity to aid charge transfer through the electrolyte. This theory assumes that the short periods of high current produce enough water to humidify the stack by a significant amount, and that humidification is responsible for the improved performance. This theory was popular until recent years, and is the basis for much of the published work.[4, 47] However, as raised by Gupta *et al*,[50] the time for which the current is significantly above the norm is so short as to be an insignificant amount of water produced. Silva *et al* present humidity data for incredibly long current perturbations[4] (10s, compared to the usual 10's or 100's of ms), and show a significant rise following a perturbation. However, because this humidity is calculated from the dew point the change in calculated humidity is likely entirely due to the change in gas pressures and temperatures, and so does not demonstrate that there is significant humidity generated by the procedure. It also assumes that the gases passing through the cell achieve equilibrium (with respect to humidity), which is unlikely given the short time they spend in the device. This conclusion is backed by the incredible similarity

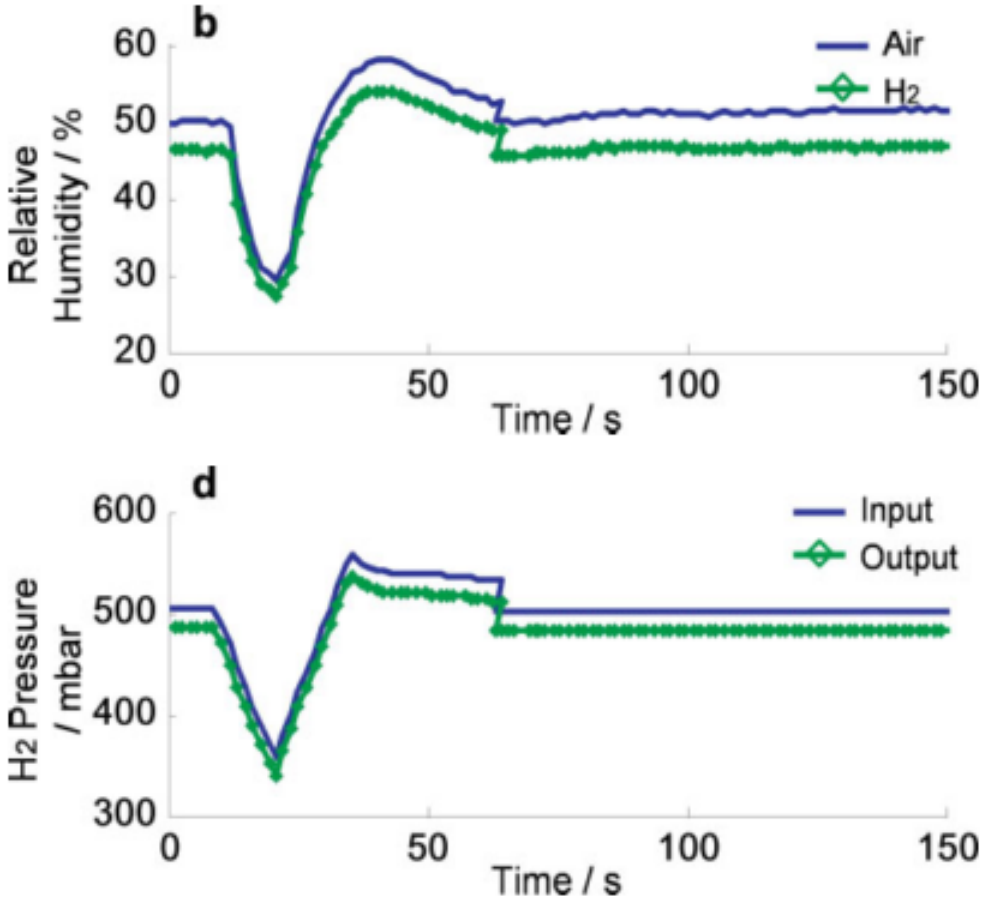


Figure 5.2: Data as presented by Silva *et al*[4]

between the humidity curve and hydrogen pressure during a perturbation event, as can be seen in Fig 5.2.

Another reason that this theory seems dubious is that the ohmic resistance of the stack changes very little, and in the wrong direction. Having measured the stack resistance via electrochemical impedance spectroscopy (EIS), they found that it *increased* by 1.23%, the data for which can be seen in Table 5.3 and Fig 5.3. These tests were performed for the purpose of studying the effect of system faults on the degradation of the PEMFC, as opposed to testing an *in operando* technique and as such the short circuit in these tests was significantly longer than usual, and so would be expected to produce a significant amount of humidity. These tests represent an extreme scenario which should produce significant increases in humidity, and therefore if the humidification theory were accurate, significant changes in both ohmic resistance and a subsequent increase in voltage. However the humidity barely changes once the impact

Table 5.3: Ohmic Resistances (Rm/mΩ) before and after 10s short circuit. Reproduced with permission from Silva *et al*[4]

Cell No.	[5-8]	[9-12]	[13-16]	[17-20]	[21-24]	[25-28]	[29-32]	[33-36]	[37-40]	Stack
Before	1.83	1.76	1.74	1.71	1.70	1.66	1.61	1.57	1.63	17.03
After	1.87	1.80	1.72	1.71	1.72	1.68	1.62	1.58	1.63	17.24

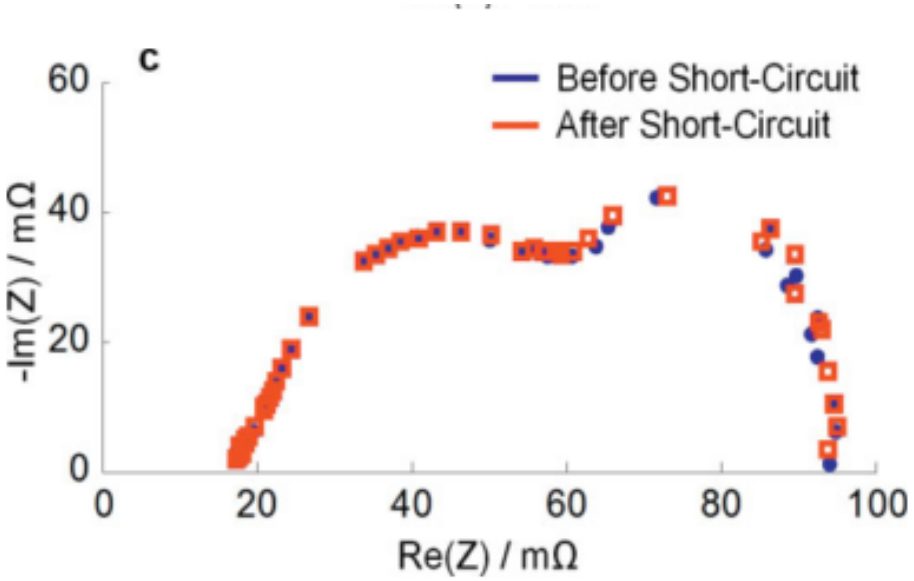


Figure 5.3: EIS data presented by Silva *et al* in[4]

of changes in pressure are accounted for, and the voltage responds as we would expect from a significantly shorter perturbation. Alone, this would not be sufficient to disregard the theory, as the voltage response increases to approximately 0.95V per cell (specific values are not given in the test, the data is only presented graphically as part of a figure with multiple plots). This occurs under the nominal current density of 0.5 Acm^{-2}, and is a very healthy voltage response. Its rise is somewhat slower and smoother compared to the data presented by Gupta *et al*,[50] but this can be explained by the fact that such a long short circuit will have consumed all the available reactants in the GDL, and so the response immediately after the short is dominated by the diffusion characteristics of the GDL. These same mass transport effects may be the cause of the increased resistance measured following their short circuit. Alternatively, the extended duration of the short may have actually dried out the anode, as a sudden spike of current through the electrolyte will likely transport more water through the membrane via osmotic drag than can be immediately replaced via diffusion.

Oxide stripping theory

The second theory posits that oxidation of the catalyst surface (typically Pt or Pt alloy nanoparticles) occurs due to high potential, and these oxidised sites are therefore no longer available for the ORR to take place. These oxide and hydroxide filled sites are stripped by the lowered potential and increased charge throughput of the device shifting the position of the active area on the Butler-Volmer curve. It is posited that the voltage response immediately after is due to there being more active sites for the ORR to occur, decreasing the overpotential required for the operating current density (the A and (E-E_{eq}) terms in Eqn 5.4). These sites are then slowly occupied by the oxidation of the surface under normal operation, hence the need for periodic perturbations to undo reversible degradation of the catalyst surface. While this theory holds merit, it ignores the fact that oxides and hydroxides present at surface sites of the catalyst are intermediate species in many pathways of the ORR, and so do not inherently prevent the ORR from occurring at those sites, but are part of it occurring. As such, a modified version of this theory is presented by Gupta *et al.*[50] This does not contradict any specific statements regarding the oxide stripping theory, but is rather a more clear explanation of the proposed mechanism, with more detail regarding how the ORR occurs on Pt nanoparticle surfaces.

Equilibrium theory

There are several steps in the ORR pathway, and limited reaction sites on the surface of the platinum. Each step has an equilibrium point in a given set of operating conditions, as described using the Butler-Volmer equation above. Of course, in actuality the device is more complex, and will have inhomogeneities in flow fields, temperature and electrical gradients, variations in transfer resistances in the materials etc. As such, the operating point is also impacted by device history, factors such as humidity (which will impact both the transport of charge carriers as well as the overpotential available at the electrode/catalyst/electrolyte tri-phasic boundary). While there are multiple different pathways the ORR can take (some of which are presented above), imaging or measuring the species during reaction is incredibly difficult, and as such it is not possible to say with certainty which pathway is dominant. It is expected that all pathways occur to some extent, with some being dominant under any given set of operating conditions,

but the focus of this work is agnostic with regard to pathway, as they all have rate limiting steps. Should the equilibrium of the intermediate reactions be mismatched, they will lower the concentration of certain species present, or of reaction sites, therefore limiting one or more of the reactions and limiting the current density at that potential. This therefore lowers the cell potential, and this is what causes the stack potential to be reversibly lowered. However, as these equilibria do not form instantaneously, it may be that the current perturbations shift the reactions, freeing up active sites and shifting the balance of the species concentrations, increasing the voltage at a given current density.[50] The reactions will of course begin to settle towards their equilibrium point, and this can be observed in the voltage response of a device immediately following a perturbation. Periodic perturbations are therefore required to keep the reaction in the artificially raised state, and a better understanding of these mechanisms can allow for this period to be better tuned to achieve the optimal outcome. This is somewhat distinct from the oxide stripping theory as the presentation of that theory implies that active sites occupied by oxide and hydroxide species are no longer participating in the desired reaction, as opposed to being the species involved in the rate limiting steps of said reaction.

5.2.5 Discussion

For the reasons laid out above, this project is more swayed by the equilibrium theory than by the humidification one, but there is still insufficient data in the literature to assert a dominant mechanism with any confidence. As such, one of the primary research questions selected was regarding this mechanism, and gathering data to help identify the mechanism through which this phenomena acts. PEMFC's are complex devices, and the current perturbations typically used in this technique are aggressive (usually via a very low resistance short circuit), which means that they impact the FC in multiple ways. While the humidity will technically be raised, the amount of additional water produced seems negligible compared to the total amount produced by normal operation. It could be that a rapid change in the humidification at a very specific part of the devices structure has an exaggerated influence, or that the increased reaction rate is selectively heating the catalyst nanoparticles, producing a lowered activation voltage until that

heat diffuses into the supporting structure. However, given the data available in the literature, it is expected that the dominant mechanism by far is the shifting of intermediate species equilibria. It is very difficult to separate the oxide stripping and equilibria theories in practice, as they will produce very similar if not identical responses, being so closely related. Indeed the separation of these 2 explanations is an act of rigour on the part of the author, as opposed to them being opposed.

One of the reasons it is difficult to discern the mechanism with clarity is lack of applicable data. As mentioned, there are very few publications on this phenomena, even fewer with data, and fewer still with directly applicable data. All available data is either gathered for a different purpose (e.g. Silva *et al*[4] doing extended shorts to investigate response to system failures, as opposed to as a control technique), is performed in a potentially suspect way (e.g. accelerated ageing studies which leave the device at OCV,[52] increasing the impact of the technique by increasing the voltage window it cycles the device through), or on devices that are not representative of modern systems (e.g. open cathode graphite plate designs with out of date MEA materials[4, 47, 49–51]). All of this data is useful and important, but does not answer the most pressing research questions, namely "Does it work?" and "How does it work?". While some work has been done to show that the technique can be optimised,[50, 51] and is likely convex with respect to the duration and period of the perturbations, the control theory applied is relatively primitive, and none of these tests were performed with any kind of live or adaptive operation.

5.3 Focus of work

5.3.1 Research questions

Based on the information presented in the literature and the known response of a PEMFC to the implementation of current perturbation, two primary research questions were selected:

Is it possible to dynamically optimise current perturbations for PEMFC applications, and what are the quantified benefits of the technique for modern automotive style PEMFC's?

Current perturbation is a well known phenomena in PEMFC's, but is little discussed in literature. It has been utilised in a range of forms, and on a range of stack designs in the past, but little work appears to have been done on optimising it's performance. Instead, fixed perturbation lengths and intervals are typically selected, regardless of operating conditions and stack design. While some preliminary work has been done to show that the technique can be optimised, and is likely convex with respect to the perturbation lengths and intervals, there is no available work on optimising the technique dynamically in operando. As it is expected to be convex from the little available literature, it is a fairly simple task to optimise dynamically using very simple methods such as perturb-and-observe.

Does the transient behaviour of a PEMFC under current perturbation lend credence to a particular theory as to the mechanism through which the technique gives improved performance, and can this then be used to find an analytical solution for the optimal time between perturbations?

While the technique is well known, the underlying mechanism through which it improves performance is poorly understood. While there are some competing theories, the most dominant, and in the authors mind most scientifically likely is to do with the Oxygen Reduction Reaction (ORR). Namely, by perturbing the current, the equilibrium of the slower steps of the multi step reaction occurring on the cathode catalyst surface is disrupted, and takes a time to settle (with a time constant on the order of 10-1000 seconds). Therefore, by repeatedly disrupting the equilibrium from what it would settle to given a set of steady-state conditions, and artificial equilibrium can be imposed upon these steps which has a higher mean voltage (at a given current density), and therefore higher output power and efficiency. If these equilibria follow one or more Arrhenius relationships (i.e. $\alpha e^{-\beta t}$), then an analytical solution can be found for the optimum time between perturbations. By taking high frequency measurements of the voltage

behaviour during a perturbation, as well as comparing this analytical solution to the value found using perturb-and-observe optimisation, we can investigate in more detail the underlying mechanism through which the technique acts.

These research questions were selected as it is felt they were under-represented area in the literature, as well as being key for the implementation of the technique in modern PEMFC systems. The impact of the technique long term was not selected for study due to the time and resource constraints of the project. It is also felt that a developed implementation of the technique may mitigate many of the concerns raised in the literature regarding the impact on long term performance, which will be discussed further in Section 6.3. While some work has been done on optimising the technique by Gupta *et al*[50] and Zhang *et al*,[51] this work suffers some shortcoming in it's use for modern systems. First of all, the devices tested are not representative of modern systems. Their current densities are significantly lower, so the position on the Butler-Volmer curve is different. It may be that at higher current densities (achievable only with modern MEA materials, stack designs and liquid cooling), the technique offers significantly less benefit than on older devices. The more modern systems will have significantly reduced losses from things such as transport of charge carriers, and will likely have a significantly higher concentration of active sites due to advances in catalyst nanoparticle synthesis. They will also have better humidity management, both from intrinsic properties, such as being closed cathode, and from plant design. Their thermal management is more uniform, they would be expected to have a lower cell to cell variability, and they even have lower calendar ageing, as it is easier to seal the system during storage, preventing species that would poison the MEA or catalysts from diffusing in. For example, an open cathode system stored open on a shelf near any VOC's (particularly those containing amines, benzene or sulphur components) could easily lose performance due to this proximity, and will likely do so in a non-uniform way, whereas a closed cathode system can be more easily protected from this source of degradation.

5.3.2 Experiment parameters

With these factors in mind, the proposed experiments begin to take shape. They should be performed on a device that is representative of modern systems, with a balance of plant that allows repeatable and measurable operating conditions. In particular the device should be capable of operating near 2 Acm^{-2}, be liquid cooled and utilise modern plate designs (in terms of both flow fields and materials). Active area of the cells should be significant, I.e. while it is not necessary for the device to pass as much current as a modern high power system, it should at least be large enough to capture some of the nuances that come with such a scale, as opposed to being something that could fit in the palm of a hand. Number of cells in the stack should be large enough to potentially see any inhomogeneities, but need not be so large as to enter HV regions where additional dangers are posed, I.e. in the region of $10<N<40$. In addition, a current perturbation unit (IPU) must be created that can be accurately controlled to provide perturbations of any length between approximately 1-500 mS, and must be very repeatable as well as enough instrumentation and processing power to perform simple optimisation strategies. Finally, the voltage response in particular must be measured at a reasonably high frequency (>1 kHz) in order to attempt to capture the behaviour of the device during a perturbation. How each of the criteria were addressed shall be detailed in Section 5.4 .

Next, the operating conditions must be selected in order to produce clear results for both the optimisation and for scientific purposes. As such, the system will be operated at fixed currents and the improvement delivered by the technique will be determined by the voltage delivered. While it would be interesting to investigate the performance of these optimisations at either constant power or over an indicative drive cycle, the complexity introduced is not beneficial to the focus of this work: a proof of concept, as opposed to the presentation of a mature or developed technique. It would be beneficial to determine the benefit of the technique across the operating window of the device, but the complexities and increase in amount of testing required pushes this beyond the resources of this project. Ideally, a range of current densities would be tested, chosen to give a fair representation of operating points at near maximum efficiency, a typical balanced load, and near maximum load. In order to limit the impact of other factors

the system must be allowed to reach a reliable steady state at the selected current and then have the technique applied in various forms, with the system returning to steady state between tests to identify any variations.

5.3.3 Optimisation

Optimisation of this technique requires a cost function. Because the system is isolated on the bench, and thanks to the selection of appropriate experimental conditions (such as constant load current outside of a perturbation), the cost function is very simple. The intention is to maximise power output, specifically the increase in power above the baseline (when no perturbation is applied to the FC). As DC electrical power is simply the product of the voltage and current applied to a load and the current is fixed, we are attempting to maximise the mean voltage output by the FC. Our input values are the duration and period of the perturbations, which must be real positive numbers. In order to protect the FC from damage, these must be bounded to prevent damage from extended over-current (which could cause localised hotspots in the MEA, causing irreversible loss of performance), extended under-voltage (which can cause corrosion of the GDL materials) or extreme voltage cycling (which can cause both corrosion of GLD materials and agglomeration of the catalyst nano-particles). As such, the duration of the perturbations was limited to be <400 ms, which is expected to be well above any optimal values found. In addition, the period must be >5s, in part to prevent damage to the FC, but also to allow the CVM to perform it's required handshake protocols with the perturbation controller to enter and exit it's high frequency acquisition mode during each perturbation. In practice, this handshake required only 10s of ms to perform, but this had not yet been demonstrated at the point that these bounds were set. This bound was again expected to be significantly outside the region of optimal values so that the optimisation controller never actually hit this limit. While this limit was not found during the course of testing, it was also not as far from optimal values found as expected, as will be discussed in Section 5.4.2. Because the system has a closed anode operating at constant pressure, and is not purged during the gathering of data, the fuel consumed by the device is proportional to the current drawn. This is more convenient

than directly measuring gas input to the system via a flowmeter (although this was also done), as it excludes any contributions from minor leaks or crossover in the FC. The data from the flowmeter is noisy thanks to the flow controllers oscillatory behaviour, and as such accurate measurement of the fuel consumption directly is difficult. As this work is an early proof of concept under fixed conditions and not the demonstration of a fully optimised and integrated system, it is sufficient to treat the fuel consumption of the FC to be proportional to the current, and therefore fixed. As discussed in Section 5.2.4, the charge transferred during a perturbation (when current is not fixed) is negligible in relation to the amount transferred during normal operation, so the impact of current perturbations on both water generated and additional fuel consumed is also negligible.

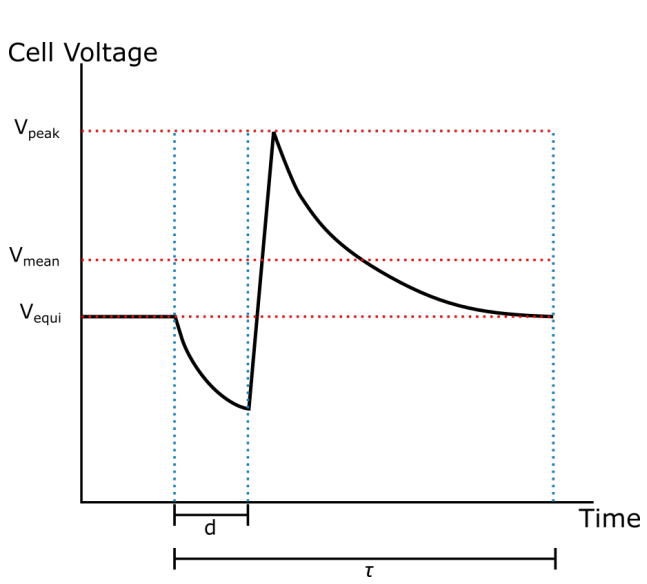
As the system is expected to be convex with respect to the parameters controlled (perturbation duration and period), with no local minima, there are myriad optimisation techniques that could be easily applied. It was felt that the simplest, most robust technique applicable should be chosen for these tests in order to prevent complications that could cloud the scientific process, but also that some attempt at an analytical solution should be made, to contrast with some purely numerical methods. As such, the optimisation technique selected was "Perturb and Observe". This is one of the simplest protocols possible, and is extremely robust, as well as widely used. It does, however, suffer from being somewhat slow to converge, although this allows it to be resistant to transient effects and temporary influences from external factors. It is commonly used for a range of purposes, for example in control of solar panels (both angle of incidence and to find the maximum power point on the Voltage-Current curve), or wind turbine blade angle of attack. Its basic premise is that for a convex system with no local minima (with respect to the parameter being controlled), any sufficiently large perturbation of the input parameter should yield an observable change in the output parameter. If this response is in the direction desired, the next perturbation (or step) should continue in the same direction until the output response changes direction. This will cause the system to settle around the optimum point. It can be enhanced slightly by allowing the system to adjust the size of the step changes in the input parameter to perform a golden section search, where each time the direction of the steps is changed, the step sized is multiplied by $1/\phi$. This allows the system to begin with

fairly coarse steps, increasing speed of convergence, but rapidly decrease the size of these steps once the region in which the desired minima lies. It is important to set a minimum step size for a system like this, as too small a perturbation will yield a change in the output parameter that is smaller than the noise of the system, yielding unreliable results. This is trivial to do in practice, however, and can also be mitigated by taking an average of multiple samples, albeit at further cost to the speed of convergence. This is particularly true of this system, as the observed parameter is average voltage over the course of an entire period (which are expected to be on the order of 10s of seconds), so taking multiple samples will require each observation period to be minutes, if not 10's of minutes long. However, once some initial tests have been performed, a multi-dimensional system map can be produced for the parameters, similar to that used in the ECU's of modern ICE's. This map will give much better initial estimates of the parameters optimum points for a given scenario, and can easily be interpolated (or if need be extrapolated) to account for scenarios it has not previously encountered. As such, it should improve the time required to converge significantly, as well as improving the average performance of the system and technique during convergence.

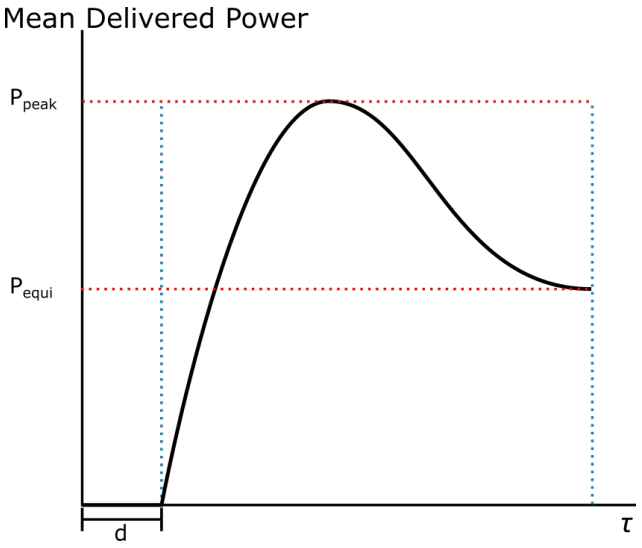
While P&O is designed for the optimisation of systems with a single degree of freedom it can be adapted for use in 2DOF systems, provided both parameters are convex. This is done by simply freezing one parameter, optimising the other until a desired minimum step size is reached, then freezing that parameter and optimising the other in a similar way. This produces a somewhat "zig-zag" path in the plane of freedom, and so is once again not as fast to converge as some other techniques, but once again benefits from being incredibly robust. In particular, it does not require the input parameters to be completely independent of one another, provided they remain convex with respect to the output (as is expected from the available literature, described in Section 5.2.4).

5.3.4 Analytical Solution

An analytical solution is difficult to determine without more detailed and applicable data to base it on, but some work can be done provided some assumptions are made. Namely, we are going to assume that the mechanism through which this technique acts is predominantly the equilibrium theory presented in Section 5.2.4. As such, the expected voltage response of the system to a current perturbation is as seen in Fig 5.4. It should be noted that the time-scales are intentionally not correct in this figure, but have been expanded for clarity of communication. I.e. the duration of the perturbation (d) is significantly larger relative to the period (τ) in the figure than in reality.



(a) Cell voltage during a current perturbation



(b) Mean power delivered against time between perturbations

Figure 5.4: Graphical representations of the expected voltage and delivered power response of a PEMFC while using current perturbations. Times are not to scale, but have been expanded for clarity of communication.

The voltage during a perturbation is expected to act in a rapid exponential decay as the double layer capacitance of the electrochemical system is expected to be dominant. The rates and time constants of this behaviour are difficult to estimate, but the peak current is expected to be on the order of 5-10 Acm^{-2} based on data presented in the literature.[4, 50] The time required for the voltage response to increase to it's maximum is expected to be negligible and will be aided by the back emf produced by the inductive characteristics of the current carriers, so is represented as being instantaneous. The voltage is then expected to decay exponentially,

asymptotically approaching the steady state voltage reached prior to any perturbation. The whole of this decay is shown here, although in use of this technique it would not typically be allowed to decay fully, but instead be periodically perturbed to achieve a higher mean voltage. It is also assumed that the load is disconnected during a perturbation, or at least that the perturbation device has an impedance significantly smaller than the load, and so the power delivered to the load during a perturbation is null. The mean power is a cumulative average taken from the beginning of the perturbation and assumes current is constant (outside the perturbation), and power delivered is therefore a function of voltage only. Because of this initial time where no power is delivered, the mean begins by being weighted at 0, rises to an optimum, and then decays over time to also asymptotically approach the value found at steady state. A very simple model of the voltage delivered to the load is presented in Eqn 5.11, and is based on the voltage decaying exponentially from a maximum immediately after the perturbation to the voltage found at steady state prior to the perturbation.

$$V(t) = \begin{cases} 0 & \text{for t} \in [0, d] \\ V_{equi} + (V_{peak} - V_{equi})\beta e^{-\alpha(t-d)} & \text{for t} \in [d, \tau] \end{cases} \tag{5.11}$$

This can be expanded to n exponential terms to include an unknown number of steps in the reaction pathway, each with a different time constant. In practice, more than 2 or 3 terms here would be superfluous, as any impact of the additional terms would be increasingly insignificant. It is presented in this form in Eqn 5.12 to demonstrate that this simplified expression has sufficient breadth to be fitted to any voltage decay we can realistically expect to observe, and so the later conclusions drawn from it are insensitive to the details of the initial expression. This will also be demonstrated in an intuitive sense, but this is insufficiently rigorous to be trusted scientifically, and is mostly of benefit for ease of understanding.

$$V(t) = \begin{cases} 0 & \text{for t} \in [0, d] \\ V_{equi} + (V_{peak} - V_{equi}) \sum_n \{\beta_n e^{-\alpha_n(t-d)}\} & \text{for t} \in [d, \tau] \end{cases} \tag{5.12}$$

Remember that the parameter we are attempting to maximise in this scenario is the mean

voltage delivered to the load, which is the definite integral of the function presented above between 0<t< τ, divided by τ. The beginning of the perturbation is defined as t=0 and t=τ is an undetermined time at which the next perturbation will begin. We are attempting to find an analytical solution for the value of τ at which V_{mean} is maximised. As the response is zero while 0<t<d, so is the integral, leaving V_{mean} as:

$$V_{mean} = \frac{1}{\tau}\int_{d}^{\tau} V_{equi} + (V_{peak} - V_{equi}) \sum_{n} \{\beta_n e^{-\alpha_n (t-d)}\} dt \tag{5.13}$$

This assumes that both α_n and V_{peak} are independent of τ, which fits with our initial assumptions. α_n would be determined by the device and its immediate operating conditions, and V_{peak} a function of d. It is expected that too short a duration would lead to a minimised voltage peak, as it does not push enough charge through the device to shift the species concentrations, and too long would have no added benefit (i.e. will not increase the voltage response significantly above a maximum). This is supported by data from the literature, in particular Gupta *et al*[50] and Zhang *et al.*[51] A plot from Gupta *et al* can be seen in Fig 5.5 demonstrating that above 200 ms short circuit duration there is no significant increase in voltage response immediately after the event.

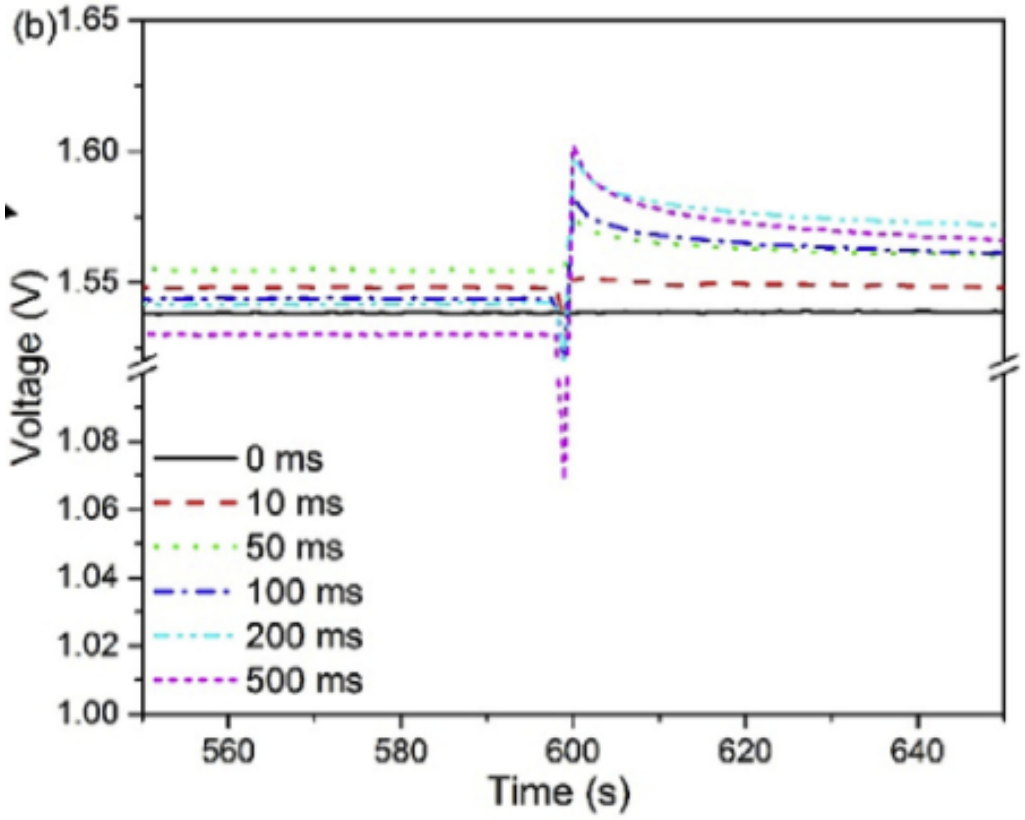


Figure 5.5: Fig 2. from Gupta *et al* showing the voltage response of a 2 cell stack to short circuits of various lengths. Reproduced with permisison.

To continue our derivation of an analytical solution for the optimum value of τ, we must evaluate the integral presented in Eqn 5.13. Doing so, and expanding out the terms gives Eqn 5.14.

$$V_{mean} = \frac{1}{\tau}\left(V_{equi}(\tau - d) - (V_{peak} - V_{equi})\sum_{n}\frac{\beta_n}{\alpha_n}\{e^{-\alpha_n(\tau-d)} - 1\}\right) \tag{5.14}$$

As we are trying to find the value of τ for which V_{mean} is at its maximum, we must take the partial differential with respect to τ. This is given in Eqn 5.15.

$$\frac{dV_{mean}}{d\tau} = V_{equi}\left(\frac{\tau-(\tau-d)}{\tau^2}\right) + \frac{V_{peak}-V_{equi}}{\tau}\sum_n\{\beta_n e^{-\alpha_n(t-d)}\} + \frac{V_{peak}-V_{equi}}{\tau^2}\sum_n\frac{\beta_n}{\alpha_n}\{e^{-\alpha_n(\tau-d)}-1\} \quad (5.15)$$

We can re-arrange this expression to give us 2 familiar terms, as in Eqn 5.16.

$$\frac{dV_{mean}}{d\tau} = \frac{1}{\tau}\left(V_{equi}+(V_{peak}-V_{equi})\sum_n\{\beta_n e^{-\alpha_n(\tau-d)}\}\right) - \frac{1}{\tau^2}\left(V_{equi}(\tau-d)-(V_{peak}-V_{equi})\sum_n\{\frac{\beta_n}{\alpha_n}e^{-\alpha_n(\tau-d)}-1\}\right) \quad (5.16)$$

Setting $\frac{dV_{mean}}{d\tau}$ to 0, multiplying through by τ and rearranging:

$$(V_{equi}+(V_{peak}-V_{equi})\sum_n\{\beta_n e^{-\alpha_n(\tau-d)}\} = \frac{1}{\tau}\left(V_{equi}(\tau-d)-(V_{peak}-V_{equi})\sum_n\frac{\beta_n}{\alpha_n}(e^{-\alpha_n(\tau-d)}-1)\right) \quad (5.17)$$

Substituting in Eqns 5.12 and 5.14 gives:

$$V(t) = V_{mean} \quad (5.18)$$

This gives us a very easy analytical solution for the optimum value of τ. All we must do is measure the voltage, take a cumulative average of it's value from the beginning of a perturbation, and compare these 2 values. When the momentary value meets (or falls below) the average value, then the average has reached it's turning point and will begin to fall, so in order to maximise it, another perturbation should be triggered. This also makes intuitive sense, as the average begins weighted at 0, rises due to the increased voltage immediately following a perturbation, and will only fall once the voltage falls below it, as the new values being added to the cumulative total are smaller than the average. This solution requires no complex electrochemical models to be reduced and packaged onto an embedded chip, or even a significant amount of processing power, just a moderate amount of memory ($<$ 1 kB at a sensible sample rate of 1-10 Hz). It is insensitive to the terms used to model the voltage decay, and does not require any paramaterisation prior to use. The same code can be used on any PEMFC system and should perform optimally, and can be expanded or adjusted in order to account for real world considerations with ease.

This obviously does not find us an analytical solution for the value of d, but by reducing the degrees of freedom we have to optimise, we cut the convergence time in half at a minimum, and realistically by a significantly larger margin. An analytical solution for d would most likely require some careful and complex calibration or parameterisation for any given system, for which there was unfortunately insufficient time and resources in this project. Given that we expect it to be fairly insensitive to the value of d above the optimum, however, we can speed our convergence (as well as the system performance during initial optimisation) by beginning with a value of d we expect to be higher than the optimum and aggressively reducing it until the peak voltage immediately after a perturbation begins to suffer significantly. In addition, we could also use the values of τ to help, as we would expect these to also fall rapidly once d is reduced below it's optimum value.

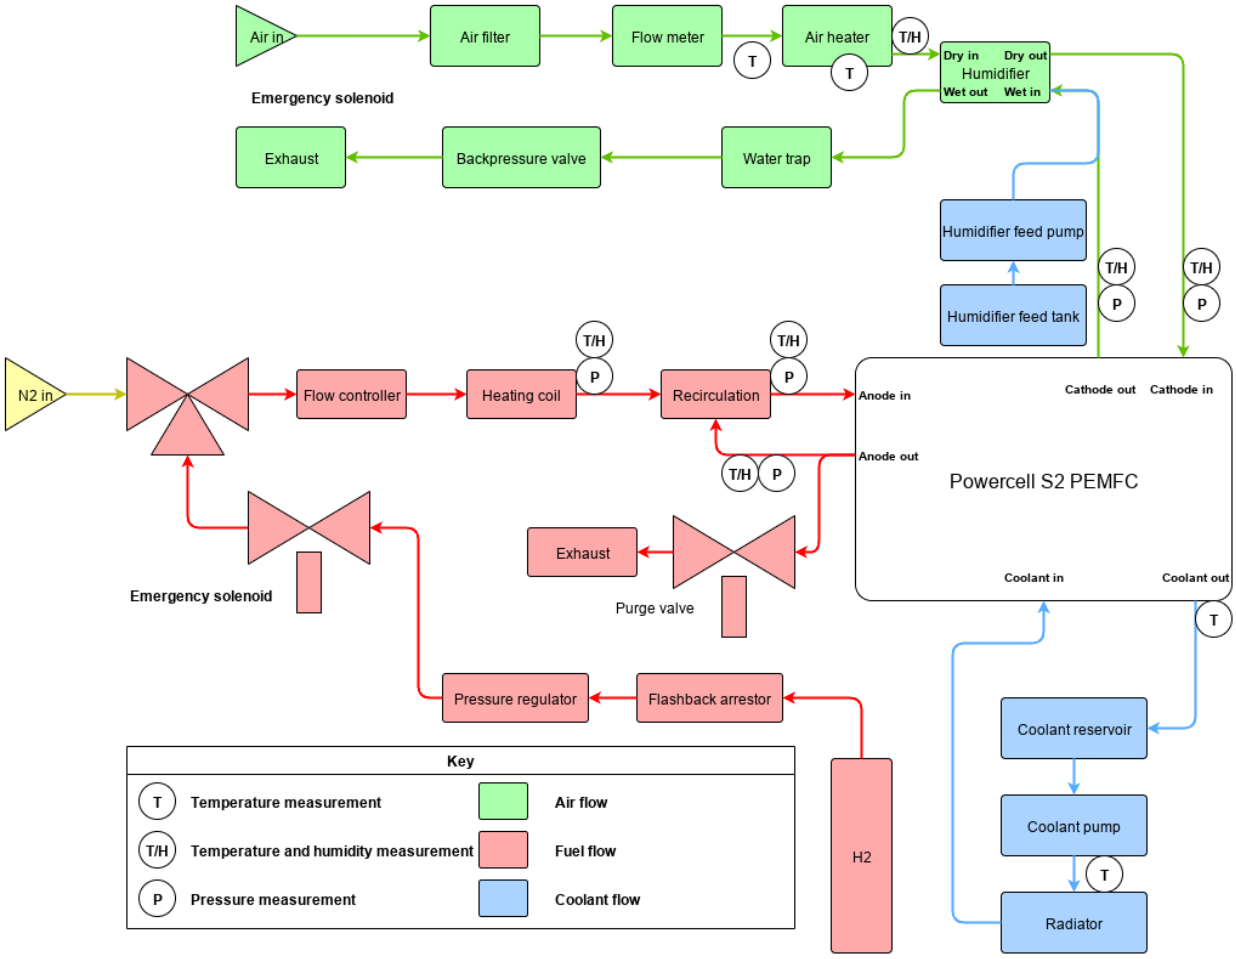


Figure 5.6: A diagram of the PEMFC balance of plant designed and built for this project

5.4 Methods

5.4.1 Selection of materials

PEMFC stack

In order to investigate the use of current perturbations for modern systems, it was necessary to test a PEMFC that is representative of those implemented in commercially available systems. As such, the device must be manufactured with modern methods (namely metal plates as opposed to machined graphite), utilise modern active materials (in particular catalyst and exchange membranes), and operate at current densities that make it viable for use in systems such as automotive and aerospace. These are typically closed cathode systems, although some open cathode options do exist for low altitude aerospace applications where minimising the weight by

removing the cathode flow conditioning balance of plant is preferred over maximising electrochemical performance by controlling cathode pressure, temperature and humidity. However, as it would be much easier to produce reliable and comparable test data with a system that allowed these parameters to be controlled, and due to the interest in how the technique would perform at higher current densities, a closed cathode system was selected. Specifically, the Powercell S2 stack, configured with 24 cells. This stack design is used in a variety of commercial systems, including commercial aerospace APU's, marine and automotive, and as such is representative of the current industrial market. The 24 cell system is the smallest sold, and was selected in order to make the balance of plant and cost of the experimental setup as small and simple as possible, while still being a viable example of a modern system. It is a closed cathode system, and while the active area of each cell is not given, based on the physical measurements and current delivering capabilities, a current density approaching 2 Acm^{-2} is possible. It is liquid cooled, another key difference between modern commercial devices and those presented in the literature on this topic. Finally it has been designed with the ability to operate using reformate gas as fuel, which is typically high in impurities that can poison the active materials in a PEMFC, and does so by utilising modern catalyst and MEA materials.

Balance of plant

The plant required to operate a PEMFC of this size is not readily available in a form that lends itself to academic research. While full systems are sold, they tend to be sealed units, and the manufacturers are wary of selling them to purchasers who intend to open them up, either for data gathering or reverse engineering purposes. As such, it was decided to build the plant in house to allow more rigorous control during testing, as well as facilitate gathering of all pertinent data from the system. The layout of this plant can be seen in Fig 5.6. This gives sufficient control of all relevant flow parameters, namely temperature, pressure, flow rate and humidity. Components marketed for scientific purposes were avoided, except where necessary, as they tend to have significantly increased cost, as well as being more complex. This is usually done to make them perform very predictably for the use in which they were intended, but given that there are no available solutions for the use intended here, unfortunately serves only to make

them more unpredictable the rest of the time. As such, common industrial solutions were used wherever possible in order to reduce cost and increase reliability.

Control electronics

In order to control this balance of plant, as well as gather, display and record the relevant data for scientific purposes, a set of custom control electronics was designed and fabricated. This was formed of multiple subsystems in order to make the system more robust, easily tested, repaired and altered. It was also intended to simplify the programming process. These systems were:

Safety System Responsible for identifying failure modes, measuring safety devices (e.g. H_2 sniffers) and responding automatically, as well as offering the operator optional response protocols

Process controller Responsible for operating the bulk of the balance of plant, including all fluid flows and associated devices, as well as reading sensors related to operating conditions

Cell Voltage Monitor A custom unit designed in order to read 24 stacked voltages accurately, as well as at a high sample frequency (up to 7 kHz) when triggered

Load controller/Perturbation unit Responsible for controlling the electronic load applied to the FC, as well as the perturbation circuitry. Also applies the desired perturbation and optimisation protocol, including performing any calculations necessary.

Data logger/Interface Responsible for collating data, displaying it for operator, and passing commands to necessary subsystems

All systems read at 10 Hz, with data piped over serial communications and plotted live using SerialPlot, an open source data visualisation and capture tool.[180]

5.4.2 Design and Fabrication

As mentioned, a rig capable of performing the tasks required for these experiments is not available for purchase. As such, one fit for purpose was created. This section will enumerate the

parts of this process that are pertinent to the data presented. This will be done on a sub-system by subsystem basis.

Cathode

The cathode subsystem is required to provide conditioned air to the stack, under moderate pressure. The operating window desired is 0.6-1 bar, 30-60°C, 50-200 SLPM, 90-95% humidity. Consistency is of most importance, provided these parameters are within acceptable windows, as the aim is to fix all variables that could impact operating performance apart from those under investigation. Unfortunately, the lab in which this work was conducted was not furnished with a supply of compressed air, so portable industrial compressors were used. This meant that the electronic system could not control the inlet pressure, as the compressors required manual setting of their single stage regulators. A variable inlet valve, similar to that used to control the backpressure and flow rate at the outlet of the cathode subsystem, would have allowed the system some input, but unfortunately due to time limitations it was not feasible to implement this. In addition, it would have required very careful tuning and the implementation of extra safety protocols, as multiple control loops would be controlling linked variables (i.e. the input valve would help control system pressure, but via a variable choke, in which pressure and flow rate are intrinsically linked). As such, it was very feasible that an error could have either caused a rapid spike in pressure, damaging the MEA's of the stack beyond repair, or starved the cathode of oxidiser, causing rapid and irreversible degradation of the cathode catalyst and GDL. While these are certainly surmountable problems, unfortunately their implementation required time that was not available. As such, the cathode pressure varies somewhat during operation, primarily due to the non-linear behaviour of single stage pressure regulators. This variance is <100 mbar under normal operation, but the impact of this on system performance can easily be seen in the data, given the sensitivity of the measurement devices used and the high repeatability of the PEMFC's behaviour.

Temperature is also vital, in particular in achieving higher current densities. This is primarily because colder fluids will cause water vapour (the product of the electrochemical reaction in the

PEMFC) to condense at much lower concentrations. This is problematic because when it occurs inside the flow field of a cell, it will block the gaseous oxidiser from following that flow field, starving some of the active area and causing a severe and immediate loss of performance, as well as longer term, irreversible degradation. While an air heater was manufactured and installed in the cathode system, as well as the electronics to control it, it was not used for the testing because there was insufficient time to overcome the safety considerations. Namely, the heater required mains (240V AC) power, but was present in a very humid environment. While most of the time sufficient grounding and fusing of the circuit would be sufficient (and was of course designed in), the potential for burns, or thermal damage to other subsystems was more difficult to mitigate. While the air supplied to the cathode was reasonably uniform in temperature, it was not warm enough to allow operation at higher current densities. This is unfortunate, as the impact of current density on the performance of the technique under investigation is of great interest, and would undoubtedly give insight into the mechanism at work. However, because of this shortcoming of the system, all testing was performed at 50A. While the exact active area of the Powercell S2 stack is not disclosed, it is expected to be on the order of 100cm^2 per cell (based on performance characteristics, physical dimensions and available literature), meaning testing was performed at approximately 0.5 Acm^{-2}.

Consistency of flow rate is slightly less vital, provided sufficient oxidiser is supplied to the catalysts surface. This would not be the case in an open cathode system, where the cathode flow is also required for cooling, and this was one of the reasons for selecting a closed cathode system. An inlet stoichiometry of approximately 2 is advised, although it was found that a slight overestimate aided stability of performance, and was easy for the system to achieve at the reduced operating current used. Flow was measured by an SMC PFMC7501-F04-F, which is marketed as a "Flow Controller", but does not have any actuators to effect the flow itself, but rather is intended to be linked to an actuator using an analogue circuit. However, for the purposes of this testing it was required to record the data, so the output was sampled digitally. This also allowed it to be used as the input for a tuned PID control loop, the output of which was a variable solenoid valve placed at the outlet of the cathode flow. This controlled the flow

rate of the system by varying the size of the outlet aperture, and once tuned was easily capable of maintaining the required setpoint to within the error of the measurement device (1 SLPM). The valve used for this was taken from a Ford Focus Exhaust Gas Recovery system, as this was robust, cheap, easily replaced, gave sufficient resolution and repeatability of the actuator position, and was also low voltage. This was controlled via a 12V supply, MOSFET, PWM signal and a simple low pass filter. The PID control was tuned using the Ziegler-Nichols method, with the FC removed from the loop and replaced with a pipe with constriction to mimic the pressure drop it produces. This was done to ensure that the FC was not damaged due to an overpressure during tuning, as Ziegler-Nichols requires steady state oscillations to be induced in the system, resulting in rapid pressure fluctuations. The constriction was not tuned to very accurately mimic the FC, as this was not necessary. The pressure drop across the FC is not significantly larger than that from the humidifier or flowmeter, and so the total system behaviour (with respect to cathode flow rate and pressure drop) is not significantly different without it. It was sufficient to bypass it, perform tuning and validating the stability of the calculated values by demanding step changes in the target value. The control system met the largest expected step change in target value with no observable overshoot and within 2s, which is significantly faster than the FC can meet a similar step change in load, even under ideal conditions.

Finally, the flow needed to be dry from any oils or VOCs, but near saturated humidity. The importance of humidity to the performance of PEMFC's has already been elucidated in Section 5.2.3, and was achieved using an exchange membrane humidifier, namely the FumaTech H10. This consists of two counter flows (a "dry" and a "wet"), separated by a membrane permeable to water via osmosis, but not gaseous flow. This membrane is wetted by the "wet" fluid flow, and passes this water to the "dry" by evaporation. The wet flow is connected to the output of the stack, where the water produced during operation can be collected. This water can be either gaseous or liquid, allowing the system to be rapidly conditioned for reliable operation by the injection of liquid DI water on startup. While the unit selected was slightly oversized for this application (designed for 10kW systems), the next unit down would have been undersized, and it was deemed that an oversized humidifier was preferable to the added

complexity of multiple smaller ones, or a single undersized one. Indeed, humidity control was not troublesome, with the exception of flooding at higher current densities, but as has been explained this is due to other shortcoming of the system, not the humidifier. The oils and VOCs can be incredibly damaging to a PEMFC, however, especially those produced by silicone or sulphur cured vulcanised rubbers. As such, no silicone products of any kind could be used in the system, especially curable sealants, and fibre reinforced EPDM tubing had to be sourced for connecting the components. In addition, the compressors used were required to be "dry" oil free, and an SMC AF20-N02C-CZ-A 5μm filter was installed between them and the rest of the system to catch any errant lubricants aerosolised in the flow.

Anode

The anode system was slightly more complex, not because of more demanding conditions, but because of the dangers associated with pure hydrogen fuel in a laboratory setting. All equipment must be rated for containing pressurised hydrogen (which limits selection of polymers, metal alloys as well as mechanical seals), and no source of ignition can be present in the system at any point. This made heating the flow challenging, as an electric heater would have the potential to fail in such a way as to produce a source of ignition. To solve this, trace heaters were installed outside the pipework (designed for outdoor, agricultural use to prevent pipes from freezing in inclement weather, these locally self limit, and are incapable of heating over 65°C, including in any known failure mode). Additionally, coils of 316 stainless steel pipe were fabricated and submerged in the coolant reservoir, allowing heat exchange between the warm (typically 65°C) coolant and the fuel flow. All plumbing connections utilised tapered socket connections, which has been the industrial standard for repeatable and secure hermetic connections since its use in steam locomotion, and were made similarly of 316 stainless steel. The majority of the pipework was Nylon as its flexibility increased ease of integration and reduced risk of leakage by relieving strain on the plumbing fittings. Custom brass adaptors were turned to connect the 1/4" pipework to the significantly larger inlet port of the stack. This port is oversized because the manifold is designed to accommodate a much larger number of cells (in the hundreds for automotive use), and is common to both the anode and cathode flows, hence

must be sufficiently large to flow 1000's of SLPM, despite this particular case requiring 10's SLPM.

While the cathode system can be simply exhausted (albeit into laboratory extraction, to mitigate any safety risks due to fuel crossover), it would be exceptionally wasteful to do so with the anode, as this is supplied with fuel, as opposed to oxidiser in the form of air. This is especially true because the anode requires a stoichiometry of >1.5 in order to ensure uniformity in the flow field, a common design criteria of PEMFC stacks. Therefore the anode was operated as a closed loop, with a diaphragm pump installed on the outlet of the stack in order to recirculate the fuel ad achieve the required inlet stoichiometry. Fuel was supplied to the system via an Omega FMA5528A-ST flow controller (as well as a number of safety and control valves to shut off the flow of fuel in an emergency and allow the system to be flushed with inert nitrogen for safe shutdown and maintenance). The intention was that this flow controller would supply a uniform flow rate, and the pressure of the system would be managed via occasional purging (required to prevent the build up of impurities and liquid water anyway). However, the flow controllers feedback proved to be too unstable; oscillating wildly, as well as causing rapid fluctuations in the anode pressure. To overcome this, its flow rate setpoint was set to a nominal value greater than the expected fuel consumption of the system, and a separate control loop was implemented to supersede the devices internal control which acted based on pressure. Due to the significantly lower flow rate required compared to the cathode, as well as the closed nature of the system, this proved to be both much more stable, as well as much more reliable, and ensured that the anode was pressurised to 800m bar at all times. Higher pressures could have been used, but given the devices limit of 1 bar as well as the limitations of the cathode systems pressure regulation it was felt that some headroom in case of emergencies was warranted, and the increased performance would have no significant impact on the research questions under investigation.

The initial design of the rig utilised a glass bubble humidifier, which could be refilled without depressurising the system to open it by injecting DI water into it's inlet, and would be heated by being placed in a water bath (such as the coolant reservoir). However, due to the difficulty in achieving a hermetic seal, and the safety considerations, this was omitted and some initial tests

of the rigs operation were performed to determine the necessity of humidification at the anode. It was found that, perhaps due to it's closed system, no added humidification was required. In fact, it also suffered from flooding - even at moderate currents - despite its flow being heated by the coolant. Part of this was because these tests were done without raising the temperature of the coolant into it's usual operating conditions, meaning the fuel flow was colder than it typically would be, but as no water had been injected into the flow and there was a considerable amount of condensate - even at low current densities - it was clear that it would need to be removed during operation for reliable performance. As such, an SMC condensate trap was installed on the stack outlet. It is expected that an automated purging protocol would also be necessary to ensure good operation at the highest current densities by periodically sending a pressure wave through the flow fields of the stack to dislodge any potential blockages due to condensation. Unfortunately, this would also add significant complexity and/or require a good amount of testing to tune appropriately and so would have required more time than was available.

Coolant

The coolant system was relatively simple, and consisted of an intentionally oversized reservoir (to prevent rapid changes in coolant temperature supplied to the stack), 24V DC pump, a computing radiator and fans, and a few thermistors. The thermistors were placed in the fluid flow at connections in the plumbing, and were selected to have very thin connection wires with sufficient length to allow hoses to seal over them. The pump and fans were controlled by the Process Controller (via PWM and MOSFET), using basic hysteric control based on the reading at the input of the stack.

Cell Voltage Monitor

The Powercell S2 stack is fitted with a D-SUB 25 connector which allows connection to the voltage taps for each cell. The CVM designed for this rig connects via the same connector, and uses 6 MAX22531 isolated Analogue to digital converters (ADC's), each reading 4 consecutive cells, their voltages dropped to within it's voltage resolution via potential dividers using high

accuracy resistors (<0.1% accuracy, <10PPM temperature drift). These IC's were chosen as they have an SPI interface, which is straightforward to access, but also has a sufficiently high rate of data transmission to achieve the high sample rates desired. The IC's continuously sample at $20kSs^{-1}$, and can be polled for either the most recent set of readings, or a moving average of the last 4, which was of use in initial debugging and noise rejection. For the data gathered for scientific (as opposed to rig fabrication) purposes, only the single readings were used. The SPI interface was managed by an RP 2040 MCU, which also communicated with the Load Controller in order to synchronise the high frequency acquisition. This was necessary due to the rate at which this generated data, as it would not be feasible to record at this sampling rate continuously. The MCU also recorded directly to an onboard SD card (in plaintext during it's normal acquisition rate and binary during high frequency acquisition), as well as sending the data to the central data logger during normal acquisition. The maximum rate of acquisition was limited by the speed at which the SD card could be written to, which was found to be stable up to $7kSs^{-1}$, but was set to $5kSs^{-1}$ for testing. Onboard flash storage could have increased this, but was deemed unnecessarily complex to implement. The RP 2040 was specifically utilised because it has 2 cores, and could therefore be set up to have one core continuously reading data into it's internal memory buffer, and the second writing it out to the SD card. This hugely minimises executive overhead as each core needs only check that it should continue, and not decide which of multiple tasks to do. This more than doubles the data rate that could be achieved by a single core.

Load Controller

This system controlled the Kikusui PLZ 1205 electronic load via it's external control connector, as well as controlling the current perturbation circuits, including implementing the various optimisation strategies utilised. While these strategies were kept simple, intending to demonstrate a robust system that was agnostic to hardware, they are closed loop and so require an input. For this purpose, the Load Controller could also measure the total stack voltage which, as explained in Section 5.3.2, under constant current operation is directly proportional to power. This yields a lower resolution than, for example, using the CVM's reading and summing them,

but also doesn't suffer from the potential noise, reliability and added complexity issues, and the resolution is more than sufficient to encompass the response from the technique under investigation.

The perturbations were achieved using a large IGBT as a solid state switch, in series with a resetable mechanical circuit breaker. The former allowed very accurate and precise switching of the current path, while the latter protected the system (primarily the FC) from damage in the event that the IGBT somehow failed in the closed state. It was selected such that it would not activate under short bursts well above its nominal limit (250A), but would for sustained currents. This was important because the expected current through this circuit was composed of an initial high peak (from the discharge of the double layer capacitance of the PEMFC) followed by a rapid decay to an unknown current above the PEMFC's nominal maximum current (between 200-250A, depending on operating conditions). Because of these unknowns, a mechanical system provided the most appropriate response without additional added complexity. The system was tested for timing accuracy, and found that the maximum deviation from the duration requested was 50 nS, with a random error on the order of 1-2nS. This deviation is due to the length of the clock cycles in the controlling MCU, and is well below the threshold of influencing the response from the PEMFC. These results were measured using an oscilloscope, which unfortunately did not have the capability of digitally exporting the data.

The current perturbation control could operate in three separate modes. The first, fixed parameter mode, applied perturbations according to duration and period values that were manually set, and did not vary them in any way. Secondly, two degrees of freedom perturb and observe (P&O), which began from the duration and period values it was manually given and optimised them each in turn using the P&O technique. Step size was multiplied by $\frac{-1}{\phi}$ if the gradient was negative (i.e. if the most recent steps average voltage was lower than the previous step), thus reducing it and allowing more rapid initial convergence, but more accurate a final result. Once a minimum step size was reached, the parameter being optimised would be fixed, and the other DoF would be optimised. This would repeat indefinitely, allowing the

system to adapt to changing operating conditions. Finally, adaptive τ mode uses the analytical solution found in Sec 5.3.4 to optimise the period of the duration, and 1 DoF P&O to optimise the duration. This allows them to be optimised simultaneously, and significantly reduces the convergence time.

Chapter 6

Current Perturbations for PEMFC Systems - Results and discussion

6.1 Introduction

This chapter will present and examine the data gathered using the apparatus described in the previous chapter. This will begin with a brief determination of the quality of the data output by the custom test equipment designed for the purpose, before examining the behaviour of the PEMFC in response to a perturbation. Firstly, on the time scale of a perturbation, examining what occurs during and immediately afterwards in order to glean some insight into the mechanism through which the technique acts. A variety of proposed theories have been presented and examined in the previous chapter, and the data presented here is intended to be the next step in investigating these theories. Next, the behaviour of the PEMFC on longer timescales will be investigated by evaluating the use of a variety of optimisation strategies that hope to improve the benefit gained from an implementation of this technique. These strategies will be implemented, and the quantified benefit of their application demonstrated before being compared to a baseline (i.e. no current perturbation applied), as well as to each other. It should be noted that there is no publication (that the author could find) which has attempted to achieve this live, *in operando* optimisation, and that the strategies here are intended primarily for scientific inquiry and as an initial step. Actual commercial or industrial

implementations are likely to be significantly different, and some speculation as to these more mature implementations based on the data gathered is presented at the end of this chapter.

6.2 Results

The first question to answer in this study is "What impact does current perturbation have on a modern automotive style PEMFC?". Primarily, does it produce a voltage increase, and of what magnitude, but also is the shape of the response as expected. In addition, the capabilities of the measurement equipment and stability of the rig should be examined, due to their bespoke and otherwise untested nature. Fig 6.1 shows the voltage response of the stack (at constant current equilibrium) to a perturbation of 100ms in length. The individual voltages of each cell have been plotted, overlaid with a mean of these measurements, as well as a mean of the maximum voltage the cells reached directly following the perturbation (V_{max}). It can be seen that the technique does produce a significant voltage increase (of 48.5mV, or 6.3% of V_{equi}). Before examining the shape of the response, first the capabilities of the measurement equipment should be evaluated.

6.2.1 CVM Data Validity

Individual data markers have not been displayed on this plot for clarity purposes, but a uniform time-step of 200 μs (a sampling frequency of 5 kHz) is achieved, with the greatest deviation being 1μs (the smallest unit of time measured by the MCU in this application). This gives us adequate resolution to measure the smallest features exhibited, including multiple points on the ohmic response immediately following the perturbation (which have been labelled for demonstration purposes in Fig 6.2). Perhaps a small amount of detail is lost in the Ohmic response at the end of the perturbation, but it is not expected to be of significance, and the error from other sources (such as the impact of inductance in the electrical connections to the stack) is expected to be much greater. In addition, no smoothing or filtering has been applied to this data, and it can be seen that the noise is significantly smaller than the voltage response we are examining. The noise is greater than was intended when designing the measurement devices, with the main source of error expected to be electromagnetic interference in the signal cable, but is still

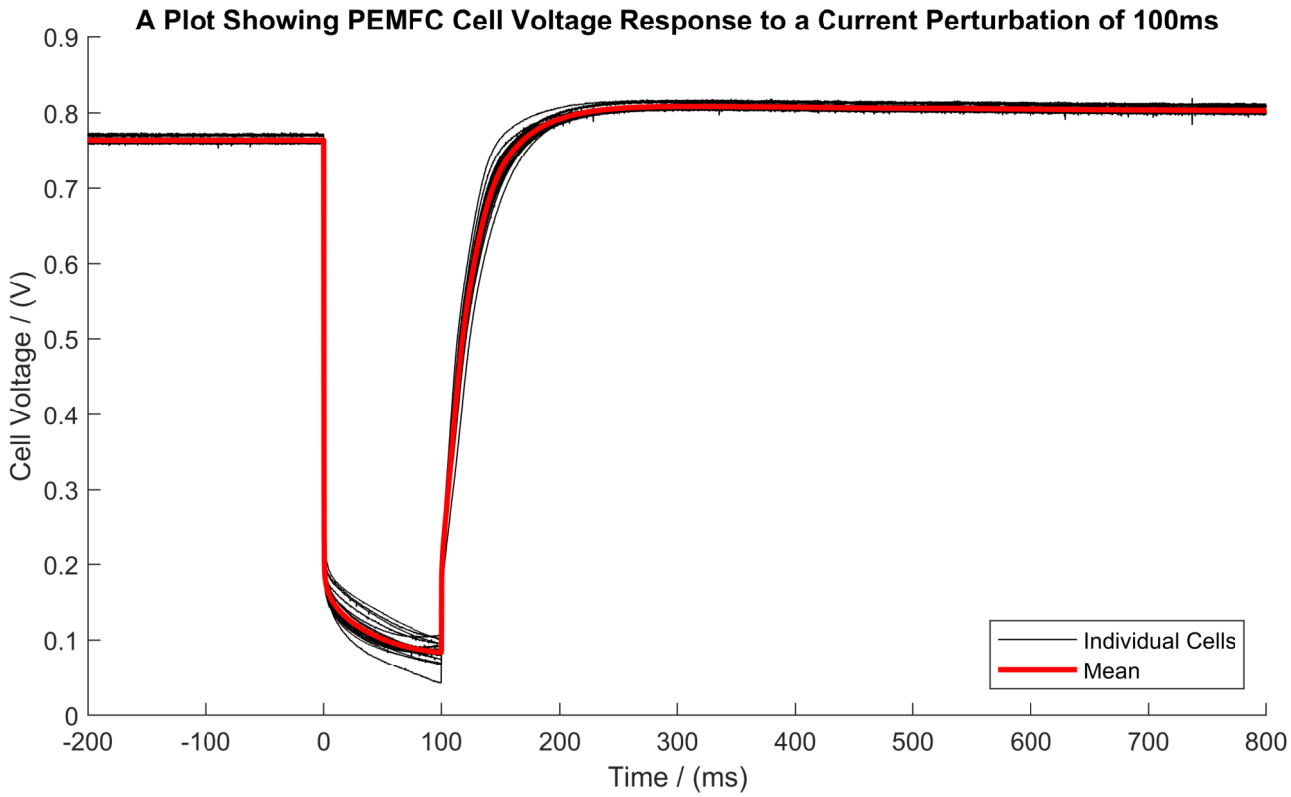


Figure 6.1: A plot of the voltage response of the stack from constant current equilibrium to a perturbation of 100ms. Operating conditions: 50A (estimated to be approximately 0.5 Acm^{-1}), P_{anode} = 0.8 bar, $P_{cathode}$ = 0.7 bar, $T_{coolant} \approx 60°C$.

small enough to be of no concern. The Peak-to-Peak value of the noise is approximately 4mV, which is smaller than the cell to cell variability of the stack, measured to be 13mV under these operating conditions. This was measured during the constant current equilibrium immediately preceding the perturbation, as the cell to cell variability during a perturbation is significantly higher and is not indicative of the devices general operation. As such, the author is pleased with the performance of the CVM, and is satisfied that any shortcomings it has are sufficiently minor as to not significantly impact the validity of the data gathered.

6.2.2 Voltage Response During a Perturbation

Turning attention to the scientific questions, as opposed to metrologic, we can see that at 50A (estimated to be approximately 0.5 Acm^{-1}), operating at $P_{anode} = 0.8$ bar, $P_{cathode} = 0.7$ bar, $T_{coolant} \approx 60°C$, the Powercell S2 stack will produce a mean of 0.764 V per cell (for a total stack power of 458 W). Immediately following a perturbation of 100ms, this increased to an average V_{max} of 0.812V, an increase of 51 mV, or 6.7% of V_{equi}. This average V_{max} was taken from 32 perturbations, with a τ of 30s. We will examine the response post-perturbation below, but it should be noted here that 30s is insufficient time for the voltage to decay back to V_{equi}.

Upon activation of the current perturbation (Time = 0 in Fig 6.2), the voltage drops immediately and rapidly to approximately 0.2 V per cell, before exponentially decaying and approaching an asymptote of 78 mV. Admittedly, the ohmic drop is larger than initially expected, but given that these expectations were not based on any significant data, this deviation from predictions does not invalidate our initial assumptions. Indeed, the shape is otherwise as expected, with a smooth exponential decay. It can be seen that the cell to cell variability is significantly larger during the perturbation, which is unsurprising, but the variation appears to be primarily in ohmic drop, with the shape (i.e. time constants and asymptote) of the voltage response for each cell being very similar, and the difference between cells appearing to be roughly constant. While some of the individual cell lines do cross each other, it is true that those that exhibit the smallest initial voltage drop end the perturbation at the highest voltage, and those with the

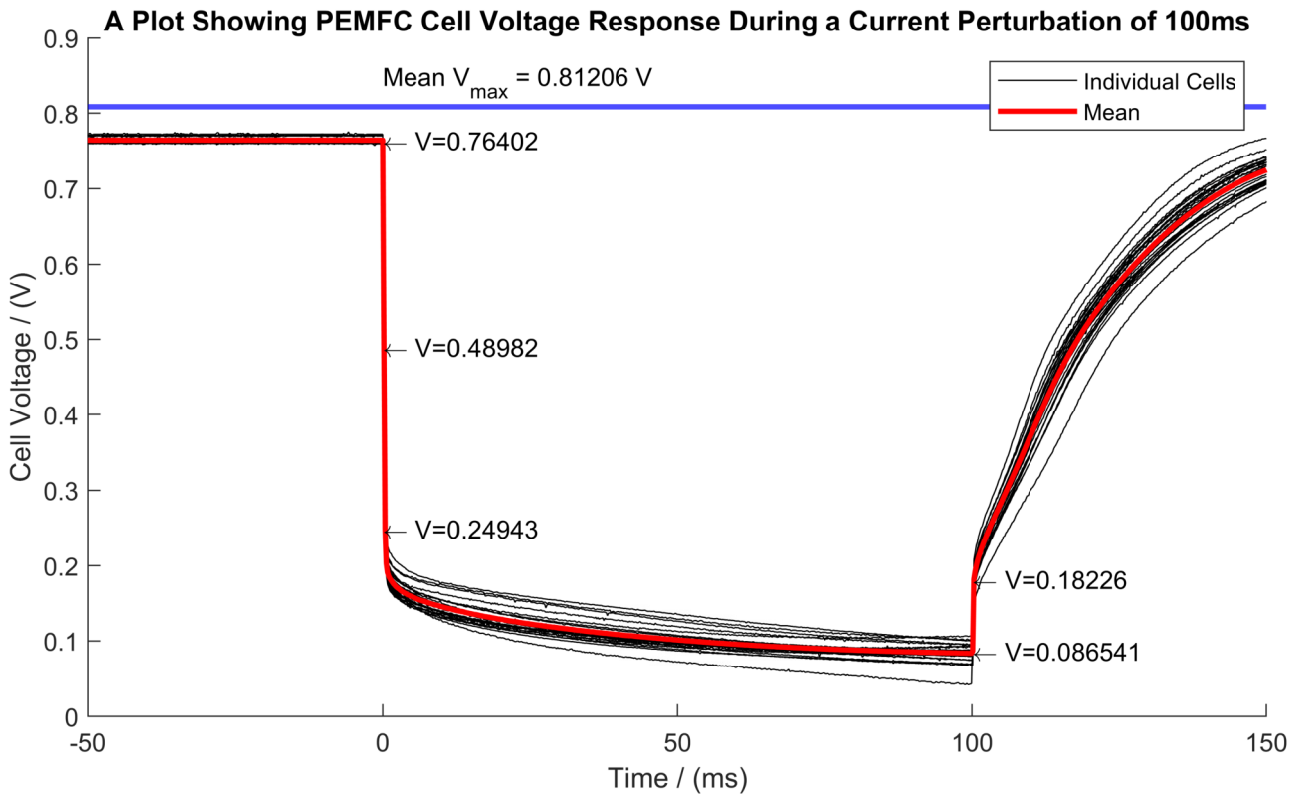


Figure 6.2: A plot of the voltage response of the stack during a perturbation of 100ms

largest drop end at the lowest. There is one particular cell (Cell 24) which has a significantly lower voltage than the others. This difference was measured repeatably, but it is unclear if it is due to fluidic effects (i.e. mass transport of fuel or oxidiser to this cell is more restricted, either due to cumulative pressure loss or specific flow geometries, such as a sharp turn in the connection between the flow field and the manifolds), electrical effects (by virtue of being the cell closest to the external electrical connections of the stack), or electrochemical (this cell may have experienced higher degradation during storage by virtue of being closest to one of the manifolds external fluid connections). Alternatively, it could be the first signs of flooding, although this is less likely, as all cells were capable of operating at over 80A with this configuration of the rig before any significant signs of flooding were exhibited. It was decided not to exclude this cell from the results, as without a reasonable insight into the source of the anomaly it is impossible to do so with any confidence.

To further examine the voltage response during the perturbation, a model has been fitted. This takes a similar form to that presented in Section 5.3, Eqn 5.12, and can be seen in Eqn

6.1. The generic exponential coefficient "α_n" has been replaced with $\frac{1}{RC_n}$, so that the model approximates n capacitances discharged over a single resistance. In addition, it is taken that t=0 at the beginning of the perturbation. For the data presented in Fig 6.2 (specifically the calculated mean cell voltage), 2 capacitances (i.e. 2 separate time constants) produced a reasonable fit, with 3 producing an excellent fit. The major difference is only for values of t<0.2 ms, and given that both models struggle to fit to the data for values of t<5 ms, this is not deemed enough of a significant improvement to justify the extra terms. It is preferable to use the fewest terms possible while achieving an accurate fit, as given sufficient terms it is possible to make any model fit nearly any data, and so the confidence in the validity of the model decreases with increasing n in this scenario. This would not be true if we had reason to ascribe each RC pair to a real physical process, but that is not the case here. We are merely using a common model to approximate the behaviour of a complex electrochemical system, the sub-reactions of which are expected to have reaction rates governed by an Arrhenius relationship. It is posited that the reason for the poor fit at values t<5 ms is due to the rapidly fluctuating current as well as the influence of factors external to the PEMFC, such as the inductance of the connections between it and the perturbation circuit. For currents of this magnitude and time scales of this length it is not unreasonable to expect significant influence from these factors given the current path takes a round trip of a few metres through 70mm^2 cable comprised of hundreds of fine strands of copper alloy. By excluding the measurement points below 5 mS, a somewhat closer fit can be found, but the improvement is not significant enough to explore further.

$$V(t) = V_{asym} + V_0 \sum_n e^{-\frac{t}{RC_n}} \tag{6.1}$$

These fits gave time constants of 38.7 ms and 1.3 ms for n=2, and 66.9, 2.3 and 0.0006 ms for n=3. This demonstrates how minor the difference between the two is, and also why adding increasingly more terms is unnecessary. In order to fit a larger order model (such as n≥3), we would need an order of magnitude less noise, as well as at least 2 orders of magnitude higher sampling rate (our current time step is 200 μS, compared with our smallest time constant of 0.6μS). Even then, the impact of external factors (such as the impedance of the perturbation circuit, details about the active material and GDL in the device etc.) would be significantly

larger than the contribution of the smallest time constant term of the model. For reference, in the current time constant calculated for this term, light in a vacuum travels 177m, which is just slightly higher than St. Pauls Cathedral. Given the processes we are interested are electrochemical, and so many orders of magnitude slower, it is felt that in this case, more than 2 terms is superfluous. While we could manually force the time constants to be closer together, and potentially get a small increase in the goodness of fit, this rapidly becomes an exercise in overfitting. This is because we would be trying to reduce the size of error in the model well below the measurement error of the input data, and ultimately produces a less rigorous scientific proposition.

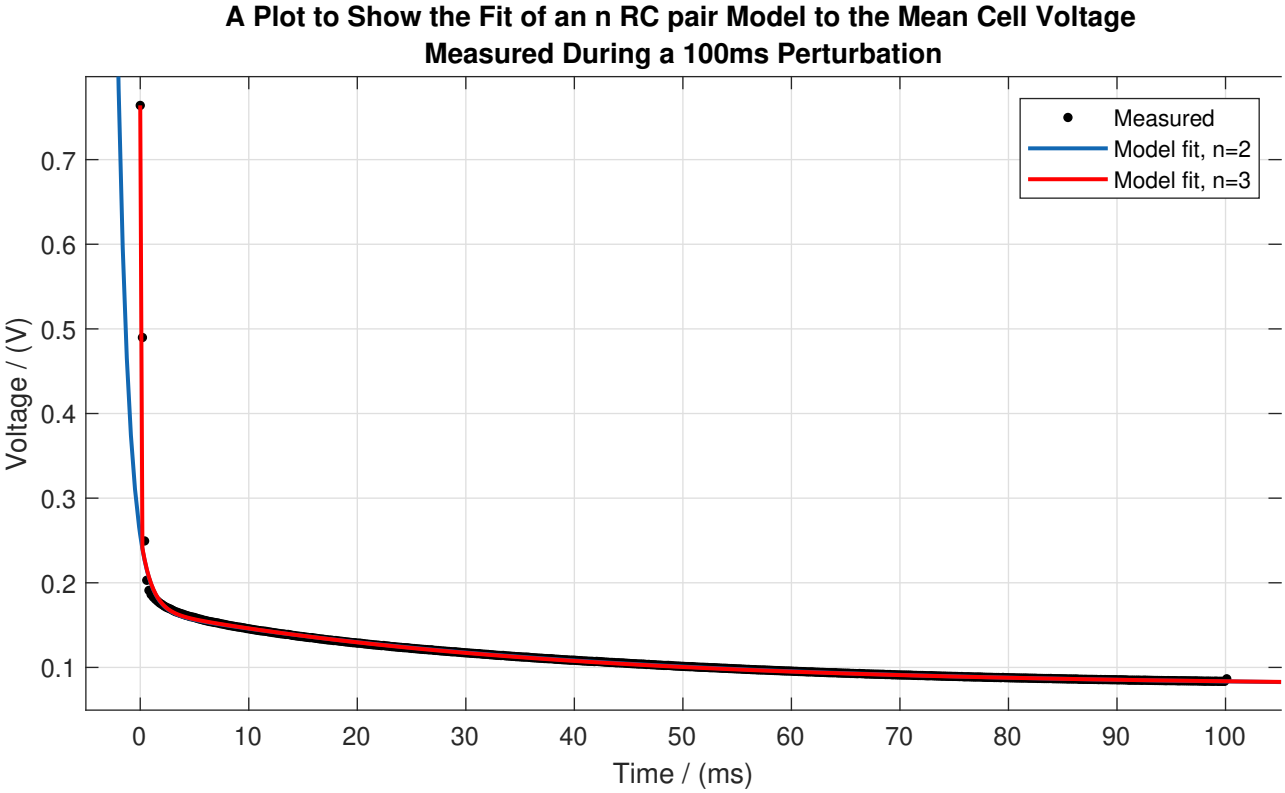


Figure 6.3: A plot demonstrating the fit of the proposed model to the voltage response during a perturbation for values of n=2,3

6.2.3 Voltage Response Immediately After a Perturbation

As can be seen in Fig 6.4, the response immediately after a perturbation has a much smaller ohmic region. This is likely due to two factors. One, the step change in current at the end of this perturbation is much smaller than at the beginning. Immediately before the perturbation

the FC is operating at a relatively low current density (compared to it's designed capability), and the mass transport losses are therefore minimal, and far from limiting the power output. This means that when the perturbation is applied the current jumps from 50A to between 800-1200A within micro-seconds, and hence gives a very large immediate drop in voltage. At the end of the perturbation, however, the double layer capacitance of the cells is nearly fully discharged, and the current demand is approaching infinite (as the perturbation circuit has a very small resistance in comparison with the rest of the system). As such, the current is limited by mass transport of reactants to the reaction sites, and so the current delivered is closer to 300-500A, switching back to just 50A at the end of the perturbation. In addition, the fact that the reaction is mass transport limited means that overcoming this limitation has a longer time constant, as the aggressive diffusion gradients produced by running the electrochemical system in essentially a short circuit must relax back to their usual states during standard operation. It is this relaxation we can see in the voltage response post-perturbation.

Once again, during this transient phase, the cell to cell variability is slightly higher, but it quickly returns to normal as the cell voltages recover back to their maximum. The shape of the response is largely as expected, with the recovery time being insignificant with respect to the τ of these perturbations. The mean time required for the cells to reach maximum voltage after the cessation of the perturbation under these conditions was 213 ms. While this does increase the effective time the device cannot deliver the requested power to the load, it does not do so significantly. As can be seen in Fig 6.1, the time taken for the voltage to recover to a value greater than V_{equi} is closer to 50-75 ms. We can attempt to calculate this more accurately, although some care must be taken in doing this. Remember that these perturbation were done with a 30s period, and so the cells did not have sufficient time to fully relax back to V_{equi}. Calculating the time it took each cell to recover to the voltage measured immediately prior to a perturbation slightly over-estimates recovery time, as this voltage threshold is higher than V_{equi}, (the voltage that would be delivered without the use of perturbation). This was calculated to be an average of 78 ms for these perturbations. However, the time taken to recover to the V_{equi} measured prior to applying any perturbations was found to be 66ms. While the difference

between these number is large in relation to their total value, it is somewhat inconsequential to the reason for calculating them. They are both significantly smaller than the duration these perturbations were applied for, and so the time for which power is not delivered to the load is not increased by more than a factor of 2, which is still an insignificant amount of time in relation to the τ applied.

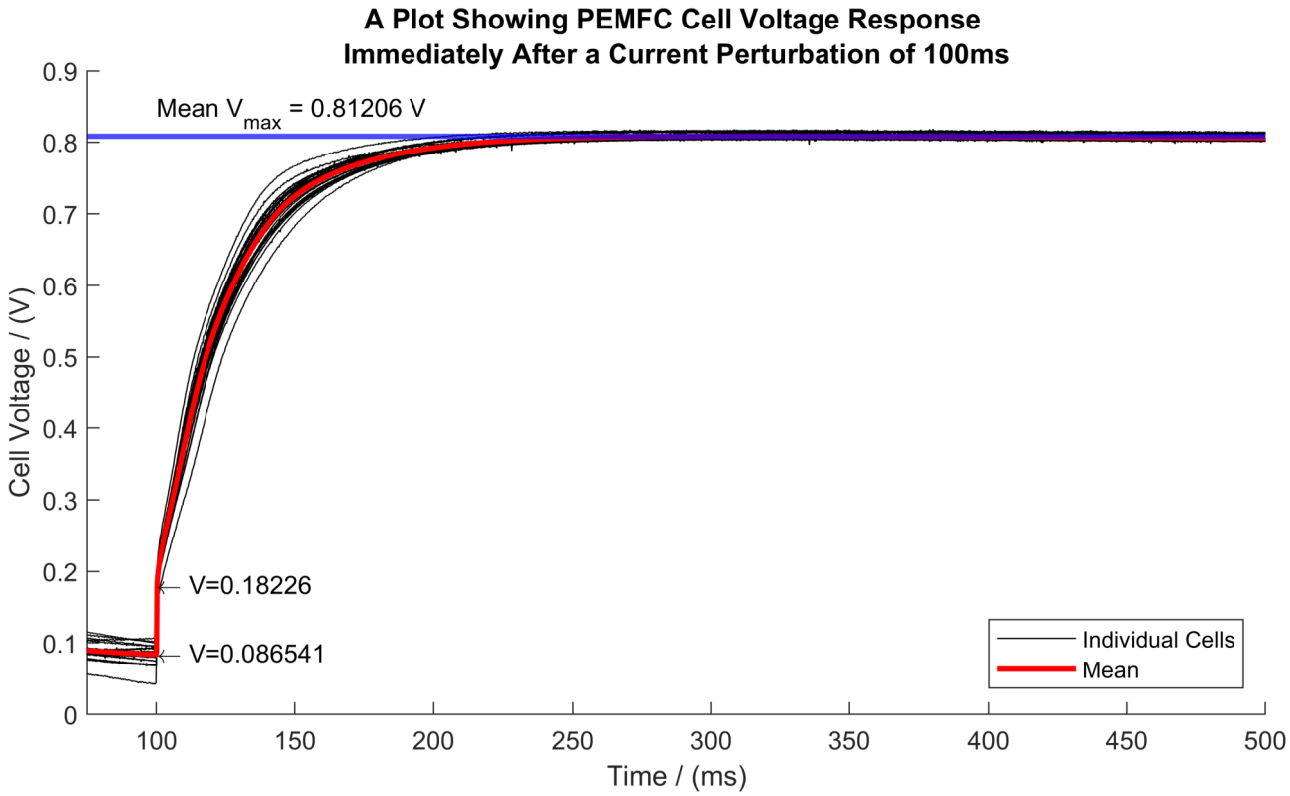


Figure 6.4: A plot of the voltage response of the stack immediately after a perturbation of 100ms

6.2.4 Voltage Response Between Perturbations

In order to examine the voltage response after a perturbation we will use a different set of data, where the PEMFC was allowed to relax for a much longer period after a perturbation. While this long relaxation is not directly relevant in the use of this technique, as it represents a τ orders of magnitude larger than would be applied in operation, it is very useful for the validation of the model used to find an analytical optimum for τ. This model, presented in Eqn 5.12, does not necessarily need to fit well for the optimisation to work, as it has been shown that

the conclusion drawn at the end of the derivation is true as long as the function $V(t)$ fulfils the criteria that $\frac{f_{mean}}{d\tau} = \frac{1}{\tau}f(\tau) - \frac{1}{\tau^2}\int_0^\tau f(t)dt$. However, a good fit of the model to the experimental data would certainly lend confidence to the validity of the assumptions made in the creation of the model, which are therefore used in the assertion that the optimum τ occurs at $f(t) = f_{mean}$. In addition, while the models terms do not represent any specific physical processes occurring, they do act to capture the relative time constants of the many processes that contribute to the output voltage, and so can give some qualitative insight based on their orders of magnitude.

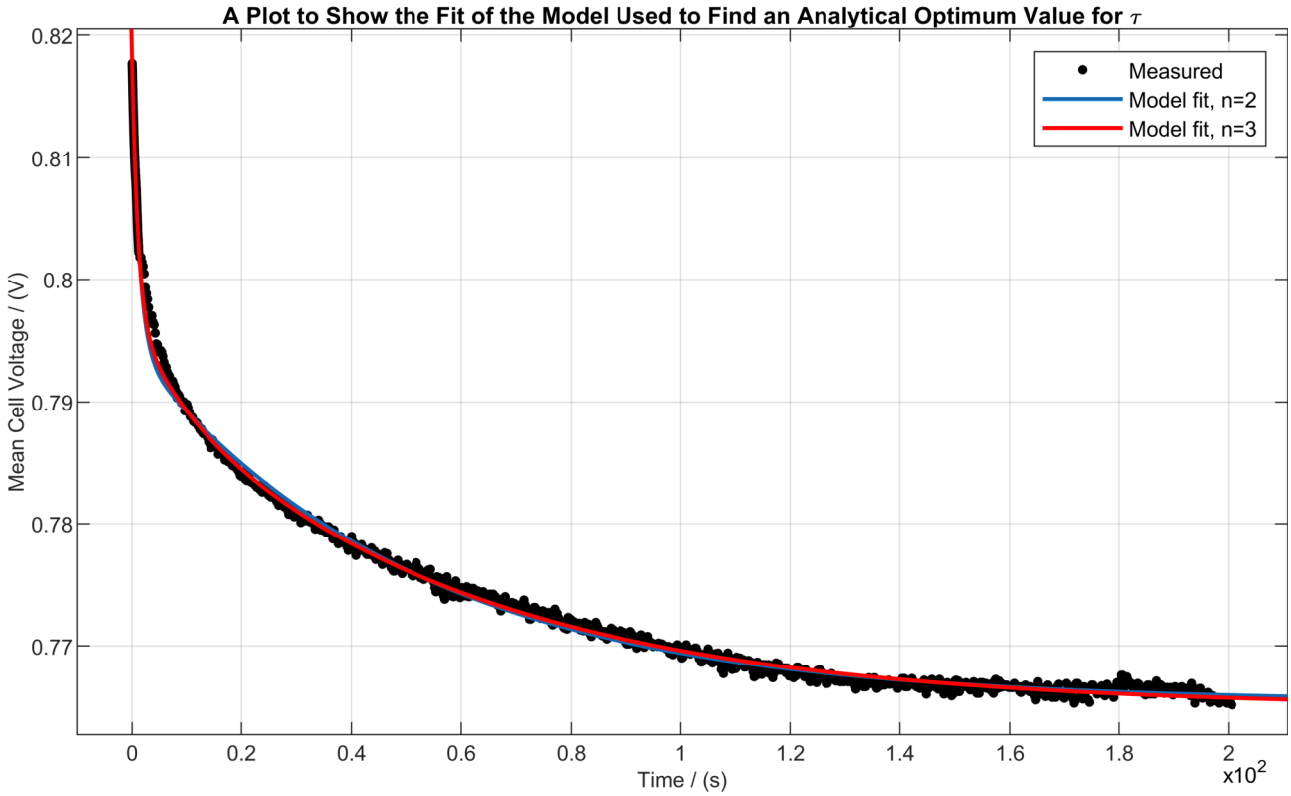


Figure 6.5: A plot of the voltage relaxation of the stack after a perturbation

It can be seen in Fig 6.5 that a good fit has been achieved. Once again, 2 terms is enough to capture the shape reasonably well. It is possible that a better fit for 2 terms is achievable, but as the test here is a question of if it can fit, as opposed to using how well it fits to quantify something (for example model error), additional time to fine tune is not necessary. A 3 term model does capture the shape of the curve somewhat better, and could be justified here as this occurs over a much longer time period than the fit in Fig 6.3. The smallest time constant term in the 3 term fit presented in Fig 6.5 is 11 ms, which is sufficiently long to represent the

behaviour of real processes occurring, compared to the previous fitting examples 600 ns, which is verging on too small to be capturing any bulk processes. I.e., while there are of course processes occurring at that time scale and smaller, they are not sufficiently small for that time scale to be capturing their aggregate behaviour. In addition, the difference between the 2 and 3 term fits in Fig 6.5 is not a smaller time constant term. The smallest time constants for the n=2 and n=3 cases are 1.3 s and 1.1 s respectively, and the largest are 50.6 s and 57.4s, with the additional term in the n=3 case being 10.5 s. The addition of this term does not attempt to expand the time scale captured by the model, but instead acts to increase the goodness of fit within an acceptable time scale that is already captured. It does widen the time scale by a very small amount (the smallest and largest time constants are pushed slightly apart by the inclusion of one in the middle of their range), but the scale of the change is negligible. As such, for this case it is believed that a 2 term model does capture the bulk of the behaviour here, but a 3 term model can aid in capturing the detail and getting a smoother and overall better goodness of fit. If this were an attempt to actually parameterise an accurate model for simulation or prediction purposes more effort would be warranted here, with an examination of the residuals, confidence bounds and initial values for the coefficients fitted. However, as the intention is merely to show that the proposed model can fit well to lend credence to its use for finding an analytical optimum for τ, no further analysis is necessary.

While some further analysis could also be done attempting to ascribe these terms to physical processes (mainly by grouping processes based on time scale), this requires a great amount of care to be done with adequate rigour, and for the purposes of this project, is not expected to yield significant insight. While the relaxation is important to the use of this technique, of primary concern is the scale of the overall time constant, as this determines the sensitivity of the mean power increase achieved to the value of τ. It does not set the magnitude of τ, which is primarily determined by the duration (d), peak voltage reached (V_{peak}), and the time scale of the rise to this voltage (as shown in Fig 6.4). Instead, the rate of decay of the stack voltage following a perturbation simply determines the magnitude of the benefit gained by optimising τ. A sharper decay would necessitate optimisation, whereas a shallower decay would be more

forgiving of sub-optimal values of τ. The speed of this decay is posited to be determined by 2 main factors: the design and frequency characteristics of the stack, and the current density or overpotential applied. The former will have some effect, but in a comparison between two mature stack designs (for example between the stack tested here, the Powercell S2 and another automotive grade stack, such as those sold by Elring-Klinger) there is expected to be little difference. Instead, the major difference would be between an open and closed cathode design, or between designs where one has an inferior flow field, and is therefore less capable of dealing with load fluctuations (such as may be used for an emergency backup generator, where the design considerations are focused on long term reliability, as opposed to a high number of operating hours, and where the system can be aggressively hybridised without issues associated with size or weight). Instead, it is expected that the latter factor, namely the position on the Butler-Volmer curve at which the stack is operating during the decay, will drive it's time scale. This is, itself, for 2 main reasons. Firstly, the difference between the steady state equilibrium voltage (V_{equi}) and the peak voltage immediately following a perturbation (V_{peak}) is expected to be larger at higher current densities. Secondly, the higher the operating current density (i.e. overpotential), the faster the aggregate forward rate of the reactions occurring, and therefore the reaction turnover at any given site must be higher. Therefore, as this project is operating under the assumption that the equilibrium theory (presented in Sec 5.2.4) is the dominant mechanism responsible for this voltage response, a higher forward rate will cause the relevant species concentration to return to their steady state equilibrium values more rapidly.

An important area for future work would be analysing the the impact of different current densities on the voltage response of the system during and after a perturbation. This was unfortunately beyond the scope and resources of this project, but the rig developed should be capable of performing the tests required. This is certainly the area of further study the author would recommend be tackled first, as it is key to the implementation of this technique in real world applications and might also yield further insight into the mechanism through which it functions, especially in quantifying and mapping it's behaviour across the expected operating range of the device. The impacts of other factors (such as pressure, temperature and

optimisation parameters) are likely to be much less significant than the current density applied. This means that current density is predicted to be the primary factor in determining where and when this technique can or should be applied, which in turn will dictate the form such an application would take. Further discussion on real world application of the technique will be done later, in Sec 6.3.

6.2.5 Optimisation

In order to fully examine the performance of each optimisation strategy a large amount of data would need to be gathered. This data would consist of a statistically significant number of runs in which the optimisation strategy is utilised in identical ways, and the time required to find an optimum (i.e. for the step size to fall below a threshold value for both parameters simultaneously) measured, as well as the position of this optimum. Unfortunately, this was beyond the resources of this project, and beyond the scope of the research questions identified. The purpose of this initial work is to determine if these optimisation techniques function, finding optima during live operation and to determine if they are sufficiently robust to be applied indiscriminately, with little to no tuning required. The data required to demonstrate this is significantly less.

Near the optimum values for both period and duration the sensitivity of the technique to variation in the input values is less than the voltage changes due to shifts in parameters such as coolant temperature and reactant pressures. While this makes it difficult to achieve clinical examples of an optimisation, it demonstrates the ease with which these optimisation techniques can be applied. Essentially, the initial estimates for minimum step size were too small, and the benefit of refining the values in step sizes $< 10\%$ of their total size is less than the benefit of constantly optimising them to account for changes in operating conditions. If the computational tasks or physical implementation of this technique was significantly higher, this might negate the benefit of live optimisation, instead suggesting that a simple multi-variable map produced by the system integrator during calibration would be preferred.

Perturbation Duration

In order to demonstrate that the perturbation controller can successfully optimise the duration of the perturbations and to better understand the speed at which it converges, reasonable start conditions were selected. These are not representative of a system that has been calibrated(e.g. with a multi-variable map giving informed guesses for initial conditions), but instead intended to represent the worst reasonable scenario, which is a blind system being applied to a new device with only the minimum required input (stack voltage). In addition, the operating conditions of this system were not fully stable. For example, partway through the optimisation the cathode pressure dropped from ≈80 kPa to ≈70 kPa, therefore reducing the stack voltage. This occurred because of the single stage pressure regulators combined with the hysteric control of the compressors. As the inlet pressure of a single stage regulator rises, it's output falls. Therefore, while the compressors are running, increasing the pressure in the reservoir, the outlet pressure will fall. The voltage drop caused by this change is impossible to accurately measure due to the dynamic perturbation protocol running, but based on other tests of the system it is expected to be on the order of 30-70 mV.

In order to demonstrate the systems function, Fig 6.6 shows a selection of measured variables. First, the instantaneous stack voltage measured by the perturbation controller. Next, a mean of this value, reset at the beginning of each perturbation. Next, a mean of the voltage over 5 consecutive perturbations (termed a "step"), required to ensure that the system is reliably reacting to changes primarily impacted by the input change, as opposed to random noise. In addition, in a live system approaching optimum values f τ and d the system does not fully relax back to its equilibrium state between each perturbation, and so there will be some cumulative effect. Namely, on a system that is repeatedly perturbed, a shorter d will be required to reach the maximum peak voltage following a perturbation when compared with a system that is allowed to fully relax back to equilibrium between each perturbation. For these reasons, as well as general rigour and good scientific practice, 5 perturbations are performed, the mean response taken, and this value is used as the input for the optimisation algorithm. It is this value that is plotted on the third axes of Fig 6.6. It should also be noted that this value lags

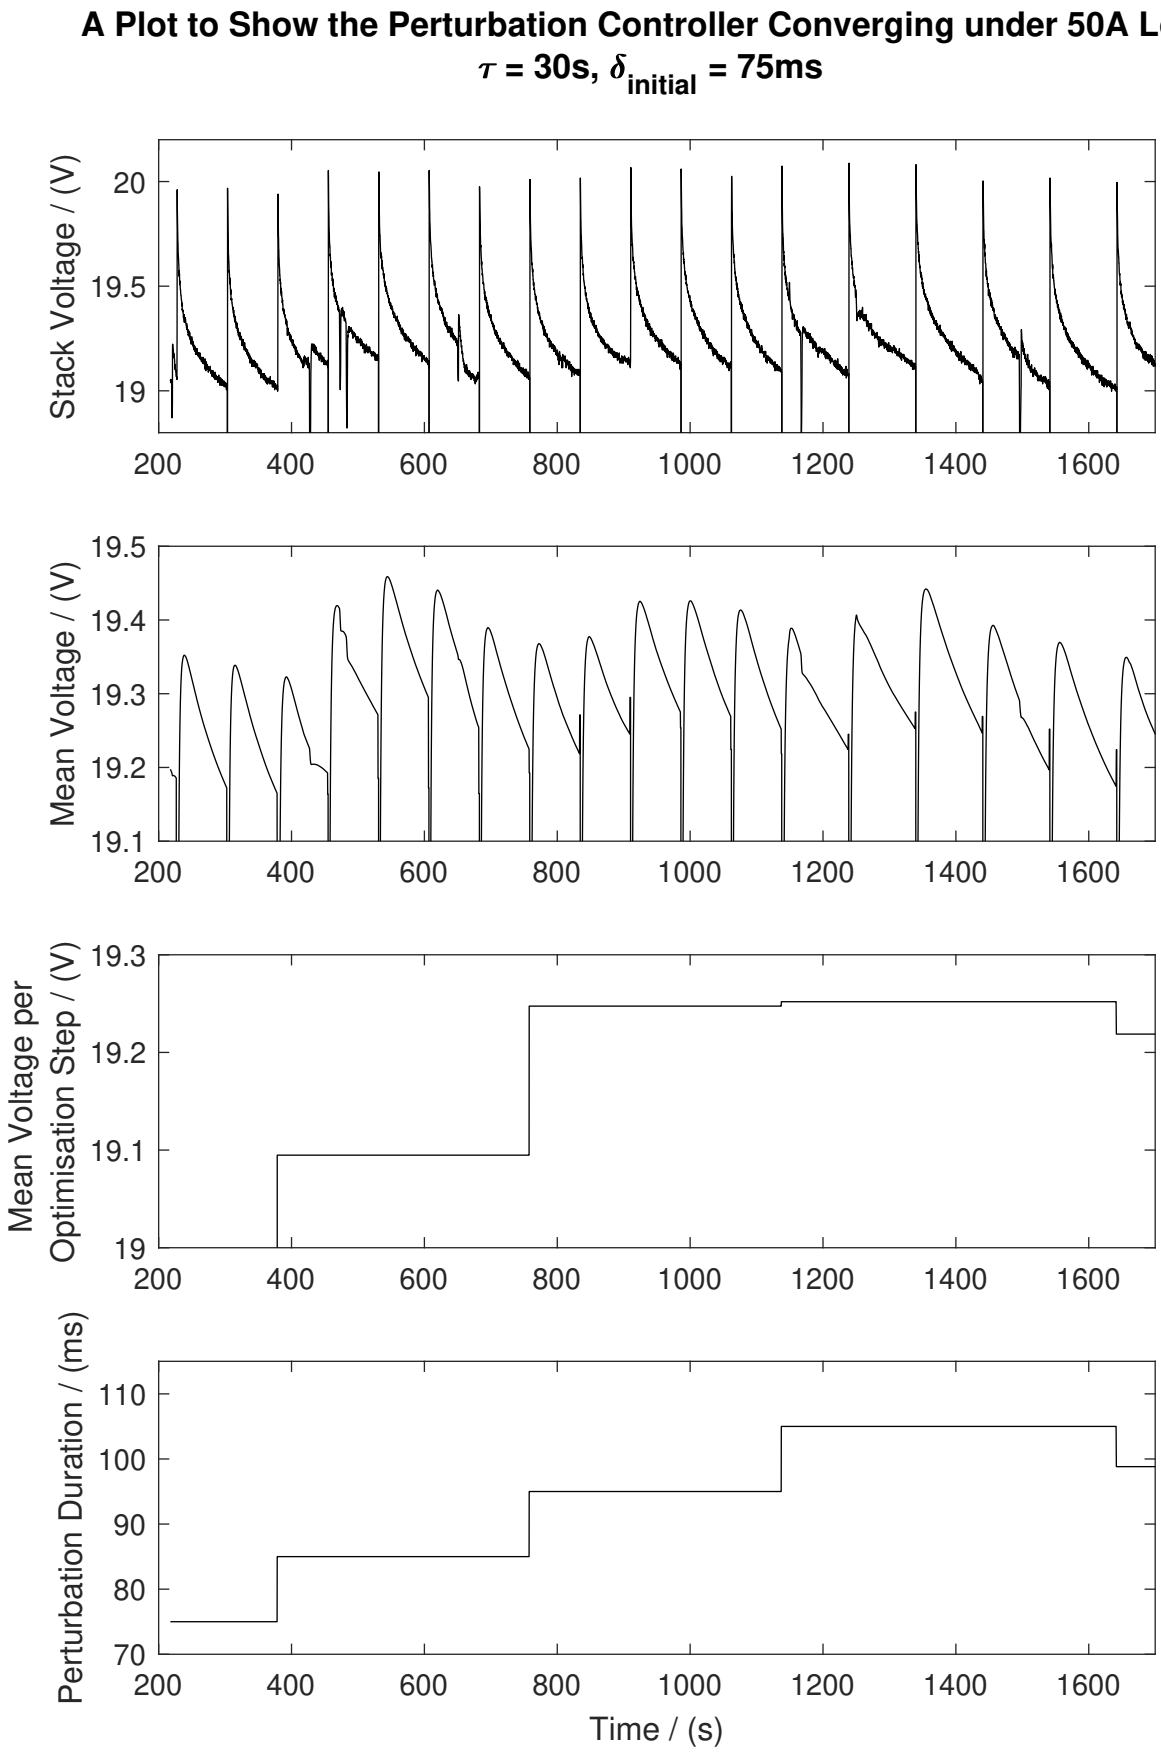


Figure 6.6: A plot showing the system optimising perturbation duration

behind the others by one step. Namely, the value plotted at time x is the value measured from the last step, and which was used to do the calculations that determine the value of d used at time x. The final axes in Fig 6.6 show the value of d calculated and applied by the system. An initial step size of 10 mS was used and starting point $\delta_{initial}$ was intentionally chosen to be an underestimate, as it was felt it demonstrated the system better than the optimum strategy, which would be a very large step size and a $\delta_{initial}$ greater than the anticipated optimum. In the case presented here, multiple steps are required to reach the optimum, which allows us to see how the sensitivity of the output changes as this optimum is approached. In addition, we can see the input parameter rise to its maximum in steps, as opposed to approach it in an uneven scatter as would happen if the step direction was constantly changing. Finally, the voltage plots are scaled to give better resolution of the voltage during the optimisation and do not include the steady state equilibrium voltage, which was 18.2±0.15 V. The wide variation in this quoted voltage is again due to variations in the operating conditions.

As can be seen in Fig 6.6 the mean voltage was increased to a value of 19.25V, representing an improvement in power and efficiency of 5.8% of the equilibrium value (i.e. with no current perturbation applied). As the system is running with a closed anode and is not being purged during the data collection periods, the fuel consumption is directly proportional to the current through the device, which is fixed. Efficiency gain is therefore directly proportional to the delivered voltage, as with power, as discussed in Section 5.2.4. This is lower than some of the values found in literature, the possible reasons for which will be evaluated in Section 6.3. The optimum value for d found was 95 ms, although this could have been further refined with a more stable system. The drop in cathode pressure did reduce the maximum voltage reached at this value, but only by a small amount. Given the purpose of this testing was primarily qualitative, it is not expected to have moved the optimum value by a significant amount (i.e. more than the step size used here). It took approximately 20 minutes to find this value, although this could be significantly reduced with a shorter period. The impact of these various factors, while somewhat frustrating from a data presentation perspective, serve to demonstrate how robust the system is to variations in operating conditions. The

author considers this a good demonstration of a reasonable worst case scenario, with very limited integration with the system managing the BOP and delivery of power to the load. As the system was easily able to improve the output by >150 mV (representing ≈15% of the total benefit gained by applying the perturbations), it is also felt that the optimisation is justified.

Perturbation Period

Having determined that the duration of a perturbation can be optimised successfully using P&O, the final thing to evaluate is the optimisation of the period of the perturbations. While this can be optimised with P&O, an analytical solution for the optimum value of τ was found in Section 5.3.4. This has two main advantages: firstly, it is much faster to find the optimum, requiring no convergence time. It instead calculates the optimum based on the instantaneous voltage readings of the stack, with no steps or tests required. The benefit of this brings us to the second advantage, which is that this allows both d and τ to be optimised simultaneously, using P&O for d and the analytical solution for τ. Fig 6.7 shows an example of this system optimising live, under the same operating conditions as the previous section. There were again some cathode pressure fluctuations due to the single stage regulator and hysteretic reservoir pressure based control of the compressors used, primarily in the 450-550s and 800-900s periods. Despite this, the system continued to converge, achieving a final mean voltage of 19.46V. Again, while it might produce a more easily digestible example were these not to occur their presence does demonstrate how robust this solution is to fluctuations in operating conditions, as well as the benefit of having a system that can adapt and optimise *in operando*. This is in contrast to a system with fixed parameter values regardless of operating conditions, load and degradation, as has previously existed.

Interestingly, despite the more rapid perturbations (i.e. shorter period), a longer duration was settled on. This could be for a number of reasons, but it is felt the most likely is small variations on reactant flow temperatures between the results of the previous section and those presented here. A much larger sample size would be required to thoroughly compare the two protocols, and as has been stated this is both beyond the scope of this project, as well as

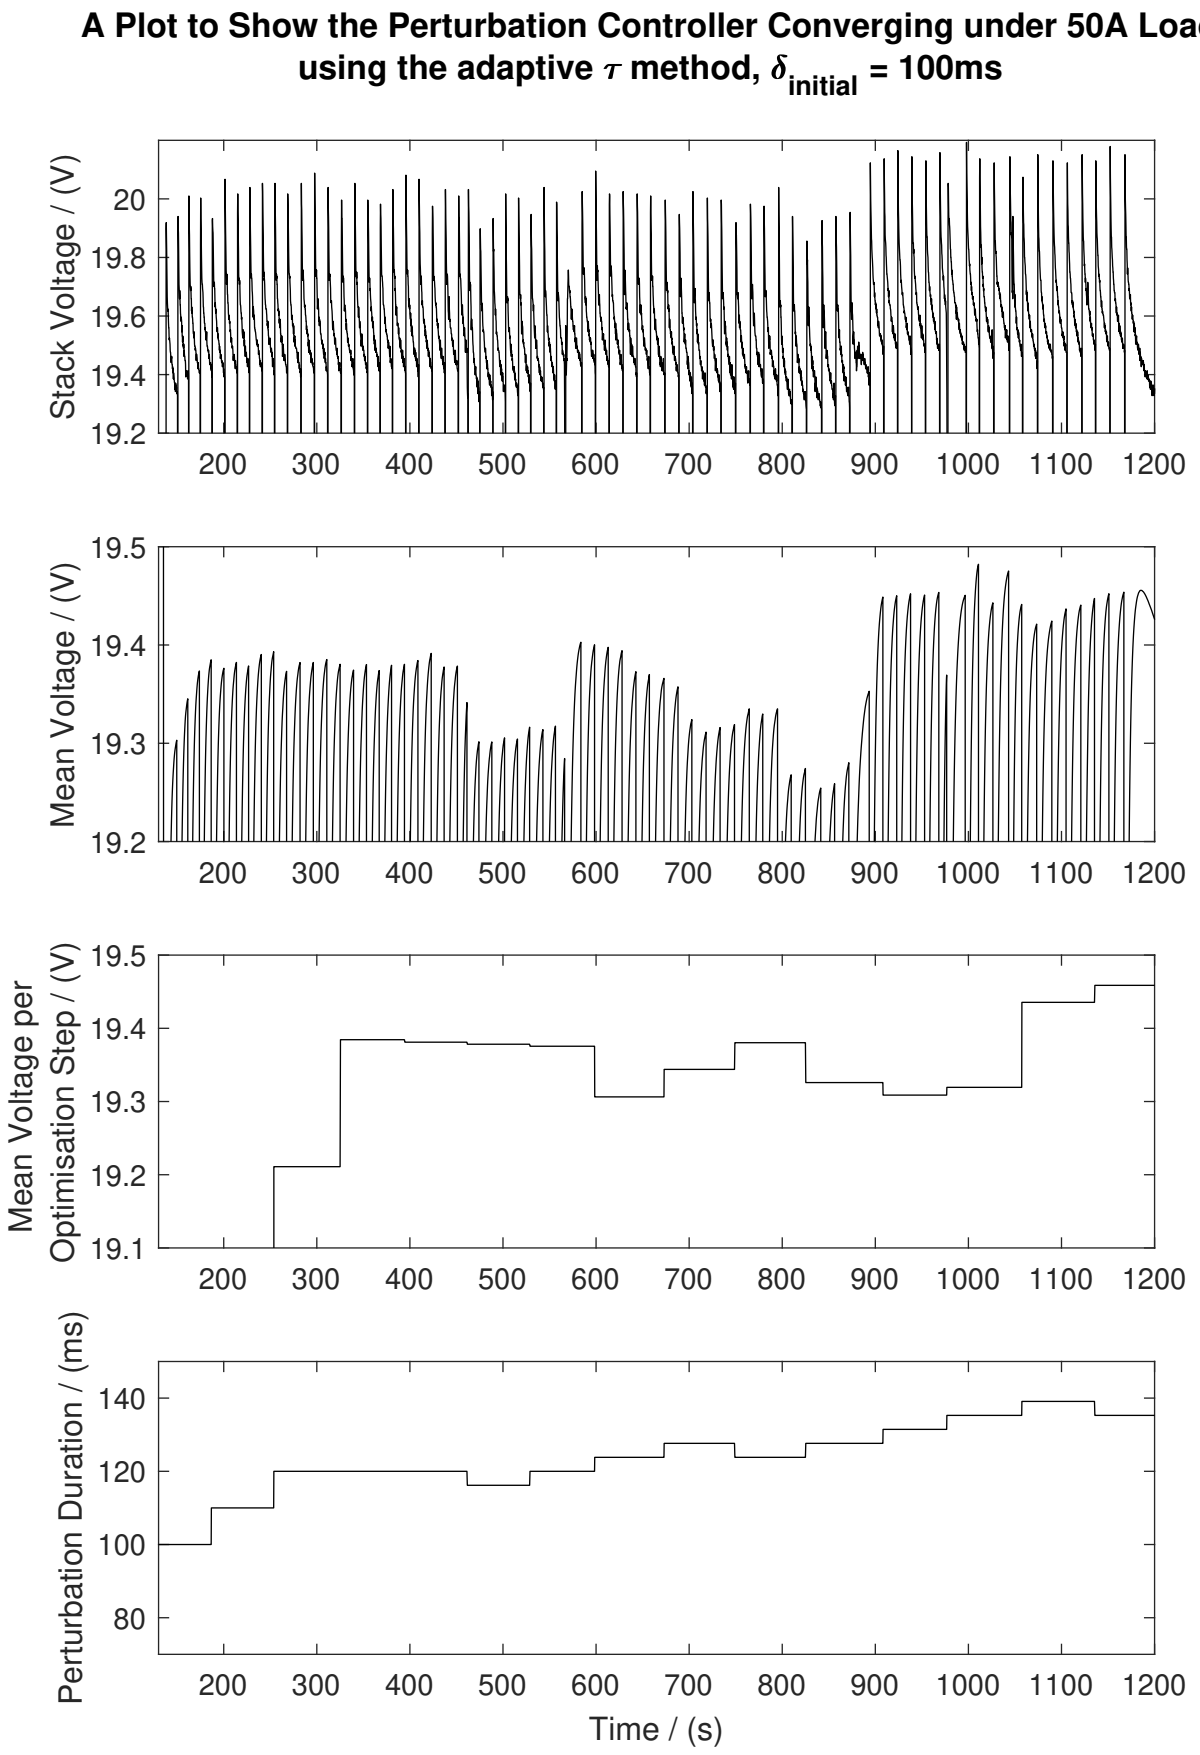


Figure 6.7: A plot showing the system optimising both τ and d, using the adaptive τ method and P&O respectively

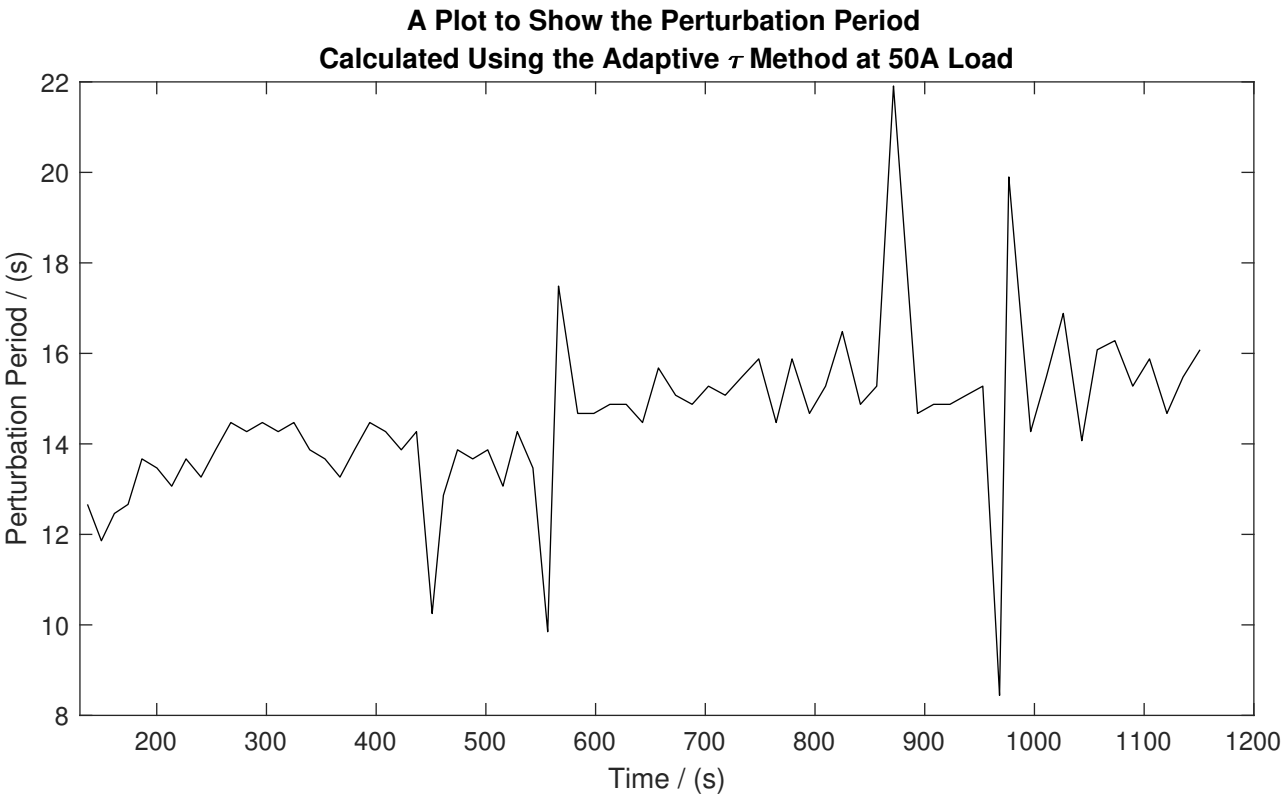


Figure 6.8: A plot showing the values of τ calculated by the adaptive τ method at 50A load

beyond the resources available. However, for the data available, the benefit of the adaptive τ technique is clear, as it was capable of optimising both parameters in less time than the P&O alone required to optimise just one. In addition, the benefit of optimising both τ and d can be seen, as this produced and increase of 0.2V, representing an additional 1.2 percentage points of additional power and efficiency, or increasing the benefit of applying current perturbations by another 15% compared to only optimising d.

One factor that was unexpected by the author, although clear in hindsight, was the optimum value found for τ. It was expected that this would be on the order of 30-100s based on the values reported in literature, but actual values calculated were in the 12-17s range for this system. A plot of τ during the optimisation presented above can be seen in Fig 6.8. When applying this method, the values calculated are primarily impacted by the gradient of the voltage decay following a perturbation. If the voltage never decayed, there would be no need to perturb, as the mean voltage would asymptotically approach the instantaneous value, but never reach it (in theory, in practice measurement noise would cause false triggers). This gradient is determined

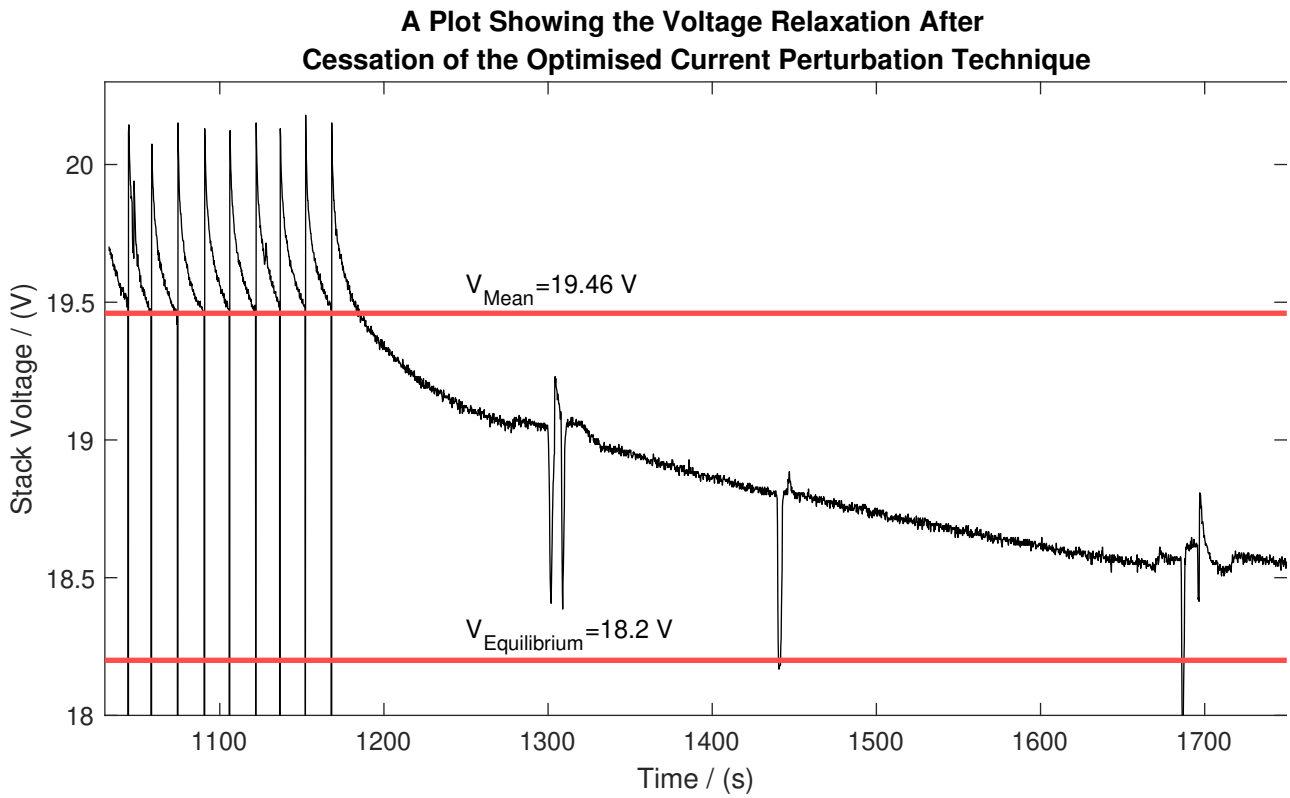


Figure 6.9: A plot showing the relaxation of the stack voltage following an optimisation using the adaptive τ method

by a wide range of parameters, including the design and manufacture of the device and its components, the applied load, the operating conditions etc. This variation is another example of the benefit of optimisation in general, as well as the advantage this proposed solution has in being agnostic to the system it is applied to. In a more stable system, it might be possible to use the value of τ calculated to optimise d, as the optimum delta should produce the largest voltage response immediately after the perturbation, and therefore the longest periods. This would be an elegant solution, but given the sensitivity of the adaptive method to other factors, may not be the most stable option.

6.3 Discussion

Prior to this work, there was no available literature to determine if current perturbations would work on modern automotive style PEMFC's. The available knowledge did imply that it should, and the favoured theory regarding the mechanism of current perturbations lends credence to

this, but electrochemical systems such as this are sufficiently complex that it could not be determined with certainty without some kind of experimental implementation of the technique. This implementation was produced and tested, albeit on a limited window of operating conditions, and the question does remain as to its benefit at higher current densities. As the technique acts to not only increase specific power but the efficiency of the device, it definitely has some applicability to systems that occasionally need to run with maximum efficiency as the goal, as opposed to maximum power. Good examples of this are range extenders on EV's with significant downtime, where the ESS needs to be charged by the PEMFC range extender over longer periods, or balancing of micro-grids to increase the utilisation of renewable sources. The benefit of the technique at the current densities tested here was lower than presented in the literature, but the current densities presented there are below the threshold that would make PEMFCs viable for many applications, whereas those used here are sufficiently high as to be of use for system designers, while maximising the efficiency of the device. The maximum increase in power and efficiency found here was 5.8% at approximately 0.5Acm^{-2}, compared to 10-12% reported by Gupta *et al* at 0.133 Acm^{-2}[50] (without humidified reactant flows and with an open cathode design). Zhang *et al* reported that their system produced the greatest benefit at the highest power,[51] although it is the opinion of this project that their measurement of the benefit uses a flawed technique in that they are measuring the peak voltage reached compared with the equilibrium voltage, as opposed to the mean voltage produced. A surface reading of this implies that at higher current densities the system tested here might yield increased benefit from application of current perturbations, however it is felt that more likely that the benefit is convex with respect to current density, increasing up to a point (the specific value of which is so dependent on system and device design that the author will not attempt to speculate), but then reducing in benefit as the operating point shifts along the Butler-Volmer curve. Indeed Gupta *et al* saw a largely uniform benefit across the current densities tested. Given that the modern style, water cooled closed cathode stacks are capable of reaching current densities an order of magnitude higher than those tested by Gupta *et al*, and the exponential shape of the Butler-Volmer relationship, it is likely that at these high current densities the intermediate reactions on which this technique acts are occurring at a rapid enough rate as to

gain reduced (or no) benefit from being perturbed. It is unfortunate that this project was not able to determine this, but it is hoped that the rig produced may be utilised to do so in the future.

6.3.1 Optimisation

The first research question set out for this section of the project is "Is it possible to dynamically optimise current perturbations for PEMFC applications, and what are the quantified benefits of the technique for modern automotive style PEMFC's?". Data presented in Section 6.2.5 demonstrates that the system designed for this purpose does allow the current perturbations to be optimised, and that with the proposed adaptive τ method this optimisation converges rapidly enough to be easily implemented on the vast majority of systems. Only a system where the PEMFC is expected to meet a rapidly fluctuating load demand with no hybridisation would it see diminished ability to converge, but such a system would either be for an application where the added efficiency is of minimal concern to the designer (such as a single use weapons platform), or would be poorly designed in the opinion of the author. In addition, the system tested here was shown to be robust, with no significant instability caused by fluctuations in operating conditions, required very little input (just the delivered voltage to the load, which would either already be measured by almost any implementation for system control and safety reasons, or could be retrofitted with ease) and is agnostic to the system performance (i.e. can be attached to any form of PEMFC and provide benefit). The only parts of the system that would need to be altered to suit a different implementation are the resistors used to step down the stack voltage to within the ADC's range (which is of negligible difficulty and cost) and potentially the switching of the current perturbation circuit itself. However, this implementation used a high power IGBT (and the associated hardware to drive such a device) to achieve this, which is realistically the most complex such a system would get, and is functionally identical to the hardware used to drive high voltage, high power rapid switching in systems such as EV motor inverters, DC-DC and DC-AC converters for connecting renewable energy sources to an electrical grid etc. While the components may need to be sized up in order to apply them to significantly larger systems (although they were intentionally oversized for safety reasons in this work), this

requires very little engineering effort and only a small added cost, which is negligible compared to the cost of the rest of the system. For example, the IGBT used here (easily the most expensive part of the current perturbation system, comprising >60% of the total cost) cost on the order of £150, which is less than the cost of the cabling required to connect the PEMFC to the load.

6.3.2 Future Implementations

Despite this, this implementation is not what the author would recommend to system designers/integrators. It was the most appropriate for the first scientific investigations applying this technique to modern devices, but the design criteria and resources of an academic study are vastly different to those of an industrial entity producing systems for use. Firstly, no consideration for system lifetime was given in this work, as it was outside the scope of the project, and it is felt that the implementation here may have a negative impact on the systems longevity. Trogadas *et al*[52] present data that demonstrates the technique negatively impacts device longevity, although the magnitude of the impact presented in that study is felt to be misleading. In order to accelerate degradation of the device and reduce the resources and cost required to perform the study they were operated at no current density (i.e. OCV). This would likely exaggerate the impact of the technique as it significantly increases the voltage cycling seen by the device it is applied to. This voltage cycling is a key factor in PEMFC degradation, as it significantly accelerates both the corrosion of the catalysts support structure[28] (typically carbon nano-fibre mat) as well as catalyst particle agglomeration. In addition, the device tested by this study utilised MEA's built in house, which are not fully representative of industrially produced components such as used in this project, especially with regards to their longevity and resistance to these specific degradation modes. Finally, the testing was performed for only a very short period (one week), and this project agrees with their conclusion that a significantly longer testing period would be required to determine the realistic impact of current perturbations on device longevity. This project does not agree with their conclusion that the technique should only be applied at the very end of a devices lifetime in order to extend it's usable life, as it is felt that the impact on the devices health is significantly less. Indeed (and it is

acknowledged that this is anecdotal evidence, and does not represent a rigorous scientific observation) this project applied current perturbations to a PEMFC for >50 hours over the course of the experiments (compared to Trogadas *et al*[52] 144 hours) and so no measurable impact on the systems performance, despite much of this consisting of perturbations with sub-optimal parameters that would accelerate degradation (such as significantly longer durations than would realistically be applied). In addition, during the purchasing negotiations with the manufacturer, the project was told in no-uncertain terms that it would require thousands of hours of operation (with current perturbations) at a minimum before any significant performance loss was exhibited.

Fig 6.10 demonstrates the impact of various operating conditions on PEMFC longevity. The most important of these causes with respect to current perturbations are voltage cycling, starvation and high temperature, with the former being by far the most significant in the implementation presented here. This is because this implementation utilised no active tuning of the impedance of the perturbation circuit. Doing so was beyond the scope of the project, and would have been premature without first determining if the technique functioned at all, as well as the response of the PEMFC stack during and immediately after the application of a perturbation. Some basic calculations based on the little data available in the literature were done, and it was found that a resistance on the order of 1-10mV would be desirable, which is approximately what the ohmic resistance of the circuit would be from the componentry and contact resistances of the connections. Starvation and high temperature are only of a concern in terms of localised effects, as the high stress and rapid fluctuations in both current and voltage applied by a perturbation are likely to exacerbate any inhomogineities in the device, disproportionately impacting specific areas. These areas could be areas of the flow field that are less capable of meeting the reactant flow rates required to recover following a perturbation, or areas where the ohmic resistance between opposing bipolar plates is significantly different from the bulk values. This would lead to localised starvation or hotspots in the MEA, and could therefore lead to any or all of the degradation mechanisms linked to these conditions. This is especially of concern with regard to those degradation mechanisms that lead to significant loss of performance to the system, even if they are only present in small fractions of the active

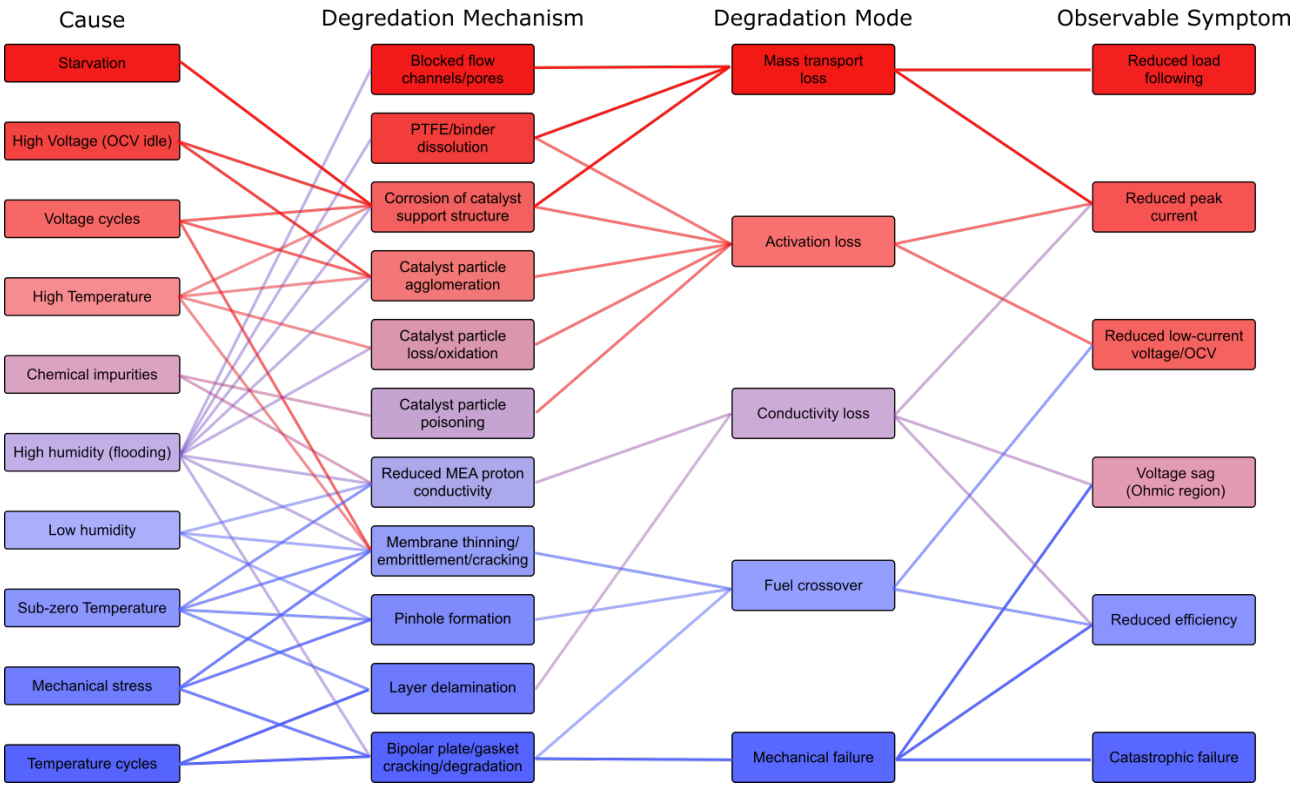


Figure 6.10: A diagram to show the impact of various operating conditions on PEMFC degradation, including the relationship between mechanisms and modes as well as the observable consequences.

area, such as pinhole formation or membrane embroilment, as even a single area suffering from these effects can lead to fuel crossover, poisoning large areas or even the entire cathode of the device. A more mature implementation could utilise a basic electrical filter to limit the impact of all of these mechanisms by reducing how aggressively the perturbation is applied. The effectively instantaneous step change in current applied here was not selected because it was the optimal implementation, but because it was the most easily achieved with the limited data available, and the most scientifically illuminating. In addition, the peak current reached is not necessarily required, provided sufficient charge is passed over a sufficiently short period. The peak current measured in these experiments was on the order of 1200A (a more precise measurement was not possible due to time constraints), but quickly fell to 400-600A for the majority of the perturbation. If the maximum current of the perturbation was limited to 200-300% of the maximum continuous current the PEMFC is rated for (in this case 240A), the potentially damaging effects of the technique would be significantly reduced.

While a perturbation circuit with more complex and careful consideration regarding the impedance would undoubtedly be beneficial compared to the rudimentary application of current used here, it is still not the recommended form of implementation in a mature system design. To understand why, we must first remember the form that mature implementations take. In order to allow the PEMFC to operate at the desired voltage, current or power regardless of the SOC/load or electrical characteristics of the load, they are invariably connected to the larger system via some form of power electronics, most typically a high power DC-DC converter. Given that such an electrical subsystem is already capable of switching at both high frequency and high loads, it would be far simpler to add an excitation signal to whatever operating point it was already imposing on the FC. This would require no additional hardware and only a modest increase in the complexity of the software. Such a system must already be measuring the voltage and current of the stack, as it is used to control these, and realistically would be more than capable of withstanding the very short periods of increased current, as the limit on such a systems capability is invariably thermal. As the magnitude of the perturbation can be controlled, it could be limited to whatever the extant system was capable of withstanding, and in the worst case scenario would only require a modest increase in size/capability of this subsystem. With this would also come the ability to apply the perturbation more intelligently, using voltage limits that prevent the device from crossing the relevant thresholds that accelerate degradation, and applying current signals that are more complex than can be achieved by a single switch, such as sinusoidal or sawtooth profiles. While the available data is woefully inadequate for the author to speculate as to what profile might prove optimal, it is felt that there is likely benefit in terms of device longevity to a more carefully design perturbation. Finally, such an implementation would negate the primary drawback of the technique as it is applied here, namely that the energy produced by the PEMFC during a perturbation is dissipated as oppose to captured, negatively impacting system efficiency. While this impact is minor, and certainly less than the benefit provided by the application of the technique, it is not a necessary loss for a mature and well engineered implementation in an industrial setting.

6.3.3 Mechanism theory

Turning to our second research question, regarding the behaviour of the PEMFC during and immediately after a perturbation and it's use for finding an analytical optimum, we shall split this into two sections. Firstly, we must define the excitation applied. In an ideal world, the excitation would have been a step change to a constant current, as would be possible with the more complex implementation discussed above. Alternatively, a high power potentiostat or battery cycler could have performed this task, but finding such a device that is capable of the currents required here is difficult. As mentioned previously, the peak current reached was on the order of 1200A, but the mean current during a perturbation was closer to the 400-600A range. This is still a very large current for instrumentation such as this, and the equipment available to do so is limited. The only equipment available to the project that could manage such a profile would be Maccor high power battery cyclers, which are limited to 200A per channel. These channels can be applied in parallel, but only in pairs, reaching a maximum current of 400A. More problematically the instruments capability of delivering precise pulses is notoriously bad, with a typical switching times being on the order of 100-150ms, longer than the majority of the durations required for this application, the voltage limits are too low (typically a maximum of 5-10V) and cannot be applied in series. "Hyperswitching" modules of the same equipment are available, but cannot be put in parallel. Pack or module cyclers are also available, but suffer from the same limitation in slew rate, and typically do not go to low enough voltages. The final option would be an electronic load similar to the Kikusui PLZ1205W used for the steady state load in these experiments, but it would require the purchase of around £10k worth of boosters to reach the required current.

As such, instead of a constant current load for the perturbation, it was effectively a constant resistance load. On the time scales we are observing there are undoubtedly some inductive effects in the impedance of this circuit which will impact the behaviour of the discharge, but attempting to meticulously measure these effects is not only very difficult, but far lower priority than other improvements to the rig. The only effort made to reduce these effects was good practice in the installation of these current paths. 75mm^2 cable was used (in order to withstand the largest continuous current the FC was rated to produce, 240A), and in order to both reduce

cable strain and limit the impact of the skin effect on rapid transients, cabling with a larger number of extra fine conductors was used. These were routed >100mm away from each other wherever possible (the primary exception being when passing through the safety enclosure) to limit any coupling they might experience.

Obviously, to conclusively prove one theory regarding the mechanism of current perturbation over the others would require a large body of work, presumably including some kind of in-situ imaging of both the species at the cathode tri-phasic boundary during an excitation, as well as an accurate measurement of the humidity profile through a cross section of the MEA. This is not only well beyond the scope and capabilities of this project, it is not readily achievable with any current metrology equipment or techniques known to the author. Instead, we must rely on the available data to lend credence one way or the other. The most compelling argument for the validity of the equilibrium theory over the humidification theory is that the voltage response reaches a maximum peak after a single perturbation on the order of 100ms. The excess water produced by such a perturbation is approximately 0.2g in the entire 24 cell stack, a clearly negligible amount when the scale of the device is considered. In addition, the time required for this water to permeate the membrane is not insignificant, as it relies on the diffusion gradient in the pores to overcome the osmotic drag of the protons moving from anode to cathode. In this testing, it was found that the maximum voltage of the cells was reached between 200-250ms after the cessation of the perturbation, too short a time for the bulk of any water produced to fully permeate the membrane. Iranzo *et al* performed EIS on a single PEM cell, and determined that the time constants for the electrochemical process in an air breathing setup to be between 1-2.5 ms, and for the diffusion processes (of reactants through the GDL, to the reaction sites on the catalyst surface) to be between 30-100 ms for $0.5 Acm^{-2}$.[39] It is gratifying to find that the time constants fitted to the voltage response during a perturbation in Section 6.2.2 fall comfortably into these windows, even if ascribing such physical meaning to such a fit is tenuous at best.

As for whether or not this data fits with the analytical model presented, it has been shown

that a very good fit can be easily achieved using the data following a perturbation. While the model is inaccurate for values of t<0.25s, this is both to be expected and of little to no concern when using it for the application it is intended, namely identifying the optimum value of τ. The reason it does not fit well in this initial region is that it makes no attempt to capture the high frequency transient behaviour of the FC, and instead assumes that this behaviour occurs on sufficiently small time constants as to be negligible when attempting to optimise τ, an assumption that is upheld by the data gathered by the CVM in its high sample rate acquisition mode. The fit of the model is not vital for the conclusions drawn from it to successfully optimise the system, but the fact that it does fit so readily to the recorded data gives confidence that the system behaves as expected, and therefore can be extrapolated accordingly to determine if a longer value of τ would yield a higher mean voltage. Doing so, we find that the adaptive τ method correctly identified the optimum value without fail (to within measurement error).

6.4 Conclusion

Current perturbations are an established technique for increasing the specific power and efficiency of PEMFC's, but are poorly represented in the available literature. The mechanism through which they act is not fully understood, with competing theories present and little proof or rigorous data lending credence one way or the other. The potential gains are significant, however, with the most representative literature boasting up to 12% increase in power and efficiency for no added fuel and very minimal added system complexity. No publication has ever applied the technique to modern, industrial PEMFC designs, such as those used in the automotive sector, and it's use tends to be a blind selection of parameters. What little data there is does suggest that the technique can be optimised, and that the two main parameters (duration of the perturbation and period in which it repeats) are convex.

In order to demonstrate that current perturbations do work on device representative of those used in industry (a 24 cell Powercell S2 stack), and that the parameters of the perturbation can

be optimised live and in situ a custom rig was produced. This rig both controls and instruments the operation of the PEMFC, as well as allowing the application of the current perturbation technique to nano-second accuracy, and optimisation of it's parameters *in operando*. To aid in this endeavour, a robust and well established optimisation algorithm was initially used, before being superseded where possible with an analytical solution to decrease convergence time. This application was demonstrated to be robust to changes in operating conditions and yield a quantified benefit. It is posited to be agnostic to the system to which it is applied, although demonstration of this would require application to multiple systems, and was beyond the resources of this project. The analytical solution found is derived from a simple model based on the most credible theory regarding the mechanism through which the technique acts, but the solution found remains valid regardless of the validity of the model.

As well as demonstrating an implementation of the technique, custom metrology electronics were designed and manufactured to allow a more detailed insight into the behaviour of the PEMFC during and immediately after a perturbation, measuring every cell at a high sampling frequency (>7kHz). To the knowledge of the author, this kind of data has never been published, and is strong evidence to the validity of the equilibrium based theory presented in this work. It also allows validation of the model used to produce the analytical solution for the optimum perturbation period. Finally, a discussion of future implementations is presented, positing that for most system minimal additional hardware (and therefore system complexity) would be needed, if any. While the impact of the technique on the device longevity is not considered in the body of this work, the more mature implementations discussed may have the potential to limit any negative effects by controlling the application of the perturbations more accurately, such as keeping the cells above voltages known to accelerate irreversible degradation, or limiting peak currents to prevent formation of hot-spots or other inhomogineities.

If this technique were implemented in the design of new hybrid FC systems such as those used in emergent automotive designs it has the potential to significantly reduce the size, cost and weight of the FC system required to meet a given load. As mentioned, this will require further

work to develop the implementation, but requires very little in terms of additional complexity or hardware. Indeed, the net weight and size of the system should be reduced, while also improving system efficiency, and therefore range on a given amount of fuel. By understanding and manipulating the phenomena occuring at the catalyst level, we can take advantage of their behaviour to improve the system level performance.

Chapter 7

Conclusion

Two areas of study were identified for investigation based on their relevance to the application in question, the design and implementation of HFCEV's. These were the effect of thermal gradients on the degradation of LIBs, and the use of current perturbations on PEMFC's for increasing the specific power and efficiency of the FC stack. The former has been studied a great deal, but work has been limited by the inability to decouple the C-rate and temperature effects due to the multiple layers of active material present in commercial LIBs, as well as the non-isotropic thermal conductivity of this assembly. In order to overcome this, custom pouch cells were commissioned, consisting of one symmetrical set of positive and negative electrodes. This allows there to be no significant thermal gradient across the active material in the cell, and therefore the decoupling of temperature and C-rate in a degradation study.

The latter area of study, current perturbations in PEMFC's, is by contrast much neglected in the available literature. There are very few publications investigating the technique, and fewer still attempting to refine its use and implementation. This is surprising, as the technique is far from new, being the subject of a patent in 1997,[181] and has been shown to significantly increase both specific power and efficiency (up to 12% in the right conditions[50]). In addition, anecdotal evidence suggests that it is well known about in commercial use, but the mechanism through which it acts is poorly understood. Because of this, implementations tend to be rudimentary, with parameters chosen essentially blindly, and no optimisation performed, either at BOL or

throughout use.

7.1 Summary of Thesis Achievements

A custom test rig was designed to hold the single layer cells at constant temperature, as well as control the pressure applied to their surfaces. The performance of this fixture was validated by applying a worst case scenario load, consisting of 4C continuous charge/discharge on already degraded cells (and therefore highest resistance and heat generation). It was found that even with a stack of these cells, the greatest difference in the mean temperature between the cell surface temperature was 0.81°C (compared to the same setup with no current passing). While this was greater than initially expected, it is still less than the variation caused by the thermal chambers control scheme, and deemed a small enough change to be considered isothermal for the purposes of the testing carried out here.

The single layer cells were then cycled, with the discharge rate and temperature during cycling varied from cell to cell in order to investigate their relative impacts on the degradation of the cells. In addition, multi-layer pouch cells of the same footprint, but approximately 24 times the capacity due to a more typical number of layers of active material were mounted in a modified version of the test rig, designed to implement active control of the thermal boundary conditions using water-cooled peltier elements. The cells were characterised periodically, and the results from these characterisations were analysed using a variety of techniques. It was found that the primary degradation mode was initially LAM at the NE (the capacity of the PE in the single layer cells was approximately double that of the NE due to manufacturing constraints, reducing the rate at which the PE lost capacity, as well as preventing htis loss from impacting the capacity of the cell as a whole). However, partway through the testing, LLI became dominant and the LAM plateaued. While it was always found to plateau at the same value to within experimental error (20% of the initial capacity), the point in the cells life at which this occurred was strongly dependent on temperature. In addition, the impact of discharge rate on the rate at which the cells lost capacity was found to be strongly non-linear,

being minimal up until a threshold value, and having an exaggerated impact at discharge rates higher than this threshold. The specific value was between 2-3C at 30°C for these cells, but was also strongly impacted by temperature, with higher temperature being significantly more rate capable without suffering as extreme increases in the rate of capacity loss.

The results were somewhat obfuscated by some unfortunate real world complications, primarily the eruption of the Covid-19 global pandemic interrupting the testing schedule. When testing was resumed, it was found that the remaining cells had significantly calendar aged, likely due to improper storage in the interim. As such, it was not possible to discern some of the finer details regarding the decoupled impact of discharge rate and temperature, but careful analysis did allow the conclusions enumerated above to be found with reasonable certainty. This also impacted all of the multi-layer pouch control cells tested, limiting the comparison between these cells and their single layer counterparts. However, the results from these cells did generally support the conclusions drawn from the single layer cells, as well as demonstrating the benefit of such an approach for a study of this nature. While the cell to cell variation at the beginning of testing was lower in the multi-layer cells compared to the single layers, the spread due to degradation effects was significantly higher, and the impact of the two variables under investigation was much more difficult to differentiate. This is to be expected, and is the reason for the commissioning of the custom cells.

Despite the limited findings from the gathered data, the primary intention of the work was achieved. The method for decoupling the impacts of C-rate and temperature on the degradation of LIBs was produced and validated. It was shown that not only can the impacts of C-rate and temperature be decoupled at the device level (albeit with custom devices), but the differences in their degradation pathways was quantitatively demonstrated. While the quantitative findings of this study are limited primarily to the chemistry used, the broader implications are valuable, and the roadmap for much more targeted studies on other chemistries has been established.

For the investigation of current perturbations, similarly a custom test rig was produced. Specifically, the investigation of the techniques application to modern automotive style PEMFC stack designs (closed cathode, stamped metal bipolar plates, water cooled, capable of cur-

rent densities >1.2Acm^{-2}), as well as the live optimisation if the perturbations parameters *in operando*. Data regarding the use of this technique on such a category of devices is not publicly available prior to this work to the best of the authors knowledge. In addition to this, a cell voltage monitor capable of sampling at a much higher frequency than those available commercially was produced in order to investigate the voltage response of the device both during and immediately after a perturbation. This was done to better help identify the mechanism through which the technique acts, as there are competing theories in the literature. This was capable of taking samples from all 24 cells simultaneously at >7kHz, and lent strong credence to the theory that the perturbations disrupt the equilibria present in the multi-step ORR taking place at the system cathode. By periodically perturbing these equilibria, and not allowing them to fully relax to their steady state values, they can be artificially pushed in the desired direction, reducing the overpotential required to achieve a set current density.

Finally, a system for the live optimisation of the perturbations key parameters (it's duration and the period at which it is repeated) was produced. An initial implementation utilising a 2DoF perturb and observe algorithm was demonstrated, as this is a well understood, robust and popular optimisation strategy. Based on the favoured theory regarding the mechanism of the technique, an analytical solution for the optimum period of the perturbations was derived, one which requires very little input data (only continuous monitoring of the stack voltage, a parameter that is always measured in any mature system implementation) and is agnostic to both the system it is applied to, as well as the accuracy of the model used in its derivation. The model was shown to be very accurate for timescales >0.5s, and the analytical solution was shown to greatly increase the speed at which the system converges on the optimum solution. This was partly due to the analytical solutions ability to find the optimum period for every individual perturbation, with no trial and error required, as well as the fact it allowed both parameters to be optimised simultaneously. In addition, it was shown that in the region of the optimum values, the magnitude of the benefit gained is significantly less sensitive to changes in the input parameters.

While this testing was not able to demonstrate the breadth of data that was initially desired, specifically regarding the application of the technique at the highest current densities achievable with the device in question and the impact of perturbation duration on the peak voltage reached immediately following the event, it is felt that the data presented here is both unique in the available literature, as well as sufficient to answer the research questions identified by the project.

It is hoped that the procedures and equipment produced by this project may be as useful as the data gathered with them. The fixture for the testing of single layer cells has been shared with the wider community already, and is believed to be useful for the decoupling of various parameters of the behaviour of LIB's, including temperature, load profile and magnitude, as well as pressure. In addition, it can allow a more rigorous comparison of cell chemistries, both of the electrodes and the electrolyte, as it can both capture the complexities of a full cell (compared to half cells tested in coin setups), but also tightly control parameters whose values it is desirable to fix. The PEMFC rig produced is of more use as an example tool, as opposed to for dissemination, but the live optimisation of the parameters of the current perturbation technique is hoped to prove useful to system integrators and manufacturers.

7.2 Future Work

The implementation of the current perturbation presented here is both somewhat rudimentary, while also being a significant advancement in the use as it is represented in publicly available literature. A brief discussion of more developed implementations is presented in Section 6.3.2. Quantifying the benefit of these implementations would be of great interest, but was unfortunately beyond both the scope and resources of this project. More detailed investigation of the mechanism through which it acts would also be of great interest, the first steps of which would be an investigation of the impact of perturbation duration on the peak voltage response, as well as the relationship between the mean voltage/power increase with respect to operating current density. Finally, the benefit of the technique, and specifically of the live, *in operando* optimisation of the technique throughout the life of the PEMFC, as well as it's impact on the

longevity of the device would be of great interest, as well as its performance on a variety of PEMFC designs, both mechanical and with respect to catalyst and electrolytes used.

The results found here do demonstrate the large potential gains from this technique. A well designed implementation of this technique would likely be able to significantly minimise any negative impacts on device longevity, and could have just as significant gains in the performance of the system. This would allow reductions in the size of system required for given applications. Given that the gains found in this rudimentary implementation at lower-than-desired current densities found an improvement of nearly 8% in specific power, and what little literature available found improvements of up to 12% (albeit on out of date hardware), it is possible that a mature implementation could allow a system that is similarly 10-15% smaller for a given application. This would require the benefit of the technique to increase with both more refinement and increasing current density, both of which are to be determined. However, given our limited understanding of the mechanism involved, both of these are comfortably in the realms of possibility. A 10-15% increase in power density is a significant change and, while not a paradigm shifting amount, is certainly a large enough performance increase that it might tip the scales in determining if PEMFC's are a viable solution for a given application. For example, the cost of the FC system is a significant factor in determining the economic viability of HFCEV's for consumer automotive products. While this will primarily be reduced by better automation, scaling and other technological improvements in manufacturing, being able to scale down the system size due to such an increase in its performance would be a notable contributing factor in that techno-economic evaluation.

For the LIB work presented here, future work would be the repetition of the degradation study without the impact of a global pandemic on the test results, as well as some refinement of the procedure based on the lessons learned. Specifically, single layer cells without the extreme overhang of the PE, as well as testing the cells to failure, with a post mortem analysis, and even ideally disassembly of cells partway through the degradation process. Of course, this requires a significantly larger sample size, as cells must be removed from the sample pool periodically for disassembly. This therefore requires a considerable commitment of resources. Work will

continue based on the output of this project, as the rig and experience have been passed on to the Faraday Institutes Multi-Scale Modelling consortium, who are using it to test a variety of variables under a standardised set of conditions.

Ultimately, because of the other factors, this body of work did not succeed in finding the limitations of the experimental rig from a hardware standpoint, but rather from an operational one. While it would be nice to assume that this indicates that there are no shortcomings of the mechanical design (or rather, the shortcomings are so minor as to be inconsequential), that is an obviously spurious conclusion. While there are some ways the mechanical fixture could be changed to address the operational difficulties found here, such as greater instrumentation to help identify sources of anomalous readings or behaviour or quicker swapping of cell units, these would require significant increases in the costs associated. It is the opinion of the author that while attempting to improve the mechanical design of the rig is not the most pressing area of further work based on this project, the optimal path for such work would be to focus on the strengths of the design. For example, while it is relatively cheap to manufacture at the small to medium scale currently (i.e. 10s to 100s of units), this could be even more significantly reduced through use of mass manufacture (for example, injection moulding the tab plates instead of machining them) or reduction in consumable costs. The cost of the thermal interface used in this body of work was a significant portion of the rigs cost, and could be greatly reduced by validation of specific conditions using lower conductivity (and therefore cheaper) materials, re-useable interface materials, or a re-design allowing other cheaper solutions. For example, the thermal interface could be achieved with a paste, as opposed to a pad material, which tend to be much cheaper. This would require better surface finish on the compression plates, which would be possible at higher production volumes. The design presented here uses the mill finish of cold rolled stock, as the added expense of precision flat surfaces was determined to yield no significant improvement for a large added expense, but were the plates to be produced by stamping from sheet material at larger volumes (as opposed to individually milled in a prototyping environment), a much finer surface finish would not add such cost. Additionally, a refined method for applying the paste in an even and repeatable coating, such as is used for the

application of solder paste to mass manufactured PCB's, would make this solution more viable. On a large enough scale, even the development of a re-useable thermal interface bonded to the surface of the plates could be achieved for minimal cost, and could significantly improve the cost-per-use of the hardware, not to mention remove the tedious operational task of carefully preparing and applying delicate thermal interface material to every mating surface in a large stack of components.

Bibliography

[1] L. Wei, Z. Lu, F. Cao, L. Zhang, X. Yang, X. Yu, and L. Jin, "A comprehensive study on thermal conductivity of the lithium-ion battery," *International Journal of Energy Research*, vol. 44, no. 12, pp. 9466–9478, 2020.

[2] N. Nitta, F. Wu, J. T. Lee, and G. Yushin, "Li-ion battery materials: present and future," *Materials Today*, vol. 18, no. 5, pp. 252–264, 2015.

[3] M. H. Hossain, M. A. Chowdhury, N. Hossain, M. A. Islam, and M. H. Mobarak, "Advances of lithium-ion batteries anode materials—a review," *Chemical Engineering Journal Advances*, vol. 16, p. 100569, 2023.

[4] R. E. Silva, F. Harel, S. Jemeï, R. Gouriveau, D. Hissel, L. Boulon, and K. Agbossou, "Proton Exchange Membrane Fuel Cell Operation and Degradation in Short-Circuit," *Fuel Cells*, vol. 14, pp. 894–905, dec 2014.

[5] X. Feng, Y. Merla, C. Weng, M. Ouyang, X. He, B. Y. Liaw, S. Santhanagopalan, X. Li, P. Liu, L. Lu, X. Han, D. Ren, Y. Wang, R. Li, C. Jin, P. Huang, M. Yi, L. Wang, Y. Zhao, Y. Patel, and G. Offer, "A reliable approach of differentiating discrete sampled-data for battery diagnosis," *eTransportation*, vol. 3, p. 100051, 2020.

[6] K. Xu, "Nonaqueous liquid electrolytes for lithium-based rechargeable batteries," *Chemical Reviews*, vol. 104, no. 10, pp. 4303–4418, 2004. PMID: 15669157.

[7] M. Dubarry and D. Anseán, "Best practices for incremental capacity analysis," *Frontiers in Energy Research*, vol. 10, 2022.

[8] 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. Bravo 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," *Phys. Chem. Chem. Phys.*, vol. 23, pp. 8200–8221, 2021.

[9] M. Ji and Z. Wei, "A review of water management in polymer electrolyte membrane fuel cells," *Energies 2009, Vol. 2, Pages 1057-1106*, vol. 2, pp. 1057–1106, 11 2009.

[10] K. O. Yoro and M. O. Daramola, "Chapter 1 - co2 emission sources, greenhouse gases, and the global warming effect," in *Advances in Carbon Capture* (M. R. Rahimpour, M. Farsi, and M. A. Makarem, eds.), pp. 3–28, Woodhead Publishing, 2020.

[11] R. Ahluwalia, X. Wang, and A. Rousseau, "Fuel economy of hybrid fuel-cell vehicles," *JOURNAL OF POWER SOURCES*, vol. 152, pp. 233–244, DEC 1 2005. Fuel Cell Seminar 2004, San Antonio, TX, 2004.

[12] C. C. Chan, "The state of the art of electric, hybrid, and fuel cell vehicles," *Proceedings of the IEEE*, vol. 95, no. 4, pp. 704–718, 2007.

[13] C. C. Chan, A. Bouscayrol, and K. Chen, "Electric, hybrid, and fuel-cell vehicles: Architectures and modeling," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 2, pp. 589–598, 2010.

[14] G. Offer, D. Howey, M. Contestabile, R. Clague, and N. Brandon, "Comparative analysis of battery electric, hydrogen fuel cell and hybrid vehicles in a future sustainable road transport system," *Energy Policy*, vol. 38, no. 1, pp. 24–29, 2010.

[15] K. S. Jeong and B. S. Oh, "Fuel economy and life-cycle cost analysis of a fuel cell hybrid vehicle," *Journal of Power Sources*, vol. 105, no. 1, pp. 58–65, 2002.

[16] S. Vengatesan, A. Jayakumar, and K. K. Sadasivuni, "Fcev vs. bev — a short overview on identifying the key contributors to affordable clean energy (sdg-7)," *Energy Strategy Reviews*, vol. 53, p. 101380, 2024.

[17] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 8, pp. 261–273. Wiley, 2009.

[18] O. Z. Sharaf and M. F. Orhan, "An overview of fuel cell technology: Fundamentals and applications," *Renewable and Sustainable Energy Reviews*, vol. 32, pp. 810–853, apr 2014.

[19] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 1, pp. 12–14. Wiley, 2009.

[20] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 8, pp. 275–7. Wiley, 2009.

[21] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 8, pp. 278–80. Wiley, 2009.

[22] M. Li, *Electrocatalysts for Ethanol Oxidation in Direct Ethanol Fuel Cell (DEFC)*. PhD thesis, Stony Brook University, 2011.

[23] P. Boldrin and N. P. Brandon, "Progress and outlook for solid oxide fuel cells for transportation applications," *Nature Catalysis 2019 2:7*, vol. 2, pp. 571–577, 7 2019.

[24] R. T. Leah, A. Bone, E. Hammer, A. Selcuk, M. Rahman, A. Clare, L. Rees, N. Lawrence, A. Ballard, T. Domanski, S. Mukerjee, and M. Selby, "Development of high efficiency steel cell technology for multiple applications," *ECS Transactions*, vol. 78, p. 2005, may 2017.

[25] R. Leah, A. Bone, M. Lankin, A. Selcuk, R. Pierce, L. Rees, D. Corcoran, P. Muhl, Z. Dehaney-Steven, C. Brackenbury, M. Selby, and S. Mukerjee, "Low-cost, redox-stable, low-temperature sofc developed by ceres power for multiple applications: Latest development update," *ECS Transactions*, vol. 57, p. 461, oct 2013.

[26] R. T. Leah, A. Bone, M. Lankin, A. Selcuk, M. Rahman, A. Clare, L. Rees, S. Phillip, S. Mukerjee, and M. Selby, "Ceres power steel cell technology: Rapid progress towards a truly commercially viable sofc," *ECS Transactions*, vol. 68, p. 95, jun 2015.

[27] W. Schmittinger and A. Vahidi, "A review of the main parameters influencing long-term performance and durability of PEM fuel cells," *Journal of Power Sources*, vol. 180, pp. 1–14, may 2008.

[28] F. A. de Bruijn, V. A. T. Dam, and G. J. M. Janssen, "Review: Durability and Degradation Issues of PEM Fuel Cell Components," *Fuel Cells*, vol. 8, pp. 3–22, feb 2008.

[29] C. of the Desert, "Hydrogen properties."

[30] I. Energy, "Useful conversion factors."

[31] J. Kurtz, S. Sprik, and T. H. Bradley, "Review of transportation hydrogen infrastructure performance and reliability," *International Journal of Hydrogen Energy*, vol. 44, no. 23, pp. 12010–12023, 2019.

[32] Envestra, "Common properties of commercial fuels."

[33] E. West, "Product listing - tesla model s lithium ion battery 18650 ev module."

[34] P. Bénard and R. Chahine, "Storage of hydrogen by physisorption on carbon and nanostructured materials," *Scripta Materialia*, vol. 56, no. 10, pp. 803 – 808, 2007. Viewpoint set no. 42 "Nanoscale materials for hydrogen storage".

[35] J. Wishart, "Fuel cells vs batteries in the automotive sector," 08 2014.

[36] "The business impact of off-route fleet refueling — geotab."

[37] M. Ben-Akiva, H. Meersman, E. Van de Voorde, M. A. Abate, and O. Kveiborg, *Capacity Utilisation of Vehicles for Road Freight Transport*. United Kingdom :: Emerald,, 2013-05-14.

[38] D. Wu and S. Williamson, "Status review of power control strategies for fuel cell based hybrid electric vehicles," *IEEE Canada Electrical Power Conference*, vol. 2, pp. 55–61, 2007. 2 basic control strategies presented, compared based on efficiency and performance. Could be interesting to compare based on degredation of components.

[39] A. Iranzo, S. J. Navas, F. Rosa, and M. R. Berber, "Determination of time constants of diffusion and electrochemical processes in polymer electrolyte membrane fuel cells," *Energy*, vol. 221, p. 119833, 4 2021.

[40] M. Huang, P. Wen, Z. Zhang, B. Wang, W. Mao, J. Deng, and H. Ni, "Research on hybrid ratio of fuel cell hybrid vehicle based on advisor," *International Journal of Hydrogen Energy*, vol. 41, no. 36, pp. 16282–16286, 2016. Special Issue: 16th China Hydrogen Energy Conference (CHEC 2015), November 2015, Zhenjiang City, Jiangsu Province, China.

[41] 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, 12 2011.

[42] J. Ogle, R. Guensler, and V. Elango, "Georgia's commute atlanta value pricing program," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1931, pp. 28–37, 1 2005.

[43] S. Schönfelder, H. Li, R. Guensler, J. Ogle, and K. W. Axhausen, "Analysis of commute atlanta instrumented vehicle gps data," *Arbeitsberichte Verkehrs- und Raumplanung*, vol. 350, 2006.

[44] J. Wolf, R. Guensler, S. Washington, W. Sarasua, C. Grant, S. Hallmark, M. Simas Oliveira, M. Koutsak, R. Thittai, R. Funk, and J. Hsu, "Development of a comprehensive vehicle instrumentation package for monitoring individual tripmaking behavior: Final report," 04 1999.

[45] D. D. Battista, C. Villante, R. Cipollone, D. D. Battista, C. Villante, and R. Cipollone, "Optimal components design of a fuel cell electric vehicle," *SAE Technical Papers*, vol. 2015-September, 9 2015.

[46] C. Cosse, M. Schumann, F. Grumm, D. Becker, and D. Schulz, "Numerical investigation of pemfc short-circuit behaviour using an agglomerate model approach," *Energies*, vol. 13, no. 16, 2020.

[47] J. Kim, D.-M. Kim, S.-Y. Kim, S. W. Nam, and T. Kim, "Humidification of polymer electrolyte membrane fuel cell using short circuit control for unmanned aerial vehicle applications," *International Journal of Hydrogen Energy*, vol. 39, pp. 7925–7930, may 2014.

[48] J. A. Flach and F. A. Farret, "Circuit for monitoring and reestablishment of pem fuel cell voltage using controlled short circuits," *2019 IEEE PES Conference on Innovative Smart Grid Technologies, ISGT Latin America 2019*, 9 2019.

[49] Y. Zhan, Y. Guo, J. Zhu, and L. Li, "Current short circuit implementation for performance improvement and lifetime extension of proton exchange membrane fuel cell," *Journal of Power Sources*, vol. 270, pp. 183–192, 12 2014.

[50] G. Gupta, B. Wu, S. Mylius, and G. J. Offer, "A systematic study on the use of short circuiting for the improvement of proton exchange membrane fuel cell performance," *International Journal of Hydrogen Energy*, vol. 42, pp. 4320–4327, feb 2017.

[51] C. Zhang, H. Liu, T. Zeng, J. Chen, P. Lin, B. Deng, F. Liu, and Y. Zheng, "Systematic study of short circuit activation on the performance of pem fuel cell," *International Journal of Hydrogen Energy*, vol. 46, pp. 23489–23497, 7 2021.

[52] P. Trogadas, J. I. Cho, N. Kapil, L. Rasha, A. Corredera, D. J. Brett, and M. O. Coppens, "Effect of extended short-circuiting in proton exchange membrane fuel cells," *Sustainable Energy & Fuels*, vol. 4, pp. 5739–5746, 10 2020.

[53] S. Ma, M. Lin, T.-E. Lin, T. Lan, X. Liao, F. Maréchal, J. Van herle, Y. Yang, C. Dong, and L. Wang, "Fuel cell-battery hybrid systems for mobility and off-grid applications: A review," *Renewable and Sustainable Energy Reviews*, vol. 135, p. 110119, 2021.

[54] K. Field, "Toyota mirai hydrogen fuel cell vehicle review," 2018. Accessed on 01/06/2024.

[55] V. K. Kasimalla, N. S. G., and V. Velisala, "A review on energy allocation of fuel cell/battery/ultracapacitor for hybrid electric vehicles," *International Journal of Energy Research*, vol. 42, no. 14, pp. 4263–4283, 2018.

[56] A. Khaligh and Z. Li, "Battery, ultracapacitor, fuel cell, and hybrid energy storage systems for electric, hybrid electric, fuel cell, and plug-in hybrid electric vehicles: State of the art," *IEEE Transactions on Vehicular Technology*, vol. 59, pp. 2806–2814, 2010.

[57] P. Bubna, S. G. Advani, and A. K. Prasad, "Integration of batteries with ultracapacitors for a fuel cell hybrid transit bus," *Journal of Power Sources*, vol. 199, pp. 360–366, 2 2012.

[58] T. M. Bandhauer, S. Garimella, and T. F. Fuller, "A critical review of thermal issues in lithium-ion batteries," *Journal of The Electrochemical Society*, vol. 158, p. R1, jan 2011.

[59] T. Waldmann, M. Wilka, M. Kasper, M. Fleischhammer, and M. Wohlfahrt-Mehrens, "Temperature dependent ageing mechanisms in lithium-ion batteries – a post-mortem study," *Journal of Power Sources*, vol. 262, pp. 129–135, 2014.

[60] J. R. Belt, C. D. Ho, T. J. Miller, M. A. Habib, and T. Q. Duong, "The effect of temperature on capacity and power in cycled lithium ion batteries," *Journal of Power Sources*, vol. 142, no. 1, pp. 354–360, 2005.

[61] I. A. Hunt, Y. Zhao, Y. Patel, and G. J. Offer, "Surface cooling causes accelerated degradation compared to tab cooling for lithium-ion pouch cells," *Journal of The Electrochemical Society*, vol. 163, p. A1846, jul 2016.

[62] M. Naylor Marlow, *Understanding the effects of heterogeneities on battery pack lifetime*. Phd thesis, Imperial College London, Jan 2021. Available at `http://hdl.handle.net/10044/1/92251`.

[63] P. Taheri and M. Bahrami, "Temperature rise in prismatic polymer lithium-ion batteries: An analytic approach," *SAE International Journal of Passenger Cars - Electronic and Electrical Systems*, vol. 5, p. 164176, 05 2012.

[64] Y. Zhao, Y. Patel, T. Zhang, and G. Offer, "Modeling the effects of thermal gradients induced by tab and surface cooling on lithium ion cell performance," *Journal of The Electrochemical Society*, vol. 165, pp. A3169–A3178, 10 2018.

[65] N. Meddings, M. Heinrich, F. Overney, J.-S. Lee, V. Ruiz, E. Napolitano, S. Seitz, G. Hinds, R. Raccichini, M. Gaberšček, and J. Park, "Application of electrochemical impedance spectroscopy to commercial li-ion cells: A review," *Journal of Power Sources*, vol. 480, p. 228742, 2020.

[66] W. H. Zhu, R. U. Payne, and B. J. Tatarchuk, "Pem stack test and analysis in a power system at operational load via ac impedance," *Journal of Power Sources*, vol. 168, no. 1, pp. 211–217, 2007. 10th EUROPEAN LEAD BATTERY CONFERENCE.

[67] E. B. andJ. Ross Macdonald, ed., *Impedance Spectroscopy Theory, Experiment, and Applications*. John Wiley & Sons, Inc., Hoboken, New Jersey: John Wiley & Sons, Inc., 2 ed., 2005. https://onlinelibrary.wiley.com/doi/pdf/10.1002/0471716243.fmatter.

[68] B. S. Instruments, "Impedance compensation," Application Note 04, 2 2019. https://www.biologic.net/wp-content/uploads/2019/08/impedance-compensation_material-science-an4.pdf.

[69] H. Li, Y. Wang, J. Huang, Y. Zhang, and J. Zhao, "Microwave-assisted synthesis of cus/graphene composite for enhanced lithium storage properties," *Electrochimica Acta*, vol. 225, pp. 443–451, 2017.

[70] W. Shi, X. Hu, J. Wang, J. Jiang, Y. Zhang, and T. Yip, "Analysis of thermal aging paths for large-format lifepo4/graphite battery," *Electrochimica Acta*, vol. 196, pp. 13–23, 2016.

[71] E. Madej, S. Klink, W. Schuhmann, E. Ventosa, and F. La Mantia, "Effect of the specific surface area on thermodynamic and kinetic properties of nanoparticle anatase tio2 in lithium-ion batteries," *Journal of Power Sources*, vol. 297, pp. 140–148, 2015.

[72] M. Luo, X. Lin, H. Lan, H. Yu, L. Yan, S. Qian, N. Long, M. Shui, and J. Shu, "Lithiation-delithiation kinetics of bali2ti6o14 anode in high-performance secondary li-ion batteries," *Journal of Electroanalytical Chemistry*, vol. 786, pp. 86–93, 2017.

[73] J. Ma, C. Wang, and S. Wroblewski, "Kinetic characteristics of mixed conductive electrodes for lithium ion batteries," *Journal of Power Sources*, vol. 164, no. 2, pp. 849–856, 2007.

[74] J. Ding, Z. Lu, M. Wu, C. Liu, H. Ji, and G. Yang, "Preparation and performance characterization of alf3 as interface stabilizer coated li1.24ni0.12co0.12mn0.56o2 cathode for lithium-ion batteries," *Applied Surface Science*, vol. 406, pp. 21–29, 2017.

[75] I. Bloom, A. N. Jansen, D. P. Abraham, J. Knuth, S. A. Jones, V. S. Battaglia, and G. L. Henriksen, "Differential voltage analyses of high-power, lithium-ion cells: 1. technique and application," *Journal of Power Sources*, vol. 139, no. 1, pp. 295–303, 2005.

[76] A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures.," *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.

[77] N. Kirkaldy, M. Samieian, G. Offer, M. Marinescu, and Y. Patel, "Lithium-ion battery degradation: Measuring rapid loss of active silicon in silicon–graphite composite electrodes," *ACS Applied Energy Materials*, vol. 5, 11 2022.

[78] R. Prosser, G. Offer, and Y. Patel, "Lithium-ion diagnostics: The first quantitative in-operando technique for diagnosing lithium ion battery degradation modes under load with realistic thermal boundary conditions," *Journal of The Electrochemical Society*, vol. 168, p. 030532, mar 2021.

[79] T. Horiba, "Lithium-ion battery systems," *Proceedings of the IEEE*, vol. 102, no. 6, pp. 939–950, 2014.

[80] J. T. Warner, "Chapter 2 - electrochemistry basics," in *Lithium-Ion Battery Chemistries* (J. T. Warner, ed.), pp. 32–34, Elsevier, 2019.

[81] J. T. Warner, "Chapter 2 - electrochemistry basics," in *Lithium-Ion Battery Chemistries* (J. T. Warner, ed.), pp. 36–39, Elsevier, 2019.

[82] J. T. Warner, "Chapter 3 - lithium-ion battery operation," in *Lithium-Ion Battery Chemistries* (J. T. Warner, ed.), pp. 44–52, Elsevier, 2019.

[83] H. A. H. A. Kiehne, *Battery technology handbook*. Electrical engineering and electronics ; 60, New York: Marcel Dekker, 2003.

[84] C. Rahn and C.-Y. Wang, *Introduction*, ch. 1, pp. 1–10. John Wiley Sons, Ltd, 2013.

[85] Y. Liu, R. Zhang, J. Wang, and Y. Wang, "Current and future lithium-ion battery manufacturing," *IScience*, vol. 24, no. 4, 2021.

[86] K. Murashko *et al.*, "Thermal modelling of commercial lithium-ion batteries," 2016.

[87] J. T. Warner, "Chapter 9 - lithium-ion cell manufacturing," in *Lithium-Ion Battery Chemistries* (J. T. Warner, ed.), pp. 219–252, Elsevier, 2019.

[88] M. Winter, J. O. Besenhard, M. E. Spahr, and P. Novak, "Insertion electrode materials for rechargeable lithium batteries," *Advanced materials*, vol. 10, no. 10, pp. 725–763, 1998.

[89] H. Schmidt, B. Jerliu, E. Hüger, and J. Stahn, "Volume expansion of amorphous silicon electrodes during potentiostatic lithiation of li-ion batteries," *Electrochemistry Communications*, vol. 115, p. 106738, 2020.

[90] M. Li and J. Lu, "Cobalt in lithium-ion batteries," *Science*, vol. 367, no. 6481, pp. 979–980, 2020.

[91] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, "Commentary: The materials project: A materials genome approach to accelerating materials innovation," *APL Materials*, vol. 1, no. 1, 2013. Cited by: 5596; All Open Access, Gold Open Access, Green Open Access.

[92] G. Che, K. B. Jirage, E. R. Fisher, C. R. Martin, and H. Yoneyama, "Chemical-vapor deposition-based template synthesis of microtubular tis2 battery electrodes," *Journal of The Electrochemical Society*, vol. 144, p. 4296, dec 1997.

[93] J. Cho, Y.-W. Kim, B. Kim, J.-G. Lee, and B. Park, "A breakthrough in the safety of lithium secondary batteries by coating the cathode material with aipo4 nanoparticles," *Angewandte Chemie - International Edition*, vol. 42, no. 14, p. 1618 – 1621, 2003. Cited by: 344.

[94] T. Ohzuku, A. Ueda, and M. Nagayama, "Electrochemistry and structural chemistry of linio2 (r3m) for 4 volt secondary lithium cells," *Journal of The Electrochemical Society*, vol. 140, p. 1862, jul 1993.

[95] P. G. Bruce, A. R. Armstrong, and R. L. Gitzendanner, "New intercalation compounds for lithium batteries: Layered limno2," *Journal of Materials Chemistry*, vol. 9, no. 1, p. 193 – 198, 1999. Cited by: 221.

[96] F. Lin, I. Markus, D. Nordlund, T.-C. Weng, M. Asta, H. Xin, and M. Doeff, "Surface reconstruction and chemical evolution of stoichiometric layered cathode materials for lithium-ion batteries," *Nature communications*, vol. 5, p. 3529, 03 2014.

[97] S. K. Martha, O. Haik, E. Zinigrad, I. Exnar, T. Drezen, J. H. Miners, and D. Aurbach, "On the thermal stability of olivine cathode materials for lithium-ion batteries," *Journal of The Electrochemical Society*, vol. 158, p. A1115, aug 2011.

[98] R. Wang, X. He, L. He, F. Wang, R. Xiao, L. Gu, H. Li, and L. Chen, "Atomic structure of li2mno3 after partial delithiation and re-lithiation," *Advanced Energy Materials*, vol. 3, no. 10, pp. 1358–1367, 2013.

[99] M.-J. Lee, S. Lee, P. Oh, Y. Kim, and J. Cho, "High performance limn2o4 cathode materials grown with epitaxial layered nanostructure for li-ion batteries," *Nano Letters*, vol. 14, no. 2, pp. 993–999, 2014. PMID: 24392731.

[100] S. Choi and A. Manthiram, "Synthesis and electrochemical properties of lico2o4 spinel cathodes," *Journal of the Electrochemical Society*, vol. 149, no. 2, p. A162–A166, 2002. Cited by: 60.

[101] A. Yamada, S. Chung, and K. Hinokuma, "Optimized lifepo4 for lithium battery cathodes," *Journal of the Electrochemical Society*, vol. 148, no. 3, p. A224–A229, 2001. Cited by: 1879.

[102] D. Choi, D. Wang, I.-T. Bae, J. Xiao, Z. Nie, W. Wang, V. V. Viswanathan, Y. J. Lee, J.-G. Zhang, G. L. Graff, Z. Yang, and J. Liu, "Limnpo4 nanoplate grown via solid-state reaction in molten hydrocarbon for li-ion battery cathode," *Nano Letters*, vol. 10, no. 8, p. 2799 – 2805, 2010. Cited by: 370.

[103] J. Lloris, C. Pérez Vicente, and J. Tirado, "Improvement of the electrochemical performance of licopo4 5 v material using a novel synthesis procedure," *Electrochemical and Solid-State Letters*, vol. 5, no. 10, p. A234–A237, 2002. Cited by: 132.

[104] A. Sobkowiak, M. R. Roberts, R. Younesi, T. Ericsson, L. Häggström, C.-W. Tai, A. M. Andersson, K. Edström, T. Gustafsson, and F. Björefors, "Understanding and controlling the surface chemistry of lifeso4f for an enhanced cathode functionality," *Chemistry of Materials*, vol. 25, no. 15, p. 3020 – 3029, 2013. Cited by: 43.

[105] J. Barker, R. Gover, P. Burns, A. Bryan, M. Saidi, and J. Swoyer, "Structural and electrochemical properties of lithium vanadium fluorophosphate, livpo4f," *Journal of Power Sources*, vol. 146, no. 1-2, p. 516 – 520, 2005. Cited by: 149.

[106] M. Yoshio, H. Wang, K. Fukuda, T. Umeno, T. Abe, and Z. Ogumi, "Improvement of natural graphite as a lithium-ion battery anode material, from raw flake to carbon-coated sphereelectronic supplementary information (ESI) available: colour versions of figs. 6, 8 and 9. see http://www.rsc.org/suppdata/jm/b3/b316702j/," *J. Mater. Chem.*, vol. 14, no. 11, p. 1754, 2004.

[107] J. Zhu, D. Yang, Z. Yin, Q. Yan, and H. Zhang, "Graphene and graphene-based materials for energy storage applications," *Small*, vol. 10, pp. 3480–3498, Sept. 2014.

[108] V. Meunier, J. Kephart, C. Roland, and J. Bernholc, "Ab initio investigations of lithium diffusion in carbon nanotube systems," *Phys. Rev. Lett.*, vol. 88, p. 075506, Feb. 2002.

[109] C. M. Schauerman, M. J. Ganter, G. Gaustad, C. W. Babbitt, R. P. Raffaelle, and B. J. Landi, "Recycling single-wall carbon nanotube anodes from lithium ion batteries," *J. Mater. Chem.*, vol. 22, no. 24, p. 12008, 2012.

[110] K. Nishidate and M. Hasegawa, "Energetics of lithium ion adsorption on defective carbon nanotubes," *Phys. Rev. B Condens. Matter Mater. Phys.*, vol. 71, June 2005.

[111] J. Zhao, A. Buldum, J. Han, and Ping Lu J, "First-principles study of li-intercalated carbon nanotube ropes," *Phys. Rev. Lett.*, vol. 85, pp. 1706–1709, Aug. 2000.

[112] C. A. Bridges, X.-G. Sun, J. Zhao, M. P. Paranthaman, and S. Dai, "In situ observation of solid electrolyte interphase formation in ordered mesoporous hard carbon by small-angle neutron scattering," *J. Phys. Chem. C Nanomater. Interfaces*, vol. 116, pp. 7701–7711, Apr. 2012.

[113] H. Fujimoto, K. Tokumitsu, A. Mabuchi, N. Chinnasamy, and T. Kasuh, "The anode performance of the hard carbon for the lithium ion battery derived from the oxygen-containing aromatic precursors," *J. Power Sources*, vol. 195, pp. 7452–7456, Nov. 2010.

[114] J. Yang, X.-Y. Zhou, J. Li, Y.-L. Zou, and J.-J. Tang, "Study of nano-porous hard carbons as anode materials for lithium ion batteries," *Mater. Chem. Phys.*, vol. 135, pp. 445–450, Aug. 2012.

[115] Z. Chen, I. Belharouak, Y.-K. Sun, and K. Amine, "Titanium-based anode materials for safe lithium-ion batteries," *Adv. Funct. Mater.*, vol. 23, pp. 959–969, Feb. 2013.

[116] P. G. Bruce, B. Scrosati, and J.-M. Tarascon, "Nanomaterials for rechargeable lithium batteries," *Angew. Chem. Int. Ed Engl.*, vol. 47, no. 16, pp. 2930–2946, 2008.

[117] Z. Wang, L. Zhou, and X. W. D. Lou, "Metal oxide hollow nanostructures for lithium-ion batteries," *Adv. Mater.*, vol. 24, pp. 1903–1911, Apr. 2012.

[118] J. Yang, "SiOx-based anodes for secondary lithium batteries," *Solid State Ion.*, vol. 152-153, pp. 125–129, Dec. 2002.

[119] C.-M. Park, J.-H. Kim, H. Kim, and H.-J. Sohn, "Li-alloy based anode materials for li secondary batteries," *Chem. Soc. Rev.*, vol. 39, pp. 3115–3141, Aug. 2010.

[120] J. R. Szczech and S. Jin, "Nanostructured silicon for high capacity lithium battery anodes," *Energy Environ. Sci.*, vol. 4, no. 1, pp. 56–72, 2011.

[121] N. G. Rudawski, B. R. Yates, M. R. Holzworth, K. S. Jones, R. G. Elliman, and A. A. Volinsky, "Ion beam-mixed ge electrodes for high capacity li rechargeable batteries," *J. Power Sources*, vol. 223, pp. 336–340, Feb. 2013.

[122] A. M. Chockla, K. C. Klavetter, C. B. Mullins, and B. A. Korgel, "Solution-grown germanium nanowire anodes for lithium-ion batteries," *ACS Appl. Mater. Interfaces*, vol. 4, pp. 4658–4664, Sept. 2012.

[123] C.-H. Lai, M.-Y. Lu, and L.-J. Chen, "Metal sulfide nanostructures: synthesis, properties and applications in energy conversion and storage," *J. Mater. Chem.*, vol. 22, no. 1, pp. 19–30, 2012.

[124] S. Boyanov, K. Annou, C. Villevieille, M. Pelosi, D. Zitoun, and L. Monconduit, "Nanostructured transition metal phosphide as negative electrode for lithium-ion batteries," *Ionics (Kiel)*, vol. 14, pp. 183–190, June 2008.

[125] L. Ji, Z. Lin, M. Alcoutlabi, and X. Zhang, "Recent developments in nanostructured anode materials for rechargeable lithium-ion batteries," *Energy Environ. Sci.*, vol. 4, no. 8, p. 2682, 2011.

[126] J. Jiang, Y. Li, J. Liu, X. Huang, C. Yuan, and X. W. D. Lou, "Recent advances in metal oxide-based electrode architecture design for electrochemical energy storage," *Adv. Mater.*, vol. 24, pp. 5166–5180, Oct. 2012.

[127] R. Holze, "Self-discharge of batteries: Causes, mechanisms and remedies," *Advanced Materials Science and Technology*, vol. 4, no. 2, 2022.

[128] C. R. Birkl, M. R. Roberts, E. McTurk, P. G. Bruce, and D. A. Howey, "Degradation diagnostics for lithium ion cells," *Journal of Power Sources*, vol. 341, pp. 373–386, 2017.

[129] M. Dubarry and G. Baure, "Perspective on commercial li-ion battery testing, best practices for simple and effective protocols," *Electronics*, vol. 9, no. 1, 2020.

[130] A. Barai, K. Uddin, M. Dubarry, L. Somerville, A. McGordon, P. Jennings, and I. Bloom, "A comparison of methodologies for the non-invasive characterisation of commercial li-ion cells," *Progress in Energy and Combustion Science*, vol. 72, pp. 1–31, 2019.

[131] W. Weppner and R. A. Huggins, "Determination of the kinetic parameters of mixed-conducting electrodes and application to the system li3sb," *Journal of The Electrochemical Society*, vol. 124, no. 10, p. 1569, 1977.

[132] Q. Wang, H. Li, X. Huang, and L. Chen, "Determination of chemical diffusion coefficient of lithium ion in graphitized mesocarbon microbeads with potential relaxation technique," *Journal of the Electrochemical Society*, vol. 148, no. 7, p. A737, 2001.

[133] A. Nazari and S. Farhad, "Heat generation in lithium-ion batteries with different nominal capacities and chemistries," *Applied Thermal Engineering*, vol. 125, pp. 1501–1517, 2017.

[134] A. A. Holland, *The effects of compression on lithium-ion batteries.* Phd thesis, Imperial College London, October 2018. https://doi.org/10.25560/68739.

[135] J. O. Almen and A. Laszlo, "The uniform-section disk spring," *Transactions of the American Society of Mechanical Engineers*, vol. 58, pp. 305–314, 02 1936.

[136] "Disc springs - Quality specifications - Dimensions," standard, Deutsches Institut für Normung e.V., Am DIN-Platz, Burggrafenstr. 6 10787 Berlin, Sept. 2016.

[137] H.-J. Noh, S. Youn, C. S. Yoon, and Y.-K. Sun, "Comparison of the structural and electrochemical properties of layered li[nixcoymnz]o2 (x=1/3, 0.5, 0.6, 0.7, 0.8 and 0.85) cathode material for lithium-ion batteries," *Journal of Power Sources*, vol. 233, pp. 121–130, 2013.

[138] D. D. MacNeil, Z. Lu, and J. R. Dahn, "Structure and electrochemistry of li[nix co1-2x mnx]o2 ($0 \leq x \leq 1/2$)," *Journal of The Electrochemical Society*, vol. 149, p. A1332, aug 2002.

[139] X. Zhao, Y. Yin, Y. Hu, and S.-Y. Choe, "Electrochemical-thermal modeling of lithium plating/stripping of li(ni0.6mn0.2co0.2)o2/carbon lithium-ion batteries at subzero ambient temperatures," *Journal of Power Sources*, vol. 418, pp. 61–73, 2019.

[140] W.-J. Zhang, "A review of the electrochemical performance of alloy anodes for lithium-ion batteries," *Journal of Power Sources*, vol. 196, no. 1, pp. 13–24, 2011.

[141] M. t. Pearson, "Apparatus for improving the performance of a fuel cell electric power system," 2003. US Patent US7632583B2, https://patents.google.com/patent/US7632583.

[142] H. F. C. technologies, "Horizon 5000w fuel cell stack user manual," 2013. http://www.fuelcellstore.com/manuals/horizon-pem-fuel-cell-h-5000-manual.pdf, Accessed 12/04/2018.

[143] A. Parekh, "Recent developments of proton exchange membranes for pemfc: A review," *Frontiers in Energy Research*, vol. 10, 2022.

[144] C. Sishtla, G. Koncar, R. Platon, S. Gamburzev, A. Appleby, and O. A. Velev, "Performance and endurance of a pemfc operated with synthetic reformate fuel feed," *Journal of Power Sources*, vol. 71, no. 1, pp. 249–255, 1998.

[145] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 3, pp. 103–107. Wiley, 2009.

[146] Y.-T. Pan, D. Li, S. Sharma, C. Wang, M. J. Zachman, E. C. Wegener, A. J. Kropf, Y. S. Kim, D. J. Myers, A. A. Peterson, D. A. Cullen, and J. S. Spendelow, "Ordered copt oxygen reduction catalyst with high performance and durability," *Chem Catalysis*, vol. 2, no. 12, pp. 3559–3572, 2022.

[147] Z. Wan, F. Liu, H. Xu, S. Zhao, Z. An, Z. Ma, Z. Zhang, Y. Wu, and X. Wang, "Recent advances and trends of single-atom catalysts for proton exchange membrane fuel cell cathodes," *ChemPhysMater*, vol. 3, no. 2, pp. 143–156, 2024.

[148] L. Osmieri and Q. Meyer, "Recent advances in integrating platinum group metal-free catalysts in proton exchange membrane fuel cells," *Current Opinion in Electrochemistry*, vol. 31, p. 100847, 2022.

[149] O. Savadogo, "Emerging membranes for electrochemical systems: Part ii. high temperature composite membranes for polymer electrolyte fuel cell (pefc) applications," *Journal of Power Sources*, vol. 127, no. 1, pp. 135–161, 2004. Eighth Ulmer Electrochemische Tage.

[150] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 9, pp. 304–307. Wiley, 2009.

[151] D. Cheddie and N. Munroe, "Mathematical model of a pemfc using a pbi membrane," *Energy Conversion and Management*, vol. 47, no. 11, pp. 1490–1504, 2006.

[152] N. F. Asri, T. Husaini, A. B. Sulong, E. H. Majlan, and W. R. W. Daud, "Coating of stainless steel and titanium bipolar plates for anticorrosion in pemfc: A review," *International Journal of Hydrogen Energy*, vol. 42, no. 14, pp. 9135–9148, 2017. Special Issue on Sustainable Fuel Cell and Hydrogen Technologies: The 5th International Conference on Fuel Cell and Hydrogen Technology (ICFCHT 2015), 1-3 September 2015, Kuala Lumpur, Malaysia.

[153] P. R. Pedapati, S. R. Dhanushkodi, R. K. Chidambaram, D. Taler, T. Sobota, and J. Taler, "Design and manufacturing challenges in pemfc flow fields—a review," *Energies*, vol. 17, no. 14, 2024.

[154] M. H. Bargal, M. A. Abdelkareem, Q. Tao, J. Li, J. Shi, and Y. Wang, "Liquid cooling techniques in proton exchange membrane fuel cell stacks: A detailed survey," *Alexandria Engineering Journal*, vol. 59, no. 2, pp. 635–655, 2020.

[155] A. Jahn, B. Kessler, A. Berger, M. Seifert, and T. Roch, "Innovative welding conception for thin-walled sheet metal components," in *Conference on Lasers in Manufacturing 2023*, 2023.

[156] S. Park, J.-W. Lee, and B. N. Popov, "A review of gas diffusion layer in pem fuel cells: Materials and designs," *International Journal of Hydrogen Energy*, vol. 37, no. 7, pp. 5850–5865, 2012. XII International Symposium on Polymer Electrolytes: New Materials for Application in Proton Exchange Membrane Fuel Cells.

[157] G. Athanasaki, A. Jayakumar, and A. Kannan, "Gas diffusion layers for pem fuel cells: Materials, properties and manufacturing – a review," *International Journal of Hydrogen Energy*, vol. 48, no. 6, pp. 2294–2313, 2023.

[158] W. Vielstich, A. Lamm, H. A. Gasteiger, and H. Yokokawa, *Principles of MEA preparation*. Major Reference Works, Oxford :: Wiley-Blackwell,, 1 ed.

[159] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 9, pp. 310–319. Wiley, 2009.

[160] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 9, pp. 307–309. Wiley, 2009.

[161] J. Lei, S. Zhao, Y. Liu, X. Deng, G. Song, and H. Peng, "Effect of cathode channel structure on the performance of open-cathode proton exchange membrane fuel cell in dry environments," *Energy Reports*, vol. 9, pp. 3022–3034, 2023.

[162] X. Yu, C. Zhang, M. Li, B. Chen, M. Fan, X. Sheng, and F. Ma, "Experimental investigation of self-regulating capability of open-cathode pemfc under different fan working conditions," *International Journal of Hydrogen Energy*, vol. 48, no. 68, pp. 26599–26608, 2023. Special Issue on HYPOTHESIS XVI.

[163] R. B. Ferreira, D. F. Santos, A. Pinto, and D. Falcão, "Development and testing of a pem fuel cell stack envisioning unmanned aerial vehicles applications," *International Journal of Hydrogen Energy*, vol. 51, pp. 1345–1353, 2024.

[164] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 2, pp. 50–54. Wiley, 2009.

[165] P. Moldrik and R. Chvalek, "Pem fuel cells — the basic characteristics," in *2011 10th International Conference on Environment and Electrical Engineering*, pp. 1–4, 2011.

[166] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 3, pp. 86–89. Wiley, 2009.

[167] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 3, pp. 89–94. Wiley, 2009.

[168] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 3, pp. 94–97. Wiley, 2009.

[169] M. Jouin, R. Gouriveau, D. Hissel, M.-C. Péra, and N. Zerhouni, "Degradations analysis and aging modeling for health assessment and prognostics of PEMFC," *Reliability Engineering & System Safety*, vol. 148, pp. 78–95, apr 2016.

[170] J. J. Hwang and W. R. Chang, "Characteristic study on fuel cell/battery hybrid power system on a light electric vehicle," *Journal of Power Sources*, vol. 207, pp. 111–119, 2012.

[171] W. Bi and T. F. Fuller, "Modeling of pem fuel cell pt/c catalyst degradation," *Journal of Power Sources*, vol. 178, pp. 188–196, 3 2008.

[172] A. B. Anderson*, , J. Roques, S. Mukerjee‡, , V. S. Murthi, N. M. Markovic§, , and V. Stamenkovic, "Activation energies for oxygen reduction on platinum alloys: theory and experiment," 2005.

[173] L. Khotseng, *Oxygen Reduction Reaction*. IntechOpen, 12 2018.

[174] K. Li, Y. Li, Y. Wang, F. He, M. Jiao, H. Tang, and Z. Wu, "The oxygen reduction reaction on pt(111) and pt(100) surfaces substituted by subsurface cu: a theoretical perspective," *Journal of Materials Chemistry A*, vol. 3, pp. 11444–11452, 5 2015.

[175] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 3, pp. 101–101. Wiley, 2009.

[176] . J. K. Nørskov, J. Rossmeisl, , A. Logadottir, L. Lindqvist, J. R. Kitchin, T. Bligaard, and H. Jónsson, "Origin of the overpotential for oxygen reduction at a fuel-cell cathode," 2004.

[177] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 3, p. 128. Wiley, 2009.

[178] A. S. Rao, K. Rashmi, D. Manjunatha, A. Jayarama, S. Prabhu, and R. Pinto, "Pore size tuning of nafion membranes by uv irradiation for enhanced proton conductivity for fuel cell applications," *International Journal of Hydrogen Energy*, vol. 44, no. 42, pp. 23762–23774, 2019.

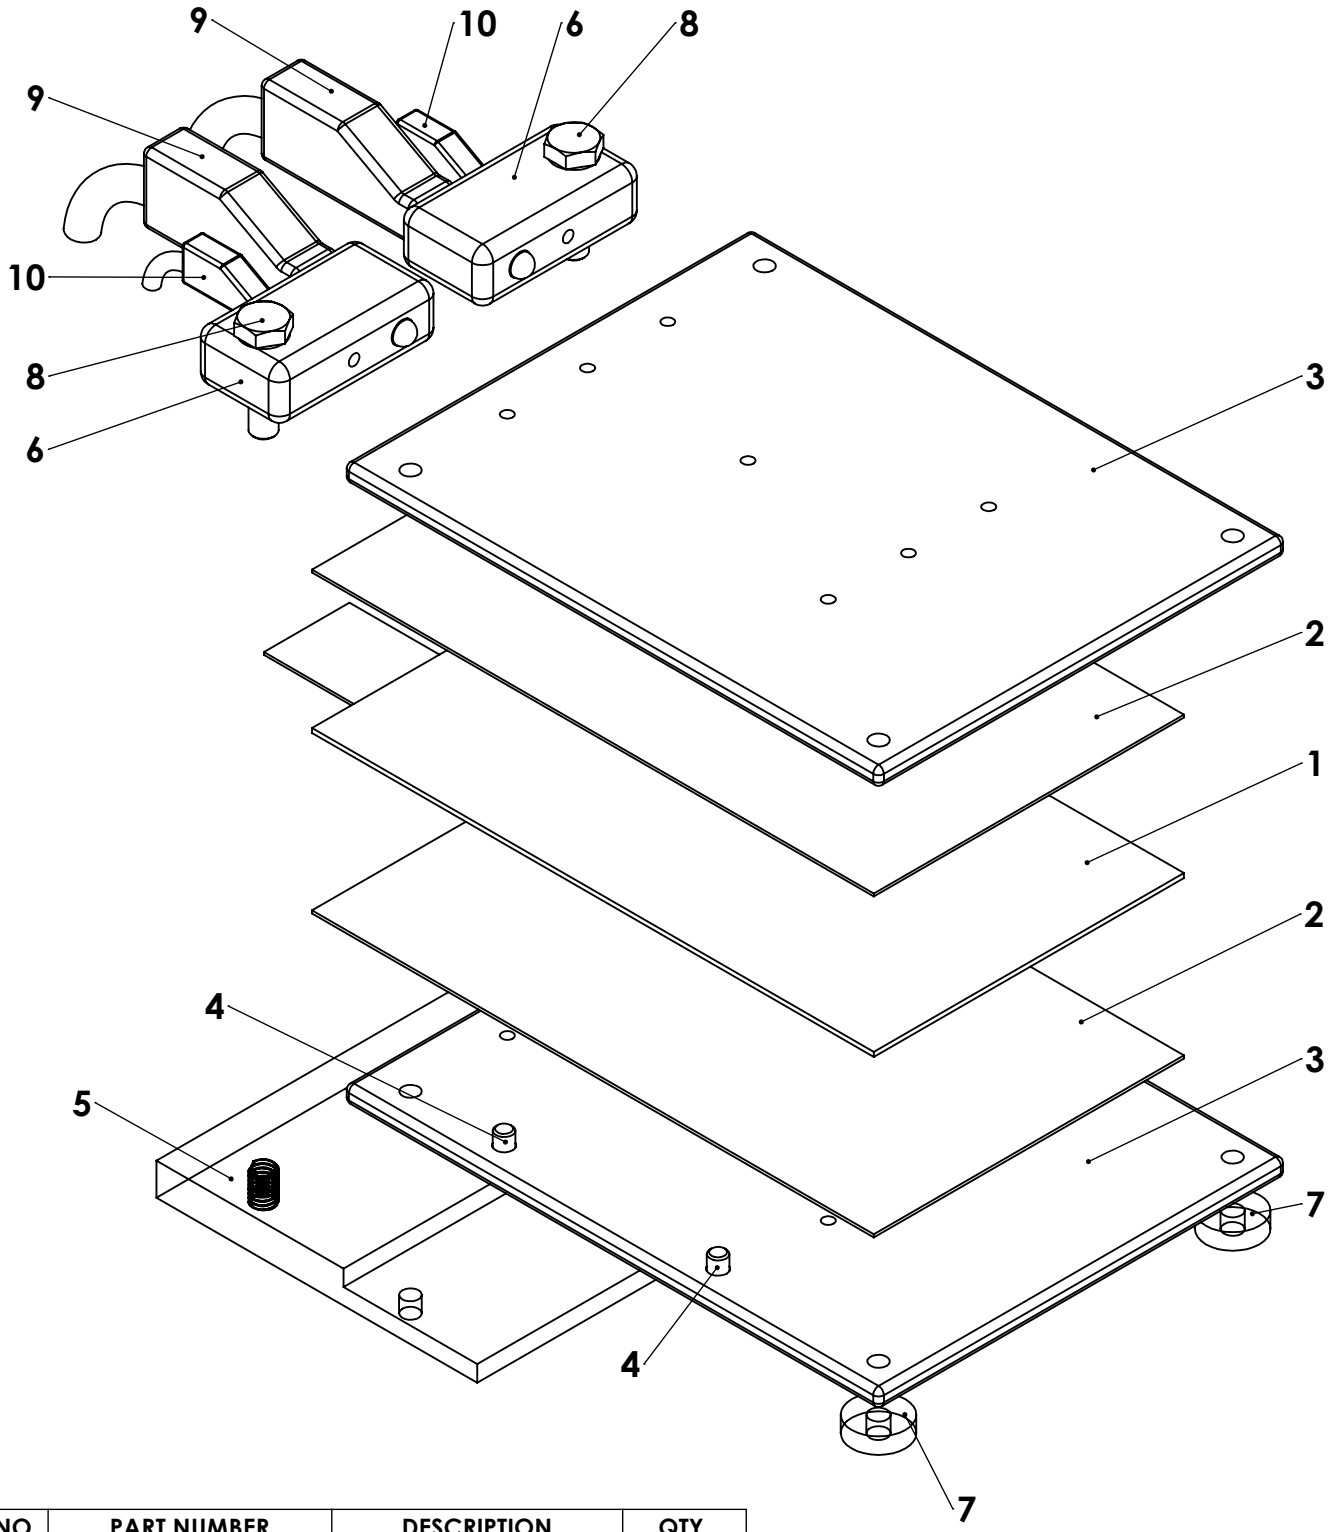
[179] R. O'Hayre, S. won Cha, W. Colella, and F. B. Prinz, *Fuel Cell Fundamentals*, ch. 3, pp. 136–145. Wiley, 2009.

[180] hyOzd, "serialplot," 2024.

[181] W. A. F. D. D. A. L. R. L. P. Scartozzi, "Fuel cell and method for controlling same," 1997. US Patent US6096449A, https://patents.google.com/patent/US6096449A/en.

Appendix A

Diagrams of Single Layer Rig



ITEM NO.	PART NUMBER	DESCRIPTION	QTY.
1	Cell		1
2	Thermal interface material	> 5W/mK, 0.5mm thick	2
3	Compression plate	Mill finish bar stock	2
4	Dowel Pin	For alignment within unit, optional	4
5	Tab Plate	Sets unit spacing in stack	1
6	Tab clamp	Linished to reduce contact resistance	2
7	Spacer	Sets unit spacing in stack	2
8	Clamp bolt - M4 x 16mm		2
9	Current plug (4mm)		1
10	Sense plug (2mm)		1

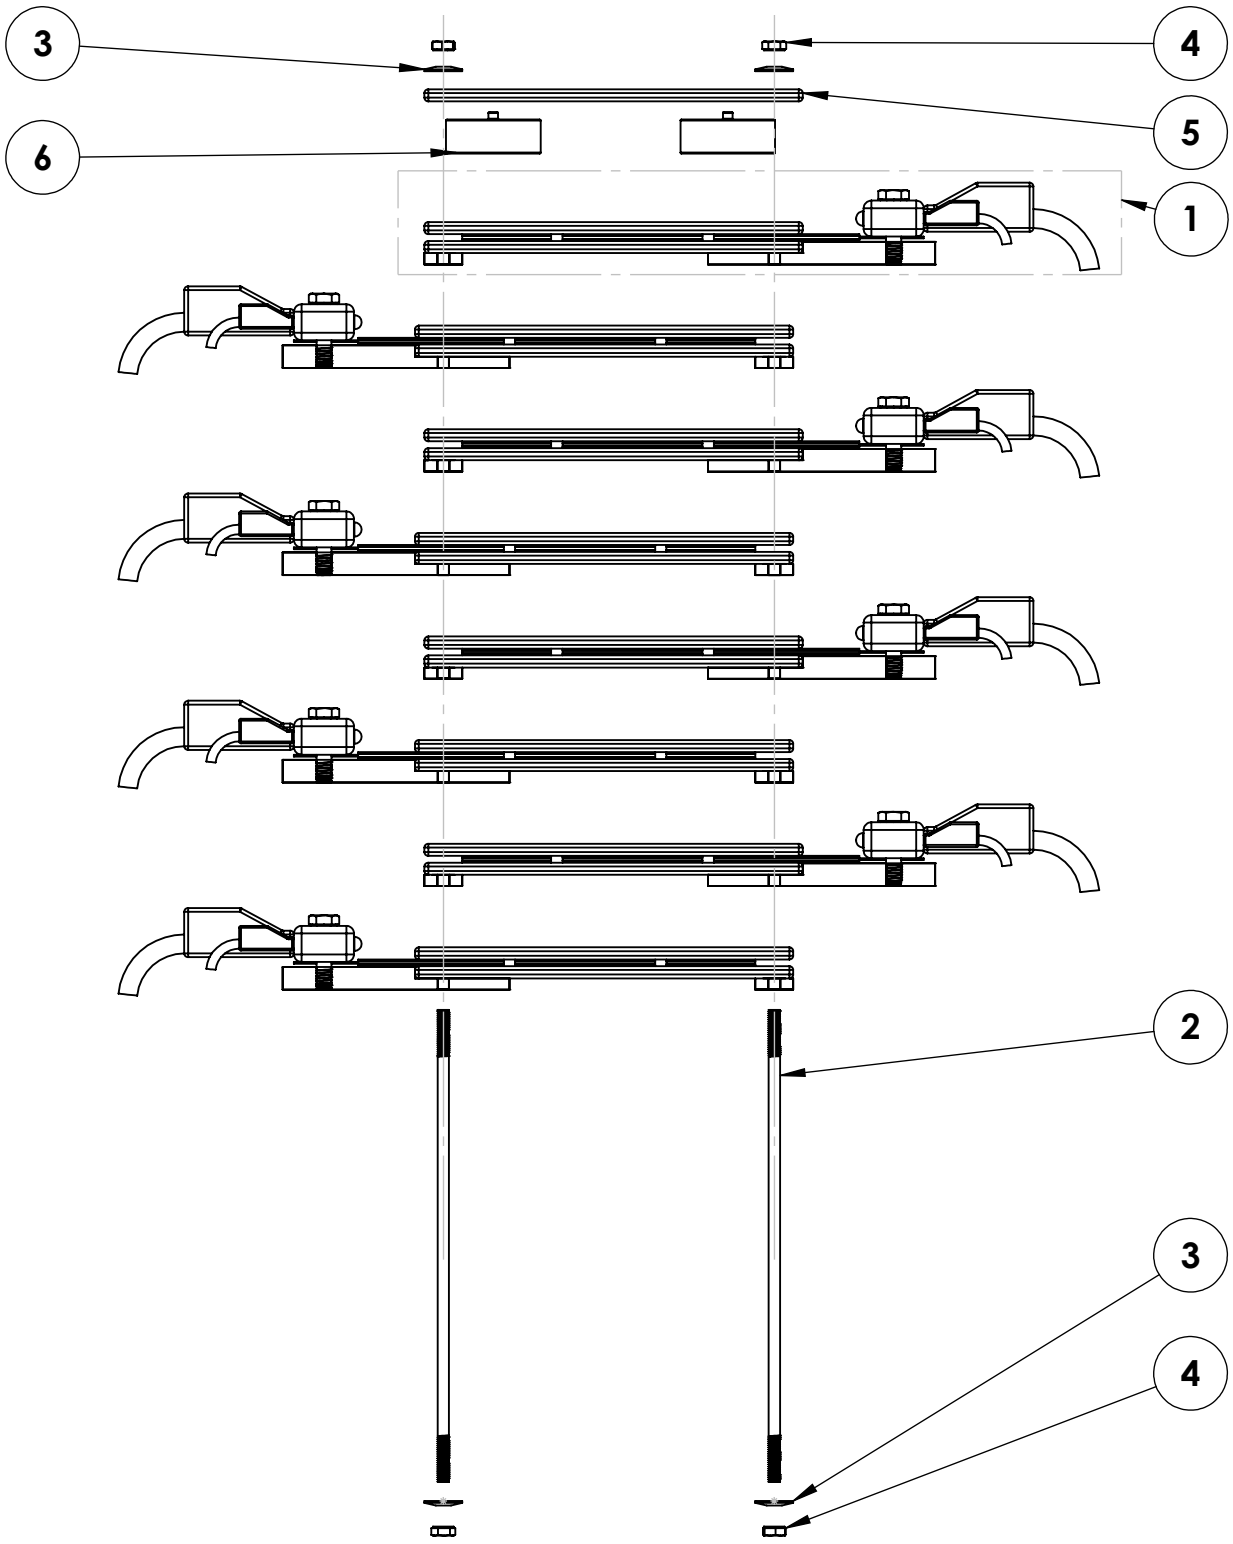
DO NOT SCALE DRAWING | REVISION

TITLE:

Stackable single layer rig cell unit

DWG NO. 1 | A4

SCALE:1:1 | SHEET 1 OF 1



ITEM NO.	PART NUMBER	DESCRIPTION	QTY.
1	Cell unit	Pre-assembled, or assembled on stack	4
2	Compression rod	Length determined by number of units	4
3	Belville washer	Multiples can be used depending on desired load	8
4	M3 nut		8
5	Compression plate	Identical to those used in cell unit	1
6	Load Cell	For compression instrumentation	4

DO NOT SCALE DRAWING | REVISION

TITLE:

DWG NO. **Stacked rig** A4

SCALE:1:2 | SHEET 1 OF 1

Appendix A

Permissions for reproduced works

ELSEVIER LICENSE
TERMS AND CONDITIONS

Jun 10, 2024

This Agreement between Oisin Shaw, Imperial College London ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	5805300248547
License date	Jun 10, 2024
Licensed Content Publisher	Elsevier
Licensed Content Publication	International Journal of Hydrogen Energy
Licensed Content Title	A systematic study on the use of short circuiting for the improvement of proton exchange membrane fuel cell performance
Licensed Content Author	Gaurav Gupta,Billy Wu,Simon Mylius,Gregory J. Offer
Licensed Content Date	Feb 16, 2017
Licensed Content Volume	42
Licensed Content Issue	7
Licensed Content Pages	8
Start Page	4320
End Page	4327
Type of Use	reuse in a thesis/dissertation

Portion	figures/tables/illustrations
Number of figures/tables/illustrations	1
Format	electronic
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title of new work	The Effect of Micro and Nano Scale Phenomena on System Level Design and Engineering of Hybrid PEMFC Systems
Institution name	Imperial College London
Expected presentation date	Jul 2024
Portions	Fig 3
The Requesting Person / Organization to Appear on the License	Oisin Shaw, Imperial College London
Requestor Location	Imperial College London Imperial College London Exhibition Road London, County SW7 2AZ United Kingdom Attn: Imperial College London
Publisher Tax ID	GB 494 6272 12
Total	0.00 USD
Terms and Conditions	

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and

conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your RightsLink account and that are available at any time at https://myaccount.copyright.com).

GENERAL TERMS

2. Elsevier hereby grants you permission to reproduce the aforementioned material subject to the terms and conditions indicated.

3. Acknowledgement: If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source, permission must also be sought from that source. If such permission is not obtained then that material may not be included in your publication/copies. Suitable acknowledgement to the source must be made, either as a footnote or in a reference list at the end of your publication, as follows:

"Reprinted from Publication title, Vol /edition number, Author(s), Title of article / title of chapter, Pages No., Copyright (Year), with permission from Elsevier [OR APPLICABLE SOCIETY COPYRIGHT OWNER]." Also Lancet special credit - "Reprinted from The Lancet, Vol. number, Author(s), Title of article, Pages No., Copyright (Year), with permission from Elsevier."

4. Reproduction of this material is confined to the purpose and/or media for which permission is hereby given. The material may not be reproduced or used in any other way, including use in combination with an artificial intelligence tool (including to train an algorithm, test, process, analyse, generate output and/or develop any form of artificial intelligence tool), or to create any derivative work and/or service (including resulting from the use of artificial intelligence tools).

5. Altering/Modifying Material: Not Permitted. However figures and illustrations may be altered/adapted minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of Elsevier Ltd. (Please contact Elsevier's permissions helpdesk here). No modifications can be made to any Lancet figures/tables and they must be reproduced in full.

6. If the permission fee for the requested use of our material is waived in this instance, please be advised that your future requests for Elsevier materials may attract a fee.

7. Reservation of Rights: Publisher reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

8. License Contingent Upon Payment: While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by publisher or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.

9. Warranties: Publisher makes no representations or warranties with respect to the licensed material.

10. Indemnity: You hereby indemnify and agree to hold harmless publisher and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

11. No Transfer of License: This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.

12. No Amendment Except in Writing: This license may not be amended except in a writing signed by both parties (or, in the case of publisher, by CCC on publisher's behalf).

13. Objection to Contrary Terms: Publisher hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and publisher (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

14. Revocation: Elsevier or Copyright Clearance Center may deny the permissions described in this License at their sole discretion, for any reason or no reason, with a full refund payable to you. Notice of such denial will be made using the contact information provided by you. Failure to receive such notice will not alter or invalidate the denial. In no event will Elsevier or Copyright Clearance Center be responsible or liable for any costs, expenses or damage incurred by you as a result of a denial of your permission request, other than a refund of the amount(s) paid by you to Elsevier and/or Copyright Clearance Center for denied permissions.

LIMITED LICENSE

The following terms and conditions apply only to specific license types:

15. **Translation**: This permission is granted for non-exclusive world **English** rights only unless your license was granted for translation rights. If you licensed translation rights you may only translate this content into the languages you requested. A professional translator must perform all translations and reproduce the content word for word preserving the integrity of the article.

16. **Posting licensed content on any Website**: The following terms and conditions apply as follows: Licensing material from an Elsevier journal: All content posted to the web site must maintain the copyright information line on the bottom of each image; A hyper-text must be included to the Homepage of the journal from which you are licensing at http://www.sciencedirect.com/science/journal/xxxxx or the Elsevier homepage for books at http://www.elsevier.com; Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

Licensing material from an Elsevier book: A hyper-text link must be included to the Elsevier homepage at http://www.elsevier.com . All content posted to the web site must maintain the copyright information line on the bottom of each image.

Posting licensed content on Electronic reserve: In addition to the above the following clauses are applicable: The web site must be password-protected and made available only to bona fide students registered on a relevant course. This permission is granted for 1 year only. You may obtain a new license for future website posting.

17. **For journal authors:** the following clauses are applicable in addition to the above:

Preprints:

A preprint is an author's own write-up of research results and analysis, it has not been peer-reviewed, nor has it had any other value added to it by a publisher (such as formatting, copyright, technical enhancement etc.).

Authors can share their preprints anywhere at any time. Preprints should not be added to or enhanced in any way in order to appear more like, or to substitute for, the final versions of articles however authors can update their preprints on arXiv or RePEc with their Accepted Author Manuscript (see below).

If accepted for publication, we encourage authors to link from the preprint to their formal publication via its DOI. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help users to find, access, cite and use the best available version. Please note that Cell Press, The Lancet and some society-owned have different preprint policies. Information on these policies is available on the journal homepage.

Accepted Author Manuscripts: An accepted author manuscript is the manuscript of an article that has been accepted for publication and which typically includes author-incorporated changes suggested during submission, peer review and editor-author communications.

Authors can share their accepted author manuscript:

- immediately
    - via their non-commercial person homepage or blog
    - by updating a preprint in arXiv or RePEc with the accepted manuscript
    - via their research institute or institutional repository for internal institutional uses or as part of an invitation-only research collaboration work-group
    - directly by providing copies to their students or to research collaborators for their personal use
    - for private scholarly sharing as part of an invitation-only work group on commercial sites with which Elsevier has an agreement
- After the embargo period
    - via non-commercial hosting platforms such as their institutional repository
    - via commercial sites with which Elsevier has an agreement

In all cases accepted manuscripts should:

- link to the formal publication via its DOI
- bear a CC-BY-NC-ND license - this is easy to do
- if aggregated with other manuscripts, for example in a repository or other site, be shared in alignment with our hosting policy not be added to or enhanced in any way to appear more like, or to substitute for, the published journal article.

Published journal article (JPA): A published journal article (PJA) is the definitive final record of published research that appears or will appear in the journal and embodies all value-adding publishing activities including peer review co-ordination, copy-editing, formatting, (if relevant) pagination and online enrichment.

Policies for sharing publishing journal articles differ for subscription and gold open access articles:

Subscription Articles: If you are an author, please share a link to your article rather than the full-text. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help your users to find, access, cite, and use the best available version.

Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

If you are affiliated with a library that subscribes to ScienceDirect you have additional private sharing rights for others' research accessed under that agreement. This includes use for classroom teaching and internal training at the institution (including use in course packs and courseware programs), and inclusion of the article for grant funding purposes.

Gold Open Access Articles: May be shared according to the author-selected end-user license and should contain a CrossMark logo, the end user license, and a DOI link to the formal publication on ScienceDirect.

Please refer to Elsevier's posting policy for further information.

18. **For book authors** the following clauses are applicable in addition to the above: Authors are permitted to place a brief summary of their work online only. You are not allowed to download and post the published electronic version of your chapter, nor may you scan the printed edition to create an electronic version. **Posting to a repository:** Authors are permitted to post a summary of their chapter only in their institution's repository.

19. **Thesis/Dissertation**: If your license is for use in a thesis/dissertation your thesis may be submitted to your institution in either print or electronic form. Should your thesis be published commercially, please reapply for permission. These requirements include permission for the Library and Archives of Canada to supply single copies, on demand, of the complete thesis and include permission for Proquest/UMI to supply single copies, on demand, of the complete thesis. Should your thesis be published commercially, please reapply for permission. Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

Elsevier Open Access Terms and Conditions

You can publish open access with Elsevier in hundreds of open access journals or in nearly 2000 established subscription journals that support open access publishing. Permitted third party re-use of these open access articles is defined by the author's choice of Creative Commons user license. See our open access license policy for more information.

Terms & Conditions applicable to all Open Access articles published with Elsevier:

Any reuse of the article must not represent the author as endorsing the adaptation of the article nor should the article be modified in such a way as to damage the author's honour or reputation. If any changes have been made, such changes must be clearly indicated.

The author(s) must be appropriately credited and we ask that you include the end user license and a DOI link to the formal publication on ScienceDirect.

If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source it is the responsibility of the user to ensure their reuse complies with the terms and conditions determined by the rights holder.

Additional Terms & Conditions applicable to each Creative Commons user license:

CC BY: The CC-BY license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article and to make commercial use of the Article (including reuse and/or resale of the Article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by/4.0.

CC BY NC SA: The CC BY-NC-SA license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article, provided this is not done for commercial purposes, and that the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. Further, any new works must be made available on the same conditions. The full details of the license are available at http://creativecommons.org/licenses/by-nc-sa/4.0.

CC BY NC ND: The CC BY-NC-ND license allows users to copy and distribute the Article, provided this is not done for commercial purposes and further does not permit distribution of the Article if it is changed or edited in any way, and provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, and that the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by-nc-nd/4.0. Any commercial reuse of Open Access articles published with a CC BY NC SA or CC BY NC ND license requires permission from Elsevier and will be subject to a fee.

Commercial reuse includes:

- Associating advertising with the full text of the Article
- Charging fees for document delivery or access
- Article aggregation
- Systematic distribution via e-mail lists or share buttons

Posting or linking by commercial companies for use by customers of those companies.

20. **Other Conditions**:

v1.10

Questions? customercare@copyright.com.

JOHN WILEY AND SONS LICENSE
TERMS AND CONDITIONS

Jun 10, 2024

This Agreement between Oisin Shaw, Imperial College London ("You") and John Wiley and Sons ("John Wiley and Sons") consists of your license details and the terms and conditions provided by John Wiley and Sons and Copyright Clearance Center.

License Number	5805271137729
License date	Jun 10, 2024
Licensed Content Publisher	John Wiley and Sons
Licensed Content Publication	Fuel Cells
Licensed Content Title	Proton Exchange Membrane Fuel Cell Operation and Degradation in Short-Circuit
Licensed Content Author	K. Agbossou, L. Boulon, D. Hissel, et al
Licensed Content Date	Sep 15, 2014
Licensed Content Volume	14
Licensed Content Issue	6
Licensed Content Pages	12
Type of use	Dissertation/Thesis
Requestor type	University/Academic
Format	Electronic
Portion	Figure/table

Number of figures/tables	3
Will you be translating?	No
Title of new work	The Effect of Micro and Nano Scale Phenomena on System Level Design and Engineering of Hybrid PEMFC Systems
Institution name	Imperial College London
Expected presentation date	Jul 2024
Portions	Table 6, Fig 11, Fig 12
The Requesting Person / Organization to Appear on the License	Oisin Shaw, Imperial College London
Requestor Location	Imperial College London Exhibition Road London, County SW7 2AZ United Kingdom Attn: Imperial College London
Publisher Tax ID	EU826007151
Total	0.00 USD
Terms and Conditions	

TERMS AND CONDITIONS

This copyrighted material is owned by or exclusively licensed to John Wiley & Sons, Inc. or one of its group companies (each a"Wiley Company") or handled on behalf of a society with which a Wiley Company has exclusive publishing rights in relation to a particular work (collectively "WILEY"). By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the billing and payment terms and conditions established by the Copyright Clearance Center Inc., ("CCC's Billing and Payment terms and conditions"), at the time that you opened your RightsLink account (these are available at any time at http://myaccount.copyright.com).

Terms and Conditions

- The materials you have requested permission to reproduce or reuse (the "Wiley Materials") are protected by copyright.

- You are hereby granted a personal, non-exclusive, non-sub licensable (on a stand-alone basis), non-transferable, worldwide, limited license to reproduce the Wiley Materials for the purpose specified in the licensing process. This license, **and any CONTENT (PDF or image file) purchased as part of your order,** is for a one-time use only and limited to any maximum distribution number specified in the license. The first instance of republication or reuse granted by this license must be completed within two years of the date of the grant of this license (although copies prepared before the end date may be distributed thereafter). The Wiley Materials shall not be used in any other manner or for any other purpose, beyond what is granted in the license. Permission is granted subject to an appropriate acknowledgement given to the author, title of the material/book/journal and the publisher. You shall also duplicate the copyright notice that appears in the Wiley publication in your use of the Wiley Material. Permission is also granted on the understanding that nowhere in the text is a previously published source acknowledged for all or part of this Wiley Material. Any third party content is expressly excluded from this permission.

- With respect to the Wiley Materials, all rights are reserved. Except as expressly granted by the terms of the license, no part of the Wiley Materials may be copied, modified, adapted (except for minor reformatting required by the new Publication), translated, reproduced, transferred or distributed, in any form or by any means, and no derivative works may be made based on the Wiley Materials without the prior permission of the respective copyright owner.**For STM Signatory Publishers clearing permission under the terms of the STM Permissions Guidelines only, the terms of the license are extended to include subsequent editions and for editions in other languages, provided such editions are for the work as a whole in situ and does not involve the separate exploitation of the permitted figures or extracts,** You may not alter, remove or suppress in any manner any copyright, trademark or other notices displayed by the Wiley Materials. You may not license, rent, sell, loan, lease, pledge, offer as security, transfer or assign the Wiley Materials on a stand-alone basis, or any of the rights granted to you hereunder to any other person.

- The Wiley Materials and all of the intellectual property rights therein shall at all times remain the exclusive property of John Wiley & Sons Inc, the Wiley Companies, or their respective licensors, and your interest therein is only that of having possession of and the right to reproduce the Wiley Materials pursuant to Section 2 herein during the continuance of this Agreement. You agree that you own no right, title or interest in or to the Wiley Materials or any of the intellectual property rights therein. You shall have no rights hereunder other than the license as provided for above in Section 2. No right, license or interest to any trademark, trade name, service mark or other branding ("Marks") of WILEY or its licensors is granted hereunder, and you agree that you shall not assert any such right, license or interest with respect thereto

- NEITHER WILEY NOR ITS LICENSORS MAKES ANY WARRANTY OR REPRESENTATION OF ANY KIND TO YOU OR ANY THIRD PARTY, EXPRESS, IMPLIED OR STATUTORY, WITH RESPECT TO THE MATERIALS OR THE ACCURACY OF ANY INFORMATION CONTAINED IN THE MATERIALS, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, ACCURACY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, USABILITY, INTEGRATION OR NON-INFRINGEMENT AND ALL SUCH WARRANTIES ARE HEREBY EXCLUDED BY WILEY AND ITS LICENSORS AND WAIVED

BY YOU.

- WILEY shall have the right to terminate this Agreement immediately upon breach of this Agreement by you.

- You shall indemnify, defend and hold harmless WILEY, its Licensors and their respective directors, officers, agents and employees, from and against any actual or threatened claims, demands, causes of action or proceedings arising from any breach of this Agreement by you.

- IN NO EVENT SHALL WILEY OR ITS LICENSORS BE LIABLE TO YOU OR ANY OTHER PARTY OR ANY OTHER PERSON OR ENTITY FOR ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL, INDIRECT, EXEMPLARY OR PUNITIVE DAMAGES, HOWEVER CAUSED, ARISING OUT OF OR IN CONNECTION WITH THE DOWNLOADING, PROVISIONING, VIEWING OR USE OF THE MATERIALS REGARDLESS OF THE FORM OF ACTION, WHETHER FOR BREACH OF CONTRACT, BREACH OF WARRANTY, TORT, NEGLIGENCE, INFRINGEMENT OR OTHERWISE (INCLUDING, WITHOUT LIMITATION, DAMAGES BASED ON LOSS OF PROFITS, DATA, FILES, USE, BUSINESS OPPORTUNITY OR CLAIMS OF THIRD PARTIES), AND WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY PROVIDED HEREIN.

- Should any provision of this Agreement be held by a court of competent jurisdiction to be illegal, invalid, or unenforceable, that provision shall be deemed amended to achieve as nearly as possible the same economic effect as the original provision, and the legality, validity and enforceability of the remaining provisions of this Agreement shall not be affected or impaired thereby.

- The failure of either party to enforce any term or condition of this Agreement shall not constitute a waiver of either party's right to enforce each and every term and condition of this Agreement. No breach under this agreement shall be deemed waived or excused by either party unless such waiver or consent is in writing signed by the party granting such waiver or consent. The waiver by or consent of a party to a breach of any provision of this Agreement shall not operate or be construed as a waiver of or consent to any other or subsequent breach by such other party.

- This Agreement may not be assigned (including by operation of law or otherwise) by you without WILEY's prior written consent.

- Any fee required for this permission shall be non-refundable after thirty (30) days from receipt by the CCC.

- These terms and conditions together with CCC's Billing and Payment terms and conditions (which are incorporated herein) form the entire agreement between you and WILEY concerning this licensing transaction and (in the absence of fraud) supersedes all prior agreements and representations of the parties, oral or written. This Agreement may not be amended except in writing signed by both parties. This Agreement shall be binding upon and inure to the benefit of the parties' successors, legal representatives, and authorized assigns.

- In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall prevail.

- WILEY expressly reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

- This Agreement will be void if the Type of Use, Format, Circulation, or Requestor Type was misrepresented during the licensing process.

- This Agreement shall be governed by and construed in accordance with the laws of the State of New York, USA, without regards to such state's conflict of law rules. Any legal action, suit or proceeding arising out of or relating to these Terms and Conditions or the breach thereof shall be instituted in a court of competent jurisdiction in New York County in the State of New York in the United States of America and each party hereby consents and submits to the personal jurisdiction of such court, waives any objection to venue in such court and consents to service of process by registered or certified mail, return receipt requested, at the last known address of such party.

WILEY OPEN ACCESS TERMS AND CONDITIONS

Wiley Publishes Open Access Articles in fully Open Access Journals and in Subscription journals offering Online Open. Although most of the fully Open Access journals publish open access articles under the terms of the Creative Commons Attribution (CC BY) License only, the subscription journals and a few of the Open Access Journals offer a choice of Creative Commons Licenses. The license type is clearly identified on the article.

The Creative Commons Attribution License

The Creative Commons Attribution License (CC-BY) allows users to copy, distribute and transmit an article, adapt the article and make commercial use of the article. The CC-BY license permits commercial and non-

Creative Commons Attribution Non-Commercial License

The Creative Commons Attribution Non-Commercial (CC-BY-NC)License permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.(see below)

Creative Commons Attribution-Non-Commercial-NoDerivs License

The Creative Commons Attribution Non-Commercial-NoDerivs License (CC-BY-NC-ND) permits use, distribution and reproduction in any medium, provided the original work is properly cited, is not used for commercial purposes and no modifications or adaptations are made. (see below)

Use by commercial "for-profit" organizations

Use of Wiley Open Access articles for commercial, promotional, or marketing purposes requires further explicit permission from Wiley and will be subject to a fee.

Further details can be found on Wiley Online Library
http://olabout.wiley.com/WileyCDA/Section/id-410895.html

Other Terms and Conditions:

v1.10 Last updated September 2015

Questions? customercare@copyright.com.

ELSEVIER LICENSE
TERMS AND CONDITIONS

Jun 10, 2024

This Agreement between Oisin Shaw, Imperial College London ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	5805270779892
License date	Jun 10, 2024
Licensed Content Publisher	Elsevier
Licensed Content Publication	Materials Today
Licensed Content Title	Li-ion battery materials: present and future
Licensed Content Author	Naoki Nitta,Feixiang Wu,Jung Tae Lee,Gleb Yushin
Licensed Content Date	Jun 1, 2015
Licensed Content Volume	18
Licensed Content Issue	5
Licensed Content Pages	13
Start Page	252
End Page	264
Type of Use	reuse in a thesis/dissertation
Portion	figures/tables/illustrations

Number of figures/tables/illustrations	1
Format	electronic
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title of new work	The Effect of Micro and Nano Scale Phenomena on System Level Design and Engineering of Hybrid PEMFC Systems
Institution name	Imperial College London
Expected presentation date	Jul 2024
Portions	Table 1
The Requesting Person / Organization to Appear on the License	Oisin Shaw, Imperial College London
Requestor Location	Imperial College London Exhibition Road London, County SW7 2AZ United Kingdom Attn: Imperial College London
Publisher Tax ID	GB 494 6272 12
Total	0.00 USD

Terms and Conditions

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your RightsLink account and that are available at any time at https://myaccount.copyright.com).

GENERAL TERMS

2. Elsevier hereby grants you permission to reproduce the aforementioned material subject to the terms and conditions indicated.

3. Acknowledgement: If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source, permission must also be sought from that source. If such permission is not obtained then that material may not be included in your publication/copies. Suitable acknowledgement to the source must be made, either as a footnote or in a reference list at the end of your publication, as follows:

"Reprinted from Publication title, Vol /edition number, Author(s), Title of article / title of chapter, Pages No., Copyright (Year), with permission from Elsevier [OR APPLICABLE SOCIETY COPYRIGHT OWNER]." Also Lancet special credit - "Reprinted from The Lancet, Vol. number, Author(s), Title of article, Pages No., Copyright (Year), with permission from Elsevier."

4. Reproduction of this material is confined to the purpose and/or media for which permission is hereby given. The material may not be reproduced or used in any other way, including use in combination with an artificial intelligence tool (including to train an algorithm, test, process, analyse, generate output and/or develop any form of artificial intelligence tool), or to create any derivative work and/or service (including resulting from the use of artificial intelligence tools).

5. Altering/Modifying Material: Not Permitted. However figures and illustrations may be altered/adapted minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of Elsevier Ltd. (Please contact Elsevier's permissions helpdesk here). No modifications can be made to any Lancet figures/tables and they must be reproduced in full.

6. If the permission fee for the requested use of our material is waived in this instance, please be advised that your future requests for Elsevier materials may attract a fee.

7. Reservation of Rights: Publisher reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

8. License Contingent Upon Payment: While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by publisher or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.

9. Warranties: Publisher makes no representations or warranties with respect to the licensed material.

10. Indemnity: You hereby indemnify and agree to hold harmless publisher and CCC, and their respective officers, directors, employees and agents, from and against any and all

claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

11. No Transfer of License: This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.

12. No Amendment Except in Writing: This license may not be amended except in a writing signed by both parties (or, in the case of publisher, by CCC on publisher's behalf).

13. Objection to Contrary Terms: Publisher hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and publisher (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

14. Revocation: Elsevier or Copyright Clearance Center may deny the permissions described in this License at their sole discretion, for any reason or no reason, with a full refund payable to you. Notice of such denial will be made using the contact information provided by you. Failure to receive such notice will not alter or invalidate the denial. In no event will Elsevier or Copyright Clearance Center be responsible or liable for any costs, expenses or damage incurred by you as a result of a denial of your permission request, other than a refund of the amount(s) paid by you to Elsevier and/or Copyright Clearance Center for denied permissions.

LIMITED LICENSE

The following terms and conditions apply only to specific license types:

15. **Translation**: This permission is granted for non-exclusive world **English** rights only unless your license was granted for translation rights. If you licensed translation rights you may only translate this content into the languages you requested. A professional translator must perform all translations and reproduce the content word for word preserving the integrity of the article.

16. **Posting licensed content on any Website**: The following terms and conditions apply as follows: Licensing material from an Elsevier journal: All content posted to the web site must maintain the copyright information line on the bottom of each image; A hyper-text must be included to the Homepage of the journal from which you are licensing at http://www.sciencedirect.com/science/journal/xxxxx or the Elsevier homepage for books at http://www.elsevier.com; Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

Licensing material from an Elsevier book: A hyper-text link must be included to the Elsevier homepage at http://www.elsevier.com . All content posted to the web site must maintain the copyright information line on the bottom of each image.

Posting licensed content on Electronic reserve: In addition to the above the following clauses are applicable: The web site must be password-protected and made available only to bona fide students registered on a relevant course. This permission is granted for 1 year only. You may obtain a new license for future website posting.

17. **For journal authors:** the following clauses are applicable in addition to the above:

Preprints:

A preprint is an author's own write-up of research results and analysis, it has not been peer-reviewed, nor has it had any other value added to it by a publisher (such as formatting, copyright, technical enhancement etc.).

Authors can share their preprints anywhere at any time. Preprints should not be added to or enhanced in any way in order to appear more like, or to substitute for, the final versions of articles however authors can update their preprints on arXiv or RePEc with their Accepted Author Manuscript (see below).

If accepted for publication, we encourage authors to link from the preprint to their formal publication via its DOI. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help users to find, access, cite and use the best available version. Please note that Cell Press, The Lancet and some society-owned have different preprint policies. Information on these policies is available on the journal homepage.

Accepted Author Manuscripts: An accepted author manuscript is the manuscript of an article that has been accepted for publication and which typically includes author-incorporated changes suggested during submission, peer review and editor-author communications.

Authors can share their accepted author manuscript:

- immediately
  - via their non-commercial person homepage or blog
  - by updating a preprint in arXiv or RePEc with the accepted manuscript
  - via their research institute or institutional repository for internal institutional uses or as part of an invitation-only research collaboration work-group
  - directly by providing copies to their students or to research collaborators for their personal use
  - for private scholarly sharing as part of an invitation-only work group on commercial sites with which Elsevier has an agreement
- After the embargo period
  - via non-commercial hosting platforms such as their institutional repository
  - via commercial sites with which Elsevier has an agreement

In all cases accepted manuscripts should:

- link to the formal publication via its DOI
- bear a CC-BY-NC-ND license - this is easy to do
- if aggregated with other manuscripts, for example in a repository or other site, be shared in alignment with our hosting policy not be added to or enhanced in any way to appear more like, or to substitute for, the published journal article.

Published journal article (JPA): A published journal article (PJA) is the definitive final record of published research that appears or will appear in the journal and embodies all value-adding publishing activities including peer review co-ordination, copy-editing, formatting, (if relevant) pagination and online enrichment.

Policies for sharing publishing journal articles differ for subscription and gold open access articles:

Subscription Articles: If you are an author, please share a link to your article rather than the full-text. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help your users to find, access, cite, and use the best available version.

Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

If you are affiliated with a library that subscribes to ScienceDirect you have additional private sharing rights for others' research accessed under that agreement. This includes use for classroom teaching and internal training at the institution (including use in course packs and courseware programs), and inclusion of the article for grant funding purposes.

Gold Open Access Articles: May be shared according to the author-selected end-user license and should contain a CrossMark logo, the end user license, and a DOI link to the formal publication on ScienceDirect.

Please refer to Elsevier's posting policy for further information.

18. **For book authors** the following clauses are applicable in addition to the above: Authors are permitted to place a brief summary of their work online only. You are not allowed to download and post the published electronic version of your chapter, nor may you scan the printed edition to create an electronic version. **Posting to a repository:** Authors are permitted to post a summary of their chapter only in their institution's repository.

19. **Thesis/Dissertation**: If your license is for use in a thesis/dissertation your thesis may be submitted to your institution in either print or electronic form. Should your thesis be published commercially, please reapply for permission. These requirements include permission for the Library and Archives of Canada to supply single copies, on demand, of the complete thesis and include permission for Proquest/UMI to supply single copies, on demand, of the complete thesis. Should your thesis be published commercially, please reapply for permission. Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

Elsevier Open Access Terms and Conditions

You can publish open access with Elsevier in hundreds of open access journals or in nearly 2000 established subscription journals that support open access publishing. Permitted third party re-use of these open access articles is defined by the author's choice of Creative Commons user license. See our open access license policy for more information.

Terms & Conditions applicable to all Open Access articles published with Elsevier:

Any reuse of the article must not represent the author as endorsing the adaptation of the article nor should the article be modified in such a way as to damage the author's honour or reputation. If any changes have been made, such changes must be clearly indicated.

The author(s) must be appropriately credited and we ask that you include the end user license and a DOI link to the formal publication on ScienceDirect.

If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source it is the responsibility of the user to ensure their reuse complies with the terms and conditions determined by the rights holder.

Additional Terms & Conditions applicable to each Creative Commons user license:

CC BY: The CC-BY license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article and to make commercial use of the Article (including reuse and/or resale of the Article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by/4.0.

CC BY NC SA: The CC BY-NC-SA license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article, provided this is not done for commercial purposes, and that the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. Further, any new works must be made available on the same conditions. The full details of the license are available at http://creativecommons.org/licenses/by-nc-sa/4.0.

CC BY NC ND: The CC BY-NC-ND license allows users to copy and distribute the Article, provided this is not done for commercial purposes and further does not permit distribution of the Article if it is changed or edited in any way, and provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, and that the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by-nc-nd/4.0. Any commercial reuse of Open Access articles published with a CC BY NC SA or CC BY NC ND license requires permission from Elsevier and will be subject to a fee.

Commercial reuse includes:

- Associating advertising with the full text of the Article
- Charging fees for document delivery or access
- Article aggregation
- Systematic distribution via e-mail lists or share buttons

Posting or linking by commercial companies for use by customers of those companies.

20. **Other Conditions**:

v1.10

Questions? customercare@copyright.com.

ELSEVIER LICENSE
TERMS AND CONDITIONS

Jun 10, 2024

This Agreement between Oisin Shaw, Imperial College London ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	5805290537985
License date	Jun 10, 2024
Licensed Content Publisher	Elsevier
Licensed Content Publication	Journal of Power Sources
Licensed Content Title	PEM stack test and analysis in a power system at operational load via ac impedance
Licensed Content Author	Wenhua H. Zhu,Robert U. Payne,Bruce J. Tatarchuk
Licensed Content Date	May 25, 2007
Licensed Content Volume	168
Licensed Content Issue	1
Licensed Content Pages	7
Start Page	211
End Page	217
Type of Use	reuse in a thesis/dissertation

Portion	figures/tables/illustrations
Number of figures/tables/illustrations	1
Format	electronic
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title of new work	The Effect of Micro and Nano Scale Phenomena on System Level Design and Engineering of Hybrid PEMFC Systems
Institution name	Imperial College London
Expected presentation date	Jul 2024
Portions	Fig 1
The Requesting Person / Organization to Appear on the License	Oisin Shaw, Imperial College London
Requestor Location	Imperial College London Imperial College London Exhibition Road London, County SW7 2AZ United Kingdom Attn: Imperial College London
Publisher Tax ID	GB 494 6272 12
Total	0.00 USD
Terms and Conditions	

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and

conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your RightsLink account and that are available at any time at https://myaccount.copyright.com).

GENERAL TERMS

2. Elsevier hereby grants you permission to reproduce the aforementioned material subject to the terms and conditions indicated.

3. Acknowledgement: If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source, permission must also be sought from that source. If such permission is not obtained then that material may not be included in your publication/copies. Suitable acknowledgement to the source must be made, either as a footnote or in a reference list at the end of your publication, as follows:

"Reprinted from Publication title, Vol /edition number, Author(s), Title of article / title of chapter, Pages No., Copyright (Year), with permission from Elsevier [OR APPLICABLE SOCIETY COPYRIGHT OWNER]." Also Lancet special credit - "Reprinted from The Lancet, Vol. number, Author(s), Title of article, Pages No., Copyright (Year), with permission from Elsevier."

4. Reproduction of this material is confined to the purpose and/or media for which permission is hereby given. The material may not be reproduced or used in any other way, including use in combination with an artificial intelligence tool (including to train an algorithm, test, process, analyse, generate output and/or develop any form of artificial intelligence tool), or to create any derivative work and/or service (including resulting from the use of artificial intelligence tools).

5. Altering/Modifying Material: Not Permitted. However figures and illustrations may be altered/adapted minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of Elsevier Ltd. (Please contact Elsevier's permissions helpdesk here). No modifications can be made to any Lancet figures/tables and they must be reproduced in full.

6. If the permission fee for the requested use of our material is waived in this instance, please be advised that your future requests for Elsevier materials may attract a fee.

7. Reservation of Rights: Publisher reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

8. License Contingent Upon Payment: While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by publisher or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.

9. Warranties: Publisher makes no representations or warranties with respect to the licensed material.

10. Indemnity: You hereby indemnify and agree to hold harmless publisher and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

11. No Transfer of License: This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.

12. No Amendment Except in Writing: This license may not be amended except in a writing signed by both parties (or, in the case of publisher, by CCC on publisher's behalf).

13. Objection to Contrary Terms: Publisher hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and publisher (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

14. Revocation: Elsevier or Copyright Clearance Center may deny the permissions described in this License at their sole discretion, for any reason or no reason, with a full refund payable to you. Notice of such denial will be made using the contact information provided by you. Failure to receive such notice will not alter or invalidate the denial. In no event will Elsevier or Copyright Clearance Center be responsible or liable for any costs, expenses or damage incurred by you as a result of a denial of your permission request, other than a refund of the amount(s) paid by you to Elsevier and/or Copyright Clearance Center for denied permissions.

LIMITED LICENSE

The following terms and conditions apply only to specific license types:

15. **Translation**: This permission is granted for non-exclusive world **English** rights only unless your license was granted for translation rights. If you licensed translation rights you may only translate this content into the languages you requested. A professional translator must perform all translations and reproduce the content word for word preserving the integrity of the article.

16. **Posting licensed content on any Website**: The following terms and conditions apply as follows: Licensing material from an Elsevier journal: All content posted to the web site must maintain the copyright information line on the bottom of each image; A hyper-text must be included to the Homepage of the journal from which you are licensing at http://www.sciencedirect.com/science/journal/xxxxx or the Elsevier homepage for books at http://www.elsevier.com; Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

Licensing material from an Elsevier book: A hyper-text link must be included to the Elsevier homepage at http://www.elsevier.com . All content posted to the web site must maintain the copyright information line on the bottom of each image.

Posting licensed content on Electronic reserve: In addition to the above the following clauses are applicable: The web site must be password-protected and made available only to bona fide students registered on a relevant course. This permission is granted for 1 year only. You may obtain a new license for future website posting.

17. **For journal authors:** the following clauses are applicable in addition to the above:

Preprints:

A preprint is an author's own write-up of research results and analysis, it has not been peer-reviewed, nor has it had any other value added to it by a publisher (such as formatting, copyright, technical enhancement etc.).

Authors can share their preprints anywhere at any time. Preprints should not be added to or enhanced in any way in order to appear more like, or to substitute for, the final versions of articles however authors can update their preprints on arXiv or RePEc with their Accepted Author Manuscript (see below).

If accepted for publication, we encourage authors to link from the preprint to their formal publication via its DOI. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help users to find, access, cite and use the best available version. Please note that Cell Press, The Lancet and some society-owned have different preprint policies. Information on these policies is available on the journal homepage.

Accepted Author Manuscripts: An accepted author manuscript is the manuscript of an article that has been accepted for publication and which typically includes author-incorporated changes suggested during submission, peer review and editor-author communications.

Authors can share their accepted author manuscript:

- immediately
  - via their non-commercial person homepage or blog
  - by updating a preprint in arXiv or RePEc with the accepted manuscript
  - via their research institute or institutional repository for internal institutional uses or as part of an invitation-only research collaboration work-group
  - directly by providing copies to their students or to research collaborators for their personal use
  - for private scholarly sharing as part of an invitation-only work group on commercial sites with which Elsevier has an agreement
- After the embargo period
  - via non-commercial hosting platforms such as their institutional repository
  - via commercial sites with which Elsevier has an agreement

In all cases accepted manuscripts should:

- link to the formal publication via its DOI
- bear a CC-BY-NC-ND license - this is easy to do
- if aggregated with other manuscripts, for example in a repository or other site, be shared in alignment with our hosting policy not be added to or enhanced in any way to appear more like, or to substitute for, the published journal article.

Published journal article (JPA): A published journal article (PJA) is the definitive final record of published research that appears or will appear in the journal and embodies all value-adding publishing activities including peer review co-ordination, copy-editing, formatting, (if relevant) pagination and online enrichment.

Policies for sharing publishing journal articles differ for subscription and gold open access articles:

Subscription Articles: If you are an author, please share a link to your article rather than the full-text. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help your users to find, access, cite, and use the best available version.

Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

If you are affiliated with a library that subscribes to ScienceDirect you have additional private sharing rights for others' research accessed under that agreement. This includes use for classroom teaching and internal training at the institution (including use in course packs and courseware programs), and inclusion of the article for grant funding purposes.

Gold Open Access Articles: May be shared according to the author-selected end-user license and should contain a CrossMark logo, the end user license, and a DOI link to the formal publication on ScienceDirect.

Please refer to Elsevier's posting policy for further information.

18. **For book authors** the following clauses are applicable in addition to the above: Authors are permitted to place a brief summary of their work online only. You are not allowed to download and post the published electronic version of your chapter, nor may you scan the printed edition to create an electronic version. **Posting to a repository:** Authors are permitted to post a summary of their chapter only in their institution's repository.

19. **Thesis/Dissertation**: If your license is for use in a thesis/dissertation your thesis may be submitted to your institution in either print or electronic form. Should your thesis be published commercially, please reapply for permission. These requirements include permission for the Library and Archives of Canada to supply single copies, on demand, of the complete thesis and include permission for Proquest/UMI to supply single copies, on demand, of the complete thesis. Should your thesis be published commercially, please reapply for permission. Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

Elsevier Open Access Terms and Conditions

You can publish open access with Elsevier in hundreds of open access journals or in nearly 2000 established subscription journals that support open access publishing. Permitted third party re-use of these open access articles is defined by the author's choice of Creative Commons user license. See our open access license policy for more information.

Terms & Conditions applicable to all Open Access articles published with Elsevier:

Any reuse of the article must not represent the author as endorsing the adaptation of the article nor should the article be modified in such a way as to damage the author's honour or reputation. If any changes have been made, such changes must be clearly indicated.

The author(s) must be appropriately credited and we ask that you include the end user license and a DOI link to the formal publication on ScienceDirect.

If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source it is the responsibility of the user to ensure their reuse complies with the terms and conditions determined by the rights holder.

Additional Terms & Conditions applicable to each Creative Commons user license:

CC BY: The CC-BY license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article and to make commercial use of the Article (including reuse and/or resale of the Article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by/4.0.

CC BY NC SA: The CC BY-NC-SA license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article, provided this is not done for commercial purposes, and that the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. Further, any new works must be made available on the same conditions. The full details of the license are available at http://creativecommons.org/licenses/by-nc-sa/4.0.

CC BY NC ND: The CC BY-NC-ND license allows users to copy and distribute the Article, provided this is not done for commercial purposes and further does not permit distribution of the Article if it is changed or edited in any way, and provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, and that the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by-nc-nd/4.0. Any commercial reuse of Open Access articles published with a CC BY NC SA or CC BY NC ND license requires permission from Elsevier and will be subject to a fee.

Commercial reuse includes:

- Associating advertising with the full text of the Article
- Charging fees for document delivery or access
- Article aggregation
- Systematic distribution via e-mail lists or share buttons

Posting or linking by commercial companies for use by customers of those companies.

20. **Other Conditions**:

v1.10

Questions? customercare@copyright.com.