Control of Active Cell Balancing Systems

Innovation Report

Thomas Michael Bruen

EngD (International)

WMG, University of Warwick

September 2017

Abstract

Lithium-ion battery packs are increasingly being used for high power and energy applications such as electric vehicles and grid storage. These battery packs consist of many individual cells connected in series and/or parallel. Manufacturing tolerances and varied operating conditions mean that each cell will be different one from another, being able to store different amounts of energy and deliver different amounts of power. This also means some cells will finish charging or discharging before others, resulting in unutilised energy in the remaining cells.

Passive balancing systems are often used in multi-cell battery packs to ensure that all of the cells can be fully charged. However, this does not account for differences in cell capacity, meaning that not all cells will be fully discharged. Active balancing systems have been developed to transfer energy between the cells, in theory allowing for stronger cells to compensate for weaker ones. However, their perceived cost and complexity have prevented them from being widely adopted in commercial applications.

In this work, an innovative control strategy was developed to determine how and when to energy balance a set of battery cells, with the aim of maximising battery pack energy utilisation. A model-based control system was designed, using state of charge to evaluate the level of energy imbalance between cells. Real-time implementation using second-hand electric vehicle cells and commercial balancing hardware demonstrated that the control strategy can decrease the amount of unused charge in the battery pack from 8% with passive balancing to 1% with active balancing, which has significant impact for battery pack energy throughput, physical size, mass, and long-term health.

Acknowledgments

Firstly, I'd like to thank James Marco for supervising this research. Your advice, regular feedback and, most importantly, enforcement of deadlines made this work possible.

Thanks also to Paul Haney, Miguel Gama, Sue Slater and everyone in the JLR team for their support, opportunities and experience. I hope you can make the most of this research. I'd like to acknowledge the EPSRC and JLR for funding the project. The wider WMG centre HVM Catapult (refrerred to throughout this report as Catapult) has been instrumental in gathering the large volumes of test data used throughout this work.

I'm grateful to everyone who put aside time for technical discussions, notably Dhammika Widanalage and Mark Tucker, and to Andrew Moore and John Palmer for their assistance with laboratory work.

To my friends and family: thanks for your encouragement and distraction this past four years.

Finally, and most importantly, Claudia, I couldn't have done this without you.

Declarations

I declare that the work in this document is my own unless otherwise stated.

Thomas Bruen, September 2017

Table of Contents

List of Figures

List of Tables

List of Abbreviations

Abbreviation	Full term
ABS	Active Balancing System
ADC	Analogue-Digital Converter
BEV	Battery Electric Vehicle
BMS	Battery Management System
DOD	Depth of Discharge
DSIC	Dynamic Steady-Input Control
ECM	Equivalent Circuit Model
EIS	Electrochemical Impedance Spectroscopy
EOC	End of Charging
EOD	End of Discharging
EV	Electric Vehicle
GITT	Galvanostatic Intermittent Titration Technique
HEV	Hybrid Electric Vehicle
LQR	Linear-Quadratic Regulator
JLR	Jaguar Land Rover
MPC	Model Predictive Control
OCV	Open Circuit Voltage
PBM	Physics Based Model
PCB	Printed Circuit Board
PHEV	Plug-in Hybrid Electric Vehicle
PWM	Pulse Width Modulation
RC	Resistor Capacitor
SIL	Software-in-Loop
SOC	State of Charge
SOE	State of Energy
SOH	State of Health
TMETC	Tata Motors European Technical Centre

Mathematical Notation

Notation	Description	Units
$\boldsymbol{A}$	State-space system state matrix	-
$\boldsymbol{B}$	State-space system input matrix	(Ah^{-1})
$\boldsymbol{B_c}$	State-space system controlled input matrix	(Ah^{-1})
$\boldsymbol{B_d}$	State-space system disturbance input matrix	(Ah^{-1})
$\boldsymbol{B_1}$	State-space augmented system input matrix	(Ah^{-1})
$\boldsymbol{b_2}$	State-space augmented system disturbance vector	(Ah^{-1})
$\overline{\boldsymbol{B}}_{\boldsymbol{1}}$	State-reduced state-space augmented system input matrix	-
$\overline{\boldsymbol{b}}_{\boldsymbol{2}}$	State-reduced state-space augmented system disturbance vector	-
$\boldsymbol{C}$	State-space system output matrix	-
$\boldsymbol{C_p}$	Cell polarization capacitance, as part of an impedance model	Farads (F)
$\boldsymbol{d}$	Disturbance (uncontrolled) input to a system	-
$\boldsymbol{\delta}$	The error between a measured and simulated/estimated value	-
$\boldsymbol{e}$	The error between a measured and simulated/estimated value	-
$\boldsymbol{E_{ps}}$	Parallel cell current calculation state matrix	-
$\boldsymbol{\eta_{chg}}$	The efficiency of a balancing system for a charging current	-
$\boldsymbol{\eta_{dch}}$	The efficiency of a balancing system for a discharging current	-
$\boldsymbol{f}$	Cost function scalar term	-
$\boldsymbol{f_{ps}}$	Parallel cell current calculation input vector	Ohms (Ω)
$\boldsymbol{F}$	Balancing current coupling matrix	-
$\boldsymbol{g}$	Cost function linear vector	-
$\boldsymbol{G_d}$	Disturbance-output transfer function	-
$\boldsymbol{G_{ps}}$	Parallel cell current calculation current matrix	-
$\boldsymbol{G_{tf}}$	Input-output transfer function	-
$\boldsymbol{H}$	Cost function quadratic matrix	-

i	An arbitrary current	Amperes (A)
i_{app}	The current applied to a battery pack	Amperes (A)
i_{bal}	The current generated by balancing	Amperes (A)
i_{cell}	The current through a cell	Amperes (A)
i_{fb}	The feedback current generated through active cell balancing	Amperes (A)
i_{loop}	The current within an electrical circuit mesh	Amperes (A)
i_{norm}	The current through a cell, normalised against applied current	%
I	Identity matrix	-
j	Angular frequency	$Rads^{-1}$
J	Optimization cost function value	-
K	State feedback gain matrix	-
L	Kalman gain matrix	-
lb	Optimization lower inequality constraint vector	-
m	Difference of cell SOC from the battery pack mean	%
M	Balancing s-to-m transfer matrix	-
N_m	The number of cells in series in a module	-
N_s	The number of states in a system	-
$ones$	A Matrix composed of elements of the value 1	-
q_{cap}	Cell charge capacity	Ampere-hours (Ah)
q_{bal}	Balancing charge	Coulombs
R_{int}	Cell internal resistance, found by applying a 10s current pulse to the cell	Ohms (Ω)
R_{DC}	Cell internal resistance, at the point where the complex component of impedance is zero	Ohms (Ω)
R_{conn}	Connection resistance between cells	Ohms (Ω)
R_0	Cell internal resistance, as part of an impedance model	Ohms (Ω)

R_p	Cell polarization resistance, as part of an impedance model	Ohms (Ω)
R_{ps}	Parallel cell current calculation resistance matrix	Ohms (Ω)
s	Cell SOC	%
SOH_E	State of Health with respect to energy	%
SOH_P	State of Health with respect to power	%
SOC	State of charge	%
t	time	Seconds (s)
τ_p	impedance time constant	Seconds (s)
T	m-to-z transfer matrix	-
u	Controlled input to a system	-
ub	Optimization upper inequality constraint vector	-
UF	Battery pack energy utilisation factor	-
v_{oc}	Cell open circuit voltage	Volts (V)
v_p	Cell polarization voltage, for an impedance model	Volts (V)
v_t	Cell terminal voltage	Volts (V)
W_r	The winding ratio of a transformer	-
v	Generic dynamic system process noise vector	-
w	Generic dynamic system input noise vector	-
x	Generic dynamic system state vector	-
y	Generic dynamic system output vector	-
z	Balancing control system state vector	-
ω	Angular Frequency	$Rads^{-1}$

1 Introduction

Lithium-ion cells are increasingly being used for a variety of technologies, ranging from smartphones through to electric vehicles (EVs) and grid storage for renewable energy sources [1]. They offer around twice the energy and power density than nickel metal-hydride cells, and about four times that of lead-acid cells, opening up battery power to applications which were not feasible until recently. Lithium-ion cells can be comprised of a variety of different chemistries depending on requirements such as energy, power, lifetime and cost [2]. The main thing these cells all have in common is that the voltage is quite low. Generally when a cell is fully charged it will be circa 4V, and when it's fully discharged it will be slightly under 3V. For some chemistries, the operating range is even lower. This is insufficient for an EV whose powertrain requires 300V or above [3, 4], and closer to 1000V for Formula E racing cars [5]. To reach these sorts of voltages, hundreds of cells need to be connected in series. Furthermore, cells connected in parallel may be required in order to store more energy or increase the power capability of the battery system. For example the 2016 Tesla Model S 100KWh battery pack contains 8256 cells [6]. This introduces a number of problems:

- All of the cells need to be packaged, electrically connected, monitored, thermally regulated and controlled. The more cells there are, the greater the cost, complexity and weight.
- Manufacturing tolerances and varying operating conditions mean that each of these cells will be slightly different. The amount of energy each one can store will differ, as will the amount of heat generated internally and the amount of charge lost over time through self-discharge.

Multi-cell battery packs can become quite complex and require a battery management system (BMS) to control them. This uses sensor information such as cell voltages, currents and temperatures to ensure the battery pack is in a safe state, as well as predict the cells' energy and degradation levels.

For cells connected in series, each cell is charged or discharged with the same current, meaning that a lower capacity cell will complete charging or discharging

before a higher capacity one. Conventionally there is no control over how individual cells are loaded which means when one cell is fully charged or discharged, the entire pack has to cease operation. Different combinations of capacities, and variations in energy levels, means that most cells will not be fully charged or fully discharged, resulting in unused areas of energy as shown in grey in Figure 1. The height of the cell represents how much energy it can store, and the alignment shows the differences in energy levels: at the point in time where one cell has 100Wh remaining, another cell may only have 80Wh remaining and as such is closer to the end of discharge (EOD). In other words, the battery pack is not storing as much energy as it could, and the energy it does store is not being completely used. If this battery pack could fully utilise its energy, it could be made smaller without compromising EV performance.

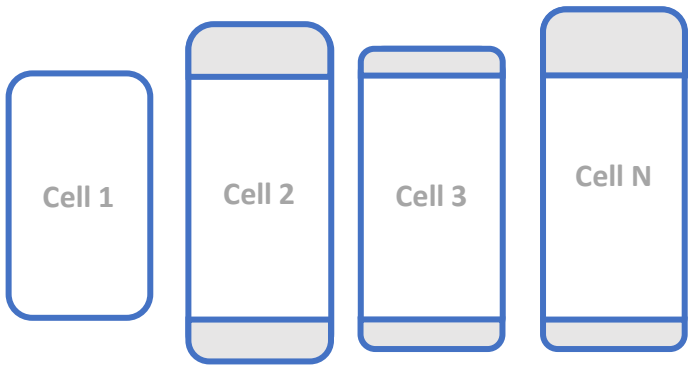

Figure 1: Unavailable energy owing to imbalance in series connected cells

Cells at different energy levels are described as “imbalanced”, and balancing systems have been developed to address this issue. Prior to li-ion technology this was straightforward, as the cells could be overcharged: the current was applied until the last cell reached the maximum voltage, and the other cells would settle at this maximum. However, it is dangerous to over-charge li-ion cells and so another means of balancing must be sought. The simplest way is to apply passive balancing: connecting a resistor over the cell. This will gradually discharge the cell according to Ohm’s law. Cells at higher energy levels can be discharged down to the same level as the cell with the lowest energy, and then charging can be completed. This simplicity comes with a number of drawbacks:

- The energy from passive balancing is wasted as heat dissipated by the resistor. This heat also has to be removed from the battery pack.

- The balancing current is generally low in order to limit the heat generation. This means that balancing can take a long time, increasing the overall charging time of the battery pack.
- Even if the cells are all completely balanced by end of charge (EOC), the differences in capacity mean that some cells will reach EOD before others. The best case for passive balancing is that all of the energy of the weakest cell is utilised.

Active balancing is an alternative strategy which could address the downsides to passive balancing by using power electronics to move energy between cells. If one cell is going to reach EOD before another, it can be charged up by the other cells and similarly if one cell is going to reach EOC before the others, it can discharge into them. The energy is utilised by other cells rather than being dissipated as heat. While this can potentially offer significant performance benefits, there are some downsides. The cost is higher, and the electronics takes up extra space and adds new failure modes to the battery pack [7].

For these reasons, active balancing systems (ABSs) have not, to date, been adopted for any mass-production vehicle applications. However, Jaguar Land Rover (JLR), with whom this EngD is in partnership with, has not looked in-depth at the current state of active balancing systems, and more generally there is a lack analysis of active balancing systems within the context of a complete battery pack. The Research Problem that this EngD addresses is: **while active cell balancing technology is available, there is little understanding of, or agreement on, how it should be used, what specific problems it is addressing, and how much benefit it can provide within a future EV application.**

This research problem is broken down further into Research Questions and Research Objectives in the next section.

1.1 Research Questions and Objectives

Several open research topics in the area of imbalance were identified after a preliminary analysis of the literature, including:

- What are the fundamental causes of cell-to-cell variation, and how can they be reduced?
- Using large datasets from aged EV battery packs, can statistical models be generated which accurately predict the spread of imbalance over time for different applications and operating conditions?
- How can balancing hardware be further improved, particularly regarding scaling up to a high voltage battery packs containing hundreds or thousands of cells?
- Can cost and size reductions be made to make load-regulating smart cells feasible for near-future battery pack integration?

While these questions are apposite, the major observation from the literature review was that there does not appear to have been a thoroughly designed balancing control system, which factors in a high-level strategy, a low-level energy management scheme and adaptability to the many types of balancing hardware. The lack of a comprehensive active balancing strategy means that the potential benefits, such as increased energy output, have not been fully realised and so active balancing appears less cost effective than it could otherwise be. This EngD therefore aims to address the following Research Questions:

1. **In what context can active balancing provide the greatest benefit?**
2. **What kind of gains in energy utilisation can be achieved using active balancing, compared to passive balancing?**

Question 1 considers the higher-level aspects of active balancing, such as if it is only suited to certain applications and whether parallel cells also require balancing. Question 2 uses the information from Question 1 to formulate a control strategy to maximise energy utilisation, which can be verified experimentally to quantify its performance. To achieve this, the following objectives were defined:

1. **Define the sources of imbalance within a battery pack and the affect each of these has on overall battery pack performance.**

2. **Understand the current status of active balancing hardware and evaluate their suitability for different applications.**
3. **Create a modelling framework to analyse imbalance and aid in designing a control system.**
4. **Develop a generic balancing control system with the specific goal of maximising the energy utilisation of the battery pack.**
5. **Implement the control system in real-time with specific balancing hardware and evaluate its performance, including the gain in energy utilisation.**

Figure 2 presents the relationship between the two Research Questions and the Research Objectives. In addition, Figure 2 highlights the relationship between each Research Question and the corresponding Portfolio Submission and Chapter within this Innovation Report. The structure of the Portfolio and Innovation Report are further explained in sections 1.2 and 1.3 respectively.

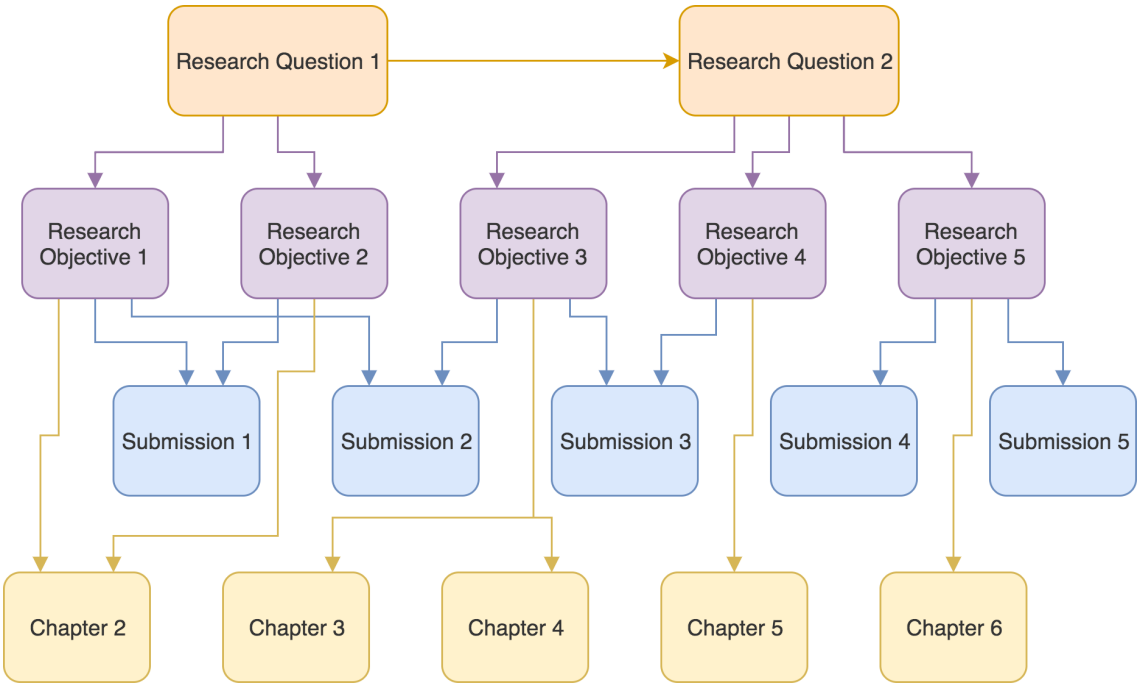

Figure 2: Structure of the EngD Portfolio and Innovation Report

1.2 Structure of the Portfolio

As well as the personal profile and Innovation Report, there are 5 main Submissions in the portfolio, plus a Submission covering the international placement. Submission 1 contains literature reviews of cell-to-cell variation and

balancing systems, and some initial modelling of balancing systems. An IFAC conference paper was produced from this work, concerning high-level modelling and control of balancing systems [8]. Submission 2 documented the modelling and experimental work on cells connected in parallel. From this, a journal paper and conference paper were published [9, 10]. This work was also combined with some further modelling work related to [8], to publish another journal paper on multi-cell BMS applications [11] Submission 3 introduced the control model and used this to design and critically assess various options for feedback design. In Submission 4, the cells and balancing hardware used for experimental work was detailed, including the various iterations of the test set-up. The commissioning tests performed to analyse the system were documented. In Submission 5, experimental balancing tests were performed. Two of the controllers from Submission 3 were developed further, and state of charge (SOC) estimators were also designed to use as part of the control system. The controllers were run in real-time and their performance evaluated. The research Submissions should be read in their order of writing, as they follow a simplified systems "V" diagram as shown in Figure 3. The literature review in Submission 1 helped outline the performance criteria which would ultimately be validated via the experimental work in Submission 5.

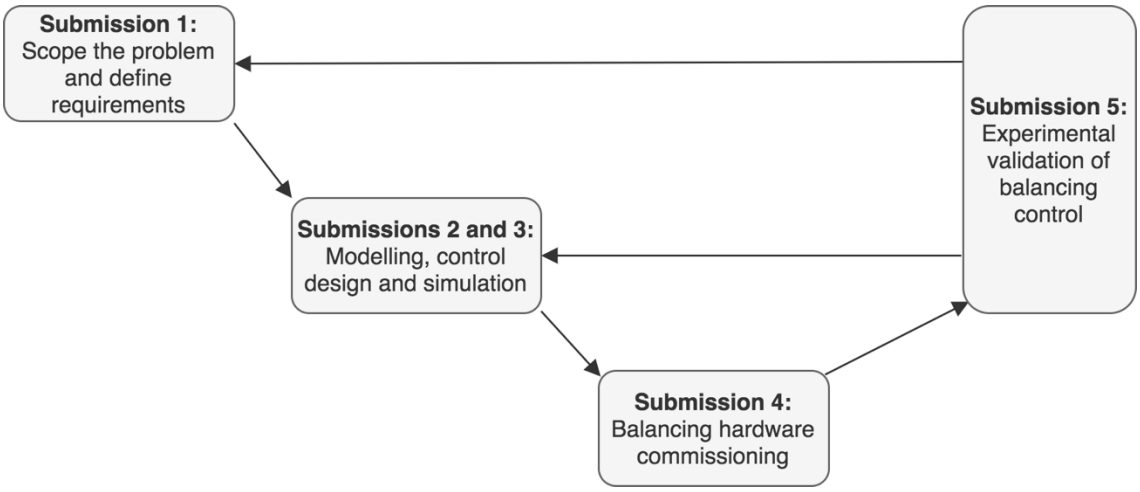

Figure 3: Portfolio Submissions as a systems "V" diagram

1.3 Structure of the Innovation Report

The Innovation Report is comprised of five main Chapters. Chapter 2 contains two literature reviews. The first is on cell-to-cell variations and the implications

this has on battery pack performance. The second is a critical survey of the active balancing systems proposed within the literature, along with discussion on future trends on cell and balancing technology. Chapter 3 introduces cell models, including the various types of model and the data required to parameterise and validate the models. These are then applied in Chapter 4 to show how models were used in three main areas: the implications of connecting cells in parallel, connecting cells in series as a basis for running simulations of ABSs, and finally SOC estimation. Chapter 5 describes the control of ABSs. This covers higher-level aspects such as a strategy for when to balance, and specifics of control algorithms. A control model is introduced and used to design two control systems, which are simulated for a generic ABS. The balancing hardware used for experimental work is detailed in Chapter 6. The theoretical operation is described along with the practicalities of instrumenting and operating the equipment. The control systems introduced in Chapter 5 are implemented in real-time, using the SOC estimators from Chapter 4, and analysed. Finally, conclusions are drawn in Chapter 7, including areas of future work emerging from the results of this Innovation Report.

2 A Review of Imbalance and Balancing

Two distinct literature reviews were conducted to understand the many aspects and research areas covered by energy imbalance. The first literature review, in section 2.1, addresses the first Research Objective of understanding which aspects of cell performance a balancing system can address, and what the benefits will be for the overall battery pack. It covers how cell-to-cell variations arise and what this means for battery pack performance, followed by an analysis of the impact each of these variations has on imbalance, and how imbalance should be defined. A second review was conducted in section 2.2 on the various means of transferring energy between cells, and existing control strategies for managing this energy transfer. This directly meets the second Research Objective of surveying the current status of balancing hardware. Conclusions are drawn in section 2.3.

2.1 Cell-to-Cell Variation

Each cell within a battery pack will have different properties to one another. These can be caused by intrinsic factors, for example limitations in the manufacturing process, or extrinsic factors such differences in operating conditions [12].

Fabricating cells is a complex, multi-step process [13, 14], and variations in properties arise from tolerances in these steps. Improving this process will reduce the size of the variations, which could reduce the amount of balancing required. However, it is not possible to remove variations completely, especially when there are often conflicting objectives such as the need to reduce the cost of manufacturing [14, 15]. In addition, new cell chemistries require different manufacturing techniques, and so there would likely be an increase in variation before these new processes are refined.

Screening the cells prior to battery pack assembly can ensure that the pack is comprised only of cells with similar impedances and capacities [16–18]. This is

beneficial for excluding particularly weak cells but, ageing is open loop[1] so even a very small initial imbalance will increase over time. There is also the growing second-life sector, in which aged cells are repurposed for other applications. For example, a cell which no longer has sufficient performance for its automotive application is still capable of delivering energy, and can be used as part of a grid storage system [19]. Nissan [20] and Renault [21] are both trialling energy storage units comprised of used automotive cells for homes with solar panels. This is essentially the opposite of screening: as the purpose is to make use of old cells, there will inherently be large variations in cell properties, and possibly even the types of cell.

Factors such as electrode thickness and porosity, and the amount of active material will affect cell properties [19, 22, 23]. There is generally a positive correlation between cell mass and cell capacity, as heavier cells should contain more active material [12, 24], but there is evidence that some variations in manufacturing affect mass but not capacity [25]. There is a lack of data on cell-to-cell variability, especially for automotive grade cells, but figures for the initial spread in capacity for cells from established manufacturers include 2-6% [26], 4% [25], 4.5% [19], 8% [27] down to about 1% after screening [28].

2.1.1 Cell Ageing

One of the key drivers of cell-to-cell variation over time is how the cells degrade (also known as ageing). Understanding and reducing cell degradation is one of the key challenges to overcome in the lithium-ion battery field. There are many causes of degradation, which are often complex and nonlinear in [29, 30]. Figure 4, reproduced from [30], shows how system-level inputs such as current and temperature cause a variety of chemical interactions which ultimately affect the two main systems-level metrics of cell degradation: capacity fade and power fade.

[1] There is some control over how cells age, such as the charging regime, specification of power limits and thermal management. However, there is typically not the ability to manage individual cells, for example to increase the cooling on a specific cell if it is weaker than the others.

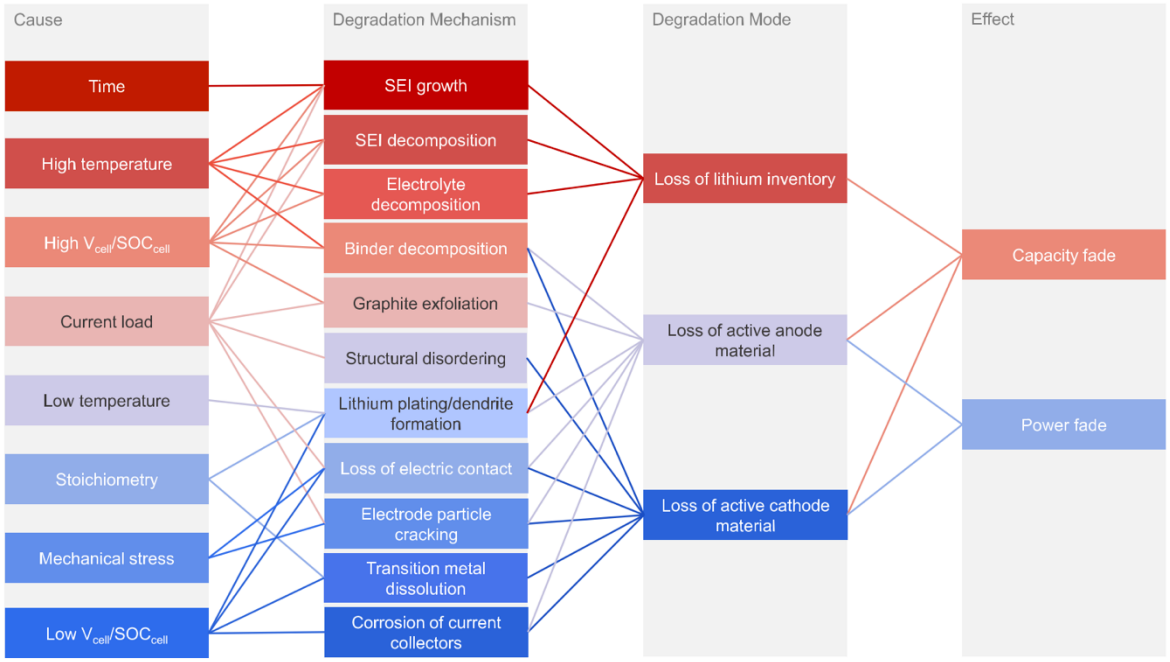

Figure 4: Lithium-ion cell ageing mechanisms- from [30]

Capacity fade reduces the energy availability of the cell and is caused by a reduction in the amount of lithium available for cycling and a reduction in the space in the electrodes to store it. Power fade is caused by an increase in impedance: more of the power generated by the cell will be dissipated as heat, meaning that the efficiency drops and the cell cannot deliver as much power to the wider system. This leads to two common definitions of state of health (SOH): a 0-100% of cell degradation relative to when the cells were new. The energy SOH is defined by (1), and the power SOC defined by (2). Q_{cap} is the cell capacity in ampere-hours, R_{int} is the internal resistance of the cell in ohms, and t_f is the present time. Capacity and resistance are defined and discussed further in Chapter 3.

$$SOH_E = 100\frac{q_{cap}^{t=t_f}}{q_{cap}^{t=0}} \quad (1)$$

$$SOH_P = 100\frac{R_{int}^{t=t_f}}{R_{int}^{t=0}} \quad (2)$$

For automotive applications, a cell's end of life (EOL) criteria are typically defined as when SOH_E drops below 80% [31–34] or when SOH_P rises above 200% [33, 34]. For Hybrid Electric Vehicles (HEVs), the battery pack provides an additional

source of power alongside an internal combustion engine system. The latter ensures that range requirements are met, and the battery helps deliver sufficient power, as well as absorbing power through regenerative braking [35]. The battery pack must therefore be capable of high power, rather than delivering a large amount of energy. Conversely, for battery electric vehicles (BEVs) the battery pack is the only power source, and the dominant requirement is being able to delivery large amounts of energy (driving range), while the large pack size means the power demand is relatively low. Plug-in-hybrid electric vehicles (PHEVs) lie between the two, where there is still an engine, but unlike a HEV, the vehicle can be driven using battery power alone, and recharged from the grid as well as the engine. For HEVs, the SOH_P criterion is more likely to be reached first, whereas SOH_E is the dominant factor for BEVs [33, 34]. Experimental results such as in [36, 37] demonstrate a correlation between SOH_E and SOH_P due to the overlap and interaction between the various ageing mechanisms. However, other studies show little correlation between resistance and capacity fade [38] which emphasises the complex nature of ageing, and the importance of analysing both SOH_P and SOH_E and using the appropriate metric depending on the application.

Generally, higher temperatures increase ageing [39], with a general relationship that the rate of ageing doubles every increase of 10°C [34]. When the cells are being cycled, cold temperatures also increase ageing via different mechanisms. Experimental results from [31] show that for cells under load, the primary cause of ageing below 25°C is lithium plating during charging, whereas degradation reactions and solid electrolyte interphase (SEI) growth occur above 25°C. High currents rates also increase ageing [40], with particularly high magnitudes inducing electrode fatigue and fracturing as the electrode expands and contracts during lithium insertion and removal [29]. High currents will also increase the cell temperature, potentially driving additional ageing as discussed above.

Depth of Discharge (DOD) is also a factor in ageing. A large DOD (wide SOC window) is generally considered to accelerate the ageing process. Therefore a BEV pack will most likely age quicker than a HEV pack even for the same energy throughput [41], although this is dependent on other factors such as current

magnitude and temperature, so the BMS strategy may work to limit ageing in other ways (e.g. more conservative power limits). However, grid charging of BEVs offers more opportunity for cell balancing rather than a HEV, so cycle life may be improved. This DOD relationship is disputed by Peterson et al. [42] and Choi et al. [43] The latter compared different DODs but all started at 100% SOC rather than being centred about 50% SOC, so the tendency for a high SOC to accelerate ageing may factor into their results. The former results are based on vehicle drive cycles rather than typical laboratory cycling so the effects of the current profile on ageing may be tied to DOD ageing. This suggests that results from laboratory testing do not necessarily extrapolate to real-world vehicle usage. A higher mean SOC also tends to accelerate ageing due to the increased difference in potential between the electrodes and electrolyte [34, 39].

2.1.2 Battery pack ageing

While the causes of cell degradation are useful to understand, it is the cell-to-cell variation which is most relevant to balancing: even if the values are all poor from a performance point of view (e.g. a very high self-discharge rate), if all cells have the same value then the cells will still be balanced. One of the lesser studied areas of particular relevance for balancing high-voltage battery packs is a statistical analysis of how large numbers of cells age.

Baumhofer et al cycle-aged 48 pre-screened commercial Li-ion cells under the same operating conditions (current profile and ambient temperature) and measured cell capacity periodically [28]. Not only did capacity difference amongst the cells increase from 1% to 8%, there was no correlation between the initial and aged capacities: one of the lowest capacity cells when new ended up as one of the highest capacity cells after ageing, and vice versa (a result also observed in [44]). This means that screening is not a reliable indicator of cell performance long-term. The authors also found there was little correlation between relative resistance and capacity when new [45], which means that screening for capacity might result in a wide variation in impedance, and vice versa.

Schuster et al. [19] performed capacity and impedance tests on a large number of 1.95Ah automotive cylindrical cells. 484 new cells were used as a baseline to

compare to 1908 aged cells from BEVs, taken from two different vehicles both driven regularly for three years. One important finding, reproduced in Figure 5, shows how the distribution of cell capacity changes with ageing. The distribution of the aged cells is noticeably skewed compared to the approximately normal distribution of the new cells. This means that outliers are much more likely to have a lower SOH than the mean (their results applied for both resistance and capacity). **As battery pack performance is limited by the weakest cell, there will be a greater percentage of unused energy remaining in an aged pack than a new one**. The distributions between the two sets of aged cells is also different. The cells from BEV 1 appear to have a bimodal distribution, with two peaks in the histogram, whereas BEV 2 has a more defined single peak. The authors note that BEV 2 was driven further and more aggressively, covering 26,745 km and consuming 19.0kWh/100km, compared 21397km and 16.9kWh/100km for BEV 1. This emphasises that long-term battery pack performance needs to be understood from a statistical point of view, as well as the physical basis behind the variation.

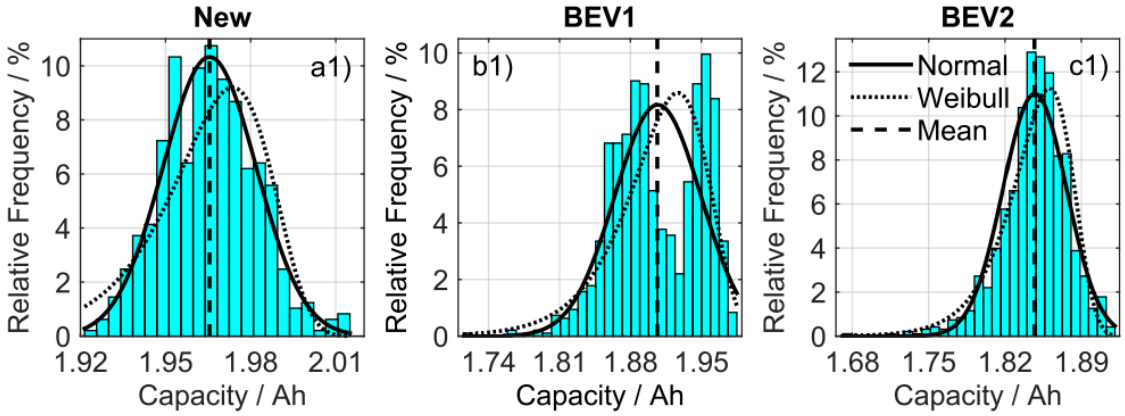

Figure 5: Capacity variation within battery packs - from [19]

An analysis of 20000 LFP cells [27] showed the new cell capacities had a standard deviation of 1.3%. A 96-cell module was then aged using a HEV drive cycle derived from customer driving data, which maintained the cells in a relatively narrow SOC window of 30-50%, with the cooling system active. After the ageing test, there was a mean of 18% capacity degradation, with maximum and minimum of 25% and 14% respectively. This suggests agreement with the skewed results in Figure 5: the least-aged cell was only 4% above the mean, but most-aged cell was 7% below.

A Weibull distribution for cell capacity was also observed in [38], in which the authors disassembled a battery pack from a vehicle driven for 32500km over 3 years in a city in southern China. The 95s5p (ninety-five units in series, each unit comprising five cells connected in parallel) battery pack had a nominal cell capacity of 12Ah, though the authors considered each parallel stack of series as a single cell with a 60Ah nominal capacity. The mean SOH_E was 82.4%, with the maximum 4.4% above the mean and the minimum 8.7% below. In [46], an EV battery pack comprised of 50Ah pouch cell was disassembled after 30000km of driving. The authors found that the SOH_E varied from 91 to 98% and, as Figure 6 shows, there was significant variation within each 10-cell module, as well as between each module. This shows that cell balancing just within modules is insufficient: reducing imbalance within module 13 has no impact on total driving range since module 11 would still limit battery pack operation.

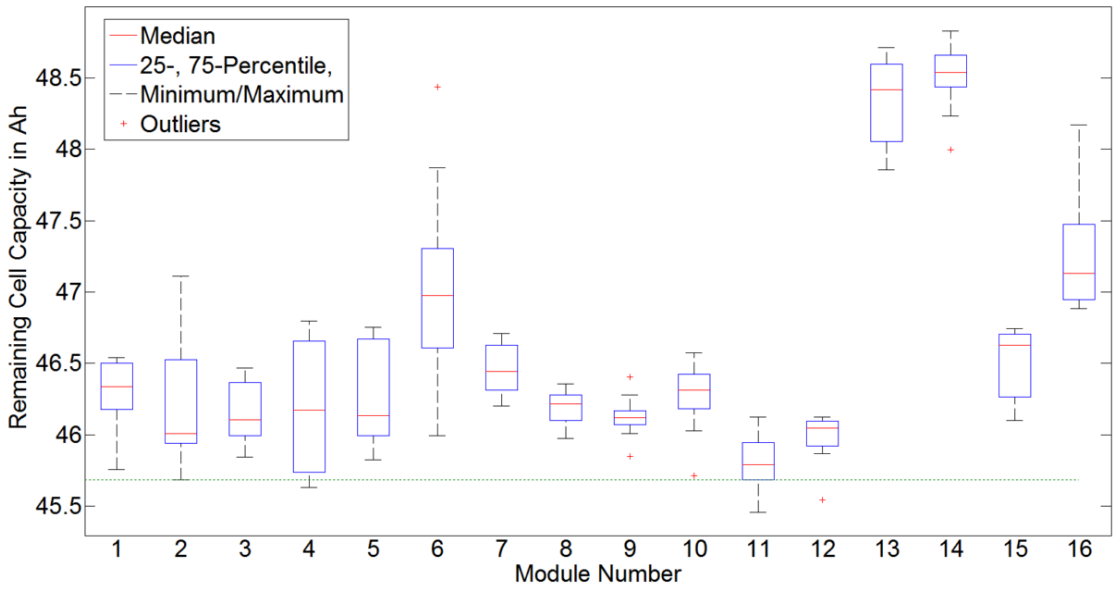

Figure 6: Variations in cell ageing within EV battery modules – from [46]

In contrast, Campestrini et al. [47] found little variation within modules after ageing. Two 8s14p modules were built using commercial Panasonic cells, with a nominal capacity of 2.8Ah. The cells were screened by a capacity measurement prior to assembly, and passive balancing was employed. The modules were then cycled for 12 months using a standard US06 highway profile, which the authors claim equates to 120000km of real-world driving. Both modules were contained in thermal chambers, one with an ambient temperature of 25°C, the other at 40°C. After the cycle ageing, there was less than 1% difference in capacity within each

module. However, there are several caveats to these results. Work co-authored by the author [37] demonstrated that the SOHs of cells connected in parallel will converge over time (and as such, not diverge if they were similar to begin with). As such, each module can be considered to contain 8 units with potentially different SOHs. This is a low number from a statistical point of view, especially considering the importance of outliers on overall performance, as discussed above. The same repeated drive cycle is not representative of real-world operation, and compressing ageing into 12 months does not account for how cells degrade with time. There is no information on the thermal design of the pack, such as any cooling system or ventilation. Ultimately, the statistical data from [27] and [46] give a better indication of the imbalance expected from EV operation, as they are obtained from actual EV modules. However, the results in [47] are still significantly different to other laboratory testing such as in [28], which may suggest that manufacturing standards are improving.

There is a consensus within the literature above that differences in cell temperature is a key driver of imbalance. As well as changing the rates of self-discharge between cells, it also causes greater variations in cell ageing, setting up a larger amount of imbalance in the future.

2.1.3 Drivers of Imbalance

Section 2.1 discussed the ways in which variations between cells arise. These variations can be considered within the context of imbalance specifically. There are three main properties which vary between cells: self-discharge rate, impedance and capacity [17, 48]. These variations are best considered at the different time-scales in which they create imbalance.

2.1.3.1 Long Term: Self-discharge

Self-discharge is a reduction in stored charge in the cell owing to internal chemical side reactions [49]. Unlike cell ageing, self-discharge is a reversible process: re-charging the cell recovers this lost energy [50]. As well as from manufacturing tolerances, difference in self-discharge can arise through temperature gradients across the pack, with higher temperatures increasing the

rate of self-discharge [51–54] owing to the Arrhenius relationship [54]. Additionally, self-discharge is known to be a function of SOC [49].

Self-discharge can be modelled as a small current perpetually discharging the cell [55]. However, a physical current is not actually generated external to the cell and so is not measured by a BMS. Depending on battery pack hardware design, some of the electronics powered by the cells themselves (such as monitoring and safety circuitry) may draw different amounts of current, thus acting in a similar manner to self-discharge when the battery pack is active [51].

Differences in self-discharge rates between cells mean that over time the SOCs of cells will diverge even if they were previously balanced. One of the advantages of Li-ion cells over other chemistries such as nickel-metal-hydride is the low self-discharge [1, 56]. Figures quoted include 1-10% SOC per month [49], 0.5-8% SOC per month [48], 6-10% SOC per month at 40°C, and 3-5% per month [57]. There is very little data concerning the spread of self-discharge between cells, but the consensus is that it will only induce significant imbalance over many days, meaning that cell balancing to counter this only has to be performed every few weeks [48].

2.1.3.2 Short Term: Capacity

Cells which were balanced at the end of charging may well be imbalanced by the time the battery pack has finished discharging. This is because of variations in cell capacity. When the balancing system is not operating, all of the cells in a series string undergo the same current, meaning at any point in time the same amount of charge has been removed from each cell. The cell which can store the least amount of charge will therefore reach EOD first. At this point it is unsafe to continue discharging this cell, so the entire series string must cease operating. Because of this, the battery pack capacity is determined by the lowest cell capacity in the series string [48]. The energy available from the battery pack will depend on the voltage levels of the other cells [23], which is why it is beneficial to balance at EOC rather than EOD [7]: the average cell voltage will be higher, increasing the amount of energy available.

Differences in capacity can arise through manufacturing variations, and the spread increases with ageing [27, 28] The imbalance induced through capacity differences is reversible, in that once the cells are recharged they will return to approximately the same state of balance as when discharging began. If balancing was performed during discharging, it would then have to be performed again during charging, which wastes energy and potentially increases the time to complete charging if balancing is slow.

The impact of capacity on overall performance is dependent on the application. For example, a HEV typically maintains the cells within a narrow SOC range, e.g. 40-60% SOC [41, 58, 59]. Differences in capacity will not affect battery pack performance significantly, because there is scope to widen the DOD window slightly if necessary. However, a BEV application might require using almost the entire SOC window. The Tesla Model S reportedly utilises 96% of the total available energy [60]. In this event, a relatively small divergence in cell capacities could reduce the energy availability of the battery pack to less than is required. This also increases the chances of the driver running out of range: SOC estimates are subject to some uncertainty, and that becomes more apparent when there is only a few percent SOC remaining.

2.1.3.3 Instantaneous: Impedance

The voltage across a cell's terminals changes with the applied current. The greater the cell's impedance, the more the voltage will change for a given current magnitude. For example, two cells at rest with the same voltage appear to be balanced. But once a load is applied differences in impedance mean that they appear imbalanced. Removing this load and letting the voltages settle will make the cells appear balanced again (assuming only a small change in SOC has occurred or the cells have the same capacity). This makes voltage a difficult metric for assessing imbalance. However, it can be useful when considering operation near EOC and EOD. Cells have upper and lower voltage limits which must not be exceeded for safety reasons [58]. Near EOD, a cell with a higher impedance may hit the lower voltage limit even though it is not at the lowest SOC.

In this sense, the weakest cell (from a capacity point of view) is not currently limiting pack performance. Derating strategies are often required to limit currents in low and high SOC regions to avoid cells undergoing loads which cause them to exceed these voltage [7, 61].

As well as increasing with ageing, a cell's impedance is a function of temperature and SOC [58, 62]. A higher temperature reduces the impedance, improving efficiency at the expense of increasing ageing [63, 64]. Impedance also tends to increase at SOCs above 80-90%, and in particular at SOCs below 20-30% [64, 65]. An application which requires a wide DOD, such as a BEV, is likely to enter these regions. This reduces efficiency as more energy is dissipated as heat.

Cell voltage and impedance are at the core of imbalance within parallel connected cells. It is commonly assumed that energy balancing is only required for cells in series [17, 66] as the cells in a parallel unit are inherently balanced through being at the same terminal voltage [58, 67]. However, until recently there was little experimental data to explore this further. Following the work published by the researcher [10, 37], detailed in section 4.1, there has been more research into parallel connected cells, generally looking at the impact of temperature variations on parallel currents, the implications for the BMS, and better understanding the underlying cell properties which cause different currents [68–71].

2.1.4 Quantifying imbalance

The literature review above has made clear that cell-to-cell variations have implications for overall battery pack performance, and a balancing system can mitigate against several aspects. In order to do this, imbalance must be quantified as a useable metric. The BMS can then analyse the imbalance status of each cell and decide on the nature of the corrective action to be taken.

2.1.4.1 Imbalance Metric

From a review of battery system operation, there are five potential measures which can form the basis of imbalance quantification:

- Voltage

- Charge
- Energy
- SOC
- State Of Energy (SOE)

Charge level is perhaps the most logical reference point to evaluate imbalance. As discussed in section 2.1.3.2, imbalance arises during operation because cells in series can store different amounts of charge, but all undergo the same current. However, there are some issues with using charge. Firstly, it is an estimate, although this is true for all measures apart from voltage. A poor estimate means that the perceived imbalance is incorrect, and implementing the balancing strategy could lead to more imbalance. Secondly, as it is an absolute value, there is only one reference point. When all cells are fully discharged, they will all be at 0Ah charge, which means they are balanced. However, if all cells are fully charged, they will be at varying Ah levels owing to their different capacities. This implies a large amount of imbalance, but is actually desirable, so the imbalance metric needs to be recalculated depending on the direction of charge.

To work around this, SOC can be used instead. As this is a 0-100% metric which factors in the different capacities, it provides a consistent quantification of imbalance. Removing imbalance means bringing all cells to the same SOC, whether that is 0%, 100% or anywhere in between. However, a relative metric like SOC can potentially result in over-balancing (balancing when not necessary). Figure 7 shows two cells with different capacities but containing the same amount of charge. Cell B's lower capacity means that it is at a higher SOC than cell A, but this also means its SOC will decrease faster than cell A's. The cells will be balanced at 0% SOC, but their SOCs indicates they are presently out of balance which would result in (unnecessary) balancing. While SOC is also an estimate, it is a more reliable estimate than absolute charge level, as to estimate charge, the SOC estimate is combined with an estimate of the cell capacity, compounding two sources of uncertainty. There is a large body of literature around the topic of SOC estimation, with a wide range of proposed solutions. This is discussed further in section 2.2**.**

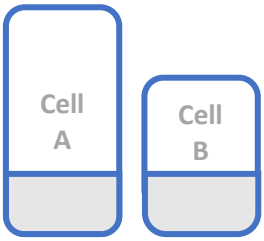

Figure 7: Comparison of SOC imbalance and charge imbalance

While charge and current are often the focus of discussions about cell performance, energy and power are ultimately what are required of a battery system. To achieve a certain power value, equation (3) shows that a higher current is required to compensate for a lower open circuit voltage (OCV). This means that for a BEV driver following a certain route, the battery pack SOC will decrease slower when the pack is fully charged compared to if it is only partially charged, even though the amount of energy used is the same[2]. The SOE would change the same amount both times.

$$P = v_{oc} i \tag{3}$$

From this perspective, state of energy (SOE) is a more representative metric for describing imbalance. However, SOE adds further complexity compared to SOC. The equation for SOC is linear, as it is simply accumulating current over time, and scaling by cell capacity (see Chapter 3 for a thorough derivation of cell state equations). SOE calculation must factor in how the OCV changes with the energy level, which is a nonlinear relationship. This nonlinearity makes control design more complex, while still ultimately being a similar imbalance metric to SOC.

Often, differences in cell voltage between cells is used to gauge imbalance [52, 72–74]. Voltage is a direct measurement, avoid issues with estimates. However, as mentioned in section 2.1.3.3, cells at the same charge or energy level can be at different voltages, because of their impedance differences. If the current direction changes, the cell at the highest voltage would become the cell at the lowest voltage, which from an imbalance perspective is confusing as it also reverses the state of imbalance. This also hints at another problem with voltage:

[2] In practice there will be slight differences in energy usage because the efficiency of the battery pack is a function of SOC and current rate.

unlike the others metrics discussed here, it is not a state and changes instantaneously with current. This again complicates the control strategy. Consider some cells at rest, where one cell is at a slightly higher voltage than the others. The balancing controller sets it to discharge, to bring its voltage down to the same as the other cells. However, the application of the balancing current drops the cell voltage to approximately the same voltage as the other cells. This would imply it no longer needs to be balanced, but removing the balancing current would return the cell voltage to close to its original voltage, higher than the others.

Often balancing, especially passive balancing, takes place at or near EOC [7], where the current magnitude is relatively low and will not change direction. This can be considered a coarse approximation of OCV and so SOC (the relationship between OCV and SOC is discussed further in Chapter 3). However it is clear that under a dynamic load, voltage is an unusable metric for gauging imbalance, and unsuitable for control purposes. One relatively simple means of improving the voltage metric is to subtract the resistive voltage from the measured voltage v_t to obtain a more stable value v_s using (4). This removes the most dynamic fluctuations in cell voltage. However, this has some drawbacks. The internal resistance estimate R_{cell} does not account for the full cell impedance and so some voltage dynamics are still present. R_{cell} is also an estimate, and as noted above varies significantly with SOC, temperature and ageing. This is similar to a common method of SOC estimation in which the impedance is modelled to predict v_t [75, 76], but with a very simple impedance model and no corrective feedback to improve performance. As the BMS requires a SOC estimate for other purposes, it would be logical to use that SOC estimate rather than produce what is essentially a poorer version specifically for the balancing system.

$$v_s = v_t - R_{cell} i_{cell} \tag{4}$$

The authors of [57] used voltage as a reference of imbalance, stating that using SOC is unreliable, although they do not provide evidence for this claim or quantify the impact. As such, their balancing control system requires hysteresis (upper and lower voltage limits) to start and stop balancing, leaving a dead zone in the middle to account for voltages changing dynamically under load. This means there will always be some level of imbalance: narrower threshold limits reduce

this imbalance, but could mean the balancing system repeatedly being switched on and off without having an impact on imbalance. SOC offers a suitable compromise, with minor drawbacks being outweighed by its simplicity and estimation reliability compared to other metrics.

2.1.4.2 Relative imbalance

Whatever metric is used, its value has to be assessed relative to the other cells. There are two main options for doing this. The global difference for an SOC vector $\boldsymbol{s}$ (5) can be used, which gives a good indication of overall imbalance, but does not give any indication of the state of imbalance amongst the cells, which then raises the question of how each cell should be controlled. Secondly, the distance from the mean $\boldsymbol{m}$ can be calculated, with the ith mean given by (6). This is useful when trying to determine a path for how all cells should be balanced, and the mean is a convenient reference. Alternatively, the distance of each cell from the maximum/ minimum cell could also be used. This would limit the balancing speed, as it would only occur in one direction, but may be applicable when the module of cells is near EOC/ EOD. For example, if the cells are being charged and are approaching EOC, then charging the strongest cells is not required: the priority is to bring all the cells up to the same level as the strongest. Taking the mean of the SOCs can be achieved with a constant matrix, whereas comparing to the maximum or minimum would require a change in function as different cells become the minimum/ maximum, causing discontinuities in the control process. Using this imbalance vector, the total imbalance can also be defined using the sum of squares (7). However, this is less relevant than (5) for describing total imbalance as it is the maximum and minimum that govern the overall pack performance.

$$\text{imbalance} = \max(\boldsymbol{s}) - \min(\boldsymbol{s}) \tag{5}$$

$$m_i = \frac{1}{\mathrm{N}} \sum_{n=1}^{N} \boldsymbol{s}_N - \mathrm{s_i} \tag{6}$$

$$\text{imbalance} = \sum_{i=1} [m]^2 \tag{7}$$

Consider the three cell SOC vector in (8). The ΔSOC vector based on the mean, maximum and minimum are shown in Figure 8. Table 1 shows these different measures for the example 3-cell SOC vector, $\boldsymbol{s}$. The sum of squares is different depending on the measure, so is not as effective as an overview of total imbalance and would have to be recalibrated based on operating conditions.

$$\boldsymbol{s} = [26 \quad 33 \quad 31] \tag{8}$$

Table 1: Summary of total imbalance measures for example cells

Imbalance Reference	Max-Min	Sum of Squares (relative to mean)	Sum of Squares (relative to maximum)	Sum of Squares (relative to minimum)
Value	7	26	53	74

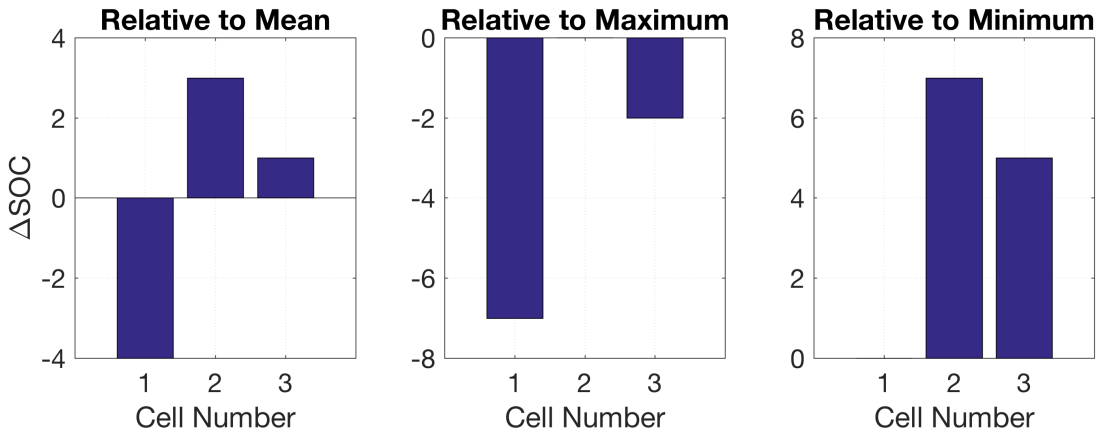

Figure 8: ΔSOC metrics based on the mean, maximum and minimum SOC

Based on the above discussion, the distance from the mean is best suited to an imbalance controller, and max-min imbalance used in the higher-level balancing control strategy, to assess the total state of imbalance.

2.1.5 SOC Estimation

For the SOC-based analysis in 2.1.4, an accurate SOC estimate is required. As discussed within [7, 58, 77], an SOC estimate is also critical for several aspects of battery pack operation. It identifies how much charge is remaining (which can be extrapolated to other metrics such as time remaining or driving range). It may also be part of the vehicle's energy management strategy, for example HEVs typically maintain the battery pack SOC inside a narrow window. SOC cannot be measured directly. If the cell is fully relaxed, then the terminal voltage is equal to the OCV, and so SOC can be directly obtained from a lookup table using data as

in section 3.3.1. However, even this can be subject to error if the OCV-SOC data is incorrect or simplified owing to interpolation or averaging the charge and discharge OCV-SOC curves. SOC can be calculated using (17), also known as coulomb counting, but there are several sources of error in this equation as summarised in Table 2**Error! Reference source not found.**.

Table 2: Sources of error from coulomb counting

Source	Errors
Initial SOC *s*	SOC is an integrator and so accumulates charge relative to a starting point. If the initial SOC is incorrect, the future values will also be incorrect.
Current i_{cell}	Errors in current can arise from incorrect calibration and quantisation. Additionally, discrete sampling means that the sample taken is assumed constant for the following time-step, when in practice it may differ.
Cell capacity Q_{cap}	The cell capacities well generally be an estimate and as such subject to some uncertainty. Incorrect estimates will incorrectly scale conversion from charge to SOC.

The above methods can be combined, by coulomb counting until the cells come to (or close to) a relaxed state, at which time the cell voltage is used to look-up SOC. However, this is not a viable solution unless the cells reach a near-relaxed state frequently, which is not guaranteed with an EV, particularly a BEV which is designed to cover a wide DOD. These methods can more tightly integrated using an estimation algorithm. There have been many methods proposed for cell SOC estimation, ranging from data-driven methods such as neural networks [78] and support vector machines [79] to model-based methods such as the Kalman Filter (KF) [75, 80, 81]. They all have the same premise, which is to continually use voltage measured to improve (correct) the SOC estimate, even when the cell is not relaxed. The KF was chosen because data-driven methods require a large amount of training data, which is time-consuming to generate and must be carefully chosen such that it represents the final application, whereas the KF integrates well with standard Catapult characterisation tests mentioned in section 3.4. Additionally, data-driven methods are more complex to develop and more

computationally expensive to execute. The literature suggests there is little difference in accuracy between the KF and data-driven methods.

2.1.5.1 Kalman Filter Framework

The KF is a commonly used algorithm for estimating states in physical systems and has been proven over time to be efficient and robust for a wide variety of applications [82–84]. The general data fusion method for a KF is outlined below. The ^ accent denotes that the variable is an estimate not a measurement, the - superscript indicates that the estimate is prior to correction, and the + superscript indicates the estimate is post correction. The *k* subscript denotes the relative point in time. The input vector ***u***, state vector ***x***, and process noise vector ***w*** are passed into the state update function to obtain the uncorrected state estimate at the next step using (9)**Error! Reference source not found.**. This is stored and used at the next time step. The uncorrected state vector, input vector, and output noise vector ***v*** are passed through the output function (10) to obtain an estimate of the measured output. The Kalman gain matrix L_k provides a systematic way of correcting the state estimates using the measurement data (11)**Error! Reference source not found.**.

$$\hat{\boldsymbol{x}}_{\boldsymbol{k+1}}^{-} = f(\hat{\boldsymbol{x}}_k^{+}, \boldsymbol{u}_k, \hat{\boldsymbol{w}}_k) \tag{9}$$

$$\hat{\boldsymbol{y}}_k^{-} = g(\hat{\boldsymbol{x}}_k^{-}, \boldsymbol{u}_k, \hat{\boldsymbol{v}}_k) \tag{10}$$

$$\hat{\boldsymbol{x}}_k^{+} = \hat{\boldsymbol{x}}_k^{-} + L_k(\boldsymbol{y}_k - \hat{\boldsymbol{y}}_k^{-}) \tag{11}$$

The specific calculation of *L* depends on the type of filter used. It is calculated according to two matrices: the state process noise matrix and output noise matrix. The process noise matrix indicates the accuracy of the model: small matrix values imply confidence in the model, meaning that the states will not be corrected much for a given measurement error, whereas larger values acknowledge that there are a number of model errors or unmodelled effects in **Error! Reference source not found.** [82]. Similarly, the output noise matrix indicates noise on the output measurements. In theory, these matrices are the covariance of ***w*** and ***v***

respectively, but in practice often become tuning factors whose values are adjusted until the desired performance characteristics are reached. One of the difficulties of the KF is that the ***v*** and ***w*** values are generally unknown and can vary with time, and often may not be Gaussian white noise as the filter assumes [82]. Implementation of a nonlinear KF for real-time SOC estimation is detailed in section 4.3.

2.2 Balancing Methods

A literature review was conducted on balancing hardware to understand the state of the art as well as future directions. There have been many proposed methods for active balancing, reviewed in section 2.2.1. Research into control of balancing systems is covered in section 2.2.2. The majority of the review focused on ABSs as there is little research to be done on passive balancing. Passive balancing involves discharging a cell by connecting a resistive load across its terminals. This is commonly used for commercial applications because of its low cost and [7, 73, 85]. However, there are several drawbacks:

- The discharging process wastes the energy – it is dissipated as heat.
- Balancing currents are typically low to avoid generating too much heat. Daowd et al. [73] suggest a maximum passive balancing current of C/100. Andrea [7] suggests 10-100mA for passive balancing depending on battery pack size. This means that balancing can take a long period of time. Even so, passive balancing may be limited owing to excessive heat generation [86]. Often the resistive loads for several cells are in a small, poorly ventilated box within the battery system where temperature rises can be significant. Moore and Schneider [74] deem the heat power generated by passive balancing large enough to cause costly thermal management requirements, even for a C/100 current rate.
- As discussed in section 2.1.3, passive balancing can only account for energy difference caused by self-discharge, not capacity differences between cells. This means that while passive balancing can ensure that cells are fully balanced at EOC, they will be out of balance by EOD, with the weakest cell still limiting the entire pack.

Despite these drawbacks, there does not appear to be much use of active balancing at a production level. The Toyota Prius and McLaren P1 both use some form of active balancing, though few details are available [87]. Both vehicles are HEVs, and it may be that the battery packs are small enough to make integrating an active balancing system easy enough compared to a larger BEV pack. Linear Technology Corporation have developed a production-ready active balancing chip, the LTC 3300-1 [88], which can be used on battery packs over 1000V. This chip was used for experimental work in Chapter 6. General Motors patented an active balancing method which uses a DC-DC converter using the strongest cell to the weakest cell [89], though it is not clear if it has been used on a production vehicle.

2.2.1 Active balancing methods

One alternative method is to bypass the cell completely, using transistors to switch the cell(s) out of the series string [17]. For example, during discharging, a cell at a lower SOC could be regularly disconnected from the series string so that over time not as much current is drawn from it as other cells. This has a low component count as no energy is stored or transferred. However, regularly switching during high loads is a challenging problem, and ensuring this switching can be achieved over the lifetime of a battery pack may be difficult. Temporarily removing some cells from the series string also reduces the power and energy density of the battery pack, and increases the load on the remaining cells [72], although a secondary benefit of such a system is that in the event of cell failure the cell could be switched out, allowing the pack to remain functional. Using transistors to bypass current has been theorised [57, 90], effectively using them as a variable resistor to regulate the proportion of current flowing through the cell, relative to the total applied load. This has the benefit of being able to have control over cell currents, particularly useful when cells are approach EOC and some cells will be closer to the maximum cell voltage than others, but still suffers from dissipating the current as heat, and not being able apply current in both directions. Andrea [7] concluded that ABSs which do not distribute charge (i.e. those which use a similar algorithm to passive balancing to balance during charging) are generally not worth adopting unless particularly high energy efficiency or a large balancing current is required, presumably for a BEV where the larger battery pack will require a higher balancing current and makes use of a wide SOC window.

One of the simplest methods of energy redistribution is using capacitive systems. A capacitor can act as a temporary energy storage medium between two cells: charged up by a higher voltage cell, and then discharged into the lower voltage cell. In [91], the authors recommend using ultracapacitors, whose low resistance can speed up balancing times. Despite this, the experimental results still show balancing taking 4 hours to remove 15% imbalance, which is relatively slow compared to other systems covered in this review. Sheng et al. [92] used cell-to-cell balancing using supercapacitors. As well as being used for equalization, the supercapacitors are also employed to absorb charge during a high-power event. By being in parallel with a cell, the total current is distributed between the cell and supercapacitor, reducing the load through the cell, which will help limit ageing. It is not clear from the paper how this is regulated; specifically, if the supercapacitors are maintained in a partially charged state such that they can accept and deliver large currents. An inductor can be used for cell-to-cell balancing in a similar way to capacitive systems, but the energy is stored in a magnetic field rather than an electric field [17]. The inductor size and switching frequency dictates how much charge is stored and the transfer rate between the cells.

This cell-to-cell connection is less effective in a large battery pack. Preindl et al. [93] showed how equalization time varies as a function of cell count, for different hardware topologies. Some schemes such as a dissipative system easily scale with the number of cells, whereas systems that use neighbouring cell connections will take proportionally longer to balance. As the number of cells increases, the likelihood that a low SOC cell is near to a high SOC cell reduces and so the imbalance reduction must ‘filter’ through the cells. Figure 9 shows a schematic of two cases of imbalance distribution. On the left-hand plot, there is a gradual increase in SOC from left to right.

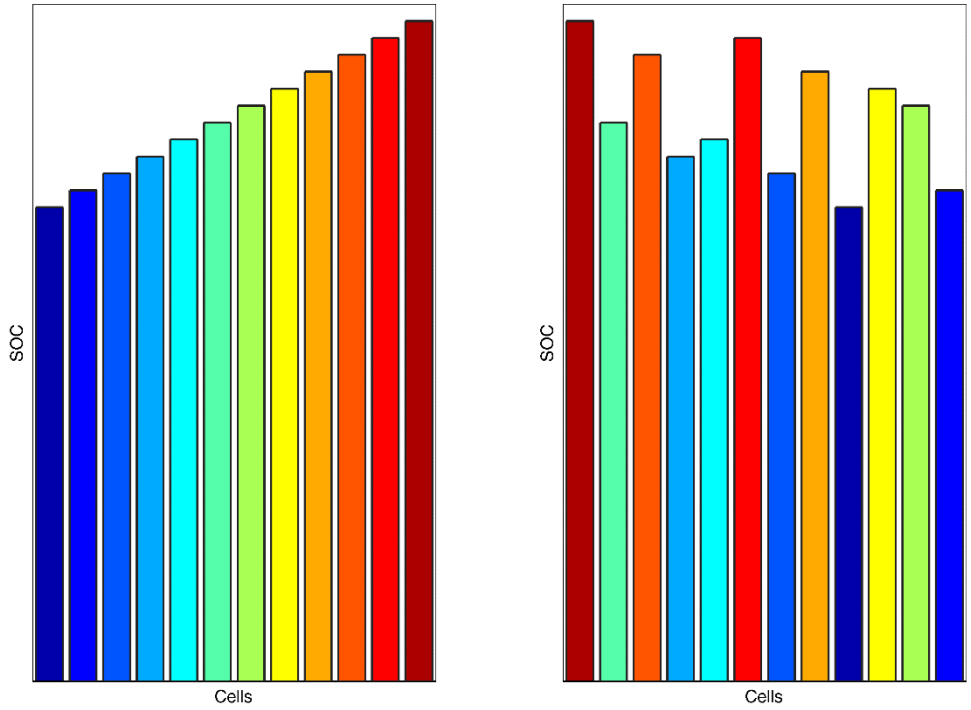

Figure 9: Schematic of two cell imbalance distributions

A cell-to-cell balancing system would have to perform a series of balancing operations to pass the excess charge from the right-hand cells along to the weaker left hand cells. Each balancing operation between neighbouring cells will not reduce pack imbalance by much, and increases energy inefficiency. However, on the right-hand plot the imbalance distribution is random. In this case, there are generally high SOC cells next to low SOC cells and so there can be a significant reduction in imbalance through a single balancing operation between two neighbouring cells. Kim et al. [94] proposed adding additional switches and/ or capacitors to speed up the process, but this still lacks flexibility with regard to where energy can be moved to and from, and the amount of charge transferred is dependent on the relative cell voltages.

A more flexible system involves using a cell-module connection, for example using transformers as in Figure 10, proposed by Bonfiglio and Roessler [95]. This bi-directional flyback converter system can either be charged by the module and discharged to individual cells (boosting the charge level of weaker cells), or utilise the energy of an above-average cell to energise the transformer. The left-hand diagram of the figure shows the module driving the primary winding of the transformer, which creates a magnetic flux in the transformer core. One of the switches (“sec2” in the right-hand diagram) is then closed, creating a circuit for

the secondary current to be established. In the paper only equalization of single cells is discussed, which could result in a slow balancing time.

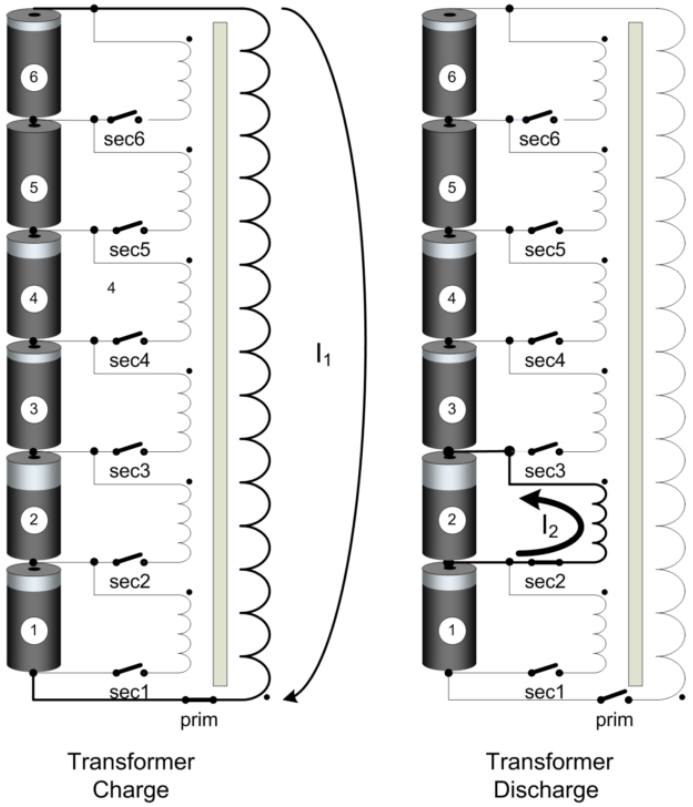

Figure 10: A transformer-based pack-cell system - from [95]

A time-shared system is proposed in [96] in which a group of cells share a single cell-module flyback converter. Multiplexing is used to connect each cell to the converter for a short period of time. The authors emphasise that the system is made up of fewer components than most other ABSs. Despite this, it is was able to remove a 33% SOC difference between four 2.65Ah cells in 90 minutes. The main issue is how this is scaled up to a battery pack. As the converter is shared, the balancing speed will generally decrease as the number of cells increases, unless the power is increased.

Kim et al. [97] developed a balancing system which uses a DC-DC converter to create a voltage source equal to the mean cell voltage, powered by the pack. This is used to charge or discharge individual cells until they are all at this mean voltage. The proposed benefits are primarily speed of balancing and low cost. A prototype was developed and tested using eight 7Ah cells and the results show a 13% absolute decrease in SOC in 30 minutes. However, these results are for a prototype with 8 cells in series. Since the balancing works sequentially on each

cell, the balancing time will increase in proportion with the number of cells, making it much slower for an EV battery pack which can use over 100 cells in series.

Dual balancing systems have also been proposed, which separates imbalance at the cell and module level. Liye et al. [98] employ an inductive system to balance modules, and cell-to-cell capacitor system between the cells in each module. The schematic show in Figure 11 shows the modules E_1 to E_n which can each be switched into the transformer circuit which in turn will charge the pack. The authors see it as being particularly suitable for a HEV drive cycle due to its ability to work during charge and discharge and not requiring a long period of time to balance, although there is no further justification of why it suits an HEV as opposed to other EVs. The results are not described in terms of balancing time or current so comparison with other systems is difficult.

Lin et al. [99] uses an inductive system to actively balance neighbouring modules, followed by dissipative balancing of the cells within each module. In [95], an extension is proposed to balance between modules in a similar way that cells within each module are balanced. None of the papers contains analysis on how much each system contributes to reducing the amount of imbalance or how much energy each system wastes through heat dissipation or inefficiency. Furthermore, analysis of the effectiveness of the dual balancing system as a function of pack and module size would offer greater insight into the benefits and limitations of such a system.

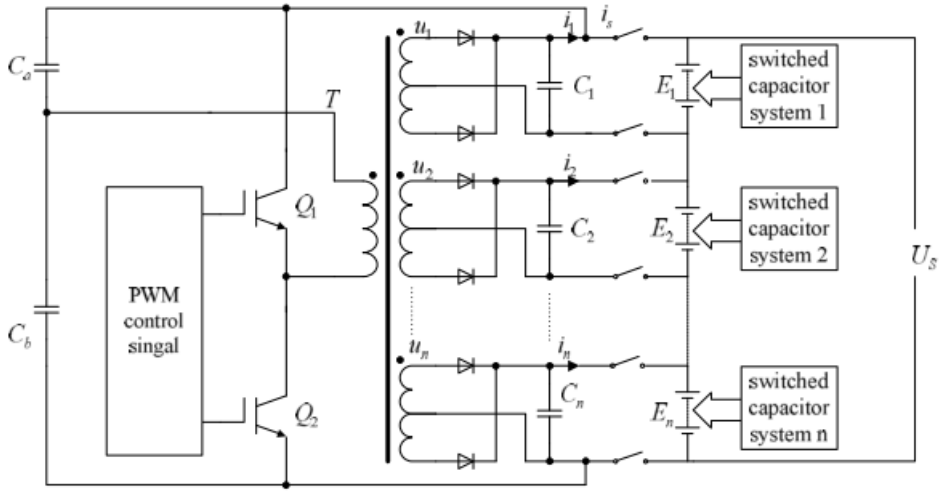

Figure 11: Dual balancing system – from [98]

A similar two-layer approach is used in [100], where cell-to-cell balancing takes place within modules, and then the modules are balanced. As discussed above, cell-to-cell balancing becomes slower as the number of cells increases, so the performance depends on how the battery pack is subdivided into modules. In [101], the authors avoid having to split the system into modules. They use a DC-DC converter on each cell which all link to a common output bus for the low voltage system of an EV. The total applied load current can be distributed amongst the cells, with weakest cells undergoing less of the load. The hardware appears promising from a control perspective as it offers variable current allocation of each cell, but the problem of scaling the system up from the 3 cell demonstration to a full battery pack remains open.

Gallardo et al. [17] summarised 24 balancing configurations, covering passive (dissipative) and active systems. Active systems are categorised into shunting (changing the current path), shuttling (moving energy form one cell to another) and energy conversion (using a DC-DC converter to transfer energy). The configurations were given a score of 1-3 (poor, average, good) on factors such as efficiency and balancing speed. Complete cell shunting was deemed the best method, primarily due to its relative simplicity. Transformer based systems tended to rank poorly due to their complexity. The primary exception is for a full bridge converter, which is considered by the authors and Omar et al. [73] to be an effective system, especially with regard to efficiency and balancing speed. This comes at the expense of control complexity and cost. Cell bypassing scores well, despite the issues raised above. Reliability of active balancing systems has not been explored in the literature, likely because of the difficulty in setting up realistic, long-term experiments. However, this is a critical future area of research for assessing the potential of ABS for applications such as EVs, where several years of operation is expected. The authors note that the scores are subjective and may be weighted depending on system requirements. For example, when considering a balancing system for a production vehicle, the scoring may be changed to emphasise low cost and reliability, or for a BEV the system will need to cope with higher currents than for a HEV. Both analyses are hardware focused, not considering system-level factors such as how much control there is over the amount of charge transferred and what the requirements (such as power and

balancing time) should be for a given application. The focus tends to be on simplicity rather than considering the ultimate performance benefits such as increased energy utilisation. Cao et al. suggest that of the configurations assessed within [52], only a capacitor-based system is appropriate for use during discharge. This is largely in agreement with Gallardo et al., who add that certain inductive methods are effective during discharge. It is not clear from the papers what makes a bi-directional system effective rather than merely functional.

2.2.2 Balancing System Control

Until recently there was little work on how ABSs should be controlled: the majority of the work emphasised proof of concept for novel hardware. References [98, 99, 102–105] use simple rule-based logic to determine what cells to balance. Often, open-loop control, or occasional feedback is applied [106, 107] and the cell capacities are used to predict how much charge each cell requires. This does not account for the fact that cell capacities are estimates and the charge delivered by the balancing hardware may be different to what is expected owing to set-point error and unquantified losses in the electronic hardware. The balancing system in [100] is promising from a control point of view because of its flexibility in controlling the cells. The authors use SOC distance from the mean, as proposed in section 2.1.4.2, but their experimental results only involve one demonstration where the system is activated for a period of time to demonstrate imbalance reduction, rather than exploring the control strategy further. The balancing system in [101] uses SOC as a metric, but the authors have designed it to operate autonomously (without a dedicated control system), with it balancing perpetually. This could reduce the overall efficiency substantially, since there are many times at which balancing is not required (see section 5.1).

There has been some research focussing on balancing control, but these have generally not been validated experimentally. Danielson et al. [108] developed a model predictive control algorithm which aims to maximize pack capacity (where pack capacity is tied to the weakest cell) and reduce SOC imbalance with a capacitor based system. The system is considered as a network whereby the cells are ‘nodes’ and the interconnections (capacitors) between them are ‘edges’. Additional constraints, such as the maximum current and absolute limits on SOC

are applied. An example of SOC convergence for 16 cells with variations in capacity and initial SOC is shown in Figure 12. The y axis is SOC, scaled from 0 to 1, and the x axis is time (the units are not given). The SOCs converge to a common value, although trajectory of the SOCs is not ideal: two cells (green and teal lines) diverge significantly before they converge. There are several assumptions made, including a linear cell model with no impedance, and the pack topology and cell interconnections are unlikely to be similar to what is found in an EV. However, the paper shows the potential for a more advanced control system which uses predictive control to maximise energy utilisation.

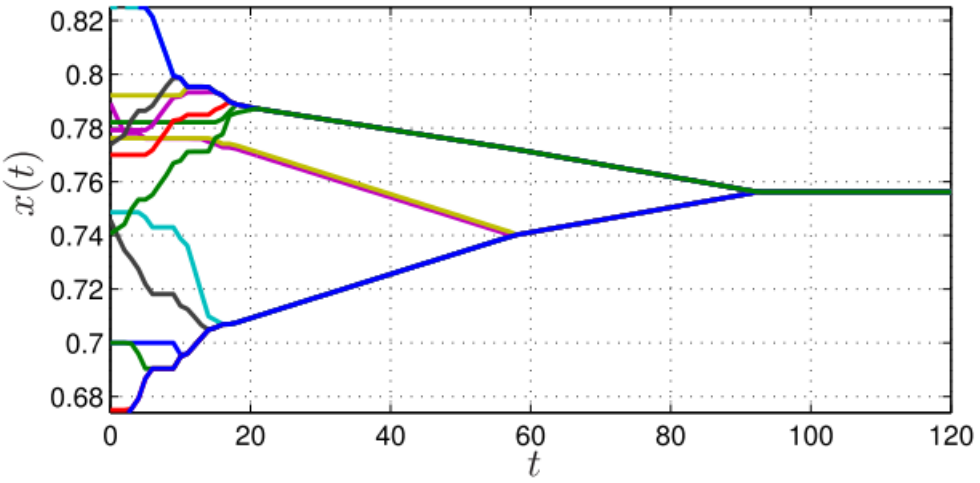

Figure 12: SOC balancing using model predictive control - from [108]

Quinn and Hartley [109] also treat the battery pack as a network, where each cell is a 'node', and balancing connections (in this case cell-to-cell transformers) are 'edges'. The pack arrangement is two series strings of four cells, connected in parallel. Different amounts of interconnects were applied, and as would be expected, a pack with a greater number of cell-to-cell interconnects can balance more quickly and settle on a lower level of imbalance. They also note that the increased number of interconnects improves the reliability as the system performance is less affected by an interconnect failure. However, more interconnects also equates to higher system complexity, size and cost and a greater chance of a failure. The type of analysis in [108] and [109] could be used in order to evaluate possible options for cell interconnects as part of the battery pack design process, for example comparing the cost of components against balance time for different configurations.

Lee and Cheng [110] applied fuzzy logic to reduce the balancing time for a modified Ćuk converter, which uses a capacitor to link neighbouring cells and inductors over each cell to control current flow. Imbalance between cells is grouped based on the voltage difference between each cell and the pack average. A duty cycle is created based on the amount of imbalance, with the aim of all cells reaching the maximum SOC simultaneously during charging. Results show the fuzzy logic scheme reduces balancing time by 32% compared to conventional methods, though the conventional method is not defined; it may refer to a simple rule-based approach. Nonlinear model predictive control (MPC) has also been applied for a Ćuk converter system [111]. While potentially offering benefits for speed and overall imbalance reduction, it is computationally expensive and so unsuitable for a control system on board a vehicle, and has not been validated experimentally. Nonlinear MPC has also been applied for a multi-level converter, which uses a DC-DC converter per cell to regulate individual cell currents. The authors propose a scheme which aims to balance cell temperature as well as charge level, which has the potential to reduce the capacity spread from ageing. However, the addition of a temperature model adds complexity. This control system has not been validated experimentally, although the simulation studies do at least consider the battery pack under load from a standardised vehicle drive cycle.

Recently McCurlie et al. [112] proposed a model-predictive control scheme which uses SOC as a metric of imbalance. It was evaluated with a 6s3p battery pack and a Linear Technology Corporation balancing development board, which utilises their LTC3300-1 flyback converter balancing chip. The control system aimed to minimise the time to remove imbalance by solving a quadratic optimisation. The optimisation finds the ideal balancing currents along with the balancing time, factoring in the constraints on the amount of current the system can develop. The control strategy proved effective in a real-time environment: a 27% SOC different was reduced to 2% in approximately 20 minutes. The authors compared it to a simple rule-based controller, which sets a cell to charge if it is below a certain SOC threshold and discharge if it above a threshold. The rule-based controller reduced imbalance in a similar amount of time, but resulted in

twice as much power loss through inefficiency, as the cells are always balancing at full power rather than being regulated by a control system. The hardware was based on the same chips as for this research (see Chapter 6), and the control strategy is also related, although derived in a different manner (see Chapter 5). However, there are some limitations. Firstly, there was no SOC estimate (coulomb counting was used which was assumed sufficiently accurate considering the laboratory conditions) and no applied load. This means the balancing controller was not subject to the noise and uncertainty expected during more realistic operation. As with much of the above discussion, the system was proven to be able to balance, but the wider question of when to balance, and what it ultimately aims to achieve is not covered.

Moura et al. [113] considered using switches to alter the current through cells with the aim of minimising film growth (which leads to impedance rise) in cells. An equation for SEI layer thickness was presented, based on electrochemical properties, and this was then turned into a cost function for an optimal control problem. For the study only parallel cells were considered, but this leads to the concept of utilising the balancing system to combine balancing and charging to reduce ageing.

2.2.3 Cost-Benefit Analysis of ABS

Much of the above discussion, and work in this Innovation Report demonstrates the potential performance benefits of ABS, there are other factors, notably cost, reliability and manufacturing. While these are outside of the scope of the project, it is important to consider these issues as they will affect the viability of ABS regardless of performance. The primary benefit has been identified as energy utilisation, so a metric such as battery pack cost relative to the available energy of the pack ($/kWh) is a suitable starting point. A further analysis considering the lifetime of the pack would also be beneficial, considering the potential benefits for ageing and long-term performance.

The primary benefits of ABS over passive balancing are as follows:

- The energy output from the same specification battery pack will be increased

- The SOC operating window of the pack can be narrowed to reduce ageing
- The same energy output can be maintained while reducing the battery pack size (volume and weight), and number of cells (each cell being a potential source of pack failure)
- There is less chance of one poor-performing cell causing the entire pack to lose performance and require replacing

There are several drawbacks which generally come about through the increased complexity compared to passive balancing:

- Increased design, development and testing resources are required
- The direct cost of hardware will increase
- There are more wires and interconnections, which add weight and design challenges, as well are more failure modes.

The specifics of the cost-benefit analysis will be heavily dependent upon the application, such as driving performance, the expected lifetime and cost constraints.

2.3 Summary

The literature review in section 2.1 clearly shows that energy imbalance is inevitable. Even if a large batch of cells is screened to ensure consistent properties within a new battery pack, significant differences will emerge with time and usage. This variation is difficult to predict, with stronger cells becoming weaker cells with ageing, and vice versa. Analyses of used cells from EV battery packs shows that the outliers which are inevitable in a large battery pack are much more likely to be low capacity rather than high capacity, increasing the proportion of unused energy in the battery pack unless active balancing is employed. A greater than 10% difference between maximum and minimum cell capacity can be expected after several years of operation.

This information was used to address Research Question 1. **The context in which active balancing can provide the greatest benefit is applications where large variations in cell capacity can be expected, and the majority of battery pack energy is required**. The main example of this is BEVs, where cells are subject to several years of operation under demanding conditions, and almost

the entire SOC window is required to meet range requirements. For applications such as HEVs, losing some battery capacity because of some weaker cells may not be critical to operation, as their full capacity is not utilised.

Research Objective 2 was addressed through the literature review in section 2.2 This showed the variety of options for active balancing and suggested that there is no single best solution, and there is little analysis of what systems is best suited to what application. It can be difficult to compare similar systems given the nuances in design and different test set-ups (for example, power capability of the hardware, the number of cells and cell specifications).

The focus of active balancing research to date has largely been on demonstrating proof of concept for the hardware, with the control system generally an afterthought. Some more advanced control strategies have been proposed, but even these often do not consider how the ABS is best utilised, such as balancing while discharging. Also, most of the control systems considered are generally tied to a specific hardware type, limiting their applicability as new hardware is developed.

There has been little work to date which covers both areas of the literature reviews. ABSs are typically designed in isolation, and evaluation of cell variation and ageing focuses is on analysing the degradation modes and root causes rather than evaluating how active balancing could have impacted the results. Existing analyses of ABS often lacks the context of how it will be used in the application. Several systems in the literature only allow a single cell to be set to balance within the pack, and many are unidirectional. The remaining Research Objectives aim to address this gap. A balancing control system has yet to be developed which fully considers the wider BMS and maximization of energy utilisation. More generally, there has been a lack of analysis of balancing control systems. For the second Research Question, a model-based design approach was used, including using cell models to evaluate battery pack performance and for simulations as part of control system development. These cell models are introduced in the next section.

3 Cell Modelling

The lack of holistic analysis of ABSs can be addressed through mathematically modelling each area of the complete system. Research Objective 3 was chosen specifically to close this knowledge gap. To better understand how a balancing system interacts with the battery pack and BMS, a model is required. A series of modelling studies was performed throughout this research, at the core of which is a cell model. This is first introduced in this Chapter, before being used for several applications in Chapter 4.

Consider the generic balancing system in Figure 13. The cells in the plant model can be simulated without a balancing system to understand cell behaviour, such as how imbalance is generated under various conditions and drive cycles. They can also be used with a balancing system model as part of the controller design and development process. Cell models can also be further integrated together to explore whether cells connected in parallel are inherently balanced, as the literature suggests. This can then feed into the selection of balancing hardware. A cell model can be used to estimate SOC, a value which cannot be directly measured but is important for several BMS functions, including determining the amount of imbalance within the battery pack.

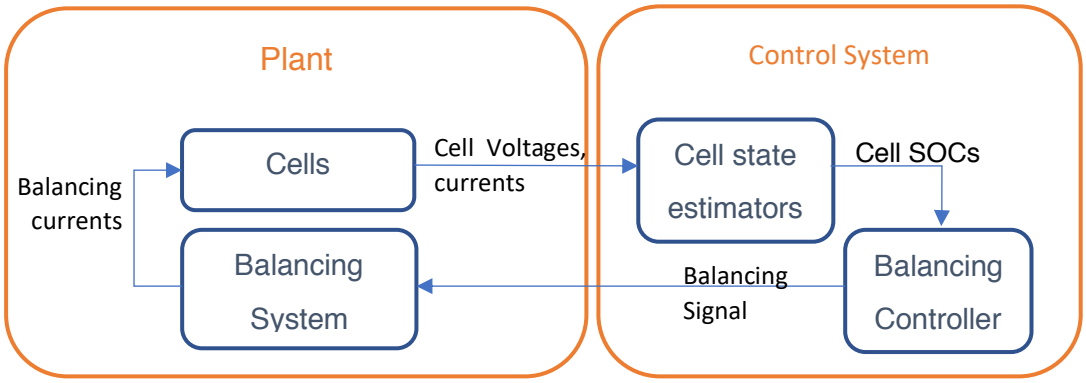

Figure 13: High-level summary of balancing control system components

These applications are implemented in Chapter 4, but before this an appropriate cell model must be clearly defined. In section 3.1, the type of model is selected, and the equations derived. In order to use these equations, the parameters must be assigned specific values, which depend on the properties of the particular cell being modelled. Section 3.4 describes what cell data are required and how this can be used to parameterise the generic cell model.

3.1 Model Selection and Derivation

For most applications, in line with its causality, a cell model accepts current as an input, and outputs the cell voltage. Depending on requirements, there may be other inputs and outputs such as temperature. Within the context of battery performance simulation, there are two main categories of cell models: equivalent circuit models (ECMs) and physics-based models (PBMs). The latter aims to capture the electrochemical processes using partial differential equations [55]. ECMs use a combination of electrical components which aim to match the impedance response of the cell. In this sense ECMs lack a physical basis, although certain components can be related to physical phenomena.

PBMs can potentially offer more insight into cell behaviour as well as better accuracy [114]. However, they are much harder to parameterise, with a full model containing typically 88 parameters, not all of which are identifiable [115] and a reduced model still containing 35 [55], many of which being difficult to measure or estimate. This makes finding the parameters for a particular cell challenging and time-consuming. Additionally, the equations are complex and computationally expensive, making them generally unsuitable for real-time applications [55, 114]. ECMs are proven to be effective for the applications considered here, such as offline performance modelling and state estimation. For this reason, ECMs have been used for the modelling work throughout this research. However, in the future as computing power increases, battery system development could be improved by using PBMs, either for control and state estimation, or as an accurate plant model to test ECM-based systems against. References such as [116–118] show various ways of reducing the model order of these equations for control and estimation purposes.

3.2 ECM Variants

Parameterisation of an ECM is a curve-fitting problem: finding a combination of components which accurately matches some experimental data whereby a current signal was applied to the cell and the resultant change in terminal voltage was measured. Because there is a lack of physical basis, there is no ‘true’ model structure of values, and the best structure and values will depend on the

application and the data collected. There are two types of dataset for generating ECMs: frequency domain and time domain. The majority of the work here considers time domain models only. Details of how frequency domain data was collected and processed is in Appendix A.

3.3 Time-domain ECM Derivation

Sample rates of 1-100Hz are common for cell systems, meaning that some high-frequency phenomena such as charge transfer, and induction will not be resolved using this data. A typical vehicle drive cycle current is primarily a low frequency (less than 1Hz) signal [119], and so modelling high frequency phenomena would add complexity without much performance benefit. Time domain data will typically offer less insight into cell behaviour than frequency domain data because of the reduced bandwidth, but is more effective at determining a model which predicts the cell voltage during current excitation**.** A wider array of input signals can be used, such as representative vehicle drive cycles, and the lower bandwidth reduces the number of elements in the model.

There are a range of ECMs proposed for time-domain solutions [114]. A typical ECM takes the form of Figure 14**.** Throughout this document, the sign convention is that a positive current means the cell is being charged, and a negative current indicates discharging. The ECM can be used alongside the OCV v_{oc}, to find the cell terminal voltage v_t from cell current i_{cell} using (12). This incorporates an internal resistance R_0, connected in series with a resistor-capacitor (RC) pair with component values R_p and C_p respectively. Several RC pairs can improve model accuracy by increasing the bandwidth, with each RC pair having a different time constant.

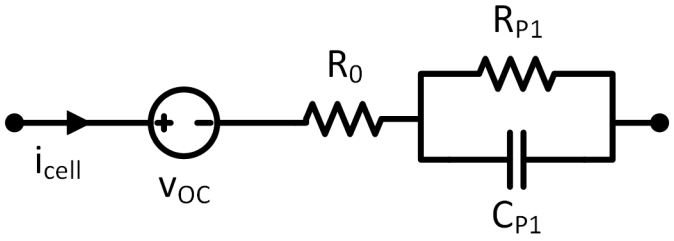

Figure 14: ECM with 1 RC pair

$$v_t = v_{oc} + \sum_{n=1}^{N_{RC}} v_{p_n} + R_0 i_{cell} \tag{12}$$

There is no limit on how many RC pairs can be added in theory, but it becomes challenging to parameterise owing to a lack of structural identifiability [120], and there may be more time constants than are required for a good fit, causing several RC pairs to have the same time constant. The internal resistance in this model is not strictly the direct current (DC) resistance of the cell as shown in Appendix A. This is because the sample rate of the data is generally not high enough to capture the DC response. Instead, this resistance encapsulates the impedance above the sample rate employed to measure cell data.

The governing equation for the RC pair state voltage v_p is given in continuous and discrete forms by (13) and (14) respectively. The time constant used in these equations is calculated from the resistance and capacitance using (15).

$$\dot{v}_p(t) = -\frac{1}{\tau_p} v_{rc}(t) + \frac{1}{C_p} i_{cell}(t) \tag{13}$$

$$v_p^{k+1} = v_p^k e^{-\frac{\Delta t}{\tau}} + i_k R_p \left(1 - e^{-\frac{\Delta t}{\tau}}\right) \tag{14}$$

$$\tau_p = R_p C_p \tag{15}$$

OCV is also a state, although generally it is obtained from SOC, *s*. The continuous and discrete SOC state equations are given by (16) and (17) respectively, where Q_{cap} is the measured cell capacity in Ampere-hours. The 36 is required to scale the numerator in As and denominator in Ah to a percentage. SOC is a linear integrator, which has implications for control, discussed in Chapter 5.

$$\dot{s}(t) = \frac{1}{36 Q_{cap}} i(t) \tag{16}$$

$$s_{k+1} = s_k + \frac{i_{cell,k} \Delta t}{36 Q_{cap}} \tag{17}$$

$$v_{oc} = f(s) \tag{18}$$

The OCV can be obtained from a typically nonlinear function (18). Often a look-up table is combined with interpolation is used. Alternatively, a nonlinear function can be parameterised, such as the variants in [121]. This will generally reduce

the accuracy of the data as, unlike with interpolation, the curve will not pass exactly through the measured data points, but it means that less data have to be stored in the memory of the control system, which can be useful for low-cost online applications. Additionally, if the function is differentiable it can simplify certain SOC estimation algorithms where a gradient is required, such as the Extended Kalman Filter approach, discussed further in section 4.3.

Variations on this basic structure of the model have been proposed and implemented. Using different values depending on the current direction can help improve accuracy, at the expense of increasing the number of parameters. For some cell chemistries such as Lithium Iron Phosphate (LFP), there can be a significant difference in the OCV curves depending on if the cell is charged and discharged. In this case, adding a hysteresis function to link the two curves can substantially improve modelling accuracy compared to using a single curve [84, 114, 122]. As discussed in Submission 5 and section 7.2.2, a key area of research on improving the ECMs is on improving their low frequency performance, as RC pairs fundamentally do not capture the diffusive aspect of cell dynamics [63, 123].

The resultant ECM can be considered as a state-space system, with current as an input, terminal voltage as output, and SOC and RC pair voltages as states. The state equation are linear (but potentially parameter-varying), and the output nonlinear because of the OCV-SOC relationship. For a 1RC pair continuous-time ECM, the state-space equations are given by (19).

$$\begin{bmatrix} \dot{s} \\ \dot{v_p} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & -\frac{1}{\tau_p} \end{bmatrix} \begin{bmatrix} s \\ v_p \end{bmatrix} + \begin{bmatrix} \frac{1}{36 q_{cap}} \\ \frac{1}{C_p} \end{bmatrix} i_{cell} \qquad (19)$$

$$[v_t] = [f_{sv}(s) \quad 1] \begin{bmatrix} s \\ v_p \end{bmatrix} + [R_0] i_{cell}$$

As noted in Submission 2, the OCV can be used as a state instead of SOC, as the two are equivalent. This moves the nonlinearity from the output equation to the state equations. SOC is used as a state in almost all of the literature, because it is more intuitive, and of more interest than OCV.

Controllability and observability are two important characteristics of a system. The former determines whether all of the states can be guided to a given target state through the control inputs, and the latter determines whether the initial state can be obtained from the time history of the inputs and outputs [124]. For a generic linear state space system of the form in (20), the two concepts are related and quantified using the same method for a linear system:

- Obtain the controllability and observability matrices $\mathcal{C}$ and $\mathcal{O}$ according to (22) and (23) respectively, where N_s is the number of states.
- Calculate the rank of the matrices;
- The difference between the total number of states and the rank equals the number of uncontrollable/ unobservable states.

As noted above, the cell model output equation is not linear, but can be linearised about a given SOC s_{lin}, with the resultant output matrix calculated using (21).

$$\begin{aligned} \dot{\boldsymbol{x}}(t) &= A\boldsymbol{x}(t) + B\boldsymbol{u}(t) \\ \boldsymbol{y}(t) &= C\boldsymbol{x}(t) + D\boldsymbol{u}(t) \end{aligned} \tag{20}$$

$$\begin{aligned} C &= [g_{sv} \quad 1] \\ g_{sv} &= \frac{df_{sv}(s_{lin})}{ds} \end{aligned} \tag{21}$$

$$\mathcal{C} = [B \quad AB \quad \cdots \quad A^{N_s-1}B] = \begin{bmatrix} \frac{1}{36q_{cap}} & 0 \\ \frac{1}{C_p} & \frac{-1}{C_p^2 R_p} \end{bmatrix} \tag{22}$$

$$\mathcal{O} = [C \quad CA \quad \cdots \quad CA^{N_s-1}]^T = \begin{bmatrix} g_{sv} & 1 \\ 0 & \frac{-1}{\tau_p} \end{bmatrix} \tag{23}$$

The states are completely controllable, and generally observable. The system becomes unobservable if the OCV-SOC gradient g_{sv} (21) is zero, that is if the OCV-SOC is completely flat. In this case, a number of different SOCs can result in the same voltage, so it is not possible to obtain a specific SOC from voltage. The impact of OCV-SOC curve gradient on SOC is discussed further in [125], in which the authors evaluate performance bounds of SOC estimators based on the

curve, and also note the importance of voltage measurement resolution on the estimation accuracy. The system remains controllable and observable if more RC pairs are added, as long as each time constant is distinct.

3.3.1 SOC-OCV Data

To calculate the terminal voltage using equation (12), the OCV is also required. As mentioned in section 3.3 experimental data are necessary, and often used directly in the model. There are two main methods for obtaining OCV Data: galvanostatic intermittent titration technique (GITT) and slow discharge, and procedures for both are presented in Appendix B. GITT increments the SOC by a certain amount (typically 3%) then rest for 3 hours so that the cell is relaxed, and a measurement of OCV is made. This is repeated, starting from fully charged, until the cell is fully discharged, and then charged with the same increment. This produces different OCV-SOC curves for charge and discharge, which is expected based on theoretical knowledge of cell behaviour [126]. The difference between the two curves depends on material and chemical properties. LFP cells are known to exhibit significant hysteresis [63, 114], but the nickel-cobalt-aluminium cell in Figure 15 shows little hysteresis, with a maximum of 9mV.

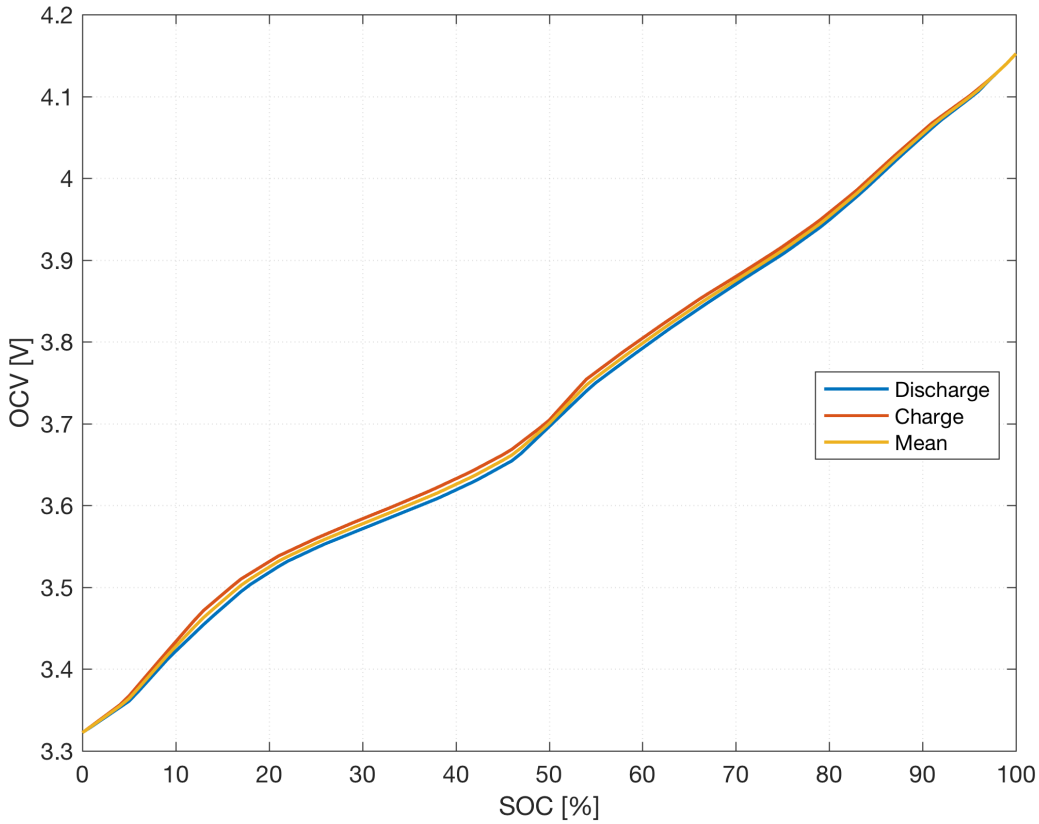

Figure 15: Example OCV-SOC curves obtained using GITT for a nickel-cobalt-aluminium cell

An example SOC-OCV curve can be obtained by discharging the cell at a low constant current, such as C/25. This low current aims to minimise the impact of the cell's impedance making the terminal voltage close to the OCV, hence sometimes being referred to as a pseudo-OCV-SOC curve [127]. This curve takes less time to produce, and can be useful for other purposes such as incremental capacity analysis [128], but is generally less accurate because only discharge data are collected, and there is still an impedance contribution, especially towards low SOC as more charge is built up and impedance increases.

3.4 Parameterisation and validation

Pulse power tests are often used to generate data to fit an ECM. This consists of a 10 second pulse of constant current followed by a 30-minute rest, shown in Figure 16**,** taken from the cylindrical cell data discussed in section 4.1. The rest period ensures that the cell has fully relaxed before the next pulse begins. Typically, five different magnitudes of current are applied, before moving to another SOC level to repeat the test. These are determined such that the highest magnitude equals the maximum safe current according to the manufacturer's specifications. The bandwidth of this signal is not representative of a typical load profile: improvements have recently been proposed [119, 129] to address this.

As mentioned above, finding ECM parameters is a curve-fitting exercise[3]. Non-linear least squares (NLLS) algorithms are commonly used for this purpose as it is relatively straightforward to implement and a variety of constraints can be applied [130]. For much of the offline parameterisation, as discussed within Submssions 2 and 5, the MATLAB function `lsqcurvefit` was used. However, an alternate NLLS algorithm was written for real-time use, as the MATLAB implementation was not compatible with Simulink or C-code generation to improve its speed. To solve this, a MATLAB class was written to solve a constrained NLLS function, which was shown to perform to a similar level of

[3] It is also possible to fit an ECM recursively rather than to a batch of data. This is particularly useful for online applications, and is discussed in section 4.3. However, it will generally be less accurate because there is less data available. The higher computational expense and memory requirements of batch-fitting are not a concern for offline data processing, but can inhibit their use online.

accuracy as the MATLAB algorithm. This algorithm is based on a quadratic optimisation, and is derived and explained fully in Submission 4.

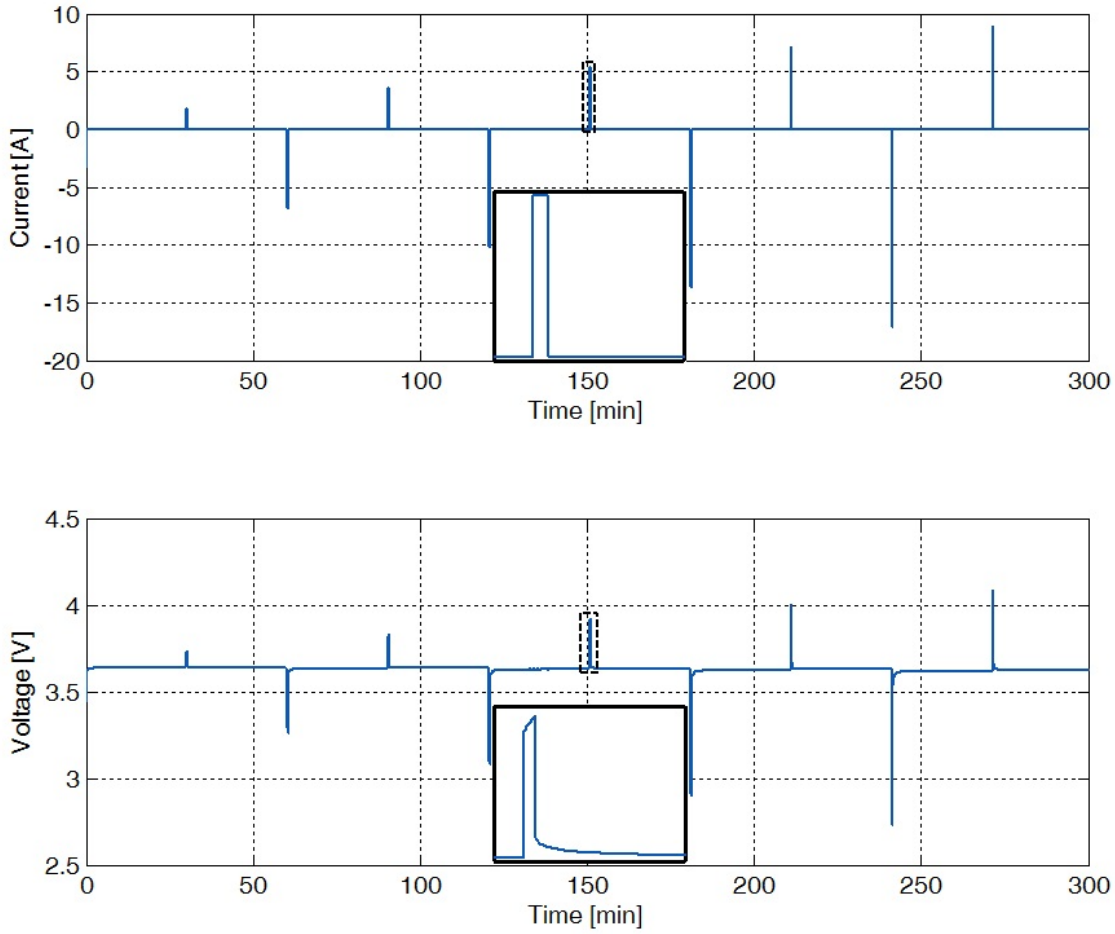

Figure 16: Example pulse power test current and voltage

For time domain data, a capacitor C_0 is introduced in series with the ECM. This capacitor is used to account for the small change in OCV which occurs during the test. Since the change in SOC should only be small, typically a fraction of a percent, this linear approximation of OCV-SOC is appropriate.

The NLLS algorithm is based around minimising the error between a series of measurements *y* and a nonlinear function *f*, which is dependent on the parameter vector $\boldsymbol{\Theta}$. Other known parameters can be passed to *f*, such as the input current. For each measurement, an error *e* is defined according to (24). The NLLS problem, defined by (25), aims to find $\boldsymbol{\Theta}$ such that the sum-of-squares error of $\boldsymbol{e}$ is minimised. An N^{th} order RC pair ECM has the parameter vector defined in (26); this can be generalised to any model.

$$e_i = y_i - f(\boldsymbol{\theta})_i \tag{24}$$

$$\arg\min_{\theta} \sum_{i=1}^{N} e_i^2 \tag{25}$$

$$\boldsymbol{\theta} = \begin{bmatrix} C_0 & R_o & R_{p_1} & C_{p_1} & \cdots & R_{p_N} & C_{p_N} \end{bmatrix}^T \tag{26}$$

For time domain data, the cell terminal voltage is used as the output equation (27). If the cell is not in a relaxed state when the time-domain fit begins, a NLLS fit can still be applied. This parameter vector must also include an OCV parameter and initial RC pair state voltages (these are not required when the cell starts at rest as the OCV is known from the terminal voltage, and the RC pairs are at zero volts). This will typically reduce the accuracy of the fit because there are more unknowns, but is useful for identifying model parameters while the cell is under load, and this method is used in section 4.3.

$$y_i = v_{t_i} \tag{27}$$

3.4.1 Example Parameterisation Results

Example results are shown in Figure 17 for a 3C discharge pulse. This was taken from a 3Ah cylindrical cell discussed further in section 4.1. Not all 30 minutes of the pulse data were used for the fit. For the majority of the test, the voltage is approximately at its OCV, and at this point only C_0 affects the voltage within the model. As a result, the NLLS algorithm will see a small error as long as it finds an accurate C_0 value, as poor parameters for other ECM components only create errors for a few data points. By restricting the parameterisation data to 1 minute after the current pulse ends, the dynamic region is given more weight, and better parameter estimates are obtained. The model parameters and associated simulation errors are given in Table 3. In this case the 2RC model performs similarly to the 1RC model, but captures the fast dynamics better, as can be seen when the current is applied and stopped (at 0 and 10s respectively). The 1RC model has a 4.4mV root mean square (RMS) error and 51mV peak error, whereas the 2RC model has a 3.0mV RMS error and 18mV peak error. The peak error occurs when the current stops and fast dynamics occur, which is why there is a substantial reduction in error with the 2RC pair model. The 2RC model has a slightly lower internal resistance, because the faster dynamics capture some of

the charge transfer region, meaning the ohmic resistance is closer to the DC resistance. Accordingly, R_0 in the 1RC model is approximately equal to the sum of R_0 and R_1 in the 2RC model. The bandwidth requirements differ for parallel connected cells, as discussed in section 4.1.

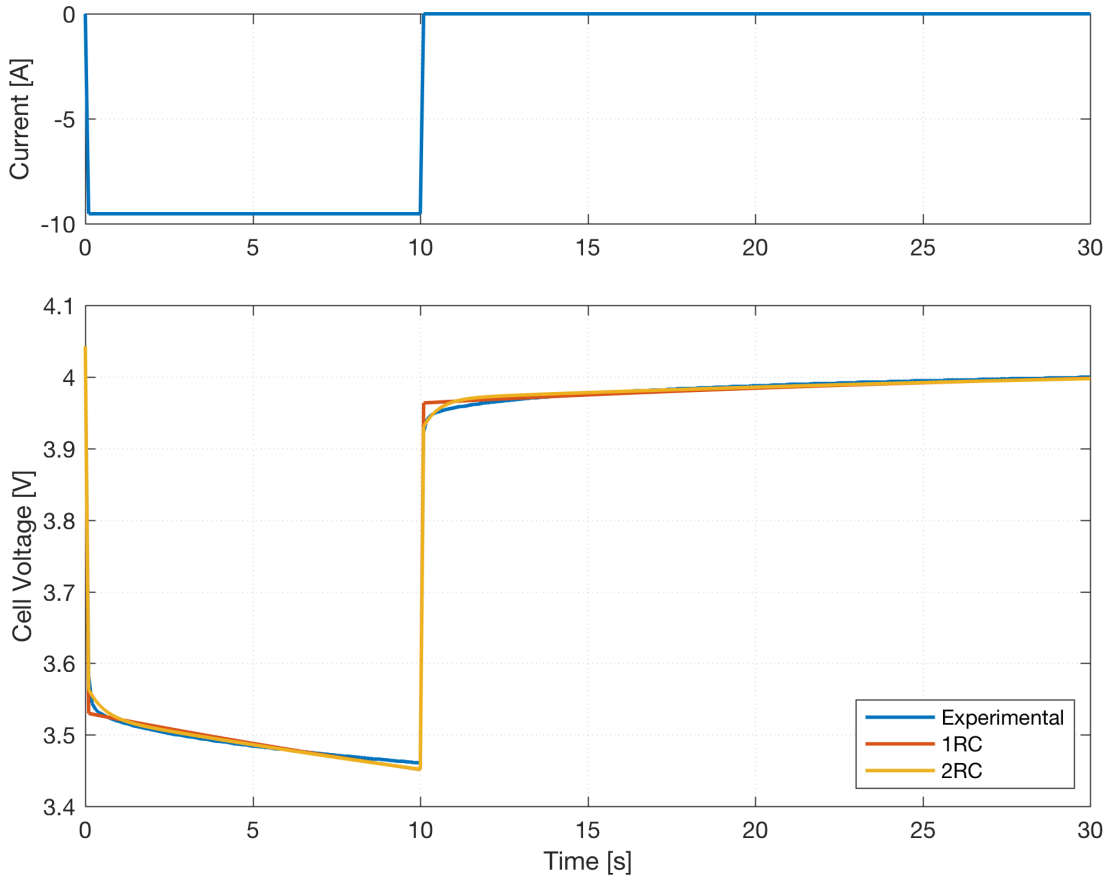

Figure 17: Example pulse power fit results for 1 and 2 RC ECMs

Table 3: Summary of ECM parameterisation results

Model Order	RMS Error (mV)	Peak Error (mV)	C_0 (F)	R_0 (mΩ)	R_{p1} (mΩ)	τ_{p1} (s)	R_{p2} (mΩ)	τ_{p2} (s)
1	4.34	51.2	5340.3	53.9	19.0	24.2	-	-
2	3.02	17.8	6660.1	49.9	4.1	0.48	22.3	30.6

3.4.2 Example Validation Results

The standard Catapult[4] validation profile is used to evaluate the performance of the ECMs, shown in Figure 18. This is a dynamic drive cycle over a narrow SOC window, which is derived from EV driving data and as such is representative of how cells are used within a battery pack. The ECM is used to simulate the cell

[4] Catapult referring to the WMG centre HVM Catapult, a joint agreement between WMG, JLR and TMETC

voltage using the current profile, and voltages compared. Example results are shown in Figure 19.

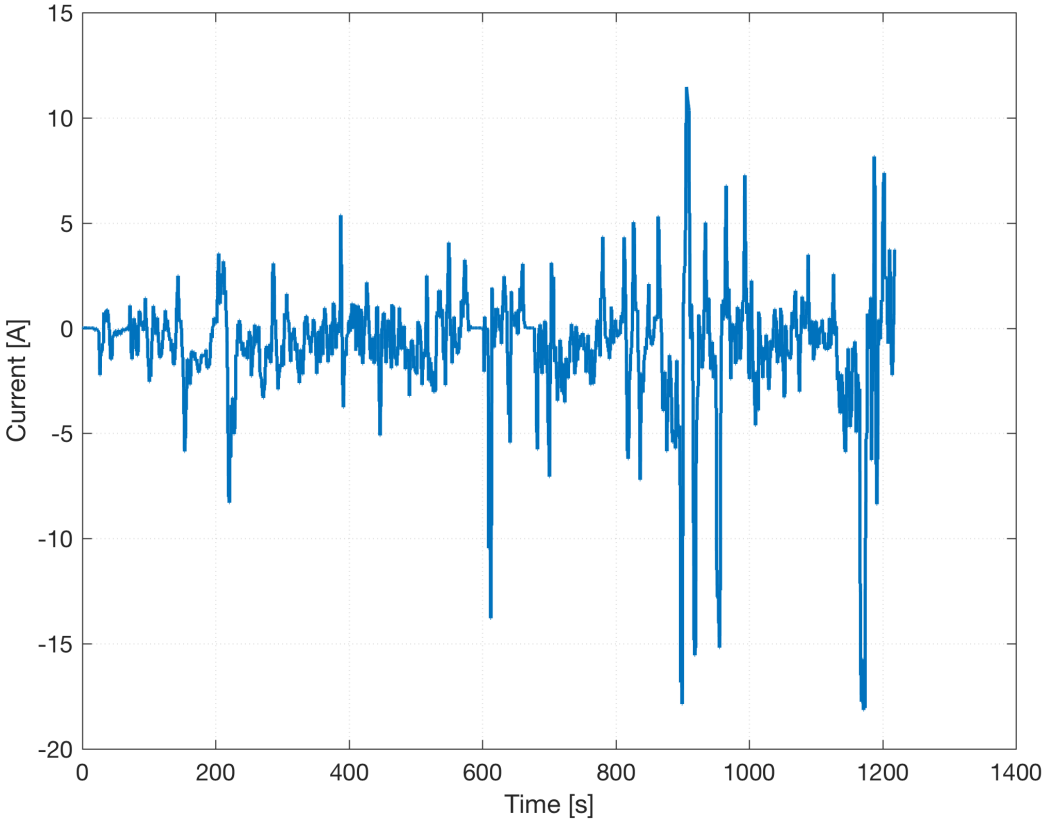

Figure 18: Catapult Validation Current Profile

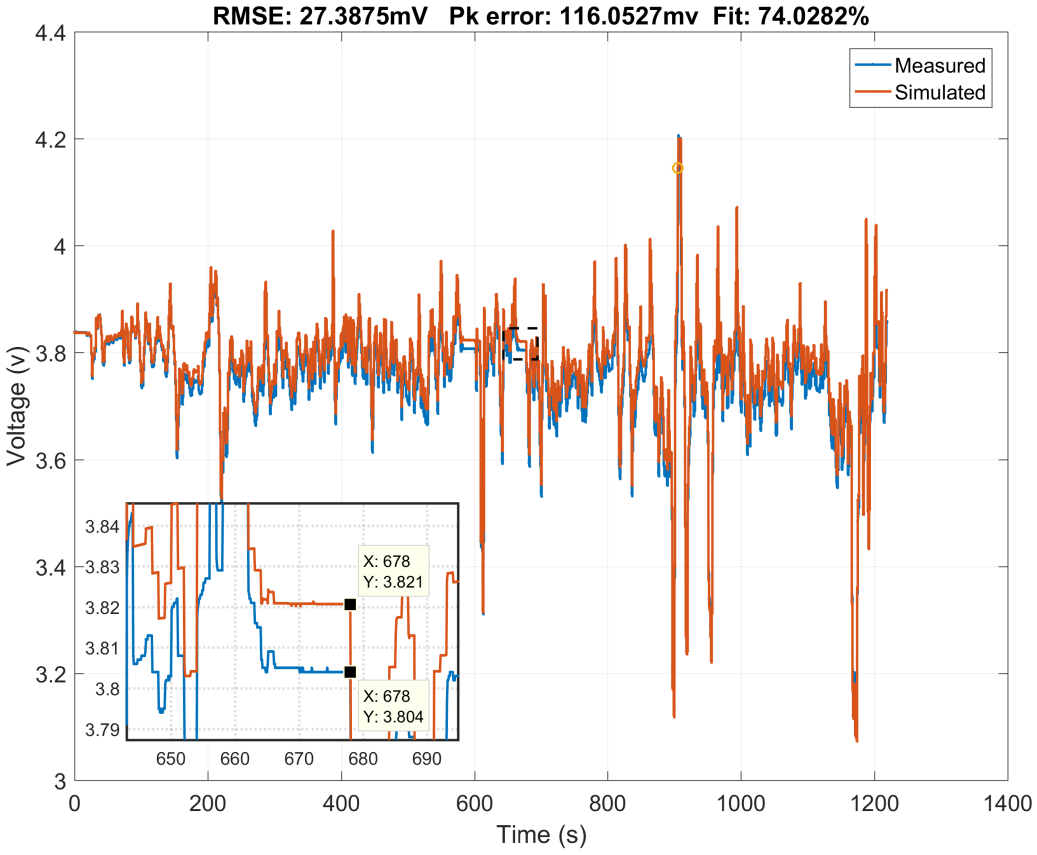

Figure 19: Measured and simulated cell voltage during validation profile

One observation from this data is that a constant between measured and simulated voltage is apparent: the orange simulated line is generally about 20mV above the blue measured line, when the dynamics are filtered out. This appears to be an error with the OCV-SOC data, and is discussed further in Submission 5. The simulation was initialised by using the OCV-SOC curve and relaxed cell voltage prior to the test to initialise the SOC. It appears that the OCV values interpolated during simulation are higher than the actual cell's OCV. An averaged OCV-SOC curve was used for the simulation, but factoring in separate OCV-SOC curves only removes a small amount of the error. It suggests that further work is required on improving SOC-OCV methodology and verifying the data.

3.5 Summary

This section has shown the ways in which ECMs can be created. The parameterisation method is flexible: arbitrary current and voltage signals can be used to fit a model. This is well-suited to on-vehicle applications, where the parameters can be updated during operation without causing any disruption. However, ECMs also have some flaws. They fundamentally do not capture the cell response in the low frequency region of operation, resulting in poor accuracy over wide DOD operation. There also appears to be some errors generated during the OCV-SOC curve procedure. This only became apparent when running the SOC estimators in real-time for analysing imbalance, where requirements of the model are different to much of the work within Catapult to date. This is an important area of further work, and should be quantified experimentally as well as using the data as part of SOC estimators to evaluate the impact the data has on estimator performance. Nevertheless, ECMs have proved useful for a number of applications throughout this research, as discussed in the next section.

4 System Modelling

The ECM variants in section 3.2 are well established within the literature. However, there are few studies simulating several cells in series or parallel to better evaluate imbalance, which Research Objective 3 addresses by building a modelling and simulation framework around the balancing system. Models of balancing systems in the literature tend to be at a very high level [108] or, conversely, are in-depth models of specific high-frequency components to study particular hardware elements such as converters [91, 94, 105, 131]. The derivation and use of control-orientated models has received significantly less attention.

As noted in Chapter 3, there are three main application areas considered for this research where ECMs are central to wider system modelling. In section 4.1, the question of imbalance within parallel-connected cells is addressed through a combination of experimental and simulation work. Evaluating whether parallel-connected cells require a balancing system closes off Research Objective 1. Section 4.2 shows how a battery pack can be modelled, enabling closed loop balancing simulations which are essential to meeting Research Objective 4. A means of estimating SOC is introduced in section 4.3. Unlike for simulations, an exact SOC measurement is not available during real-time experimental work. An accurate SOC estimate is critical for meeting Research Objective 5 and so answering Research Question 2.

4.1 Imbalance within Parallel-Connected Cells

Cells in a battery pack may be electrically connected in parallel to increase the pack capacity and meet requirements on power and energy [132, 133]. For example, the Tesla Model S 85kWh battery pack uses 74 3.1Ah cylindrical cells to create a parallel unit, and 96 of these units in series. The Nissan Leaf 24kWh battery pack consists of 33Ah cells, with 2 in parallel and 96 in series [134]. The nature of a parallel connection means each cell terminal voltage is the same, and the applied current equals the sum of the individual cell currents.

One of the open questions from the literature review in Chapter 2 is whether cells connected in parallel are balanced. There was very little research into this at the time, with the consensus being that energy balancing is only required for cells in series [17, 66] since the cells in a parallel unit are inherently balanced due to the common voltage [58, 67]. However, while the voltage over each cell is the same, there can be differences in the OCV, which equates to differences in SOC. The same variations which cause imbalance between cells in series (such as impedance and capacity) also have an effect on battery pack performance when cells are connected in parallel.

Gogoana et al. [135] cycle-aged two cylindrical LFP cells connected in parallel. They found that a 20% difference in internal resistance resulted in a 40% reduction in the useful life of the pair of cells compared to if the cells had approximately equal internal resistances. The authors attribute this due to the uneven current distribution between the cells: each cell will go through periods where it sees high currents which age the cell more quickly. Gong et al. [132] drew similar conclusions from their experimental work with 32Ah cells. When two cells with a 20% impedance difference were connected in parallel, the peak current seen was 40% higher than the nominal current (which would be seen if the cells were identical). The authors also performed some simulation studies, using the Simscape toolbox in Simulink to connect two ECMs in parallel. This is one of the few examples of parallel cell modelling in the literature. Typically a parallel unit 'lumped' model is created, where the parameters of a single cell model are scaled to create an effective parallel unit [16]. While this allows for high-level simulation of a battery pack, it assumes the cells within a parallel unit are identical and as such undergo the same current. It is argued that this approach is not sufficient given the cell-to-cell variations which can be expected through ageing and temperature differences discussed in section 2.1.2).

4.1.1 Parallel Equivalent Circuit Model

An ECM can be expanded to incorporate any number of cells in parallel using mesh analysis. For this derivation, four cells with one RC pair each is used, but the solution can be extended to any number of cells or RC pairs. Figure 20 is a schematic of four cells in parallel with an interconnection resistance (R_{conn})

between each terminal. Currents i_1 to i_4 are the currents around each loop; i_1 is known since it is the applied current and so an input to the model.

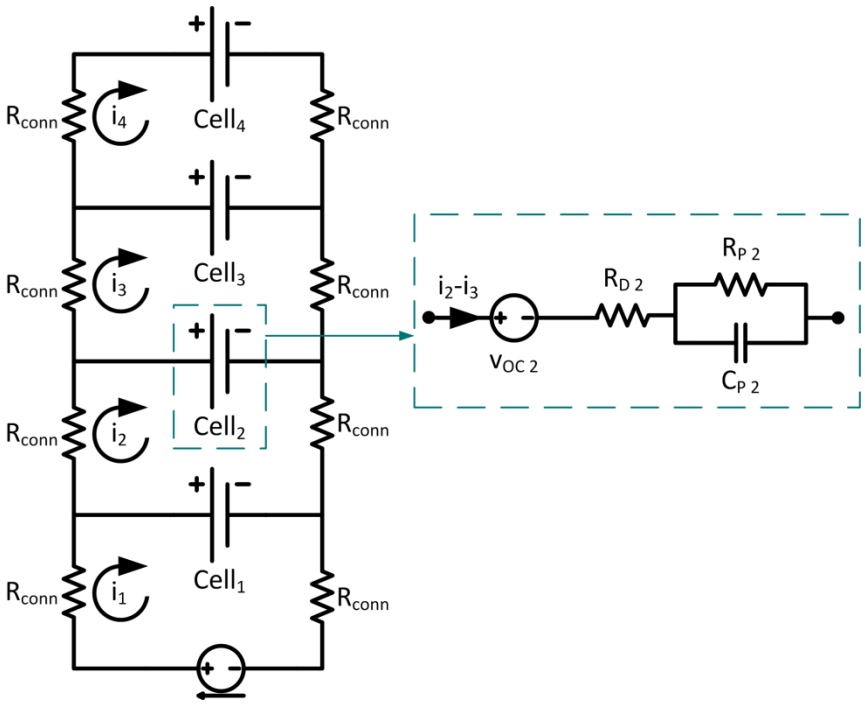

Figure 20: Schematic of four cells connected in parallel

In addition to the state equation for the single cell from Chapter 3, there are algebraic constraints on the parallel cell system, based on Kirchoff's laws for current and voltage. The currents at a junction must sum to zero, which for the parallel cell system occurs at the cell connections, as shown in the left-hand diagram of Figure 21. For *N* cells, the cell current is given by (28). The two cases arise due to there being no i_{n+1} for the final cell. Similarly, the voltages around a loop must sum to zero. A typical loop is shown in the right-hand diagram of Figure 21; there are *N-1* loops to solve for, as the first loop incorporates the current source and is therefore trivial.

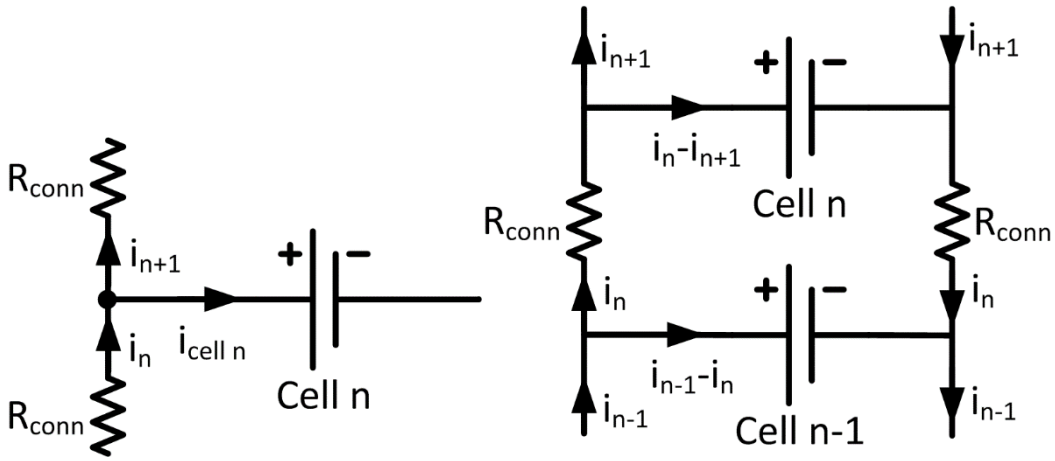

Figure 21: Current junction (left-hand) and voltage loop (right-hand) in the parallel cell model

The voltage loop equation is given by (29), and is expanded to (30) by decomposing the cell voltage into its components using (28). The right-hand side of (30) consists only of unknown currents, and the left side is comprised of state voltages (whose value at any given time-step does not depend on the current values at that step, avoiding an algebraic loop). With one equation per unknown current, (30) can be arranged into a system of equations and solved using matrix algebra.

$$i_{cell\,n} = \begin{cases} i_{loop,n} - i_{loop,n+1}, & n < N \\ i_{loop,n}, & n \equiv N \end{cases} \tag{28}$$

$$v_{t\,n} - v_{t\,n-1} + 2R_{conn} i_{loop,n} = 0 \tag{29}$$

$$\begin{aligned} &v_{OC\,n} - v_{OC\,n-1} + v_{p\,n} - v_{p\,n-1} \\ &= \begin{cases} -R_{o\,n}\left(i_{loop,n} - i_{loop,n+1}\right) + R_{o\,n-1}\left(i_{loop,n-1} - i_{loop,n}\right) - 2R_{conn} i_{loop,n}, & n < N \\ -R_{o\,n} i_{loop,n} + R_{o\,n-1}\left(i_{loop,n-1} - i_{loop,n}\right) - 2R_{conn} i_{loop,n}, & n \equiv N \end{cases} \end{aligned} \tag{30}$$

A state vector $\boldsymbol{x}$ is defined in equation (31), containing all the cell voltage states. Equation (32) uses this vector to rewrite (30) in matrix form. This requires three more coefficients to be defined. The R_{ps} matrix contains the resistances from the right-hand side of (30), the E_{ps} matrix collates the left-hand elements of (30), and the $\boldsymbol{f_{ps}}$ vector factors the input current to the loop current calculation.

$$\boldsymbol{x} = [v_{oc_1} \quad v_{p_1} \quad v_{oc_1} \quad v_{p_1} \quad \cdots \quad v_{p_N}]^T \tag{31}$$

$$R_{ps} \begin{bmatrix} i_1 \\ i_2 \\ \vdots \\ i_N \end{bmatrix} = E_{ps}\boldsymbol{x} + \boldsymbol{f_{ps}} i_1$$

$$R_{ps} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -\left(R_{o_1} + R_{o_2} + 2R_{conn}\right) & R_{o_2} & 0 \\ 0 & R_{o_2} & \ddots & R_{o_{N-1}} \\ 0 & 0 & R_{o_{N-1}} & -\left(R_{o_{N-1}} + R_{o_N} + 2R_{conn}\right) \end{bmatrix} \tag{32}$$

$$E_{ps} = \begin{bmatrix} 0 & 0 & \cdots & & & & \cdots & 0 \\ -1 & -1 & 1 & 1 & 0 & & \cdots & 0 \\ 0 & \cdots & -1 & -1 & 1 & 1 & 0 & 0 \\ 0 & \cdots & & 0 & -1 & -1 & 1 & 1 \end{bmatrix}$$

$$\boldsymbol{f_{ps}} = \begin{bmatrix} 1 \\ -\left(R_{o_1} + 2R_c\right) \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Equation (33) introduces another matrix G_{ps}, which applies (28) in matrix form. This is combined with (32) to produce (34).

$$\begin{bmatrix} i_{cell\ 1} \\ i_{cell\ 2} \\ \vdots \\ i_{cell\ N} \end{bmatrix} = G_{ps}\boldsymbol{i_{loop}}$$
$$G_{ps} = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & \ddots & 1 & -1 \\ 0 & \cdots & 0 & 1 \end{bmatrix} \tag{33}$$

$$\boldsymbol{i_{cell}} = G_{ps}R_{ps}^{-1}\left(E_{ps}\boldsymbol{x} + \boldsymbol{f_{ps}}i_{loop,1}\right) \tag{34}$$

This demonstrates that each cell current can be calculated using only the measured current input and the cell states. The cells are now effectively decoupled from one another: each current can now be used an input to an individual ECM, rather than solving for all cell states in one system. The current and voltage constraints have already been enforced, meaning that any cell model within the parallel stack can be used to calculate terminal voltage as they are all equal. This means parallel cells can easily be integrated into existing models which consider a single cell, or cells in series. The cells can be combined into one complete state-space system, which can be useful for analysing system dynamics and stability. A full derivation of this model is provided within Submission 2 and forms the foundation of the author's publication [10].

4.1.2 Experimental Data

The purpose of this experimental work was to evaluate the behaviour of cells with different properties when connected in parallel, and validate the model in section 4.1.1. For this, four new 3Ah commercial lithium-ion cells were employed. There were three stages to this work. Firstly, the differences in properties had to be created by ageing the cells. Secondly, characterisation tests were performed so that the model could be parameterised. Finally, test data of the cells connected in parallel was collected to validate the model and for additional insight into cell

behaviour. A full description and analysis of the tests is contained with Submission 2 and the research paper [10].

To create differences in cell properties, each of the four cells was aged by a different amount. Each ageing cycle consists of a 1C (3A) discharge to 2.5V followed by a CC-CV charge at C/2 (1.5A) and 4.2V. The four cells were each aged by different amounts in order to ensure that each cell would have significantly different impedance and capacity values. After every 50 cycles, a 1C capacity check and electrochemical impedance spectroscopy (EIS, explained further in Appendix A) tests at 20%, 50% and 90% SOC were performed. Additionally, pulse power tests were performed on the new and aged cells at 20%, 50%, 80% and 95% SOC and are summarised in Table 4. Full ECMs and their component values are discussed in section 4.1.3. The cell capacity decreases linearly with the number of ageing cycles, and the resistance increases linearly, which are both expected from the ageing process. This demonstrates that the cells have substantially different characteristics and so when connected in parallel will not behave in a homogenous manner.

Table 4: Parallel Cell Ageing Results

Cell	Ageing Cycles	1C Discharge Capacity (Ah)	Resistance (mΩ)
Nominal	-	3.07	48.1
1	0	2.94	46.3
2	50	2.86	51.2
3	100	2.77	58.2
4	150	2.65	63.1

4.1.3 Model Validation

Several different ECMs were parameterised, as discussed later in this section, with the aim of analysing the impact of model order on overall accuracy. For all parallel cell models, the resistance was not a function of current direction. This is because resistance is used to calculate current (and so its direction), creating an algebraic loop.

As with single cell ECMs in Chapter 3, the simulated voltage was compared to the measured data to evaluate performance. However, the model's cell current

prediction also required evaluation. To do this, the validation test was performed on all four cells connected in parallel, with the individual cell currents measured. To measure the cell currents, a 10mΩ shunt resistor was connected in series with each cell, and its voltage logged using a dSPACE MicroAutobox[5], which can sample at 100Hz. The cell voltage was also logged, since it would be synchronised with the currents, rather than having to combine the dSPACE and cell cycler datasets. The shunt resistors were calibrated by passing a series of known currents through them using a bench power supply, and measuring the corresponding voltages over the resistors. An example voltage trace for various currents is shown in Figure 22.

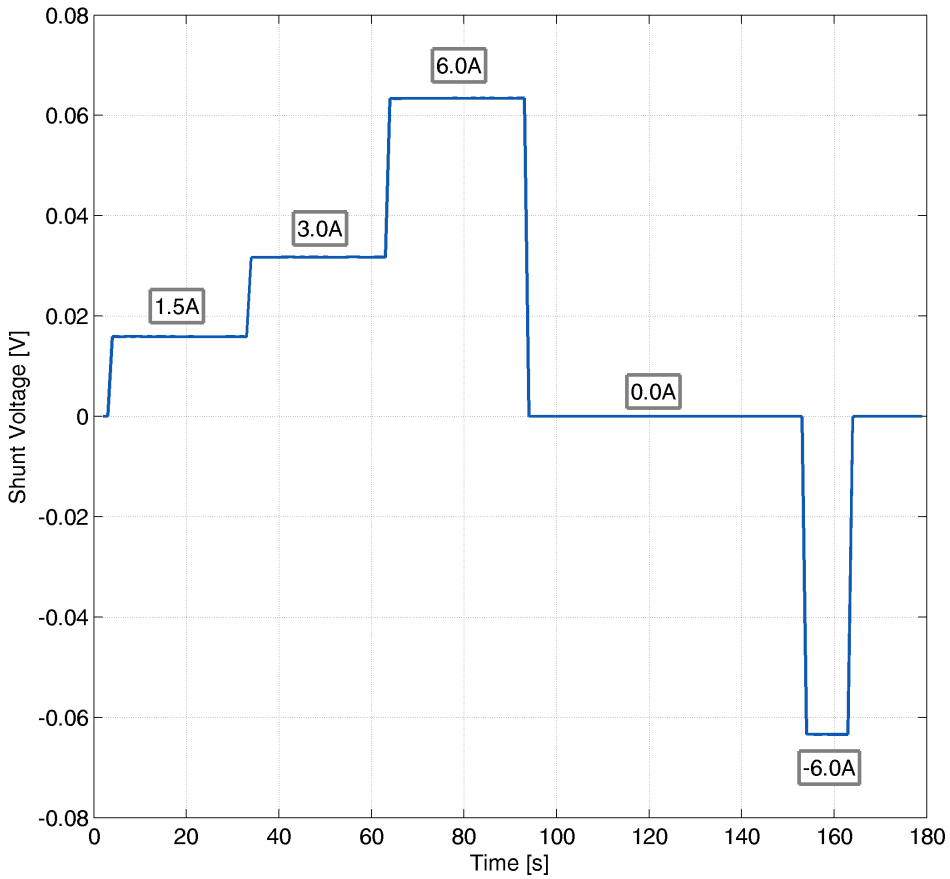

Figure 22: Example profile used to calibrate shunt resistors

[5]The dSPACE data acquisition accepts a 0-5V input voltage, whereas the resistor produces a ±250mV voltage. This voltage had to be offset and amplified for it to be measured, the details of which are in Submission 2.

Table 5: Validation of current shunt measurements

Applied Current (A)	6	12	24	0	-24	-12	-6
Estimated current (A)	5.9981	11.998	24.011	4.2E-4	-24.005	-12.007	-6.0053
Error (%)	-0.03	-0.02	0.05	-	0.02	0.06	0.09

A linear last squares fit was performed for each current shunt, and validated by applying currents to all four cells, and comparing the measured current to the sum of each current estimated from the fit. The validation results, in Table 5 show the shunt measurements were accurate to a worst case of 0.1%.

The standard ECM discussed in section 3.3 was used, with different numbers of RC pairs used as discussed in the remainder of this section. The cells were parameterised according to the method described in section 3.4, and the standard Catapult validation drive cycle from Figure 18 was used. Originally, the individual cell pulse power obtained from the Bitrode MCV16-100-5 cell cycler [136] was used. The simulation results with 2 RC pairs showed that while the RMS voltage error of 25mV was in line with typical Catapult data [137], the current estimation was inaccurate. Figure 23 shows a snapshot of the measured and simulated currents during the validation drive cycle. The current for cells 3 and 4 (the green and grey lines) are significantly different from the measured values.

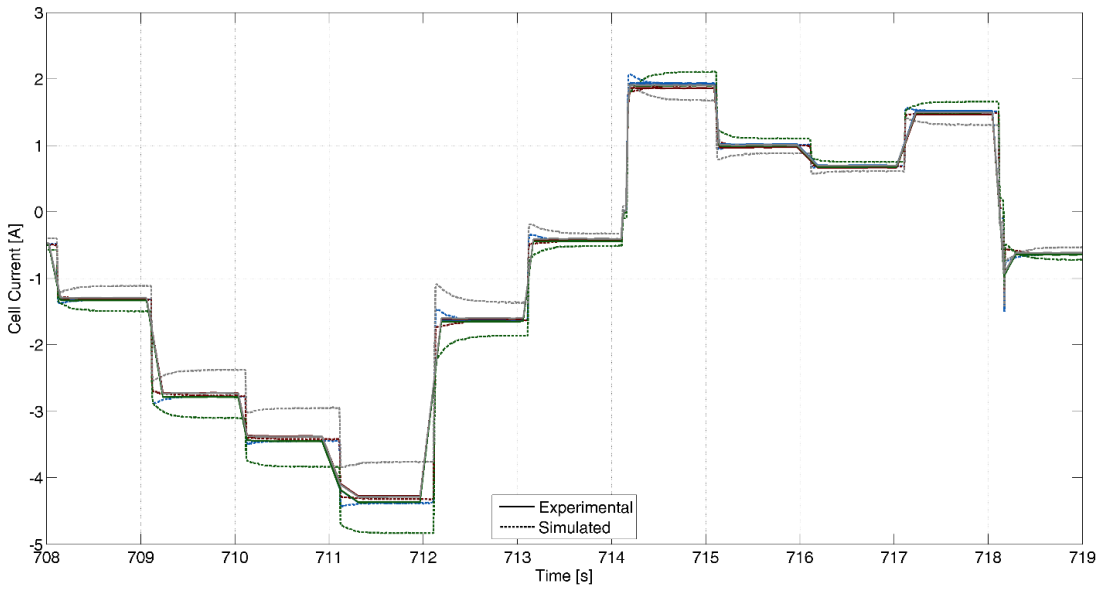

Figure 23: Experimental and simulated current for four cells in parallel using 10Hz parameterisation data

Section 4.1.1 shows that the current calculated by the model is strongly linked to the ohmic resistance value. As discussed in Appendix A, the lower the frequency, the further the ohmic resistance is from the DC resistance. A higher sample rate should result in a more accurate ohmic resistance estimates and so a more accurate current calculation.

Limited availability of test equipment at this point meant that a full re-characterisation could not be performed. Instead, Two pulse power tests were performed on the parallel unit of cells, at 10% SOC and 30% SOC[6] and used to fit the models. While this is a limited dataset with missing high SOC data, it still performed well as discussed below. It is believed that this is because the ECM parameters are largely consistent at mid-high SOCs, but do change considerably at lower SOCs, e.g. 20% and below. This limited dataset essentially captures the two main regions of impedance of the cell and as such was prove to be adequate for the tests performed here.

Re-parameterisation from pulse power tests at 100Hz performed much better. The Bitrode cell cycling equipment, used to collect the original test data, can only sample at up to 10Hz, To sample at 100Hz, a dSPACE Autobox was used to log cell voltage and current, and the Bitrode used only to control the applied load. Figure 24 is a repeat of Figure 23 but using a 3 RC model. The higher bandwidth allowed for more time constants: it was found that fitting a 3RC model to 10Hz data did not significantly improve performance.

[6] In Submission 2, it was erroneously stated that EIS data was used as a basis to extrapolate the new pulse power data across the 20%, 50% and 90%. This was used for a period of time but it was difficult to reliably extrapolate as an ECM from EIS data is not equivalent to an ECM from pulse power data because of the differences in bandwidth and signal.

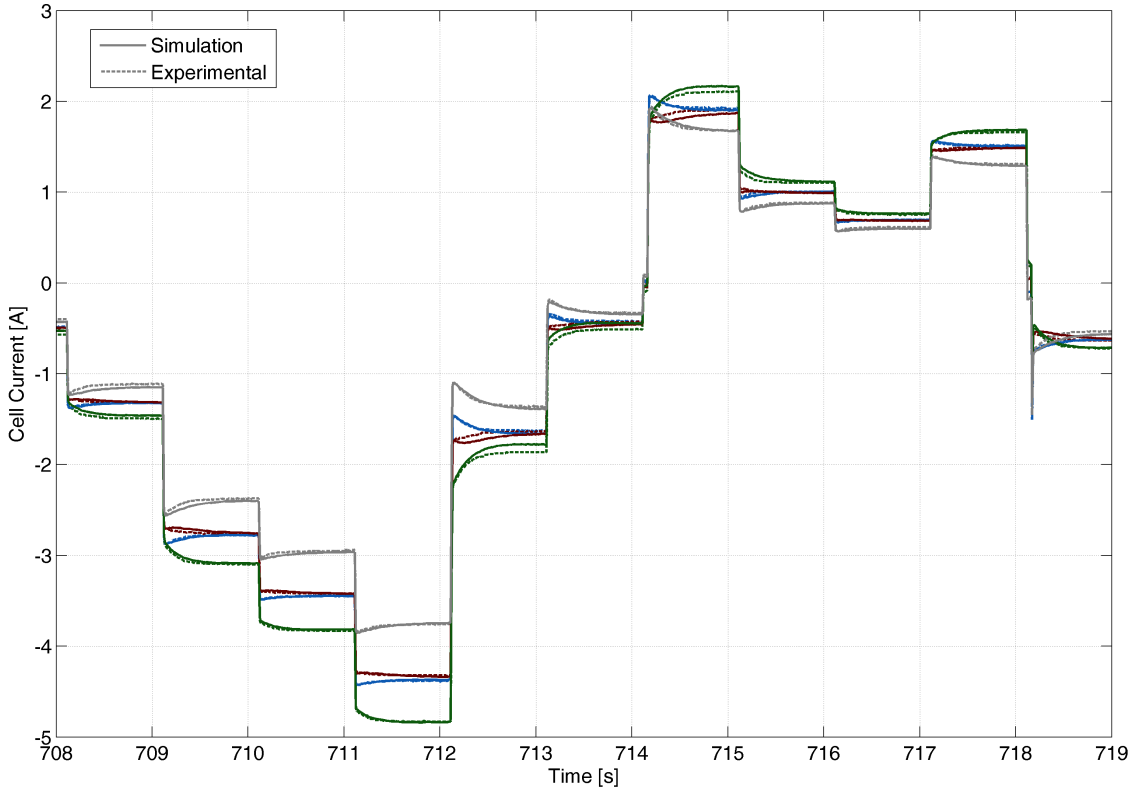

Figure 24: Experimental and simulated current for four cells in parallel using 100Hz parameterisation data

A complete summary of model performance for different sample rates and RC pairs is given in Table 6. The 2RC model using the 10Hz data is less accurate than a 1RC model using the 100Hz data, so the sampling rate of the parameterisation data is important as well as model fidelity. For current, the RMS and peak error both decrease as the number of RC pairs increases. The third RC pair does not reduce RMS error much, but does considerably reduce peak error. Terminal voltage error appears to be largely independent of the number of RC pairs. This may be because the limiting factor is to do with other aspects such as incorporating charge-discharge hysteresis and OCV-SOC data errors. The complete parameterisation datasets for all variants of the model are given in Table 7. Generally, for a given model order and similar SOC, the Bitrode data results in a similar R_0 for each cell, with the RC pair values creating larger impedance differences, whereas the dSPACE data produces more distinct R_0 values for each cell.

Table 6 : Comparison of error between simulated and experimental data for various models

Sample Rate	100Hz			10Hz	
RC Pairs	1	2	3	1	2
RMS Error- Current (mA)	60.43	40.24	34.59	98.74	68.72
Peak Error- Current (mA)	519.5	403.7	266.7	730.9	538.6
RMS Error- Voltage (mV)	16.38	17.90	17.88	25.32	25.72
Peak Error- Voltage (mV)	92.8	100.8	89.3	154.5	145.8

Table 7: Parameterisation Dataset for parallel cell ECMs

Source	Cell	SOC (%)	RC Pairs	R_0 (mΩ)	R_{p1} (mΩ)	τ_{p1} (s)	R_{p2} (mΩ)	τ_{p2} (s)	R_{p3} (mΩ)	τ_{p3} (s)
dSPACE	1	10	1	81.7	42.2	8.9	-	-	-	-
dSPACE	1	10	2	71.8	92.0	1.8	21.1	19.6	-	-
dSPACE	1	10	3	71.8	132.5	2.4	0.3	18.7	37.1	36.1
Bitrode	1	20	1	62.9	36.7	14.7	-	-	-	-
Bitrode	1	20	2	57.1	57.6	2.3	29.0	59.6	-	-
dSPACE	1	30	1	67.7	12.0	15.3	-	-	-	-
dSPACE	1	30	2	60.7	5.9	0.0	11.8	17.1	-	-
dSPACE	1	30	3	60.2	6.0	0.0	4.9	8.1	14.2	46.0
Bitrode	1	50	1	47.0	16.2	26.1	-	-	-	-
Bitrode	1	50	2	46.3	8.4	12.3	10.3	73.1	-	-
Bitrode	1	90	1	53.6	27.4	33.4	-	-	-	-
Bitrode	1	90	2	52.0	7.7	7.4	23.9	68.2	-	-
dSPACE	2	10	1	66.6	165.7	5.7	-	-	-	-
dSPACE	2	10	2	59.6	23.8	0.4	16.5	12.5	-	-
dSPACE	2	10	3	61.6	20.9	0.6	26.3	20.5	43.0	222.8
Bitrode	2	20	1	62.6	61.2	17.0	-	-	-	-
Bitrode	2	20	2	58.6	89.8	3.3	29.7	54.0	-	-
dSPACE	2	30	1	60.1	20.0	0.5	-	-	-	-
dSPACE	2	30	2	55.9	19.9	0.2	13.9	12.3	-	-
dSPACE	2	30	3	55.2	16.9	0.2	5.4	2.4	17.6	30.0
Bitrode	2	50	1	55.0	22.1	27.2	-	-	-	-
Bitrode	2	50	2	52.7	7.9	5.5	20.9	67.9	-	-
Bitrode	2	90	1	66.8	44.2	33.4	-	-	-	-
Bitrode	2	90	2	47.8	13.4	3.0	42.4	68.1	-	-
dSPACE	3	10	1	79.9	210.7	6.1	-	-	-	-
dSPACE	3	10	2	68.2	27.9	0.4	30.1	21.7	-	-
dSPACE	3	10	3	67.1	27.8	0.4	28.8	13.2	1.7	24.2

Bitrode	3	20	1	61.4	44.9	13.0	-	-	-	-
Bitrode	3	20	2	55.9	76.7	2.8	31.9	66.9	-	-
dSPACE	3	30	1	73.4	11.4	6.6	-	-	-	-
dSPACE	3	30	2	62.2	12.9	0.1	14.0	13.0	-	-
dSPACE	3	30	3	61.5	11.8	0.1	4.5	3.3	17.7	35.4
Bitrode	3	50	1	50.2	17.5	24.3	-	-	-	-
Bitrode	3	50	2	48.6	9.0	9.8	16.7	82.1	-	-
Bitrode	3	90	1	59.2	32.5	32.3	-	-	-	-
Bitrode	3	90	2	50.2	9.2	4.0	30.5	56.4	-	-
dSPACE	4	10	1	72.3	94.2	3.2	-	-	-	-
dSPACE	4	10	2	65.7	28.6	0.5	33.9	18.4	-	-
dSPACE	4	10	3	68.7	33.3	0.9	0.9	15.9	185.4	422.6
Bitrode	4	20	1	62.3	66.4	15.2	-	-	-	-
Bitrode	4	20	2	58.0	99.5	3.6	35.2	57.8	-	-
dSPACE	4	30	1	67.5	30.5	0.7	-	-	-	-
dSPACE	4	30	2	61.8	28.2	0.3	15.4	8.6	-	-
dSPACE	4	30	3	60.7	22.6	0.2	15.4	7.8	20.3	33.0
Bitrode	4	50	1	58.4	26.3	27.2	-	-	-	-
Bitrode	4	50	2	52.9	9.9	3.8	25.3	63.6	-	-
Bitrode	4	90	1	71.0	48.0	34.8	-	-	-	-
Bitrode	4	90	2	49.3	22.1	0.5	63.8	58.5	-	-

4.1.4 Analysis of Cells connected in Parallel

This section uses the validated model to simulate various scenarios. These results, along with experimental data, are used to evaluate the implications for battery performance and design when using cells connected in parallel.

4.1.4.1 Current Distribution

Figure 25 shows the experimentally measured cell currents and temperatures during a 1C discharge from fully charged until the lower voltage cut off was reached. The discharge rate is calculated using the capacity of a new cell, making it over 1C for the aged cells, and as a result the discharge time is less than one hour. The temperatures were measured using a T-type thermocouple attached to cell exterior, midway along its length. There is a substantial change in current distribution after 2800s, as the cells enter the low SOC region. By the end of discharge, the most aged cell (cell 4) is undergoing the largest current, and the least aged cell (cell 1) is undergoing the smallest current.

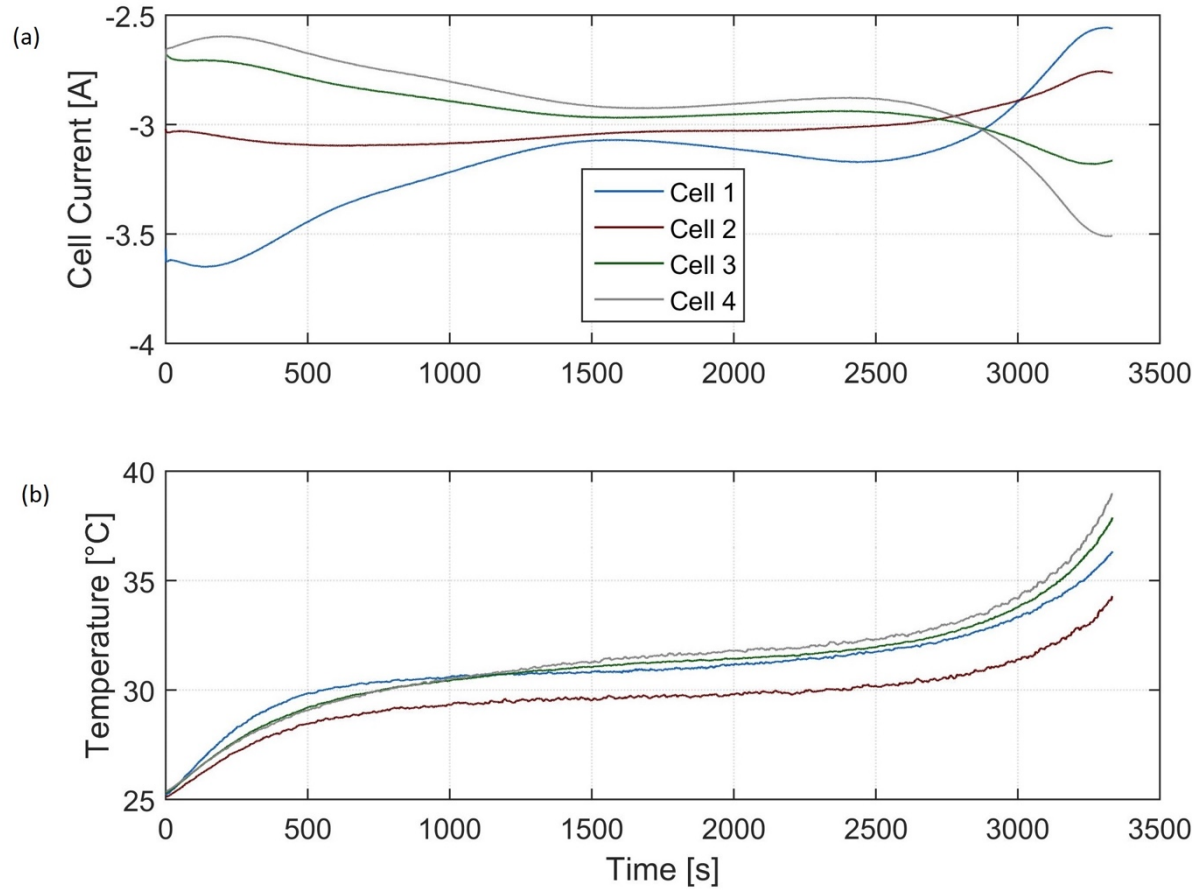

Figure 25: Experimentally measured current and temperature during a 1C constant discharge

This phenomenon has been previously reported in [135], where the authors attributed it to the increase in impedance that occurs at low SOCs. Less aged cells, which have previously undergone more current, will be at a lower SOC (although this in turn depends on their relative capacities) and so eventually their impedance will increase beyond the impedance of the more aged cells. This causes the more aged cells to see larger currents towards the end of the discharge. In addition to this, the OCVs of some cell chemistries decrease more significantly at low SOCs. This means that during discharge the OCV of lower SOC cells will be significantly closer to the terminal voltage and so less current is able to flow. Equation (35) approximates the cell current by rearranging (12) with RC pair voltages removed for simplicity. This shows the resistance rise and lower OCV both work to reduce the current through low SOC cells.

$$i_{cell} = \frac{v_t - v_{OC}}{R_0} \tag{35}$$

Clearly, the current is quite dynamic and it is an oversimplification to assume that the distribution is fixed and that the most aged cell will always undergo the

smallest current. These results highlight that this is application dependent and relative to the magnitude, duration and frequency content of the electrical load current. For example, a HEV battery pack is typically maintained within a narrow SOC window centred about 50% SOC, meaning it won't undergo the same low-SOC current switching as a BEV may do.

The cell's temperature change during discharge is known to be a function of impedance and current [138]. The lower plot of Figure 25 shows the measured cell surface temperature corresponding to the discharge profile in the upper plot. The ambient temperature was regulated at 25°C. At first, cell 1 undergoes the greatest temperature rise due to the larger cell current. However, after 1000s the temperature of cell 4 surpasses that of cell 1. By this point there is not as much difference in cell currents, and so the higher impedance of cell 4 causes greater heat generation despite the lower current. The temperature of cell 2 is consistently the lowest of all of the cells because it has a relatively low impedance without undergoing the large currents experienced by cell 1, which combine to produce relatively low heat generation.

4.1.4.2 Drive Cycle Type

As mentioned above, a typical of drive cycle will impact the current distribution. A simulation of the validation drive cycle, which is typical of HEV usage, is shown in Figure 26. A simulation of the ageing cycle from section 4.1.2 was performed, but with the current magnitude doubled. The drive cycle finishes at 1219s, but the simulation is continued with no applied load to observe the cell balancing.

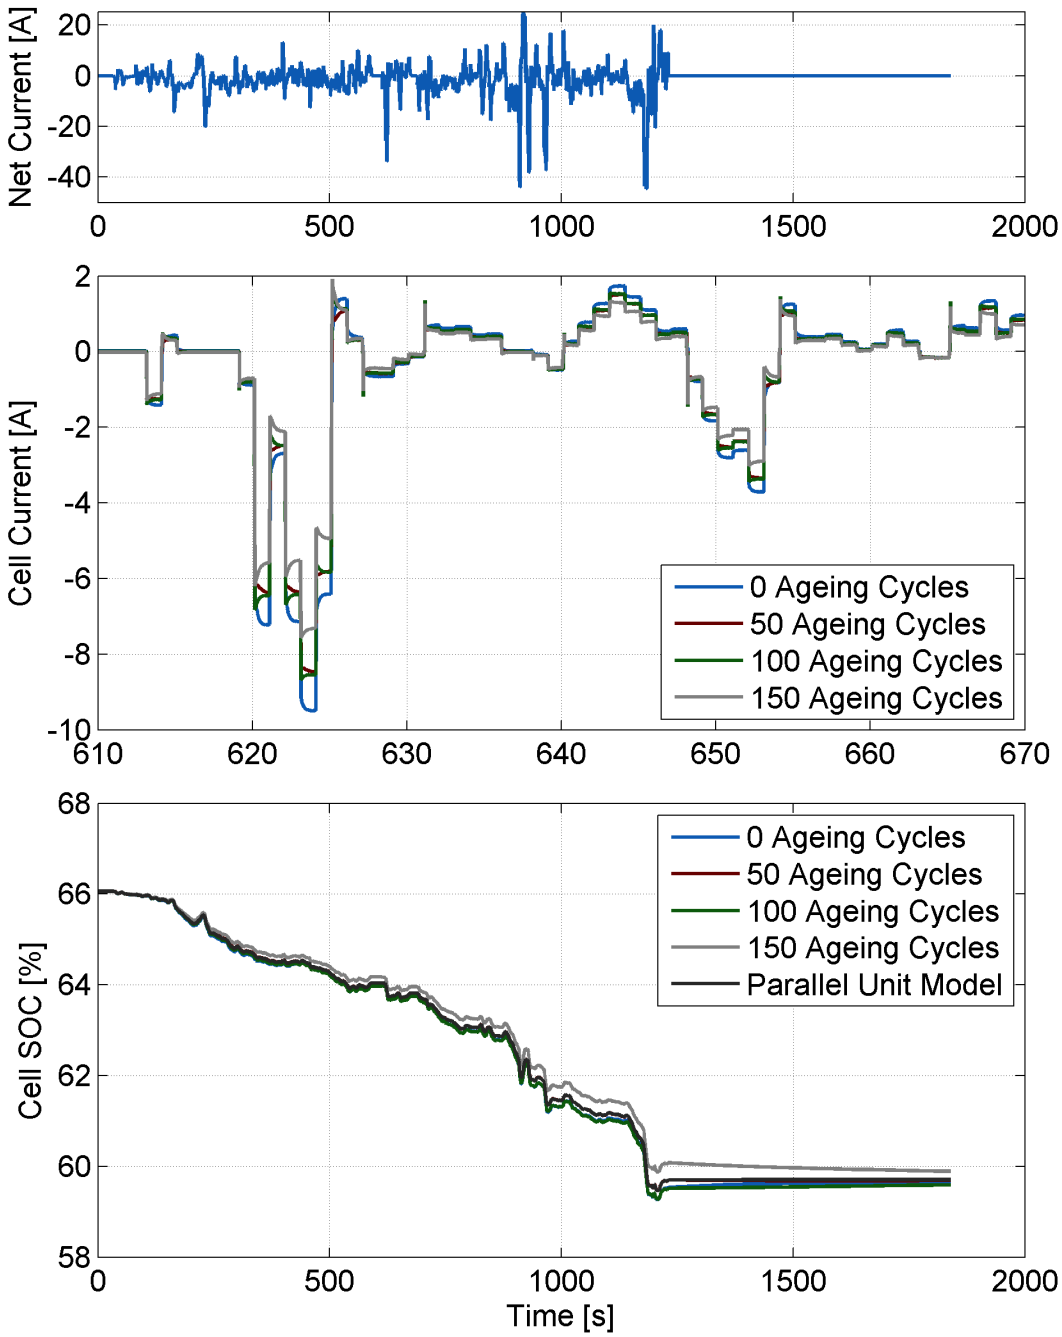

Figure 26: Simulation results for the validation drive cycle

A discharging-charging profile with the moderate current and wide DOD is representative of a BEV was also simulation, and the results are shown in Figure 27.

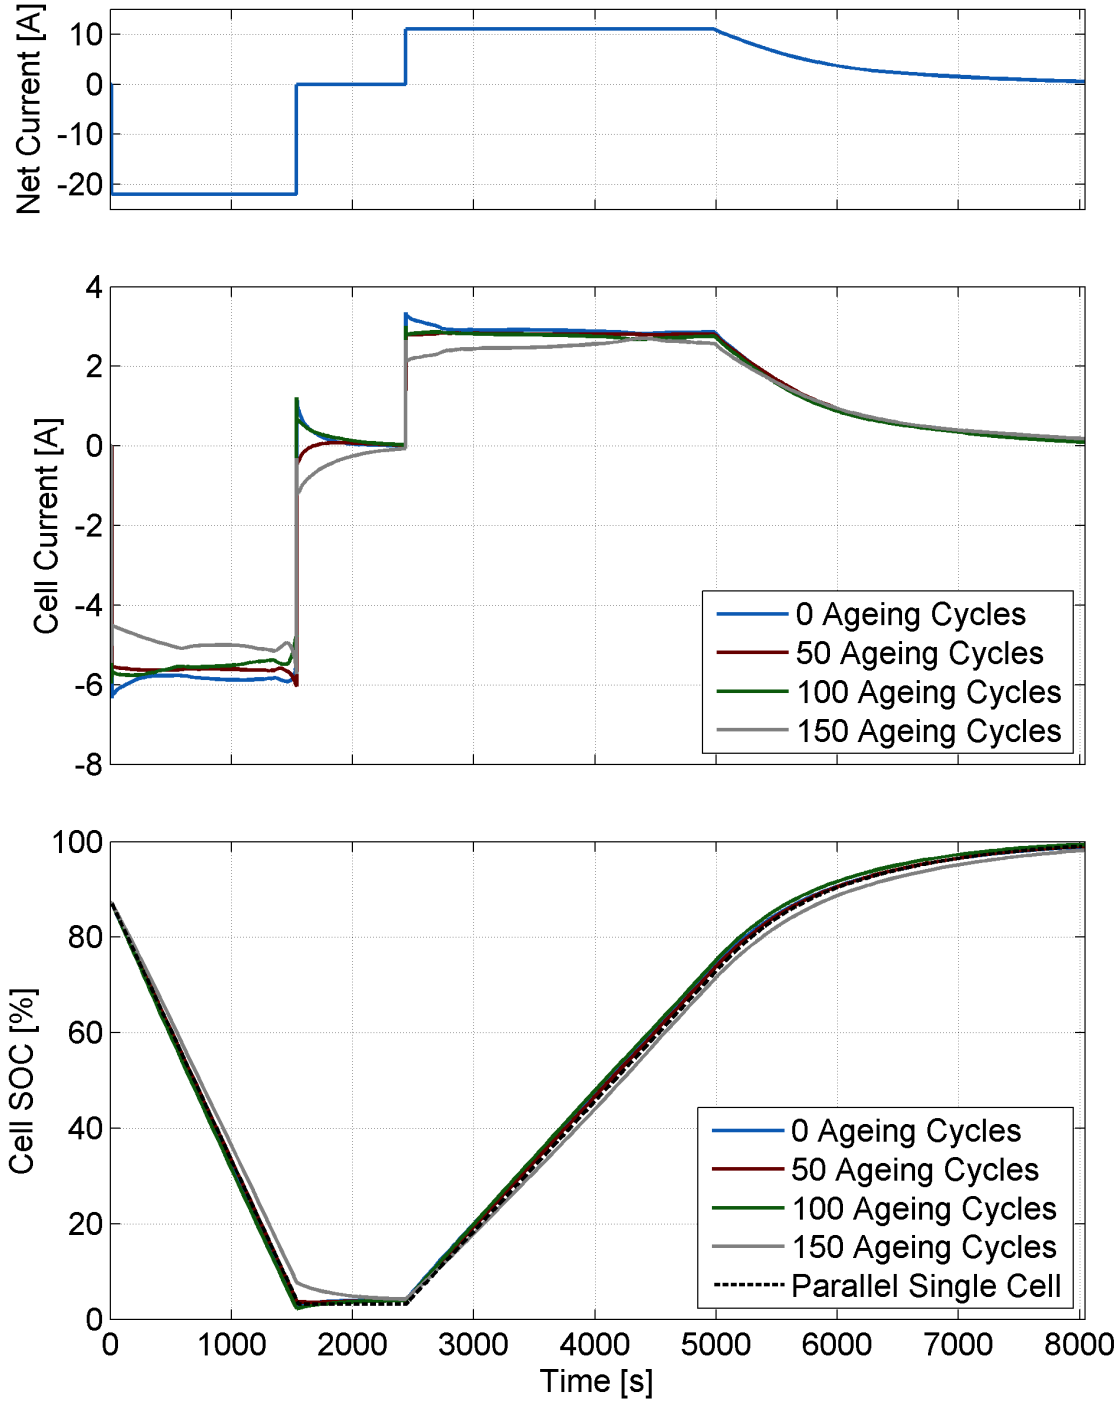

Figure 27: Simulation results for the ageing drive cycle

Comparing the two, the BEV application will result in greater SOC difference between the cells, with approximately 5% difference between the highest and lowest SOCs. This will potentially reduce the energy utilisation of the pack as EOD will be reached with some cells within the parallel stack having energy remaining. The rest after the discharge in Figure 27 show the cells self-balance in an exponential nature when the applied current is zero (or relatively low). The time taken to self-balance will depend on the initial SOC difference, internal

resistance and slope of the OCV-SOC curve (a flat curve means more charge throughput is required to bring the cells to the same voltage). The mean cell current loading can be calculated using (36), which normalises the cell charge throughput against the net applied charge. Table 8 shows the results of applying (36) to both drive cycles. The results show that the BEV application results in a more even loading between cells because of the lower magnitude. The relative loadings of the most and least aged cells are further apart for the HEV cycle. This is likely to increase the ageing of the least aged cell, and reduce the ageing of the most-aged cells, bringing the SOHs closer together [37].

$$i_{norm} = 100\frac{\int|i_{cell}|}{\int|i_{app}|} \tag{36}$$

Table 8: Summary of parallel cell current distribution simulation results

Cell Number	Ageing Cycles	Normalised Current (BEV)	Normalised Current (HEV)
1	0	26.61	27.77
2	50	25.64	25.09
3	100	25.25	25.45
4	150	23.47	22.5

4.1.4.3 Imbalance Through Inter-cell Connections

The experimental set-up of four cells in parallel consisted of each of the cells connected to two common terminals. However, in some battery packs, such as the Tesla Model S, parallel cells are connected in a "ladder" form, shown in *Figure 20*, where only the first cell is connected to the terminals. The interconnection resistance R_{conn} gradually increase the effective resistance of the cells further away from the terminals. The resistances have a cumulative effect and so will be most noticeable in a highly parallelized battery pack. As such, this simulation uses 70 cells in parallel, with a 10$\mu\Omega$ connection resistance between each cell. This is less than 1000th of the impedance of the cell and could represent the resistance of the bus bars connecting the cells or imperfect welds. The validation current profile was used, but the current magnitude was scaled up by a factor of 70:4 to account for the increase in capacity. The 30% SOC dSPACE parameters from Table 7 were used for all 70 cells, meaning without the interconnection resistance, all of the cell currents would be identical. The results showed that the

increased effective cell resistance of the further away cells means that they draw a lower magnitude current than the cells closer to the terminals. This is quantified in Figure 28, which shows the charge throughput of each cell over the drive cycle, scaled relative to what the cell current would be if there were no interconnection resistance. The plot shows that the cell closest to the terminals (with the lowest effective resistance) undergoes over 50% more charge throughput compared to the cell furthest from the terminals. This will increase the amount of energy imbalance, the rate of ageing and heat generation.

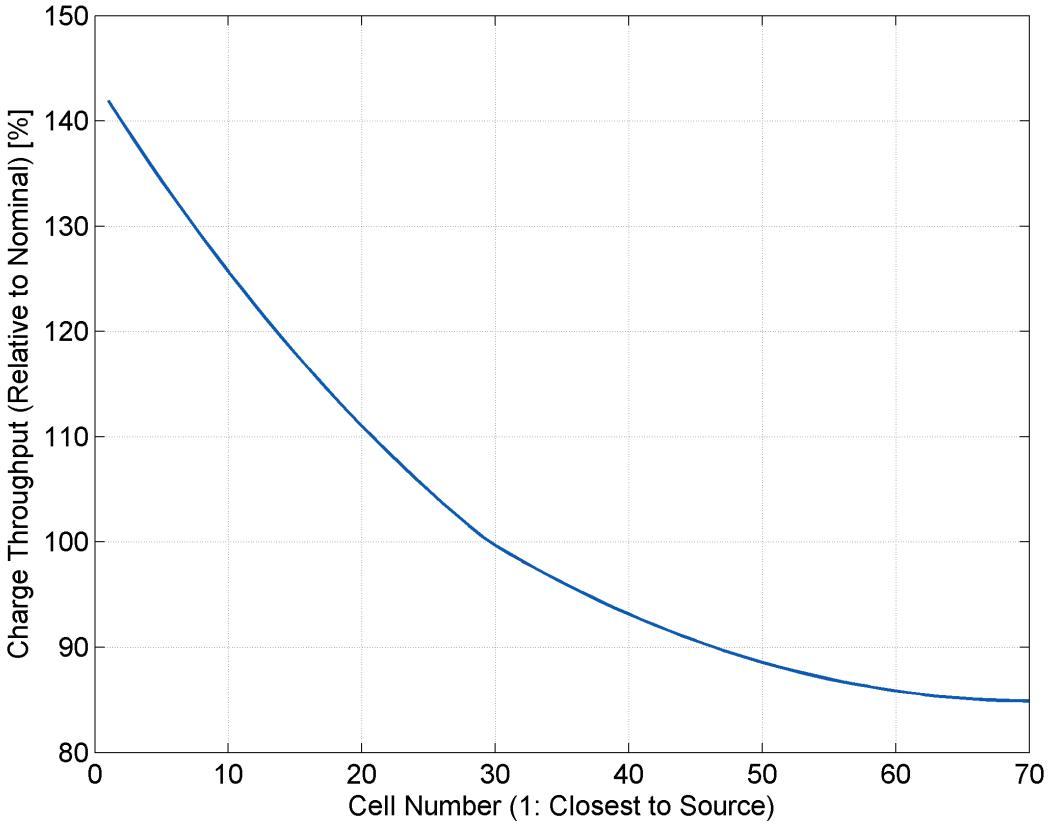

Figure 28: Relative charge throughput for a 70p1s pack with interconnection resistance

4.1.5 Conclusions

The experimental and simulation work shows that cells with different impedances and capacities connected in parallel do not behave in a uniform manner and can see significantly different currents. As a BMS does not typically monitor current within parallel units, some cells may be taken above their intended operating current, or be aged more quickly due to increased charge throughput and ohmic heat generation, shortening the lifespan of the pack. Similarly, a cell losing electrical contact with the battery pack could result in poor SOC estimation, leading to the driver being misled about the available range of the vehicle. A typical HEV drive cycle does not appear to cause a significant amount of SOC

variation between cells in parallel due to its dynamic nature. However, steady currents and drive cycles which cover a large DOD, where cells can enter the region where impedance increases and the OCV-SOC curve gradient increases, can cause significant variation and give rise to cases where higher impedance cells can take larger currents. These effects, combined with the cumulative impact of cell connection resistance, may lead to premature ageing of some cells, weakening the pack as a whole. Variation in current can also cause uneven heat generation within a pack which may require a more advanced thermal management system.

SOC imbalance within a parallel unit of cells is induced during a drive cycle. The cells do self-balance, but this occurs relatively slowly and so the rate at which imbalance is induced will often be greater than the rate at which the cells naturally self-balance (this effect will be more pronounced during more aggressive drive cycles). From this analysis balancing cells in series is the more pertinent issue. Additionally, hardware options for balancing are currently designed for cells in series [17, 48, 139], and designing for parallel-connected cells would shift the focus of this work from control algorithm design to hardware electronics and manufacture. As such, the decision was made to focus on series cells for the remainder of this research.

4.2 Cell Connected in Series

Connecting ECMs in series can scale up simulations from the cell level to the pack. This can be used to evaluate imbalance and form the basis of a software-in-loop (SIL) plant model, which can be combined with a model of a balancing system and used to evaluate performance and design of a control system.

Truchot et al. [140] connected three cells in series and compared individual cell models with a combined pack model, evaluating different ways of extrapolating from the single cell SOCs to the pack SOC, for example taking the average of each cell's SOC. The authors found that no method of extrapolation accurately represented the pack SOC across the whole SOC range. However, the authors did not consider the impact of imbalance on usability. Even if the pack SOC is at

for example 5%, if one of the cells within that pack has reached 0% then the pack has to cease operation, and so the 5% is a misleading figure for the end user. In this case, the pack SOC could be seen as the minimum cell SOC during discharge, and maximum cell SOC during charging. In the case of active balancing, the actual usable SOC could approach the pack SOC if all the energy in each cell is utilised (they would not strictly be equal owing to the energy lost during balancing). Other approaches are often similar, with a nominal cell model created in [141] and statistical methods applied to model the strongest and weakest cells. This could reduce computational demand compared to having an estimator per cell, but the specifics of which cell is at which SOC is lost, which is required for improved balancing control. Similarly, a battery model is proposed in [142] lacks the individual cell level information required for balancing.

Combining cells in series is straightforward, as (without balancing), all cells undergo the same current and are independent of one another. Some initial modelling utilising cells in series was published by the author [8], combining the cell models with high-level models of different balancing systems to perform closed-loop simulations. This type of modelling and simulation work is built upon here. From Chapter 2, the key properties which affect cell behaviour in the short-term are capacity and internal resistance. To test the balancing system in Chapter 6, seven 15Ah pouch cells were obtained from a second-hand electric vehicle which underwent four years of real-world driving. The cells were characterised as part of the standard Catapult procedure, using procedures A-C from Appendix B, and the resultant capacity and resistance measurements are shown in Table 9. As Table 9 shows, there is significant variation amongst the cells. The capacity shows a 0.67Ah standard deviation, which is 4.6% relative to the mean capacity. Similarly, the internal resistance has a 0.14mΩ standard deviation, which is 4.9% relative to the mean. This has the potential to create a significant amount of imbalance.

Table 9: Summary of ageing characteristics of automotive pouch cells

Cell Number	1C Discharge Capacity (Ah)	Average resistance (mΩ)		
		90% SOC	***50% SOC***	***20% SOC***
1	13.44	1.89	3.02	3.73
2	13.87	1.80	2.81	3.20
3	15.16	1.76	2.71	2.87
4	15.13	1.78	2.75	2.91
5	14.90	1.76	2.64	2.72
6	14.86	1.76	2.64	2.71
7	14.78	1.77	2.66	2.72

Example pulse power fit is shown in Figure 29, and a comparison of the fit results is given in Table 10. While a 1RC model gives acceptable results, a significant improvement can be made with 2RC model, and a further smaller improvement with a 3RC model.

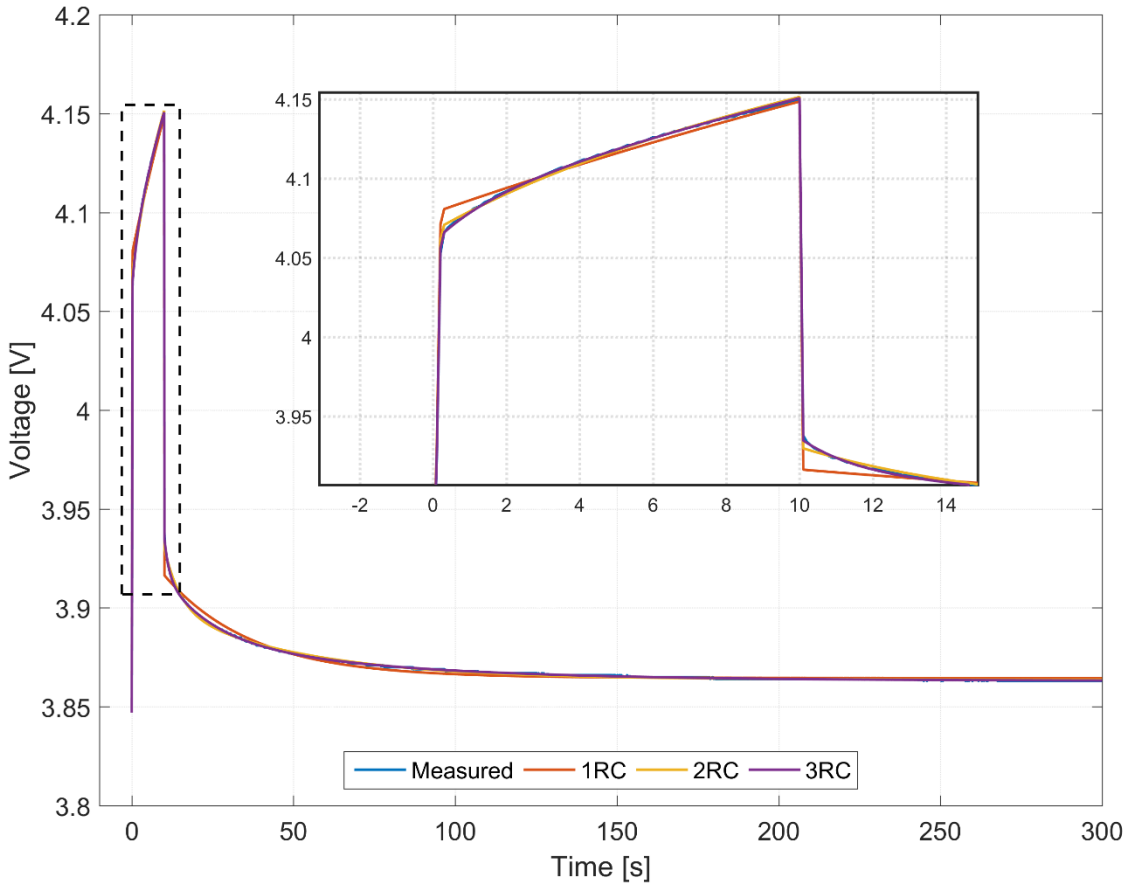

Figure 29: Example pulse power fit for an automotive pouch cell

The OCV-SOC data, obtained using GITT, used are shown in Figure 30. Note that there is no charging data above 92%, and no discharging data below 8%, so

outside of these respective limits the yellow average line is equal to the respective remaining data source. For this test work there was not sufficient availability in the laboratory to run the standard Catapult validation profile. Instead, test data collected during the experimental work were used to validate the ECM. For example, Figure 31 shows a current for one cell profile used when setting the cells to balance in different combinations, which is explained further in section 6.6. Note that in this case, only 4 cells were connected due to the issues mentioned in section 6.1. The parameterisation results are shown in Table 11. It is noteworthy that the voltage errors are lower compared to the Catapult validation profile because the currents are lower.

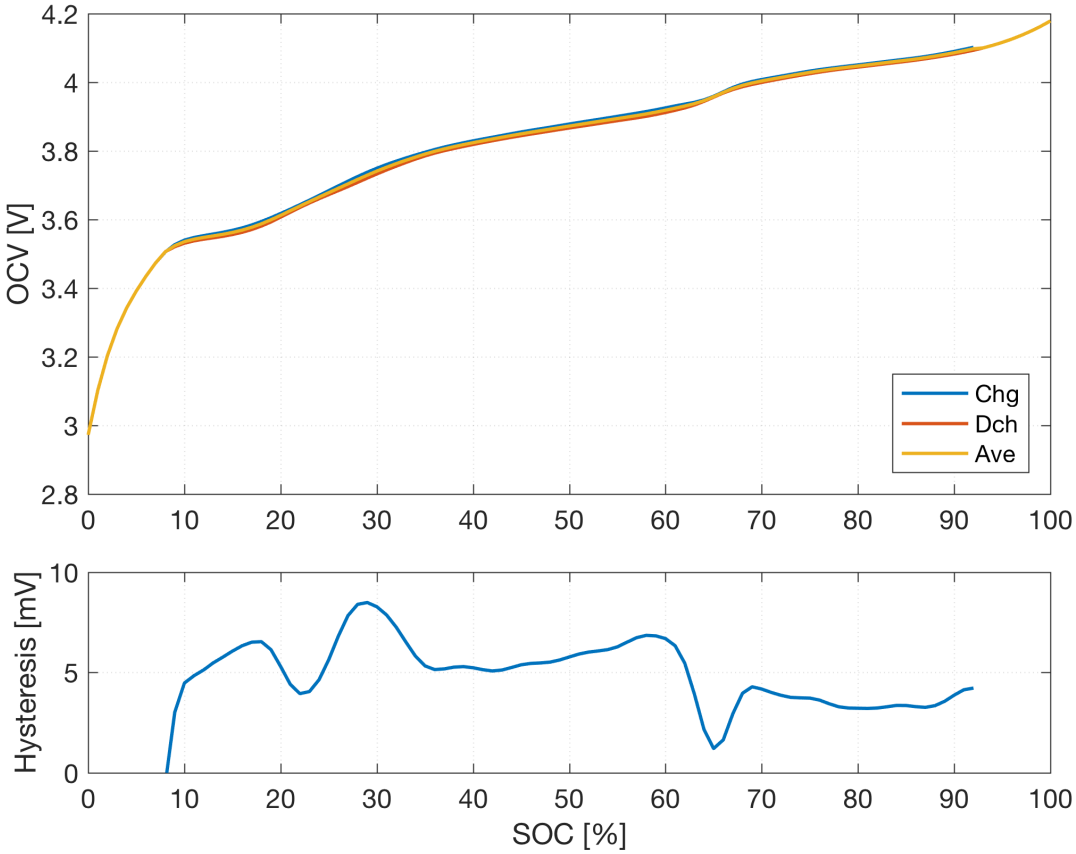

Figure 30: OCV-SOC data for an automotive pouch cell

Table 10: Pulse power ECM fit results for an automotive pouch cell

Order	RMS error (mV)	Peak error (mV)
1	1.77	21.52
2	0.63	8.16
3	0.37	3.11

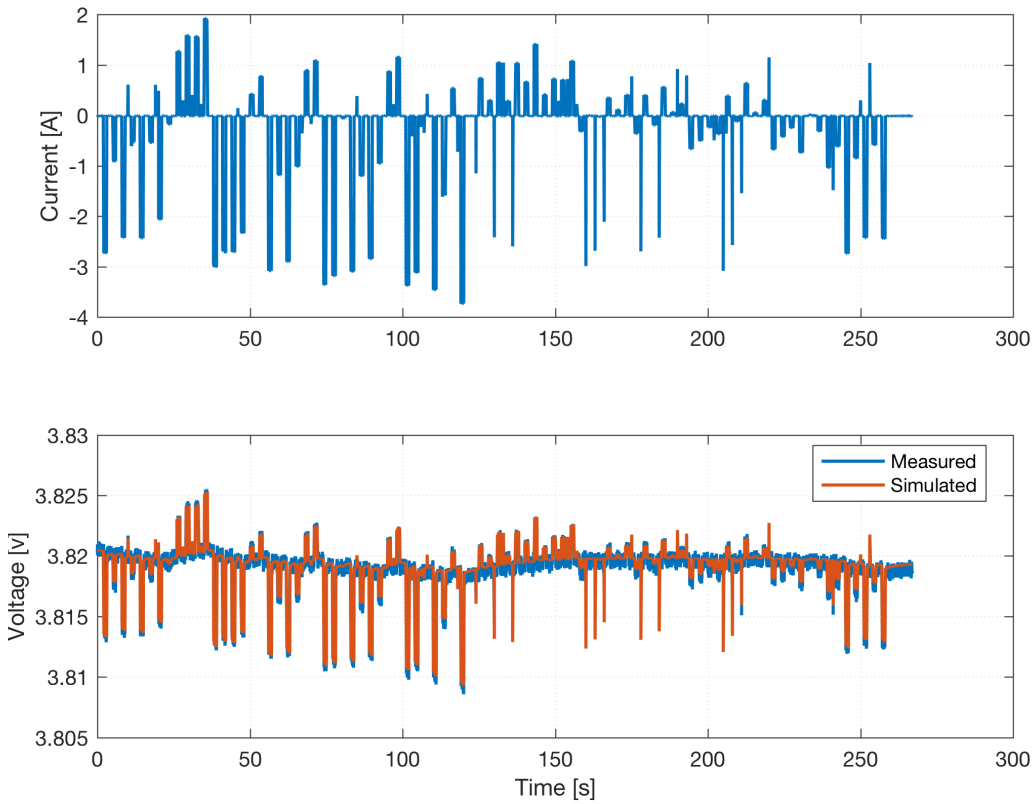

Figure 31: Example model validation for an automotive pouch cell

Table 11: Pouch Cell model validation results

Cell	RMS error (mV)	Peak error (mV)
1	0.80	2.09
2	0.97	4.25
3	3.91	5.31
4	0.59	2.62

Pulse power data were used to parameterise the cells at 20,50 and 90% SOC as described in section 3.4. Figure 32 shows the cell voltages and SOCs during a 2C discharge pulse. The initial conditions were set such that the cells would be exactly balanced at the end of the test. The cell voltages at the end of the pulse were 12mV different owing to impedance differences. The difference is purely because of the impedance differences between the cells, and is an example of how voltage can be a misleading metric of imbalance.

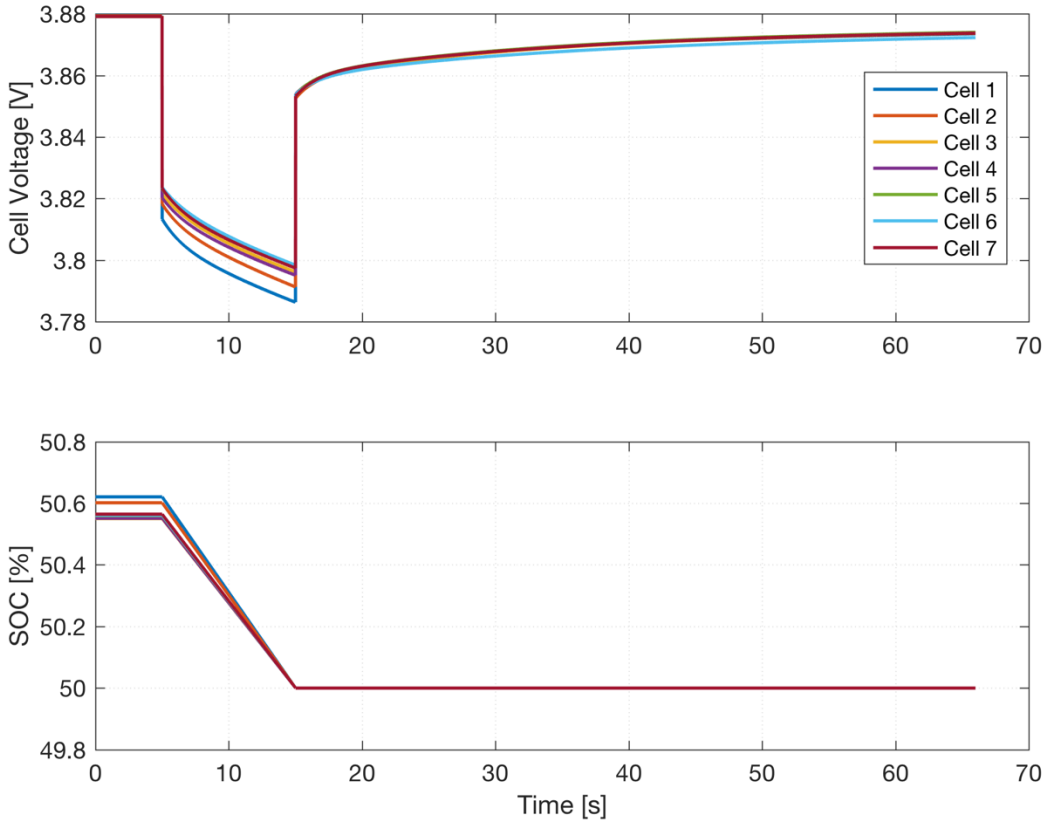

Figure 32: Cell Voltages and SOCs during a 10 second pulse simulation

Vehicle drive cycles can be used to simulate the battery pack during a variety of typical driving scenarios. Often the drive cycle is for vehicle speed, and can be translated to a load current using a 1D vehicle model, which requires some basic vehicle information such as drag coefficient, vehicle frontal area and vehicle mass [143]. Vehicle speed for the Artemis Combined drive cycle along with the battery pack current calculated using the Catapult 1D model [143] is shown in Figure 33. Figure 34 shows the SOCs of the cells when subject to the drive cycle in Figure 33 with no balancing. As can be seen, the cells start balanced at 90% SOC, but by the end of discharge the cell SOCs have diverged, with an 8.4% difference between the highest and lowest SOCs.

Battery pack simulations such as these help to assess performance and quantify imbalance for various scenarios. They are also used later in this research for closed loop simulations, where state estimators and controllers can be designed and developed offline.

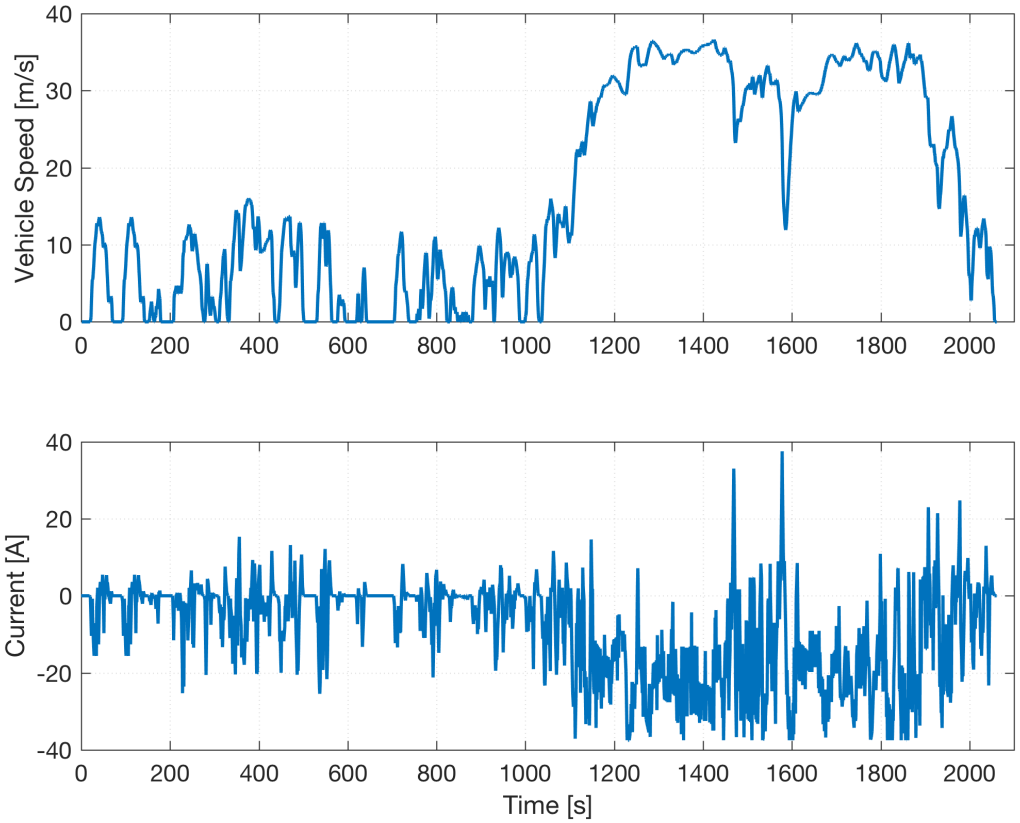

Figure 33: Artemis Combined drive cycle vehicle speed and derived battery current

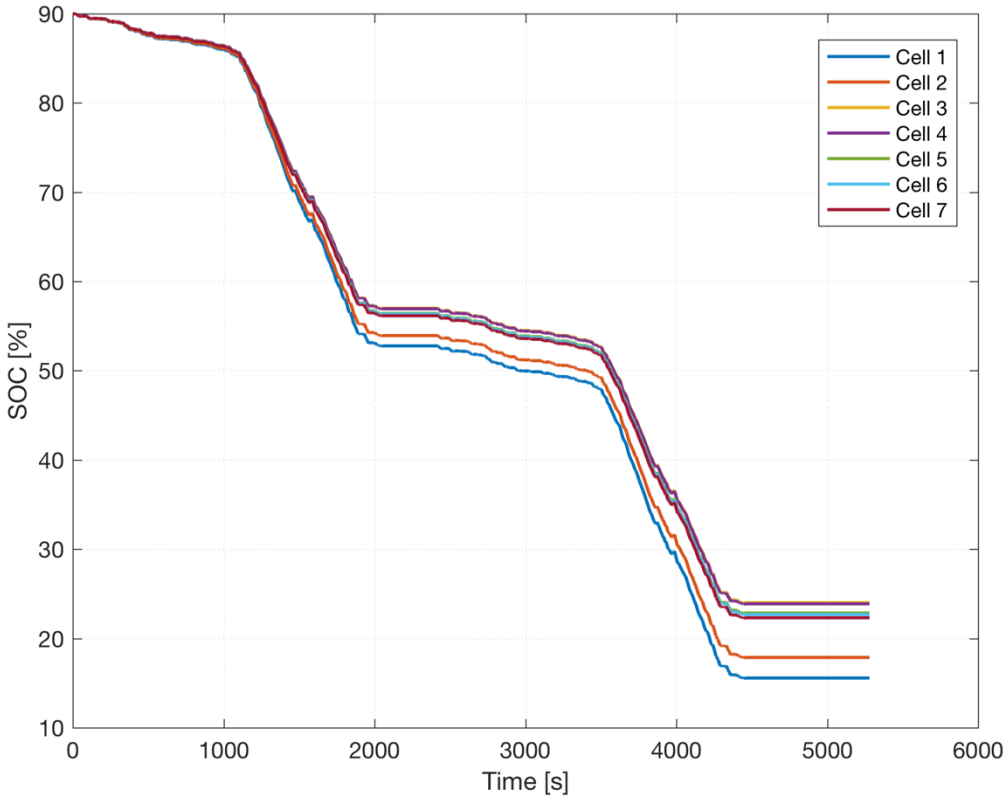

Figure 34: Simulation of Artemis Combined drive cycle

4.3 SOC Estimation

In section 2.1.5, the fundamentals for cell state estimation using a Kalman Filter were introduced. The basic premise is as follows:

1. Use coulomb counting to predict the SOC at the next time step.
2. Use an ECM to predict the cell's terminal voltage.
3. Compare this to the measured cell voltage.
4. Use the difference between the two to update the states with the Kalman gain matrix. For example, if the estimated voltage is too high, it is likely that the OCV estimate, and so SOC estimate, is also too high.

As the cell model is nonlinear, a nonlinear variant of the KF is required[7]. The Sigma Point Kalman Filter (SPKF) was chosen because of its balance between performance and complexity. It offers superior performance for highly nonlinear systems compared to an extended Kalman Filter (EKF) without much higher computational complexity [75]. There is no specific definition of 'highly nonlinear', but parameter estimation using a KF can cause large nonlinearities. For example, (37) shows the gradient for internal resistance (the partial derivative of (12)) is equal to the cell current, which can often be highly discontinuous and dynamic rather than a smooth profile.

$$\frac{\partial v_t}{\partial R_0} = i \tag{37}$$

The EKF also requires Jacobians to be calculated for the state and output equations. The Jacobians contain the partial differentials of the state and output equations with respect to each state [82]. These Jacobians are not required with the SPKF, making it easier to implement and change between different models. For these reasons, the SPKF provides a suitable framework for this work and future developments. The specifics of the SPKF algorithm and its implementation are contained within Submission 5. The data and experimental results from this

[7] Another alternative is to linearise the model and use gain scheduling to switch between models when appropriate (for example at fixed SOC and temperature points). Linearisation was not pursued as it lacks the flexibility of nonlinear methods. Different models with differing degrees of nonlinearity were used in developing the estimator, and each one would require evaluation to determine whether linearization was feasible and how performance would be affected.

section use the automotive pouch cells from section 4.2 to help illustrate the application of KFs within this context. These were ultimately used as part of the closed-loop real-time system detailed in section 6.5.

During development, it was found that tuning the estimator using SIL had limited applicability to when it was run online. This is because ECMs in the plant model are missing dynamics compared to the actual cell, particularly with regard to relaxation. The plant and estimator models are similar (both missing the same dynamics), meaning the estimators appear to perform much better in SIL than during testing. This is discussed further in Submission 5. While SIL simulations were used to get a functional solution, a considerable amount of online tuning was required. This typically involved utilising all test data: for example when loading the cells without balancing (for example in preparation for a balancing test), different configurations could be evaluated since poor estimator performance is not a concern for that particular test. As more test data were collected, offline tuning could be performed using previous test results. This has the downside of not having a reference SOC as a SIL model has, but it means that unmodelled cell affects could be better accounted for.

4.3.1 Estimator Design

A full discussion of the estimator structure is contained within Submission 5. It was found that estimating internal resistance online allowed for the estimator to be more robust to uncertainty in temperature, ageing and a wide SOC operating range. However, the other ECM parameters and cell capacity were harder to estimate using a KF and their values has less of an impact on SOC estimation performance, so the offline parameterisation values were used. For much of the experimental work in Chapter 6 , a 1RC ECM was used to simplify the estimator. It is generally recommended [137, 144] that a 1-2RC pairs provide a suitable balance between simplicity and performance. The internal resistance was also estimated. The estimator states are given by (38).

$$\boldsymbol{x} = [s \quad v_{RC} \quad R_o]^{\mathrm{T}} \tag{38}$$

This generally performed well, but there were two areas identified for improvement. Firstly, it was difficult to tune: a good resistance estimate could be

achieved, but sometimes this was at the expensive of SOC state accuracy or noise. Secondly, only having one time constant resulted in poor performance under wider DODs– particularly if the current was constant. A wide DOD means there is a significant low-frequency component to the drive cycle which requires a wider model bandwidth. An ECM with 3 RC pairs contains lower frequency time constants which can help improve on the first issue. However, the additional states made tuning the system even harder. This is especially the case for SIL, as now the cell and estimator models are largely identical. To simplify the estimator, resistance was removed as a state, so the SPKF states were defined by (39).

$$\boldsymbol{x} = [S \quad v_{RC_1} \quad v_{RC_2} \quad v_{RC_3}]^{\mathrm{T}} \tag{39}$$

To estimate resistance, the NLLS solver from section 3.4 was run every 1 second on the previous 20s of test data. Tests were run using experimental data to try and solve for a 3RC model, which could mean all of the SPKF parameters could be updated based on changes in operating conditions. However, there was a lack of consistency in the parameter estimates from one second to the next, apart from internal resistance. To capture slower dynamics, a longer time history would be necessary, increasing the execution time and memory requirements, making it difficult to run online in real-time[8]. A 1RC model was fitted instead, and only the resistance estimate from this NLLS fit used in the SPKF model.

4.3.2 Performance

The relative accuracy of estimators is difficult to ascertain as much of the error comes from dynamics not modelled by the reference model, and so only show up during experimental testing. The current sensors and cell capacity estimates used for this testing were not sufficiently accurate to be able to use coulomb counting as a reliable metric. As discussed further in Submission 5, it was found in practice that the 3RC estimator performed better than the 1RC version, but still suffered from many of the same issues, notably sensitivity under relaxation and a lack of excitation during constant current operation. SOC estimation accuracy

[8] A longer time frame could be fitted by downsampling the data. As the focus is on low frequency dynamics, downsampling should not affect the parameterisation too much, although it may impact the internal resistance estimate. More research is required to explore this further.

could be inferred from observing how much the estimate changes once a load profile has finished and the cell relaxes, as the voltage correction becomes much stronger (after complete relaxation, the filter will just act like an OCV-SOC lookup table). The estimators appeared to consistently operate with less than 2% error.

4.4 Summary

This section has demonstrated the various uses of ECMs throughout this work. Section 4.1 showed how cells in parallel can be modelled, and that contrary to popular belief, can become significantly imbalanced. Other effects such as differences in current and temperature were also discussed. These have implications for battery pack design as well as imbalance and there is a considerable amount of possible research around this area. Section 4.2 used ECMs in series to analyse imbalance and create a plant model, which is used further in chapters 5 and 6 for developing a balancing control system and perform closed-loop simulations. Cell state estimators were developed in section 4.3 to estimate SOC so that imbalance can be evaluated. These are used in Chapter 6 as part of the real-time experimental set-up. With cell and battery modelling addressed, the next Chapter focuses on modelling and control of the balancing system itself.

5 Balancing System Control

In Chapter 2, a number of conclusions were drawn regarding the requirements of a balancing control system:

- The high-level control goal is to ensure that all of the cells are balanced by EOC and EOD. This allows for all of the energy within the battery pack to be utilised.
- The low-level control goal is to bring all cells to the same SOC. This will generally be the mean SOC, but could also be the maximum or minimum;
- A generic control framework is desirable, as it avoids limiting the control system to one type of hardware.
- This generality means it must be easy to adapt to specific hardware, e.g. factoring in current limits and inter-cell connections.

This Chapter addresses the fourth Research Objective of the EngD: to develop a generic balancing control system with the specific goal of maximising the energy utilisation of the battery pack. A control-oriented model of the cells with respect to the balancing system is derived and used to implement controllers to remove energy imbalance. At this point, a generic, idealised balancing system is considered in order to focus on the more fundamental aspects of the control problem. An implementation with specific hardware is detailed in Chapter 6**.** The generic system considered for development is shown in Figure 35.

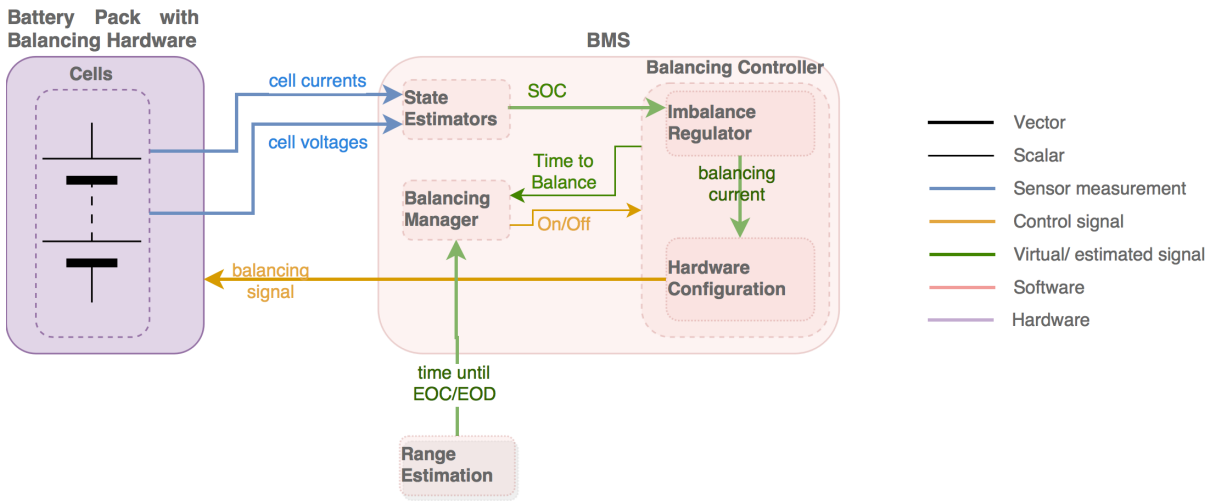

Figure 35: Overview of the battery system in the context of balancing

Details about each signal are presented in Table 12, and a description of each subsystem in Table 13. This encapsulates the balancing system itself, along with the other aspects of the battery pack and BMS that are required for balancing.

Table 12: Battery subsystem descriptions

Signal	Range	Accuracy	Description
Cell Voltage	0 to 5V	1mv	The voltage values used by the BMS are assumed to be post filtering and processing. Higher accuracy voltage measurements can improve SOC estimation, especially for cells with flat regions in the OCV-SOC curve.
Load Current	±2C	0.1% full scale	The current values used by the BMS are assumed to be post filtering and processing. The specific range and resolution will be dependent on the application and performance specifications. HEVs typically operate at much higher C-rates than a BEV, but the cell capacity is smaller, so the measured current magnitude might be lower.
SOC	0 to 100%	0.5%	The SOC estimation accuracy depends on cell properties and algorithm performance, which in turn is determined by the requirements of the system.
Balancing Current	±0.2C	0.01A	The specific range of current is dependent on the application and hardware capability.
Time to EOD/EOC	0-7200s	60s	The remaining mileage/ time remaining is estimated by the energy management system, and sent to the BMS.
Time-to-balance	0-7200s	60s	The time-to-balance estimate is calculated by the balancing controller, and can be compared with the time to EOC/EOD to assess whether balancing should be activated.
Balancing Signal	0-1	0.01	Depending on the balancing system interface, cell current may be required to be translated into a normalised value (or some other abstract signal). It will also likely need to be quantised

			and transmitted via a digital communications bus such as CAN or SPI.
Balancing Management Signal	Boolean	-	An overall signal to enable or disable balancing. If disabled, the balancing algorithm should still run in the background to give an up to date time-to-balance.

Table 13: Subsystems relating to cell balancing

Subsystem	Description
Battery Pack	This contains the cells, balancing hardware, battery management system microcontrollers, safety circuitry and thermal management.
Cells	Each series unit of cells requires a voltage measurement. For systems with passive balancing, typically only the applied current is measured but for ABSs, cell current measurements could improve SOC estimation accuracy, though may be inferred from the load and the balancing current calculation. Temperatures sensors are also positioned throughout the pack (typically less than one sensor per cell). As well as for safety and thermal management, these can also feed into the estimation algorithms for improved accuracy.
State Estimators	Estimates of SOC and SOH (cell impedance and capacity) are required by the BMS. A joint estimator may be used to estimate all signals simultaneously, or SOH may be decoupled from SOC owing to the different timescales.
Balancing Manager	This makes the decision whether balancing should take place, given the status of the cells and the load profile and history. See section 5.1 for details.
Imbalance regulator	This is the core algorithm to determine what balancing current each cell requires based on the cell SOCs. The details of the proposed solutions are in section 5.4.
Hardware Configuration	This translates the balancing current request into a signal required by the balancing hardware. See “Balancing Signal” in Table 12.
Range Estimator	An estimate of the remaining driving time/distance of the vehicle is performed by the energy management system, and is outside of the scope of the balancing system.

The following assumptions are made regarding this generic architecture of the balancing system:

- Every cell can be set to balance, and can be operated simultaneously. If this is not directly feasible (for example, the Texas Instruments EV1401EVM balancing development board [145] can only set one cell to balance per module), the hardware configuration layer could implement multiplexing, which cycles through each cell, setting it to balance for a short period of time.
- There are limits on balancing current magnitude. These are considered here as inequality constraints, but other forms of constraints could also be applied. Each cell can have different current limits, and different values for charging and discharging (by setting one of these to zero, uni-directional balancing can be implemented). These constraints could be a function of operating conditions such as cell voltage, but are considered time-invariant here.
- While SOC is fundamental to the balancing system, the cell state estimators are separate to the balancing controller. The cell SOC estimates are treated as measurements (with some uncertainty) by the balancing controller. The cell capacity estimates from the SOH estimator are also required by the controller. However, these will change slowly (for example, 20% over the course of 5-8 years [146]) and can be assumed as constant during any balancing operation, and periodically updated between balancing.

The high-level strategy for deciding when to balance is proposed in section 5.1. The low-level controller structure is outlined in section 5.2, and the model used to design the controllers derived in section 5.3. The specific controllers for regulating imbalance are derived in section 5.4, based on this model. Simulations are presented in section 5.5, with conclusions drawn in section 5.6. The control system development is primarily covered in Submissions 3 and 5.

5.1 Balancing Management

Before addressing the specifics of how balancing currents should be determined, there is the wider question of when to balance. While active balancing aims to redistribute charge rather than dissipate it as heat, the system will not be completely efficient and some energy will still be wasted. Additionally, there will

be some power consumed by the electronics when active. The energy loss during balancing means that unless balancing explicitly contributes to additional energy utilisation, it could actually reduce the energy output of the battery pack.

In Chapter 1, it was noted that when using SOC to quantify imbalance, there is a potential to balance when not necessary. Figure 7 showed that depending on the cell capacities and initial SOCs, it is possible that the two currently imbalanced cells would naturally reach a balanced state without any active balancing required. In this case, the fact that the cells are imbalanced in the central SOC region did not matter because they would be balanced by EOD. There are other situations when this also applies. Take for example, a BEV driver. Often, the car will be charged overnight, which gives a long time for the balancing system to ensure that all cells are fully charged. The next day, as the battery pack is discharged through driving, the cells will start to become imbalanced owing to capacity differences. However, it is quite likely that the driver will return the vehicle to charge before the SOC gets too low, e.g. when the cells are around 30% SOC. When the cells are charged, they will naturally return to an approximately balanced state[9], as all of the cells have undergone the same current and the amount of charge removed during driving and replenished during charging is the same. If balancing had been applied during driving to ensure the cells were continuously balanced, balancing would again be required during charging. And since the driver stopped driving at a moderate SOC, this balancing has not contributed to any extra energy utilisation: the battery pack was capable of supplying the required energy, despite the imbalance.

There are clearly many times when balancing is of no benefit. However, if the pack is brought down to a low SOC, it could be limited by a weaker cell (or that cell being at a very low SOC is undesirable for the vehicle strategy, worsening its efficiency and ageing). In this case, active balancing can offer quantifiable gains. The proposed active balancing management strategy for a BEV (or other wide DOD application) is shown in Figure 36 and detailed below.

[9] There will be differences which have emerged through self-discharge and electronics power consumption, but as discussed in 2 this will be small compared to the amount of imbalance which emerged during driving.

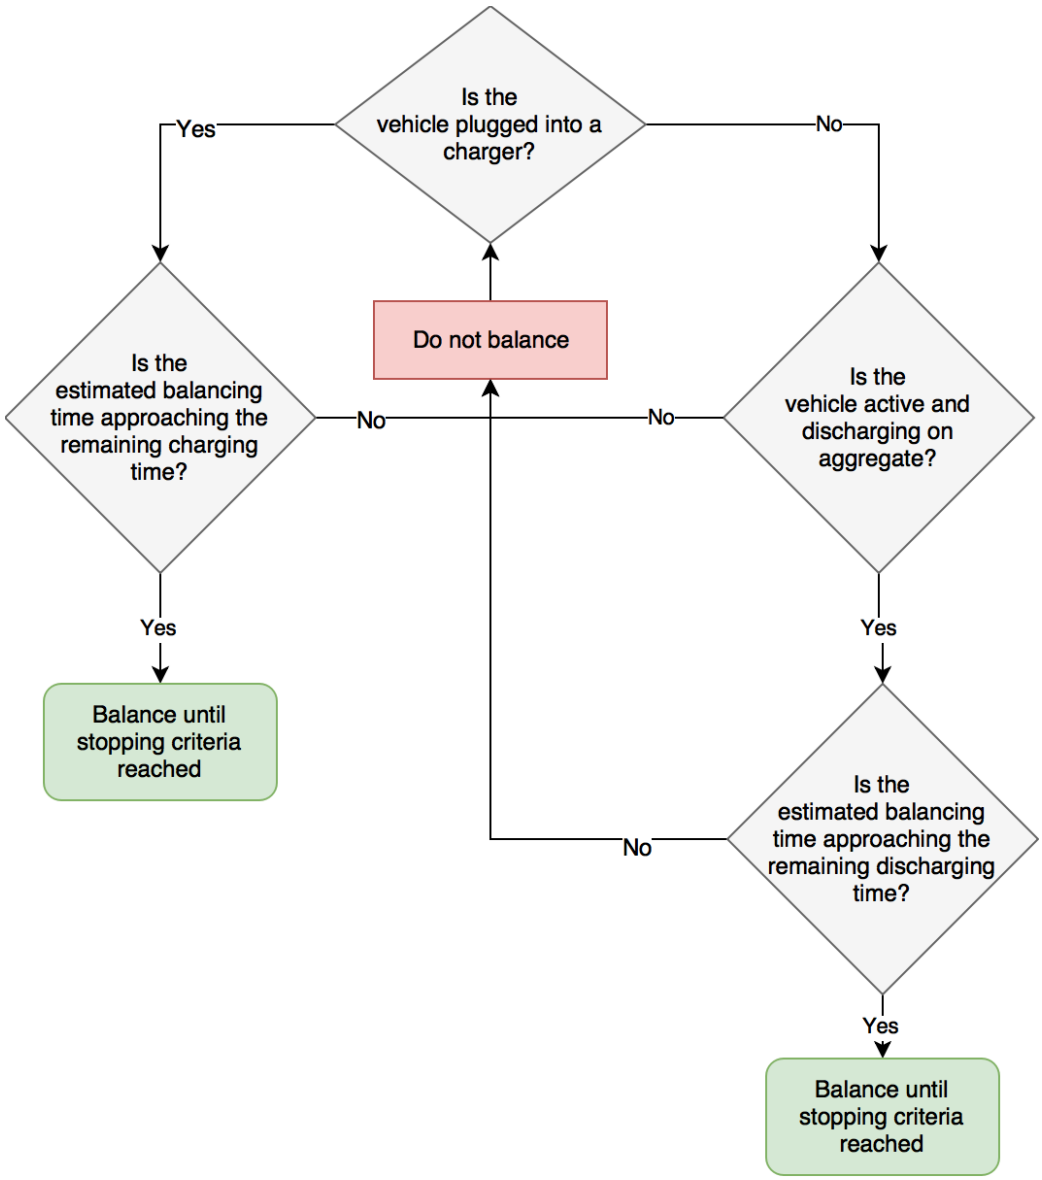

Figure 36: Balancing Decision Flowchart

1. Ensure the cells are fully balanced by EOC (some balancing after charging has finished would be acceptable in some circumstances such as the car being charged and parked overnight, but other times the driver will want to continue driving as soon as charging is complete).
2. Do not balance while the vehicle is being discharged. Having the active balancing hardware in a standby mode will ensure it drains a minimal amount of current. However, imbalance should be monitored and used to estimate the amount of time required to remove this imbalance.
3. If it appears the driver will require the full energy of the battery pack, then activate the balancing system.
    a. The point at which this occurs will depend on how fast imbalance can be removed, and so is hardware dependent.

b. At a simple level, this could be when the cell with the lowest SOC goes below a threshold. A more effective strategy would be to use the vehicle's energy management system. This should predict a remaining driving time (assuming that the cells are balanced). When the time-to-balance estimate starts to approach the time-to-EOD, balancing should be activated.
c. Further improvements could be made by employing wider vehicle data. For example, if the driver has programmed a charging point (such as their home address) into the navigation system, that journey may be achievable without balancing, even though the time-to-balance might be greater than the time-to-EOD.

4. When the vehicle is charging, stop balancing, even if the cells are still imbalanced. When the cells are approaching EOC, return to step 1.

An arbitrary SOC or voltage difference is often used in the literature. Values such as 2% [112] to 3% [48] are often used. However, this can still limit pack performance. In a BEV, where the cells may be charged to close to 100%, and discharged to near 0%, reducing imbalance by a further 2% could make a big difference and avoid the cells having to be discharged quite so far. The danger with trying to remove imbalance completely is that the cells will be over-balanced: some cells will be close to the mean before others. These cells may then be frequently set to charge/ discharge as they hover around the mean. A hysteresis value can be applied around the threshold, although this can mean that global imbalance becomes relatively large, as discussed in 5.5. This wastes energy without achieving any imbalance reduction. The control system should be able to minimise this effect while still reducing imbalance as much as possible. In section 2.1.4, the absolute SOC difference (the maximum minus the minimum) was chosen to use as the high-level imbalance metric. A value of 0.05% is used in the simulation studies in this work to determine when balancing should finish. This is chosen as a small value, but greater than zero, which could cause balancing to continue indefinitely, wasting energy. No limit was applied for the experimental studies, in order to analyse close-to-balanced behaviour.

5.2 Imbalance regulation

With the high-level strategy defined, the low-level controller can be designed. The low-level strategy deals with the specifics of moving energy between cells. A logical control goal is to bring all cells to the same SOC simultaneously. Having some cells finish balancing before others means that a higher than necessary balancing current has been used, resulting in more power loss [112], while the battery pack is still being limited by the other cells which are out of balance. Creating a trajectory for each cell in order for them to reach the mean simultaneously requires knowledge of the cells in order to predict forward in time, rather than just react to the current state. To do this, a model-based approach is used for the control system design. This uses information about the cells to understand how SOC will be affected by balancing currents, so that the ideal balancing currents can be determined.

5.3 Derivation of control model

As discussed within Chapter 2, SOC is used as the metric of imbalance. The SOC state equation (16) first introduced in section 3.3 forms the basis of the control system. However, with balancing, the cell current is comprised of two components. There is the balancing current, which is controllable, and the battery pack applied load current, a disturbance which the balancing system cannot control but may account for. Controlling the balancing system also adds an interaction between the cells, as all SOCs are required to determine the value of the balancing currents that should be applied.

5.3.1 Cell SOC System

The combined system for N cells is given by (40). The SOC vector $\boldsymbol{s}$ is calculated using the control input $\boldsymbol{u}$ (a vector of each cell's balancing currents) and the scalar disturbance d (the battery pack load current). The B_c matrix and $\boldsymbol{b_d}$ vector are both similar, containing the scaling factor which translates current into a percentage. At this point, the cells are still independent, as there is no control law linking the cells states to the input vector.

$$\dot{\boldsymbol{s}} = B_c\boldsymbol{u} + \boldsymbol{b_d}d \tag{40}$$

$$B_c = \begin{bmatrix} b_1 & & & \\ & b_2 & & \\ & & \ddots & \\ & & & b_N \end{bmatrix}$$

$$\boldsymbol{b_d} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix}$$

$$b_n = \frac{1}{36Q_{cap_n}}$$

Even without this coupling, the balancing system may cause interactions. For example, a cell-module balancing system connected to *N* cells will load the module with approximately 1/*N*th the balancing current of each cell (including itself), in the opposite direction. For example, if a 4-cell system set the first cell to charge-balance at 2A, each cell would undergo -0.5A, meaning the first cell current is 1.5A. This interaction can be incorporated into the control model using (41), which calculates the modified control input matrix B_c'. The *F* matrix for the above example is also shown, though this can change depending on the physical connections between each cell. It also assumes 100% system efficiency. In practice, more than 0.5A would be required from the cells to result in 2A at cell 1, and if the first cell was discharging instead, they would be charged at less than 0.5A. However, an algebraic loop can occur within the control system if *F* is updated based on the balancing direction, as it is itself used to calculate balancing direction. The additional complexity of factoring in the efficiency is not worth the relatively low increase in cell current accuracy that the model provides. The change in current from factoring in efficiency will typically be in the region of milliamperes, a fraction of the total cell current.

$$\dot{\boldsymbol{s}} = B_c'\boldsymbol{u} + \boldsymbol{b_d}d$$

$$B_c' = B_c(I + F)$$

$$F = \begin{bmatrix} -\frac{1}{N} & -\frac{1}{N} & \cdots & -\frac{1}{N} \\ -\frac{1}{N} & -\frac{1}{N} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ -\frac{1}{N} & \cdots & \cdots & -\frac{1}{N} \end{bmatrix} \quad (41)$$

5.3.2 State transformations

As noted in section 5.1, the goal is to bring all cells to the same SOC. However, the SOCs will be continually changing because of both *u* and *d*, meaning the set-point reference will also be moving. The distance from the mean of the SOC vector (ΔSOC) was used as the target, as discussed in section 2.1.4. To achieve this, a linear state transform was applied to create the ΔSOC vector $\boldsymbol{m}$, as shown in (42). *M* subtracts each cell's SOC from the mean of all cells.

$$\boldsymbol{m} = M\boldsymbol{s}$$
$$M = I - \frac{1}{N}\begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ 1 & \cdots & \cdots & 1 \end{bmatrix} \tag{42}$$

The Matrix *M* is idempotent [147], which also means it is singular and so cannot be inverted. As such, the transformation cannot be reversed: $\boldsymbol{s}$ cannot be obtained from $\boldsymbol{m}$. This is because the mean is factored out, so there are an infinite amount of possible *s* vectors for any *m* vector. For example, if $\boldsymbol{m}$ is [-1 , 0 , 1], *s* could be [9 , 10 , 11], or [85 , 86 , 87] and so on. The transformed state model is given by (43), which is obtained by applying (42) to (41). In standard linear state-space notation, the system is written as (44). Since the *C* state-space matrix is the identity matrix, there is full state feedback and no state estimation is required[10].

$$\dot{\boldsymbol{m}}(t) = MB_c'\boldsymbol{u}(t) + Mb_d d(t) \tag{43}$$

$$\dot{\boldsymbol{m}}(t) = A\boldsymbol{m}(t) + B_1\boldsymbol{u}(t) + \boldsymbol{b}_2 d(t)$$
$$\boldsymbol{y}(t) = C\boldsymbol{m}(t)$$

$$A = 0 \tag{44}$$
$$B_1 = MB_c'$$
$$\boldsymbol{b_2} = M\boldsymbol{b_d}$$
$$C = I$$

[10] This vector is actually estimated by the cell state estimators, but this is assumed as a measurement to the control system. In this sense, no further estimates of any controller states is required.

The system poles are found by calculating the eigenvalues of the *A* matrix as per (45). As the *A* matrix is zero, the poles are all zeros, meaning the system is conditionally stable. The control system must therefore introduce negative eigenvalues to stabilise the system. In the complex frequency ($j\omega$) domain, the transfer function G_{tf} for a general state-space system is given by (46). The transfer function G_d which relates the disturbance to the output is calculated using (47). As frequency ω tends to zero, the magnitude of the transfer function becomes infinite. This represents the divergence in SOCs driven by differences in capacity when a constant load current is applied. If the capacities are all identical, the transfer function has a zero magnitude with respect to the disturbance (i.e. the disturbance exactly maintains whatever ΔSOC state the system is currently at). This also emphasises that DC and low frequency disturbances are of most relevance for balancing, since high frequency inputs do not result in a large SOC change.

$$\boldsymbol{\lambda} = eig(A) \tag{45}$$

$$G_{tf}(j\omega) = C(j\omega I - A)^{-1}B \tag{46}$$

$$G_d(j\omega) = \frac{I}{j\omega}\boldsymbol{b_d} \tag{47}$$

The controllability and observability matrices from equations (22) and (23) are straightforward to calculate because of the A matrix being equal to zero. They are given by (48) and (49) respectively.

$$\mathcal{C} = B_1 \tag{48}$$

$$\mathcal{O} = I \tag{49}$$

As the observability matrix is just the identity matrix, there are no unobservable states. However, there is one uncontrollable state, which arises because the *M* matrix which is part of B_1 does not have full rank. The system can be reduced to a fully controllable *N-1* state system by performing another linear transformation. A new matrix *T* is the new transformation matrix to create the state vector *z* from *m* in (50). The bar accent denotes that the matrices are for the transformed system. *T* is an *N-1*x*N* matrix calculated such that the overall input-output transfer function (46) is not modified even though the states are different.

$$\begin{aligned} \boldsymbol{z}(t) &= T\boldsymbol{m}(t) \\ \bar{A} &= TAT^T \end{aligned} \tag{50}$$

$$\bar{B}_1 = TB_1$$
$$\bar{\boldsymbol{b}}_{\boldsymbol{2}} = T\boldsymbol{b_2}$$
$$\bar{C} = CT^T = T^T$$

Further details on the transformation and its derivation are given in Submission 3. The transformed system is given by (51), and the dynamics are in discrete form by (52). The state vector $\boldsymbol{z}$ lacks a physical meaning compared to $\boldsymbol{m}$ or $\boldsymbol{s}$.

$$\begin{aligned} \dot{\boldsymbol{z}}(t) &= \bar{B}_1\boldsymbol{u}(t) + \bar{\boldsymbol{b}}_{\boldsymbol{2}}d(t) \\ \boldsymbol{m}(t) &= \bar{C}\boldsymbol{z}(t) \end{aligned} \tag{51}$$

$$\boldsymbol{z}_{k+1} = \boldsymbol{z}_k + \Delta t\bar{B}_1\boldsymbol{u}_k + \Delta t\bar{\boldsymbol{b}}_2 d_k \tag{52}$$

With the uncontrollable state removed, there is now one more control input than there are states, meaning the system is over-actuated [148]. As a result there is more than one combination of inputs which can achieve the same control objective, including one solution whereby any one of the control inputs can be maintained at zero. A sub-problem can be developed to select the best option [149], for example by finding the current which results in the greatest imbalance reduction without exceeding the maximum balancing current. It also means that there is redundancy in the event of a failure.

5.3.3 Summary

A number of steps are required to take the conventional cell equations and formulate an appropriate system model, given by (51), for balancing controller design. These steps are summarised below:

- The control input to the system is a vector of balancing currents, and the disturbance is the scalar load current applied to the pack. The output has been chosen as ΔSOC (the distance of each SOC from the mean). This means that ΔSOC is assumed to be a measured output, although in reality it is an estimate from the BMS.
- The states have also been chosen as ΔSOC since this is what will dictate the magnitude of each balancing current. The target vector for the states is always zero and so does not have to be continually updated by the balancing control system as the cell SOCs change through pack operation.

- The formulation of this system model is not fully controllable and so feedback algorithms may not be able to drive the system to the appropriate state. A state transformation has been derived which results in a fully controllable and observable system with one fewer states.

5.4 Proposed control systems

The ultimate aim of the controller is to achieve some form of feedback with the module of cells. In generic form, this is defined by (53), where the balancing current vector $\boldsymbol{u}$ is a function of the transformed imbalance vector $\boldsymbol{z}$, and time *t.*

$$\boldsymbol{u} = f(\boldsymbol{z}, t) \tag{53}$$

In this section, three methods of achieving closed-loop feedback are presented, two of which were ultimately used in the experimental test work in Chapter 6. A variety of established control methods were explored in Submission 3, but most were not developed further and so are not considered here. A summary of the other methods from Submission 3 is given below:

- Linear quadratic regulation (LQR). This creates a static feedback matrix *K* based on minimising the quadratic cost function (54) according to the control law (55). The cost function is dependent on the two weighting matrices *Q* and *R*, which change the contribution of $\boldsymbol{x}$ and $\boldsymbol{u}$ to the total cost, which changes *K*. This is a form of state feedback, discussed below, and the implementation is similar to pole placement, one of the selected control methods. However, one of the stated aims of the control system is to ensure all cells finish balancing at the same time, which requires their poles be equal. Unlike pole placement, LQR does not explicitly set specific pole values, making it less suitable for this particular problem[11].

$$J = \int_0^\infty \boldsymbol{x}(t)^T Q \boldsymbol{x}(t) + \boldsymbol{u}(t)^T R \boldsymbol{u}(t)\, dt \tag{54}$$

$$\boldsymbol{u}(t) = -K\boldsymbol{x}(t) \tag{55}$$

- Pole placement with disturbance gain. The standard state feedback approach does not account for a disturbance with a non-zero mean. Such

[11] In Submission 3, a method was derived for selecting the weighting matrices such that all poles are equal. This provides the same end result as pole placement, but the method is more involved and there is no reason to adopt it over conventional pole placement methods.

a disturbance will prevent the states from reaching their target values, resulting in a steady-state error. An additional static gain matrix can be used to apply feedforward control based on the measured disturbance input [150]. In theory it can cancel out the disturbance's contribution to imbalance. This is a suitable extension to pole placement to improve its performance if necessary. However, its performance is limited by the balancing current constraints, and assumes a perfect measurement rather than accounting for more general model uncertainty.

- LQR with integral action. In general, Integral action can be used to avoid steady-state errors which arise through model uncertainty [151]. This method adds additional states to an LQR system as a means of tracking the error history, and removing the steady state error. However, the additional model states increase complexity. Additionally, the integral poles have to be selected as well as the imbalance pole. These can be difficult to tune, and there will be a trade-off between speed and stability.
- H-infinity output feedback [152]. H-infinity involves minimising the infinity-norm of the transfer function between the disturbance and the output. Minimising for the 2-norm (such as in LQR) means minimising the transfer function gain across all frequencies (suitable for a white-noise input [153]), whereas for the infinity-norm, the peak gain is minimised. The formulation incorporates a matrix for the disturbance input as well as for the control input, removing the assumption that the noise is a zero-mean signal. This makes it possible to design a controller specifically to achieve the best reaction to the worst disturbance case. Simulation results from Submission 3 suggest this should perform similarly to pole placement but with a faster initial imbalance reduction and lower steady state error. It is more difficult to implement, requiring more tuning, so is a suitable choice to consider if pole placement performs reasonably well but needs improvement.

Additionally, including an estimator in the system was also considered. This could act to refine the SOC estimate from the BMS, and is inherent to some controllers, such as the standard H-infinity formulation [154]. However, the results presented

within Submission 3 show that it did not appear to improve performance, while still adding complexity to the system.

5.4.1 Rule-based Method

As a baseline, a non-model based system is considered, where the cells are simply switched on according to the logic in (56). The variable δ is a small tolerance to prevent indefinite switching, in this case 0.01. A larger δ will reduce the amount of additional switching but prevent the system converging to a low state of imbalance.

$$u_n = \begin{cases} u_{max}, & m_n > \delta \\ u_{min}, & m_n < -\delta \\ 0, & -\delta < m_n < \delta \end{cases} \tag{56}$$

As the model is not required it is straightforward to implement. However, by only switching the system on or off, it means all cells balance at the same rate even if some are much further from the mean SOC than others. Some cells will reach the mean SOC prematurely, which has two disadvantages. Firstly, i^2R power losses will be higher and secondly, the cells will undergo some additional switching as they are maintained at the mean, which also increases energy loss.

5.4.2 Pole Placement Output Feedback

State feedback is a commonly used approach for multiple-input-multiple-output systems due to its relative simplicity and inherent stabilising nature [153]. This uses the feedback law (57), where K is a constant matrix to be calculated. For this system, the output $\boldsymbol{m}$ is required to drive the control signal, not $\boldsymbol{z}$. $\boldsymbol{m}$ can be factored into by using the transform relationship (50).

$$\boldsymbol{u}(\boldsymbol{t}) = -K\boldsymbol{z}(\boldsymbol{t}) \tag{57}$$

$$\boldsymbol{u}(\boldsymbol{t}) = -KT\boldsymbol{m}(\boldsymbol{t}) \tag{58}$$

The steady state error when a step disturbance is applied can be calculated by considering the point at which the derivative of $\boldsymbol{z}$ is zero, which is given by (59). The resultant steady state output error $\boldsymbol{m}_{ss}$ is given by (60)[12]. As is expected, there is no steady state error when there is no disturbance.

[12] K, $\bar{B}_1$ and $\bar{B}_2$ are all rectangular so do not strictly have an inverse. However, unlike with the pseudoinverse, the equations can be solved using similar methods as for square matrices, and so the superscript $^{-1}$ is still used.

$$\mathbf{0} = -\bar{B}_1 K \boldsymbol{z_{ss}} + \bar{\boldsymbol{b}}_{\boldsymbol{2}} d \tag{59}$$

$$\boldsymbol{m_{ss}} = T^T K^{-1} \bar{B}_1^{-1} \bar{\boldsymbol{b}}_{\boldsymbol{2}} d \tag{60}$$

Two common methods of designing a state feedback regulator are pole placement and LQR [124]. Both methods are considered in Submission 3, but pole placement was carried forward as the tuning parameters are more intuitive for this system than LQR and it is easier to ensure the cells reach EOB simultaneously. Pole placement is used to calculate a gain matrix such that a desired vector of poles is achieved for the closed loop system. The time constant of a pole is obtained using (61), so it is straightforward to select a pole based on the time available for balancing. In their standard form, both methods assume there is no disturbance (or that the disturbance has a mean of zero), which means there will be steady state error in the control solution. As noted above, three methods for accounting for the disturbance were considered in Submission 3: integral action, H-infinity output feedback, and a disturbance gain. These were not implemented experimentally and so are not discussed here.

$$\tau = -\frac{1}{\lambda} \tag{61}$$

Iterative algorithms are available [155] to find *K* such that each state has the desired pole. However, for this case, it can be simplified. Two important aspects of the balancing system model are that the *A* matrix is zero and the desired poles are the same for each cell (so they will all reach the mean SOC simultaneously). The closed loop dynamics are given by (62) and the poles are found by solving the eigenvalue problem in (63). The points above simplifies the problem considerably compared to many other linear state space systems, and (64) is the most straightforward solution to (63). In this case λ is a scalar and real, and should be negative to ensure stability.

$$\dot{\boldsymbol{z}}(t) = -\bar{B}_1 K \boldsymbol{z}(t) \tag{62}$$

$$|-\bar{B}_1 K - \lambda I| = 0 \tag{63}$$

$$\bar{B}_1 K = -\lambda I \tag{64}$$

The steady state error can be calculated using (65), which substitutes (64) into (60). This shows that a slower balancing pole results in a larger steady state error. As such the fastest pole should be chosen that the balancing hardware is capable of. The same applies with the disturbance magnitude.

$$\boldsymbol{m_{ss}} = -\frac{1}{\lambda} T^T I \bar{\boldsymbol{b}}_{\boldsymbol{2}} d = -\frac{1}{\lambda} M \boldsymbol{b_d} d \tag{65}$$

The most direct way of obtaining *K* is to solve (64) by finding the inverse of $\bar{B}_1$ to give (66). Since it is rectangular it does not have an inverse: there are essentially *N* unknowns but only *N-1* equations, making it an under-determined system. As derived and discussed within Submission 3, there are two main ways of solving this: a least-squares solution, and singular value decomposition (SVD). The least-squares solution results in the *K* matrix containing a row of zeros, making the balancing input corresponding to that row redundant. The SVD approach distributes the control action across all inputs. MATLAB's pole-placement algorithm `place` [156], which calculates a robust solution resulting in a non-zero *K* matrix, calculates the same matrix as the SVD approach.

$$K = \bar{B}_1^{-1}(-\lambda I) \tag{66}$$

For these solutions, the inverse of B_c and its associated matrices can be stored as a static matrix (and periodically updated if the elements of B_c change from the SOH estimation). *K* can then be simply calculated based on the desired pole value. This pole can be adjusted based on the estimated time until EOD or EOC, and updated regularly because of the simple calculation required to obtain the new *K*. These methods are generally reliable solutions to obtain since they are based around well-known and robust matrix techniques rather than more complex algorithms. It is also clear that the gain is proportional to the chosen pole, which means if the pole results in excessive balancing currents, it can easily be adjusted to bring the currents to within a feasible range.

The two methods show how there can be multiple *K* matrices which produce the same poles. Using the least squares method means that one cell will also be at zero balancing current, which offers redundancy in the event of one balancing channel failing. The SVD solution distributes the control effort across all of the cells, which means that for a given pole, the maximum balancing current will be lower and there is less chance of saturating one of the control inputs.

The pole can be chosen based on the estimate of time to EOD. Alternative, the fastest pole which does not saturate any inputs can be found. A simple bisection

algorithm can be used to search for the pole. This uses the current pole to calculate the state feedback matrix. This, along with the initial ***m*** vector, is used to calculate the initial balancing currents. If any of the currents are outside of the limits imposed by the hardware, the pole is decreased for the next iteration. If none of the currents are above 99% of the limit, then the pole could be faster and so is increased. If neither of these conditions are met, then a suitable pole has been found, for which at least one of the cells is set to balance at its maximum value. Example code to demonstrate the algorithm is given in Figure 37.

```
lambdaMin = -1/10000; %less likely to cause 100%duty
lambdaMax = -1/10; %more likely to cause 100% duty
while nIter < nIterMax % avoid searching indefinitely
    nIter = nIter + 1; % count number of iterations

    lambda = (lambdaMax+lambdaMin)/2; %update lambda as
halfway between max and min

    % ...
    %Here, K is solved for using lambda value
    % ...
    u0 = K*T*m0; %the initial current requested by controller
    if any(u0 < uMin || u0 > uMax)
      %Input saturated: reduce the maximum to give a slower
      pole next iteration.

        lambdaMax = lambda;
    elseif (max(u0) < 0.99*uMax) || (min(u0) > 0.99*uMin)
        %pole is too slow: no inputs are at the maximum.
Increase the
        %minimum so next iteration a faster pole is used.
        lambdaMin = lambda;
    else
        break; %a suitable pole has been found.
end
```

Figure 37: Example code for bisection algorithm

5.4.3 Model Predictive Control and Dynamic Steady-Input Control

MPC is a framework which uses a plant model to calculate future control inputs over a finite time horizon, by minimising an objective function [157]. MPC also uses a similar quadratic cost function as LQR (54), but there are two main differences in their implementation:

1. MPC optimises over a receding time-window, rather than the fixed (and infinite, for the LQR controller calculated in Submission 3) time-window.
2. LQR uses the optimal solution for the entire implementation: a static gain matrix calculated offline. MPC re-calculates the solution dynamically during online operation.

Compared to a regulator such as state feedback, MPC has the potential to calculate more efficient and faster imbalance reduction as it can incorporate knowledge of disturbances and adapt to changes in operating conditions. It is also easier to account for constraints to states and inputs [157]. However, it is generally more complex, more computationally expensive to solve for and more sensitive to model errors if a long time horizon is considered [157]. Active cell balancing will typically take a long time, within the order of 30 minutes or more. For a one second time-step, this means calculating at least 1800 control vectors, which results in a large optimisation problem (compounded as the number of cells increases) not suited to online calculation.

For the model in section 5.3 there are five sources of error when planning a trajectory, shown in Table 14. The combination of these errors means that solving a complete trajectory then applying the control input vectors is unlikely to result in an optimal or reliable solution.

Table 14: Sources of error in control model

Source	Description
Disturbance current *d*	The future inputs are not known, adding uncertainty to any trajectories. For example, if a BEV was at rest when balancing started, but was then used for motorway driving, the impact of the disturbance has changed compared to when the trajectory was calculated
Control currents *u*	Set-point errors or variations in current mean that the requested currents may not actually be applied, meaning the calculated trajectory is not exactly applied.
Cell capacities Q_{cap}	The cell capacities well generally be an estimate and as such subject to some uncertainty. The model calculates ΔSOC using coulomb counting, which is proportional to capacity. As such, incorrect capacity estimates will result in incorrect ΔSOC projections.
Feedback matrix *F*	This matrix models how the balancing currents impact the other cells. Inaccuracies in this matrix, for example ignoring the balancing efficiency, will result in incorrect trajectories.
ΔSOC estimate *m*	The ΔSOC estimates determine the initial conditions for the MPC scheme. If these initial conditions are incorrect then the control vectors calculated will not actually remove imbalance as intended.

Instead of solving for a time-varying solution, a single control vector can be solved for to apply until balancing is complete. This can then regularly be updated based on the latest SOC data. The objective function used for this controller is the sum-of-squares quadratic cost function (67). Quadratic cost functions are often used for control systems because there are a variety of robust methods to solve them [157], and its value approaches zero as the state vector approaches zero (i.e. when there is no imbalance). The optimisation problem statement is given by (68) and aims to find the balancing current vector which minimises total cost *J*, subject to the lower and upper inequality constraints ***lb*** and ***ub*** that limit the balancing currents to what the hardware is capable of. While (67) does not explicitly contain the ***u*** vector, given the discrete system dynamics in (52), (67) can be rewritten as (69). This can be grouped to create a standard quadratic programming problem (70). This is broken into three components: the quadratic component by matrix *H*, the linear component from vector ***g*** and the constant term *f*.

$$J = \boldsymbol{m}_{k+1}^T \boldsymbol{m}_{k+1} = \boldsymbol{z}_{k+1}^T \boldsymbol{z}_{k+1} \tag{67}$$

$$\arg\min_{\boldsymbol{u}_k} J(\boldsymbol{u}_k) \qquad s.t.\, \boldsymbol{lb} \leq \boldsymbol{u} \leq \boldsymbol{ub} \tag{68}$$

$$J = \left(\boldsymbol{z}_k + \Delta t \bar{B}_1 \boldsymbol{u}_k + \Delta t \bar{\boldsymbol{b}}_{\boldsymbol{2}} d_k\right)^T \left(\boldsymbol{z}_k + \Delta t \bar{B}_1 \boldsymbol{u}_k + \Delta t \bar{\boldsymbol{b}}_{\boldsymbol{2}} d_k\right) \tag{69}$$

$$J = f + \boldsymbol{g}^T u + \boldsymbol{u}^T H \boldsymbol{u}$$

$$\begin{gathered} f = \boldsymbol{z}_k^T \boldsymbol{z}_k + \Delta t^2 d_k^T \bar{\boldsymbol{b}}_{\boldsymbol{2}}^{\boldsymbol{T}} \bar{B}_2 d_k + 2\Delta t d_k \bar{\boldsymbol{b}}_{\boldsymbol{2}}^{\boldsymbol{T}} \boldsymbol{z}_k \\ \boldsymbol{g} = 2\Delta t \bar{B}_1^T \left(\boldsymbol{z}_k + \Delta t \bar{\boldsymbol{b}}_{\boldsymbol{2}} d_k\right) \\ H = \Delta t^2 \bar{B}_1^T \bar{B}_1 \end{gathered} \tag{70}$$

The optimisation problem means that a balancing vector will be found which removes the largest amount of imbalance feasible from time step *k* to *k*+1. However, this period of time is only 1 second, meaning that it will be impossible to reach a completely balanced state unless there is only a very small amount of imbalance. It means that generally all cells will be saturated permanently, which is not the aim of the control system. Some cells will have been balanced faster than necessary, reducing efficiency through increased power loss from the higher

currents. The end result is that it is essentially the same as using a rule-based approach where all cells below the mean are set to charge, and all cells above the mean are set to discharge. For example, McCurlie et al. [112] found that using a rule-based scheme resulted in almost 50% higher energy loss through balancing compared to a similar MPC scheme to the one presented here. The exact amount of energy saved will depend on the initial SOCs and cell capacities.

To avoid the saturation issue, the time-step itself is adjusted here to project further forward in time (extending the time in which the charge is delivered, thereby reducing the average current) . Changing Δt in (52) changes the balancing window: a shorter time will result in faster balancing, with increased likelihood that the inputs are all saturated. A longer time will bring the cells to the mean SOC simultaneously, but can increase the balancing time. The time window can be maintained while continually recalculating the control vector based on new SOC and disturbance current. Alternatively, the time window can be rescaled based on current performance. Equation (71) shows the scheme employed to adjust the time-step. If any of balancing currents are at the maximum, and the average current is within 95% of the maximum, then the window is increased by 10%. Similarly, if no inputs are at the maximum it is possible to balance faster than at present, so the window can be reduced by 10%. This avoids all inputs being saturated while still allowing for saturation to occur, which is useful for if one cell is particularly far from the mean compared to the others.

$$\Delta t_{k+1} = \begin{cases} 1.1\Delta t_k, & \exists \boldsymbol{u}(\boldsymbol{u} = u_{max}) \wedge \sum_{n=1}^{N} \boldsymbol{u}_n \geq 0.95 N u_{max} \\ 0.9\Delta t_k, & \forall \boldsymbol{u}(\boldsymbol{u} < u_{max}) \end{cases} \tag{71}$$

The disturbance current is part of the control system, but is not known in advance, so an estimate of what the current will be is required. As this is a single-step solution, a detailed prediction of a time-varying current profile is not required, only the average current during the following time window. This could be obtained from the vehicle's energy management system, which makes predictions of how the vehicle will be driven to give an estimate of how many miles can be driven before the battery pack reaches EOD [158]. Alternatively, the average current over a previous time window can be used for the following time step.

In summary, there are three key differences from this method compared to MPC:

1. The quadratic cost function has been altered to maximise imbalance reduction.
2. A single input vector is solved for, rather than a series of vectors over the time horizon. This significantly reduces the optimality of the solution, while being more feasible to solve online. This input vector is found on the basis that it would be applied for the rest of the time input.
3. The solution is updated every time-step, not just when the time horizon of the current solution is reached. This ensures that any sudden changes in SOC estimate are accounted for as early as possible.

While this control strategy was inspired by the principles of MPC, it is distinct and is referred to here as Dynamic Steady-Input Control (DSIC). Figure 38 shows an example process for DSIC. A change in SOC estimates means that imbalance can no longer be removed by the requested window. A longer window is calculated, and from then on the control vector is updated at each time-step to account for small changes in SOC estimation and model errors.

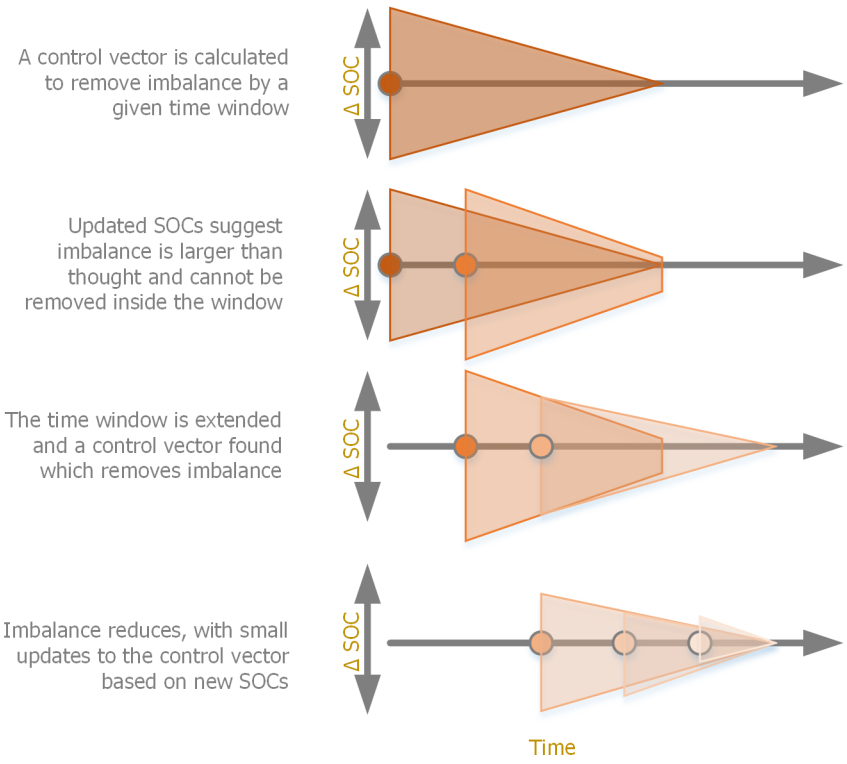

Figure 38: DSIC imbalance removal

5.4.4 Feed-forward Control

There are some downsides to the DSIC scheme. It is relatively complex, requiring the solution to an optimisation problem every time-step. Similarly, the maximum imbalance removal is not always required. The overall balancing speed will always be limited by one cell, typically the cell whose SOC is furthest from the mean (though this is dependent on the relative cell capacities and the nature of the interconnections). Balancing will be fastest when setting this cell to balancing at maximum current for the duration of balancing. An efficient but quick balancing system would therefore aim to do this, while minimising the energy loss for the other cells. As long as these cells also reach EOB by the time the slowest cell does, then there is no overall reduction in battery pack performance.

Feed-forward was first introduced within a conference paper published by the researcher on high-level balancing modelling, where it was used to show how passive balancing could be made more efficient [8]. It can be further developed using the control model, and also considering the time window from DSIC. Using the discrete state equation (52), a target can be formed such that there is no imbalance at the next step (72). There are two unknowns to solve for: Δt and $\boldsymbol{u}$. A direct solution can be found if the right-hand disturbance term is removed. This adds some inaccuracy as the disturbance is not accounted for, but this is difficult to account for anyway as the future trajectory is not known, and as the simulation results for DSIC in section **Error! Reference source not found.** showed, the disturbance contributes only a small amount to the balancing current calculation. The total balancing charge vector $\boldsymbol{Q}_{bal}$ can then be calculated using the SVD inversion method (73). This charge needs to be decomposed into the balancing currents and the time-step. The time-step can be found using (74). This scales the charge against the lower and upper balancing current limits ***lb*** and ***ub*** to find the point at which one of the balancing currents is equal to the maximum charge or discharge amount permitted. This cell governs overall balancing speed.

$$0 = \boldsymbol{z_k} + \Delta t \bar{B}_1 \boldsymbol{u_k} + \Delta t \bar{\boldsymbol{b}}_{\boldsymbol{2}} d_k \tag{72}$$

$$\boldsymbol{Q_{bal}} = -\bar{B}_1^{-1} \boldsymbol{z}_k + \bar{B}_1 \tag{73}$$

$$\Delta t = \max\left[\max\left(\frac{\boldsymbol{Q}_{bal}}{\boldsymbol{lb}}\right) \quad \min\left(\frac{\boldsymbol{Q}_{bal}}{\boldsymbol{lb}}\right)\right] \tag{74}$$

$$\boldsymbol{u_k} = \frac{1}{\Delta t}\boldsymbol{Q}_{bal} \tag{75}$$

As with DSIC, while this is in theory the current to apply for the duration of balancing, it is regularly updated to account for uncertainty in the SOC estimates and applied load. One of the drivers behind both controllers is that the SOC estimation is a significant source of error and the main limiting factor in projecting forward in time. As can be seen, the feed-forward implementation is much simpler than DSIC.

5.5 Simulation results

To compare these control methods, the automotive pouch cells first introduced in section 4.2 are used within a simulation study. As these controllers are designed for a generic balancing system, there is no explicit modelling of any hardware. However, the following assumptions are made, based on typical balancing performance from the literature:

- The balancing system can produce balancing currents of up to 2A charge or discharge per cell simultaneously.
- The balancing system is an idealised cell-to-module system, with the same *F* matrix as in equation (41).

One of the key issues identified in the literature review is that there has been very little consideration of balancing under load, which is necessary to account for capacity differences between cells. For the simulation studies, the US06 Highway drive cycle was used to generate a load current, which was chosen as it is more aggressive than many of the other standardised drive cycles, such as the New European Drive Cycle and Federal Test Procedure drive cycles, and so provides a more thorough test of the balancing controller. A battery current profile was derived from the speed profile using a vehicle model of a JLR BEV in pre-production. The current from that model was rescaled to give the same C-rate for these cells as for the battery pack in the vehicle model. The speed profile and current are shown in Figure 39. The average C-rate is 0.13, which is low compared to many other applications, but is necessary to ensure the vehicle can be driven for several hours. The drive-cycle is looped several times to achieve the necessary length of simulation necessary and the required DOD.

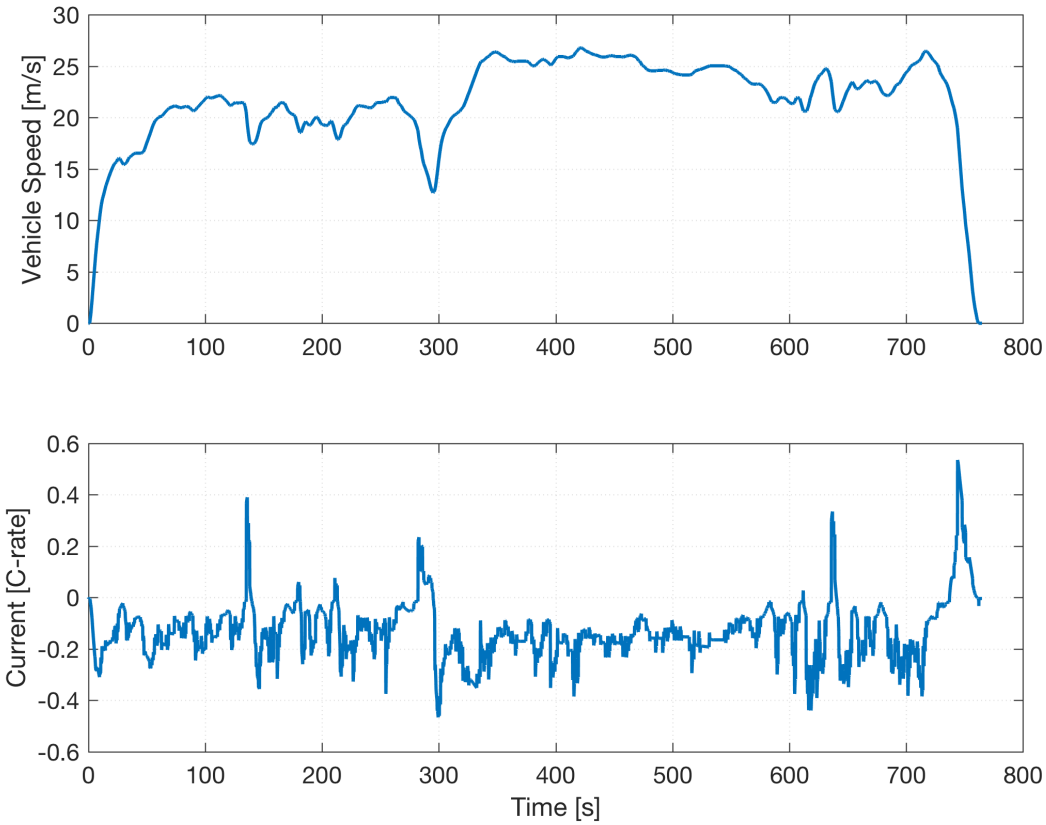

Figure 39: US06 drive cycle speed and current

The initial cell SOCs are chosen by assuming the cells were all fully balanced at 100% SOC, and have then been discharged a certain DOD. This is expressed as a function of the weakest cell's SOC, as this will govern the pack SOC. Each cell's initial SOC is calculated using (76), and example SOC vectors at different DODs are shown in Table 15.

$$s_n(0) = 100 - DOD\frac{\min(\boldsymbol{Q_{cap}})}{Q_{cap_n}} \tag{76}$$

Table 15: Initial SOC vectors at different DODs

DOD (%)	SOC Vector (%)	Absolute SOC difference (%)
40	$[\mathbf{60}\ 61.2\ 64.5\ 64.5\ 63.9\ 63.8\ 63.6]$	4.5%
60	$[\mathbf{40}\ 41.9\ 46.8\ 46.7\ 45.9\ 45.7\ 45.4]$	6.8%
80	$[\mathbf{20}\ 22.5\ 29.0\ 28.9\ 27.8\ 27.6\ 27.3]$	9.0%

SOC estimators are not used as the aim of these simulations is to evaluate the potential of the controllers to remove imbalance and SOC estimation could result in an unequal comparison. The impact of SOC estimation on controller

performance is discussed further in Chapter 6. All of the simulations were performed with the 80% DOD vector from Table 15, again to ensure an equal comparison between controllers.

5.5.1 Rule-based method

The results for the rule-based approach are shown in Figure 40. Note that the simulation was halted at the point where the balancing currents would be set to zero. As expected, the other cells, particularly cells 5-7, reach the mean prematurely, and then undergo some additional switching as they are maintained at the mean. This particular set of conditions is fortuitous, as the two cells furthest from the mean are cells 1 and 2, which are also the only two cells below the mean. The remaining cells are all discharged, which significantly boosts the charging current of cells 1 and 2, reducing balancing time. Other imbalance distributions may not be as amenable to imbalance reduction as this.

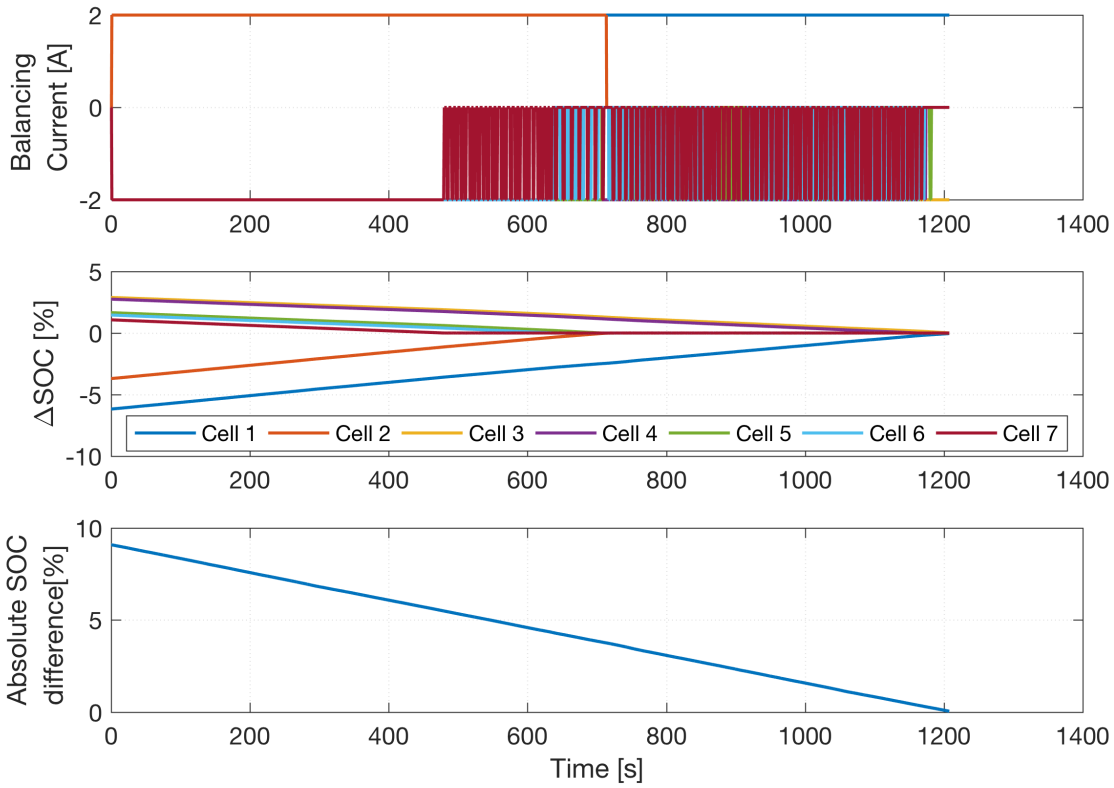

Figure 40: Rule-based logic while under a drive cycle load

5.5.2 Comparison of Pole Placement Design Methods

As noted in section 5.4.2, there are two ways of calculating the pole placement matrix: SVD and least squares. These are compared below, without the drive cycle in order to focus on how they reduce imbalance. To validate that they can

both reduce SOC in the same way, a pole with a time constant of 2400s was used to ensure that neither controller saturates the inputs. The results using SVD are shown in Figure 41, and the same test using least squares in Figure 42.

As can be seen from the ΔSOC plots, the imbalance vector follows the same trajectory. However, the balancing current plots are quite different. The least squares method has set cell 3 (the cell with the largest capacity and so highest initial SOC) to zero current, and the result is that all balancing currents are positive. The different balancing current magnitudes combined with the cell-module feedback ensures that the cell currents are a mix of positive and negative, despite all the balancing currents being positive, thus still bringing all cells to the mean SOC. The SVD method produces negative balancing currents for cells 3-7, and the nature of the feedback means these will charge cells 1 and 2. As a result, these cells do not need as high a charging current as the least squares method. For least squares, cells 1 and 2 start at 1.9A and 1.4A respectively, as opposed to 1.3A and 0.8A for the SVD method.

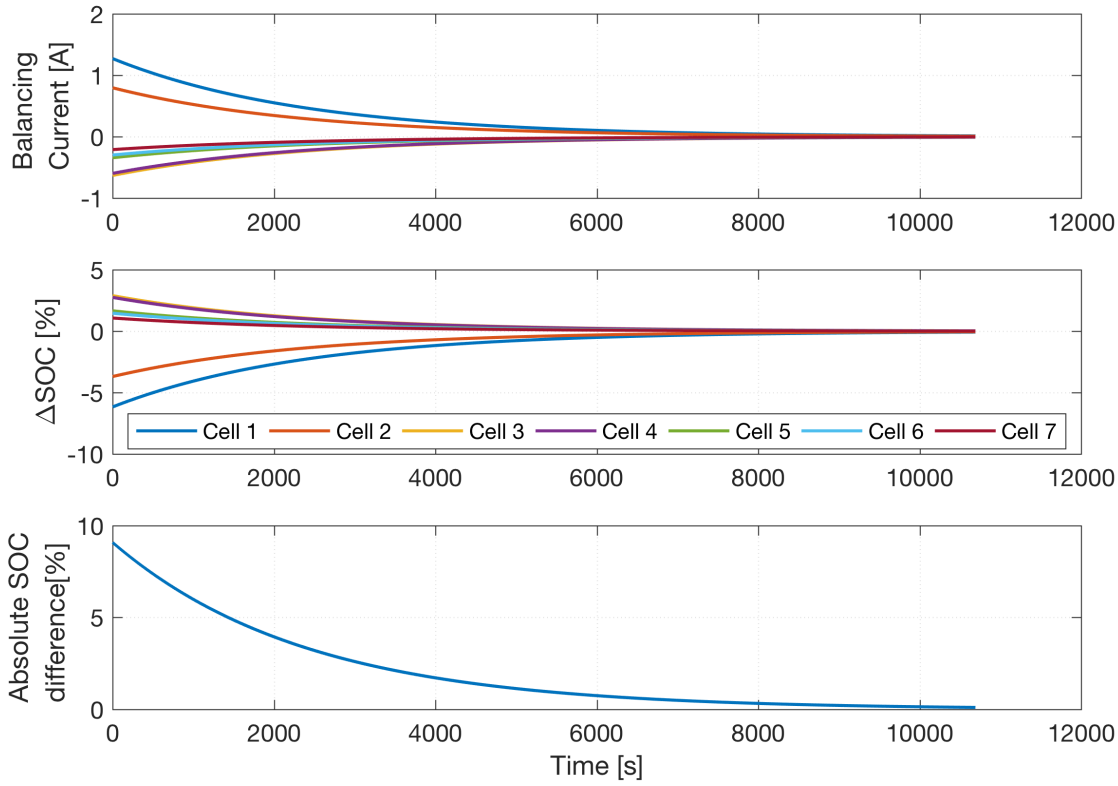

Figure 41: Pole placement using SVD under no load

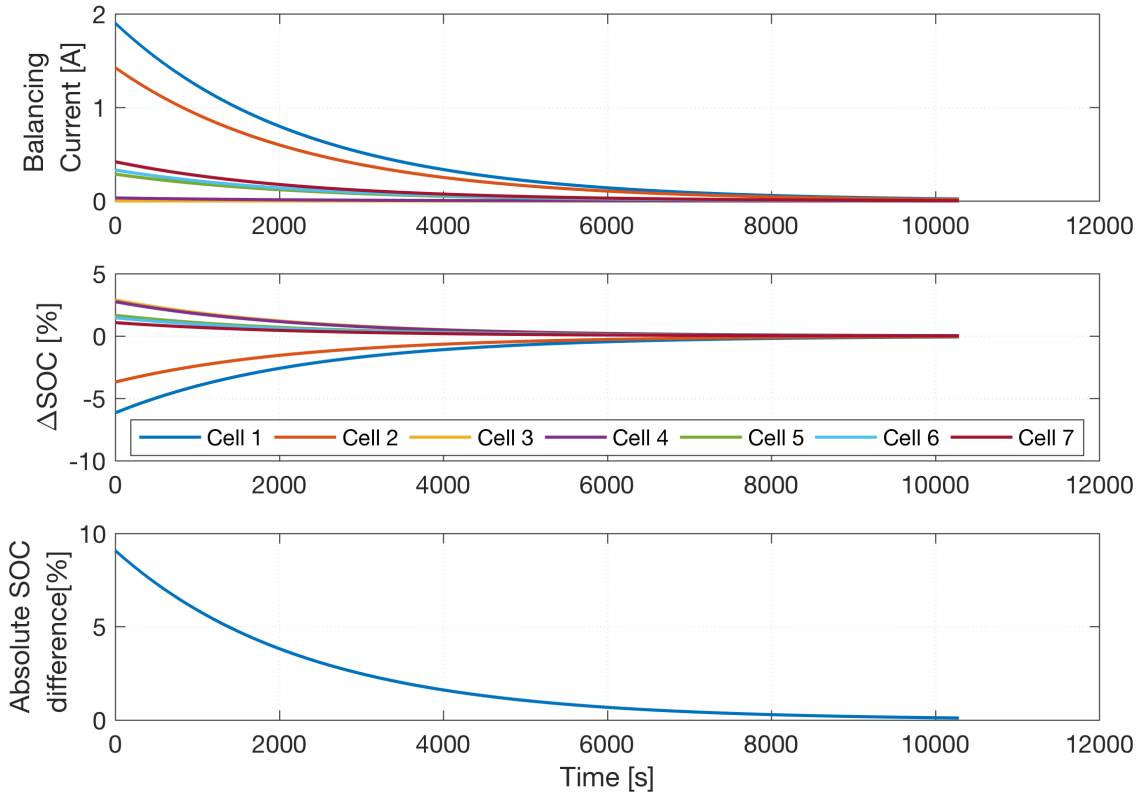

Figure 42: Pole placement using least squares under no load

Using the bisection method in 5.4.2, the maximum time constant achievable for this simulation is 1541s with SVD, but 2297s with least squares. Using SVD spreads the load more evenly across the cells and allows greater imbalance reduction per unit of time.

5.5.3 Pole placement under load

With SVD being chosen as the means of generating the feedback matrix, its performance under load is now evaluated. The time constant, found using bisection, was 1541s. The results in Figure 43 show that by 6800s, the imbalance has reached 0.83%. The weakest cell had an SOC of 1.6%, meaning that that it took a DOD of 18.4% to complete balancing. As (65) shows, this is proportional to the disturbance, meaning that the error would be greater in applications with a higher C-rate.

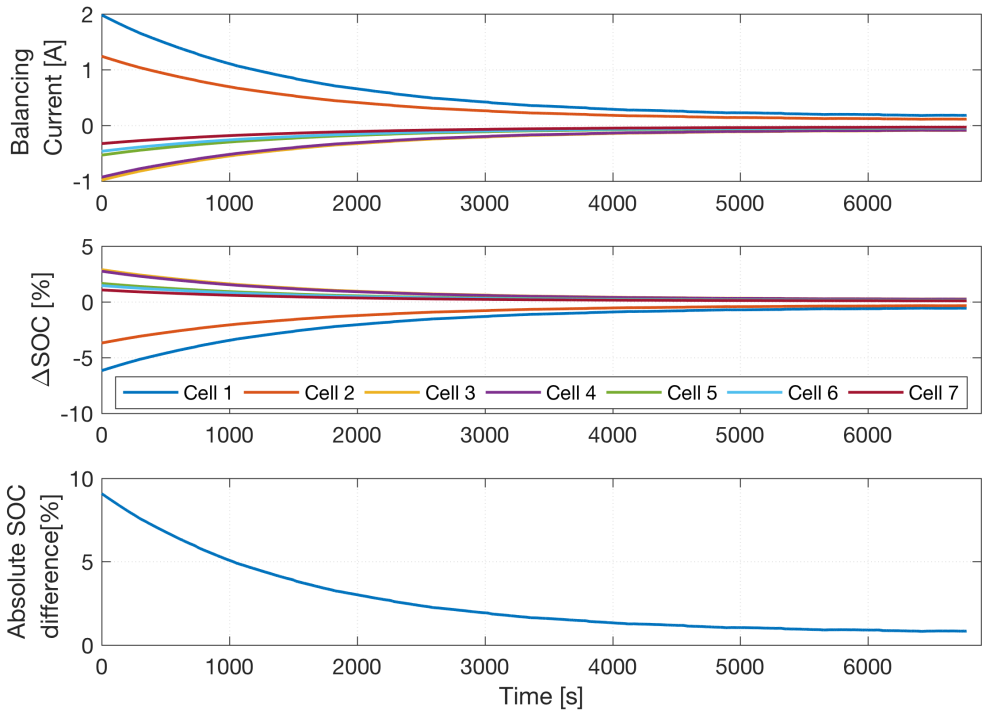

Figure 43: Pole placement while under a drive cycle load

Methods for reducing or eliminating this steady state error are outlined in Submission 3 and section 5.4. However, these results show that while pole placement is straightforward to implement, it creates a smooth trajectory at the expense of balancing speed. Its exponential nature means that the speed of imbalance reduction will reduce as imbalance itself reduces. The speed could be increased by periodically rescaling the controller to finding the new fastest pole as imbalance reduces, but this is counter to the fundamental negative feedback behaviour of the controller.

5.5.4 DSIC

The MPC simulation was initialised with a window of 200s. The results are shown in Figure 44. Note that the simulation was halted at the point where the balancing currents would be set to zero. The jagged changes in current, for example at 600s arise from the time window changing according to equation (71). The cells do not quite reach EOB simultaneously because of the simplifications that come with the method. However, the EOB difference is small compared to the balancing time. In this time, the minimum cell SOC has only decreased by 0.2%, while the total imbalance is at 0.07%, so effectively no steady state error.

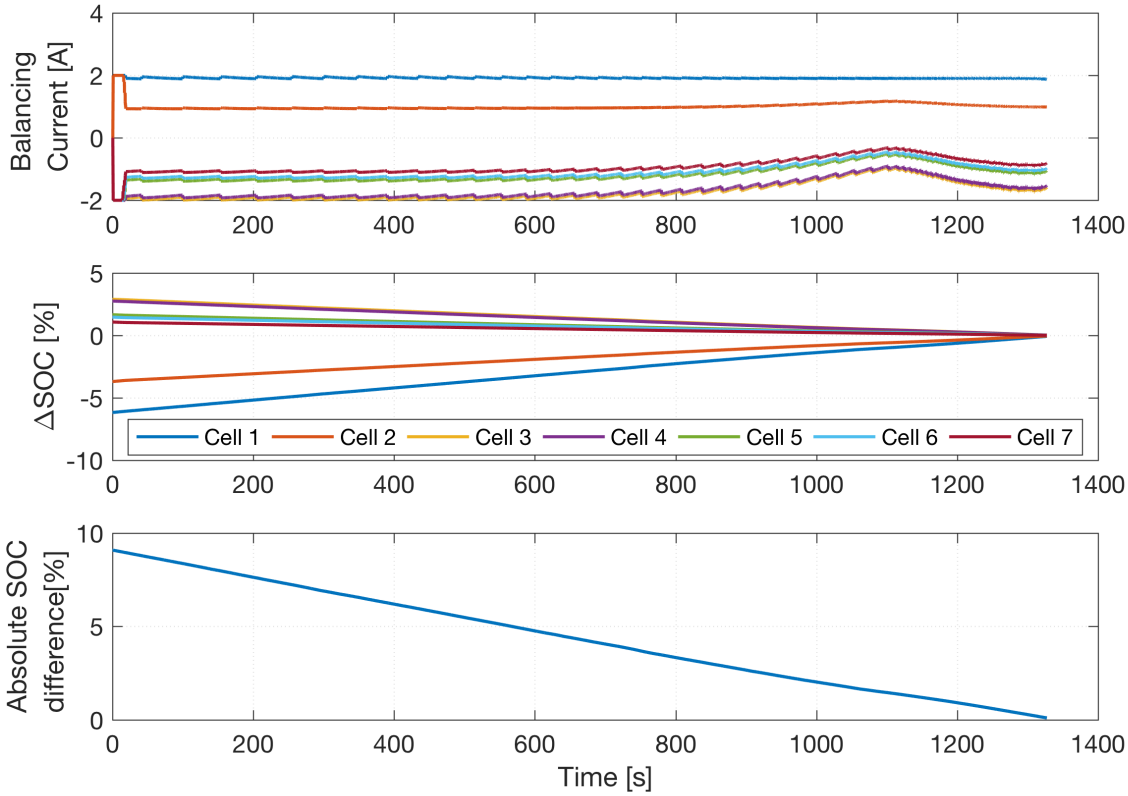

Figure 44: DSIC while under a drive cycle load

This scheme also used the load current directly as a source, rather than using an algorithm to predict the future current. It was found that an accurate disturbance current had little impact on controller performance, as it only contributes a small amount to the cost function. Removing the disturbance from the cost function resulted in a final imbalance of 0.08%. The ***g*** vector in (70) is the only term in the optimisation which contains the disturbance term. For these initial conditions, the disturbance only contributed to 0.5% of the total *J* value. Even when imbalance is very low, the disturbance component is still smaller than the ΔSOC term.

5.5.5 Feed-forward Control

The results for the feed-forward controller are given in Figure 45. Balancing takes longer than the rule based system. This is because feedback currents calculated by this scheme comes out as zero, meaning none of the balancing currents affects one another. For the rule-based scheme, the charging currents for cells 1 and 2 are bolstered by the remaining cells discharging. This means that cell 1, which takes longest to balance, undergoes a larger charging current for the rule-based system. This effect will vary depending on the distribution of cell currents and the connections between cells.

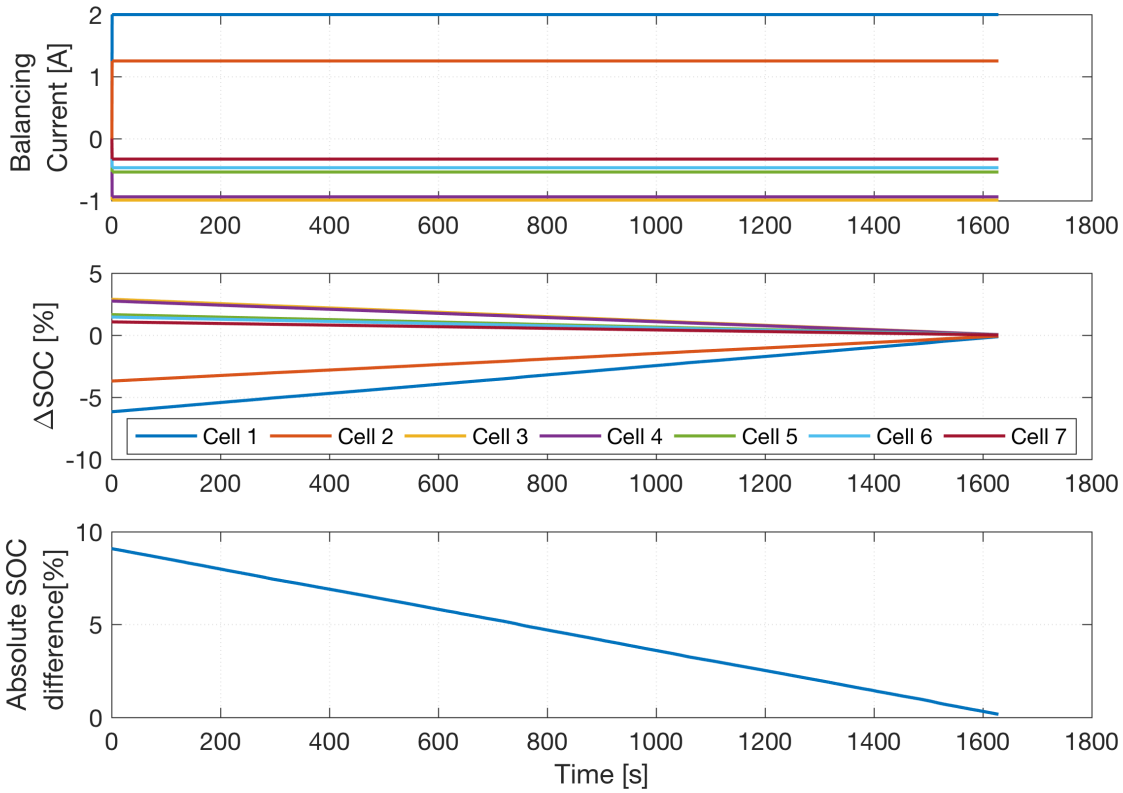

Figure 45: Feed-forward control while under a drive cycle load

5.5.6 Comparison

As noted, the balancing time for pole placement is slow, with the remaining three controllers all balancing in similar times. In reality, the amount of energy loss will be heavily dependent on the specific hardware used. For this generic system, it is assumed there are two components to the losses. There is a transformer efficiency which causes a power loss P_c, (77), and resistive losses P_s, given by (78). The transformer efficiency η_{tf} is assumed to be 95%, and the effective transformer resistance R_{tf} is set as 0.6Ω, which was chosen to give a total efficiency of 80% at maximum current.

$$P_c = \left(1 - \eta_{tf}\right) i_{bal} v_t \tag{77}$$

$$P_s = i_{bal}^2 R_{tf} \tag{78}$$

The energy usage for the four controllers implemented here are given by Table 16. As expected, the rule based system has consumed the most energy, with the feed-forward control approach only dissipating 78% as much energy. Pole placement used more energy despite operating a low currents, because of the longer balancing time. The energy wasted by the rule-based system is 9.88% of the total energy remaining in the cells at the start of the simulation, compared to

7.78% for the feed-forward controller. This allows for over 2% greater energy extraction, while is the ultimate aim of the balancing controller, and important considering how important relatively small gains in energy utilisation can be when a battery pack is operating across its full SOC window.

Table 16: Energy loss during balancing simulations

	Balancing Energy Loss (J)			
Cell	*Rule-based*	*PP*	*MPC*	*Linear*
1	2187.1	3311.8	2301	2942.7
2	1492.1	2082.6	1210.7	1849.5
3	2212.5	1659.1	2100.4	1472
4	2136.3	1571.3	2023	1394.2
5	1542.9	898.76	1428.8	797.57
6	1438.4	781.98	1325.7	693.97
7	1230.7	548.68	1119.3	486.94
Total	12240	10854	11509	9637

5.6 Summary

This section has outlined a means of controlling a generic balancing system based on cell SOC knowledge. The control model is linear and only dependent on one parameter (cell discharge capacity) per cell. Using cell capacity adds knowledge of whether the cell SOCs will naturally diverge or converge owing to the applied load. A large battery pack may not transfer energy from all cells, and instead by split into submodules of a few cells. The interaction matrix *F* was introduced to specify which cells connect to which.

The model has been used to create three different controllers, in addition to a non-model based controller. The results presented in section 5.5 show that pole placement takes a long time to balance and results in a steady-state error. This is at odds with the requirements of an effective balancing system: to balance fast and minimise the final SOC difference. The DSIC was derived which has target of removing the most amount of imbalance, while ensuring better convergence of balancing times. This avoids bringing some cells to the mean before others, at which point there will be extra switching on those cells as they are maintained at the mean. The feed-forward controller has the express aim of ensuring the cells reach EOB simultaneously, and was shown to use the least energy. Furthermore,

it has a slightly longer balancing time than the rule-based and DSIC approaches, but the fastest possible balancing is not crucial to performance: there is a compromise between time and wasted energy.

It is possible to refine the DSIC and feed-forward approaches to increase balancing time without significantly increasing energy usage. There are also more control methods which were not explored. For example, sliding mode control [159] could be effective, but was not considered here because of the considerable amount of design work required and implementation challenges with respect to “chattering” – continual oscillations about the target.

The purpose of this Chapter was to meet Research Objective 4: the design of a generic balancing control system. Many controllers were considered, and three have been described in detail and simulated to compare performance. This leads into the experimental work defined by Research Objective 5. For these controllers to be validated, they need to be applied to a specific ABS. The hardware used to do this, and the experimental results, are detailed in the next Chapter.

6 Hardware Implementation

The control framework in Chapter 5 was designed considering a non-specific balancing system. As per Research Objective 4, this was to demonstrate how a common model can be applied to a wide variety of balancing hardware, rather than becoming obsolete when the hardware changes or is updated. However, to validate the proposed control solution it is essential to demonstrate that the control system can be implemented with specific hardware, which is why Research Objective 5 was proposed. As well as testing the controller, this was also an opportunity to evaluate third-party hardware in more detail and from this learning experience produce recommendations for future balancing equipment. Finally, it meant implementing the control system in real-time subject to imperfect measurements and SOC estimation errors. This brings the research back up the systems engineering "V" in Figure 3. Using these experimental results, Research Question 2 can be answered by using the experimental results to quantify the benefits of active balancing.

Some issues with the experimental equipment are first highlighted in section 6.1. The balancing hardware used for this test work is described in section 6.2, followed by the instrumentation required to evaluate and monitor the system in section 6.3 and the load to apply a current to the series string of cells in section 6.4. The control system implementation and interfacing with the hardware is discussed in section 6.5. The balancing system is evaluated in section 6.6, and then the closed loop balancing test results presented and discussed in 6.7. Simulation studies to complement these results are contained in section 6.8, and the total system performance is analysed in section 6.9. Based on this test work, recommendations for hardware design and specification are detailed in section 6.10, followed by a summary of the Chapter in 6.11.

6.1 Experimental Background and Context

The balancing hardware was designed and built by an engineering consultancy which has previously worked JLR on Battery Management System (BMS) related development, and had previously developed a similar active balancing system for another automotive client. On using the equipment several issues were identified,

which were not all resolved before the company went into administration, which also meant that replacement parts were no longer available. When several cells were set to balance simultaneously, the firmware would go into an error state and prevent the balancing system from working. It was not possible to diagnose this fault condition given the level of documentation and support provided by the supplier. Only balancing 4 cells at a time avoided the issue apart from at low and high SOCs (typically lower than 20% and higher than 80%). In these regions, the cell voltages may have gone above or below limits imposed by the firmware. This limitation meant that it was not possible to test the high-level strategy described in Chapter 5, whereby cells are balanced towards the end of charging, and then only set to balance during discharge at low SOCs However, the system was suitable to demonstrate proof-of-concept for the active balancing system. Tests could still be run over a relatively wide SOC window (between 80 and 20% SOC), and there is no reason to suggest that if the control system can maintain a state of balance over this window, it could not achieve the same over a different SOC window, as long as the SOC estimator can maintain accuracy in these regions.

As testing progressed, another one of the balancing channels ceased operation (the reason for which is not known), which required rewiring and recalibration of the equipment. Eventually the entire circuit board ceased to function. This meant it was not possible to perform as many tests as intended, but the author asserts that the tests which were performed demonstrate proof of concept of the balancing control system.

6.2 Balancing Hardware

While there have been many types of balancing hardware proposed in the literature, very few have been implemented commercially. As the focus of this work is on control strategy rather than hardware design, an off-the-shelf system was preferred as it should be more robust and be quicker to set-up. The main commercial balancing system is the LTC3300-1 [88] by Linear Technology Corporation (LTC). This is a chip which is designed to drive one flyback converter per cell, with between 3 and 6 cells per chip. All cells can be set to balance (charge or discharge) simultaneously, which makes it more flexible and better suited to the control framework in Chapter 5.

The specifics of how energy is transferred between cells are covered in section 6.2.1. An implementation of this system was commissioned with JLR to an engineering consultancy with previous experience of the LTC3300-1 chip. This was designed with two chips, allowing for 6-12 cells to be connected. A 2A balancing current was specified, which is configurable by setting a resistor in the circuit. The intention was to be able to replace resistors with different values based on what cells are connected, but this proved unfeasible because the resistance required is in the region of a few milliohms, and as such is sensitive to any contact resistance and additional circuitry. This circuitry was complemented by a LTC6811 voltage measurement chip, designed for BMS applications. This measures up to 12 cell voltages as well having 7 analogue channels, which in this case were used to connect thermistors for temperature measurement.

The principle of operation means that the cell is only set to charge or discharge: there is no ability to regulate current directly. To add some level of current control, a duty cycling scheme was implemented within the firmware. It is not possible to update the chip faster than 10ms [88], and so the scheme is based on a new balancing command being sent over Controller Area Network (CAN) every 1s, and each percentage point then equates to 10ms of balancing. After the time specified by the duty cycle has elapsed, balancing for that cell is switched off, and remains off until another balancing command is sent. The average current over the second can be calculated using (79) where *d* is the duty cycle and $\hat{\imath}_{bal}$ is the current while balancing is in progress.

$$i_{bal} = \frac{\hat{\imath}_{bal}}{100} \tag{79}$$

The board was initially configured for 7 cells, which results in the layout in Figure 46. Module A consists of one chip and cells 1-4, and module B consists of one chip and cells 5-7. The primary side of the transformer is connected across a cell, and the secondary side connects back to the module. The secondary side of module A is connected between ground and the top of the cell stack (i.e. the positive terminal of cell 7), whereas module B is connected between the bottom of module B (i.e. the positive terminal of cell 4/ negative terminal of cell 5) and

the top of the cell stack. This means that module A interacts with all 7 cells, whereas module B only interacts with cells 5-7. This connection is recommended by LTC [88] and avoids requiring a separate module-module balancing system, as discussed in section 2.2.

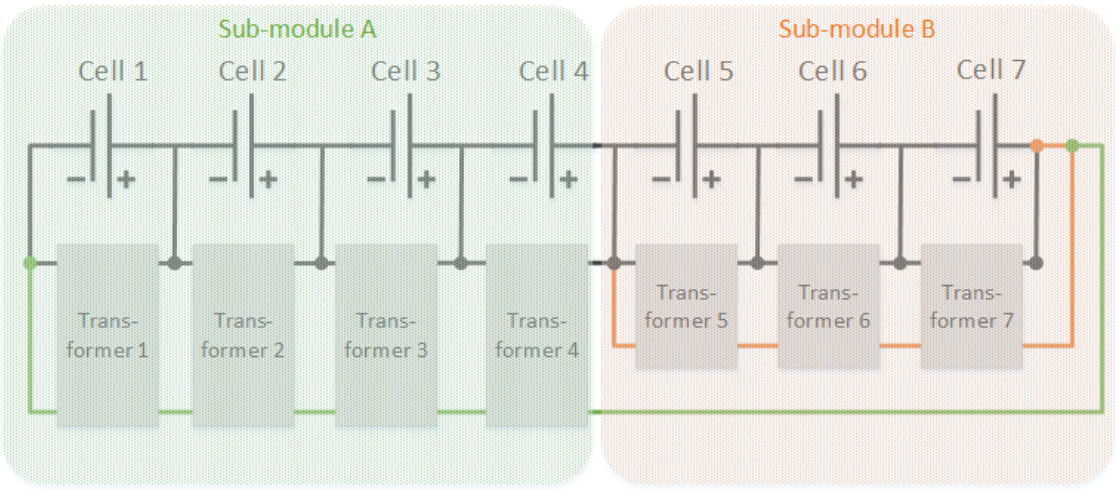

Figure 46: Arrangement of balancing sub-modules

6.2.1 Theory of operation

A simplified circuit diagram of one converter is shown in Figure 47. This replaces the MOSFET with an ideal switch and removes the smoothing filters across the windings and shunt resistor.

When a cell is set to discharge through balancing, the primary winding is first connected over the cell by closing the switch. This creates a first-order exponential rise in current, which can be approximated as a linear increase in current over time as the time constant of the system is much larger than the switching period. The current passes through the shunt resistor connected in series with the primary side windings of the transformer. The voltage over this shunt resistor (R_{sense} in Figure 47) is connected to a comparator in the LCT3300-1 chip: when this voltage is over 50mV (this value is fixed and is the reference for comparator in the chip), the chip opens the primary side switch and closes the secondary side switch. The energy stored in the primary side windings is now released into the secondary side windings. When the voltage across the secondary side shunt drops to 0V the secondary side switch is opened and the primary side switch closed again, and the process is repeated. To charge a cell through balancing the process is similar, but instead the secondary side draws a discharging current from the stack of cells until the voltage over the secondary

shunt resistor reaches 50mV. Then, this energy is released into the primary side windings (and so the cell) until the primary side shunt resistor voltage drops below 0V.

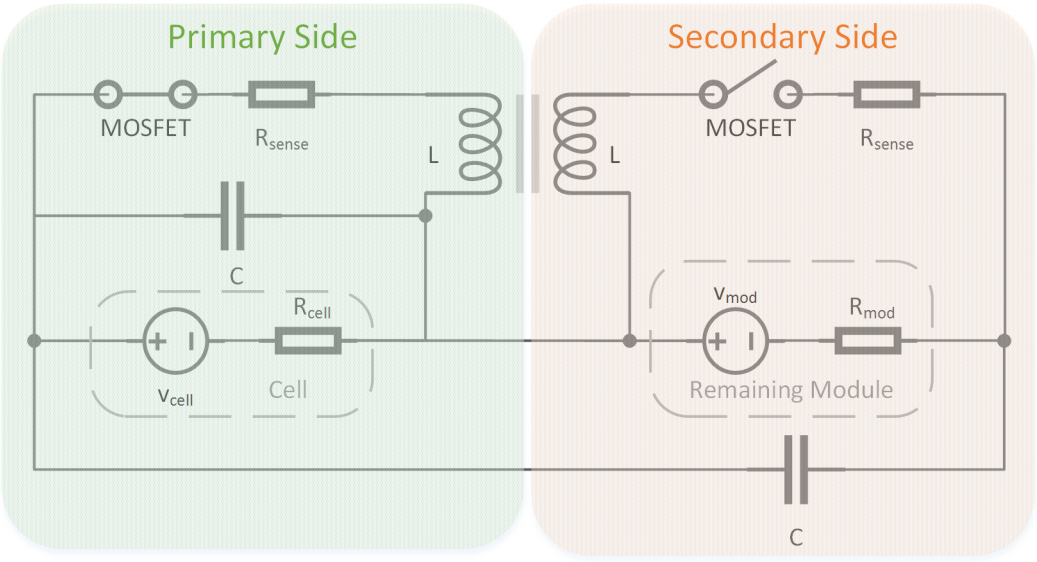

Figure 47: Simplified circuit diagram of one transformer

The switching period employed depends on the values of these circuit components as well as the cell properties, but is typically in the order of microseconds. Over longer periods of time (in the order of milliseconds-seconds), this produces a time-averaged current over the cell and module connections, which is also smoothed by the capacitors (*C*) on either side of the transfer winding. The value of the shunt resistors dictates the amount of charge transferred during balancing: a smaller resistor values means a larger current will be drawn by the windings before the energy is released to the other windings.

In Submission 5, a detailed transformer model was derived from this circuit diagram, and validated against experimental data as being accurate to under 3%. Figure 48 shows an example simulation using this model, for the currents generated by one cell discharging into a module of seven cells. This could represent any of the cells in sub-module A of Figure 46. The balancing current is larger in magnitude than the cell current because the cell-to-module method of balancing means that the cell feeds back into itself. In this case, a steady-state balancing current of -2.5A results in around -2.2A through the cell because of the 0.3A feedback current. The balancing current varies by over 1A peak-to-peak, and the feedback current by 0.8A. The cell current sees less variation – about

0.5A – because of the larger smoothing capacitors across the cell. The total switching period is approximately 5μs, which equates to 200kHz. The current averaged over each switching cycle is given by the dashed lines. The simulation results show what is likely to be measured by sensors sampling at relatively slow BMS frequencies (1-100Hz). The average current reaches a steady state by 80μs. The shortest time the transformer will be active for is 10ms, meaning that the system reaches the steady state in less than 1% of the on-time and so the control system can ignore these dynamics.

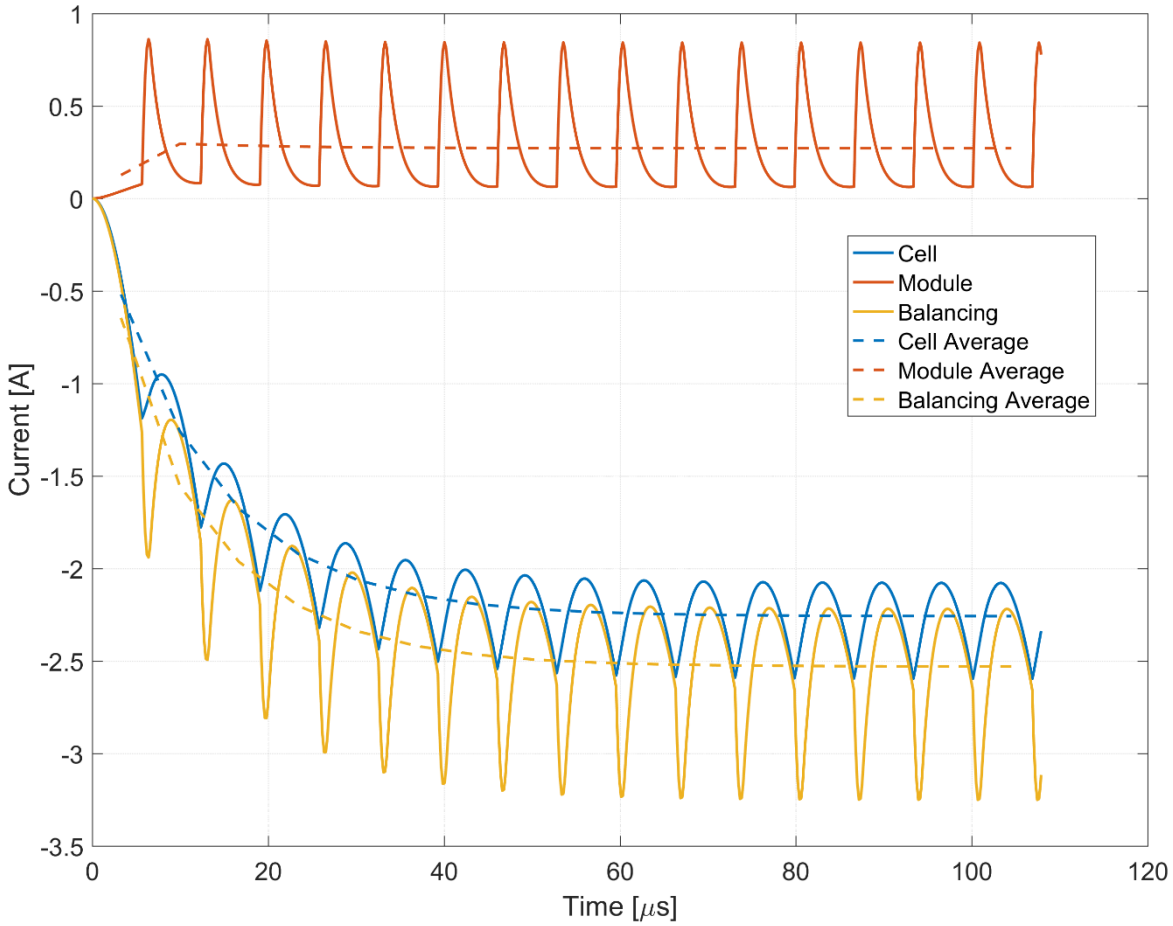

Figure 48: Simulation of the primary and secondary currents arising from converter switching

To assess the consistency of current generation, a parameter sweep of the simulation was performed with the following vectors:

$$\boldsymbol{v_{cell}} = [3.0 \quad 3.2 \quad 3.4 \quad 3.6 \quad 3.8 \quad 4.0]\ v$$

$$\boldsymbol{v_{mod}} = 6 * [3.0 \quad 3.2 \quad 3.4 \quad 3.6 \quad 3.8 \quad 4.0]\ v$$

$$\boldsymbol{R_{cell}} = [5 \quad 15 \quad 25 \quad 35 \quad 45]\ m\Omega$$

The results showed that all parameters have an impact on balancing current generated. A cell internal resistance of 45mΩ generated a 15mA higher

magnitude balancing current than at 5mΩ internal resistance, which is less than 1% relative change. As such, the impact of the cells' resistance changing with SOC and temperature can be ignored. The balancing current can also change as a function of the cells' voltages. However, as Figure 49 shows, this change primarily occurs when the cell which is set to balance has a significantly different voltage to the remaining cells in the module. The black line represents the points where the cell and module voltages are the same, and the current across this line only varies by 17mA. In practice the maximum expected difference between the balancing cell voltage and the average module voltage is around 100mV. As such, the balancing current will be assumed constant with voltage for modelling and control purposes.

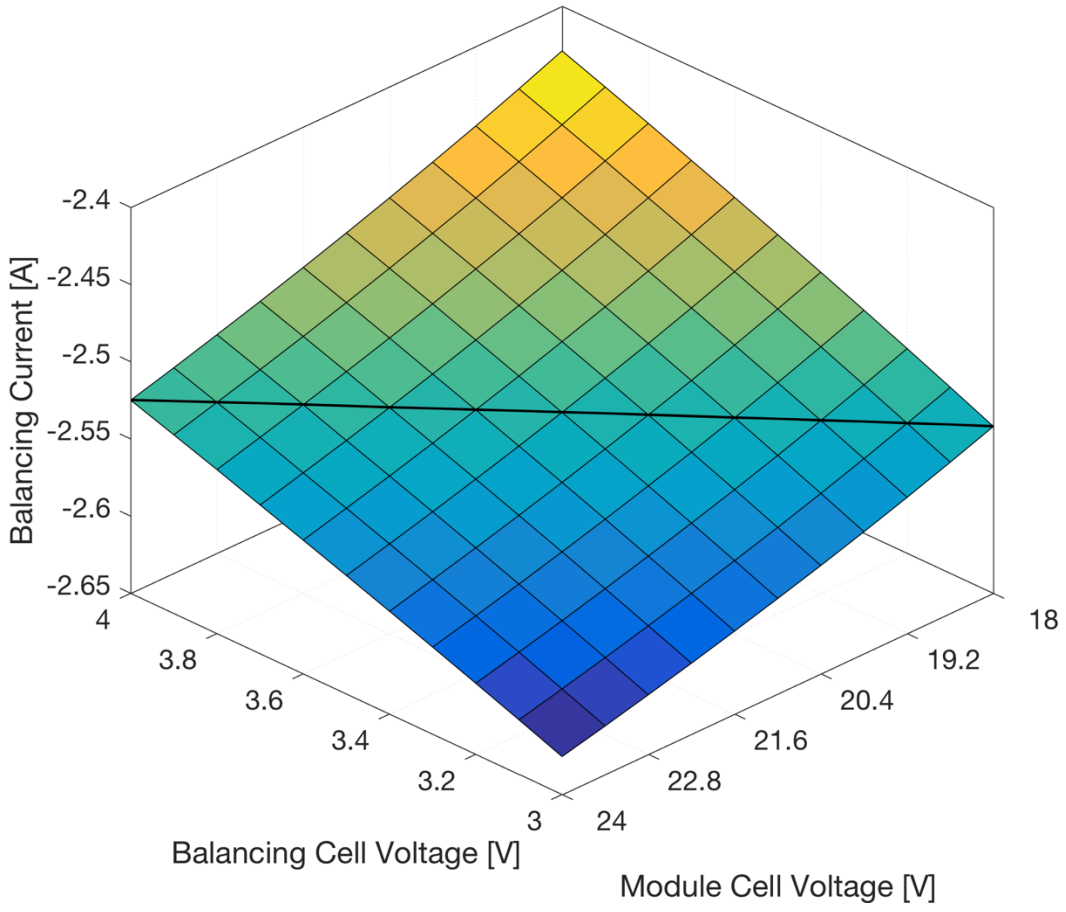

Figure 49: Balancing current as a function of cell and module voltages

The LTC3300-1 datasheet [88] provides the high-level equations for calculating cell and feedback balancing currents. If a cell is set to discharge then (80) and (81) give the balancing and feedback currents respectively. R_{sp} is the current sense resistor on the primary side (8mΩ nominally), N_m refers to the number of

cells on the series side (7 for submodule A and 3 for submodule B), W_r is the ratio of windings between the primary and secondary sides, (in this case, 2) and η_{dch} is the efficiency. This makes the assumption that the current through the inductor increases linearly. In reality, the increase is exponential, but as shown in Submission 5, the time constant of the exponential function is sufficiently large that the slope is approximately linear before the 50mV comparator limit is reached. As such, the average voltage over the on period is 25mV.

$$i_{bal} = \frac{0.025}{R_{sp}} \frac{N_m}{N_m + W_r} \tag{80}$$

$$i_{fb} = \frac{0.025}{R_{sp}} \frac{1}{N_m + W_r} \eta_{dch} \tag{81}$$

Transformer efficiency can be calculated by combining (80) and (81) to produce equation (82). Similarly, equations (83)-(85) apply for when a cell is charged. The parameters N_m and W_r are the same as for charging, but the sense resistor R_{ss} is 15mΩ nominally. Equation (85) is the inverse of (82).

$$\eta_{dch} = N_m \frac{i_{fb}}{i_{bal}} \tag{82}$$

$$i_{bal} = \frac{0.025}{R_{ss}} \frac{N_m W_r}{N_m + W_r} \eta_{chg} \tag{83}$$

$$i_{fb} = \frac{0.025}{R_{ss}} \frac{T}{N_m + W_r} \tag{84}$$

$$\eta_{chg} = \frac{i_{bal}}{N_m i_{fb}} \tag{85}$$

These equations show that if a cell in module A is being charged, the cell balancing current will be much higher than if a cell in module B is being charged. This is because all cells have the same secondary sense resistor and so draw the same current from the stack. However, the energy into the windings will be much larger because there is the voltage of 7 cells being applied, rather than 3. Conversely, if a cell is set to discharge, the feedback current into the module will be higher for module B than module A because the same amount of energy is only distributed among 3 cells rather than 7.

6.2.2 Practical operation

The Balancing PCB requires a 12V power supply, which supplies the balancing chips and microcontroller running the firmware. The voltage measurement board is connected to the cells and powered directly by them. During operation the board drew about 100mA, giving a total power consumption of 1.2W. The high frequency switching is controlled by the LTC300-1 chip, and the lower frequency duty cycling managed by the firmware on the balancing PCB. The firmware requires information as to what these duty cycles should be. This information is provided via CAN messages. The full list of CAN messages sent and received by the balancing PCM is given in Submission 4. As well as the balancing information, other messages include cell voltages and temperatures, and status information about the LTC300-1 chips such as their temperature and errors detected in the transformers or transistors. A balancing command is sent over CAN bus which, as shown in Table 17, specifies a percentage duty cycle and the direction of balancing. There is one message per cell, and all messages are sent together. The firmware, upon receipt of a message, sends a signal over a serial peripheral interface (SPI) bus to the LTC3300-1 chip to enable balancing, and then disable it after the period of time specified by the duty cycle.

Table 17: CAN message to specify balancing for one cell

Signal	Length	Description
Active Balancing Direction	2	Integer to determine whether to charge, discharge, or not balance
Active Balancing Duty Cycle	8	Integer from 0-100 to specify the duty cycle to apply (switch balancing on from 10-1000ms)

As mentioned earlier, the firmware would sometimes move into an error state and prevent balancing. The cause is not known but was mitigated to some extent by controlling four cells rather than all seven. However, this error state would still sometimes occur. The error state means that balancing is disabled and cannot be resumed until the board has been reset, by switching the power supply off and back on. While this is disruptive, if it is a one-off event and if detected quickly it only halts balancing for a short period of time and does not significantly affect the results. However, if the error state resumes then the test must be aborted.

6.3 Instrumentation

Voltage, current and temperature measurements are required for each cell. Temperature was only used for safety monitoring, but voltage and current were also used for further processing within the control system, primarily for SOC estimation. While every cell voltage is commonly measured in a battery pack, often only the applied load current is measured. For passive balancing, the current magnitude is sufficiently low that all cell currents can be considered equal during balancing. Alternatively, the current drawn by passive balancing can be reliably calculated from the cell voltage and balancing resistance by using equation (86). Depending on the system, it may also be possible to accurately estimate the balancing currents and so avoid measuring current on each cell. However, measuring the current through each cell allows for more detailed analysis of the balancing system's operation and performance.

$$i_{bal} = \frac{v_t}{R_{bal}} \tag{86}$$

Voltage and current were sampled at 100Hz. This is at the higher limit of typical BMS sampling rates. For many applications, 1-10Hz would be adequate [75, 118, 160]. However, in this case higher frequency measurements were useful for analysing the performance of the balancing system in operation. The duty cycling means that the cell currents are particularly dynamic, being switched at up to 100Hz (ignoring the very high frequency switching performed by the chip itself). Temperature was sampled at 10Hz as it changes much slower than current and voltage.

6.3.1 Voltage

The cells connected in series have a small wire connecting the positive terminal of the lower cell to the negative cell of the upper cell, as shown in Figure 50 (the balancing system is omitted for clarity). The wire has a finite resistance, which ideally is as low as possible, but for this test setup was in the same magnitude as the cell, as discussed in section 6.3.2.In a production battery pack, careful design, materials selection and advanced welding techniques are used to minimise the resistance of connections between cells [161, 162]. Such design and manufacture was not feasible for a simple prototype such as this. Instead, a wire

with eyelets crimped on each end is bolted down to brass blocks, with one block bolted to each cell terminal.

The cell voltage is the voltage across the cell's terminals. However, many BMSs measure voltage using one wire per cell, which as the BMS voltage in Figure 50 shows, means the connection resistance between cells will also be included. A more accurate cell voltage measurement can be obtained using two wires per cell, which excludes the wire voltage. If the connection resistance is much lower than the cell's internal resistance, then both measurements will be approximately equal and the additional wiring has no benefit. Even if the connection resistance has the same magnitude as the cell, it would not impact SOC and SOH estimation as long as it is factored into the cell model. However, it would still result in the manufacturer specified upper and lower cell voltage limits being reached by the BMS voltage before the cell itself is at those voltages.

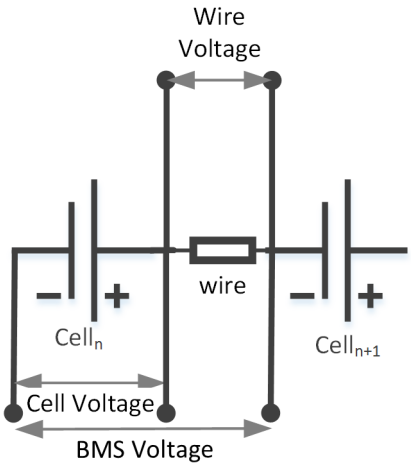

Figure 50: Schematic of voltage measurements within series-connected cells

The balancing system was designed to incorporate a Linear LTC6811 cell monitoring chip to measure cell voltage and temperature. The system was supplied with the voltage sense wires also acting as the electrical connection to the balancing system. This was also the case for the EM140EVM board [145]. In this configuration, each cell shares a common voltage sense wire. While the combined harness is convenient, it causes a noticeable voltage drop during balancing. Rather than using the combined harness for performing balancing and measuring voltage, another harness was assembled and added to the system, solely for voltage measurement. This replaced the original connection, which

used a jumper harness to link the connection wires from the balancing PCB to the voltage measurement PCB. A wiring schematic is shown in Figure 51, where $b_{0\text{-}7}$ are the wires connecting each cell to the balancing system, and $s_{0\text{-}7}$ are the wires connecting to the voltage measurement chip. Previously, $b_{0\text{-}7}$ were linked directly to the voltage measurement chip.

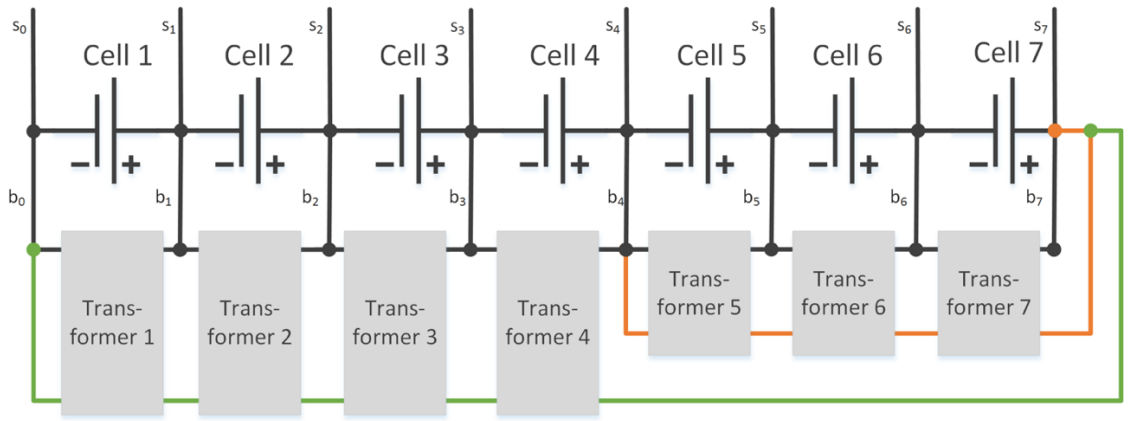

Figure 51: Balancing system wiring

Figure 52 shows an overlay of the voltage and current measurements when using the combined and separate harnesses. The voltages are slightly different because the measurements were taken at different points in time, but they show a similar balancing event with the same current generated.

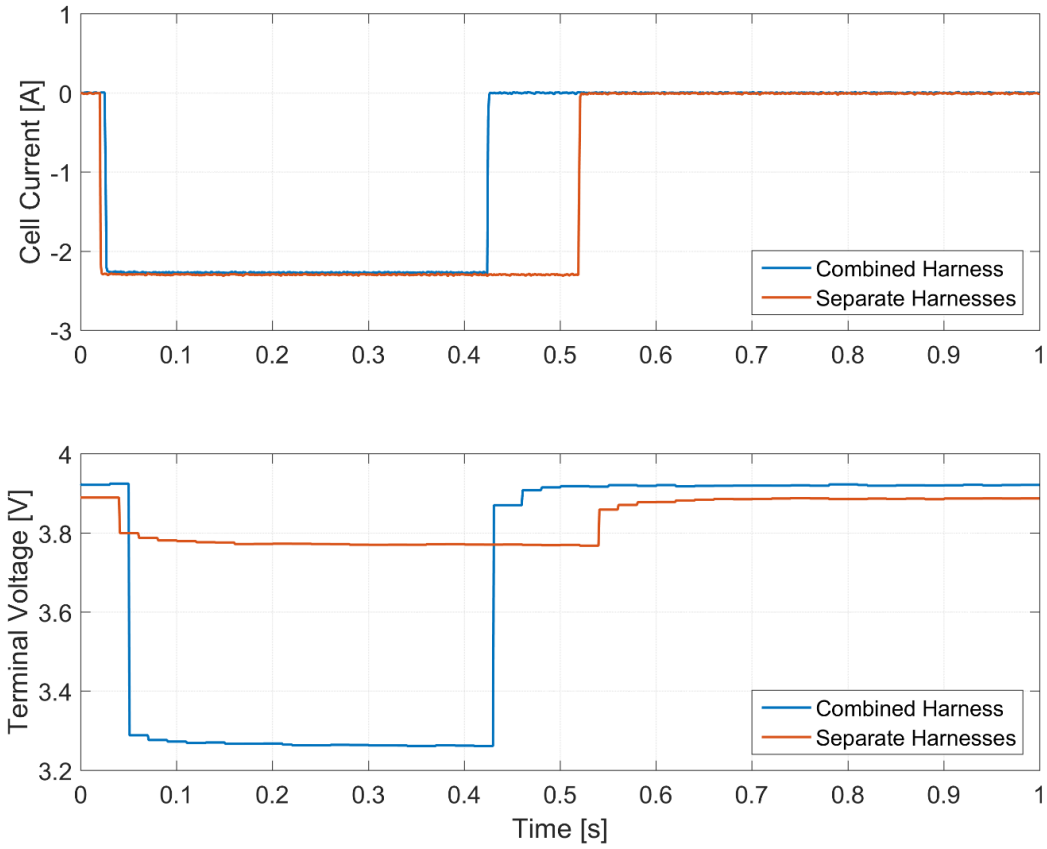

Figure 52: Comparison of Cell 1 voltage and current for the combined and separate harnesses

In both cases cell 1 was set to discharge, for a duration of 400ms for the combined harness and 500ms for the separate harnesses. The voltage plot shows that the combined harness reads an additional 0.5V drop in voltage compared to the separate harnesses. A scope meter connected directly over the cell terminals confirmed that the separate harness was providing an accurate voltage reading. The specific cause of this voltage drop is not known, and is difficult to evaluate as it requires access to specific points on the PCB. It is unlikely that this is caused solely by the resistance of the balancing wires as this should not exceed a few milliohms. There is a 5A surface-mount fuse on each cell connection to the board. This has a resistance which contributes to the voltage drop, but its resistance could not be accurately measured to determine if this is the primary cause. There may also be some resistance in the connections from the cells to the PCB. These large voltage drops could regularly imply that a cell had gone outside of its safe operating voltage window when the cell itself was within the window. This also reduces the performance of an SOC/ SOH estimator which uses voltage feedback.

Another issue caused by the combined harness is apparent changes in voltage when a neighbouring cell is set to balance. This is because the wires are positioned between cells. For example, there is a wire between the positive terminal of cell 4 and the negative terminal of cell 5. The upper plot of Figure 53 confirms that when cell 5 is set to balance, there is no current through cell 4. Despite this the shared wire creates an apparent change in voltage for cell 4, as the lower plot shows. Using separate voltages removed the balancing system voltage drop and, so significantly reduces this neighbouring cell interaction, although a voltage drop from the cell connections remained. From this analysis it is clear that using a combined harness is not suitable for monitoring cell voltages: the reduction in wiring does not compensate for the significant reduction in accuracy that occurs.

There is also some noise on the voltage measurements of the LTC6811. This was analysed by logging 60 seconds of cell voltages with no current excitation and analysing the resulting data. Histograms of these voltages are shown in Figure 54. Note that a logarithmic scale is used for the y axes so that the low

counts of outlier voltages can be seen, notably for cell 7. There is some additional noise on cell 7 in the form of occasional large spikes in the measured voltage. The reason for this is not known: it appears to only be present when the balancing system is connected, but occurs even when no balancing is taking place. There is no obvious pattern to the noise on each cell.

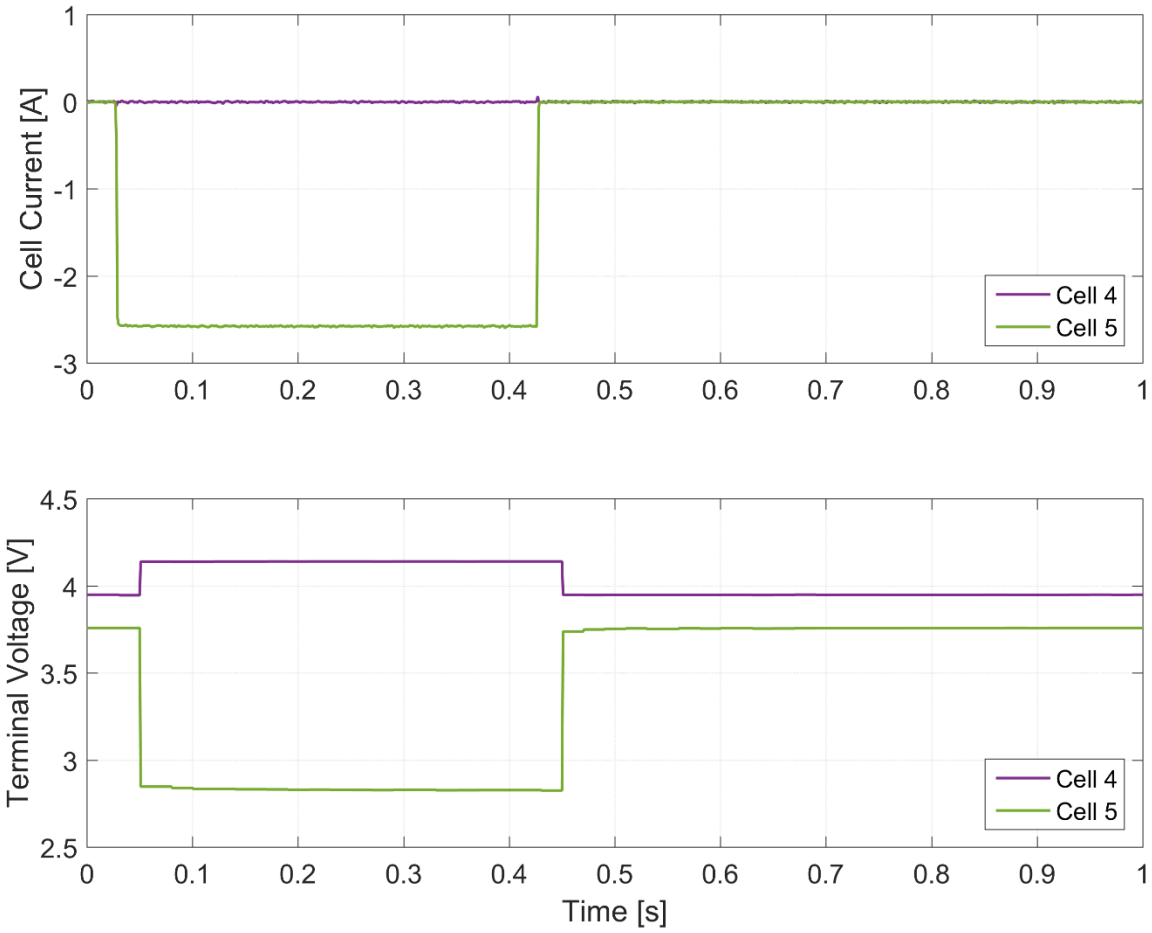

Figure 53: Apparent change in a cell voltage caused by a neighbouring cell set to balance

Private communication with a TATA Motors Senior Engineer suggested that this could be an artefact of the digital analogue converter being 16 bits in resolution. The LTC6811 has different conversion rates, and at higher rates there is a corresponding reduction in resolution. The full 16-bit voltage is still outputted, the final few bits of the voltage may just be noise from the conversion process and so redundant. Rounding the measured voltage to the nearest mV is more appropriate for this sample rate. This means that the voltage is less accurate than initially thought, and the unusual noise pattern means that averaging voltage over time does not necessarily result in the 'true' voltage, as it would with a Gaussian distribution and could therefore hinder state estimation.

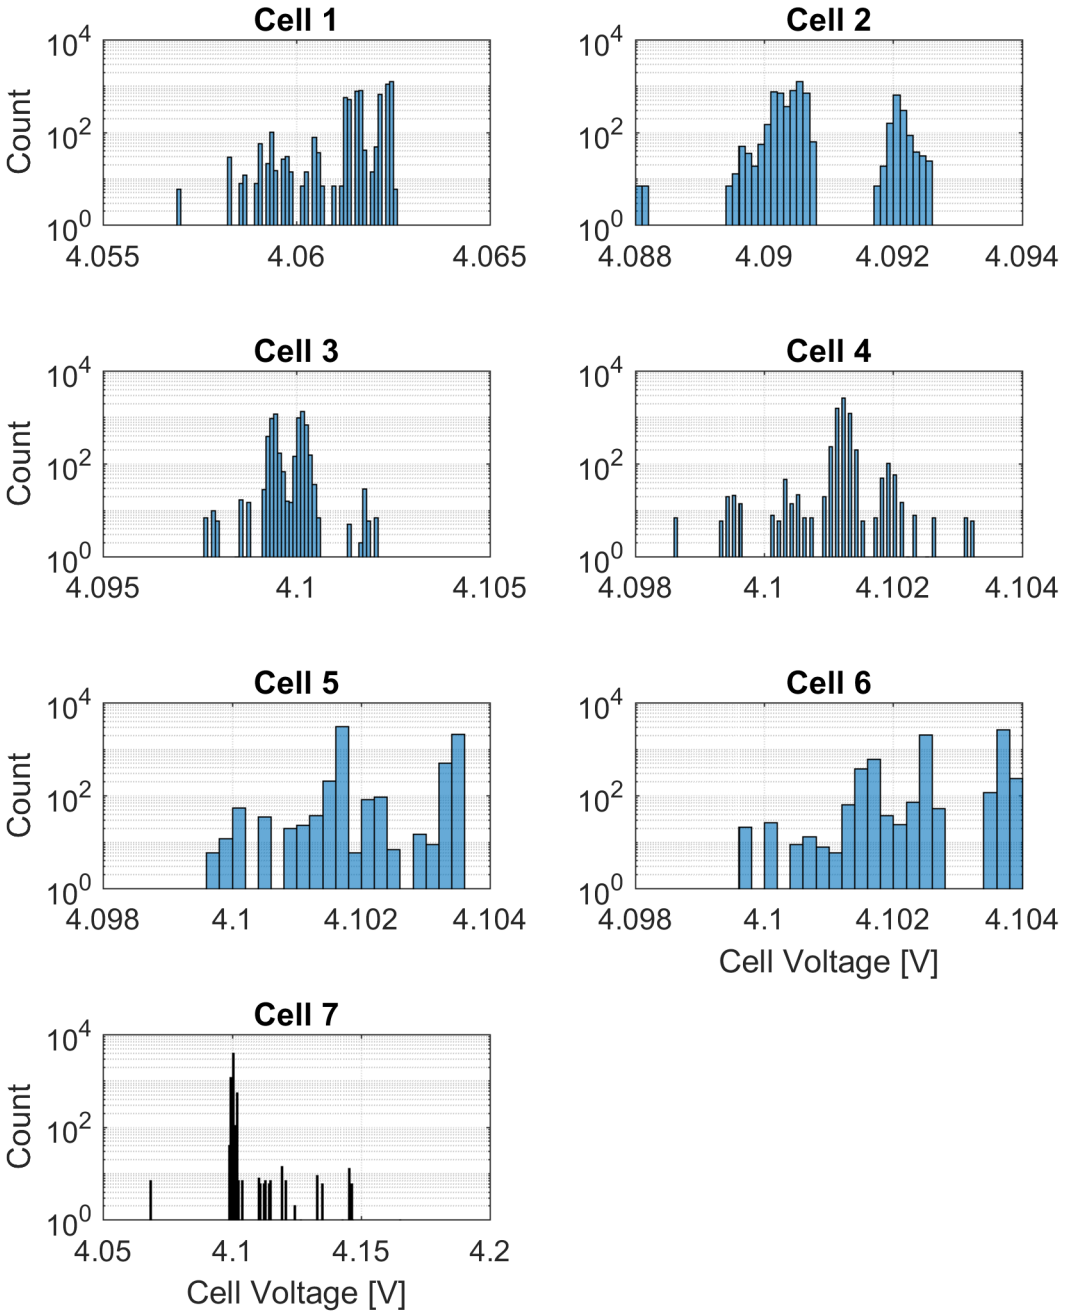

Figure 54: Histograms of cell voltage measurement noise for the LTC6811 chip

Through experimentation, it was found that there was some time-delay in the region of a few milliseconds on the voltage measurement from the LTC6811 because it is read, sent over SPI to the balancing PCB, then packaged into CAN messages where it is received by the control system. The current measurements discussed in 6.3.2 were generally more direct and so had fewer delays, making synchronising the voltage and current challenging. This has important implications for SOC estimation. A significant portion of the voltage comes from direct feedthrough from the current. The schematic in Figure 55 shows how a

delayed voltage during a dynamic current profile might appear as though the voltage increases despite the current decreasing (the cell going from charging to discharging). This suggests to a parameter estimator that the cell internal resistance is negative which is clearly impossible. If the time-delay is consistent and precisely known, the current measurement can be delayed a fixed amount to align with the voltage. However this was not always reliable, and ideally the voltage and current would be sampled from the same source. The issue is not so much the communication delay, but the delays between when the two measurements are taken. The voltage measurements are packed into several CAN messages (three voltages per cell, with 12 total voltage measurements) which each had slightly different time-stamps, which made post-processing the data and diagnosing system behaviour more difficult.

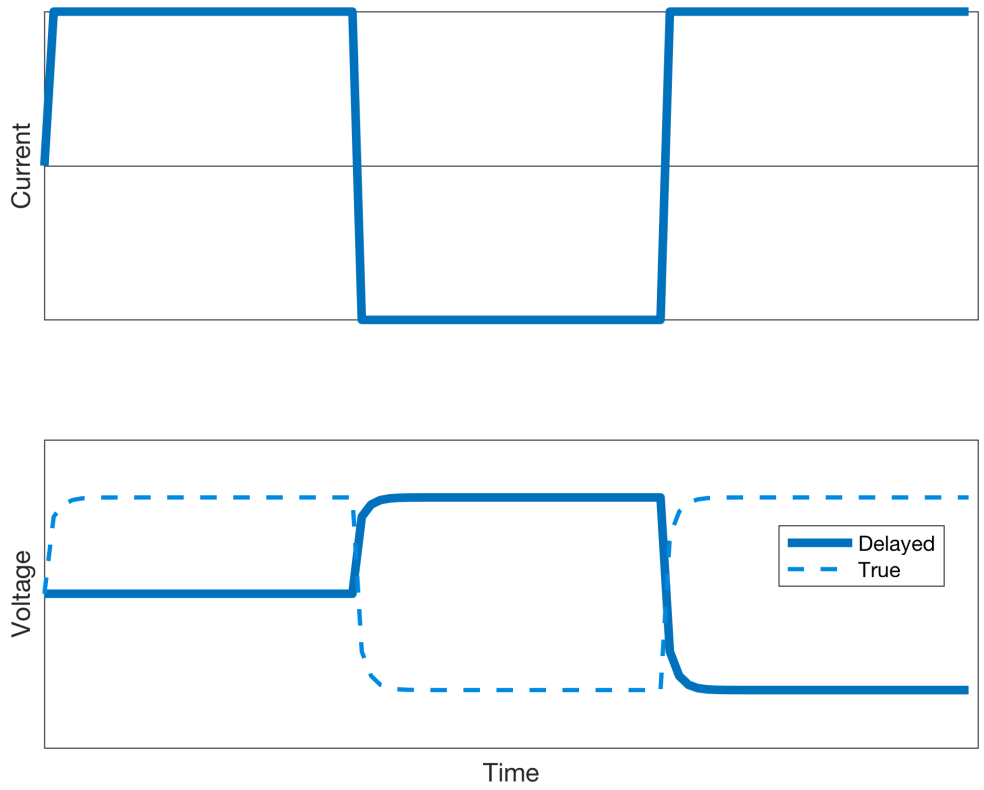

Figure 55: Schematic of delayed voltage measurements during a dynamic current profile

In summary, using the LTC6811 had the following challenges when used as part of the control system:

- The measured voltage includes the connection resistance, which in this case is not trivial, so the measured voltage is not equal to the cell voltage.
- There is some evidence of non-Gaussian noise which can hinder accuracy and SOC estimation.

- It is difficult to synchronise the voltage and current measurements.

An alternative means of measuring voltage was used instead: A National Instruments 9219 analogue-digital converter (ADC) card connected to a National Instruments cDAQ 9174 USB data acquisition chassis. Each card has four channels, a 100Hz sample rate and 24 bit resolution, which equates to 1.4µV [163] when using the ±15V voltage scale setting. It uses two wires per ADC channel, meaning that the cell voltage in Figure 50 can be measured, rather than the BMS voltage. Four cell voltages were measured in total, as by this point in time it was decided that only four of the seven cells would be controlled in order to maintain stable operation as noted in section 6.1.

Analysis of the voltage measurements shows that it has superior noise characteristics compared to the LTC6811. Figure 56 shows histograms for the cell voltage measurements using NI 9219 modules.

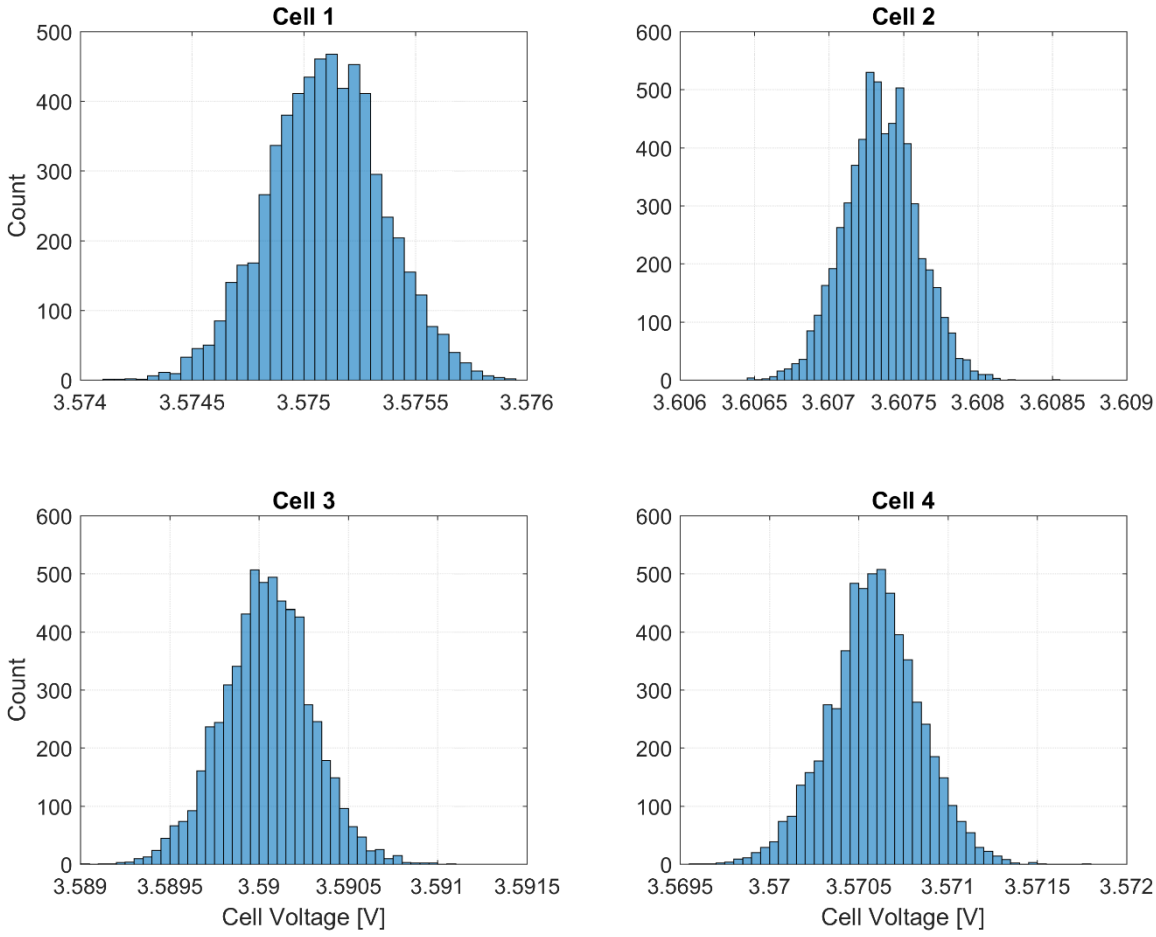

Figure 56: Histograms of cell voltage measurement noise for the NI 9219 module

A t-test confirms that the distribution is Gaussian with 99.99% certainty. The standard deviation of voltage is approximately 0.25mV per cell, meaning that a

more accurate voltage and so SOC estimation can be achieved compared to the LTC6811 solution. The same ADC was ultimately used to measure cell current, as described in section 6.3.2, and since all channels on the 9129 card can sample voltage simultaneously, the voltage and current measurements are automatically time-synchronised.

6.3.2 Current

LEM CASR-6 current sensors [164] were initially used to measure each cell current, and their integration and calibration are documented in Submission 4. However, the sensors introduced a large resistance in series with each cell, contributed by the copper tracks on the PCB the sensors were mounted on, the connection from the board to the wires, and the connnection between the wires and the cells. The total resistance added was around 35mΩ. This is an order of magnitude larger than the ~3mΩ of the automotive pouch cells used for the testing and as such interfered with system performance. While another board could have been designed which used non-intrusive hall-effect sensors, these still require careful wiring and a stable power supply, filtering, and there was still the question of synchronising with voltage measurements. The 9129 ADC card has an option of scaling the input voltage to ±125mV, which equates to a resolution of 15nV and worst-case accuracy of 271μV [163]. The total connection wire resistance between cells was measured to be about 5mΩ, meaning that it is possible to directly mesure the wire voltage in Figure 50. This is simpler to implement than the hall-effect sensors, although is reliant on accurate calibration.

The wires were calibrated by passing several known currents through the cells. Values of -15 to 15A were applied at 3A intervals to the series string using the load in section 6.4, and measured using a Fluke i30s current clamp. The wire voltages were sampled at 100Hz and averaged across each current value. A linear least squares fit was performed on these values to obtain the gain and offset required to translate wire voltage into cell current using (87). In the ideal case the offset O would be zero and the gain G the inverse of wire resistance.

$$i_{cell} = Gv_{wire} + O \tag{87}$$

The calibration results for the four wires tested are given by Table 18 and Figure 57. The gain is the inverse of the wire's resistance, suggesting that wires 1 and

4 have a slightly poorer connection than wires 2-3. Theoretically the offset should be zero (as there is no voltage over a resistor if there is no current), but in reality there are small errors in measurement of both current and voltage which result in a non-zero offset.

Table 18: Wire voltage calibration results

Sensor	Offset (O)	Gain (G)	Resistance (mΩ)
1	0.02015	141.47	7.07
2	0.01875	175.47	5.70
3	0.01935	155.22	6.44
4	0.01969	176.06	5.68

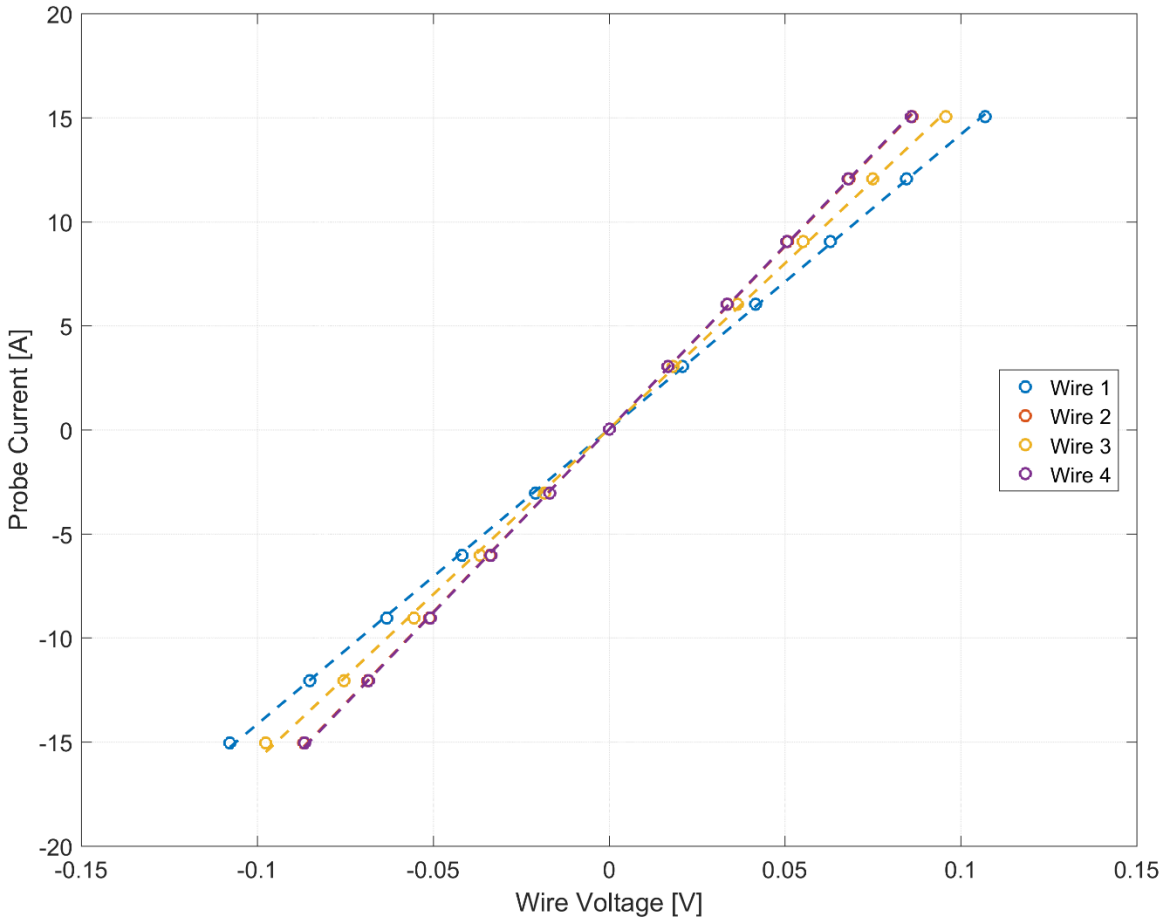

Figure 57: Wire voltage calibration results

6.4 Module cycler

A lack availability of laboratory equipment meant that it was not possible to use a module cycler to charge and discharge the series string of cells. However, an important consideration for the active balancing control system is how well it performs with the cells under a current load. A Kepco BOP-MG 1KW bipolar load

was used to apply the current. This can deliver/ accept up to 20A (at up to 50V), which is 1.3C for the automotive pouch cells and is representative of the current magnitudes expected during operation in a BEV [41, 165].

The load can be controlled externally using analogue voltage signals. These were generated using a National Instruments NI9263 module, which connected to the same National Instruments cDAQ chassis. The load current and voltage can be set dynamically using voltage signals, along with limits for current and voltage. An analogue signal of ±10V is used to set these values: a 10V control input corresponds to a 50V or 20A load output respectively. To achieve pseudo-current control, the load was set to voltage control mode, and the current limits were continually updated. If a charging current was required, the load voltage was set as 29.4V, or 18.9V for a discharging current. These values are based on there being 7 cells with a maximum and minimum voltage of 4.2V and 2.7V respectively. Unless the cells are at a very high or low SOC, it would require several hundred amps to force the module voltage to these values, which guarantees that the load will limit the current to the current limit set by the external signal. By continually setting the current limit to the current value which is actually desired as the output, a form of current control is achieved. For a given drive cycle current *d*, the control signal sent to the Kepco load to set the voltage is given by (88). *N* is the number of cells in series, and v_{min} and v_{max} are the minimum and maximum voltages defined by the cell manufacturer. The signal to set the load current is given by equation (89). The scaling factors of 0.2 and 0.5 come from 10V signals equating to the load being set to 50V and 20A respectively.

$$v_{lv} = \begin{cases} 0.2Nv_{min}, & d < 0 \\ 0.2Nv_{max}, & d \geq 0 \end{cases} \tag{88}$$

$$v_{lc} = 0.5d \tag{89}$$

The bipolar load must be initialised carefully to ensure safe operation. When switched off, it essentially acts as a resistive load which discharges the cells at a potentially damaging rate. When switched on, it defaults to voltage control mode. If the value at which it is switched on at is even 100mV different to the total cell string voltage, it will cause a large current to flow which is undesirable. To ensure safe operation a cut-off switch was connected in series with the wire connecting

to the negative cell terminal, and the following initialisation procedure followed (assuming the switch is initially open and so the cells are disconnected):

1. The bipolar load is switched on and the following settings are applied:
    a. The load is set to voltage control mode.
    b. The maximum and minimum voltage limits configured to the values in (88), and the maximum current magnitude is set to 20A.
    c. The load is set to respond to the analogue voltage control signals.
2. A signal is sent by the cDAQ chassis to set the load current to zero.
3. The cut-off switch is closed, and the bipolar load's electronic display checked to ensure that the current is in fact zero.

Similarly, to switch off the bipolar load after a test, a signal is sent to set the current to zero, and then the cut-off switch opened to ensure the cells do not discharge through the load once it is switched off.

6.5 Control bench

This section describes how the different data sources are managed and tests executed. The development of the software platform to execute the control software is discussed and the various elements of the software are detailed.

A schematic of the configuration used for experimental work is shown in Figure 58[13]. The cDAQ chassis is used to log the cell voltages and currents of the four cells being controlled, and output a signal to the bipolar load to generate the drive cycle. The CAN-USB adapter is used to read the voltage and temperature of all seven cells (used for safety purposes only), and send balancing requests. These both connect via USB to an HP EliteBook 850 laptop running Windows 10 and MATLAB R2016b. A software class was written to manage the experimental test set-up. A class was chosen as it allows for logical structuring of data, and makes it easy to control access (for example, making some data read-only to avoid accidentally overwriting or removing it). An object of the class is instantiated in the workspace, and the details of how the class is used are given in Appendix C.

[13] This set-up went through various iterations, notably with a Speedgoat real-time computer and Simulink. This is documented in Submission 5.

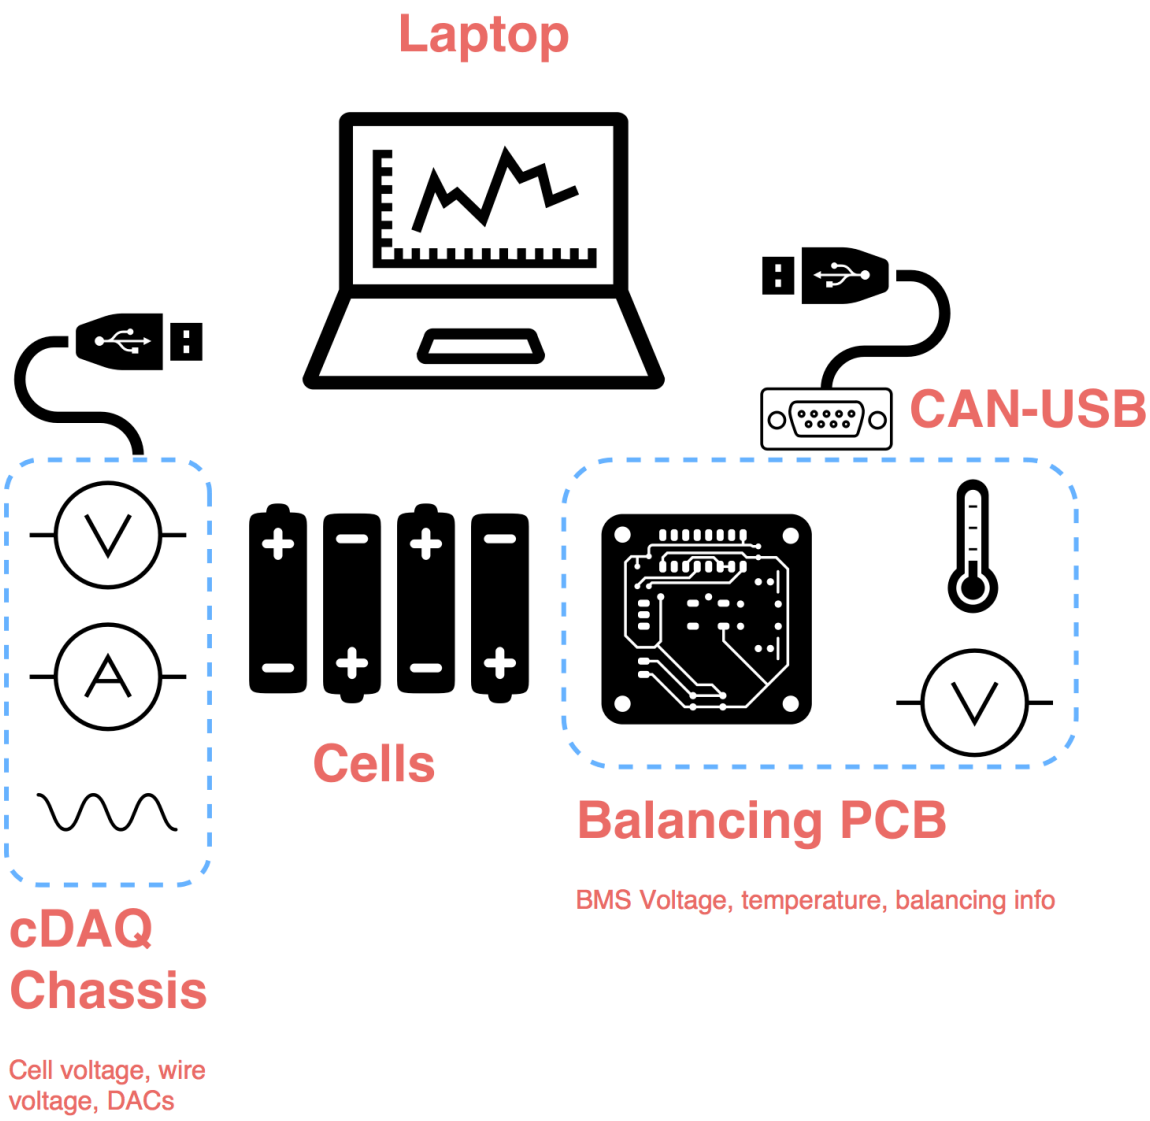

Figure 58: Schematic of final experimental configuration

The entire control system was run within MATLAB. This had the disadvantage of not strictly operating in real-time. Functions were set to trigger when one second's worth of data had been acquired, as described further below. In practice, these were triggered within 0.1% of this desired time, and this pseudo-real-time operation was not detrimental to operation of the closed loop system. This meant that the 100Hz measurement data was processed in batches every 1s, which reduced the computational overhead, especially when iterating through the state estimators. The majority of changes to the code did not require re-compilation, which sped up the development process, and meant that minor changes could be made while the code was executing (such as changing what was being plotted or printing variables to debug an issue). The SOC estimators were compiled to C-code as these were found to take up most of the execution time. After compilation, the execution was over 100 times quicker. The whole function

typically executed in around 100ms, comfortably less than the 1s window. The remaining components such as the balancing controller did not see the same performance gains (typically around a factor of 2 improvement) when compiled, and given the function as a whole was not close to requiring the full 1s to execute, compilation was not performed as it slows the development process.

The various estimators and controllers used throughout the development and test work (described in chapters 4 and 5 respectively) were written as System objects. The many matrix operations and decision logic required for these were easier to write as MATLAB code rather than using Simulink blocks. System objects can be implemented in both MATLAB and Simulink, which meant that no additional work was required between developing the controllers in simulation, and running them during tests. They also support code generation, which made the process of compiling the estimators straightforward.

The complete test set up in the laboratory is shown in Figure 59. The Kepco load is on the left-hand side, and the seven cells to its right. The red and black sense wires connect to the terminals so that the cell and wire voltages are measured by the cDAQ chassis in the foreground. The white wires connect the cells to the balancing PCB behind the cDAQ, in front of the power supply for the board. The CAN-USB adapter is on top of the power supply, and this along with the cDAQ chassis connect to the laptop on the right via USB.

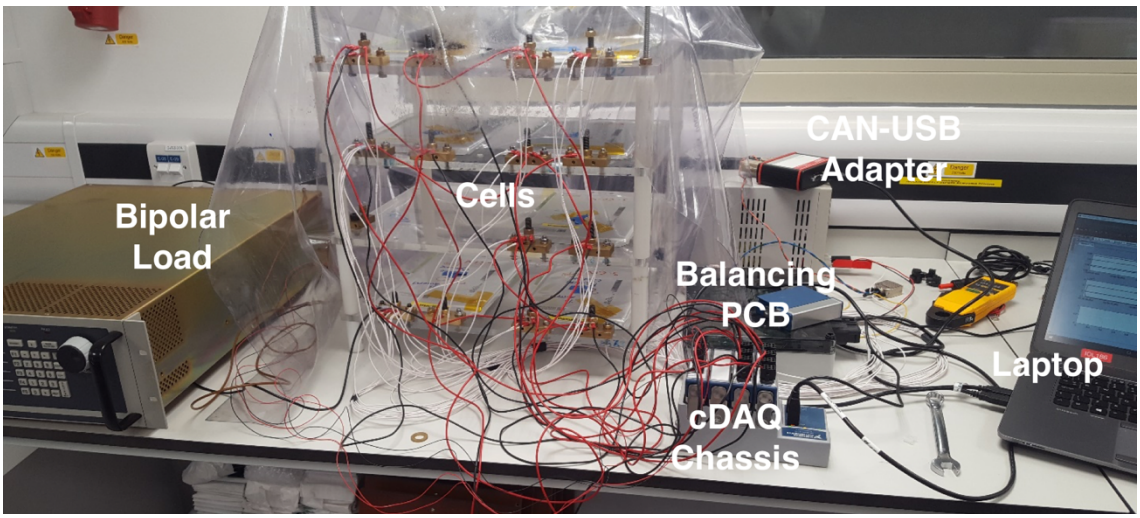

Figure 59: Equipment set-up in the laboratory

6.6 Hardware Analysis

Before using the balancing system to develop and evaluate the control systems, its performance must be assessed. The focus of the analysis was on current generation. The control system from Chapter 5 calculates a desired balancing current for each cell. This needs to be generated by the balancing hardware, but as noted in section 6.2.2, a duty cycle value is sent to the balancing board. This is a relative value, and so the current generated by the circuitry when active needs to be known so that (79) can be applied reliably. If i_{bal} is incorrect, there will be a set-point error where the current generated by the hardware is different that why is intended from the duty cycle value. There could be some uncertainty in i_{bal} owing to variances in the electronic components, the LTC3300-1 operation, or the circuitry and operation of neighbouring cells. Tests were therefore performed to answer the following questions:

- How consistent is i_{bal} on repeated operation?
- How much variation is there between cells?
- How does i_{bal} change when other cells are also set to balance?

As well as the balancing current itself, there is also the feedback current applied to the submodules in Figure 46. These also need to be evaluated to determine the relationship between balancing and cell currents, and also used to calculate the efficiency of the system based on (82) and (85). Much of the test work was undertaken on four cells because of the issues raised earlier. In this case, four cells are labelled as cells 1-4 for convenience but are actually cells 2,3,5 and 6 respectively in Figure 46. These four cells were used for all of the experimental work here unless otherwise stated. For the test work undertaken with all seven cells, the measurements were obtained using the hall-effect current sensors discussed in section 6.3.2. As noted, the current sensors were primarily replaced because they interfered with the cell voltage and state estimation, not because of the quality of the readings.

A simple test was performed on all 7 cells, where one cell was set to balance for 0.5s, every 1s. An example plot is shown in Figure 60, which shows that the balancing currents are consistent.

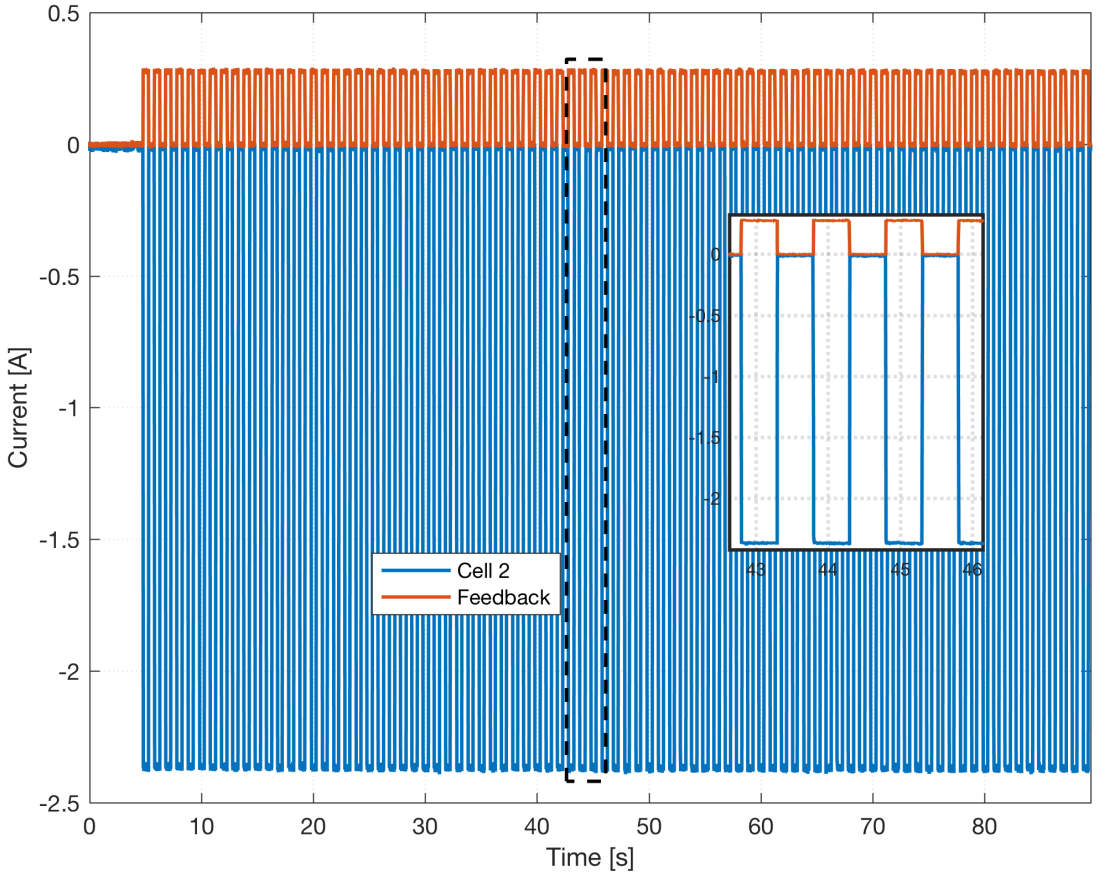

Figure 60: Repeated single-cell balancing

To test how different each cell is from one another, one cell at a time was set to discharge, and then charge as shown in Figure 61 (Note these were the cell load currents, not the balancing currents). According to Kirchoff's current law, the balancing current i_{bal}, cell current i_{cell}, feedback current i_{fb} and external load current i_d are related by (90). Equation (90) can be rearranged to (91) so that the balancing current can be calculated. In general, the feedback currents from balancing are not measured, but (90) was validated with no external load by using Fluke i30s current probes. When a cell is not set to balance and there is no external load, its cell current measurement is a direct reading of the total feedback current. If only one cell is set to balance within a sub-module, there will only be one feedback current which is now known.

$$i_{cell} = i_{bal} - i_{fb} + i_d \tag{90}$$

$$i_{bal} = i_{cell} - i_{fb} \tag{91}$$

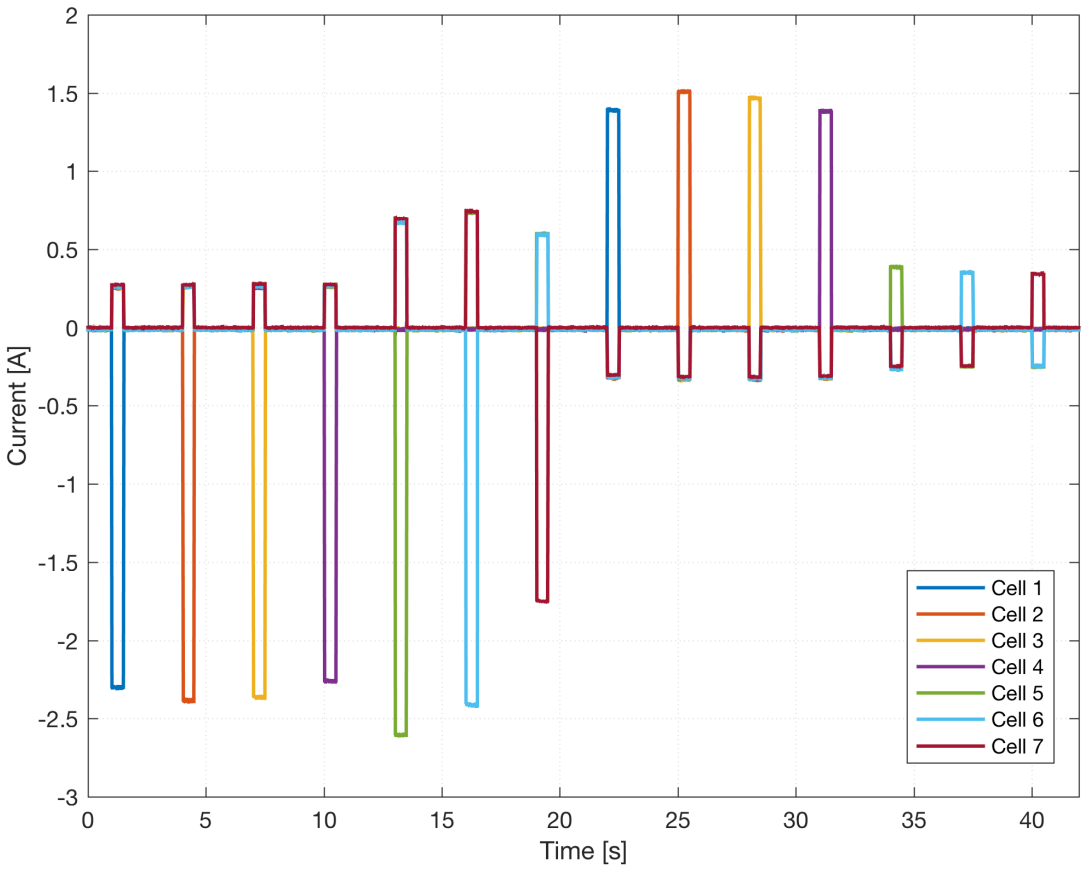

Figure 61: Single cell balancing experiment

For each balancing event, the balancing and feedback currents are shown in Table 19. The corresponding efficiency is also shown, calculated using (82) and (85) depending on the balancing direction. It is clear that the discharging currents are much larger than the charging currents. This is because of the choice of a larger secondary side sense resistor (15mΩ as opposed to 8mΩ) and a 1:2 transformer ratio. These design decisions were made by the consultancy, although there are limits defined by LTC. For example, the 1:2 ratio is deemed optimal for the LTC3300-1s chips, but other ratios can be used [88]. For charging, the balancing current is dictated by the secondary side winding. Each cell in submodule A draws about 0.3A, with a corresponding cell-side balancing current of 1.4-1.5A. At around 0.25A, the sub-module B feedback current is lower, which combined with there being only 3 cells to source energy from makes the balancing current around 0.6A. The reason for the lower feedback currents for module B may be because performance is a function of the number of cells (6-12 cells is considered an ideal number by LTC [88]). For discharging, the balancing current should be independent of the connections, with the feedback current changing instead. The feedback currents are higher for sub-module B,

which is expected because the energy generated from the balancing cell is only redistributed between three cells rather than seven. There is significant variation between the discharge balancing currents. This is likely to be because of the tolerance on the 8mΩ resistor which determines the balancing current. It demonstrates that each channel of the balancing hardware must be evaluated and calibrated: there will be large set-point errors during operating if the nominal values are assumed correct. The efficiency is higher during charge and discharge, at 82% and 73% respectively. There should be little difference between the two according to the datasheet [88]. The lower magnitude of the charging currents may mean there are fewer resistive losses, which are proportional to the square of current. This includes the large resistance responsible for the voltage drop observed in Figure 52. Both of these values are lower than the 90%+ efficiency claimed by LTC [88]. Further investigation is required to evaluate this, as there are many factors which could affect efficiency, including the type of cell used, design decisions and component selection.

Table 19: Current and efficiency for single-cell balancing

	Cell	Cell Current	Balancing Current (A)	Feedback Current (A)	Efficiency (%)
Charging	1	1.39	1.69	-0.30	80.1
	2	1.51	1.82	-0.31	83.0
	3	1.47	1.79	-0.31	81.0
	4	1.39	1.69	-0.31	78.3
	5	0.39	0.63	-0.25	86.2
	6	0.35	0.59	-0.23	83.6
	7	0.34	0.57	-0.23	82.2
Discharging	1	-2.30	-2.57	0.27	74.3
	2	-2.38	-2.66	0.28	72.4
	3	-2.36	-2.64	0.28	73.8
	4	-2.26	-2.54	0.28	76.2
	5	-2.60	-3.30	0.70	63.5
	6	-2.41	-3.16	0.75	71.2
	7	-1.75	-2.36	0.61	77.2

To evaluate the currents through the four instrumented cells in more depth, all possible permutations for switching of balancing were applied. The total number of permutations is given by (92), where the 3 arises from there being three possible states: charging, discharging or not balancing. The duty cycle of each

cell was 100% when enabled, and each permutation was applied for 1 second, followed by a 2 second rest and the next permutation.

$$p = 3^{N_{cells}} \tag{92}$$

The four cell currents during the 81 balancing events are shown in Figure 62. This test has been performed more than once and found to be very closely repeatable.

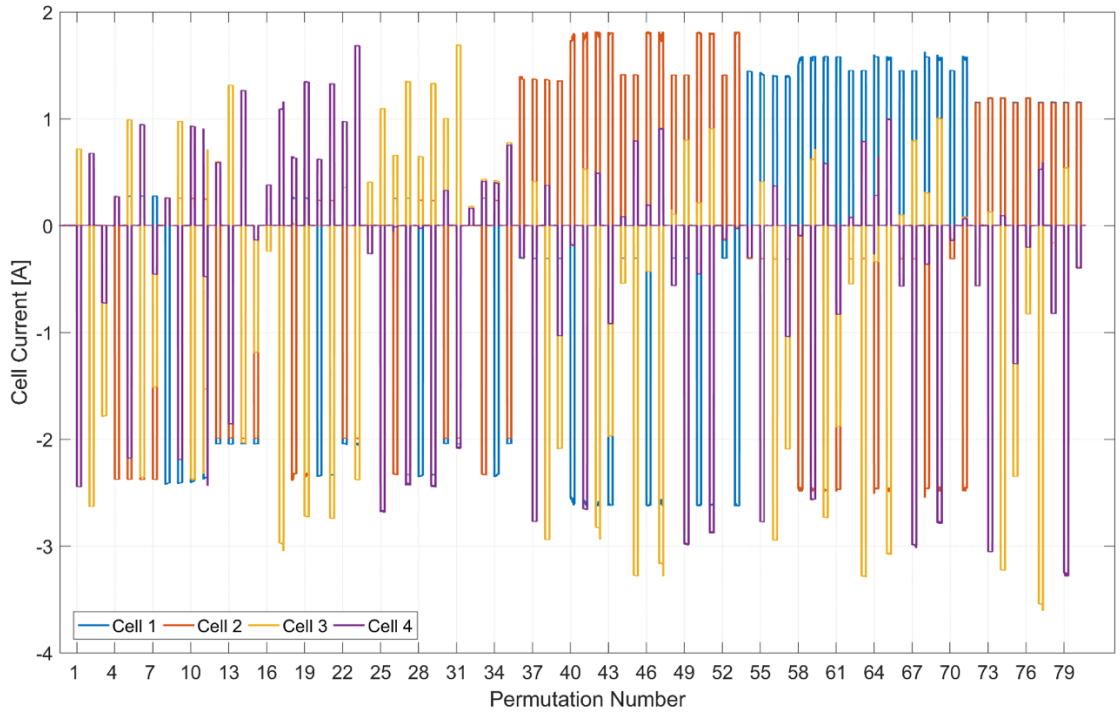

Figure 62: Cell currents measured during the balancing permutation test

There are several observations that can be made based on this dataset. The current magnitudes of cells 1 and 2 appear to be less stable when one is charging while the other is discharging, see for example permutations 40-43 on Figure 62. Cells 3 and 4 do not seem to have this problem and it is unclear what the cause is, but the variation in current is not sufficient to substantially affect performance. The balancing current of each cell is not fixed: it varies depending on the status of the other cells. This is not accounted for in equations (80)-(84), in which the balancing current only depends on the number of cells, transformer windings and sense resistance. The repeatability of the results suggests a specific interaction between transformers rather than some random fluctuations or unpredictability to the current. This has been confirmed by using a Fluke i30s current probe to directly measure the current through the balancing wires during operation.

To demonstrate this, consider the data in Table 20. The first row of data is for when only cell 3 is set to discharge, the second row for only cell 4 discharging, and the third row when both are set to discharge simultaneously. When one cell is set to discharge, the other cell current is the balancing feedback current. Observing the cell current measurements, when one cell is balancing the cell 4 current is only 0.2A lower than for cell 3, and the feedback currents are very similar in both cases. However, when they are both set to discharge, the current through cell 4 is over 1A lower than cell 3. This strongly suggests that the balancing current generated by cell 4 is much lower when cell 3 is also balancing.

Table 20: Discharge balancing results for cells 3 and 4

Cell set to discharge	Cell Current	
	cell 3	*cell 4*
3	-2.63	0.68
4	0.72	-2.44
3 and 4	-1.78	-0.72

Using (91) this means balancing currents of -3.31A for cell 3, and -3.16A for cell 4[14]. If the balancing systems were consistent and independent of one another, the cell currents when both cells are discharging should be given by (93), which is obtained by rearranging equation (91) and accounting for there now being two components to the feedback current. Using the values above, this would result in the cell currents in row 3 of Table 20 being -1.91A and -1.76A respectively. The measured value for cell 3 is relatively close at 0.13A difference, but cell 4 is 1.04A from the theoretical value. This variation in loading means that calculating a duty cycle for each cell becomes complex, as the same duty cycle value could result in a different current depending on the status of the other cells.

$$\begin{bmatrix} i_{cell_3} \\ i_{cell_4} \end{bmatrix} = \begin{bmatrix} i_{bal_3} \\ i_{bal_4} \end{bmatrix} + \begin{bmatrix} i_{fb_3} + i_{fb_4} \\ i_{fb_3} + i_{fb_4} \end{bmatrix} \tag{93}$$

The next test was performed to assess operation under load. The permutation cycle from Figure 62 was run, then repeated with a 7.5A (C/2 based on the

[14] Note that these values are almost identical to the results from Table 19 (cells 5 and 6), despite the two tests being performed months apart and with different methods of current sensing. This emphasises both the repeatable operation of the balancing system and the robustness of both current sensing methods and the associated calibration procedure.

nominal cell capacity) current applied. The current trace from this experiment is shown in Figure 63. The currents from each permutation were compared, and are summarised in Table 21. There is only a minor change in current values, which could in part be due to noise or inaccuracy in the current sensors. This means the balancing controller does not have to account for an external load in the duty cycle calculation.

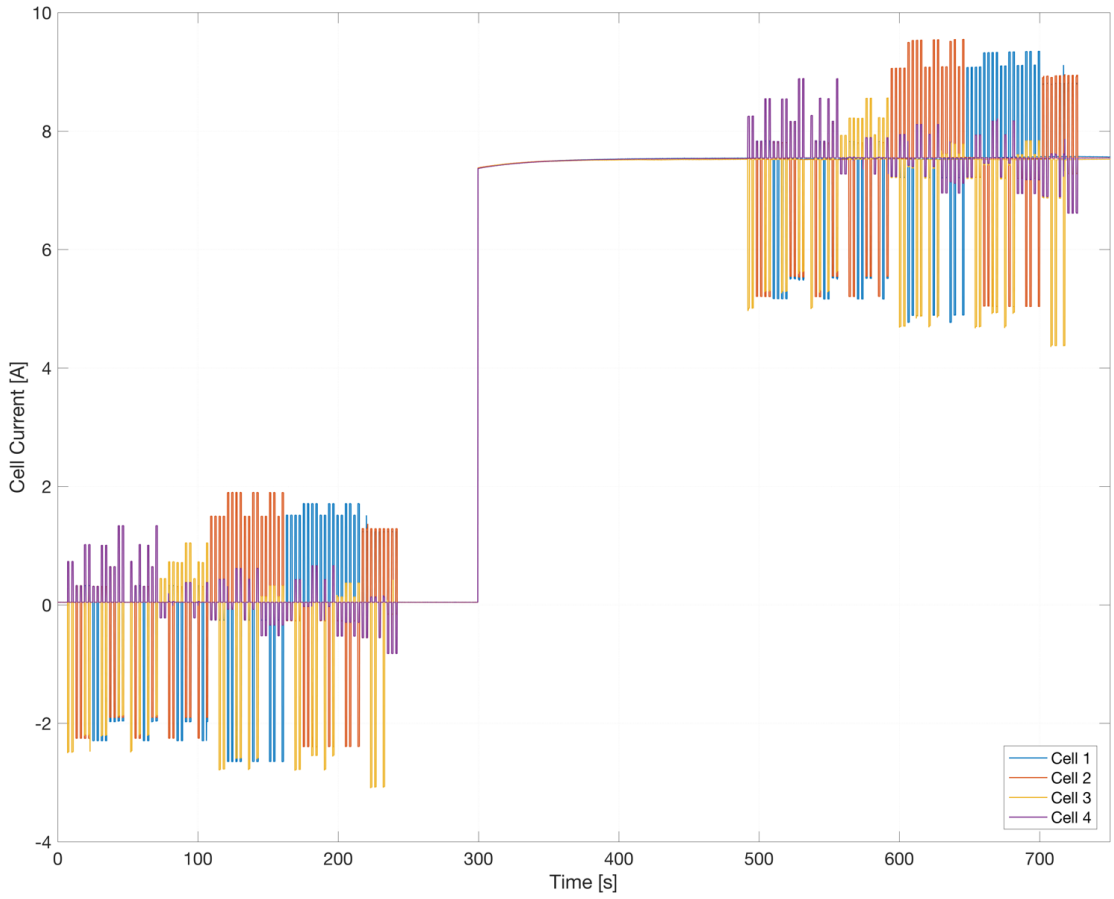

Figure 63: Evaluation of balancing system under external load

Table 21: Difference in balancing currents while under external load

Average Current	Cell 1	Cell 2	Cell 3	Cell 4
difference (mA)	-19.65	-24.99	2.81	6.46

Two approaches were used for converting the controller's current output into a duty cycle, which are fully derived in Submission 5. The first was to find average charge and discharge currents for each cell. To do this, an offline NLLS fit was performed on the permutation dataset. This gives a single reference for each balancing current at a 100% duty cycle, making it straightforward to calculate other duty cycles. However, the variations in balancing current meant that the duty cycle calculation could be inaccurate, with cell 4 showing errors of up to

almost 600mA. An improved method was found which also used NLLS, but as an online optimisation. The goal of the optimisation was to find the combination of duty cycles which most closely matches the charge transferred over the course of the following one second of balancing, explicitly factoring in the changes in current which occur as other cells are switched on and off. This method also has a secondary benefit when used with the DSIC controller from section 5.4.3. This technique also uses a similar optimisation to find balancing currents, given constraints on what currents the hardware can produce. The duty cycle optimisation also factors in these limits as the duty cycles are limited and cell data are directly used. As such, the unconstrained gradient from the DSIC objective function can be used directly as the target currents by the duty cycle optimisation, without performing its own optimisation.

6.7 Control System Experimental Test Results

The technical issues with, and eventual failure of, the balancing hardware meant that it was not possible to perform a single comprehensive set of tests with a completely consistent of experimental conditions. A summary of the tests performed is given in Table 22. The Artemis Combined drive cycle was chosen to act as a representative vehicle load, as it contains a wide range of vehicle operating modes and so varied current excitation. A standard vehicle model developed within Catapult [143] was used to convert the Artemis speed profile into a battery current profile.

Table 22: Summary of experimental test conditions

Test	Controller	Drive Cycle	Cells Used	Estimator	Duty Cycle Calculation	Comments
1	Pole Placement	None	2,3,5,6	1RC CDKF	Average	-
2	Pole placement	Artemis (15A peak)	2,3,5,6	1RC CDKF	Optimisation	The cell 6 current sensor saturated at 12A
3	DSIC	Artemis (20A peak)	2,3,6,7	CDKF With NLLS	Optimisation	Balancing currents for cells 6 and 7 were lower than usual because of a short circuit

6.7.1 Pole placement

Figure 64 shows the cell SOCs and absolute (maximum minus minimum) SOC difference for the pole placement controller with no external load. The circles at the start and end of the SOC plot give the SOCs obtained from using the OCV-SOC with the relaxed cell voltages. This is used to indicate if the estimators are performing poorly. The SOC estimation results are particularly accurate for this test for two reasons. Firstly, there is no external load, the cells are only subject to low currents which affect the cell terminal voltage by a few millivolts. Secondly, duty cycling means that some or all the cells will be undergoing no current (or very low feedback currents) for part of each one second window. This short excitation followed by a rest means the cells are at an almost relaxed state for part of the time – particularly towards the end of balancing when the duty cycles are small. Both of these factors result in strong voltage correction from the estimator. At the end of the test, the mean difference between the estimated and measured SOCs was 0.01% SOC.

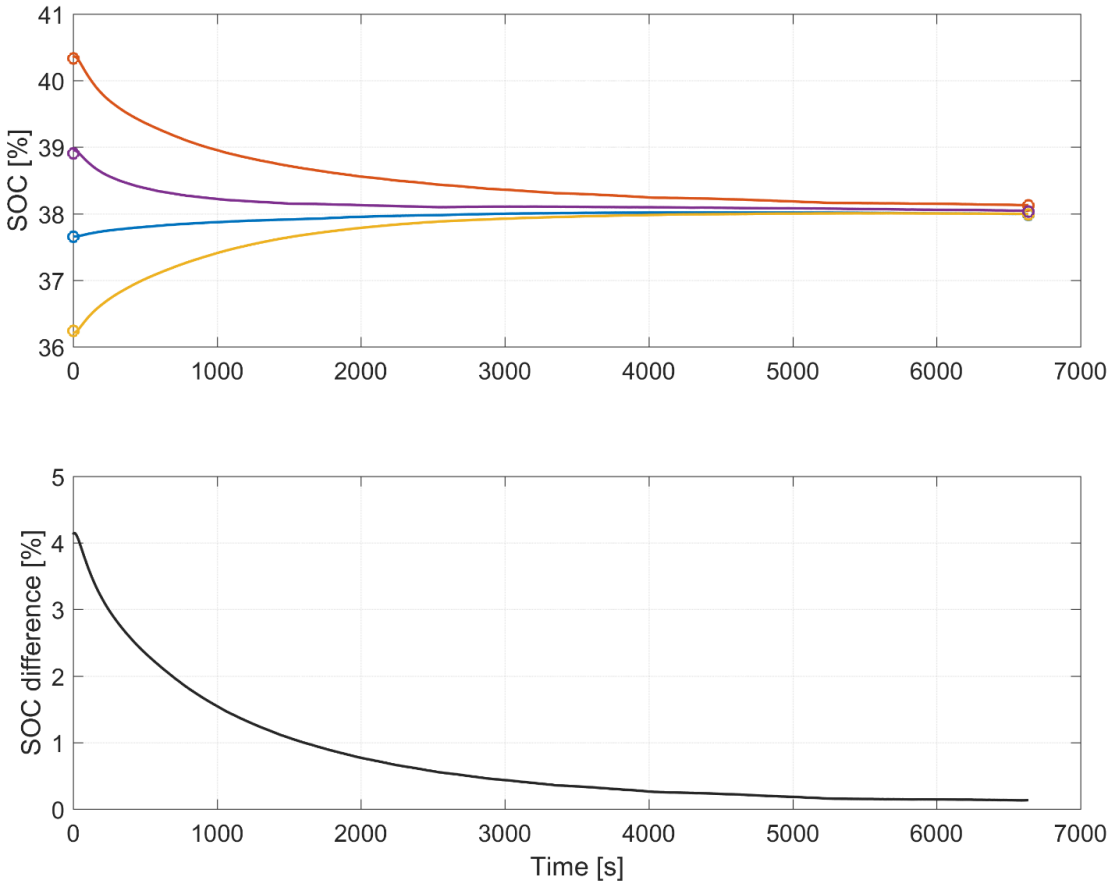

Figure 64: SOC and duty cycle results for pole placement under no load

For this test the closed-loop pole obtained using bisection was -6.06e-4 which equates to a time constant of 1650 seconds. The results show that stable

imbalance removal was achieved in line with the exponential response of a state feedback controller. The absolute imbalance started at 4.1% and was reduced to 0.14% by the end of the experiment. As expected with state feedback, the majority of the imbalance reduction takes place early on in the test. The lower plot shows that absolute imbalance which reduced to 1% after 1650 seconds. This means that the balancing system is on for a long time which, on board a vehicle, would increase the charge lost due to standby power draw, and is not as effective when imbalance has to be reduced in a short period of time.

The same pole placement controller was then applied with a discharging Artemis Combined drive cycle. The drive cycle was repeated until the cell SOCs became relatively low, at which point the balancing system stopped functioning as discussed in section 6.1. The cell currents and resultant SOC estimates are shown in Figure 65.

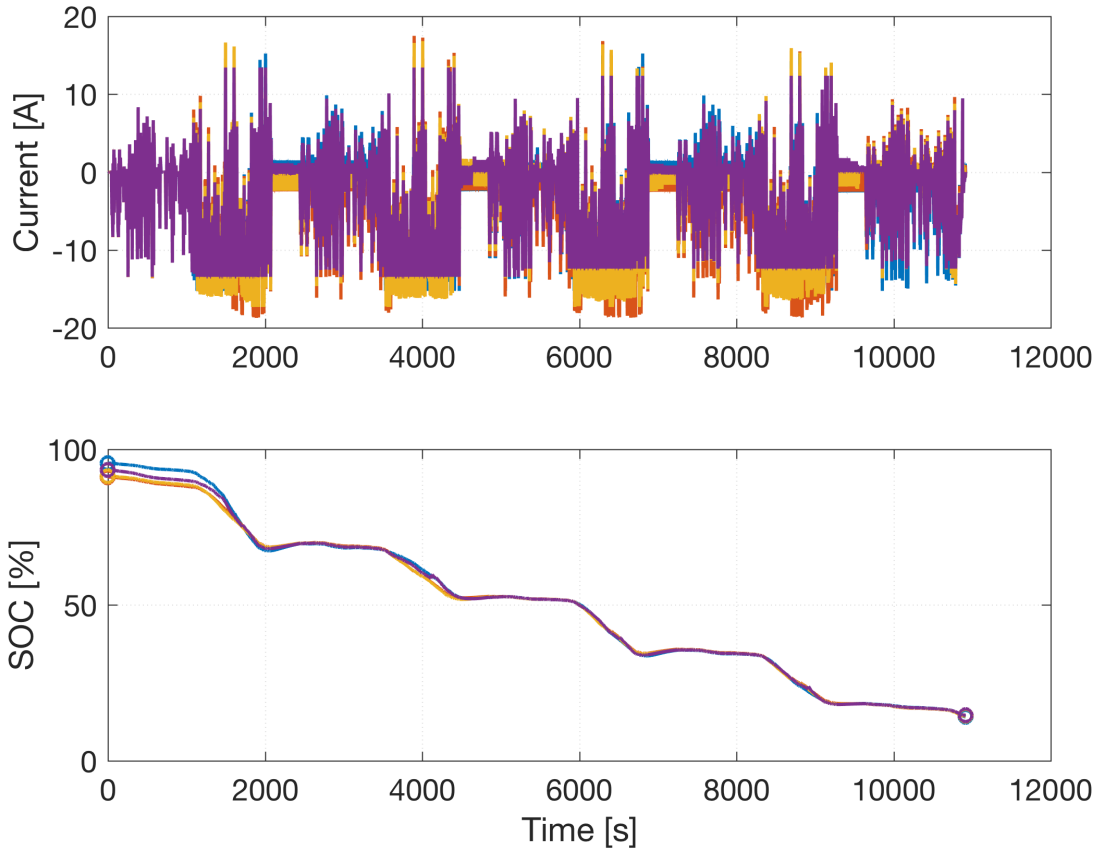

Figure 65: Cell currents and SOCs for pole placement under Artemis Combined drive cycle

They show the periodic discharging profile as the cells discharge from around 92% SOC to 14% SOC. Figure 66 shows the SOC imbalance results. The response is not exponential like the no load case, because of the SOC

uncertainty, the controller not factoring in the load current, and saturation of some of the duty cycles. Despite this, the imbalance is reduced significantly by 30 minutes. After this point, the controller would ideally operate as a regulator reacting to the drive cycle disturbance. This is achieved to some extent, but there are several periods during which imbalance increases. This is typically during the high-current region of the input profile: the 60-75 minute, 100-115 minute and 140-155 minute regions. This apparent increase in imbalance arises from SOC estimation error. The nature of the experiment means there is no baseline reference to observe how accurate the SOC estimation is apart from before and after the test. However, there are two observations on estimator performance. Firstly, after this part of the drive cycle there is a rest period, which should allow the estimator to converge on the 'true' SOC (from the OCV-SOC curve) as the cell relaxes. Secondly, the SOC estimates generally increase during this convergence period, implying that the estimates have been too low during the previous discharging phase.

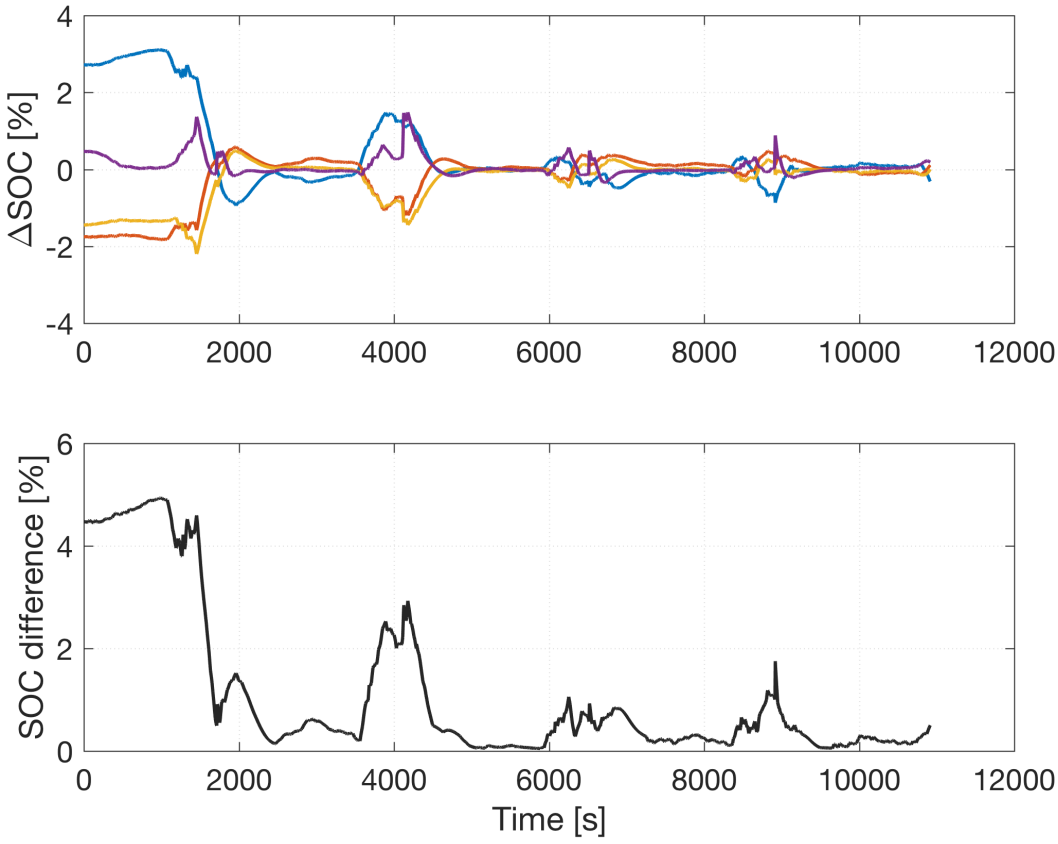

Figure 66: ΔSOC and absolute SOC difference for pole placement under Artemis Combined drive cycle

Pole placement can be slow to respond to these increases in imbalance. As mentioned above for the no load case, performance could be improved by retuning the controller. Given that there may be continual and unexpected changes to the SOC estimates, the controller should respond to them quickly because the uncertainty is not known in the future. A slow response increases the chances that the controller will not be able to reduce imbalance before the cells reach EOC/EOD. As discussed in Chapter 5, while the above pole placement algorithm could be improved to speed up this imbalance response, its fundamental operation is not well suited to fast imbalance removal.

6.7.2 DSIC

To validate the performance of the DSIC algorithm, the same Artemis profile was applied as for the pole placement tests, but starting from around 60% SOC with a larger initial imbalance. The current and SOC results are shown in Figure 67, and the relative and absolute SOC differences shown in Figure 68. The drive cycle is scaled to a higher maximum value (20A rather than 15A).

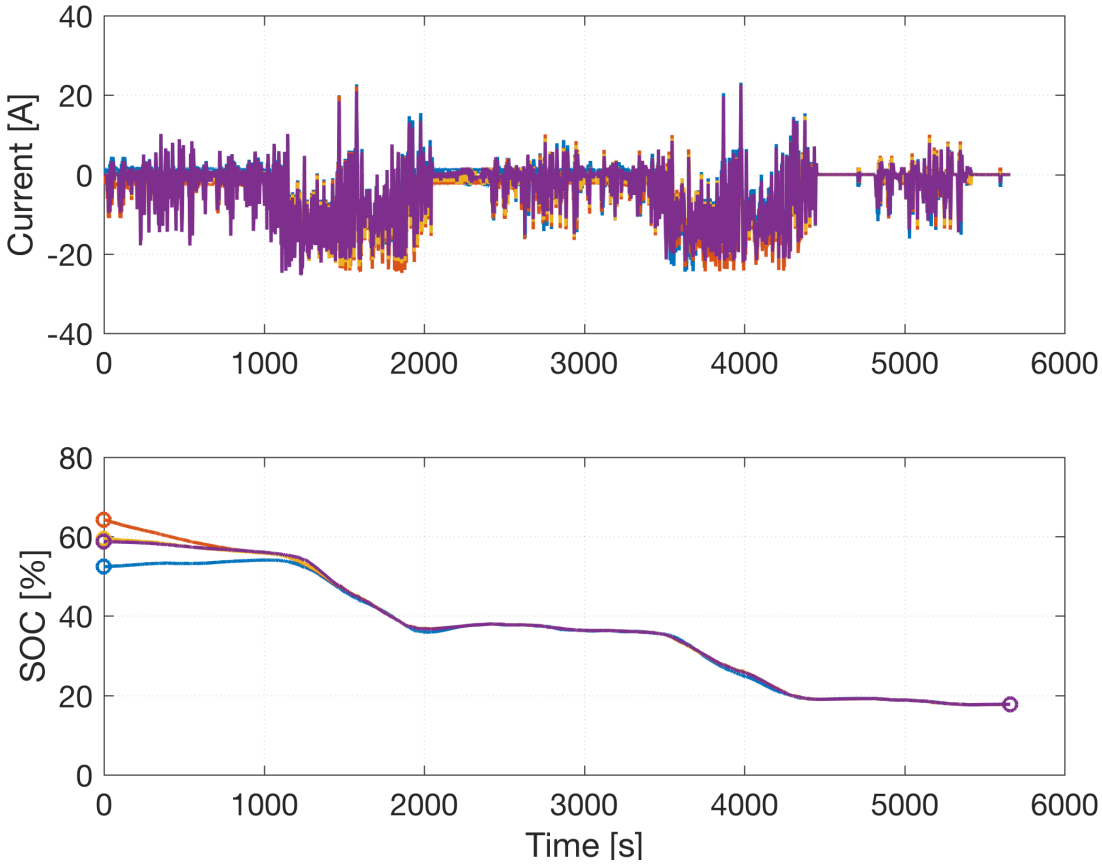

Figure 67: Cell currents and SOCs for DSIC under Artemis Combined drive cycle

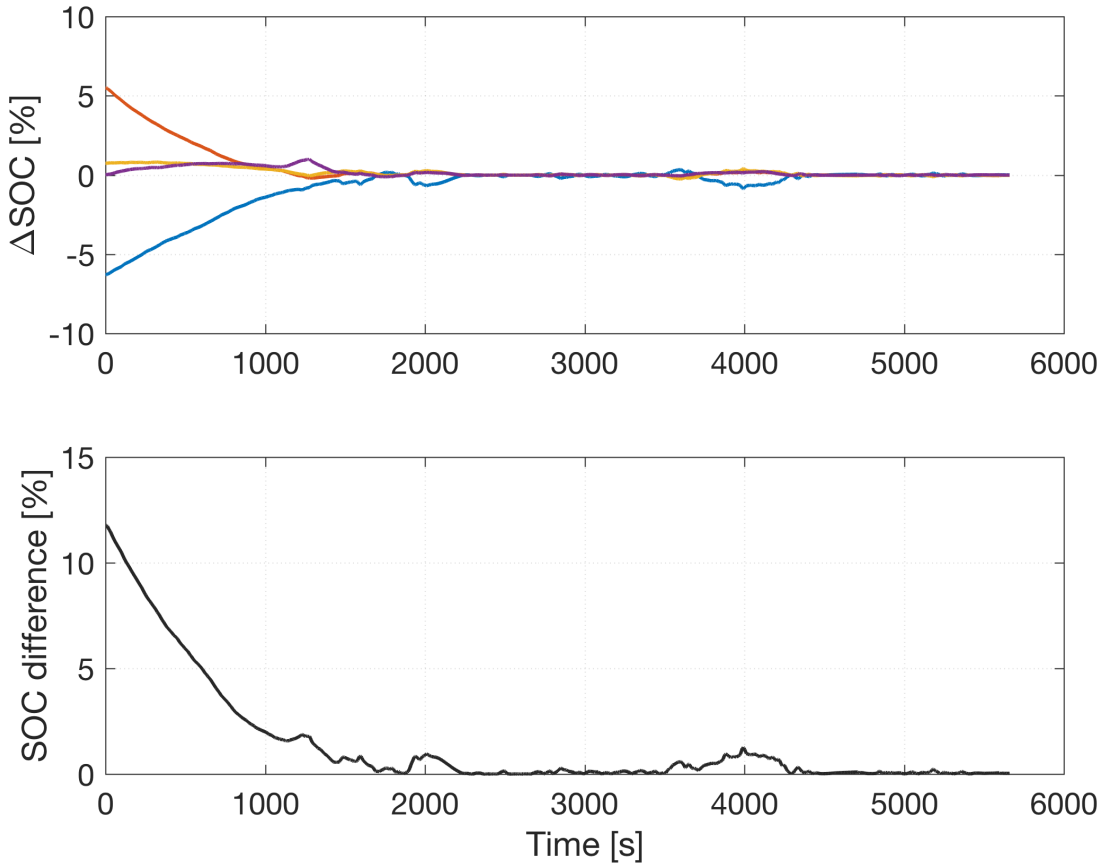

Figure 68: ΔSOC and absolute SOC difference for DSIC under Artemis Combined drive cycle

The results show that the controller can quickly reduce the SOC difference, and maintain a low difference over the drive cycle. This quadratic controller is quicker to returned the system to a balanced state in the event of a disturbance such as SOC estimations correcting during cell relaxation. In this sense, the DSIC approach effectively acts as a better regulator than pole placement in the face of the various sources of uncertainty.

6.8 Simulation Studies

To complement the experimental results above, a number of SIL studies were undertaken to improve the breadth of the validation activity for both the DSIC and Pole Placement designs. Unlike in Chapter 5, these simulations include a model of the specific balancing hardware used for this work, the details of which are contained in Submission 5. These studies allow for a comparison of different controllers under the exact same conditions, which was difficult to achieve with the experimental set-up. The 3RC plant cell model from section 4.2 and 1RC estimator from section 4.3 were used for these studies. While the 3RC estimator showed superior performance, it also uses the same model parameters as the

model, making the estimator essentially redundant. Using the 1RC estimator induces some uncertainty, for a more thorough analysis of the control systems.

The experimental work in this Chapter showed that balancing under load was more challenging than under no load, as the estimator performance was worse. These simulation studies use the same US06 Highway drive cycle in Figure 39. The current from that model was rescaled to give the same C-rate for these cells as for the battery pack in the vehicle model. The starting SOCs of the cells is the same in both cases, with the simulation using cells 2,3 5 and 6 and the SOCs in Table 15. As can be seen in Figure 69, pole placement does cause the cell SOCs to converge, but not fully, with it taking over 1 hour to reduce total imbalance to under 1%, and then converging to a steady state error of about 0.5%. The results for DSIC shown in Figure 70 show it is able to remove imbalance almost completely within 15 minutes, and maintain this state of balance despite the presence of a drive cycle.

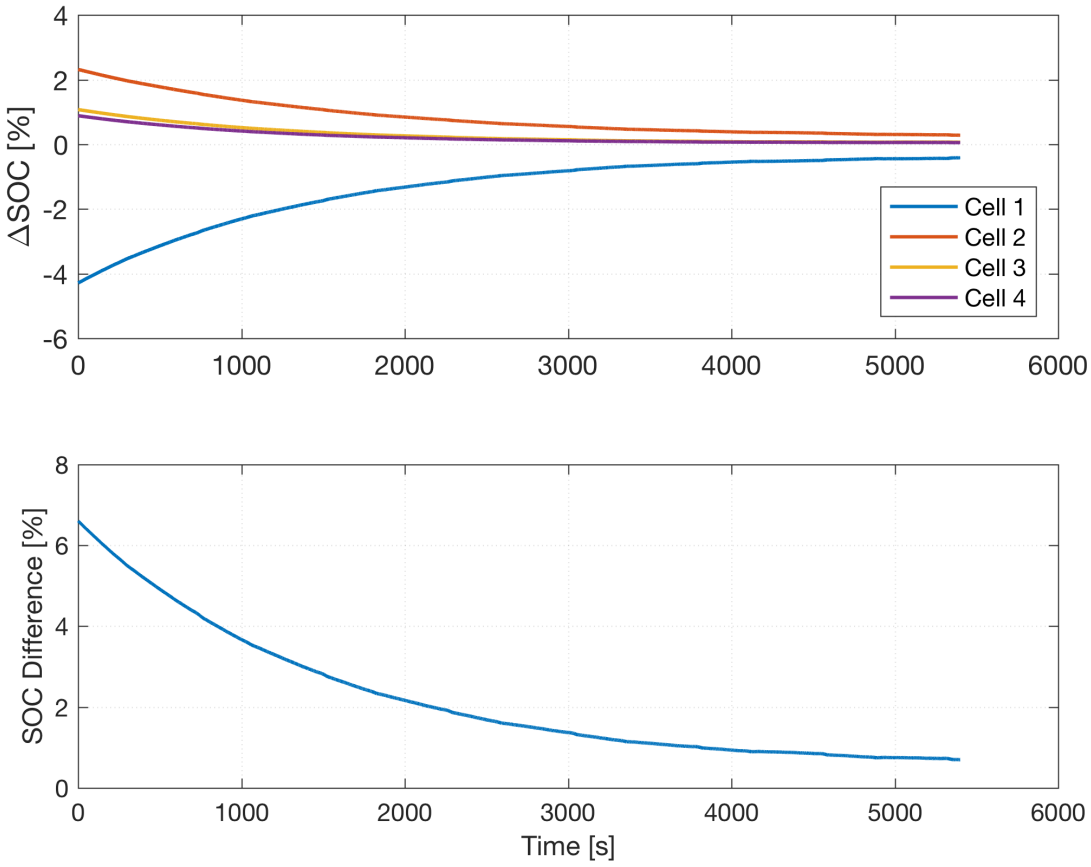

Figure 69: Simulation results for pole placement under the US06 drive cycle

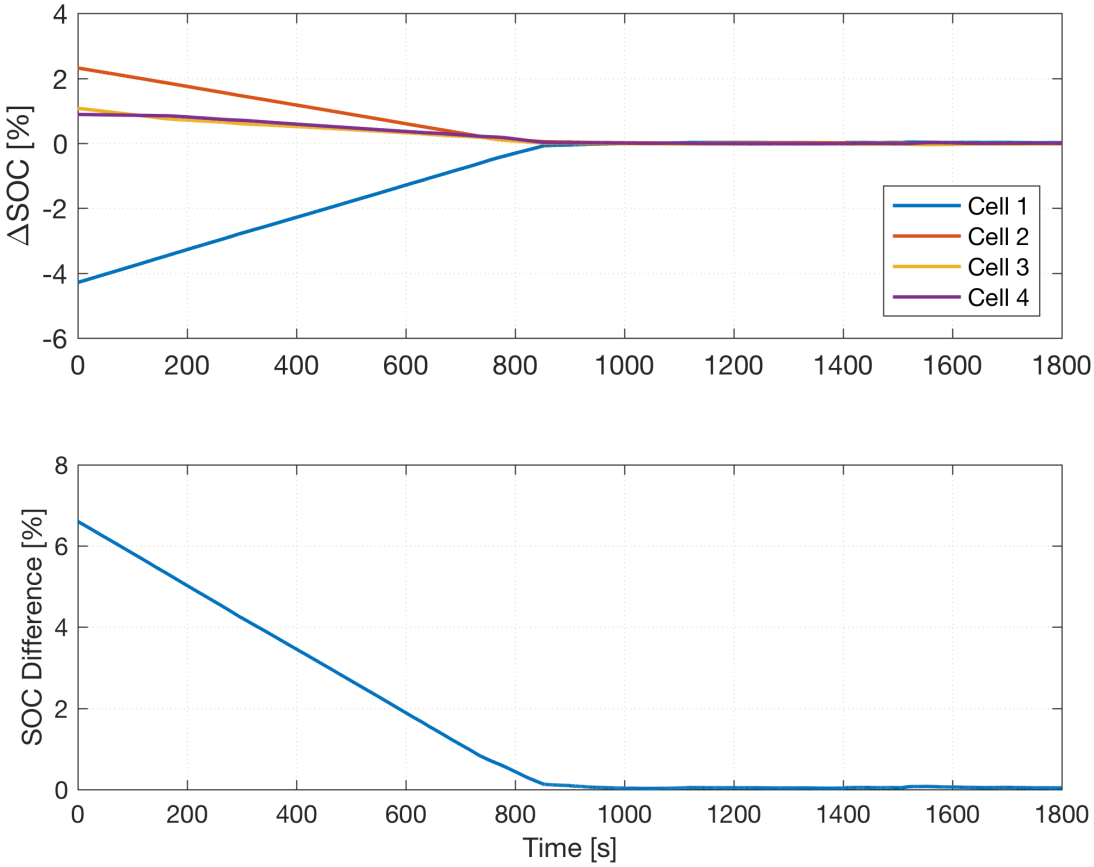

Figure 70: Simulation results for DSIC under the US06 drive cycle

6.9 Control System Analysis

The DSIC controller was the best-performing of the controllers. It removed imbalance at a much faster rate, and was more effective at maintaining balance in the presence of uncertain SOC estimates. The proposed strategy from section 5.1 is to only balance when it appears that the vehicle driver would otherwise completely discharge one of the cells. Active balancing during this phase increases the driving range by extending the time until one cell reaches end of discharge (ideally all cells reach end of discharge simultaneously). Being able to balance faster with the same hardware means that balancing can be left until later during discharge, which in turn increases the likelihood that it will not have to be used at all, improving energy efficiency.

The experimental results shown in Figure 68 brought the cells to within 0.05% SOC of each other (measured from the relaxed cell voltage not from the SOC estimates), with a standard deviation of 0.02%. This emphasises that while there has been much discussion of SOC estimation error, this error is not statistical noise but follows the same trend between cells. For example, if one cell's SOC

estimate is 2% above the true value, it is likely that all the other cells' SOC estimates are about 2% above the true value, rather than some cells being up to 2% above and others up to 2% below. The latter case would be problematic, as the spread of imbalance would be much higher.

The battery utilisation factor, *UF*, is presented in (94) as a means of assessing what proportion of the pack's charge has been utilised, where 100% implies that all cells have been fully discharged. Differences in final SOC s^f will reduce the utilisation factor as it means there is charge remaining in some cells. Assuming that the cells started off fully charged and balanced, for the cells in Table 9, the *UF* for passive balancing is 92%, owing to the capacity differences.

$$UF = \frac{\sum_n \left(Q_{cap,n} \frac{[100 - s_n^F]}{100} \right)}{\sum Q_{cap,n}} \tag{94}$$

For active balancing, the charge lost from system inefficiency during balancing must be accounted for. This can be calculated using (95), where for this system efficiency is taken as 0.8. The amount of charge lost during balancing will depend on the amount of imbalance which is to be removed. For this study it is assumed that balancing starts when one cell reaches 20%, resulting in the SOC vector from Table 15. The *UF* for active balancing can then be calculated using (96). The uncertainty from SOC estimation error can also be incorporated by considering a set of normally distributed SOC values, whose minimum value is zero (i.e. one of the cells has been completely discharged). Using the experimental final imbalance standard deviation of 0.02%, a Monte-Carlo simulation was run to understand the impact of cells finishing at different SOCs based on the uncertainty in their estimate. Ten-thousand normally distributed final SOC vectors were generated, and (96) applied each time.

$$Q_{loss} = (1 - \eta) \int |i_{bal}| \, dt \tag{95}$$

$$UF = \frac{\sum_n \left(Q_{cap,n} \frac{[100 - s_n^F]}{100} - Q_{loss} \right)}{\sum Q_{cap,n}} \tag{96}$$

From these results, the mean utilisation factor for these cells was 99.19%, with a worst-case of 99.14%. **This means that almost all the battery pack charge can be extracted, with a 7% improvement over the 92% offered by passive balancing**. This has several potential benefits. The driving range of the vehicle could be increased, which is a key selling point for BEVs. Alternatively, the same range could be achieved with a smaller pack (saving cost, weight and complexity) or a narrowed DOD window (improving efficiency and reducing ageing).

6.10 Balancing Hardware Requirements

Several observations were noted during the commissioning and validation phases of this research using the energy balancing hardware based on the LTC300-1 system. The aim of this section is to reflect on this research, highlight the challenges faced and to make recommendations for future hardware designs to improve performance and integration with a real-world BMS solution.

For simplicity and control robustness, balancing current generation would be independent of factors such as the number of cells and the cell voltages. This minimises the set-point error, ensuring that the requested current from the controller is reliably generated by the hardware. In practice this is unlikely to be feasible, but should be mitigated against as much as possible. The LTC3300-1 chip comes relatively close to achieving this, as the theory of operation and simulation studies in section 6.2 showed the current only varies a small amount with cell voltage and internal resistance. The hardware design was such that the discharging current depended on primary side connections, whereas the charging current depended on secondary side connections. This means the charging current is dependent on the number of cells, whereas the discharging current is independent. The cell connections will most likely remain unchanged during operation, so this dependency can be resolved during the design process.

If there is some dependency, it should be predictable and repeatable. That is, the factors causing the balancing current to change should be known, and so can be factored into the control system accurately calculates the balancing current with minimal set-point error. The LTC3300-1 system is sensitive to the sense resistor

values, meaning that each channel had to be individually tested and calibrated, but when setting a cell to balance, the current generated was consistent. A key issue with the LTC3300-1 system is the balancing current varying when other cells were also set to balance. The cause of this variation could not be found, and so an analytical set of equations describing it could not be derived and incorporated into the model. Despite this, it was repeatable and so accounted for.

The cell-module connections influence the effectiveness of the balancing system. A series of isolated modules, each with its own balancing system, will more than likely result in not being able to globally reduce imbalance. As was first discussed in section 2.1.2, each module can have a very different SOC distribution, and if one module's average SOC is much lower than the others, it will still limit the battery pack despite the presence of an ABS. The obvious solution would be to have one module: the entire battery pack. This is generally difficult to achieve because of the large ratio between pack and cell voltage (around 100:1 for a BEV). However, there is progress on this front. Module-module balancing has been proposed, as discussed in 2.2.1, though this adds a considerable number of components, and with it cost, complexity and new failure modes. The interleaving method recommended by LTC is it allows cells from one module to transfer energy to cells in a neighbouring module, without any additional circuitry. In this case, the number of modules should be kept to a minimum, as this will increase the amount of time to balancing modules (the logical extreme is having one cell per module, i.e. cell-to-cell balancing, where as noted in section 2.2.1 is ineffective at moving energy between areas of the pack). The control model from section 5.3 can adapt to these different cases. It contains knowledge of the cell connections and so will inherently balance neighbouring modules.

6.11 Summary

A number of integration issues became apparent while commissioning the balancing hardware. The ABS was delivered with a combined balancing and voltage measurement harness. This proved to be unsuitable for measuring cell voltage, with the voltage during balancing being hundreds of millivolts from the actual cell voltage. These types of compromise are common for cost and reliability reasons, but may have to be revised for more advanced active

balancing hardware and control strategies. Another issue raised is the additional higher-frequency dynamics induced by balancing. The impact of this will depend on the type of balancing system. In this case, it was induced by the duty cycling scheme which operated at 100Hz. This created sudden, rapid changes in voltage as balancing was switched on and off, and it was important to capture this voltage as well as ensuring the associated current measurement was synchronised as part of the state and parameter estimation process.

The experimental results demonstrated that imbalance can be removed and largely maintained in the face of SOC estimation uncertainty, with charge utilisation increased from 92% to over 99% for the cells under test. One of the issues with pole placement, first noted in Chapter 5, is that the corrective action reduces with imbalance. The experimental results showed that the main issue with this in practice is that it is slow to adjust to SOC estimation errors. The balancing system can only work with the SOC estimates it has, and so there is a chance that it has apparently removed imbalance, only to discover later that there is actually still a significant amount of imbalance because the SOC estimate was incorrect. If this happens close to EOD then the pack may reach EOD in an imbalanced state and not all the energy was utilised. Clearly, the SOC estimate is crucial to balancing system performance, and it may be that if there is some trade-off (either in terms of computational limitations or development time) then a better SOC estimate combined with a simple rule-based strategy may result in better energy utilisation than the DSIC controller with a poorer estimate.

7 Conclusions

A thorough analysis has been performed on energy imbalance within a multi-cell battery system, ranging from broad questions on the causes of imbalance through to specifics of how energy can be transferred between cells. The wider context of balancing was addressed first. An underpinning theme throughout this research is that active balancing is often not considered holistically. Many studies have focused only on hardware, and demonstrating functionality rather than showing how the technology can be used to improve overall system performance relative to passive balancing. From the literature review, there was an awareness within the sponsoring company that active balancing exists, but a lack of understanding of exactly how and when it should be used, and how to quantify its benefit and assess its commercial suitability. This challenge is further complicated by the various types of hardware and disparate claims often made by many manufacturers within the supply-chain.

The first Research Question addressed many of the issues above. It was answered by critically reviewing the literature and addressing several flaws and missing or contentious areas with an active balancing control strategy. The first Research Objective led to the conclusion that a 10% difference between the highest and lowest cell capacities is a conservative expectation for automotive cells in use for several years. For the seven cells used to test the ABS for this work, only up to 92% of the total capacity could be utilised if passive balancing is implemented. The second Research Objective emphasised that there was a wide range of balancing hardware, but little research into how it should be controlled, and how much benefit could actually be achieved through its adoption within an EV application. The focus of existing research was often on the balancing speed, rather than energy utilisation. For this reason, balancing during discharging was almost entirely neglected within the literature, which is the only means by which active balancing can utilise energy where passive balancing cannot.

The increased energy utilisation has several implications for long-term automotive applications. From a performance perspective, the reduction in the vehicle's driving range several years after purchasing will be less significant, a

key selling point for potentially customers who expect many years of consistent performance, and an important metric for retaining existing customers. Another pertinent topic for EVs is warranty claims. Active balancing could prevent one weak cell causing the entire battery pack to be replaced under warranty, which as well as being costly to the manufacturer is a negative customer experience. Increased energy utilisation means that a smaller battery pack could be used to achieve the same performance. With the battery pack being one of the largest and most costly elements of an EV.

These issues are the motivation behind the second Research Question. Given these statistics and high-level strategy, the increase in energy utilisation needs to be quantified, as this is the main benefit and justification for the use of active balancing. This was broken down into three stages. Research Objective 3 was to develop a modelling and simulation framework which could be used for balancing system design, and was essential for estimating SOC as part of the experimental process. This then allowed for Research Objective 4 to be met, whereby a number of controllers were designed and evaluated. A simple, non-model based scheme was found to be effective, but optimisation-based controllers were found to be up to 22% more energy efficient by employing lower average currents. This novel control framework was specifically designed to be generic, such that it can easily be tailored to any balancing system, maintaining its usefulness as balancing systems evolve over time and new technologies become available. To test these control systems, an ABS based on commercial hardware was commissioned, and connected to aged automotive cells. There were several learning points from the integration process, notably on obtaining accurate cell voltage measurements and minimising set-point error upon requesting balancing currents from the ABS. The experimental results showed a substantial increase in energy utilisation through active balancing. For the seven cells under test, over 99% of the pack energy could be utilised compared to only 92% when using passive balancing.

7.1 Impact and Innovation

The primary impact for the sponsoring company is that they now have the knowledge and capability to evaluate and integrate different active balancing systems. This capability is underpinned through the knowledge gained through the successful implementation of an ABS with production-intent hardware and automotive cells. The suitability of different applications has been assessed, with the conclusion that ABSs only offer a notable benefit when the application utilises close to the full pack DOD, such as for most BEV applications. A high-level strategy solves the problem of when to balance. This minimises ABS usage (which is beneficial for efficiency and reliability), while ensuring the entire battery pack capacity is available. SOC was shown to be the most suitable metric for quantifying imbalance, despite voltage being more commonly adopted, and experimental results showed that imbalance can reduced to within 0.1% SOC, even in the presence of SOC estimator uncertainty. Numerous physical integration challenges have been addressed which were not considered by the makers of balancing hardware, notably on accurate cell voltage measurement. The next time a supplier offers some evaluation hardware for a new balancing system, the sponsor company can now perform high-level calculations and specific, consistent tests to quantify its performance, and suitability for a specific vehicle platform.

This work also feeds into the expanding second-life market, in which used cells, no longer capable of meeting performance targets for their original application, can be combined to form battery packs in less intensive environments such as grid storage for wind farms. There will inherently be a lot of cell-cell variation, which is where active balancing can offer a significant performance benefit. By prolonging the useful life of cells, manufacturers can extract more profit either by selling used battery packs to a second-life specialist, or by repurposing their own battery packs for other applications. This favourably shifts the cost-benefit ratio of ABSs when considering the total lifetime of the cells.

The key innovation of this research is that a control strategy was developed for the first time with the explicit aim of maximising energy utilisation. This strategy

includes balancing while the battery pack is being discharged, which to the author's knowledge has not been performed experimentally before, but is crucial to energy utilisation. A model-based design framework has been created which can be used to design a variety of controllers independent of the type of hardware. This model addresses issues such as controllability and over-actuation, which to date have not been addressed within the literature for this domain problem. Similar MPC algorithms to those developed in this research have been recently published by other authors, but this framework has scope to underpin more control solutions. For example, the feed-forward controller has the potential to further reduce energy loss during balancing while being simpler than other proposed solutions, but was unable to be validated experimentally. As mentioned above, SOC estimators were successfully implemented as part of the real-time control loop. Concern over the reliability of SOC estimation has been raised within the literature on active balancing, but until now has not been addressed and quantified experimentally. These results demonstrate conclusively that the use of SOC estimation is suitable within the context of an overall active balancing control system.

7.2 Further Work

A number of research gaps were identified during this work which are yet to be addressed. Additionally, there are some specific aspects which arose during the experimental work. By exploring these further, improvements could be made to cell models and future BMS performance.

7.2.1 Research Gaps

One of the issues first noted in section 2.2 is that it is very difficult to compare the various types of hardware proposed. This is because they have been tested with different numbers of cell (and types of cell), are designed to operate at various power levels and are controlled in different ways. The fact that one ABS balances faster than another does not mean a great deal, as it is likely they could both operate at very different speeds if different component values are chosen. The balancing control framework proposed in this innovation report could be used as a basis of a more systematic and rigorous comparison of different topologies. Of

interest are 3 types of system: smart cells with a built-in converter, a cell-module system with simultaneous operation, such as the LTC3300-1 used in this work, and a cell-module system which uses a shared converter to reduce the component account. These three systems, highlighted by the initial literature review, should be designed such that they can balance at a similar rate according to the control system proposed here. They could then be evaluated to compare the cost, size, efficiency and complexity to give a much better idea of the suitability and feasibility of different ABSs for different applications.

The other area of ABS research is how to scale up from a few cells to a large battery pack. This is not a topic addressed within the literature, because of the difficulty involved in building a complete battery pack within a university research environment. This would better indicate how global imbalance reduction is affected by having large numbers of cells and module-to-module balancing connections, and would likely raise many integration issues that are not currently apparent. The reliability of ABSs has also not been addressed, likely because of the lack of commercial or long term applications to date, but this is important in validating the feasibility of ABSs in such applications.

There is also scope for further control system development. Issues with test equipment meant that it was not possible to continue to develop and refine the MPC controller or implement the feed-forward controller. It may be possible to adjust the cost function to reduce energy loss or improve balancing time. The rule for adjusting the time window is also quite simplified, and ideally it would be calculated as part of the optimisation routine.

7.2.2 Model Performance

While cell models were useful for several aspects of this work, they could not be used as effectively as was hoped. SOC estimator development was hindered by the limitations of ECM accuracy over wide DOD ranges, notably during cell relaxation. Accurately modelling relaxation would enhance convergence to the true SOC when the applied load has stopped even for a short period of time (such as a vehicle waiting at a junction). Using similar ECMs for the estimator and plant

model means that it is difficult to robustly tune the estimator. As such, the creation of an extended ECM to better capture these low frequency dynamics, or continuing research into making PBMs feasible for plant models, will further help to improve SOC estimator design and therefore balancing system performance. There were also some issues with the OCV-SOC data generated for this work. The procedure employed does not thoroughly validate the curves, and some errors were apparent in earlier Catapult OCV-SOC test data which was not identified as the focus was on ECM development. Developing a procedure for generating OCV-SOC curves and ECM parameterisation data specifically for SOC estimation purposes could help overcome these issues and further improve estimator performance.

7.2.3 Robustness Analysis

While the control system has been demonstrated on representative hardware, an important means of assessing its performance long term, and for different applications, is through robustness analysis. This quantifies how the control system performance is affected by various noise sources, and the properties of the noise. Of particular relevance to this control system is the impact of SOC estimation accuracy, and cell capacity estimation accuracy. The former is critical for assessing the state of imbalance, and so what the controller must do to remove this imbalance, and the latter is a key part of the model for projecting forward to calculate how much charge (and so current) must be added or removed to each cell. The robustness characteristics are also important for setting requirements for other subsystems, for example how accurate the SOC estimator needs to be, or the noise properties of the current sensors.

7.2.4 Productionisation

The test work here was based on prototype hardware. As mentioned in section 2.2.3, there are a number of additional factors in when making the decision whether to include an ABS in a battery pack. Research into minimising the cost and component count of the power electronics, as well as reducing the amount of wiring and number of connections would all improve the business and engineering cases for active balancing. Similarly, research into reliability, failure

modes and safety considerations over the operating life-time would help inform design decisions and help identify suitable applications.

7.3 Reflection

One of the recurring themes throughout this research was some unexpected relationships between different topics. The parallel cell modelling involved thinking about how interconnections between cells can be modelled, and currents calculated. This indirectly led to the balancing control model. Despite appearing quite different, in both cases there is a dependency between each cell which dictates their currents. Similarly, the quadratic optimisation/ NLLS algorithm was originally written for the online DSIC implementation, but two other online applications emerged: parameter estimation and duty cycle calculation. This emphasises the benefit of systems thinking and writing code for efficient reuse.

Much of the learning over the course of this research came from the experimental work and specifically adopting an iterative approach. The parallel cell model was only accurate because, after evaluating the results from the standard procedure, additional tests were performed with data collection at a higher bandwidth. Similarly, many of the requirements, performance evaluation and strategy elements of the balancing system only crystallised while testing the controllers, even though much of it now seems obvious in hindsight. It is unfortunate that many more iterations of the control systems could not be performed, as this would have allowed for a more thorough analysis and refinement of the controllers: the DSIC scheme and state estimators were largely designed and tested online due to time and resource constraints, and the feed-forward controller was only developed after the ABS had terminally failed. The most challenging part of the research was in having to respond to continued equipment failure and unavailability, which considerably reduced the time available for control system development. However, this is part of the research and development process, and was itself a valuable learning process, both for technical aspects (planning, continual improvement and debugging electronics) and personal development (communication, perseverance and resource management).

References

[1] A. Väyrynen and J. Salminen, “Lithium ion battery production,” *J. Chem. Thermodyn.*, vol. 46, no. 0, pp. 80–85, Mar. 2012.

[2] J. W. Fergus, “Recent developments in cathode materials for lithium ion batteries,” *J. Power Sources*, vol. 195, no. 4, pp. 939–954, 2010.

[3] C. Arcus, “A Tale of 3 Battery Packs,” *Clean Technica*, 2016. [Online]. Available: https://cleantechnica.com/2016/01/06/a-tale-of-3-battery-packs/. [Accessed: 28-Aug-2017].

[4] M. A. Delucchi and T. E. Lipman, “Chapter Two - Lifetime Cost of Battery, Fuel-Cell, and Plug-in Hybrid Electric Vehicles,” in *Electric and Hybrid Vehicles*, G. Pistoia, Ed. Amsterdam: Elsevier, 2010, pp. 19–60.

[5] FIA Formula E, “The Formula E Battery: Q&A with Craig Wilson,” 2014. [Online]. Available: http://www.fiaformulae.com/en/news/2014/september/the-formula-e-battery-a-qa-with-craig-wilson.aspx. [Accessed: 28-Aug-2017].

[6] Fred Lambert, “Teardown of new 100 kWh Tesla battery pack reveals new cooling system and 102 kWh capacity | Electrek,” 2017. [Online]. Available: https://electrek.co/2017/01/24/tesla-teardown-100-kwh-battery-pack/. [Accessed: 28-Aug-2017].

[7] D. Andrea, *Battery Management Systems for Large Lithium Ion Battery Packs*. Artech House, 2010.

[8] T. Bruen, J. Marco, and M. Gama, “Model Based Design of Balancing Systems for Electric Vehicle Battery Packs,” *IFAC-PapersOnLine*, vol. 48, no. 15, pp. 395–402, 2015.

[9] T. Bruen, J. Marco, and M. Gama, “Current Variation in Parallelized Energy Storage Systems,” *2014 IEEE Vehicle Power and Propulsion Conference (VPPC)*. pp. 1–6, 2014.

[10] T. Bruen and J. Marco, “Modelling and experimental evaluation of parallel connected lithium ion cells for an electric vehicle battery system,” *J. Power Sources*, vol. 310, pp. 91–101, 2016.

[11] T. Bruen, J. M. Hooper, J. Marco, M. Gama, and G. H. Chouchelamane, “Analysis of a battery management system (BMS) control strategy for

vibration aged Nickel Manganese Cobalt Oxide (NMC) Lithium-Ion 18650 battery cells," *Energies*, vol. 9, no. 4, 2016.

[12] M. Dubarry, N. Vuillaume, and B. Y. Liaw, "Origins and accommodation of cell variations in Li-ion battery pack modeling," *Int. J. Energy Res.*, vol. 34, no. 2, pp. 216–231, 2010.

[13] H. Bockholt, M. Indrikova, A. Netz, F. Golks, and A. Kwade, "The interaction of consecutive process steps in the manufacturing of lithium-ion battery electrodes with regard to structural and electrochemical properties," *J. Power Sources*, vol. 325, pp. 140–151, 2016.

[14] K. G. Gallagher and P. A. Nelson, "Manufacturing Costs of Batteries for Electric Vehicles," in *Lithium-Ion Batteries: Advances and Applications*, Elsevier, 2014, pp. 97–126.

[15] G. E. Blomgren, "The Development and Future of Lithium Ion Batteries," *J. Electrochem. Soc.*, vol. 164, no. 1, pp. A5019–A5025, Jan. 2017.

[16] J. Kim and B. H. Cho, "Screening process-based modeling of the multi-cell battery string in series and parallel connections for high accuracy state-of-charge estimation," *Energy*, vol. 57, no. 0, pp. 581–599, Aug. 2013.

[17] J. Gallardo-Lozano, E. Romero-Cadaval, M. I. Milanes-Montero, and M. A. Guerrero-Martinez, "Battery equalization active methods," *J. Power Sources*, vol. 246, no. 0, pp. 934–949, 2014.

[18] F. Sun and R. Xiong, "A novel dual-scale cell state-of-charge estimation approach for series-connected battery pack used in electric vehicles," *J. Power Sources*, vol. 274, pp. 582–594, Jan. 2015.

[19] S. F. Schuster, M. J. Brand, P. Berg, M. Gleissenberger, and A. Jossen, "Lithium-ion cell-to-cell variation during battery electric vehicle operation," *J. Power Sources*, vol. 297, pp. 242–251, Nov. 2015.

[20] "Nissan and Eaton power ahead with second-life battery system," 2015. [Online]. Available: http://newsroom.nissan-europe.com/eu/en-gb/media/pressreleases/140285. [Accessed: 30-Jun-2017].

[21] "Renault and Powervault Give EV Batteries a 'second-life' in Smart Energy Deal," 2017. [Online]. Available: http://media.renault.com/global/en-gb/renaultgroup/Media/PressRelease.aspx?mediaid=92203. [Accessed: 30-Jun-2017].

[22] D. Shin, M. Poncino, E. Macii, and N. Chang, "A statistical model of cell-to-

cell variation in Li-ion batteries for system-level design," *International Symposium on Low Power Electronics and Design (ISLPED)*. pp. 94–99, 2013.

[23] B. Kenney, K. Darcovich, D. D. MacNeil, and I. J. Davidson, "Modelling the impact of variations in electrode manufacturing on lithium-ion battery modules," *J. Power Sources*, vol. 213, no. 0, pp. 391–401, Sep. 2012.

[24] F. An, L. Chen, J. Huang, J. Zhang, and P. Li, "Rate dependence of cell-to-cell variations of lithium-ion cells," vol. 6, p. 35051, Oct. 2016.

[25] M. Dubarry, N. Vuillaume, and B. Y. Liaw, "From single cell model to battery pack simulation for Li-ion batteries," *J. Power Sources*, vol. 186, no. 2, pp. 500–507, 2009.

[26] V. Muenzel *et al.*, "A Comparative Testing Study of Commercial 18650-Format Lithium-Ion Battery Cells," *J. Electrochem. Soc.*, vol. 162, no. 8, pp. A1592–A1600, 2015.

[27] S. Paul, C. Diegelmann, H. Kabza, and W. Tillmetz, "Analysis of ageing inhomogeneities in lithium-ion battery systems," *J. Power Sources*, vol. 239, pp. 642–650, Oct. 2013.

[28] T. Baumhöfer, M. Brühl, S. Rothgang, and D. U. Sauer, "Production caused variation in capacity aging trend and correlation to initial cell performance," *J. Power Sources*, vol. 247, no. 0, pp. 332–338, Feb. 2014.

[29] K. B. Hatzell, A. Sharma, and H. K. Fathy, "A survey of long-term health modeling, estimation, and control of Lithium-ion batteries: Challenges and opportunities," in *American Control Conference (ACC), 2012*, 2012, pp. 584–591.

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

[31] T. Waldmann, M. Wilka, M. Kasper, M. Fleischhammer, and M. Wohlfahrt-Mehrens, "Temperature dependent ageing mechanisms in Lithium-ion batteries – A Post-Mortem study," *J. Power Sources*, vol. 262, pp. 129–135, Sep. 2014.

[32] Y. Xing, E. W. M. Ma, K. L. Tsui, and M. Pecht, "Battery Management Systems in Electric and Hybrid Vehicles," *Energies*, vol. 4, no. 12, pp. 1840–1857, 2011.

[33] A. Farmann, W. Waag, A. Marongiu, and D. U. Sauer, “Critical review of on-board capacity estimation techniques for lithium-ion batteries in electric and hybrid electric vehicles,” *J. Power Sources*, vol. 281, pp. 114–130, Jan. 2015.

[34] M. Ecker *et al.*, “Development of a lifetime prediction model for lithium-ion batteries based on extended accelerated aging test data,” *J. Power Sources*, vol. 215, pp. 248–257, Oct. 2012.

[35] B. Y. Liaw and M. Dubarry, “CHAPTER FIFTEEN - A Roadmap to Understand Battery Performance in Electric and Hybrid Vehicle Operation,” in *Electric and Hybrid Vehicles*, G. Pistoia, Ed. Amsterdam: Elsevier, 2010, pp. 375–403.

[36] S. F. Schuster, M. J. Brand, C. Campestrini, M. Gleissenberger, and A. Jossen, “Correlation between capacity and impedance of lithium-ion cells during calendar and cycle life,” *J. Power Sources*, vol. 305, pp. 191–199, Feb. 2016.

[37] C. Pastor-Fernández, T. Bruen, W. D. Widanage, M. A. Gama-Valdez, and J. Marco, “A Study of Cell-to-Cell Interactions and Degradation in Parallel Strings: Implications for the Battery Management System,” *J. Power Sources*, vol. 329, pp. 574–585, 2016.

[38] Y. Jiang, J. Jiang, C. Zhang, W. Zhang, Y. Gao, and Q. Guo, “Recognition of battery aging variations for LiFePO4 batteries in 2nd use applications combining incremental capacity analysis and statistical approaches,” *J. Power Sources*, vol. 360, pp. 180–188, 2017.

[39] A. Barré, B. Deguilhem, S. Grolleau, M. Gérard, F. Suard, and D. Riu, “A review on lithium-ion battery ageing mechanisms and estimations for automotive applications,” *J. Power Sources*, vol. 241, no. 0, pp. 680–689, 2013.

[40] N. Omar *et al.*, “Lithium iron phosphate based battery – Assessment of the aging parameters and development of cycle life model,” *Appl. Energy*, vol. 113, no. 0, pp. 1575–1585, Jan. 2014.

[41] M. Broussely, “CHAPTER THIRTEEN - Battery Requirements for HEVs, PHEVs, and EVs: An Overview,” in *Electric and Hybrid Vehicles*, G. Pistoia, Ed. Amsterdam: Elsevier, 2010, pp. 305–345.

[42] S. B. Peterson, J. Apt, and J. F. Whitacre, “Lithium-ion battery cell

degradation resulting from realistic vehicle and vehicle-to-grid utilization," *J. Power Sources*, vol. 195, no. 8, pp. 2385–2392, 2010.

[43] S. S. Choi and H. S. Lim, "Factors that affect cycle-life and possible degradation mechanisms of a Li-ion cell based on LiCoO2," *J. Power Sources*, vol. 111, no. 1, pp. 130–136, Sep. 2002.

[44] S. J. Harris, D. J. Harris, and C. Li, "Failure statistics for commercial lithium ion batteries: A study of 24 pouch cells," *J. Power Sources*, vol. 342, pp. 589–597, 2017.

[45] S. Rothgang, T. Baumhofer, and D. U. Sauer, "Diversion of Aging of Battery Cells in Automotive Systems," in *2014 IEEE Vehicle Power and Propulsion Conference (VPPC)*, 2014, pp. 1–6.

[46] M. Baumann, S. Rohr, and M. Lienkamp, "Development and Investigation of a modular stationary Second Life Storage System," *Conf. Futur. Automot. Technol.*, no. September, 2016.

[47] C. Campestrini, P. Keil, S. F. Schuster, and A. Jossen, "Ageing of lithium-ion battery modules with dissipative balancing compared with single-cell ageing," *J. Energy Storage*, vol. 6, pp. 142–152, 2016.

[48] Z. Zhang and B. Sisk, "Model-Based Analysis of Cell Balancing of Lithium-ion Batteries for Electric Vehicles," *SAE Int. J. Alt. Power.*, vol. 2, no. 2, pp. 379–388, 2013.

[49] V. Pop, *Battery Management Systems: Accurate State-of-Charge Indication for Battery-Powered Applications*. Springer, 2008.

[50] R. YAZAMI and Y. REYNIER, "Mechanism of self-discharge in graphite–lithium anode," *Electrochim. Acta*, vol. 47, no. 8, pp. 1217–1223, Feb. 2002.

[51] W. F. Bentley, "Cell balancing considerations for lithium-ion battery systems," in *Battery Conference on Applications and Advances, 1997., Twelfth Annual*, 1997, pp. 223–226.

[52] J. Cao, N. Schofield, and A. Emadi, "Battery balancing methods: A comprehensive review," in *Vehicle Power and Propulsion Conference, 2008. VPPC '08. IEEE*, 2008, pp. 1–6.

[53] T. Utsunomiya, O. Hatozaki, N. Yoshimoto, M. Egashira, and M. Morita, "Self-discharge behavior and its temperature dependence of carbon electrodes in lithium-ion batteries," *J. Power Sources*, vol. 196, no. 20, pp. 8598–8603, 2011.

[54] J. P. Schmidt, A. Weber, and E. Ivers-Tiffée, “A novel and fast method of characterizing the self-discharge behavior of lithium-ion cells using a pulse-measurement technique,” *J. Power Sources*, vol. 274, pp. 1231–1238, Jan. 2015.

[55] A. Seaman, T.-S. Dao, and J. McPhee, “A survey of mathematics-based equivalent-circuit and electrochemical battery models for hybrid and electric vehicle simulation,” *J. Power Sources*, vol. 256, pp. 410–423, Jun. 2014.

[56] N. A. Chaturvedi, R. Klein, J. Christensen, J. Ahmed, and A. Kojic, “Algorithms for Advanced Battery-Management Systems,” *Control Syst. IEEE*, vol. 30, no. 3, pp. 49–68, 2010.

[57] J. Gallardo-Lozano, E. Romero-Cadaval, M. I. Milanes-Montero, and M. A. Guerrero-Martinez, “A novel active battery equalization control with on-line unhealthy cell detection and cell change decision,” *J. Power Sources*, vol. 299, pp. 356–370, Dec. 2015.

[58] L. Lu, X. Han, J. Li, J. Hua, and M. Ouyang, “A review on the key issues for lithium-ion battery management in electric vehicles,” *J. Power Sources*, vol. 226, no. 0, pp. 272–288, 2013.

[59] A. F. Burke, “Batteries and Ultracapacitors for Electric, Hybrid, and Fuel Cell Vehicles,” *Proc. IEEE*, vol. 95, no. 4, pp. 806–820, 2007.

[60] J. Hughes, “Pics and Info: Inside the Tesla 100kWh Battery Pack,” 2017. [Online]. Available: http://skie.net/skynet/projects/tesla/view_post/20_Pics+and+Info:+Inside+the+Tesla+100kWh+Battery+Pack. [Accessed: 28-Aug-2017].

[61] A. Farmann and D. U. Sauer, “A comprehensive review of on-board State-of-Available-Power prediction techniques for lithium-ion batteries in electric vehicles,” *J. Power Sources*, vol. 329, pp. 123–137, 2016.

[62] Y. Troxler *et al.*, “The effect of thermal gradients on the performance of lithium-ion batteries,” *J. Power Sources*, vol. 247, no. 0, pp. 1018–1025, Feb. 2014.

[63] C. Fleischer, W. Waag, H.-M. Heyn, and D. U. Sauer, “On-line adaptive battery impedance parameter and state estimation considering physical principles in reduced order equivalent circuit battery models. Part 1. Requirements, critical review of methods and modeling,” *J. Power Sources*,

vol. 260, pp. 276–291, Aug. 2014.

[64] D. Andre, M. Meiler, K. Steiner, C. Wimmer, T. Soczka-guth, and D. U. Sauer, “Characterization of high-power lithium-ion batteries by electrochemical impedance spectroscopy. I. Experimental investigation,” *J. Power Sources*, vol. 196, no. 12, pp. 5334–5341, Jun. 2011.

[65] W. Waag, S. Käbitz, and D. U. Sauer, “Experimental investigation of the lithium-ion battery impedance characteristic at various conditions and aging states and its influence on the application,” *Appl. Energy*, vol. 102, no. 0, pp. 885–897, 2013.

[66] Z. J. Zhang and P. Ramadass, “Lithium-Ion Battery Systems and Technology,” in *Batteries for Sustainability*, R. J. Brodd, Ed. New York, NY: Springer New York, 2013, pp. 319–357.

[67] L. Zhong, C. Zhang, Y. He, and Z. Chen, “A method for the estimation of the battery pack state of charge based on in-pack cells uniformity analysis,” *Appl. Energy*, vol. 113, no. 0, pp. 558–564, 2014.

[68] M. J. Brand, M. H. Hofmann, M. Steinhardt, S. F. Schuster, and A. Jossen, “Current distribution within parallel-connected battery cells,” *J. Power Sources*, vol. 334, pp. 202–212, 2016.

[69] M. Dubarry, A. Devie, and B. Y. Liaw, “Cell-balancing currents in parallel strings of a battery system,” *J. Power Sources*, vol. 321, pp. 36–46, 2016.

[70] M. P. Klein and J. W. Park, “Current Distribution Measurements in Parallel-Connected Lithium-Ion Cylindrical Cells under Non-Uniform Temperature Conditions,” *J. Electrochem. Soc.*, vol. 164, no. 9, pp. A1893–A1906, 2017.

[71] L. Cheng and Q. Sun, “Discharge Capacity Estimation for Lithium-Ion Battery Packs with Cells in Parallel Connection Based on Current Prediction of In-Pack Cells,” *Energy Technol.*, vol. 410000, pp. 1–8, 2017.

[72] A. Baughman and M. Ferdowsi, “Battery Charge Equalization-State of the Art and Future Trends,” in *Future Transportation Technology Conferencee*, 2005.

[73] M. Daowd, N. Omar, P. Van den Bossche, and J. Van Mierlo, “Passive and active battery balancing comparison based on MATLAB simulation,” in *Vehicle Power and Propulsion Conference (VPPC), 2011 IEEE*, 2011, pp. 1–7.

[74] S. W. Moore and P. J. Schneider, “A Review of Cell Equalization Methods

for Lithium Ion and Lithium Polymer Battery Systems," *SAE Publ.*, vol. 2001010959, no. 724, pp. 1–7, Mar. 2001.

[75] G. L. Plett, "Sigma-point Kalman filtering for battery management systems of LiPB-based HEV battery packs. Part 1: Introduction and state estimation," *J. Power Sources*, vol. 161, no. 2, pp. 1356–1368, Oct. 2006.

[76] B. Fridholm, T. Wik, and M. Nilsson, "Robust recursive impedance estimation for automotive lithium-ion batteries," *J. Power Sources*, vol. 304, pp. 33–41, Feb. 2016.

[77] C. D. Rahn and C. Y. Wang, *Battery Systems Engineering*. Wiley, 2013.

[78] S. Tong, J. H. Lacap, and J. W. Park, "Battery state of charge estimation using a load-classifying neural network," *J. Energy Storage*, vol. 7, pp. 236–243, 2016.

[79] T. Hansen and C. J. Wang, "Support vector based battery state of charge estimator," *J. Power Sources*, vol. 141, no. 2, pp. 351–358, 2005.

[80] W. He, N. Williard, C. Chen, and M. Pecht, "State of charge estimation for electric vehicle batteries using unscented kalman filtering," *Microelectron. Reliab.*, vol. 53, no. 6, pp. 840–847, Jun. 2013.

[81] W. Zhang, W. Shi, and Z. Ma, "Adaptive unscented Kalman filter based state of energy and power capability estimation approach for lithium-ion battery," *J. Power Sources*, vol. 289, pp. 50–62, Sep. 2015.

[82] D. Simon, *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*. Wiley-Interscience, 2006.

[83] G. L. Plett, "Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 1. Background," *J. Power Sources*, vol. 134, no. 2, pp. 252–261, 2004.

[84] J. Kalawoun, K. Biletska, F. Suard, and M. Montaru, "From a novel classification of the battery state of charge estimators toward a conception of an ideal one," *J. Power Sources*, vol. 279, pp. 694–706, Apr. 2015.

[85] S. Steinhorst and M. Lukasiewycz, "Formal approaches to design of active cell balancing architectures in Battery Management Systems," *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. pp. 1–6, 2016.

[86] S. Ci, N. Lin, and D. Wu, "Reconfigurable Battery Techniques and Systems: A Survey," *IEEE Access*, vol. 4, pp. 1175–1189, 2016.

[87] R. Ratz, “BMS System Benchmark and Standardization,” 2015.

[88] L. T. Corporation, “LTC3300-1 - High Efficiency Bidirectional Multicell Battery Balancer (Datasheet),” 2013.

[89] X. Tang, L. Hao, X. MAO, and J. S. Chen, “Cell balancing system and method.” Google Patents, 2013.

[90] L. Hu, M. L. Zhao, X. B. Wu, and J. N. Lou, “Cell balancing management for battery pack,” *2010 10th IEEE International Conference on Solid-State and Integrated Circuit Technology*. pp. 339–341, 2010.

[91] I. Arasaratnam, J. Tjong, and S. Habibi, “Switched-Capacitor Cell Balancing: A Fresh Perspective.” SAE International, 2014.

[92] C. Sheng *et al.*, “Implementation of cell balancing with super-capacitor for robot power system,” in *Intelligent Control and Automation (WCICA), 2011 9th World Congress on*, 2011, pp. 468–473.

[93] M. Preindl, C. Danielson, and F. Borrelli, “Performance evaluation of battery balancing hardware,” *Control Conf. (ECC), …*, pp. 4065–4070, 2013.

[94] M.-Y. Kim, C.-H. Kim, J.-H. Kim, and G.-W. Moon, “A Chain Structure of Switched Capacitor for Improved Cell Balancing Speed of Lithium-Ion Batteries,” *Industrial Electronics, IEEE Transactions on*, vol. 61, no. 8. pp. 3989–3999, 2014.

[95] C. Bonfiglio and W. Roessler, “A cost optimized battery management system with active cell balancing for lithium ion battery stacks,” in *Vehicle Power and Propulsion Conference, 2009. VPPC ’09. IEEE*, 2009, pp. 304–309.

[96] A. Imtiaz and F. Khan, “‘Time Shared Flyback Converter’ Based Regenerative Cell Balancing Technique for Series Connected Li-Ion Battery Strings,” *Power Electron. IEEE Trans.*, vol. 28, no. 12, pp. 5960–5975, 2013.

[97] M. Kim, J.-W. Kim, C.-H. Kim, S.-Y. Cho, and G.-W. Moon, “Automatic charge equalization circuit based on regulated voltage source for series connected lithium-ion batteries,” *Power Electronics and ECCE Asia (ICPE & ECCE), 2011 IEEE 8th International Conference on*. pp. 2248–2255, 2011.

[98] W. Liye *et al.*, “Research on battery balance system applied on HEV,” in

Vehicle Power and Propulsion Conference, 2009. VPPC '09. IEEE, 2009, no. 1, pp. 1788–1791.

[99] C. Lin, C. Wang, L. Chang-Hua, C. Hsuan-Yi, W. Chien-Ming, and H. Min-Hsuan, "Battery management system with dual-balancing mechanism for LiFePO<inf>4</inf> battery module," in *TENCON 2011 - 2011 IEEE Region 10 Conference*, 2011, pp. 863–867.

[100] Z. Zhang, H. Gui, D. J. Gu, Y. Yang, and X. Ren, "A Hierarchical Active Balancing Architecture for Lithium-Ion Batteries," *IEEE Transactions on Power Electronics*, vol. 32, no. 4. pp. 2757–2768, 2017.

[101] M. M. Ur Rehman, F. Zhang, M. Evzelman, R. Zane, and D. Maksimovic, "Control of a series-input, parallel-output cell balancing system for electric vehicle battery packs," *2015 IEEE 16th Work. Control Model. Power Electron. COMPEL 2015*, 2015.

[102] F. Ran, H. Xu, Y. Ji, J. Qin, and W. Li, "An active balancing circuit for lithium battery management system with optoelectronic switches," *IEEE Reg. 10 Annu. Int. Conf. Proceedings/TENCON*, vol. 2016–Janua, 2016.

[103] P. Prakash and J. K. Kishore, "An active cell balancing algorithm for Li-Ion battery and its simulations for GEO satellite missions," *2016 IEEE 1st International Conference on Power Electronics, Intelligent Control and Energy Systems (ICPEICES)*. pp. 1–6, 2016.

[104] M. Lukasiewycz, M. Kauer, and S. Steinhorst, "Synthesis of Active Cell Balancing Architectures for Battery Packs," *IEEE Trans. Comput. Des. Integr. Circuits Syst.*, vol. 35, no. 11, pp. 1876–1889, 2016.

[105] M. Kauer, S. Narayanaswamy, S. Steinhorst, M. Lukasiewycz, and S. Chakraborty, "Many-to-many active cell balancing strategy design," *20th Asia South Pacific Des. Autom. Conf. ASP-DAC 2015*, pp. 267–272, 2015.

[106] Y. Zheng, M. Ouyang, L. Lu, J. Li, X. Han, and L. Xu, "On-line equalization for lithium-ion battery packs based on charging cell voltages: Part 1. Equalization based on remaining charging capacity estimation," *J. Power Sources*, vol. 247, no. 0, pp. 676–686, 2014.

[107] Y. Zheng, M. Ouyang, L. Lu, J. Li, X. Han, and L. Xu, "On-line equalization for lithium-ion battery packs based on charging cell voltages: Part 2. Fuzzy logic equalization," *J. Power Sources*, vol. 247, no. 0, pp. 460–466, 2014.

[108] C. Danielson, F. Borrelli, D. Oliver, D. Anderson, K. Ming, and T. Phillips,

"Balancing of battery networks via constrained optimal control," in *American Control Conference (ACC), 2012*, 2012, pp. 4293–4298.

[109] D. D. Quinn and T. T. Hartley, "Design of novel charge balancing networks in battery packs," *J. Power Sources*, vol. 240, no. 0, pp. 26–32, 2013.

[110] Y.-S. Lee and M.-W. Cheng, "Intelligent control battery equalization for series connected lithium-ion battery strings," *Industrial Electronics, IEEE Transactions on*, vol. 52, no. 5. pp. 1297–1307, 2005.

[111] M. F. Samadi and M. Saif, "Nonlinear model predictive control for cell balancing in Li-ion battery packs," *American Control Conference (ACC), 2014*. pp. 2924–2929, 2014.

[112] L. McCurlie, M. Preindl, and A. Emadi, "Fast Model Predictive Control for Redistributive Lithium Ion Battery Balancing," *IEEE Trans. Ind. Electron.*, vol. 46, no. c, pp. 1–1, 2016.

[113] S. J. Moura, J. C. Forman, S. Bashash, J. L. Stein, and H. K. Fathy, "Optimal Control of Film Growth in Lithium-Ion Battery Packs via Relay Switches," *Ind. Electron. IEEE Trans.*, vol. 58, no. 8, pp. 3555–3566, 2011.

[114] X. Hu, S. Li, and H. Peng, "A comparative study of equivalent circuit models for Li-ion batteries," *J. Power Sources*, vol. 198, pp. 359–367, Jan. 2012.

[115] J. C. Forman, S. J. Moura, J. L. Stein, and H. K. Fathy, "Genetic parameter identification of the Doyle-Fuller-Newman model from experimental cycling of a LiFePO<inf>4</inf> battery," *Proc. 2011 Am. Control Conf.*, pp. 362–369, 2011.

[116] A. M. Bizeray, S. Zhao, S. R. Duncan, and D. A. Howey, "Lithium-ion battery thermal-electrochemical model-based state estimation using orthogonal collocation and a modified extended Kalman filter," *J. Power Sources*, vol. 296, pp. 400–412, Nov. 2015.

[117] K. D. Stetzel, L. L. Aldrich, M. S. Trimboli, and G. L. Plett, "Electrochemical state and internal variables estimation using a reduced-order physics-based model of a lithium-ion cell and an extended Kalman filter," *J. Power Sources*, vol. 278, pp. 490–505, Mar. 2015.

[118] C. D. Rahn and C.-Y. Wang, "System Response," in *Battery Systems Engineering*, Wiley, 2013, pp. 89–118.

[119] W. D. Widanage *et al.*, "Design and use of multisine signals for Li-ion battery equivalent circuit modelling. Part 2: Model estimation," *J. Power*

Sources, vol. 324, pp. 61–69, 2016.

[120] T. R. B. Grandjean, A. McGordon, and P. A. Jennings, “Structural Identifiability of Equivalent Circuit Models for Li-Ion Batteries,” *Energies*, vol. 10, no. 1, 2017.

[121] B. Pattipati, B. Balasingam, G. V. Avvari, K. R. Pattipati, and Y. Bar-Shalom, “Open Circuit Voltage Characterization of Lithium-ion Batteries,” *J. Power Sources*, vol. 269, pp. 317–333, Jul. 2014.

[122] G. L. Plett, “Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 2. Modeling and identification,” *J. Power Sources*, vol. 134, no. 2, pp. 277–292, Aug. 2004.

[123] A. Jossen, “Fundamentals of battery dynamics,” *J. Power Sources*, vol. 154, no. 2, pp. 530–538, Mar. 2006.

[124] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*. John Wiley & Sons, 2005.

[125] X. Lin and A. G. Stefanopoulou, “Analytic Bound on Accuracy of Battery State and Parameter Estimation,” *J. Electrochem. Soc.*, vol. 162, no. 9, pp. A1879–A1891, Jan. 2015.

[126] W. Dreyer, J. Jamnik, C. Guhlke, R. Huth, J. Moskon, and M. Gaberscek, “The thermodynamic origin of hysteresis in insertion batteries,” *Nat Mater*, vol. 9, no. 5, pp. 448–453, May 2010.

[127] M. Dubarry *et al.*, “Evaluation of commercial lithium-ion cells based on composite positive electrode for plug-in hybrid electric vehicle applications. Part II. Degradation mechanism under 2C cycle aging,” *J. Power Sources*, vol. 196, no. 23, pp. 10336–10343, Dec. 2011.

[128] M. Dubarry, V. Svoboda, R. Hwu, and B. Y. Liaw, “Capacity loss in rechargeable lithium cells during cycle life testing: The importance of determining state-of-charge,” *J. Power Sources*, vol. 174, no. 2, pp. 1121–1125, Dec. 2007.

[129] W. D. Widanage *et al.*, “Design and use of multisine signals for Li-ion battery equivalent circuit modelling. Part 1: Signal design,” *J. Power Sources*, vol. 324, pp. 70–78, 2016.

[130] K. Madsen, H. B. Nielsen, and O. Tingleff, “Methods for non-linear least squares problems,” in *Infomatics and Mathematical Modeling*, vol. 2, 2004, pp. 1–30.

[131] T. H. Phung, A. Collet, and J.-C. Crebier, “An Optimized Topology for Next-to-Next Balancing of Series-Connected Lithium-ion Cells,” *Power Electronics, IEEE Transactions on*, vol. 29, no. 9. pp. 4603–4613, 2014.

[132] X. Gong, R. Xiong, and C. C. Mi, “Study of the characteristics of battery packs in electric vehicles with parallel-connected lithium-ion battery cells,” *Applied Power Electronics Conference and Exposition (APEC), 2014 Twenty-Ninth Annual IEEE*. pp. 3218–3224, 2014.

[133] X. Zhang and C. Mi, “Management of Energy Storage Systems in EV, HEV and PHEV,” in *Vehicle Power Management*, Springer London, 2011, pp. 259–286.

[134] U. D. of Energy, “Advanced Vehicle Testing – Beginning-of-Test Battery Testing Results,” 2011. [Online]. Available: http://www1.eere.energy.gov/vehiclesandfuels/avta/pdfs/fsev/battery_leaf_0356.pdf. [Accessed: 22-Oct-2014].

[135] R. Gogoana, M. B. Pinson, M. Z. Bazant, and S. E. Sarma, “Internal Resistance Matching for Parallel-Connected Lithium-Ion Cells and Impacts on Battery Pack Cycle Life,” *J. Power Sources*, vol. 252, no. 0, pp. 8–13, 2014.

[136] Bitrode, “Datasheet: MCV Battery Cell Test System.” St Louis, p. 2, 2017.

[137] D. Widanalage, “Summary on Equivalent Circuit and Reference Equivalent Circuit Model Usage , Structure and Validation,” 2013.

[138] M. Debert, G. Colin, G. Bloch, and Y. Chamaillard, “An observer looks at the cell temperature in automotive battery packs,” *Control Eng. Pract.*, vol. 21, no. 8, pp. 1035–1042, Aug. 2013.

[139] W. C. Lee, D. Drury, P. Mellor, and L. Wai Chung, “Comparison of passive cell balancing and active cell balancing for automotive batteries,” in *Vehicle Power and Propulsion Conference (VPPC), 2011 IEEE*, 2011, pp. 1–7.

[140] C. Truchot, M. Dubarry, and B. Y. Liaw, “State-of-charge estimation and uncertainty for lithium-ion battery strings,” *Appl. Energy*, vol. 119, pp. 218–227, Apr. 2014.

[141] J. Li, J. Klee Barillas, C. Guenther, and M. A. Danzer, “Multicell state estimation using variation based sequential Monte Carlo filter for automotive battery packs,” *J. Power Sources*, vol. 277, pp. 95–103, Mar. 2015.

[142] J. Li and M. S. Mazzola, “Accurate battery pack modeling for automotive applications,” *J. Power Sources*, vol. 237, no. 0, pp. 215–228, Sep. 2013.

[143] J. Taylor, R. Ball, A. Mcgordon, K. Uddin, and J. Marco, “Sizing tool for rapid optimisation of pack configuration at early-stage automotive product development,” in *EVS28 International Electric Vehicle Symposium and Exhibition*, 2015, vol. 28, no. May, pp. 1–8.

[144] Z. Hanlei and C. Mo-Yuen, “Comprehensive dynamic battery modeling for PHEV applications,” in *Power and Energy Society General Meeting, 2010 IEEE*, 2010, pp. 1–6.

[145] T. Instruments, “EM1401EVM Balancing Evaluation Board.” 2014.

[146] M. T. Lawder, P. W. C. Northrop, and V. R. Subramanian, “Model-Based SEI Layer Growth and Capacity Fade Analysis for EV and PHEV Batteries and Drive Cycles,” *J. Electrochem. Soc.*, vol. 161, no. 14, pp. A2099–A2108, Jan. 2014.

[147] K. B. Petersen and M. S. Pedersen, “Matrix Cookbook.” 2012.

[148] W. C. A. Kishore, S. Sen, and G. Ray, “Disturbance Rejection and Control Allocation of Over-Actuated systems H,” *Industrial Technology, 2006. ICIT 2006. IEEE International Conference on*. pp. 1054–1059, 2006.

[149] T. A. Johansen and T. I. Fossen, “Control allocation—A survey,” *Automatica*, vol. 49, no. 5, pp. 1087–1103, May 2013.

[150] A. Tewari, *Modern Control Design: With MATLAB and SIMULINK*. Wiley, 2002.

[151] K. J. Astrom and R. M. Murray, “State Feedback,” in *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed., Princeton University Press, 2012, p. 283.

[152] J. Gadewadikar, F. L. Lewis, and M. Abu-Khalaf, “Necessary and Sufficient Conditions for H-Infinity Static Output-Feedback Control,” *J. Guid. Control. Dyn.*, vol. 29, no. 4, pp. 915–920, Jul. 2006.

[153] J. M. Maciejowski, *Multivariable Feedback Design*. Addison-Wesley, 1989.

[154] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory*, vol. 134, no. 6. 1990.

[155] T. MathWorks, “Pole Placement Design,” 2016. [Online]. Available: https://uk.mathworks.com/help/control/ref/place.html. [Accessed: 11-Jul-2016].

[156] “Pole Placement Design - MathWorks UK,” 2016. [Online]. Available: http://uk.mathworks.com/help/control/ref/place.html. [Accessed: 11-Jul-2016].

[157] E. F. Camacho and C. Bordons, “Introduction to Model Predictive Control,” in *Model Predictive control*, London: Springer London, 2007, pp. 1–11.

[158] J. C. Ferreira, V. Monteiro, and J. L. Afonso, “Dynamic range prediction for an electric vehicle,” *2013 World Electr. Veh. Symp. Exhib. EVS 2014*, pp. 1–11, 2014.

[159] K. D. Young, V. I. Utkin, and Ü. Özgüner, “A control engineer’s guide to sliding mode control,” *IEEE Trans. Control Syst. Technol.*, vol. 7, no. 3, pp. 328–342, 1999.

[160] C. Campestrini, S. Kosch, and A. Jossen, “Influence of change in open circuit voltage on the state of charge estimation with an extended Kalman filter,” *J. Energy Storage*, vol. 12, pp. 149–156, 2017.

[161] M. J. Brand, P. A. Schmidt, M. F. Zaeh, and A. Jossen, “Welding techniques for battery cells and resulting electrical contact resistances,” *J. Energy Storage*, vol. 1, no. 1, pp. 7–14, 2015.

[162] M. J. Brand, E. I. Kolp, P. Berg, T. Bach, P. Schmidt, and A. Jossen, “Electrical resistances of soldered battery cell connections,” *J. Energy Storage*, vol. 12, pp. 45–54, 2017.

[163] National Instruments, “User Guide and Specifications: NI USB-9219,” 2007.

[164] LEM, “Datasheet: LEM CASR Series Current Transducers.” 2012.

[165] A. Devie, E. Vinot, S. Pelissier, and P. Venet, “Real-world battery duty profile of a neighbourhood electric vehicle,” *Transp. Res. Part C Emerg. Technol.*, vol. 25, no. 0, pp. 122–133, 2012.

[166] C. Glaize and S. Geniès, “Definitions and Measuring Methods,” in *Lithium Batteries and Other Electrochemical Storage Systems*, John Wiley & Sons, Inc., 2013, pp. 21–87.

[167] C. R. Birkl and D. A. Howey, “Model Identification and Parameter Estimation for LiFePO4 Batteries,” *Hybrid Electr. Veh. Conf. 2013 (HEVC 2013)*, pp. 1–6, 2013.

[168] D. A. Howey, P. D. Mitcheson, V. Yufit, G. J. Offer, and N. P. Brandon, “Online Measurement of Battery Impedance Using Motor Controller

Excitation," *Vehicular Technology, IEEE Transactions on*, vol. 63, no. 6. pp. 2557–2566, 2014.

[169] A. Barai, G. H. Chouchelamane, Y. Guo, A. McGordon, and P. Jennings, "A study on the impact of lithium-ion cell relaxation on electrochemical impedance spectroscopy," *J. Power Sources*, vol. 280, pp. 74–80, Apr. 2015.

[170] E. Meissner and G. Richter, "Battery Monitoring and Electrical Energy Management: Precondition for future vehicle electric power systems," *J. Power Sources*, vol. 116, no. 1–2, pp. 79–98, 2003.

[171] B. Wang, S. E. Li, H. Peng, and Z. Liu, "Fractional-order modeling and parameter identification for lithium-ion batteries," *J. Power Sources*, vol. 293, pp. 151–161, Oct. 2015.

[172] J. Sabatier, M. Merveillaut, J. M. Francisco, F. Guillemard, and D. Porcelatto, "Lithium-ion batteries modeling involving fractional differentiation," *J. Power Sources*, vol. 262, pp. 36–43, Sep. 2014.

[173] J. P. Schmidt, P. Berg, M. Schönleber, A. Weber, and E. Ivers-Tiffée, "The distribution of relaxation times as basis for generalized time-domain models for Li-ion batteries," *J. Power Sources*, vol. 221, no. 0, pp. 70–77, 2013.

Appendix A: Cell Impedance Modelling

Frequency domain data are generally obtained using electrochemical impedance spectroscopy (EIS) equipment. This generates a sinusoidal current signal[15] and measures the amplitude and phase difference between the current and voltage to calculate the impedance at that frequency. Signals over a range of frequencies (10mHz to 100KHz is typical) are applied to build a thorough picture of the cell's response. This wide frequency range enables several distinct electrochemical phenomena to be discerned. A cylindrical 3Ah lithium-ion cell was used for a variety of test work during this research, notably in Submission 2 and [10], and both time and frequency domain fits were performed. Figure A- 1 shows a Nyquist plot of the cell's impedance. A Solartron ModuLab-XM was used to perform the EIS test within the laboratory. Frequencies from 10mHz to 100kHz were applied to a partially aged cylindrical battery, with 10 logarithmically spaced frequencies per decade. A C/20 current was applied at 90% SOC and a VC3 4060 thermal chamber used to maintain the ambient temperature at 25°C.

The impedance can be broken down into four primary contributions [64, 123, 166] which are also annotated on Figure A- 1:

1. Ohmic drop: the total DC resistance of the electrodes, electrolyte, separators and electrical connection of the cell to the external circuitry. There is no reactance so the effect is instantaneous. For the cell in Figure A- 1, this occurs at 600Hz.
2. Charge transfer: due to the double layer capacitance between the electrode-electrolyte interfaces, ions are moved due to gradients in the electric field. The time constant for this is typically in the millisecond-second region. For aged cells, this process can occur in a second region. A solid electrolyte interphase (SEI) layer develops on the carbon anode which increases in thickness as the cell ages. This adds an additional impedance which is distinct from the charge transfer process.

[15] A sinusoidal voltage (potentiostatic operation) can also be applied. However, a sinusoidal current (galvanostatic operation) is typically used for cells because it ensures the overall current into the cell is zero so the cell SOC does not change over the duration of the test.

3. Diffusion/ mass transfer: concentration gradients drive the diffusion of lithium ions from the inside of each electrode to the surface, and vice versa. The time constant for this is typically in the seconds-minutes region. This is the main response that an ECM running on a BMS aims to capture.
4. Induction: the components of the cell, and wiring to the external circuit all have an inductance. This is general small and so can be neglected until the frequency is sufficiently high (generally over 1kHz). Since the bandwidth of current through cells, and the BMS sampling rate, do not typically reach close to such high frequencies, it is often ignored [167]. However, there are two sources of high frequencies in battery packs: noise from the motor inverter and switching from certain cell balancing systems.

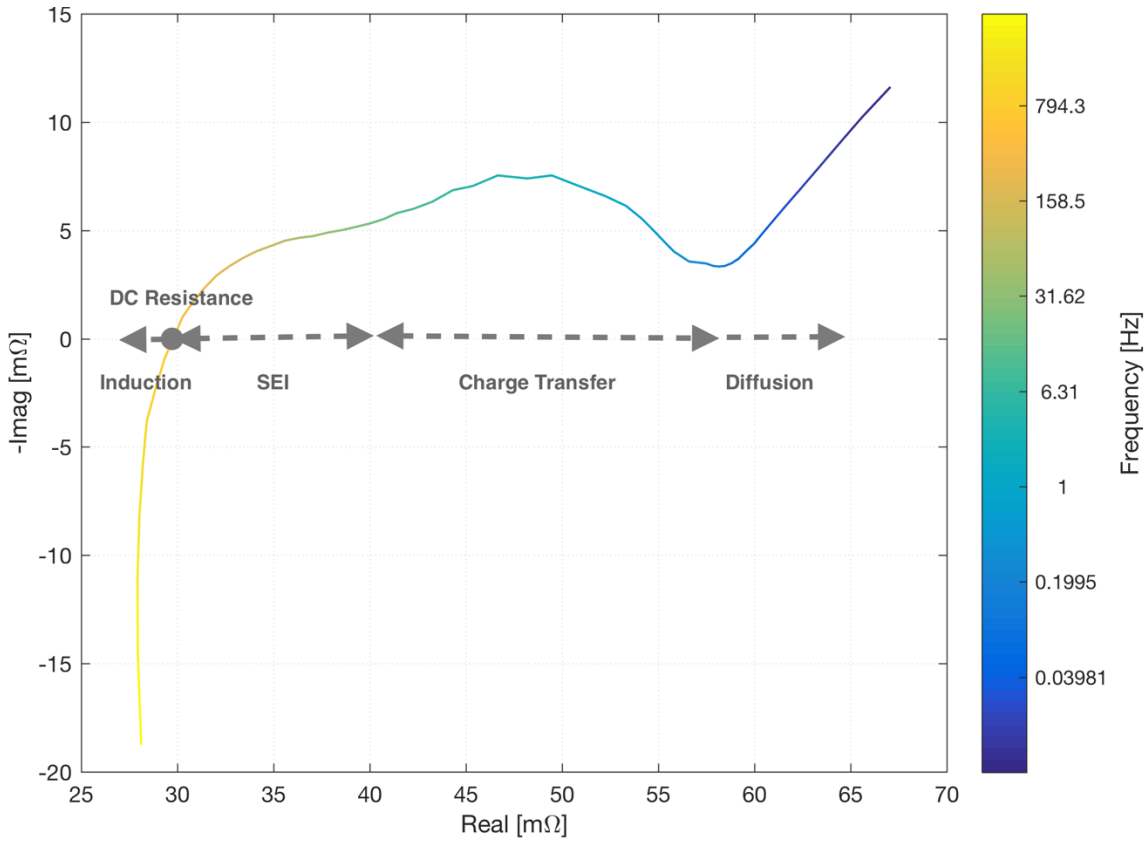

Figure A- 1: Example EIS data of a cylindrical lithium-ion cell

This level of granularity makes EIS particularly useful for comparing a cell at various states, for example, how a new cell compared to one aged for two years, or how it performs at -25°C against 45°C. However, there are several drawbacks to EIS. A low-amplitude sinusoidal current is not representative of how cells are used in any application, particularly an EV. EIS requires specialist circuitry and

so is not suited to operation on board a vehicle. EIS using a motor controller was proposed in [168], but has a limited frequency range and requires accurate calibration, which may not be possible to maintain over the lifetime of a vehicle. In order to obtain the impedance data, the cells must be have been at rest for several hours, which limits when the tests can be performed. The test itself can become time consuming if low frequencies (below 10mHz) are excited, which is an area of interest for BMS applications. As measurements cannot occur while the cells are in use, it is not possible to measure the impedance at any instance in time, limiting its use for BMS applications.

Equivalent circuits for EIS data typically preserve the four regions of operation described above. This enables changes in impedance from one test to another to be quantified and linked to a particular process (such an increase in DC resistance or a slow-down in diffusion). Figure A- 2 shows a commonly used ECM for EIS impedance data [118, 169, 170]**.** The Warburg element σ models the diffusion, and the constant phase element (CPE_{dl}) is a capacitor-like component which captures the charge transfer reactance. This circuit is best considered in the complex frequency domain, where complex frequency s is defined by (97) from frequency f. The circuit has the impedance equation (98), where n and w are {0,1} values for the constant phase element and Warburg element respectively. An additional R-CPE pairs can be added in series with this model to factor in additional SEI growth which occurs with ageing. The equation for this pair is given by (99).

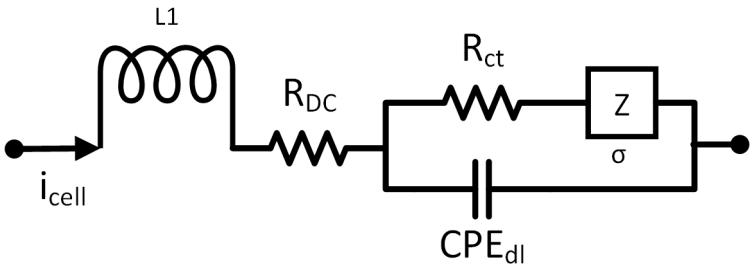

Figure A- 2: ECM for EIS data

$$s = 2\pi f j \tag{97}$$

$$Z(s) = Ls + R_{DC} + \frac{R_{ct}s^n + \sigma}{s^n + s^w CPE_{dl}(R_{ct}s^n + \sigma)} \tag{98}$$

$$Z_{sei}(s) = \frac{R_{sei}s^{n_{sei}}}{1 + R_{sei}C_{sei}s^{n_{sei}}} \quad (99)$$

These equations are problematic because they cannot easily be translated into the time domain: it results in a fractional-order system (100), which is more difficult to use for time domain simulation, as well as estimation and control. Methods have been proposed, such as [171, 172], but these add complexity. These fractional order components can be approximated using several conventional resistor-capacitor (RC) pairs [173], which form the basis for most time-domain models.

$$s^n \rightarrow \frac{d}{dt^n} \quad (100)$$

In summary, EIS data can provide useful insight about the behaviour of cells but is less relevant to time-domain simulation and for on-vehicle applications. The time required for testing can be prohibitive, especially when the low frequency impedance is required. Another approach is to collect time-domain data directly.

The NLLS parameterisation method discussed in section 3.4 can also be applied to EIS data. For frequency domain models, the measurement vector (and correspondingly what the function *f* should approximate) is defined by (101), and is the impedance measured using EIS equipment. The real and imaginary components of the impedance are separated to ensure both parts are accounted for equally in the fit (if a single value such as the absolute of the impedance is used, a poor fit could occur because several combinations of imaginary and real values can give the same absolute value).

$$y_i = [\Re(Z_i) \quad \Im(Z_i)]^T \quad (101)$$

The same cell data as in Figure A- 1 was used to fit the model from equation (98), including the SEI layer impedance from (99).The experimental data and impedance estimated from the model in additional R-CPE pair) are shown in Figure A- 3. As can be seen the model fits the data well, and has a mean error of 0.4mΩ. The resultant parameter vector is shown in Table A- 1. As expected,

the Warburg exponent is close to 0.5, but of note is the charge transfer exponent is less than 0.5, making it closer in behaviour to a resistor than a true capacitor. This can be seen in Figure A- 1 around the 10Hz region, where the impedance is a flattened arc rather than a semi-circle, as a true capacitor would create. The SEI exponent is 0.84 which is closer to a true capacitor. This demonstrates that the time-domain usage of true capacitors will does not represent the actual cell behaviour. Each of these exponents can be approximated using a series of RC pairs in series, but given that there are three CPEs in the impedance model, this would result in a large number of RC pairs which as discussed becomes difficult to parameterise in the time domain.

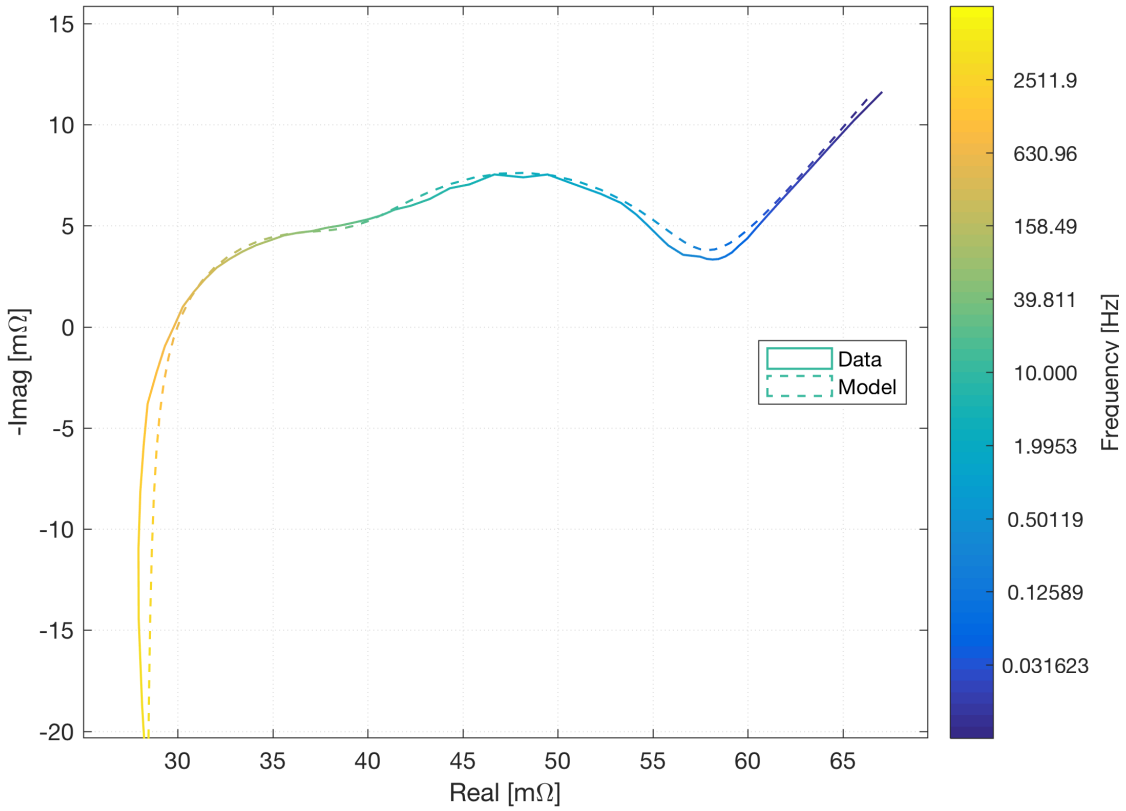

Figure A- 3: Example EIS results and model fit

Table A- 1: Parameter vector for EIS model fit

Parameter	L (μH)	R_{DC} (mΩ)	R_{sei} (mΩ)	C_{sei} (F)	n_{sei}	R_1 (mΩ)	C_1 (F)	w	σ (mΩ)	n
Value	0.77	28.2	16.8	4.72	0.84	11.5	0.45	0.79	3.38	0.54

Appendix B: Test Procedures

Several procedures were developed internally with the Catapult working group for cell testing. This appendix describes the relevant parts of these procedures so that they can be explained and replicated.

A. Cell Capacity

B. Place the cell in a thermal chamber set to 25°C, and leave for 8 hours to ensure the cell reaches the chamber temperature
C. Fully discharge the cell by discharging at 1C until the minimum voltage specified by the manufacturer is reached.
D. Let the cell relax for 3 hours, so the terminal voltage equals the OCV
E. Charge the cell using a CC-CV scheme according the manufacturer's specifications
F. Let the cell relax for 3 hours
G. Discharge at 1C until the minimum voltage specified by the manufacturer is reached

The cell capacity is calculated by integrating the current measured during step 6 over the time taken.

B. Pulse Power

1. Place the cell in a thermal chamber set to 25°C, and leave for 8 hours to ensure the cell reaches the chamber temperature
2. Fully discharge the cell by discharging at 1C until the minimum voltage specified by the manufacturer is reached.
3. Let the cell relax for 3 hours, so the terminal voltage equals the OCV
4. Using the previously calculated cell capacity, set the cell cycler to discharge at 1C until 10% of the cell's charge has been depleted (i.e. the cell is at 90% SOC)
5. Let the cell relax for 3 hours
6. Repeat the steps below for five current values, at 20%, 40%, 60%, 80% and 100% of the maximum specified by the manufacturer.
    a. Apply the current as a constant discharge for 10 seconds
    b. Let the cell relax for 30 minutes
    c. Apply the same current as a constant charge for 10 seconds
    d. Let the cell relax for 30 minutes
7. Repeat steps 4-6, but adjusting the cell to 50% SOC
8. Repeat steps 4-6, but adjusting the cell to 20% SOC

For each set of pulse tests (a set being the currents in the charging or discharging direction), the voltage difference between the start and end of the pulse (Δv_{pp}) is calculated for each current (i_{pp}). These values can be applied to (102) to perform

a linear regression (essentially Ohm's Law) to find an internal resistance value. In the ideal case R_{pp}.

$$\boldsymbol{\Delta v_{pp}} = R_{pp}\boldsymbol{i_{pp}} + o_{pp} \quad (102)$$

C. OCV-SOC (GITT)

1. Place the cell in a thermal chamber set to 25°C, and leave for 8 hours to ensure the cell reaches the chamber temperature
2. Fully discharge the cell by discharging at 1C until the minimum voltage specified by the manufacturer is reached.
3. Let the cell relax for 3 hours, so the terminal voltage equals the OCV
4. Discharge at 1C until 3% of the cell's charge has been depleted (i.e. the cell is at 97% SOC), calculated using the cell's capacity
5. Let the cell relax for 3 hours
6. Repeat steps 4-5, decreasing the SOC by 3% each time, down to 1% SOC.
7. Discharge the cell at 1C until the minimum voltage has been reached.
8. Let the cell relax for 3 hours
9. Charge at 1C until 3% of the cell's charge has been added (i.e. the cell is at 3% SOC), calculated using the cell's capacity.
10. Let the cell relax for 3 hours
11. Repeat steps 9-10, increasing the SOC by 3% each time, up to 99% SOC

This provides discrete data for separate charge an discharge curves.

D. OCV-SOC (Slow Discharge)

1. Place the cell in a thermal chamber set to 25°C, and leave for 8 hours to ensure the cell reaches the chamber temperature
2. Fully discharge the cell by discharging at 1C until the minimum voltage specified by the manufacturer is reached.
3. Let the cell relax for 3 hours, so the terminal voltage equals the OCV
4. Discharge at C/25 until the minimum voltage is reached

This provides continuous (within the sample rate of the data acquisition hardware) OCV-SOC data for discharge only.

Appendix C: Experimental System Management

This appendix describes how the real-time management software class was used to manage test work. The process of using this class to manage a test is shown in Figure A- 4 and summarised below. The dashed lines represent where the process is asynchronous.

As the object is directly in the MATLAB workspace, it can be interacted with during the test. The class contains methods to pause balancing and/or the drive cycle, prematurely end a test, send a custom balancing command rather than use the controller, take a voltage reading and find the corresponding SOC, and plot test results.

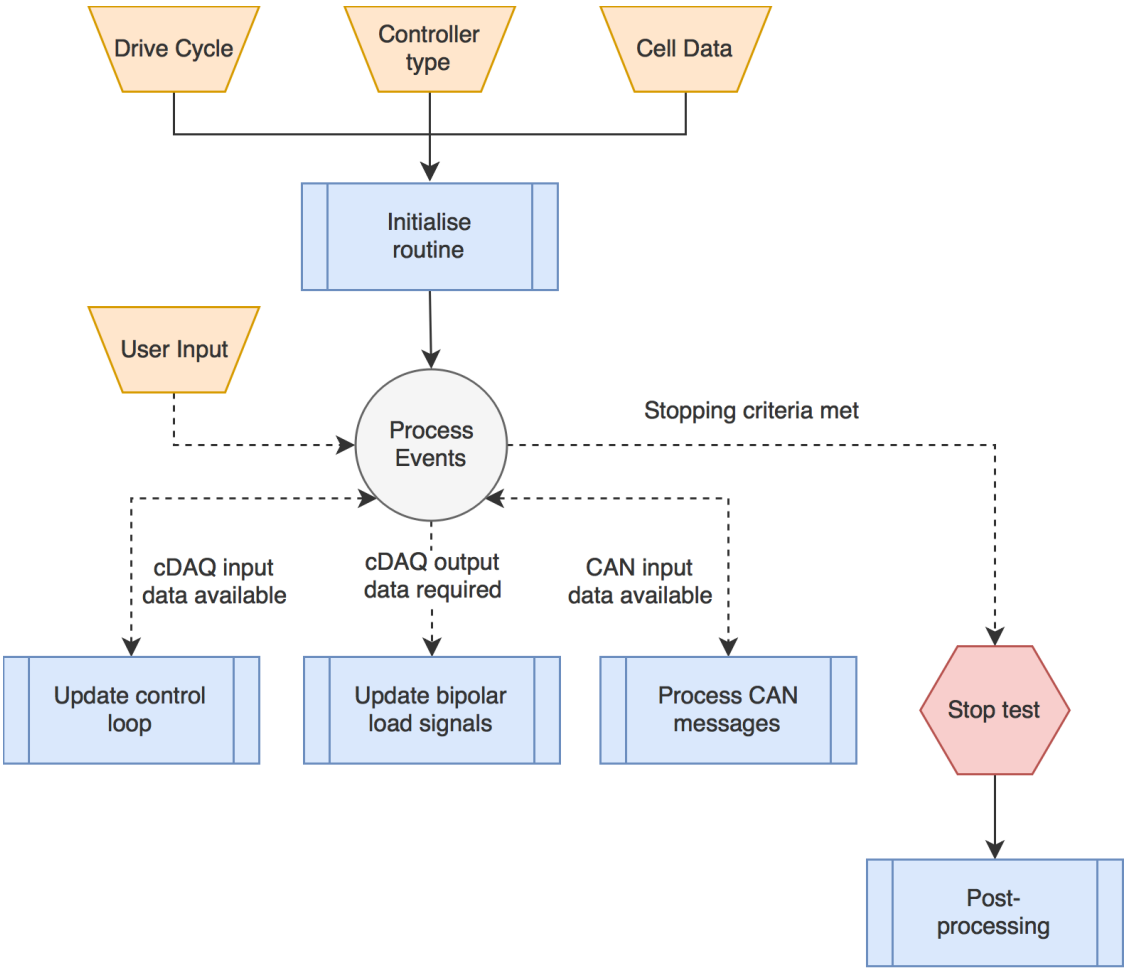

Figure A- 4: Management of real-time closed loop system

The initialisation routine performs the following steps:

1. Load in offline data required by the system. The calibration values for the current sensors are retrieved along with cell data for the estimators and controller, such as capacity, OCV-SOC data and ECM parameters, and data for the balancing system, such as the currents generated during balancing. The time and current data for the drive cycle to be applied is also loaded.
2. Open a connection to CAN-USB adapter, and specify a function to be called every 1 second to process the incoming messages. The CAN database of message information is loaded and used to create CAN message templates (using the `canMessage` data type) for sending data.
3. Open a connection to the cDAQ chassis and configure voltage inputs and outputs (the voltage range and sample frequency). The functions to be called every 1 second for the input data and output data are specified.
4. The estimators are set up using previous data to initialise internal resistance estimates, and sampling relaxed cell voltage to obtain the initial SOC estimate. They are then compiled to C code and a Mex file for each is created to allow the compiled code to be called from MATLAB.
5. An instance of the chosen controller class is created. This object is initialised, adding cell/ balancing data where necessary, and any coefficients are calculated.
6. The drive cycle is converted into the relevant voltage signals to be outputted to the bipolar load using equations (88) and (89).
7. Memory space is assigned to log data. A matrix containing 2 million rows is defined, which is sufficient to store over 5 hours of 100Hz data. Time, voltages, currents, SOCs, resistances and duty cycles are all logged.
8. A figure and axes are created and formatted so they are ready to plot the test data.
9. The CAN-USB device and cDAQ chassis are started so they begin receiving and transmitting data. Balancing is disabled and the drive cycle set to zero until the user specifies they are ready to begin.

With the experiment underway, three functions are triggered every 1s. These are outlined below. When CAN messages are received:

1. The CAN messages are unpacked to extract the cell voltages, temperatures and balancing status information.
2. If any of the temperatures are too high, the test is aborted.
3. The cell voltages are checked and the de-rate strategy in Figure A- 5 is applied if necessary.
4. If the status information shows that the PCB has stopped functioning, a message is displayed telling the user to reset the board's power supply.

When cDAQ input data are acquired:

1. If any cells are outside of the voltage window, the de-rate strategy in Figure A- 5 is applied.
2. The wire voltages are converted into cell currents using the calibration data.
3. Each cell estimator is iterated across the 100 pairs of current and voltage data.
4. Update the balancing controller using the most recent SOC estimate. This outputs a corresponding duty cycle for each cell.
5. Pack the relevant CAN messages with the duty cycle information, and transmit to the balancing PCB.
6. Plot the current, voltage and SOC data
7. Add the voltages, currents, SOCs, resistances and duty cycles to the data storage matrix

When cDAQ Output data are required:

1. The current position in the drive cycle is retrieved. If the end of the drive cycle has been reached, then the drive cycle is repeated from the start.
2. If the management class has requested to stop loading the cells (either because of a fault condition such as the cells being outside of their voltage window, the test has completed or the user has requested the drive cycle be paused).
3. Send the next 100 samples (of the drive cycle, or a zero-load signal if step 2 applies) to the cDAQ buffer for it to output over the next 1s.
4. If the drive cycle has not been paused, update the position in the drive cycle.

When the test has finished:

1. The applied load is set to zero and balancing is disabled.
2. The cells are left to relax for a period of time (from 1-30 minutes depending on the test) so they cell voltage can be assumed equal to the OCV and the final SOC can be calculated.
3. The logged data and final SOCs are saved in a `.mat` file

The class also handles the de-rate strategy shown in Figure A- 5. This was designed to maintain safe operation while avoiding a test being aborted because of a brief spike in voltage. It uses two sets of voltage limits: the regular limits defined by the manufacturer (in this case, 4.2V and 2.7V respectively), and extended limits, 0.1V outside of the regular limits. The extended limits define when a test must be aborted immediately, because of safety concerns. However, if the cell voltage is outside of the regular limits but below the extended limits, the test will only be aborted if it stays in this range for five consecutive seconds. Within this time period, the drive cycle current is reduced by 20% each second, to reduce the current through the cell. This should bring the cell voltage inside of the regular limits unless the cell OCV is outside of the limits, in which case the test will be aborted.

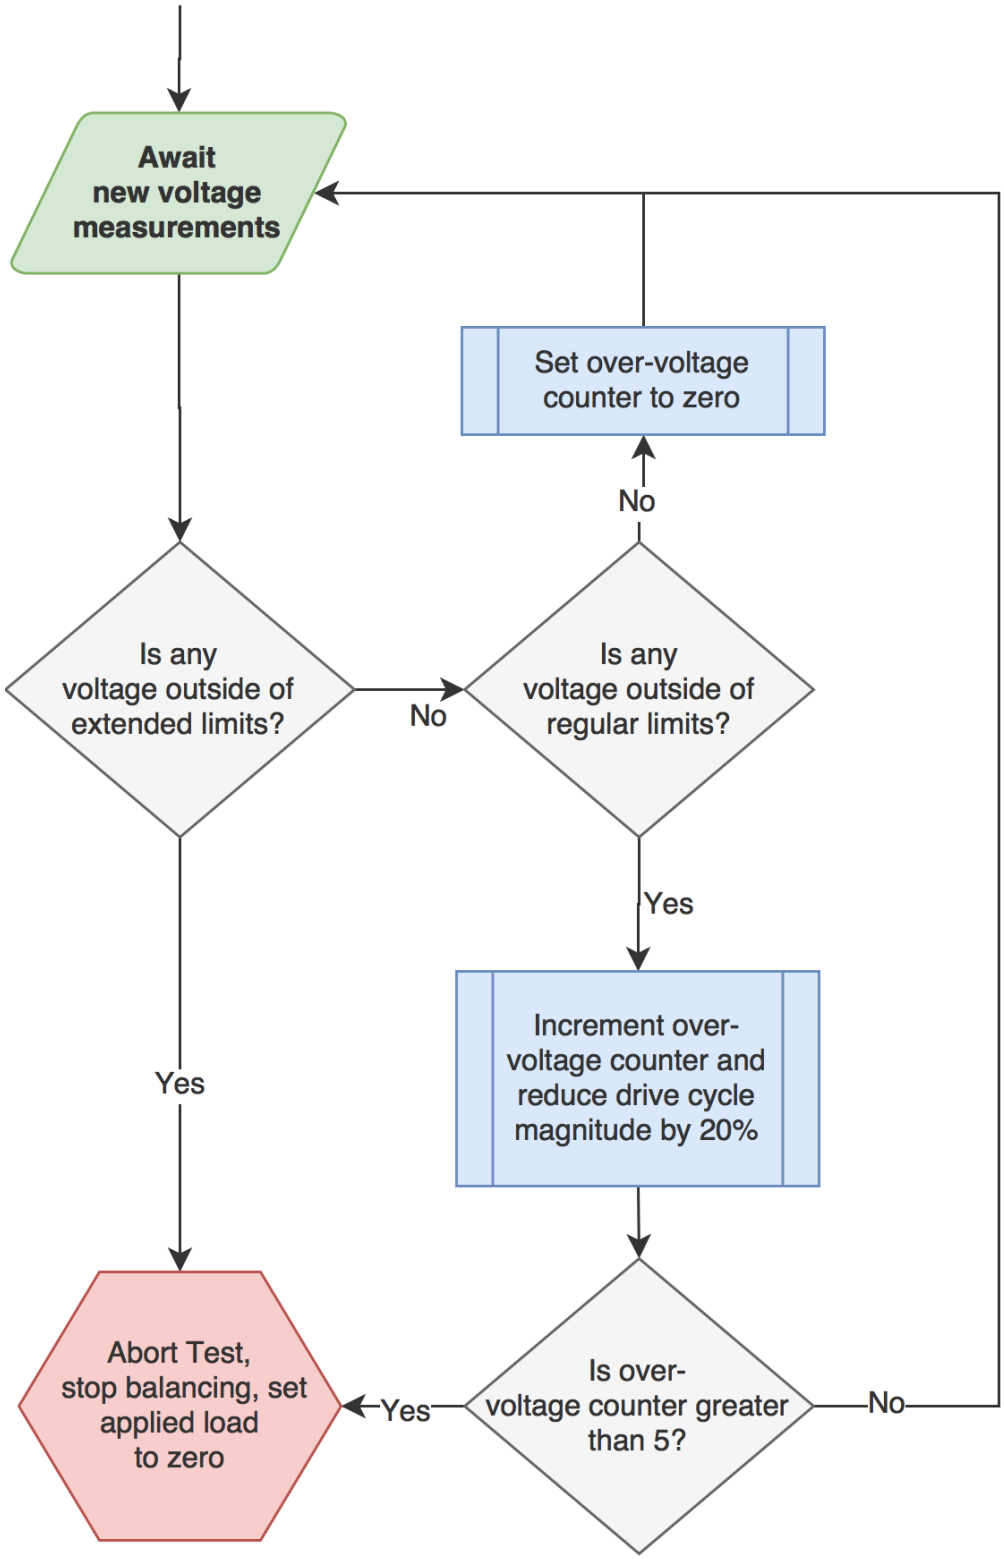

Figure A- 5: Derate strategy for cell voltage