



Scale and Bandwidth Extension of Power Converter–Based Impedance Spectroscopy

Arne Sandschulte

Thesis submitted in accordance with the requirements of the
University of Liverpool for the degree of Doctor in Philosophy.

September 2022

Abstract

The ongoing energy transition from fossil fuels to renewable energy sources has led to a substantial rise in the use of electrochemical energy storage devices. For example, electric vehicles and energy storage for distributed generation are two applications in which batteries are being widely adopted as energy storage solutions.

This research project is concerned with the development of diagnostic monitoring techniques for electrochemical energy storage devices, with a particular focus on lithium-ion batteries. Diagnostic monitoring is a critical part of systems employing such energy storage devices, as it ensures their safe, reliable, and efficient operation. This makes accurate diagnostic monitoring techniques essential for the successful adoption of the aforementioned technologies.

Electrochemical Impedance Spectroscopy (EIS) is an established technique that is used widely in the characterisation of electrochemical systems, including batteries. Recent work has focused on a novel *power converter-based* EIS implementation, which has the potential of making EIS accessible outside of research laboratories.

The aim of this work is to contribute to the development of the power converter-based EIS technique, focusing primarily on extending its scale and bandwidth.

A method is presented to compensate for variations in the state of charge, allowing the measurement bandwidth to be extended to frequencies in the millihertz range. Further, a strategy is developed to extend the upper end of the measurement bandwidth when the available sampling rate is limited. An experimental setup involving signal acquisition and real-time control capabilities is built. Experimental results allow the correct operation of the system to be validated, thus permitting for the first time a battery pack consisting of 16 series-connected lithium iron phosphate cells to be monitored using the power converter-based EIS method. With the batteries discharging at a dc current of 20 A, the controller is able to introduce the ac current perturbations in the range from 10 mHz up to 100 Hz. The measurement precision is shown to be high enough to allow meaningful impedance variations to be detected: both due to small differences between cells in a battery pack, and also due to changes in the state of charge.

Acknowledgements

I would like to extend my gratitude to the following people, without whose support this work would not have been possible:

to my second and third supervisors, Laurence Hardwick and Edoardo Patelli, for their guidance and support;

to Jill Anson and Jim Humphries for their help in and around the lab;

to Gareth Blacoe, Mark Norman, John Gillmore, and especially Mark Burnley, for all their help in getting the experimental setup up and running;

to my primary supervisor, Roberto Ferrero, for allowing me to take on this work in the first place, and the endless support that followed; and finally,

to my family, for their patience, love, and understanding;

and to my children, for the daily reminder, that patience is, indeed, a virtue.

Contents

1	Introduction	11
1.1	Objectives	14
2	Lithium-ion batteries	16
2.1	Fundamentals of electrochemical cells	16
2.1.1	The Daniell cell	17
2.1.2	Electrode naming conventions	18
2.1.3	Properties of galvanic cells	19
2.1.4	Galvanic cell structures	22
2.2	Lithium-ion batteries	22
3	Diagnostic battery monitoring	26
3.1	Rationale for monitoring	26
3.2	Monitoring techniques	29
3.2.1	State of charge	29
3.2.2	State of health	35
4	Impedance spectroscopy	39
4.1	Principles	39
4.1.1	Introduction	39
4.1.2	Procedure	40
4.1.3	Measurement perturbation types	41
4.1.4	Plotting EIS results	42
4.1.5	Equivalent circuit models (interpreting EIS results)	43
4.2	EIS-based diagnostic monitoring of LIBs	46
4.2.1	State of charge	46
4.2.2	State of health	48
4.2.3	Temperature	48
4.2.4	Discussion	49
4.3	Power converter-based EIS	49

4.3.1	Introduction	49
4.3.2	Principle of operation	51
4.3.3	State of the art	54
4.3.4	Outstanding challenges	59
5	Methodology	61
5.1	Introduction	61
5.2	Signal acquisition	61
5.2.1	Analogue-to-digital converters	62
5.2.2	Current	63
5.2.3	Voltage	64
5.3	Signal processing	66
5.3.1	Reversal of acquisition-related conversions	66
5.3.2	Signal length adjustment	67
5.3.3	Compensation for drift in voltage	69
5.3.4	Fourier transform	70
5.3.5	Signal quality verification	71
5.3.6	Analysis of drift compensation methods	73
5.3.7	Sampling below Nyquist rate	99
5.4	Power converter design and control strategy	100
5.4.1	Power converter design	100
5.4.2	Control strategy	101
6	Implementation	104
6.1	Batteries	104
6.2	Power converter design and control strategy	107
6.2.1	Power converter design	107
6.2.2	Control strategy	112
6.3	Real-time control and data acquisition system	121
6.3.1	Signal generation	122
6.3.2	Real-time current control	123
6.3.3	Signal acquisition	124
6.3.4	Data storage	126
6.4	Signal conditioning	126
6.4.1	Current	126
6.4.2	Voltage	130

7	Results	136
7.1	Verification of system operation	136
7.2	Monitoring the high-power battery pack	141
7.3	Measurement repeatability	145
7.4	Observable differences in cell impedance	147
7.5	Measurements under aliasing conditions	148
8	Conclusion	151
	References	154

List of Figures

1.1	Distribution of global greenhouse gas emissions.	12
2.1	Basic structure of Daniell cell.	18
2.2	Ragone plot showing specific energy and specific power of various energy storage devices.	21
2.3	Structure and operation of lithium cobalt oxide cell.	23
2.4	Comparison of different types of secondary cells based on specific power and specific energy.	24
3.1	Effect of discharge rate and temperature on usable cell capacity. . . .	28
3.2	EMF-SoC relationship for lithium-nickel-manganese-cobalt cell. . . .	31
3.3	OCV relaxation transient.	32
3.4	OCV-SoC relationship for lithium-iron-phosphate cell.	33
3.5	<i>Simplified Randles circuit</i> battery equivalent circuit model.	34
3.6	Current and voltage waveforms during current pulse test.	37
4.1	Nyquist plot showing a typical battery impedance spectrum.	42
4.2	Complex plane plot of simplified Randles circuit.	44
4.3	Simplified Randles equivalent circuit model.	44
4.4	Randles equivalent circuit model.	45
4.5	Regulation of power converter output voltage.	52
4.6	Principle of closed-loop power converter-based impedance spectroscopy. .	53
5.1	Common-mode voltage in a series-connection of cells.	65
5.2	Summary of signal processing steps.	66
5.3	Spectral leakage caused by incorrect signal length adjustment.	67
5.4	Spectral leakage caused by uncompensated drift in voltage waveform. . .	69
5.5	Multisine measurement perturbation in time and frequency domains. . .	71
5.6	Magnitude spectra of multisine signal — before and after drift compensation using straight line approximation method.	79
5.7	Linear least squares fit of sinusoid with drift.	81

5.8	Sinusoid with drift after compensation through the linear least squares fit method.	82
5.9	Magnitude spectra of sinusoid with and without drift, as well as compensated signal.	82
5.10	Linear least squares line of best fit for sinusoids of varying phase angles.	83
5.11	Relationship between slope of line of best fit and sinusoid phase angle.	84
5.12	Relationship between slope of line of best fit and number of cycles of input sinusoid.	85
5.13	Magnitude spectra of multisine signal — before and after drift compensation using linear least squares method.	85
5.14	The straight line component of the RPS fit for sinusoids of various phase angles.	87
5.15	Δv of ramp plus sine fit versus phase angle of input sinusoid.	87
5.16	Errors in amplitude and phase estimates after compensation of sinusoid without drift using the RPS method.	88
5.17	Relationship between the Δv of the ramp introduced by the RPS fit and the number of samples per cycle of the input sinusoid.	89
5.18	Ramp plus sine fit of sinusoid with drift.	89
5.19	Δv resulting from the slope estimated by the RPS fit when there is a mismatch between fit frequency and input frequency.	90
5.20	Magnitude spectra of multisine signal — before and after drift compensation using the ramp plus sine method.	91
5.21	Ramp plus sine fit of filtered and unfiltered multisine signal in the time domain.	92
5.22	Frequency spectra of multisine signal after drift compensation using the ramp plus sine method — with and without filtering the signal before fitting.	93
5.23	Compensation of linear drift in multisine signal (Example 1).	94
5.24	Magnitude and phase spectra of compensated and uncompensated multisine signal (Example 2).	96
6.1	Module-based battery pack.	105
6.2	Cell-based battery pack.	106
6.3	Schematic diagram of the dual-phase interleaved boost converter. . .	107
6.4	The dual-phase interleaved boost converter.	108
6.5	Input current ripple in the case of a 50 % duty cycle.	109
6.6	Model used to analyse the current transfer function of the input filter.	110
6.7	Bode plot of input filter current transfer function.	110

6.8	Model used to analyse the voltage transfer function of the output filter.	111
6.9	Bode plot of output filter voltage transfer function.	112
6.10	Bode plot of output filter impedance transfer function.	112
6.11	Circuit diagram of dual-phase interleaved boost converter used in model derivation.	114
6.12	Frequency response of control to input current transfer function. . . .	118
6.13	Frequency response of lead compensator plus integrator.	119
6.14	Closed-loop frequency response with lead compensator plus integrator.	120
6.15	Closed-loop frequency response with integral controller.	120
6.16	National Instruments CompactRIO 9035 embedded controller.	122
6.17	The current measurement board.	127
6.18	Mapping of cell voltage range to ADC input voltage range.	131
6.19	Schematic of signal conditioning circuit used in voltage acquisition. .	132
6.20	16-channel voltage conditioning PCB (top layer).	134
6.21	Output voltage of signal conditioning circuit during calibration. . . .	135
7.1	Acquired current and voltage signals before and after reversal of acquisition-related conversions.	137
7.2	Current and voltage signals after signal length adjustment.	138
7.3	Current spectrum.	139
7.4	Time domain plot of voltage signal before and after drift compensation using the RPS method.	139
7.5	Voltage before and after drift compensation using the RPS method. .	140
7.6	Nyquist plots of cell impedance based on compensated and uncompensated voltage signals.	141
7.7	The experimental setup used to monitor the high-power battery pack.	142
7.8	Time and frequency domain plots of current in high-power acquisition.	143
7.9	Time domain plot of voltage in high-power acquisition, both before and after drift compensation using the RPS method.	143
7.10	Frequency domain plot of voltage in high-power acquisition, both before and after drift compensation using the RPS method.	144
7.11	Nyquist plots of module impedance based on compensated and uncompensated voltage signals.	145
7.12	Nyquist plot with error bars showing variability in measurement set. .	147
7.13	Nyquist plot showing impedance differences between cells in pack. . .	147
7.14	Impedance spectra of one cell at different SoC values.	148
7.15	Voltage and current signals sampled at 100 kSa/s.	149
7.16	Voltage and current signals sampled at 8 Sa/s.	149

List of Tables

2.1	Theoretical specific charge capacities of the most commonly encountered positive electrode materials used in LIB construction.	20
4.1	Key characteristics of prior power converter-based EIS implementations using active measurement perturbation injection.	58
5.1	Results of compensation example 1.	95
5.2	Results of compensation example 2.	97
5.3	Results of compensation example 3.	98
6.1	Sample of battery electric vehicles available in UK in 2022.	106
6.2	Main specifications of dual-phase interleaved boost converter.	108
6.3	Component values used in the small-signal model of the converter.	118
6.4	Specifications of NI 9205 analogue input module.	125
6.5	Specifications of NI 9215 analogue input module.	125
7.1	Repeatability results.	146
7.2	Impedance estimation results from measurement under aliasing conditions.	150

Chapter 1

Introduction

The background to this research project is the international effort to reduce greenhouse gas emissions and limit the effects of global climate change. Through the Paris Agreement of 2015 the world's nations have for the first time entered into a binding treaty to limit global warming [1]. Individual nations are taking their own steps; for example, through the Climate Change Act 2008 (2050 Target Amendment) Order 2019, the United Kingdom (UK) has put legislation in place to reduce greenhouse gas emissions to net zero by 2050. Almost three quarters of the global greenhouse gas emissions come from the energy and transport sectors (see Figure 1.1). This has led to a transition in the energy landscape from using fossil fuels to renewable energy sources. As part of this transition, the UK has committed to end the sale of new petrol and diesel cars and vans by 2030, with heavy goods vehicles (HGVs) to follow by 2040 [2].

This research project is concerned with the development of diagnostic battery monitoring techniques. This is an important area of research under the current circumstances, since the number of applications in which batteries are used is steadily increasing. Electrochemical cells not only power most of the portable devices that many people have come to rely on in the last decade or so (e.g. smartphones and smartwatches), but more importantly, batteries are being widely adopted as energy storage solutions in electric vehicles (EVs) and distributed renewable energy applications — two areas of technological development which are fundamental to the energy transition. Battery diagnostics enable more efficient use of stored energy, more reliable system operation, and can help predict when essential system maintenance is required. Accurate diagnostic monitoring is therefore essential to enable batteries to be used effectively in the aforementioned applications.

Electrochemical Impedance Spectroscopy (EIS), is a well-established measurement technique in laboratories. It is used widely in various areas of research, such




Figure 1.1: Distribution of global greenhouse gas emissions in 2016 [3].

as: material characterisation, development of electrochemical devices, and analysis of corrosive processes, amongst others. As implied in the name, the measurand (or quantity of interest) in an EIS measurement is *impedance*, i.e. the ratio of voltage and current spectra under steady-state conditions. The technique consists of eliciting and measuring the response (voltage/current) of an electrochemical system to an ac signal (current/voltage) within a range of frequencies, or spectrum. Electrochemical phenomena such as diffusion and charge transfer at the electrode/electrolyte interface occur on different timescales; therefore, by varying the frequency of the excitation signal, or measurement perturbation, the physicochemical effects that rule the system can be analysed separately. In this way, EIS allows the determination of a range of system characteristics, such as kinetic parameters related to charge transfer, electrolyte resistance, and diffusion coefficients [4]. For this reason it is also finding more widespread use in battery monitoring applications, since impedance spectra are able to provide a more detailed picture of the internal state of a cell compared to other more traditional measurements [5].

There are many examples in the literature in which EIS is used as a tool for

diagnostic monitoring of batteries in research environments, e.g. [6–12]; however, the method does not appear to have been adopted in many commercial applications as of yet. An important limitation of the conventional implementation of EIS measurements is the high financial expense of the instrumentation used in laboratory settings: potentiostats/galvanostats and frequency response analysers typically command costs in the region of four to five figures dollars/sterling. Another disadvantage of conventional laboratory-based impedance spectroscopy is that it does not work *in situ*, i.e. the battery or cell must be removed from the system in which it is used, in order to perform the measurement. This is impractical at best, and in many applications simply not feasible, for example, due to battery size (as in the case of EVs), or due to device construction (as in the case of many smartphones and other portable devices), as it may not have been designed to allow easy disassembly. This also means that it is not possible to integrate the EIS measurement data in the battery management system (BMS) to properly assess the state of the system in real-time.

More recently, a new implementation of EIS has been proposed in the literature, which is based around the idea of using power converters to introduce the excitation signal required for the impedance measurement. Most battery-powered systems already employ power converters; this approach therefore has the immediate advantage of not requiring additional expensive signal generation circuitry. Also, the cell or battery does not have to be removed from the device under test and the measurement can take place *in situ*, or *online*. Power converter-based impedance spectroscopy therefore has the potential to overcome some of the limitations and drawbacks of traditional lab-based EIS, which have so far prevented this measurement method from being used more widely.

Whilst power converter-based EIS has been shown to be an effective method to acquire the impedance spectra of batteries for condition monitoring purposes, e.g. [11–19], there are still outstanding challenges that need to be overcome in order to make this a practical solution in commercial products. In particular, three areas have been identified as showing scope for development: 1) widening the measurement bandwidth; 2) verifying the feasibility of the method in larger scale systems; and 3) proposing solutions to lessen the effects of the measurement perturbation on the rest of the system.

At the low frequency end, power converter-based EIS implementations measure impedance down to hundreds of millihertz (at best), whereas commercial EIS instruments give results at frequencies as low as microhertz. The reason for this discrepancy is related to the requirement of ensuring that the system remains in steady

state for the duration of the measurement acquisition, which is more difficult when the measurement is done in situ. Whilst most power converter-based EIS implementations in the literature deal with this by restricting impedance measurements to relatively short timescales (i.e. high frequencies), there are some that try to extend measurements to lower frequencies by accounting for some of the transients that occur during acquisitions. As this work will show, there is scope for improvement in the implementation of these techniques.

Whereas batteries found in EVs and battery energy storage systems typically involve hundreds of cells, power converter-based EIS implementations are most often verified experimentally using only one or two cells. The increase in scale brings about new challenges that have not generally been considered in the literature, e.g. how to determine required transducer and analogue-to-digital converter (ADC) specifications, and also how to effectively process the measurement data in real-time.

Finally, dealing with the effects of the EIS excitation signal on the rest of the system is another area that is insufficiently addressed by the existing literature. If the perturbations in the battery current and voltage cause large power fluctuations in the output, this could challenge the adoption of power converter-based EIS in applications where strict load regulation is critical. Developing techniques to lessen these effects is therefore considered to be vital in order to widen the range of applications in which power converter-based EIS can be employed.

To summarise, power converter-based EIS is an effective method to acquire the impedance spectra of batteries; as such, it has the potential to improve the diagnostic monitoring capability of battery management systems and thus enhance the efficiency and reliability of the battery-powered applications that have become such crucial elements of the ongoing energy transition. Addressing the existing limitations of power converter-based EIS is therefore considered to be a significant area of research.

1.1 Objectives

The overarching aim of this project is to advance the development of in situ power converter-based impedance spectroscopy as a measurement technique for battery monitoring applications. The project objectives are as follows:

- To design and build a power converter-based EIS test bed with signal acquisition and real-time control capabilities that will allow the remaining objectives to be completed

- To experimentally validate the feasibility of monitoring a 16-module lithium-ion battery pack with a peak power capability of over 5 kW using the power converter-based EIS method
- To propose solutions to the technique's bandwidth limitations at both ends of the frequency spectrum and achieve a measurement range from 10 mHz to 1 kHz
- To develop strategies to reduce the effect of the excitation signal on the load, e.g. different controller designs

Chapter 2

Lithium-ion batteries

This chapter gives an overview of electrochemical cells with particular focus on lithium-ion batteries (LIBs). The first section gives a brief introduction to galvanic/voltaic cells, covering fundamental concepts and common terminology used in electrochemistry. Latter sections concentrate on LIBs: their operating principle, different types of materials used in their construction, and how they compare to each other as well as other types of rechargeable cells.

It should be noted that this chapter is not a rigorous treatise of electrochemistry. Instead, the purpose of this chapter is to give the reader an introduction to the subject, focusing on fundamental concepts, rather than detail, in order to enable better understanding of later chapters.

2.1 Fundamentals of electrochemical cells

Electrochemistry describes the relationship between electricity and chemical reactions. This interaction can happen in one of two ways that are basically the reverse of each other: chemical reactions can give rise to a flow of electric charge (i.e. an electric current), or an electric current can lead to chemical reactions. The interaction can also be described from an energy conversion point of view: chemical energy is turned into electrical energy, or the reverse, i.e. electrical energy is converted into chemical energy. The chemical reactions that are involved in these processes are known as reduction and oxidation reactions — often referred to using the portmanteau *redox* [20]. Redox reactions in electrochemistry typically involve two chemical species (could be atoms, ions, or molecules) with differing reduction potentials, i.e. the degree to which they attract electrons. The species with the higher reduction potential attracts electrons from the other species, thus leading to a transfer of electron(s). The species gaining electrons is said to be *reduced*, whilst the species

losing electrons is said to be *oxidised* [21].

This is the principle used in electrochemical cells. As stated before, electrochemistry happens in two directions, electrochemical cells therefore come in two types: galvanic/voltaic cells and electrolytic cells. Electrolytic cells are used to bring about chemical (redox) reactions through the application of an electric current. These reactions are said to be *non-spontaneous*, or thermodynamically unfavourable, since they would not occur on their own (it takes electrical energy). A simple example is the electrolysis of water to obtain hydrogen. Galvanic or voltaic cells, on the other hand, combine chemical species that produce *spontaneous* redox reactions; these are thermodynamically favourable reactions, i.e. they *do* occur on their own. The result is a transfer of electrons, or electric current.

The basic structure of a galvanic cell consists of two electrodes separated by an electrolyte. The electrodes are made from two different materials that are chosen so as to allow the spontaneous redox reactions to take place. The electrode undergoing oxidation (losing electrons) is called the *anode* and the electrode undergoing reduction (gaining electrons) is called the *cathode*. The electrolyte is an ionic conductor but electric insulator, i.e. ions can move through it, but electrons cannot. By connecting the two electrodes to each other through an external circuit, this creates a path for the electrons to flow from anode to cathode.

2.1.1 The Daniell cell

Figure 2.1 shows a very simple example of a galvanic cell that is known as the *Daniell cell*. The zinc (Zn) and copper (Cu) electrodes are immersed in beakers containing the electrolytes, which in this type of cell consist of zinc and copper sulfate solutions [20]. The two electrodes are connected through an external circuit consisting of metal wires and a load. Copper has a higher reduction potential than zinc, i.e. it exerts a stronger pull on electrons; the resulting reaction in this type of cell is that the zinc gives up two electrons, which travel through the circuit and then reduce a copper ion at the interface between the copper electrode and the copper sulfate solution. In this cell, the zinc electrode is therefore the anode, as it is oxidised, and the copper electrode is the cathode, because it is reduced. The zinc atom that was oxidised now has a positive net charge, i.e. it is now an ion, and because metal ions usually dissolve in solution, the zinc ion will fall into the solution. On the other hand, the copper ion that was reduced will now have a neutral net charge and it will no longer be soluble, i.e. it will attach to the copper electrode. As a result, as the reactions proceed, the anode will shrink/corrode, and the cathode will grow. Over time it can be seen that more and more positive ions will be added

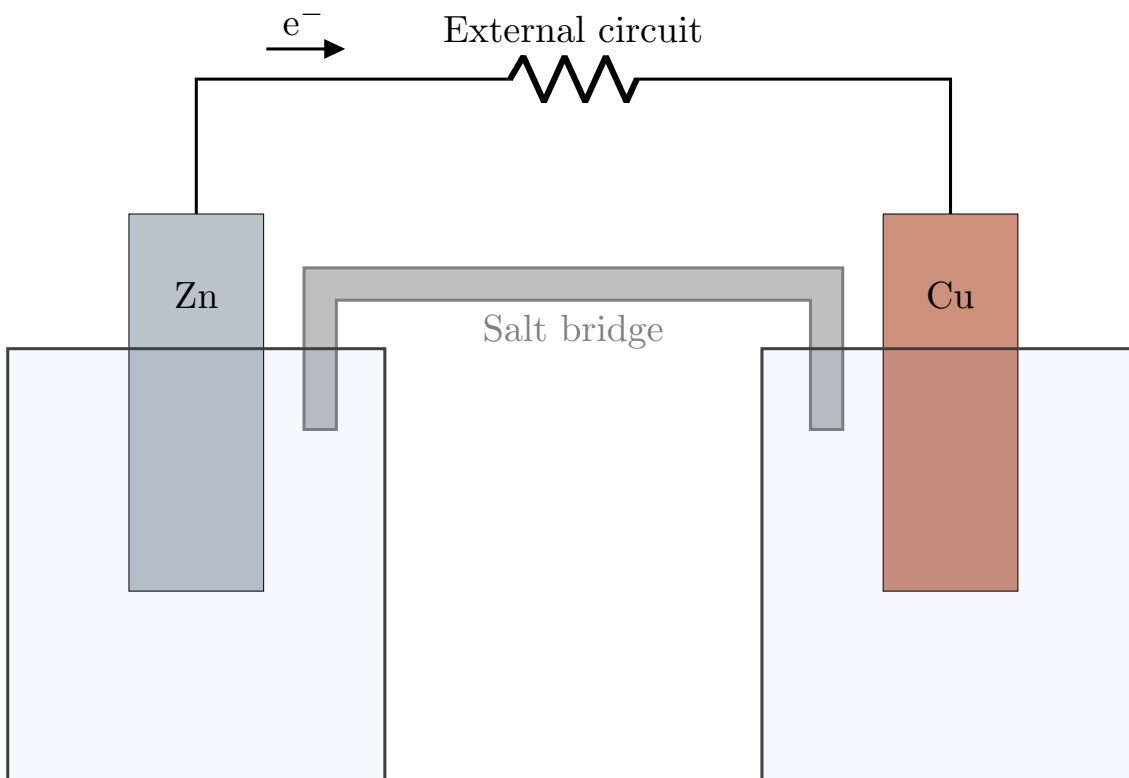


Figure 2.1: Basic structure of Daniell cell.

to the zinc solution, whereas the copper solution will continuously lose positive ions, i.e. a positive charge would be built up in the zinc solution and a negative charge would be built up in the copper solution. This build-up of opposing charges would counteract the spontaneous redox reactions and eventually cause them to stop. To prevent this from happening, a salt bridge is used. The salt bridge is immersed in both beakers and allows any built-up charge to be neutralised by letting through salt ions of opposite polarity to the metal ions in the solution.

2.1.2 Electrode naming conventions

A short note on electrode naming conventions. It was stated earlier that the electrode undergoing oxidation is known as the anode, and the electrode undergoing reduction is known as the cathode. In *primary* cells, i.e. non-rechargeable cells, this naming convention is unambiguous, because the current is always in the same direction; in *secondary* cells, i.e. rechargeable cells, however, the reaction that occurs during the discharge phase can effectively be reversed by operating the cell as an electrolytic cell. During the charge phase, when the current is reversed compared to the discharge phase, the two electrodes swap names: the electrode that was referred to as the anode during discharge is now being reduced and therefore gets the

name cathode; and likewise, the electrode that was referred to as cathode during discharge is now being oxidised and should therefore be referred to as anode during charging. This can create ambiguity because to clearly identify which electrode is being referred to as anode or cathode also requires knowledge of the state of the cell, i.e. whether it is charging, or discharging. The author has found that in the context of secondary cells the terms anode and cathode are commonly used as they apply during the discharge phase, regardless of whether the cell is being charged or discharged, but this does not appear to be consistently the case. A less ambiguous naming convention for the electrodes is to refer to them by the polarity of their relative electric potential to each other, i.e. positive and negative. For example, the cathode, i.e. the electrode undergoing reduction during the discharge phase, is always at a higher potential than the anode — both during charge and discharge phases — thus, any possible confusion can be avoided by referring to it as the *positive* electrode. Alternatively, and equally as unambiguous, the electrodes can be referred to by their active materials, e.g. ‘the zinc electrode’, or ‘the copper electrode’.

2.1.3 Properties of galvanic cells

Galvanic cells are energy transducers: they convert chemical energy into electrical energy. However, this energy conversion only takes place when the circuit between the cell’s electrodes is closed and therefore galvanic cells can act as energy *storage* devices (ESDs). How much energy is stored in a cell is primarily a function of the type and amount of active material used. The *specific charge capacity* quantifies how much charge a cell can store per unit mass and this depends on the number of electrons transferred in the redox reaction and the molecular weight of the active material [20]. The capacity calculated in this way is a theoretical figure because it does not take into account many practical considerations, such as: physical cell structure, mass and effect of other cell components (electrolyte, current collectors, etc.), operating conditions (e.g. temperature), and other limiting effects [20]. Table 2.1 shows the theoretical specific charge capacities for some of the most commonly encountered positive electrode materials used in LIBs. It should be noted that the practical specific charge capacity is generally much lower than the theoretical figure (it can be up to 75 % smaller), due to the factors listed above.

Two terms that are commonly used to describe the energy storage capability of galvanic cells are *specific energy* and *energy density*. Specific energy refers to the amount of energy a cell can store per unit mass, i.e. the higher this value, the lower the weight of a cell storing a given amount of energy. Energy density, on the other hand, refers to the amount of energy a cell can store per unit volume,

Table 2.1: Theoretical specific charge capacities of the most commonly encountered positive electrode materials used in LIB construction.

Positive electrode material	Theoretical specific charge capacity [mAh/g]
$\text{LiNi}_{0.8}\text{Co}_{0.15}\text{Al}_{0.05}\text{O}_2$	279
$\text{LiNi}_{0.33}\text{Mn}_{0.33}\text{Co}_{0.33}\text{O}_2$	278
LiCoO_2	274
LiFePO_4	170
LiMn_2O_4	148

i.e. the higher this value, the smaller the room taken up by a cell storing a given amount of energy. Sometimes specific energy is also referred to as *gravimetric* energy density, whilst energy density is referred to as *volumetric* energy density. As with specific capacity, a distinction should be made between theoretical and practical gravimetric/volumetric energy densities [20].

Whilst energy density describes how much energy an ESD can store, *power density* describes how quickly that energy can be released. Again we can talk about gravimetric power density, also known as *specific power*, i.e. power per unit mass, and volumetric power density, i.e. power per unit volume.

A useful tool to compare different ESDs is the so called *Ragone plot*. This is a plot of energy density against power density¹ that can be used to highlight differences between ESDs. For example, the Ragone plot in Figure 2.2 shows that whilst electrochemical cells (batteries) have high energy density relative to capacitors, their power density is comparatively low. In other words: batteries are good at storing large amounts of energy but not so good at releasing that energy very quickly, whereas the opposite is true of capacitors: they cannot store large amounts of energy, but they are very good at releasing the energy that they can store relatively quickly.

Also sometimes shown on a Ragone plot are straight lines which indicate the time (order of magnitude) that it will take to fully charge, or discharge, a particular ESD that falls on that line, when cycled at a constant power equal to the device's specific power rating.

The volumetric and gravimetric quantities discussed up to this point are very commonly used among electrochemists, engineers, and other scientists, particularly

¹Different variants of the Ragone plot exist: some use gravimetric quantities, and others use volumetric quantities.




Figure 2.2: Ragone plot showing specific energy and specific power of various energy storage devices. Adapted from [22].

when comparing different cell chemistries, or even different types of ESDs. Due to their “per unit mass/volume” definition, they are more frequently encountered during the research and development phases of a cell. Once development is complete, a cell’s mass and volume are fixed properties, and these terms are usually no longer the most convenient way to describe a cell’s energy or charge storage capability. Instead, manufacturer’s datasheets will most commonly list the nominal cell *capacity* in amp-hours (Ah) at a particular constant discharge rate. Sometimes they may also state the energy capacity in watt-hours (Wh), again, assuming a particular constant discharge current or power. The reason the capacity is stated in conjunction with a specific discharge rate is that the former is typically a function of the latter, i.e. the cell’s capacity varies depending on how quickly the cell is discharged (amongst other things). The discharge rate is typically stated as a coulombic rate — more widely known as *C rate* — where, a C rate of 1 C is equivalent to a discharge current equal in value to the nominal capacity of the cell [20]; therefore, for example, if a cell has a specified nominal capacity of 1 Ah at a C rate of $\frac{C}{4}$, it would take 4 hours to fully discharge the cell at a current of 0.25 A.

Secondary, or rechargeable, cells are said to *age*, as their performance, relative to that at beginning of life, will tend to get progressively worse over time. The rate of deterioration is dependent upon many factors, including: temperature (both storage and/or operating temperature), charge and discharge rates, and how many times the cell is cycled, i.e. undergoes a charge-discharge cycle. Cells are usually

rated for *cycle life*, i.e. the number of cycles they can undergo before reaching end of life, and *shelf life*, how long they can remain in storage (under specific conditions) until they reach end of life. Cells are typically deemed to have reached *end of life* (EoL) when their capacity has dropped to 80 % of the nominal full charge capacity (FCC) of a new cell [23].

2.1.4 Galvanic cell structures

The term *battery* describes an electrochemical structure composed of multiple galvanic cells. The cells can either be connected in series, in which case their electromotive forces (EMFs) add, or they can be in parallel, in which case their capacities add. Regardless of how many cells it is composed of and how they are interconnected, a battery has only two external terminals. A closely related term is *module*; a module is similar to a battery in that it may consist of many cells that are connected in series, parallel, or both, but it may also consist of multiple interconnected batteries. In large-scale battery energy storage applications, such as EVs and energy storage for electrical distribution, the complete system, consisting of potentially many thousands of cells/batteries/modules, is known as a *battery pack*. Whilst the terms ‘cell’, ‘battery’, ‘module’, and ‘pack’, are in principle all relatively well-defined, in practice, many authors use the terms interchangeably — this is particularly true for the terms ‘cell’ and ‘battery’.

Splitting large battery packs into smaller structures can simplify battery management as well as system maintenance.

2.2 Lithium-ion batteries

The term *lithium-ion battery* typically refers to a class of *rechargeable* battery in which lithium cations are shuffled from one electrode to the other during discharging, and then in the opposite direction during charging. This is in contrast to the more generic term *lithium battery*, which can refer to any type of galvanic cell with a redox reaction involving lithium (including lithium-ion batteries), but is most typically used to refer to *primary* batteries that use a lithium metal in the negative electrode, such as lithium thionyl chloride cells, for example [20]. This work only considers secondary cells, i.e. lithium-ion batteries (LIBs).

Figure 2.3 shows a high-level depiction of the structure and operation of a LIB. During discharge, the negative graphite electrode is oxidised, giving up an electron which travels through the copper current collector and the external circuit towards the aluminium current collector and the positive electrode. The other byproduct




Figure 2.3: Structure and operation of lithium cobalt oxide cell. Adapted from [23].

of the oxidation reaction is a positive lithium ion (i.e. cation) which moves via the electrolyte and through the separator before it intercalates (or inserts) in the positive electrode and is reduced again as it recombines with an electron [20]. During charging, this process is reversed and the lithium moves back from the positive electrode to the negative electrode.

The electrolyte most commonly used in LIBs is a liquid, nonaqueous solution containing an organic solvent and a lithium salt, e.g. ethylene carbonate (EC) with lithium hexafluorophosphate² (LiPF_6). The electrodes are immersed in the electrolyte, which provides a path for the conduction of lithium ions between the two electrodes, whilst acting as an insulator to the conduction of electrons. A separator, which is also permeated by the electrolyte, is placed between the electrodes to prevent them from touching each other [23].

The majority of LIBs (including that shown in Figure 2.3) use what are known as *insertion electrodes*, in which a *guest material* — lithium ions in this case — intercalates, or inserts, into specific sites provided by a *host material*. To ensure long battery cycle life, this process must be reversible with minimal change to the host material between cycles, i.e. the structural intercalation framework in the host material must stay in tact between insertion and removal of guest ions. The host material must also possess high ion and electron conductivity [20].

Graphite is a low-cost material which meets these requirements and therefore finds widespread use as host material for the negative electrode in LIBs [24]. At 372 mAh g^{-1} , the specific charge capacity of graphite is much greater than that of

²Different electrolyte compositions are in use, each with respective advantages and disadvantages, but an in-depth review is outside the scope of this work.




Figure 2.4: Comparison of different types of secondary cells based on specific power and specific energy. Adapted from [26].

the materials that have historically been used for the positive electrode in commercial cells (c.f. Table 2.1) [25]. As a result, considerable research in the field has focused on developing better materials for use in the positive electrode. This has also led to the convention of referring to different types of LIBs by the active material used in the positive electrode; for example, the cell shown in Figure 2.3 is commonly referred to as lithium cobalt oxide³ (LCO).

Other LIB types which have found use in commercial applications include: lithium nickel cobalt aluminium (NCA), lithium nickel manganese cobalt (NMC), lithium iron phosphate (LFP), and lithium manganese oxide (LMO).

The most significant advantages of LIBs over other secondary cell types, such as lead-acid and nickel metal hydride (NiMH), are much higher energy and power densities (both gravimetric and volumetric). Figure 2.4 shows a comparison of the gravimetric power and energy densities of various types of secondary cells, demonstrating the superiority of LIBs over earlier types of rechargeable cells. The figure also highlights the wide range of LIBs available and the trade-off in cell design between power and energy densities.

The aforementioned properties have made LIBs particularly common in portable electronics applications, where the associated realisable space and weight savings

³LCO was indeed the cathode material used in the first commercial LIB introduced by Sony in 1991 [23].

often offset the coupled rise in cost.

The increase in specific energy also opens up the possibility of using LIBs in applications in which the use of batteries would conventionally not have been practicable, for example, battery electric vehicles (BEVs). A typical passenger vehicle powered by an internal combustion engine uses around a gallon (approximately 4.5 litres) of petrol for every 35 miles travelled. Assuming an engine efficiency of 20 %, this would suggest an energy requirement of around 640 kJ km^{-1} . For a BEV to achieve a range of 200 km, assuming a drivetrain efficiency of 60 %, would require approximately 2,000 kg worth of lead-acid batteries (assuming a specific energy of 30 Wh kg^{-1}), but only around 600 kg of LIBs at a relatively conservative specific energy assumption of 100 Wh kg^{-1} .

LIBs also compare favourably to other secondary cell types in terms of maximum charge and discharge rates, cycle life, and how far they can be discharged, or depth of discharge (DOD) [27].

One area in which LIBs tend to be viewed slightly less positively is operational safety. Due to their relatively low thermal stability, failure of LIBs can present a fire hazard when a condition known as *thermal runaway* occurs. Most commonly caused by internal short circuits, thermal runaway describes a failure mode characterised by a self-propelling internal heat rise: a large current leads to a rise in temperature which causes even larger currents, as the exothermic reaction rates are proportional to temperature [28]. Conditions that may lead to LIB-failure can be of various types: thermal (e.g. overheating), electrical (e.g. overcharging, internal and external short circuits), or mechanical (e.g. penetration, or deformation) [28]. In addition to providing adequate mechanical protection, it is therefore essential for LIBs to be used in conjunction with a battery management system (BMS), which (amongst other things) provides the basic safety functionality to prevent failure-inducing conditions, e.g. overheating and overcharging. Depending on the application, it may also be necessary to use a suitable thermal management system that can ensure adequate environmental temperatures are maintained under all operating conditions.

This section has outlined some of the fundamental characteristics of LIBs. Their advantages over earlier types of secondary cells have been considered and how this has led to their adoption in new applications, such as BEVs. Some of their drawbacks have been considered, particularly in relation to safety of operation. Both of these points demand accurate diagnostic monitoring techniques, which are the subject of the following chapter.

Chapter 3

Diagnostic battery monitoring

As highlighted in the previous chapter, monitoring of LIBs is important to ensure their safe operation. As the following section will outline, there are other good reasons to monitor battery operating parameters, irrespective of battery chemistry. This chapter will first consider some of the main reasons for monitoring and then go on to look at a few of the predominant diagnostic battery monitoring techniques that are in use.

3.1 Rationale for monitoring

The reasons to monitor batteries can broadly be grouped into three categories:

1. safety
2. lengthening of battery life
3. provision of information to the end user

The safety requirement was explained in the previous chapter for the specific case of LIBs; due to their relatively low thermal stability, LIBs must work in operating windows that are usually very well-defined in terms of characteristics such as: operating temperature range, maximum charge and discharge voltage, and maximum charge and discharge rate. Whilst this requirement may not be quite as strict for other battery chemistries as it is for LIBs, it does still apply. For example, valve regulated lead acid (VRLA) batteries will overheat and release hydrogen gas, if the recommended float voltage is exceeded, which can also lead to thermal runaway [29]. Over-discharge of nickel metal hydride (NiMH) secondary cells can also cause venting of hydrogen gas and irreversible damage [30].

The second rationale to monitor batteries was stated above as enabling a lengthening of battery life. The term *battery life* is used here to refer to three different concepts: shelf life (also sometimes referred to as *calendar life*), cycle life, and run-time.

The *shelf life* of galvanic cells is predominantly affected by the prevailing environmental conditions at the storage location. High temperature, in particular, is a factor that can have a detrimental impact, because it generally leads to an increase in chemical reaction rates, shortening the life of most types of cells, including lead-acid [29] and LIBs [20,31], as well as increasing the self-discharge rate of NiMH cells, which can also lead to cell damage [30]. In the case of LIBs with graphite electrodes, storage at high state of charge values is also typically discouraged [20], as this can accelerate growth of the solid electrolyte interphase (SEI), leading to a reduction of both capacity and power [31].

Cycle life refers to the number of charge-discharge cycles a cell can undergo before reaching *end of life*. This is commonly defined in terms of capacity fade: when the capacity of a cell has dropped to 80 % of the nominal capacity of a brand new cell, it is considered to have reached end of life [23]. Therefore, any factor causing or promoting capacity fade can be considered as reducing cycle life. For graphite-based LIBs this includes: excessively high or low temperatures, high and low state of charge, high charging rates, and overcharging [31]. Similarly, NiMH cycle life is affected by: high temperatures, overcharging, overdischarging, as well as high charge and discharge rates [30].

The final consideration under lengthening of battery life is *run-time*. The actual *usable* capacity of galvanic cells can depend on various factors, including: discharge rate and temperature. To illustrate this, consider Figure 3.1. The plot on the left shows the relationship between capacity and discharge rate for a graphite/lithium iron phosphate cell with a rated nominal capacity of 3.2 Ah (at a discharge rate of C/5, i.e. 0.64 A), highlighting that usable cell capacity decreases as discharge rate increases. The plot on the right shows the relationship between capacity and temperature for a cell of nominal capacity 2.5 Ah (at C/2), obtained during constant current discharge at a rate of 2.5 C, demonstrating that useful capacity goes down as temperature decreases.

The dependency of capacity on temperature and discharge rate also occurs in other types of secondary cells, such as NiMH [30].

Monitoring of discharge current and temperature is therefore important to ensure cells can reach their rated performance. In some applications it may even be possible to extend run-times by dynamically adapting the discharge rate according to the




Figure 3.1: Effect of discharge rate and temperature on usable cell capacity. (a) and (b) adapted from [32] and [33], respectively.

current temperature, for example.

Another parameter that is of importance in the context of run-time is state of charge (SoC), as the two are very closely related to each other. The SoC of a battery is the ratio of its residual capacity to its full charge capacity (FCC), i.e. it describes how much charge is left in the battery at any point in time and is therefore a good indicator of *remaining* run-time. Techniques to estimate the SoC will be discussed later in this chapter, but at this point it should be noted that knowledge of the SoC can increase battery run-time by, for example, disabling power-hungry parts of a system when the SoC has dropped below a certain level. Accurate knowledge of the SoC also enables more *efficient* use of the energy available in the battery (by not underestimating the remaining capacity) and ensures more *reliable* system operation (by not overestimating remaining capacity).

The third and final motivation for monitoring batteries is to provide the end user with information about the state of the battery. The two parameters that are most likely to be of interest to the user are state of charge (SoC) and state of health (SoH); however, parameters such as battery voltage, current, and temperature, could also be of interest, depending on the application.

SoC has already been discussed from the point of view of a BMS using SoC information to make battery run-time related decisions. The end user may employ the information in a very similar manner, i.e. to adjust operation of the system in such a way as not to run out of charge at an inconvenient time or place, or in the case of BEVs, for example, to plan a long journey with adequate stop-overs to allow for recharging.

SoH, on the other hand, is a figure that aims to quantitatively describe the degradation of a battery relative to its condition when new. This could, for example, be expressed relative to its end of life. SoH is a particularly useful parameter in the context of predictive maintenance, which is an important aspect of large-scale battery packs, such as those used in grid storage applications and BEVs.

This section has demonstrated the importance of monitoring various battery operating parameters. In the next section, commonly used monitoring techniques will be considered.

3.2 Monitoring techniques

The previous section has explained the rationale for monitoring various battery operating parameters, including some, such as voltage, current, and temperature, which are directly measurable physical quantities, and others, such as SoC and SoH, which are not directly measurable, but have to be inferred from those that are. The aim of this section is not to provide a rigorous review of the techniques used to measure those quantities that *are* directly measurable¹, but instead, give an overview of the methods used to determine the latter, i.e. those quantities that are *not* directly observable, highlight some of their drawbacks and shortcomings, and explain the requirement for improved diagnostic techniques.

3.2.1 State of charge

As explained in the previous section, the state of charge, or SoC, is defined as the ratio of a battery's residual capacity to its full charge capacity (FCC). It is therefore a unitless quantity and typically stated as a percentage: a SoC of 100 % means the battery is fully charged; a SoC of 0 % means the battery is fully discharged.

A few of the most widely used SoC estimation methods are considered next.

3.2.1.1 Coulomb counting

One way to estimate a battery's SoC follows from the SoC definition given above, since residual capacity, Q_r , is the difference between the FCC, Q_{FCC} , and the charge taken from the battery, ΔQ :

$$\text{SoC} = \frac{Q_r}{Q_{FCC}} = \frac{Q_{FCC} - \Delta Q}{Q_{FCC}} \quad (3.1)$$

¹This is of course still a very relevant consideration in this project and is discussed in more detail in Chapter 5.

This can be rewritten in a more general way to allow for both charging and discharging and any starting SoC, by considering the residual capacity at the beginning of the charge/discharge transient, Q_{r0} :

$$\text{SoC} = \frac{Q_{r0} + \Delta Q}{Q_{FCC}} \quad (3.2)$$

Therefore, it is possible to estimate the SoC by measuring and integrating the cell current, I , between the beginning (t_0) and ending (t_1) of the charge/discharge transient, since:

$$\Delta Q = \int_{t_0}^{t_1} I dt \quad (3.3)$$

This is known as *coulomb counting*, or ampere-hour counting, and is one of the fundamental methods of estimating SoC.

Rewriting (3.2), we can obtain an expression for the SoC at $t = t_1$:

$$\text{SoC}_1 = \frac{Q_{r0}}{Q_{FCC}} + \frac{1}{Q_{FCC}} \int_{t_0}^{t_1} I dt = \text{SoC}_0 + \frac{1}{Q_{FCC}} \int_{t_0}^{t_1} I dt \quad (3.4)$$

From (3.4) it can be seen that to estimate the SoC using the method of coulomb counting, the SoC at the beginning of the charge/discharge transient must be known. It is also necessary to know the battery's FCC. Ideally, this would be a constant value equal to the nominal capacity given in the cell's datasheet, but due to tolerances in manufacturing, there will always be some spread in this number. More importantly, it is well-known that as a cell ages, its FCC drops. Clearly, for an accurate SoC estimation, this must be taken into account; otherwise, if the SoC is always based on the cell's nominal FCC, then as the cell ages, the reported SoC value will always be lower than the actual SoC and it will no longer be able to reach 100 %. If, on the other hand, the SoC is always calculated using an up-to-date FCC figure, then a given ΔSoC will correspond to a smaller absolute capacity compared to a fresh cell, but the reported SoC will be correct. As described in the previous section, the usable capacity also depends on the discharge rate, decreasing as the current increases; however, this effect appears to be much less significant in LIBs compared to other chemistries such as NiMH and lead-acid. Techniques to account for this effect have been proposed in the monitoring of lead-acid batteries, as for example in [34].

The main drawback of the coulomb counting method is related to the estimation of the charge taken from/added to the cell. No measurement is perfect and the measurement of current is no exception. Since the current must be integrated to

obtain the charge transferred, even small measurement errors can add up over time and lead to inaccuracies in the SoC estimation [35].

Due to its conceptual simplicity and the fact that it works for any type of secondary cell, coulomb counting is still widely employed in battery monitoring applications, particularly in combination with other techniques.

3.2.1.2 Open circuit voltage

The voltage measured across the terminals of an electrochemical cell with no load connected to it is referred to as the cell's *open circuit voltage* (OCV). The term *electromotive force* (EMF), on the other hand, refers to the value of the OCV when an internal equilibrium has been reached inside the cell. A cell's EMF is approximately equal to the potential of the electrochemical cell reaction, whose definition is thermodynamic in nature and stems from the sum of oxidation and reduction potentials of the two electrodes [20]. As a result, the EMF can be used as a direct indicator of SoC, once a cell's EMF-SoC relationship has been determined. Figure 3.2 shows an example of this relationship for a LIB with a lithium-nickel-manganese-cobalt cathode and graphite anode.

One considerable benefit of the OCV/EMF-based SoC estimation method is that the EMF-SoC relationship is relatively insensitive to both temperature and aging effects [36]. It also circumvents the problems associated with changing FCC, as described in the previous section, since knowledge of the FCC is not required in order to estimate the SoC using this method.




Figure 3.2: EMF-SoC relationship for lithium-nickel-manganese-cobalt cell [37].




Figure 3.3: OCV relaxation transient.

The two main drawbacks of this SoC estimation method are related to the fact that it requires an estimate of the cell's EMF: this can take a long time to obtain and during this time the cell must be in open circuit condition, i.e. the load cannot be powered. Whilst the second point is explained by the EMF definition given at the beginning of this section, to better understand the first point, a closer examination of the internal impedance characteristics of electrochemical cells is required, which is covered in Chapter 4; for the purpose of this section, however, it is important to understand that OCV and EMF are not the same. In practice, this is easily demonstrated by considering the voltage profile of a cell whose constant current discharge is suddenly stopped. As can be seen in Figure 3.3, when the current stops, there is an approximately exponential transient before the OCV settles; the value at which it settles is the EMF.

The process during which the cell reaches internal equilibrium (sometimes also referred to as OCV relaxation) can take many hours, thus limiting the use of this SoC estimation technique to periods of long load inactivity. Various methods have been proposed to estimate the EMF from OCV measurements in order to reduce the required measurement time, e.g. [38], however, these still require periods of load inactivity, and they add to the uncertainty in the SoC estimation, as they use approximations to estimate the EMF, and typically do not adapt to the changing dynamics of ageing batteries [35].

Two further issues that complicate the use of this SoC estimation method in practice are related to the voltage profile of certain types of secondary cells, such as lithium-iron-phosphate and NiMH, as they are a) very flat, and b) exhibit hysteresis. Compared to the voltage profile of a lithium-nickel-manganese-cobalt cell, as was

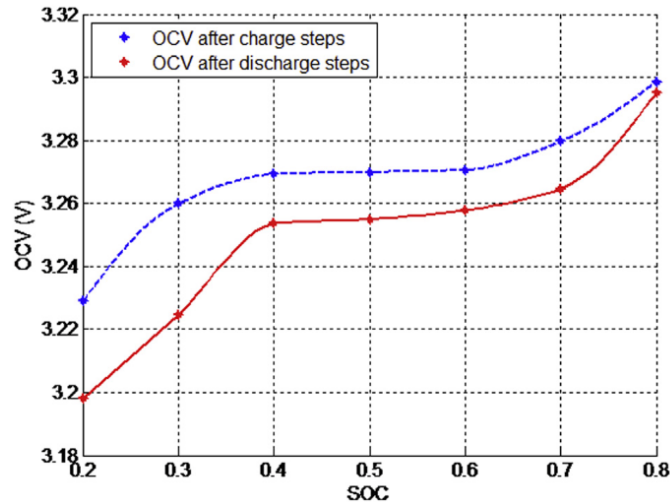


Figure 3.4: OCV-SoC relationship for lithium-iron-phosphate cell [40].

shown in Figure 3.2, it can be seen in Figure 3.4 that changes in the OCV/EMF² of a lithium-iron-phosphate cell are very small, particularly in the SoC range from 40-60 %. Figure 3.4 also shows the hysteresis effect: the OCV corresponding to a particular SoC varies, depending on whether the cell was previously charged, or discharged. Whilst hysteresis is known to affect all types of secondary cells to some extent, the effect is particularly noticeable in lithium-iron-phosphate cells [39].

3.2.1.3 Hybrid, model-based, and adaptive methods

The SoC estimation methods presented so far, i.e. coulomb counting and the SoC-EMF relationship, are fundamental methods, in the sense that they are relatively basic in how they operate, and also in that they are often combined with other techniques, or used as part of more complex systems to overcome some of their respective limitations. It is quite common, for example, for both coulomb counting and EMF measurements to be combined with each other. The basic idea in [38], for example, is to estimate the SoC using coulomb counting when the battery is being charged/discharged and to perform EMF measurements during periods of inactivity, thus correcting any error accumulated during current integration.

Often these methods are centred around the concept of a *battery model* — essentially a mathematical description of the dynamic behaviour of the battery. The nature and complexity of battery models varies greatly; the most complex do not only consider temperature and hysteresis effects, but also, as the battery ages, they can adapt to its changing behaviour by updating relevant model parameters accordingly.

²Here the OCV was measured after a one hour rest period following a charge/discharge step.

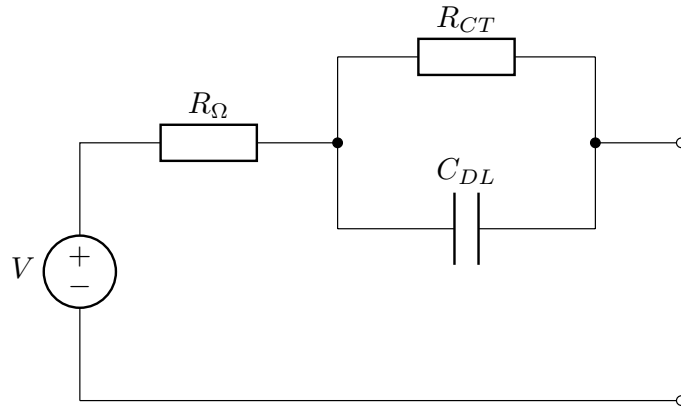


Figure 3.5: *Simplified Randles circuit* battery equivalent circuit model.

One of the most commonly encountered types of battery models is the *equivalent circuit model* (ECM). In much the same way as the behaviour of a real (non-ideal) independent voltage source can be modelled as an ideal independent voltage source with a series impedance, in battery ECMs electrical components are used to model the response of real batteries. Each component is associated with one, or several, electronic or electrochemical effects. Many of these exhibit frequency dependency; therefore it is common to speak of battery *impedance*, rather than resistance. One version of a very widely used (very simple) ECM known as the *Simplified Randles circuit* is shown in Figure 3.5. Here, the series resistance, R_{Ω} , represents the sum of all ohmic elements, such as: the ionic resistance of the electrolyte and the electronic resistance of the current collectors and the active mass of the electrode; C_{DL} and R_{CT} are used to model the behaviour of the electrochemical double layer and the charge transfer reaction³.

The SoC estimation used in connection with hybrid and model-based methods is often based on the SoC-EMF relationship [35]. This is only possible, if the following are known: a) the SoC-EMF relationship, b) the relationship between the parameters and the EMF, and c) the parameter values at different operating conditions (this could include, for example, voltage, current, and temperature). Once the model has been set up, it can be used to estimate the SoC by first finding the parameter values associated with the current operating conditions (which are obtained by measurement), then estimating the EMF, and finally the SoC. There are many different techniques in use to determine the parameter values — also known as *model parameterisation*; they all have in common that they require a dedicated cell characterisation procedure, which is also needed in order to determine the SoC-EMF

³The origin and significance of the different model parameters are considered in greater detail in Chapter 4.

relationship. This means that the parameters are determined at a particular age/state of health (SoH) and as the cell ages the model and thereby the SoC estimation becomes less accurate [35]. To overcome this problem and improve the SoC estimation accuracy, a range of adaptive techniques are described in the literature⁴ that aim to adapt the model parameters as the battery ages and its dynamic behaviour changes. Among these, the Kalman Filter finds widespread use.

3.2.2 State of health

As described in Section 3.1, the state of health, or SoH, is the concept of quantitatively describing the degradation of a battery's condition relative to when it was new. This is particularly important in applications in which reliability of operation is essential, such as grid storage and BEVs, where a battery failure could result in a power outage, or being left stranded in the middle of nowhere. A dependable SoH indication could allow the system operator to take appropriate steps to prevent such events from happening.

Although the significance of estimating SoH is beyond doubt, this is a research area that is yet to be very well-established. There is, for example, no consensus on a clear definition of SoH. In this regard, the literature can broadly be split into three groups:

1. uses the term SoH in a qualitative manner and/or does not define it
2. defines SoH in terms of capacity fade only, e.g. [41–43]
3. defines SoH in terms of both capacity and power fade, e.g. [35, 44–46]

As a battery ages, its performance or 'health' deteriorates: its FCC drops and its impedance increases — the latter resulting in a reduced output power capability. It does therefore seem justified to define SoH in terms of both capacity and power fade. Whilst a capacity fade of 20 % (relative to nominal) is often used to define 0 % SoH, it is harder to find a corresponding power fade figure that is used as commonly.

It follows, that depending on the exact definition of SoH used, it is necessary to measure FCC, impedance, or both, in order to estimate the SoH of a battery [35].

3.2.2.1 Full charge capacity

The simplest and probably most accurate way to determine the FCC of a cell, is to first fully charge and then fully discharge it at a constant current. In practice,

⁴A summary, including a wide range of references, is given in [35].

however, this method is unlikely to find much use outside the laboratory, as it takes a long time to complete — during which the cell cannot be used — and is unnecessarily wasteful.

Consider again (3.4), which, for convenience, has been restated here as (3.5):

$$\text{SoC}_1 = \text{SoC}_0 + \frac{1}{Q_{FCC}} \int_{t_0}^{t_1} I dt \quad (3.5)$$

If we let $\Delta\text{SoC} = \text{SoC}_1 - \text{SoC}_0$, we can rewrite (3.5) to obtain an expression for the FCC:

$$Q_{FCC} = \frac{1}{\Delta\text{SoC}} \int_{t_0}^{t_1} I dt = \frac{\Delta Q}{\Delta\text{SoC}} \quad (3.6)$$

From (3.6) we can see that it is possible to estimate the FCC based on the charge differential between two SoC values. The advantage is that it is not necessary to fully cycle the cell in order to estimate the FCC using this method and thus it does not take as long to complete. However, the method is subject to the same measurement uncertainties as whichever method is used to estimate the SoC, and as (3.6) shows, it suffers the same drawbacks as the coulomb counting method. A further weakness of this method is that in order for the FCC estimate to be accurate, the change in SoC needs to be large [35].

As is the case for SoC estimation, the estimation of FCC is also often based around various modelling approaches. A thorough review of these methods is outside the scope of this work; for this purpose, the interested reader is referred to [35].

3.2.2.2 Impedance

As touched upon above, as a battery is cycled and ages, its ability to supply energy at a high rate gradually decreases. This is also referred to as *power fade*. In LIBs with graphite anodes, power fade is mainly due to two processes, both of which lead to an increase in the internal cell impedance: the continuous growth of the solid electrolyte interphase (SEI), and corrosion of the current collectors [31]. It is thus possible to quantify SoH, either in terms of power fade, or impedance rise.

One of the main methods used to measure battery impedance is known as *impedance spectroscopy*, or Electrochemical Impedance Spectroscopy (EIS). This will be covered in detail in Chapter 4.

The *current pulse test*⁵ is a method that is sometimes used to specifically test a battery's power capability. It is employed in [7] to determine what is referred to therein as the direct current resistance (DCR). As shown in Figure 3.6, the test

⁵Alternatively also referred to as *pulse power characterisation* [47].




Figure 3.6: Current and voltage waveforms during current pulse test. To give an example of typical values, in [7] the direct current resistance (DCR) of a new 42.5 Ah NMC cell was determined using a current pulse of magnitude $\Delta I = 1.25$ C and duration $\Delta t_p = 10$ s. At a temperature of 25 °C and SoC of 50 %, this produced a ΔV_{DCR} of approximately 60 mV, corresponding to a DCR of around 1.2 m Ω . Figure adapted from [7].

consists of applying a pulse to the discharge current and measuring the resulting change in the terminal voltage; however, just as the terminal voltage did not instantaneously return to the value of the EMF when the discharge current was suddenly interrupted in Figure 3.3, it can be seen in Figure 3.6 that the same type of approximately exponential transient occurs in the terminal voltage when the current step is in the opposite direction. This raises a question over the timing of the voltage measurement. The aim of the measurement is for the voltage drop to include ohmic and charge transfer effects only — represented in Figure 3.6 by ΔV_0 and ΔV_1 , respectively — as they are said to be the most relevant in terms of the cell’s power capability, whilst leaving out slower effects related to the diffusion of lithium ions and the associated gradual drift in the EMF [7]. In [7] this is achieved by fitting the voltage response with a straight line during the second half of the current pulse and extrapolating the line backwards to the beginning of the pulse at $t = t_0$. The difference in voltage at t_0 between the terminal voltage and the straight line approximation is denoted ΔV_{DCR} and is used in the calculation of the DCR:

$$\text{DCR} = \frac{\Delta V_{\text{DCR}}}{\Delta I} \quad (3.7)$$

An alternative, simpler approach is used in [47], where the voltage used to calculate the DCR is taken as the terminal voltage after a specific time, measured from the beginning of the current step. This is then repeated at various times to obtain a series of DCR measurements.

One of the drawbacks of using DCR as an indicator for power fade and (by extension) SoH, is that it has a non-linear relationship with current [7]. This means that in order to get meaningful and consistent results, the same current step magnitude must always be used. The current waveform itself also limits the applications in which this method may feasibly be used in, as it demands a certain flexibility from the system. Finally, as described above, determining the appropriate timing for the voltage measurement involves some trial-and-error; this requires a dedicated laboratory-based characterisation of the specific cells being used, whilst still involving a degree of uncertainty with regards to how well the ohmic and charge transfer effects have been isolated. As the following chapter will show, a more accurate method exists to discern between different electrochemical processes occurring within galvanic cells.

Chapter 4

Impedance spectroscopy

This chapter introduces the fundamentals of impedance spectroscopy, including: basic principles, different measurement perturbation types, and equivalent circuit models. The use of impedance spectroscopy as a diagnostic monitoring tool for LIBs is considered with focus on SoC, SoH, and temperature estimation. Finally, power converter-based impedance spectroscopy is introduced and a survey of relevant literature is presented, highlighting both the outstanding challenges related to its implementation as well as the motivation for this particular work.

4.1 Principles

4.1.1 Introduction

Impedance Spectroscopy (IS) is a measurement method that enables properties of materials to be analysed. It is used in many different scientific fields and is referred to in various ways, depending on the particular research area and application, e.g. Electrochemical Impedance Spectroscopy (EIS), admittance spectroscopy, immitance spectroscopy, dielectric spectroscopy, complex conductivity, complex permittivity, and bioimpedance spectroscopy [4,48]. In this work the terms IS and EIS are used interchangeably.

In electrochemical applications EIS is used to characterise the properties of materials used in electrodes, as well as the reactions taking place in electrochemical cells. EIS is well established in laboratories (as well as some industrial settings), where it is used in a wide range of applications, such as: development of energy storage devices, and analysis of corrosive processes, amongst others.

4.1.2 Procedure

The technique consists in perturbing an electrochemical system with a voltage (or current) stimulus and measuring the system's current (or voltage) response. This allows the system's *impedance* to be determined by taking the ratio of the voltage and current spectra under steady state conditions. The stimulus — or *measurement perturbation* — that is most commonly used is a sinusoidal signal. By varying the frequency of the measurement perturbation and repeating this process, an impedance *spectrum* can be obtained. The range of frequencies covered by an impedance spectrum varies depending on the specific application and could stretch anywhere from μHz to MHz .

The electrochemical mechanisms relevant to batteries, i.e. ion diffusion in the electrolyte, charge transfer at the electrode–electrolyte interface (i.e. redox reactions), and electronic conduction in the electrodes, are all manifested in the system impedance [4]. Since they occur on different timescales, they each dominate the overall impedance in different parts of the frequency spectrum and can thus be characterised independently by selecting an appropriate frequency range for the measurement perturbation [48]. This allows the determination of a range of system characteristics, such as: kinetic parameters related to charge transfer, electrolyte resistance, and diffusion coefficients [4], which can in turn be used to infer a much more detailed picture of the inner state of an electrochemical cell compared to any of the monitoring techniques considered in Chapter 3.

Many electrochemical systems — batteries included — have non-linear current–voltage (I–V) relationships. In order for the impedance to be meaningful, EIS measurements must therefore use small-amplitude perturbations around the operating point determined by the DC conditions. Doing so enables a small portion of the I–V curve around the operating point to be considered as approximately linear. With regards to *how* small the perturbations should be, [4] states that the perturbation in the voltage should be less than the thermal voltage¹, i.e. approximately 25 mV at room temperature; an upper limit of 10 mV is commonly found in the literature. The lower limit, on the other hand, will largely depend on the noise levels in the signals being measured, the specification of the acquisition system, i.e. signal conditioning and analogue-to-digital circuitry, as well as the required signal-to-noise ratio (SNR).

Since it is possible to perturb the system either with a stimulus in the voltage, or the current, this gives rise to two basic types of EIS, named after the controlled

¹The thermal voltage, V_T is given by $V_T = kT/e$ where k is the Boltzmann constant, T is absolute temperature, and e is the electron charge.

(perturbed) variable: Potentiostatic EIS (PEIS) and Galvanostatic EIS (GEIS).

Conventional EIS measurements are conducted using instruments known as potentiostats/galvanostats in combination with frequency response analysers. These devices are most commonly used to apply a sinusoidal measurement perturbation to an isolated cell (i.e. connected solely to the measurement instrumentation), whilst sweeping the perturbation frequency through the required frequency range. Although some instruments allow the addition of a small dc component on top of the ac measurement perturbation, in conventional EIS the dc component is usually set to zero.

4.1.3 Measurement perturbation types

Another aspect in which EIS implementations differ is the type of measurement perturbation used. As stated above, the most commonly used excitation signal — particularly in laboratory instrumentation — is the sinusoidal waveform. This has the advantages of being relatively easy to implement and offering flexibility in terms of the exact signal specification used, i.e. frequency and amplitude. The main disadvantage is that each frequency is handled separately and if the impedance spectrum is required to cover a frequency range down to millihertz or even less (which can be the case in battery monitoring applications), it may take a very long time to acquire the complete spectrum, which can make it difficult to ensure that the steady state condition is adhered to. In order to speed up the spectrum acquisition it is possible to use measurement perturbations with richer harmonic content. For example, a signal consisting of multiple sine waves superimposed on each other — sometimes referred to as *multisine* — could be used to achieve this. One drawback of this type of signal is that it is harder to synthesise than a single sine wave, particularly if the components should cover a wide frequency range. Also, for a given peak-to-peak or root mean square (RMS) measurement perturbation amplitude, each component's individual amplitude will be much smaller. Another alternative is to use a square wave. This is arguably easier to implement than a sine wave excitation and has the added benefit of introducing odd harmonics of the fundamental frequency. However, the Fourier series of a square wave also tells us that the amplitude of the harmonics diminishes as $1/n$, where n is the harmonic number, i.e. the SNR quickly deteriorates. A different approach is to use a broadband noise signal, e.g. white noise. This has the advantage of exciting all frequencies (in some bandlimited range) simultaneously, thus enabling faster impedance spectrum acquisition. In practice, however, this measurement perturbation type also has drawbacks, namely: it is difficult to generate true white noise [4]; and the linearity of the response can no longer be

verified by probing for the presence of harmonics [4]. Also, due to its flat power spectral density, it is not particularly energy efficient; the signal's energy content is spread over all frequencies resulting in an unnecessarily fine spectrum resolution, and for a given signal power budget it results in a worse SNR compared to more directed perturbation types, such as sine, or multisine.

4.1.4 Plotting EIS results

Probably the most commonly encountered visualisation method of impedance spectra obtained by EIS is the *Nyquist plot*. This is a plot of the impedance in the complex plane, i.e. the imaginary part of the impedance is plotted against the real part. Since the impedance of batteries tends to have a negative reactive component, Nyquist plots depicting battery impedance spectra are conventionally drawn with an inverted imaginary axis. An example of a Nyquist plot for a typical battery impedance spectrum is shown in Figure 4.1. The plot has been annotated to indicate the dominant electrochemical mechanism in each part of the spectrum with a rough guide to the relevant frequency ranges.

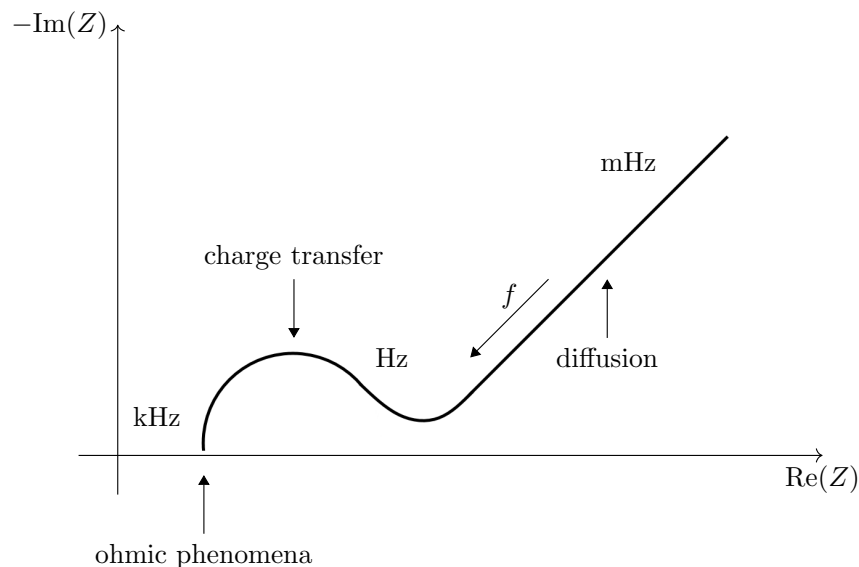


Figure 4.1: Nyquist plot showing a typical battery impedance spectrum.

Whilst Nyquist plots are useful tools to compare several impedance spectra and they also enable the impedance's real and imaginary parts to be readily determined, it can be difficult to quickly establish impedance data in magnitude and phase form. Probably the most significant drawback of the Nyquist plot, however, is that frequency is an implicit parameter. It is common to find annotations indicating in which direction the frequency increases (as in Figure 4.1); also, selected points on the plot may be annotated with their respective frequencies, but it is challenging to

show the frequency associated with *every* impedance point, as this tends to clutter the plot too much, making it hard to read. If greater detail is required — either in the magnitude, phase, or frequency data — then a Bode plot may be a more convenient way of visualising the impedance data.

4.1.5 Equivalent circuit models (interpreting EIS results)

The interpretation of impedance spectra obtained by EIS is very often based around the use of equivalent circuit models (ECMs). The fundamental idea is that the behaviour of any electrochemical cell may be modelled by an equivalent circuit consisting of combinations of passive electrical components (i.e. resistors, inductors, and capacitors), as well as a number of components without electrical counterparts that will be described in due course [20]. Each component in the model is typically associated with a particular electrochemical process occurring inside the cell. The selection of a suitable ECM for a particular application *should* therefore be based on prior knowledge of the physical and chemical characteristics of the specific cell used. The component values are determined in a model parameterisation procedure. Different parameterisation methods exist; one very commonly employed approach is the *least squares* method, in which the sum of the squared residuals — i.e. the difference between the measured impedance and that predicted by the model — is minimised. Once the component values have been determined for all impedance spectra obtained in a particular experiment, the experimental results can be analysed (and hopefully interpreted) through the relationship of the component values and the specific parameter(s) being studied.

To give a simple example, consider the complex plane plot of Figure 4.2. The ECM that matches this spectrum is the simplified Randles circuit first introduced in Chapter 3 and shown again here in Figure 4.3. The only component with a frequency-dependent impedance is the double-layer capacitor C_{DL} . At very low frequencies capacitors act as open circuits; the overall impedance thus simplifies to approximately $R_{\Omega} + R_{CT}$. At very high frequencies, on the other hand, capacitors resemble short circuits and the circuit impedance thus approximates R_{Ω} . This behaviour is reflected in the complex plane plot: at both ends of the spectrum the imaginary component of the impedance approaches zero, i.e. the impedance is resistive. At frequencies between the upper and lower limits the impedance traces out a semi-circle and has both real and imaginary components. This semi-circle corresponds to the impedance of the parallel combination of R_{CT} and C_{DL} :




Figure 4.2: Complex plane plot of simplified Randles circuit.




Figure 4.3: Simplified Randles equivalent circuit model.

$$Z_{RC}(\omega) = R_{CT} \parallel Z_{CDL} = \frac{R_{CT}}{1 + j\omega R_{CT} C_{DL}} \quad (4.1)$$

At the frequency associated with the time constant of R_{CT} and C_{DL} , i.e. $\omega = \omega_0 = \frac{1}{R_{CT} C_{DL}}$, the reactance of C_{DL} is equal in magnitude to the resistance of R_{CT} , and Z_{RC} simplifies to:

$$Z_{RC}(\omega_0) = \frac{R_{CT}}{2} - j \frac{R_{CT}}{2} \quad (4.2)$$

This is the point at the top of the semi-circle showing that R_{CT} can be determined as the semi-circle's diameter.

The overall impedance of the equivalent circuit can be stated as follows, showing that the addition of R_{Ω} in series with Z_{RC} simply shifts the semi-circle to the right by R_{Ω} :

$$Z(\omega) = R_{\Omega} + Z_{RC} = R_{\Omega} + \frac{R_{CT}}{1 + j\omega R_{CT} C_{DL}} \quad (4.3)$$

Comparing the Nyquist plot of Figure 4.1 to that shown in Figure 4.2, it can be seen that the *simplified* Randles circuit is missing the effects observed at very low frequencies, i.e. those due to very slow diffusion processes. The plot of Figure 4.1 is in fact that corresponding to the regular Randles circuit — shown in Figure 4.4; this is a model for a cell with an electrode–electrolyte interface in which the rate of electrochemical reactions is controlled by diffusion [21], i.e. the rate of charge transfer is faster than the diffusion of the reactive species in the electrolyte. Specifically, it applies when the diffusion is linear and semi-infinite, e.g. a planar electrode in an infinite electrolyte reservoir [5]. In such cases the diffusion process can be modelled using an element known as the *Warburg impedance*, Z_W , whose impedance is defined as follows:

$$Z_W(\omega) = \frac{\sigma}{\omega^{1/2}} - j \frac{\sigma}{\omega^{1/2}} \quad (4.4)$$

The term σ is known as the *Warburg coefficient* and depends on the concentrations and the diffusion coefficients of the two species involved in the redox reaction [21].

From (4.4) it is apparent that Z_W has equal real and imaginary parts at all frequencies resulting in a slope of 45° in the complex plane, which is characteristic of diffusion-controlled electrode processes [21], and can also be seen in Figure 4.1. At high frequencies Z_W has no effect on the overall impedance since the time scales involved are too short for diffusion to be able to play a role [21].

Another component without electrical counterpart that is frequently used in battery ECMs is the constant phase element (CPE). This accounts for the fact that in real impedance spectra the centre of the semi-circle often lies below the horizontal axis, i.e. the arc of the circle is shifted slightly downwards.

In-depth explanations of the Warburg impedance and the CPE are beyond the scope of this work — more detail can be found in [21] and [4].

This section has given a brief introduction in the use of equivalent circuit models for the purposes of interpreting EIS results. As well as aiding the interpretation of




Figure 4.4: Randles equivalent circuit model.

results, ECMs can be used to model the dynamic behaviour of batteries. It should be noted that other model types exist that can be used for this purpose, e.g. models that are based on the fundamental underlying electrochemical and physical principles of a particular cell type, or purely mathematical models. Whilst further consideration of modelling of battery dynamics falls outside the scope of this work, the field of *System Identification* is suggested as a more formal treatise of the principles of modelling dynamic systems. System identification techniques find use, for example, in the parameterisation of ECMs used in [19]. Further reading can be found in [49, 50].

4.2 EIS-based diagnostic monitoring of LIBs

In Chapter 3 the state of health (SoH) of galvanic cells was defined in terms of capacity fade and power fade, the latter resulting from the effect of increasing internal impedance as cells age. It follows that changes in the SoH will affect the impedance spectrum in some way. However, as this section will show, a cell's impedance spectrum is also sensitive to changes in state of charge (SoC) and temperature. It is for these reasons that EIS is being intensively investigated as a condition monitoring technique for batteries. This section introduces some of the relevant literature in this field, in order to demonstrate the versatility of EIS for the purposes of battery diagnostic monitoring.

4.2.1 State of charge

In [51] impedance spectra for a commercial LIB consisting of two series-connected LiCoO₂/graphite cells were obtained at different SoC values using an electrochemical impedance analyser covering a frequency range from 25 mHz to 100 kHz. The parameters of the identified ECM were obtained using a non-linear least squares fitting procedure in order to analyse their relationship with the SoC. Whilst the majority of parameters showed either very little or non-linear variations with the SoC, three parameters were identified in the frequency range corresponding to the low-frequency (depressed) semi-circle² (25 mHz to 5 Hz) that showed promise with regards to their use as SoC indicators, namely: the real part of the impedance, the phase angle, and the value of an equivalent series capacitance.

The same method was applied in [10] to analyse the relationship between EIS results and SoC for three different commercial lithium polymer batteries. It was found that whilst the impedance between 20 Hz and 1 kHz was largely unaffected

²such as shown in Figures 4.1 and 4.2

by changes in the SoC, the radius of the large semi-circle — residing in the range from 0.1 Hz to 20 Hz and attributed to the combination of charge transfer resistance and double layer capacitance — decreased with SoC. The charge transfer time constant was deemed to be the most accurate SoC indicator in this case.

Two different lithium-based pouch cells were investigated in [52]. The impedance spectra of a $\text{LiNiCoAlO}_2/\text{graphite}$ cell obtained at different levels of SoC revealed two semi-circles that were both attributed to the positive electrode. The low-frequency semi-circle, which was attributed to the charge transfer and double layer reactions, showed a strong non-linear dependence on SoC and increased considerably in size at both SoC extremes. The Nyquist plot of the second cell with $\text{LiFePO}_4/\text{Li}_4\text{Ti}_5\text{O}_{12}$ electrodes consisted of only one semi-circle, as well as a straight line at low frequencies (indicative of a diffusive process). With this electrode material combination the impedance spectra did *not* show strong variation with SoC, except for the spectrum at 100 % SoC, where the slope of the low frequency line increased considerably and the semi-circle showed a small reduction in size. This work highlights the need for impedance characterisation of different types of cells, since not only the shape of the spectrum changes but also its dependency on SoC.

$\text{LiNiMnCoO}_2/\text{graphite}$ pouch cells were considered in [53]. The impedance spectrum between 10 mHz and 10 kHz again resembled that of the Randles circuit and consisted primarily of a straight line at low frequencies and a (depressed) semi-circle at medium-to-high frequencies (ignoring inductive effects at high frequencies, which were attributed to measurement leads). The series ohmic resistance was found to be largely independent of SoC and the charge transfer resistance was found to decrease with SoC. It is argued that it may be possible to use the latter as an SoC indicator, however, it is worthy of note that the variation in the charge transfer resistance was less than 1 milliohm across the entire SoC range (20-100%). This highlights the requirement for high-resolution measurement equipment.

The work discussed so far involved investigations of the relationship between impedance spectra and SoC. There are also some works in the literature which go one step further and estimate the SoC from said relationship. In [54], for example, EIS was used to derive and parameterise a suitable ECM for a $\text{LiNiMnCoO}_2/\text{graphite}$ cell. The model was subsequently employed to estimate the terminal voltage and EMF of the cell. The SoC estimate was based on the EMF-SoC relationship. The error between model and measurement was used in conjunction with a Kalman filter — adapted to deal with the CPE element in the model — in order to improve the SoC estimate accuracy. The impedance model-based SoC estimation showed good agreement ($\pm 1\%$) with the Coulomb counting method that served as reference

(deemed appropriate under the circumstances, given the short experiment duration and knowledge of the initial SoC).

These works demonstrate that in many cases there is a relationship between the SoC of a battery and some of its impedance parameters, and that, even though the relationship is not the same for all batteries, EIS can be used for its determination.

4.2.2 State of health

As explained in Chapter 3, the SoH of galvanic cells can be defined in terms of capacity fade and power fade. Whilst quantification of the former would typically involve a measurement of the full charge capacity, the latter is the result of the inherent impedance growth associated with ageing cells [6] and is therefore quantifiable through impedance measurement.

The two lithium pouch cells investigated in [52] — $\text{LiNiCoAlO}_2/\text{graphite}$ and $\text{LiFePO}_4/\text{Li}_4\text{Ti}_5\text{O}_{12}$ — were subjected to an accelerated ageing procedure, after which their impedance spectra were obtained. A comparison with the EIS results obtained before ageing revealed that the impedance spectra of both types of cells evolved in a very similar fashion: there was a significant increase in both the ohmic resistance and the size of the semi-circle linked to the charge transfer reaction.

Growing ohmic resistance in response to cell ageing has been observed in various other types of LIBs, including: $\text{LiFePO}_4/\text{graphite}$ [6], $\text{LiNiMnCoO}_2/\text{graphite}$ cells [7, 53], and $\text{LiNiCoAlO}_2/\text{graphite}$ [46]. The rise in ohmic resistance was attributed by [6] to the continuous growth of the solid electrolyte interphase (SEI) layer: this causes previously free lithium ions from the electrolyte to be tied up, thus leading to a reduction in capacity and an increase in electrolyte resistance, which is considered to be the main contributing factor to ohmic resistance.

4.2.3 Temperature

The impedance spectra of LIBs are also sensitive to temperature. This is typically explained using the Arrhenius equation, which describes the relationship between temperature and reaction rates [20].

For example, [7] reports an investigation of $\text{LiNiMnCoO}_2/\text{graphite}$ cells in which the charge transfer resistance decreased approximately exponentially as temperature increased. Furthermore, this relationship was shown to be relatively insensitive to ageing. The authors also observed that as the SoH of the cell dropped, the ohmic resistance became more sensitive to changes in temperature.

An increase in both charge transfer resistance and ohmic resistance in response to decreasing temperature was also reported for $\text{LiFePO}_4/\text{graphite}$ cells in [6].

These findings demonstrate that any cell characterisation procedure should include temperature as a variable in order to determine the sensitivity of the cell's impedance spectrum to changes in temperature.

4.2.4 Discussion

The preceding sections demonstrate that in many cases dependencies exist between the impedance parameters of a battery and its SoC, SoH, and temperature, and that EIS can successfully determine the prevailing relationships. There are, however, outstanding challenges in this field. The first is a natural consequence of the fact that there are many different internal and external processes that can affect the impedance spectrum. If their timescales are of similar order, their impedance responses can overlap each other. In the first instance this makes it more difficult to determine cause-and-effect; subsequently, it can also increase the complexity of implementation, due to the added requirement for compensation. The second challenge is that the relationships between impedance and SoC, SoH, and temperature, are often not only non-linear, but also change as the cell ages — as demonstrated, for example, in [7]. The accuracy of any state estimation method that uses EIS only for an initial cell characterisation is therefore bound to decrease as the cell is cycled and/or ages.

Both of the highlighted points call for extensive cell characterisation procedures, before any of the impedance-based relationships can reliably be used for condition monitoring purposes. The second point in particular, demonstrates the requirement for *online* impedance measurements in order to enable state estimation algorithms to adapt to changing battery behaviour.

4.3 Power converter–based EIS

4.3.1 Introduction

Section 4.2 demonstrated the efficacy and versatility of EIS as a condition monitoring technique for LIBs. It was shown that EIS can be used to build behavioural models of LIBs and that it can also help reveal dependencies between the model parameters and the cell's SoC, SoH, and temperature. The fact that these dependencies can change as a cell ages was highlighted as a factor that can make it more difficult to guarantee the accuracy of the state estimation for the entirety of the cell's life. In

principle, conventional EIS — i.e. as introduced in Section 4.1 — could be used to detect these changes. In practice, this would mean — at the very least — having to isolate the batteries from the rest of the system and subsequently finding a way of connecting the measurement instrumentation to them. This is unlikely to be possible without prior partial system disassembly, thus making conventional EIS unfeasible in the majority of adaptive condition monitoring applications for LIBs. The obvious solution would be to implement EIS-capable instrumentation in the system, but this is (again) unlikely to find widespread adoption, due to the high economic cost associated with the additional hardware requirements. To summarise, conventional EIS has two major drawbacks that prevent it from being widely used in commercial battery monitoring applications: 1) it is an *offline* method, i.e. during measurement the cell-under-test is isolated and can not be used for its intended purpose in a given application, and 2) the measurement instrumentation used for EIS is typically of laboratory-grade and its associated cost is accordingly high.

Power converter-based EIS is a relatively new way of implementing impedance spectroscopy that can overcome the drawbacks of conventional EIS described above. The fundamental principle is to introduce the measurement perturbation required for the impedance determination through control of the power converter. Power converters are used in almost all electrical devices as they perform a very elementary task: they take an input with a given set of characteristics (e.g. ac/dc signal, voltage and current amplitude, frequency (if ac), regulated/unregulated) and produce from it a (typically) regulated output that matches the requirements of a particular application in terms of the stated characteristics. For example, a single-phase power supply could use a power converter that takes the fluctuating ac line voltage as input (e.g. 230 V, 50 Hz) and produces from it a regulated 5 V dc output. Power converters are also used in battery-powered systems for various purposes, e.g.: to provide a regulated dc output voltage from the unregulated battery voltage, in which case the converter may be referred to as a dc-dc converter; to provide a regulated output for battery charging; or, to convert the dc battery voltage into a single or three-phase ac output, in which case the converter may be referred to as a dc-ac converter, or inverter³.

The ubiquity of power converters gives power converter-based EIS the immediate advantage of not requiring additional signal generation circuitry, thus providing a good foundation for a cost-effective implementation. In addition, because the hardware and software required to carry out the EIS measurements can be permanently

³Often the first stage of the dc-ac conversion process also involves a dc-dc converter, which is one of the reasons why many works tend to focus on dc-dc converters.

built into the system, there is no need for the cell-under-test to be isolated, or indeed for the system to be powered down during measurements, as the measurements can be performed *online* during steady-state system operation.

Section 4.3.2 explains how power converter–based EIS works and Section 4.3.3 discusses relevant literature, including outstanding challenges in the field.

4.3.2 Principle of operation

Modern power converters make use of semiconductor devices, e.g. metal–oxide–semiconductor field-effect transistors (MOSFETs) and insulated-gate bipolar transistors (IGBTs), that operate as switches, i.e. they are either fully on or fully off [55]. The state of the switches, i.e. whether they are open or closed and for how long, determines the output the converter. Converters operating in this way are therefore also known as switch-mode converters. The switching frequencies tend to be relatively high (in the tens or hundreds of kilohertz, or higher) in order to reduce the size (and therefore cost) of the transformers and filters required for their operation [55].

Two different methods to control the output of a converter are generally distinguished. In the first, the switching frequency is kept constant, whilst the proportion of the switching period during which the switch is closed — the *duty cycle* — is varied. This is also known as *pulse width modulation* (PWM) control, since the waveforms used to turn the switches on and off are pulse-like (or rectangular), and variation of the duty cycle thus changes the pulse width. In the second method, both the switching frequency and the time the switch is closed for may be varied [55]. This method is much less commonly used and this work therefore focuses on the PWM control method.

As an example of how the converter can be controlled using PWM, consider the idealised (lossless) steady-state input-to-output voltage transfer function of a step-up dc-dc converter (also known as *boost converter*) operating in continuous conduction mode:

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{1}{1 - D} \quad (4.5)$$

This shows that the output of the converter can be controlled by varying the duty cycle, D . Equation (4.5) also demonstrates that if the input voltage varies and the duty cycle remains unchanged, the output voltage will change as a result. In more general terms: the converter output will change if there are fluctuations in either the input, or the load. As this is usually undesirable, it is common to implement a mechanism that *regulates* the converter output against such fluctuations. Typically,

this is done using a control loop with negative feedback. This concept is explained using the example of a switching dc-dc converter with a regulated output voltage — illustrated in Figure 4.5. Its operation could be described as follows: the converter's output voltage is measured and compared to the desired — or *reference* — voltage using an error amplifier. The outcome of this comparison (commonly termed the *error* signal) is then used to vary the duty cycle, which in turn will change the output voltage. By the principles of negative feedback, the error amplifier wants to drive the difference between its two inputs to zero, so that — over a number of switching cycles — the converter's output voltage will tend towards the value of the reference and the error will tend towards zero. If a time-varying reference is used, the output voltage will effectively follow the reference.

Whilst in the example above the controlled parameter was the output voltage, it is possible to base the regulation on other variables, such as the input or output current, for example.

Using the concepts introduced in this section, it is now possible to outline how the EIS measurement perturbation can be injected using a power converter. In fact, two different methods are found in the literature. In the first — which will be referred to as the *open-loop* method — the desired perturbation is introduced by directly varying the duty cycle. As an example, if a sinusoidal perturbation is desired, the duty cycle could be varied in a sinusoidal fashion (at the required frequency) around its steady-state value. In the second method — which will be referred to as the *closed-loop* method — the measurement perturbation is added to

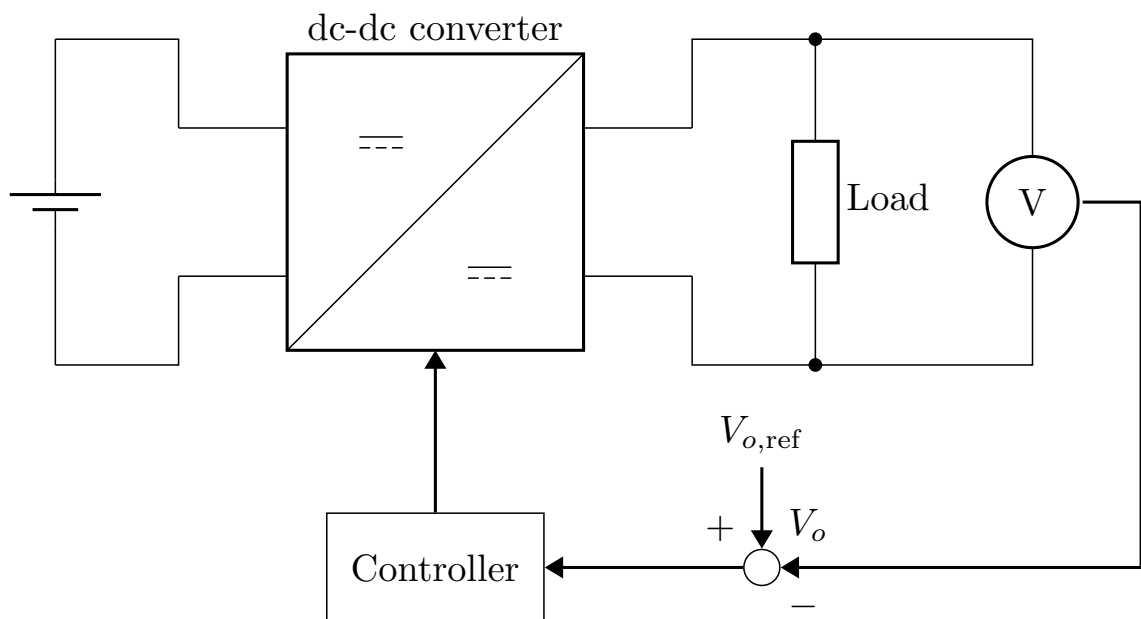


Figure 4.5: Regulation of power converter output voltage.

the error amplifier reference signal. Figure 4.6 illustrates the principle of the closed-loop method, again using the example of a sinusoidal measurement perturbation. The reference signal, i_{ref} , can be seen to consist of a dc component, I_{dc} , and a sinusoidal component with peak amplitude \hat{i}_{ac} . The figure also shows that it is the battery *current* that is being controlled, i.e. the equivalent of galvanostatic EIS is performed (the cell voltage is measured in response to a perturbation in the cell current). Due to the generally very low impedance of LIBs, this is the preferred implementation, as a very small error in a voltage reference would lead to a very large current response.

Figure 4.6 shows the implementation of power converter-based EIS in a *discharge setup*, i.e. the battery pack is connected on the input side of the converter and is discharging. It is of course also possible to use a *charge setup*; in this case the battery pack is connected to the *output* of the converter and is charging. The implications associated with the use of each setup are discussed in Section 4.3.4.

Whilst initially it may be simpler to introduce the measurement perturbation using the open-loop method, its main drawback is that it is not as stable as the

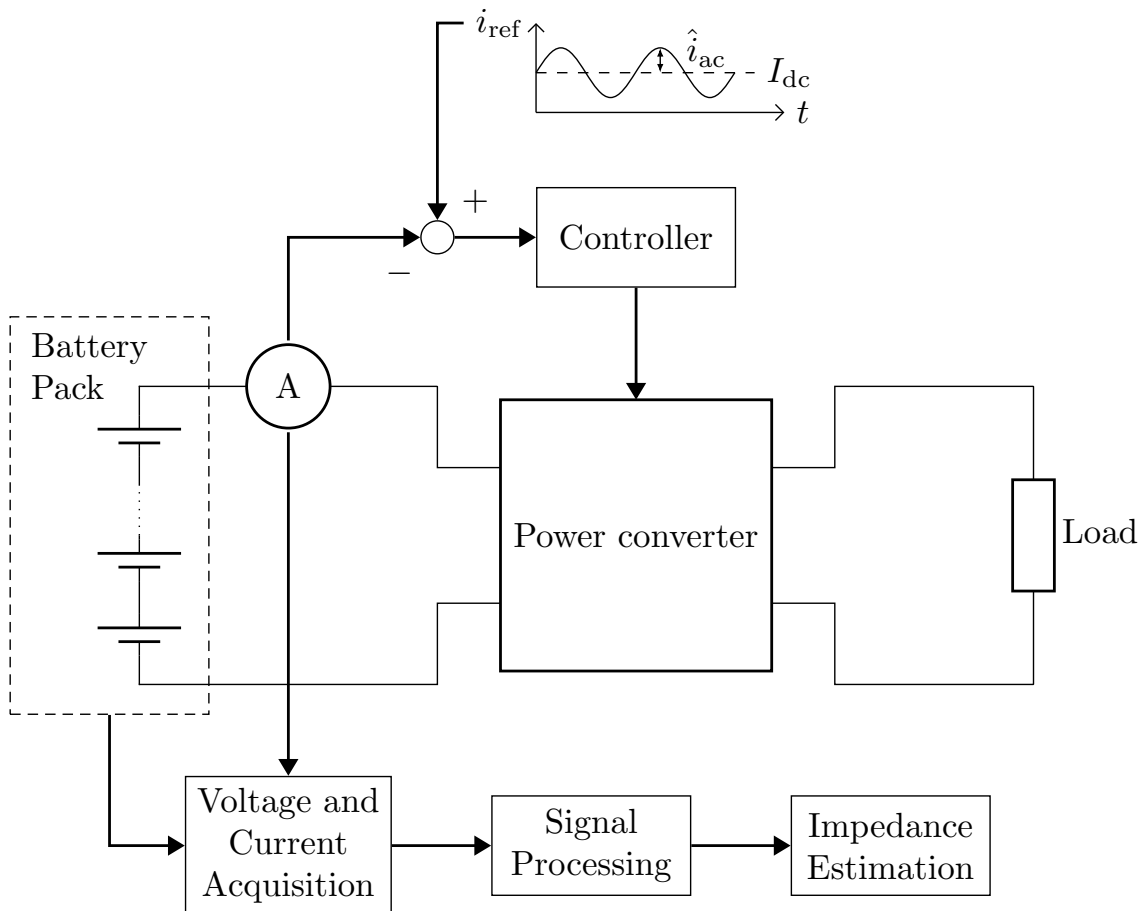


Figure 4.6: Principle of closed-loop power converter-based impedance spectroscopy.

closed-loop implementation, because it does not use feedback, i.e. it does not self-correct. Another disadvantage of the open-loop method is that it is not as accurate as the closed-loop method. As discussed in Section 4.1.2, the amplitude of the measurement perturbation should be chosen in such a way that a linear response and a good SNR are achieved. In other words: it is crucial to set the perturbation amplitude with a relatively high degree of accuracy. With the open-loop method this is not necessarily possible under all circumstances, as the relationship between the converter's duty cycle and its input current depends not only on frequency and the converter's operating point, but could also vary with time and external parameters, such as temperature. The relationship is, in general, non-linear; the open-loop method may therefore introduce significant distortion in the current waveform. The advantage of the closed-loop method is that the measurement perturbation can be defined in terms of the current, rather than the duty cycle. This has the benefit that the perturbation definition does not need to be adapted to the operating point of the converter. It may still be necessary to adapt the perturbation amplitude to frequency, depending on how much the impedance varies throughout the spectrum's bandwidth. One drawback of the closed-loop method is the increased control complexity. As described earlier, it is often desirable to regulate (for example) the output voltage of the converter; if, in order to introduce the measurement perturbation, the input current of the converter also needs to be controlled, this requires simultaneous control of two quantities. This can call for a fast control algorithm, which may increase the required processing power.

As well as the two *active* implementations, i.e. where the measurement perturbation is actively generated, *passive* implementations of power converter-based EIS can also be found in the literature, e.g. [16, 17]. In this type of implementation, instead of deliberately injecting the measurement perturbation as a new signal, the impedance estimation is based on signals that are already present in the circuit, e.g. switching ripple.

There are also other passive EIS implementations in the literature, in which the impedance is estimated from other signals present in the circuit, e.g. noise. These are not necessarily power converter-based and therefore fall outside the scope of this work.

4.3.3 State of the art

In [11] an open-loop implementation is used in conjunction with a dc-dc buck-boost converter. A single sine measurement perturbation is introduced by direct modulation of the duty cycle. Aside from the limitations associated with the open-

loop implementation (discussed in Section 4.3.2), the main constraint of this work is the limited bandwidth, as the lowest measurable frequency is 100 Hz. In addition, the experimental test setup is confined to a single lithium-ion cell.

The same open-loop approach is used in [13] with a boost converter to measure the impedance spectra for individual LiFePO_4 and NiMH cells. This work derives analytically the relationship between the duty cycle perturbation amplitude and the resulting current ripple in the boost inductor.

A boost converter is also used in [12], but in this case a closed-loop implementation based on a PI controller (proportional plus integral control) is employed to introduce the measurement perturbation into the battery current. The reference waveform consists of a multisine signal made up of three sinusoids superimposed on a dc component. This work highlights one of the difficulties associated with online EIS measurements, particularly when they involve very low frequencies. In order to accurately measure the amplitude of the voltage and current perturbations, it is necessary to capture at least one full cycle, i.e. the lower the measurement frequency, the longer the acquisition takes. Measurement perturbation waveforms tend to be charge-neutral, i.e. the charge moved out of the battery in one half-cycle is equal to the charge moved back in again in the next half-cycle; this is not the case for the dc current component present during online EIS. In other words: during online EIS there is a net change in the battery's charge, i.e. the SoC changes. Any change in SoC is accompanied by a change in battery EMF, and the magnitude of the change increases with increasing measurement acquisition duration, or decreasing measurement frequency. If this drift in the voltage is not compensated for, this can result in significant spectral leakage appearing in the frequency domain. This in turn can lead to a substantial error in the voltage estimate that propagates into the impedance estimate. In [12] the drift is compensated for by using a linear least squares fitting in the time domain. The main limitations in this work are the bandwidth (0.1 Hz to 10 Hz) and the scale of the experimental system (single lead-acid battery).

Another example of a closed-loop implementation is presented in [19]. Here a multisine perturbation consisting of 21 components is added to the output current of a buck converter. Experimental verification is limited to a single cylindrical LIB and the lowest measurement perturbation frequency used is 1 Hz.

Similarly, in [18] the output current of a synchronous buck converter is perturbed using closed-loop control. What sets this work apart is that the implementation is verified experimentally on a slightly larger scale using a 40 Ah, 13.6 V LIB. The acquired impedance spectra are compared to a commercial EIS measurement instrument with relatively good agreement (average impedance magnitude error of

less than $1\text{ m}\Omega$ and average phase error less than 2° in the range 0.1 Hz to 100 Hz). Whilst the battery employed in this work is larger in terms of its capacity and voltage than those found in other works, it does still only consist of a single module.

A different implementation of the closed-loop method can be found in [14], where the measurement perturbation is introduced into the output voltage of a buck-boost converter. A squarewave perturbation is used, which has the advantage of exciting the system at multiple frequencies simultaneously (i.e. at the fundamental and odd harmonics thereof), but, as discussed in Section 4.1, has the disadvantage that the SNR decreases with increasing harmonic number, due to the decrease in amplitude. As a consequence, multiple perturbations with different fundamental frequencies are required to cover the entire LIB impedance spectrum bandwidth. This increases the signal generation complexity, but overall the impedance spectrum acquisition is still faster than using the single sine perturbation. Whilst controlling the converter's output voltage has the advantage of simplifying the control architecture, it has a number of significant disadvantages. The first is that the relationship between the perturbation in the output voltage and the perturbation in the battery voltage (connected at the converter input side) is not straightforward. Whilst they are related by the converter's voltage transfer function, this depends on the operating point, which should therefore also be taken into account. This is further complicated by the fact that the amplitude of the perturbation in the cell voltage should be relatively small in order to fulfill the requirements of generating a linear response and good SNR (as discussed before, an upper limit of 10 mV is typical). Even if the converter steps up the voltage from input to output (and especially at lower duty ratios), this would still require a very accurate control of the output voltage. Due to the generally very low impedance of LIBs, a small error in the voltage control could lead to much larger variations in the current. Also, if the battery consists of multiple cells in series, there is no way of guaranteeing that the amplitude of the voltage perturbation across each cell will meet the required criteria, because the cell impedances will not be identical. Another disadvantage of injecting the measurement perturbation in the output voltage is that many applications require regulation of the output voltage *against* fluctuations. This would be particularly problematic for larger scale battery packs, since the fluctuations in the output voltage are proportional to the number of series-connected cells. A method is proposed in [14] to cancel the output voltage ripple by using *two* power converters and phase shifting their reference waveforms by 180° relative to each other. This solution, however, is not universally applicable. The most significant limitations in this work lie in the bandwidth — due to a low-frequency limit of 10 Hz — and the scale, as the largest

experimental setup used involves only two series-connected cells.

Due to the disadvantages associated with introducing the measurement perturbation in the voltage, it is deemed advantageous to control the battery current. Yet, (as argued above), in many applications it is also necessary to regulate the output voltage. In order to allow control of both — battery current and output voltage — more complex control architectures are required than those considered so far. A promising approach is presented in [56], where sliding mode control and a PI controller are used to both regulate the output voltage of a six-phase interleaved boost converter and introduce a single sine measurement perturbation in the input current. Whilst the suggested converter is designed for a high-power application with input and output voltages of 70 V and 350 V, the approach is only verified through simulation and has a low-frequency limitation of 1 Hz.

The key characteristics of the various active measurement perturbation injection implementations discussed above are summarised in Table 4.1.

Passive measurement perturbation injection methods were briefly introduced at the end of Section 4.3.2. One example of such an implementation is given in [16], where effectively the power converter switching ripple is used as the measurement perturbation. Due to the fact that switched converters normally operate at frequencies ranging from tens to hundreds of kilohertz, this implementation has inherent bandwidth limitations. Additionally, it mandates the use of acquisition hardware capable of high sampling rates. In [17], the impedance at the converter switching frequency is determined using the input-side switching ripple. This is then used to estimate the internal cell temperature with a reported estimation error of less than 10 °C. Whilst this demonstrates that there may be some applications in which this approach could be used, the bandwidth limitation makes passive power converter–based EIS unsuitable for most LIB condition monitoring purposes.

Table 4.1: Key characteristics of prior power converter-based EIS implementations using active measurement perturbation injection.

Ref	Type	Controlled quantity	Converter topology	Battery location	Battery type	Perturbation type	Frequency range
[11]	open-loop	duty cycle	bi-directional boost/buck	input	1x 2.6 Ah 18650 cylindrical lithium-ion cell	single sine	100 Hz to 10 kHz
[13]	open-loop	duty cycle	boost	input	1x 2.5 Ah cylindrical LFP cell	single sine	1 Hz to 2 kHz
[12]	closed-loop	input current	boost	input	1x 7 Ah lead-acid battery	multisine (7x components)	0.1 Hz to 10 Hz
[19]	closed-loop	output current	buck	output	1x 2.6 Ah 18650 cylindrical lithium-ion cell	multisine (21x components)	1 Hz to 2 kHz
[18]	closed-loop	output current	synchronous buck	output	1x 40 Ah 13.6 V LIB module	single sine	0.1 Hz to 5 kHz
[14]	closed-loop	output voltage	bi-directional boost/buck	input	2x series-connected 2.6 Ah 18650 cylindrical lithium-ion cells	square wave	10 Hz to 9 kHz
[56]	closed-loop	input current & output voltage	6-phase interleaved boost	input	n/a (fuel cell)	single sine	1 Hz to 10 kHz

4.3.4 Outstanding challenges

Section 4.3.3 has shown power converter–based EIS to be an effective method to acquire the impedance spectra of LIBs for condition monitoring purposes. It allows the impedance to be determined online, without the need to disconnect the battery from the wider system, thus opening up the possibility of applying EIS-based condition monitoring in a greater range of applications.

The review of relevant literature has also demonstrated that there are still challenges to be solved in relation to the implementation of power converter–based EIS. In particular, three areas have been identified as showing scope for development: 1) widening the measurement bandwidth; 2) verifying the feasibility of the method in larger scale systems; and 3) reducing the effects of the measurement perturbation on the rest of the system.

Whilst commercial EIS instruments deliver impedance measurements at frequencies as low as μHz , most power converter–based EIS implementations found in the literature only go down to the order of 0.1 Hz (at best). These *low* frequency bandwidth limitations are primarily associated with the difficulty of ensuring that the system remains in steady state for the duration of the measurement, which is a pre-requisite to ensure the measurement’s validity [21]. The largest influence in this respect, comes from the fact that power converter–based implementations operate *online*, i.e. whilst the battery is working. The dc current can lead not only to variations in temperature but also in SoC. Changes in either of these parameters can affect the impedance spectrum, thus disturbing the steady state and invalidating measurement results. Many of the works described in Section 4.3.3 deal with this by restricting impedance measurements to relatively short timescales (i.e. high frequencies), at which the variations can usually be neglected. Some works, e.g. [12,18], attempt to compensate for SoC-induced variations in the battery EMF in order to extend the low-frequency limit. One of the objectives of this project is to propose solutions to the bandwidth limitations associated with power converter–based EIS and the aforementioned compensation techniques have been identified as a specific area with scope for development.

Regarding the scale of the system: whereas experimental setups in the literature most commonly involve only one or two cells, batteries employed in electric vehicles and battery energy storage systems can easily consist of hundreds of cells in both series and parallel arrangements. The magnification in scale required to implement power converter–based EIS in such systems creates new challenges that have so far not been sufficiently addressed in the literature, e.g. the requirement for improved specifications of voltage and current transducers and ADCs. Such a scale-up can

also have knock-on effects on the *high* frequency bandwidth limitations, as the large number of measurement channels may need to be acquired sequentially, thus lowering the effective sampling rate. One of the objectives of this work is therefore to create an experimental test bed for power converter-based EIS that consists of a cell-arrangement in the double figures, i.e. greater than any of the existing implementations in the literature, and thus evaluate the effectiveness of the method in monitoring a battery pack of such scale.

Finally, the effects of the measurement perturbation on the rest of the system — and how to reduce them — have so far not been adequately addressed. One of the reasons behind this is that power converter-based EIS is very often implemented only in a charge setup, i.e. with the battery connected on the output side of the converter and (typically) a power supply connected on the input side, e.g. [18, 19]. In the charge setup the effect of the measurement perturbation on the rest of the system is arguably not as significant as it is in the *discharge* setup, where the battery is connected on the converter's input side and a load is connected to its output (as shown in Figure 4.6 on p. 53). This is because the discharge setup involves a separate load; regulation of the output is typically required and any remnants of the measurement perturbation appearing at the output would therefore directly interfere with this requirement. Another objective of this project is therefore to develop strategies to reduce the effect of the measurement perturbation on the load.

The previous point also highlights another challenge associated with the implementation of power converter-based EIS in the discharge setup: whereas in the charge setup it is possible to inject the perturbation in the output current and control/regulate the output current using a single control loop, the control strategy required for the discharge setup is more complex, since simultaneous control is required of both the input current (to inject the measurement perturbation) and the output voltage/current (to regulate the output). Of those works in 4.3.3 that use a discharge setup, the only one that incorporates closed-loop control of both the battery current and the output voltage is [56], however, this does not include experimental validation. A method that uses closed-loop control to both inject the measurement perturbation and regulate the output is also investigated in this work.

Chapter 5

Methodology

5.1 Introduction

The outstanding challenges of power converter-based EIS were highlighted in Section 4.3.4. This led to the formulation of the three main objectives of this project:

1. To prove the feasibility of monitoring a 16-module lithium-ion battery pack
2. To propose solutions to existing bandwidth limitations of power converter-based EIS
3. To develop strategies to reduce the effect of the excitation signal on the load

These objectives are relatively broad, both as a whole, because they address very different aspects of power converter-based EIS, as well as individually, since for each objective to be met, different parts of the system need to be developed and adapted to one another. Because the objectives address such varied parts of the power converter-based EIS implementation, an experimental methodology was deemed to be the most suitable approach. This enabled many parts of the system to be developed independently, before the integration of all sub-systems into the complete power converter-based EIS implementation.

The impedance spectrum acquisition process using power converter-based EIS can broadly be broken down into three main parts: 1) generation of the measurement perturbation (control), 2) signal acquisition, and 3) signal processing. The rest of this chapter discusses in greater detail the methods used in each of these parts.

5.2 Signal acquisition

To obtain the impedance of a battery it is necessary to acquire both the voltage across it and the current through it. Additionally, if power converter-based EIS

is implemented using closed-loop measurement perturbation injection, the battery voltage or current (depending on which one is being controlled) will also need to be acquired so that it can be fed back to the controller. Similarly, to enable the converter output to be regulated, either the output voltage or current will also need to be acquired. These signals need to be digitised to allow the signal processing steps required to determine the battery impedance, but also to enable their use in a digital controller, if this type of controller is used.

The next section covers general considerations regarding the use of analogue-to-digital converters (ADCs) in a power converter-based EIS application. This is followed by a more detailed breakdown of the considerations required in the acquisition of current and voltage signals.

5.2.1 Analogue-to-digital converters

The number of input channels that the ADC is required to have in order to allow for the impedance determination of all cells in a battery pack depends on the size and configuration of the particular battery pack in use. In a series connection of multiple cells, the current is the same in all cells, but the voltage differs; in a parallel connection, on the other hand, the voltage across all cells is the same, but the current differs. A battery pack with a n_s S n_p P connection topology, i.e. consisting of n_s series-connected cells and n_p parallel-connected cells, therefore requires n_s channels for the voltage acquisition and n_p channels for the current. In this work a battery pack consisting of 16 series-connected cells is used; the ADC must thus provide 17 input channels for the impedance determination alone. Additional channels are required to enable the implementation of the closed-loop control.

The required ADC input type, i.e. whether it is single-ended or differential, also depends on the size and configuration of the battery pack, as well as the approach chosen to prepare, or *condition*, the signals before they are acquired. This is discussed in greater detail in the two sections that follow.

The maximum sampling rate that the ADC must support largely depends on the required impedance spectrum bandwidth. Depending on the specific implementation of the converter, e.g. if multiple input channels use the same ADC via a multiplexer, the size and configuration of the battery pack may also become important.

Finally, the bit depth of the ADC must be large enough to ensure a sufficiently high measurement resolution, i.e. it depends on the impedance magnitude to be measured, the desired SNR, and similarly to the input type, it also depends to a large extent on the particular signal conditioning approach used.

The specific considerations that must be made with regards to the acquisition of

current and voltage signals in a power converter–based EIS application are examined in the following sections.

5.2.2 Current

To understand the requirements for the current acquisition, it helps to consider some of the general characteristics of the signals to be measured. The current signals will normally consist of a dc component as well as the injected ac component, i.e. the measurement perturbation. The first requirement is therefore that the current transducer must be capable of measuring both dc and ac signals. This excludes current transformers, due to the presence of the dc component. As one of the objectives of this project is to monitor a battery pack with a peak power capability of over 5 kW, the current transducer should ideally also be capable of measuring a maximum current of around 80 A. Due to the relatively high power of the battery pack, it would be advantageous if galvanic isolation was provided between the battery pack and the acquisition-and-control circuitry — for both reasons of safety and noise suppression. Finally, as EIS measurements for battery monitoring applications typically involve frequencies up to tens of kHz, the transducer should have a bandwidth that is wide enough to accommodate this.

Given the requirements listed above, closed-loop Hall effect current transducers have been determined to be a suitable choice for this application. In this type of transducer, the current that is to be measured is determined via the magnetic field that it generates. Also known as *compensated* or *zero flux* transducers, closed-loop Hall effect current transducers do not directly amplify the Hall voltage — as is done in *open-loop* Hall effect transducers — but instead, they use the Hall voltage to generate a current in a secondary winding, which exactly cancels the flux in the magnetic core generated by the current in the primary winding. This leads to a reduction in the temperature-dependent gain drift and therefore a reduced measurement uncertainty compared to the open-loop design [57].

The more commonly employed current shunt resistors have some advantages over closed-loop Hall effect transducers, namely: they are more cost-effective and in some respects are easier to implement (because they do not require any auxiliary circuitry); however, they also have several drawbacks that make them less suited to this particular application. Given the relatively high currents to be measured, shunt resistors suffer from higher power losses than Hall effect transducers. The high power dissipation can also lead to measurement inaccuracies due to the shunt’s temperature coefficient of resistance. (Whilst this could be compensated for, this may negate any supposed advantage in terms of cost-effectiveness and ease of implementation.) In

addition, the high-frequency performance of current shunts can be limited due to parasitic inductance and the skin effect. The inductance in the circuit is further increased by the need for a loop connection for voltage sensing across the shunt. Finally, current shunt resistors do not provide any electrical isolation.

Based on the above comparison of current shunt resistors and closed-loop Hall effect transducers, the latter have been determined as the preferred choice in power converter-based EIS applications.

Before feeding the current transducer output to the ADC it may be necessary to filter the signal in order to attenuate any unwanted high frequency noise. If this is done, care must be taken to design the filter in such way as not to affect either the magnitude or phase of any of the measurement perturbation components.

5.2.3 Voltage

In order to determine the impedance of a cell, it is necessary to measure the voltage across the cell. Similarly to the current, the voltage consists of a dc component with a superimposed ac component, i.e. the response to the measurement perturbation. As discussed before, the amplitude of the voltage response should be limited to approximately less than 10 mV, in order for the linearity assumption to hold. The dc component, on the other hand, will lie somewhere in the range 2 to 4 V for most types of LIBs. If this signal were to be fed directly into the ADC, in most cases (not considering high-end ADCs) this would result in a poor resolution relative to the amplitude of the ac component. In any case, the input voltage range of the ADC would not be very well utilised, because most of it would be taken up by the dc component, despite not being required to determine the impedance. Thus, before feeding the cell voltage to the ADC, it should undergo a signal conditioning process consisting of the removal of the dc component and the amplification of the ac component¹. Both of these steps could be achieved through the use of a differential amplifier.

One of the objectives of this project is to achieve a scaled-up implementation of power converter-based EIS. When cells are connected together to form battery packs, this creates an issue that complicates the voltage acquisition: with each additional cell that is connected in series, the *common-mode voltage* rises, i.e. the average potential of the cell terminals with respect to the circuit common (or ground) terminal. This is illustrated in Figure 5.1, where V_{CMi} denotes the common-mode

¹Whilst the effective resolution of the current measurement could be improved in the same way, the ratio of ac to dc components is not as small as for the voltage. Also, whilst not required for the impedance determination, the dc component of the current *is* required for the closed-loop control.

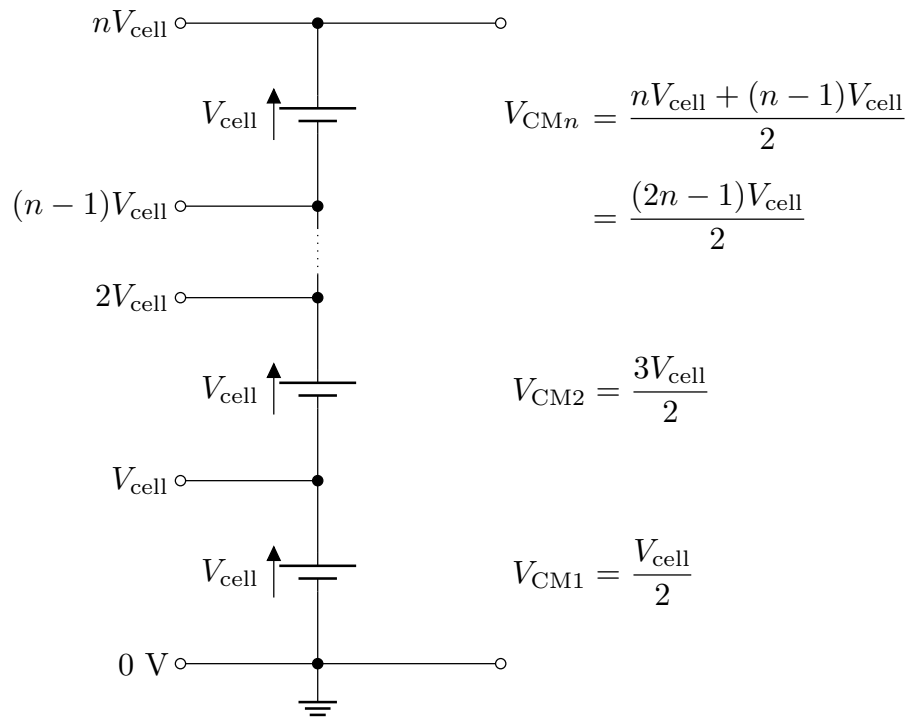


Figure 5.1: Common-mode voltage in a series-connection of cells.

voltage of cell i and, for the purposes of this example, all cell voltages are assumed to be identical and equal to V_{cell} .

Common-mode voltages can be a problem for many integrated circuits, as they are typically rated for a given maximum common-mode voltage that they are able to withstand before failure, and this rating is usually close in value to that of the supply voltage. In this project a battery pack is used that consists of 16 series-connected cells, each with an average operating voltage of 3.2 V; the common-mode voltage for the cell at the high-potential side of the battery pack will thus be almost 50 V. This is too high for most ADCs, and whilst there are some² differential amplifiers that can tolerate such high common-mode voltages, the decision was made to use *isolation* amplifiers instead. Not only are isolation amplifiers capable of handling the prevailing common-mode voltages, but they also have the additional advantage over differential amplifiers of providing galvanic isolation between their input and output. Considering that the chosen type of current transducer also offers isolation, using isolation amplifiers in the voltage acquisition enables complete separation of power and measurement common references, which helps to prevent ground loops and improves noise immunity. In addition, isolation amplifiers are capable of meeting the bandwidth requirements associated with EIS for battery monitoring applications.

Similarly to the current acquisition, it may be necessary to filter the voltage

²e.g. INA149, LT1990.

signals before feeding them to the ADC. Again, this should be done without attenuating or delaying any of the measurement perturbation components.

5.3 Signal processing

Once the raw voltages and currents have been acquired and converted into digital signals by the ADC, they have to undergo a number of signal processing steps before each cell's impedance spectrum can be determined. These steps are summarised in Figure 5.2 and described in more detail in the following sections.




Figure 5.2: Summary of signal processing steps.

5.3.1 Reversal of acquisition-related conversions

The first step is to reverse the acquisition-related conversions. In the acquisition of any measurand, the *actual* and the *acquired* values of the measurand seldomly correspond exactly to one another. This is because the device used for the acquisition will be rated to operate correctly only over a certain range of input values. Depending on how the range of the measurand compares to this input range, it may either need to be amplified, or attenuated. For example, if the measurand was a voltage ranging between -1 V and $+1$ V and the input voltage range of the ADC was -5 V to $+5$ V, it would be advantageous to amplify the voltage signal before its acquisition, in order to increase the sensitivity of the measurement. Having amplified the signal, it would then be necessary to reverse this operation in order to get back to the original signal. This is what is meant by reversing acquisition-related conversions.

In this application, a number of conversions are required during the acquisition of both current and voltage. The previous section described the need for the cell voltages to undergo two signal conditioning steps before being fed to the ADC: the removal of the dc component and the amplification of the ac component. Both of these steps need to be reversed in order to arrive back at the actual cell voltage. Similarly, the current acquisition using the closed-loop Hall effect transducer involves

a current-to-voltage conversion that needs to be reversed in order to obtain the actual current. The particular operations and operands required to reverse the voltage and current acquisition-related conversions depend on the specific implementation, e.g. the transducers and acquisition devices used, and are described in Chapter 6.

5.3.2 Signal length adjustment

To obtain the cell impedance, the Fast Fourier Transform (FFT) of the voltage and current signals needs to be determined. Before this can be done, however, it is necessary to adjust the length of the signals in order to minimise the amount of *spectral leakage* in the FFT. Spectral leakage can occur³ when the observation window, t_{obs} , i.e. the length of the time-domain data processed in the FFT, contains a non-integer number of cycles of the signal frequency. Viewed in the frequency domain, this is equivalent to the signal frequency not being a whole-number multiple of the frequency resolution, f_{res} , i.e. the inverse of the observation window ($f_{\text{res}} = 1/t_{\text{obs}}$). Under these conditions, the signal frequency does not exactly match any of the available frequency bins in the FFT and its energy is thus spread over multiple bins (those nearest in value to the signal frequency). Figure 5.3 shows an example of spectral leakage due to incorrect signal length adjustment.




Figure 5.3: Spectral leakage caused by incorrect signal length adjustment.

The occurrence of spectral leakage causes errors in the estimation of frequency, amplitude, and phase, and should therefore be minimised. To reduce spectral leak-

³Spectral leakage can be caused by other factors, e.g. noise, or a drift in the signal (as explained in the following section), however these are not considered in this section.

age, the observation window length should be chosen so that it contains an integer number of cycles of all frequencies to be analysed; however, in many situations the precise frequency content of signals is not exactly known and this can therefore not be achieved. In such cases, it is common to make use of *window functions* in order to reduce leakage. These have some drawbacks, including: they require a longer observation window to achieve the same frequency resolution as an unwindowed acquisition; they increase the required data processing volume compared to a simple signal length adjustment; and whilst they can reduce spectral leakage, they do not completely remove it and are not as effective as signal length adjustment.

In an EIS application, the frequency content of signals *is* known and spectral leakage can therefore be avoided by correct adjustment of the signal length.

The number of cycles of a given frequency component of frequency f_i and period T_i that fit into t_{obs} is given by c_i :

$$c_i = \frac{t_{\text{obs}}}{T_i} = t_{\text{obs}} \cdot f_i = \frac{f_i}{f_{\text{res}}} \quad (5.1)$$

Therefore, to reduce spectral leakage, c_i should be an integer for all frequencies to be analysed. This requirement places some restrictions on the values that can be chosen for t_{obs} and f_i .

A separate (but related) requirement is that the volume of data that needs to be handled should be kept as low as possible; this is to reduce the data storage and processing capability required by the signal processing system. Therefore, t_{obs} should be minimised (or equivalently: f_{res} should be maximised). The lowest value that t_{obs} can take, whilst still fulfilling the requirement that c_i should be an integer, is equal to the period of the lowest frequency component to be analysed, T_L . Setting $t_{\text{obs}} = T_L$ in (5.1):

$$c_i = \frac{T_L}{T_i} = T_L \cdot f_i = \frac{f_i}{f_L} \quad (5.2)$$

If in (5.2) we let $f_i = k \cdot f_L$ with $k \in \mathbb{N}$, then the requirement that c_i should be an integer is fulfilled for all values of k .

To summarise, we can fulfill the two requirements, namely: 1) minimising the volume of data that needs to be processed, and 2) minimising the error in the impedance estimation due to spectral leakage, by: a) setting the observation window length to be equal to the period of the lowest frequency to be analysed, and b) making all other frequencies at which the impedance is to be determined integer multiples of the lowest frequency component.

5.3.3 Compensation for drift in voltage

Once the signal length has been adjusted to be equal to one period of the lowest frequency, the next step is to compensate for the drift in the voltage. As explained in Chapter 4, when EIS is performed in the presence of a dc current component, this will cause a net change in the battery's SoC during the spectrum acquisition. Due to the EMF-SoC relationship, the change in SoC will lead to a change in the battery's EMF. Depending on the magnitude of the SoC change and the sensitivity of the battery's EMF with respect to changes in SoC, the resulting change in the EMF can lead to significant spectral leakage appearing in the FFT. The greater the voltage drift, the greater the spectral leakage; the issue is therefore exacerbated by large dc currents and low measurement frequencies (which require long observation windows). An example is given in Figure 5.4; this shows the magnitude spectrum of a multisine signal consisting of five frequency components (the lowest being at 10 mHz, thus requiring an observation window length of 100 s), with and without a linear drift of -1 mV/s applied to the waveform. Whilst the spectral leakage is worst at low frequencies, it can be seen to extend through several frequency decades, thus potentially affecting a substantial portion of the impedance spectrum. Compensating for drifts in the cell voltages is therefore an important step to ensure the accuracy of the impedance estimation.

Different approaches exist through which the drift in the voltage can be compensated. These can broadly be split into time and frequency domain methods, depending on which domain the compensation takes place in.

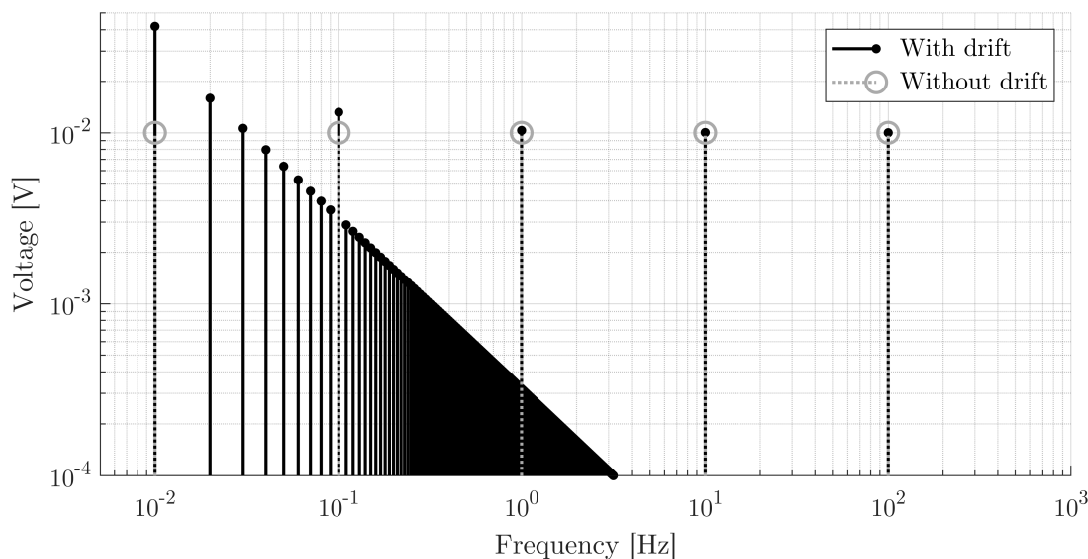


Figure 5.4: Spectral leakage caused by uncompensated drift in voltage waveform.

In [18] a time domain method is used that consists of subtracting a first order polynomial from the voltage signal. The slope of the polynomial is determined by calculating the change in both voltage and time from the first to the last sample of the observation window. Whilst this offers a very simple implementation that requires very little processing, the method has some drawbacks. For a start, it assumes that the drift in the voltage is linear, which is not necessarily always the case. In addition, whilst theoretically correct, in practice the method may not be very accurate; for example, any noise superimposed on the voltage could significantly skew the results, or introduce a drift where there was none to begin with. This is described in greater detail in Section 5.3.6.

A different time domain approach is used in [12]: here a linear least squares fitting of the voltage signal is performed. This method also assumes a linear drift and whilst it requires more processing power than the previous method, it does offer greater robustness when dealing with noisy signals.

Finally, a frequency domain compensation method is explained in [58]. Here, the assumption is made that if spectral leakage has occurred, then its contribution to both the real and imaginary values at the measurement frequency is equal to the arithmetic mean of the values at the two frequency components directly adjacent to the measurement frequency (i.e. the components directly above and below the measurement frequency). To compensate for the spectral leakage, the mean value is thus subtracted from the real and imaginary parts at the measurement frequency. The main downside of this method is that in order to compensate the lowest frequency component, i.e. the one that is worst affected by the leakage, the observation window must include at least two cycles of this frequency. When dealing with frequencies in the order of millihertz (or lower), this greatly reduces the likelihood that the system will remain in steady state for the entire measurement duration, so whilst this method may lead to satisfactory results in offline EIS, it is unlikely to do so in online EIS.

One of the objectives of this project is to propose solutions to the existing bandwidth limitations in power converter-based EIS; to this end, the various compensation methods have been analysed to determine their effectiveness and a new method is proposed that overcomes some of the limitations of the existing methods. This is covered in Section 5.3.6.

5.3.4 Fourier transform

The next step in the signal processing chain is to transform the cell's current and voltage into the frequency domain. Whilst it is possible to determine the amplitude




Figure 5.5: Multisine measurement perturbation in time and frequency domains.

and phase of the signals in the time domain, doing so in the frequency domain has several advantages. For example, any noise superimposed on the signal is much easier to handle: as long as its frequency content does not overlap with any of the frequencies of interest, it can simply be ignored. In the frequency domain it is also much easier to distinguish between different components of the measurement perturbation. This is particularly important if measurement perturbations consisting of many different frequency components are used. For example, as shown in Figure 5.5, if a multisine perturbation with five components is used, determining the amplitude and phase of the various components in the time domain is not a simple task, whereas in the frequency domain all components are easily distinguished.

Once voltage and current signals have been adjusted to the correct length and the voltage has undergone drift compensation, the Fourier transform is used to obtain the signals' spectra.

5.3.5 Signal quality verification

The final step before the cell impedance can be determined is to verify the quality of the voltage and current signals⁴. This is a crucial step — particularly in order to extend the bandwidth of power converter-based EIS measurements to very low frequencies — because the concept of *impedance* is only valid for linear elements in ac steady state. As explained in Chapter 4, electrochemical cells generally have non-

⁴Parts of this section have been adapted from the author's previous work in [59].

linear I-V curves, but by using small-amplitude perturbations it is still possible to describe the system in terms of impedance. To ensure ac steady state, it is necessary to perturb the battery current for more than one cycle. The measurement data from the first cycle can be discarded because the system will not be in ac steady state.

If subjected to a sinusoidal perturbation, the response of a linear element in ac steady state will be a sinusoid at the same frequency as the perturbation. If the steady-state and linearity conditions are *not* fulfilled, on the other hand, the response will typically also include harmonics of the perturbation frequency. In practice, assessing the signal quality therefore involves analysing the measurement perturbation components and their harmonics.

It is important to note that the drift compensation step does not circumvent the need for signal quality verification. Both the cell's EMF and its impedance are functions of SoC; a change in SoC during the measurement could therefore affect both of them. The drift compensation discussed in Section 5.3.3 deals only with changes in the EMF; in order to ensure the validity of the impedance measurement, it is also crucial to verify that any change in SoC during the measurement's acquisition do not significantly alter the impedance. Indeed, as well as validating the assumptions of linearity and steady state, the signal quality verification step can also be used to confirm the correct operation of the voltage drift compensation, since — if a drift is removed — the harmonics of the fundamental component in the compensated signal will always be of lower magnitude than those of the uncompensated waveform.

As the perturbation frequency is reduced, not only does the SoC variation due to the current's dc component increase, but so does the SoC variation due to its ac component. The SoC variation due to the current's ac component will be sinusoidal and thus cause an approximately sinusoidal oscillation in the cell EMF. Because we can only measure the voltage at the terminals of the cell, rather than the voltage across the internal impedance, the oscillation in the EMF will be indistinguishable from the perturbation across the impedance and could therefore introduce a substantial error in the impedance estimation. For this reason, the signal quality check should also involve a comparison of the magnitude of the EMF oscillation on one hand (this requires a characterisation of the cell's EMF vs SoC relationship), and the magnitude of the cell's terminal voltage on the other hand. Measurements in which the two are comparable in size should be discarded, as any impedance estimate based on such data would contain considerable errors.

5.3.6 Analysis of drift compensation methods

5.3.6.1 Simple straight line approximation

In this compensation method — originally proposed in [18] — the assumption is made that there is a linear drift (or *ramp*) present in the acquired voltage signal. The slope of the straight line is found by subtracting the first sample from the last sample and dividing the result by the length of the signal. The voltage is then compensated by subtracting the identified trendline from it.

The input voltage signal to the compensation algorithm can be denoted as $v(k)$, where k denotes the zero-indexed sample number, i.e. assuming a signal length of N samples, k ranges from 0 to $N - 1$. The slope identified by the compensation algorithm is then given by:

$$m_i = \frac{\Delta v}{\Delta k} = \frac{v(N - 1) - v(0)}{N - 1} \quad (5.3)$$

The function describing the identified ramp is thus⁵:

$$r_i(k) = m_i \cdot k \quad (5.4)$$

and the output of the compensation algorithm is given by:

$$v_c(k) = v(k) - r_i(k) \quad (5.5)$$

If the measurement perturbation consists of a single sinusoid, and the observation window length is set equal to its period (as explained in Section 5.3.2), then the acquired voltage signal that is fed into the compensation algorithm can be described as follows:

$$v(k) = A \sin\left(\frac{2\pi}{N}k + \theta\right) \quad (5.6)$$

From (5.3) we note that the change in amplitude between first and last sample is given by:

$$\Delta v = v(N - 1) - v(0) \quad (5.7)$$

Therefore, for the sinusoid of (5.6):

$$\Delta v = A \sin\left[\frac{2\pi}{N}(N - 1) + \theta\right] - A \sin(\theta) = A \sin\left[2\pi - \frac{2\pi}{N} + \theta\right] - A \sin(\theta) \quad (5.8)$$

⁵The vertical axis intercept is omitted from the equation of the ramp as it affects only the dc component and has no influence on the effectiveness of the compensation.

Since $\sin(x)$ repeats every 2π , i.e. $\sin(x + 2\pi) = \sin(x)$, the 2π term can be dropped from the argument of the first sin term:

$$\Delta v = A \left[\sin \left(\theta - \frac{2\pi}{N} \right) - \sin(\theta) \right] \quad (5.9)$$

The sum-to-product trigonometric identities can be used to rewrite (5.9), giving Δv in alternate form:

$$\Delta v = -2A \sin \left(\frac{\pi}{N} \right) \cdot \cos \left(\theta - \frac{\pi}{N} \right) \quad (5.10)$$

Substituting (5.10) back into (5.3) gives the slope that the compensation algorithm identifies when given a sinusoid as input:

$$m_i = \frac{\Delta v}{\Delta k} = \frac{-2A \sin \left(\frac{\pi}{N} \right) \cdot \cos \left(\theta - \frac{\pi}{N} \right)}{N - 1} \quad (5.11)$$

An expression for the compensated signal is found by first substituting the identified slope from above into (5.4) and then substituting the resulting $r_i(k)$ into (5.5):

$$v_c(k) = v(k) - r_i(k) = A \sin \left(\frac{2\pi}{N}k + \theta \right) - \frac{-2A \sin \left(\frac{\pi}{N} \right) \cdot \cos \left(\theta - \frac{\pi}{N} \right)}{N - 1} \cdot k \quad (5.12)$$

The compensation algorithm thus introduces the very thing it was meant to remove from the signal — a linear drift.

Figure 5.4 showed the qualitative effect that the presence of a linear drift in a signal has on the signal's magnitude spectrum. At this point it is important to consider the frequency domain representation of a ramp, in order to better understand how the ramp that is added by the compensation algorithm affects the frequency spectrum of the signal being compensated.

If a ramp were to be repeated, i.e. if it were to be turned into a periodic signal, it would be equivalent to a sawtooth waveform. To understand the composition of a ramp in terms of its constituent sinusoids, it is therefore possible to use the Fourier series of a sawtooth signal; this is a well-known result — it is the infinite sum of sine terms including both even and odd harmonics:

$$s(t) = \frac{\Delta a}{\pi} \sum_{n=1}^{\infty} \frac{\sin(2\pi n f_1 t)}{n} \quad (5.13)$$

where n is the harmonic number ($n = 1$ corresponding to the fundamental component), f_1 is the frequency of the fundamental component, and Δa is the difference between the maximum and minimum values of the sawtooth. Note that in the ramp identified by the compensation algorithm $\Delta a = |\Delta v|$. The absolute value is required because Δv can be positive or negative, whereas Δa is always positive.

The expression in (5.13) describes a sawtooth waveform with a negative slope. For a sawtooth with a positive slope, the sine term in (5.13) would be negated⁶.

A ramp that spans the length of the observation window of the FFT therefore leads to a fundamental component with magnitude $\Delta a/\pi$ at the resolution frequency of the FFT, i.e. $1/t_{\text{obs}}$, as well as harmonic components in all FFT bins that decay in magnitude with a slope of 20 dB per decade. The phase of all components is either $+90^\circ$, if the ramp has a positive slope, or -90° , if the ramp has a negative slope. Since the observation window covers one cycle of the measurement frequency, the fundamental component of the ramp and the input sinusoid both appear in the second component of the FFT (the first being the dc component), i.e. they overlap. If the sinusoid's amplitude and phase are to be estimated from the FFT, the presence of the ramp thus introduces errors in these estimates.

Since the Fourier transform is a linear operation, it is possible to write the Fourier transform of the compensated signal $v_c(k)$ as:

$$F\{v_c(k)\} = F\{v(k) - r_i(k)\} = F\{v(k)\} - F\{r_i(k)\} \quad (5.14)$$

$$V_c(f) = V(f) - R_i(f) \quad (5.15)$$

where $F\{v(k)\}$ describes the forward Fourier transform operation, i.e. $V(f)$ is the Fourier transform of $v(k)$, and f is used to denote discrete frequency.

The magnitude of the ramp's harmonics decreases as $1/n$; the error in the amplitude estimate is therefore worst at low frequencies, so to begin with consider only the fundamental component, f_1 :

$$R_i(f_1) = \frac{\Delta a}{\pi} \angle \pm \frac{\pi}{2} = \pm j \frac{\Delta a}{\pi} \quad (5.16)$$

$$V(f_1) = A \angle \left(\theta - \frac{\pi}{2}\right) = A \left[\cos\left(\theta - \frac{\pi}{2}\right) + j \sin\left(\theta - \frac{\pi}{2}\right) \right] \quad (5.17)$$

The fundamental component of the compensated signal is given by:

$$V_c(f_1) = V(f_1) - R_i(f_1) \quad (5.18)$$

$$= \text{Re}\{V(f_1)\} - \text{Re}\{R_i(f_1)\} + j [\text{Im}\{V(f_1)\} - \text{Im}\{R_i(f_1)\}] \quad (5.19)$$

$$= A \cos\left(\theta - \frac{\pi}{2}\right) + j \left[A \sin\left(\theta - \frac{\pi}{2}\right) \mp \frac{\Delta a}{\pi} \right] \quad (5.20)$$

⁶Or equivalently 180° could be added/subtracted to/from the argument of sine.

If we let $\theta = 0$, the above expression simplifies to:

$$V_c(f_1) = -j \left[A - \frac{\Delta a}{\pi} \right] \quad (5.21)$$

$$= \left(A - \frac{\Delta a}{\pi} \right) \angle -\frac{\pi}{2} \quad (5.22)$$

Similarly, if we let $\theta = \pi$, the fundamental of the compensated signal becomes:

$$V_c(f_1) = \left(A - \frac{\Delta a}{\pi} \right) \angle +\frac{\pi}{2} \quad (5.23)$$

In these two special cases, where the angle of the sinusoid is either 0° or 180° , the *absolute* error that the compensation algorithm introduces in the amplitude estimate of the input sinusoid is therefore:

$$e_a = \frac{\Delta a}{\pi} = \frac{|\Delta v|}{\pi} = \frac{2A}{\pi} \sin\left(\frac{\pi}{N}\right) \cdot \cos\left(\frac{\pi}{N}\right) \quad (5.24)$$

The *relative* error in the amplitude estimate is thus independent of the sinusoid amplitude:

$$e_r = \frac{2}{\pi} \sin\left(\frac{\pi}{N}\right) \cdot \cos\left(\frac{\pi}{N}\right) \quad (5.25)$$

As (5.25) shows, the relative error is solely a function of N , i.e. the number of samples in one cycle of the input sinusoid, which depends on the chosen sampling rate, f_s , and the frequency of the sinusoid, f :

$$N = \frac{f_s}{f} \quad (5.26)$$

It can be shown that for $N > 10$ the relationship between the relative error and N can be approximated as one of inverse proportionality:

$$e_r \approx \frac{2}{N} = 2 \frac{f}{f_s} \quad \{N > 10\} \quad (5.27)$$

Thus, for a given sampling rate, the relative error in the amplitude estimate is proportional to the input sinusoid's frequency. For example, if the input frequency was 1 Hz, then a sampling rate of at least 200 Hz would ensure that the relative error was less than 1 %. If the input frequency was increased, the size of the error would also increase, but, as explained in Section 5.3.3, drift compensation is not normally required at higher frequencies.

From the above it can be concluded that — for a single sine measurement perturbation — the error introduced by the simple linear compensation algorithm is negligible, as long as $f_s/f > 200$.

Note that in the two special cases considered above, i.e. where the angle of the sinusoid is either 0° or 180° , the signal is in phase with the ramp's fundamental component. This leads to the maximum error in the amplitude estimate and no error in the estimate of the phase. When the input signal takes on other phase angles, (5.20) shows that the compensated signal will differ from the input signal in both phase and amplitude.

Next, a multisine measurement perturbation is considered. A multisine signal can be described as the sum of its (M) components:

$$v(k) = \sum_{m=1}^M s_m(k) \quad (5.28)$$

where

$$s_m(k) = A_m \sin\left(\frac{2\pi}{N_m}k + \theta_m\right) \quad (5.29)$$

The voltage signal fed to the compensation algorithm can thus be stated as:

$$v(k) = \sum_{m=1}^M A_m \sin\left(\frac{2\pi}{N_m}k + \theta_m\right) \quad (5.30)$$

Note that the observation window is still one period of the lowest frequency component, i.e. $N = N_1$, and therefore the range of k is the same as before, i.e. 0 to $N - 1$. N_m gives the number of samples in one cycle of frequency component m , i.e. $N_m = f_s/f_m$.

The difference in amplitude between the first and last samples of the multisine signal is given by:

$$\Delta v = \sum_{m=1}^M s_m(N-1) - \sum_{m=1}^M s_m(0) \quad (5.31)$$

$$= \sum_{m=1}^M \Delta s_m \quad (5.32)$$

where each component's Δs_m can be found by using (5.10):

$$\Delta s_m = -2A_m \sin\left(\frac{\pi}{N_m}\right) \cdot \cos\left(\theta_m - \frac{\pi}{N_m}\right) \quad (5.33)$$

As mentioned before, N_m is very large at the low end of the measurement frequency range and therefore Δs_m and the related error are very small; however, as the frequency of subsequent components in the multisine is raised, the number of samples per cycle decreases, thus increasing each component's respective Δs_m , which in turn increases the Δv of the multisine as a whole.

Consider an example in which the multisine measurement perturbation consists of five components with frequencies 0.01 Hz, 0.1 Hz, 1 Hz, 10 Hz, and 100 Hz. The amplitude and phase of the components in the voltage response of course depends on the impedance of the battery at each frequency, but for the simplicity of this example, consider that the amplitude is equal to 10 mV and the phase angle is equal to zero for all components. Finally, assume a sampling rate of 1 kHz.

We can calculate Δs_m for each component by using (5.33):

$$\Delta s_1 \approx -6.3 \cdot 10^{-7} \text{ V} \quad (5.34)$$

$$\Delta s_2 \approx -6.3 \cdot 10^{-6} \text{ V} \quad (5.35)$$

$$\Delta s_3 \approx -6.3 \cdot 10^{-5} \text{ V} \quad (5.36)$$

$$\Delta s_4 \approx -6.3 \cdot 10^{-4} \text{ V} \quad (5.37)$$

$$\Delta s_5 \approx -5.9 \cdot 10^{-3} \text{ V} \quad (5.38)$$

and then sum them to obtain Δv for the multisine:

$$\Delta v = \sum_{m=1}^5 \Delta s_m \approx -6.6 \text{ mV} \quad (5.39)$$

If this multisine signal is fed into the compensation algorithm, a ramp will be identified whose fundamental frequency component is:

$$R_i(0.01 \text{ Hz}) = \frac{|\Delta v|}{\pi} \angle -\frac{\pi}{2} = 2.2 \text{ mV} \angle -\frac{\pi}{2} \quad (5.40)$$

Figure 5.6 shows the magnitude spectrum of the multisine (V), the identified ramp (R_i), and the output of the compensation algorithm (V_c). Subtracting the ramp from the multisine leads to a relative error in the amplitude estimate of the lowest frequency component of the multisine of over 20 %. In this example, the second frequency component of the multisine is one decade over the first; at this frequency, the magnitude of the ramp has decayed by 20 dB relative to its fundamental, but it still causes a relative error in the amplitude estimation of the second multisine component of over 2 %.

The errors introduced by this compensation method are not random, so there may be solutions to eliminate, or at least reduce them. Indeed, if the error ramp introduced by the compensation algorithm was known or could be estimated, it would be possible to reduce the error by adding the error ramp back into the compensated signal, v_c . Consider again the expression for the output of the compensator:

$$v_c(k) = v(k) - r_i(k) \quad (5.41)$$




Figure 5.6: Magnitude spectra of multisine signal — before (V) and after (V_c) drift compensation using the straight line approximation method. R_i is the identified ramp.

If the input signal, v , consisted of some signal (e.g. sine or multisine), v_s , and a ramp, r_s , i.e.

$$v(k) = v_s(k) + r_s(k) \quad (5.42)$$

then the ramp identified by the algorithm would be the sum of the ramp present in the signal, r_s , and the error ramp introduced in the compensation, r_c , i.e.

$$r_i(k) = r_s(k) + r_c(k) \quad (5.43)$$

and (5.41) could be rewritten:

$$v_c(k) = v_s(k) + r_s(k) - [r_s(k) + r_c(k)] \quad (5.44)$$

$$= v_s(k) - r_c(k) \quad (5.45)$$

Thus, if r_c was known, the error in v_c could be corrected:

$$v_s(k) = v_c(k) + r_c(k) \quad (5.46)$$

The error ramp is fully characterised by Δv , which was given in (5.10) for the case of a single sine, and (5.32) for a multisine input. Correcting the error would therefore require estimates of both the amplitude and phase angle of the sinusoids⁷.

⁷The number of samples in one period, N , is a known quantity, since we define the frequency components of the measurement perturbation and the sampling frequency is also known.

From the multisine example it could be observed that the largest contribution to the overall Δv was from the highest frequency component. This is expected since N decreases as the frequency of the signal increases. In cases in which the highest frequency component of the multisine is much higher than the rest of the components, it may therefore be possible to reduce the error by basing the estimate of the multisine Δv on an estimate of the Δv associated with the highest frequency component. The problem of estimating all components' amplitudes and phases, would then be reduced to estimating only the highest frequency component's amplitude and phase. It may be possible to reduce this further to only requiring an estimate of the amplitude, since the impedance of batteries often becomes purely real at high frequencies and the phase estimate could therefore be obtained from the FFT of the current. Aside from not fully correcting the error, this method places undue restrictions on the definition of the measurement perturbation and is thus not considered practical.

Since higher frequency components are responsible for the largest proportion of the error, another option could be to filter the multisine signal before feeding it into the compensation algorithm. In practice, however, this method did not considerably reduce the value of Δv .

Provided more than one cycle of the measurement perturbation is available, the simplest solution by far would involve obtaining an additional sample, i.e. the first sample of the second cycle, which (at least in theory) should be identical to the first sample of the first cycle, and using it in the estimation of the slope, thus all but eliminating Δv .

So far, only "perfect" signals have been considered; in reality, signals are never noise-free. When a signal contains noise, the amplitude difference between the first and last samples can quickly exceed the value of Δv that is solely due to the discretisation of the signal, i.e. that found through (5.10). For example, given a single sine input with amplitude 10 mV and zero phase angle, if $N = 10^5$, then $\Delta v \approx 0.63 \mu\text{V}$, which will lead to a relative error in the amplitude estimate of 0.002 %. Adding a zero-mean white noise signal with a standard deviation of 3 % of the sinusoid amplitude, i.e. 0.3 mV, is enough to increase the relative error in the amplitude estimate to over 1 % in approximately one out of three cases. Particularly problematic is the fact that it is enough for the noise (or perhaps a glitch in the data acquisition) to affect a single sample, i.e. either the first or the last sample in the observation window, in order for the compensation effectiveness to be reduced.

To summarise, this section has shown that when the voltage signal consists of a multisine, the drift compensation algorithm based on the simple straight line approximation introduces a new ramp in the signal that can lead to large errors in

both phase and amplitude estimates of the multisine components. If there is a very large linear drift present in the input, then using this compensation method can lead to more accurate results than not using it, but the residual errors are unacceptably high. It can therefore be concluded that the drift compensation method based on the simple straight line approximation is not well-suited to remove linear drifts from multisine signals. More generally, it has been shown that the method's main advantage — simplicity — causes it to lack robustness when real (noisy) input signals are considered.

5.3.6.2 Straight line linear least squares fit

This method — proposed in [12] — is based around the concept of fitting the time domain voltage signal using a linear least squares (LLS) algorithm. The algorithm attempts to find the first order polynomial (i.e. straight line) that minimises the sum of the squared residuals, where the residual is the difference between the measured voltage and the voltage estimated by the fitted line.

Consider Figure 5.7 in which a signal consisting of the sum of a sinusoid, a dc offset, and a linear drift is fitted using the LLS algorithm. Since the arithmetic mean over one cycle of a sinusoid is zero and the input contains a linear component, the expectation may be that the fitted line would be equal to the linear component present in the input signal; however, as Figure 5.7 shows, this is not the case.

Compensating the input signal using the fitted line results in the signal shown in Figure 5.8. Whilst this bears more of a resemblance to a sine wave than the uncompensated waveform, it is clear that the signal still contains a linear component.




Figure 5.7: Linear least squares fit of sinusoid with drift.




Figure 5.8: Sinusoid with drift after compensation through the linear least squares fit method.




Figure 5.9: Magnitude spectra of sinusoid with drift, V , sinusoid without drift, S , and the signal compensated using the linear least squares method, V_{LLS} .

The presence of a residual drift in the compensated signal is also apparent in its magnitude spectrum, shown in Figure 5.9 alongside the spectra of the uncompensated signal, V , and the sinusoidal component, S . Figure 5.9 highlights that whilst the straight line LLS compensation can considerably reduce the amplitude estimation error, the residual error remains high.

Similarly to the simple straight line approximation method described in the previous section, the compensation method based on the straight line LLS fit is

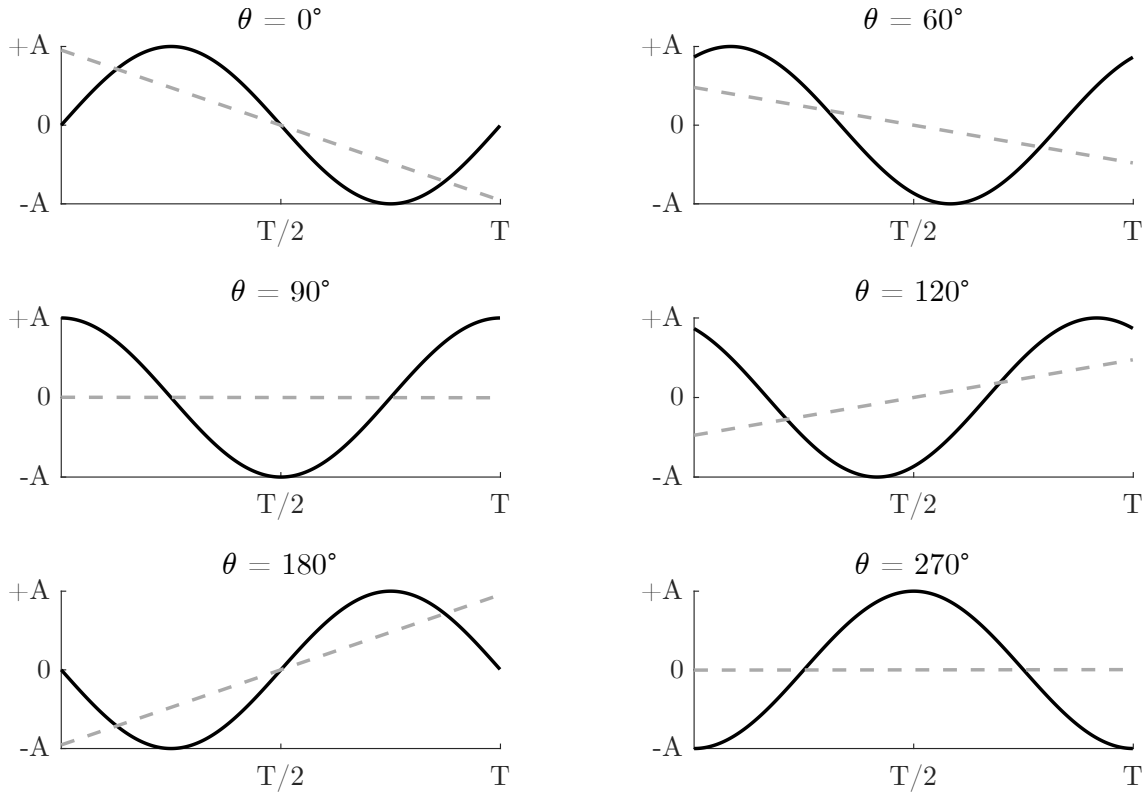


Figure 5.10: Linear least squares line of best fit for sinusoids of varying phase angles.

capable of removing a linear drift from the signal, but in doing so, adds a new ramp to the signal. The introduction of the new ramp is due to what the LLS algorithm considers to be the line of best fit for a sinusoid. As Figure 5.10 shows, when the LLS algorithm is fed a sinusoid without drift, the fitted line has a considerable slope. In compensating the signal, a ramp with this slope is subtracted from the input sinusoid, thus introducing the new drift component. Using the straight line LLS method to compensate a sinusoidal signal that does not contain a drift (or one that is much smaller than that introduced by the compensation) therefore results in greater errors than not using any compensation.

Figure 5.10 also shows that the slope of the line of best fit depends on the phase angle of the sinusoid: its absolute value is maximum at 0° and 180° and minimum at 90° and 270° . In fact, a plot of the slope of the LLS line of best fit, m_{LLS} , against the phase angle of the input sinusoid, θ , reveals that m_{LLS} shows the same phase angle dependency as the slope identified by the simple linear approximation, m_i . This is depicted in Figure 5.11. The two slopes are related by the following expression:

$$m_{\text{LLS}} = m_i \cdot \frac{N-1}{\pi} \cdot P = \frac{\Delta v}{N-1} \cdot \frac{N-1}{\pi} \cdot P = \frac{\Delta v}{\pi} \cdot P \quad (5.47)$$

where P has been determined empirically: $P \approx 0.95$ for $N \in [10^2, 10^6]$ and $P \approx 1$ when $N = 10$.

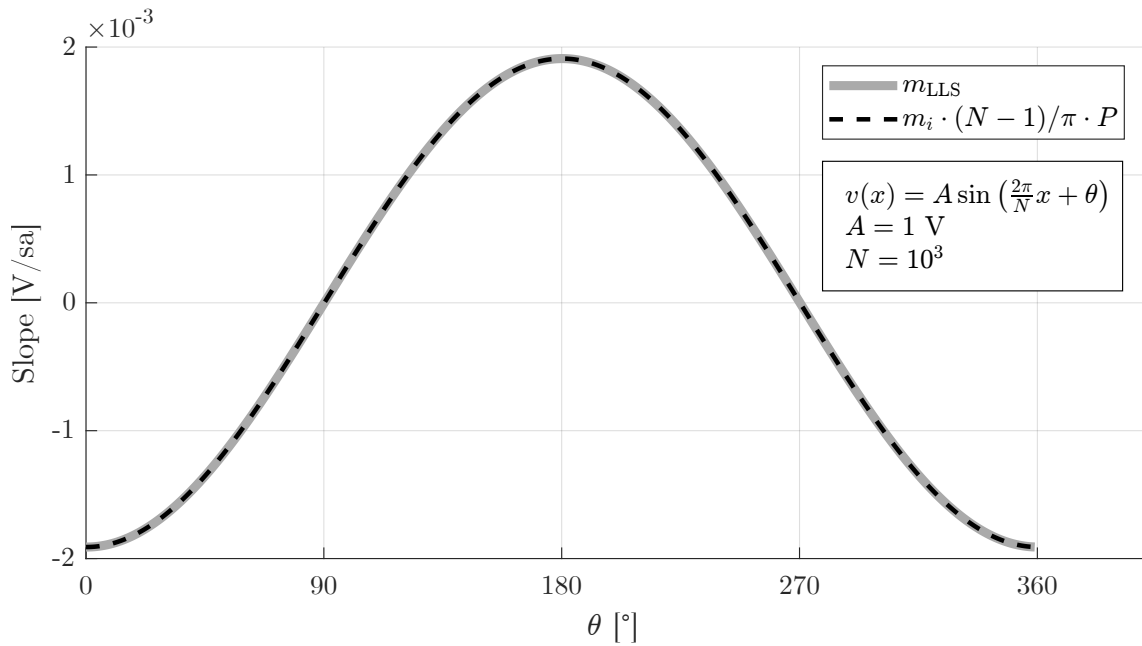


Figure 5.11: Relationship between slope of line of best fit and sinusoid phase angle.

The straight line LLS fit for a sinusoid improves as the number of cycles over which the fit is calculated increases; this is shown in Figure 5.12.

When a multisine signal is fitted using the LLS algorithm, the results are very similar to those obtained when only the lowest frequency component is fitted. The output of the fitting algorithm when given a multisine as input is found to be equivalent to the sum of its outputs when given the individual sine components as inputs (i.e. superposition applies). Hence, because the components of the multisine are of higher frequency than the fundamental, their respective slopes — and the errors they contribute — will get progressively smaller as their frequency increases.

Interestingly, whereas the slope returned by the simple linear approximation of Section 5.3.6.1 showed the greatest dependency on the *highest* frequency component of the multisine, the slope of the LLS fit depends mainly on the *lowest* frequency component.

Figure 5.13 shows the magnitude spectra of the same multisine signal used in Section 5.3.6.1, both before and after compensation with the straight line LLS method (V and V_{LLS} , respectively), as well as the spectrum of the line of best fit, R_{LLS} . Note the similarity in the scale of the amplitude estimation error v compared to the single sine input in Figure 5.9.

One of the advantages of the LLS method compared to the simple linear approximation considered in the previous section is that it is much less sensitive to noise. A comparison of the compensation effectiveness of the various methods when the input signal is subjected to noise is given in Section 5.3.6.4.

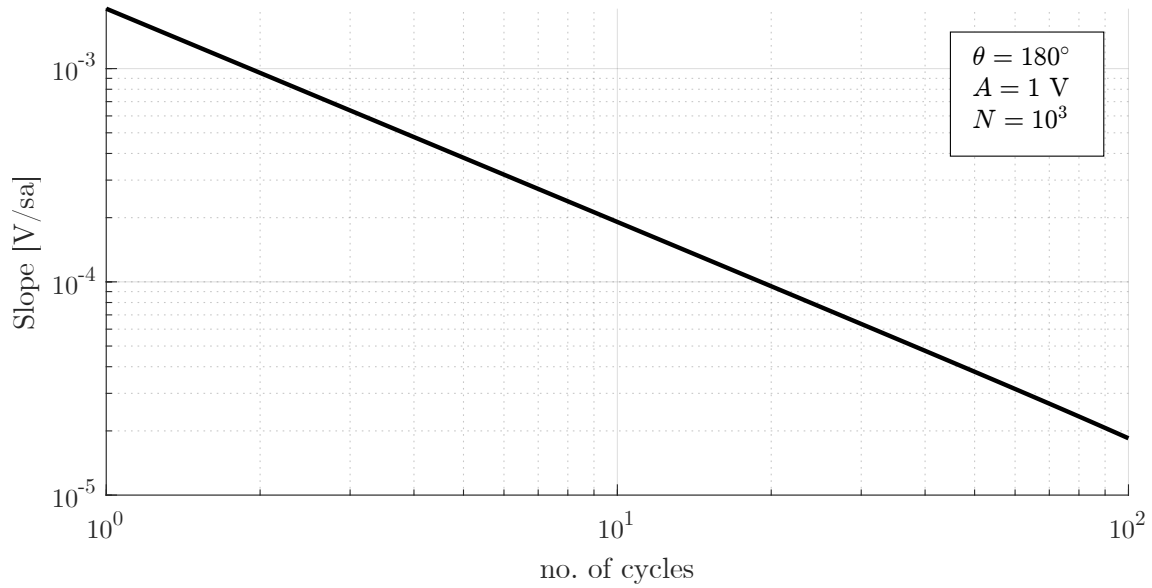


Figure 5.12: Relationship between slope of line of best fit and number of cycles of input sinusoid. θ , A , and N denote the phase angle, amplitude, and number of samples of the input sinusoid, respectively.




Figure 5.13: Magnitude spectra of multisine signal — before (V) and after (V_{LLS}) drift compensation using the linear least squares method. R_{LLS} is the identified ramp.

5.3.6.3 Ramp plus sine linear least squares fit

The ramp plus sine (RPS) linear least squares compensation method is an extension of the straight line linear least squares fit described in the previous section. It is based on the same principle of fitting the voltage signal in the time domain, but rather than using a straight line fit, it leverages the fact that the composition of the measurement perturbation is known beforehand, and thus the signal can be fitted with a function that more closely resembles the expected waveform. Specifically, the fitting function that is used consists of the sum of a ramp, r_{RPS} , and a sinusoid, s_{RPS} :

$$f(x) = r_{\text{RPS}}(x) + s_{\text{RPS}}(x) = mx + b + A \sin\left(\frac{2\pi}{N_1}x\right) \quad (5.48)$$

where N_1 is the number of samples in one cycle of the lowest frequency component in the measurement perturbation, i.e. $N_1 = f_s/f_1$, and m , b , and A are the parameters determined by the fitting algorithm⁸.

As shown in Figure 5.14 this method can handle regular (without drift) sinusoid inputs correctly, regardless of phase angle.

The relationship between the input sinusoid's phase angle and the fit accuracy is considered in more detail in Figure 5.15; this shows the Δv of the ramp identified by the RPS fit — normalised by the sinusoid's amplitude — plotted against the sinusoid's phase angle. It can be seen that the absolute value of Δv reaches a maximum of approximately 1.5 % of the amplitude, A , for phase angles of 90° and 270° . If a sinusoid without drift is compensated using the ramp identified by the RPS fit, a small error is thus introduced; however, as Figure 5.16 shows, even when $N_1 = 10^3$, the errors in amplitude and phase estimates are both small. As discussed before, N_1 will likely be much larger than 10^3 and as shown in Figure 5.17, Δv is inversely proportional to N_1 . The errors can thus be considered negligible.

The cause of the error in the estimate of the ramp's slope is the fact that (5.48) is missing a parameter to account for the sinusoid's initial phase angle. Including the phase angle in the fitting function all but eliminates the error; however, doing so also necessitates the use of a non-linear fitting algorithm, which requires much more processing power. It is also possible to use the angle sum and difference trigonometric identities to rewrite the sinusoidal term in the fit function (including the phase angle) as a linear combination of sine and cosine terms. This would allow the phase angle to be estimated using a linear fit model and would thus result in

⁸As explained before, the vertical axis intercept (denoted by b here) is not required for the compensation to be effective; it has been retained solely to better illustrate the compensation process.

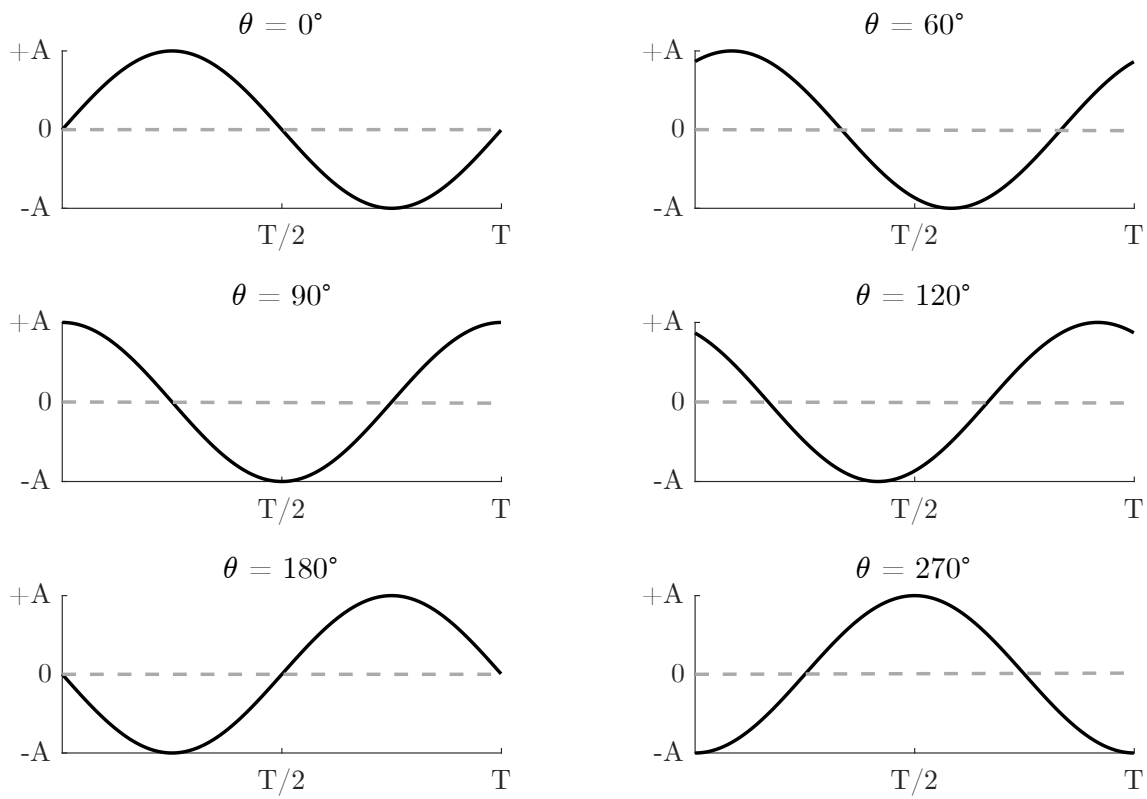


Figure 5.14: The straight line component of the ramp plus sine fit (dashed) for sinusoids of various phase angles (solid).






Figure 5.15: Δv of ramp plus sine fit versus phase angle of input sinusoid.



(a) Relative error in amplitude estimate.



(b) Absolute error in phase estimate.

Figure 5.16: Errors in amplitude and phase estimates after compensation of sinusoid without drift using the RPS method.




Figure 5.17: Relationship between the Δv of the ramp introduced by the RPS fit and the number of samples per cycle of the input sinusoid.

an improved fit without a substantial increase in processing complexity; however, as shown above, when the phase angle is omitted from the fit, the errors are very small anyway, so for simplicity this approach has been adopted.

By using the fitting function of (5.48), the RPS method is able to correctly characterise linear drifts superimposed on sinusoidal signals. Figure 5.18 shows an example of a sine wave with an added drift of -1 mV/s, as well as the ramp component of the RPS fit, r_{RPS} (cf. Figure 5.7).




Figure 5.18: Ramp plus sine fit of sinusoid with drift.

When a multisine measurement perturbation is used, (5.48) can no longer provide an accurate fit. The reason for this is that the fitting function uses N_1 as a constant parameter and this is only accurate for the multisine's fundamental component. Figure 5.19 shows the relationship between the Δv resulting from the slope estimated by the RPS fit when there is a mismatch between the value of N_1 used in the fit and the number of samples per cycle of the input sinusoid, N_i . In Figure 5.19 N_1 was kept constant whilst N_i was decreased. It can be seen that Δv is inversely proportional to N_1/N_i when $\theta = 180^\circ$. This relationship varies with the phase angle of the sinusoid. Interestingly, when the phase angle is 90° or 270° , Δv is independent of N_1/N_i , so whereas these phase angles lead to the maximum Δv when the frequency of the input sinusoid matches the frequency of the fit, when the frequencies do *not* match, they lead to the minimum Δv (for $N_1/N_i < 10^4$, approximately).

From Figure 5.19 it is clear that if the RPS method is used to compensate a multisine signal, those components that are closest in frequency to the fundamental will cause the largest proportion of the error (unless they happen to have phase angles of 90° or 270° , as described above). To illustrate this, an example is given in Figure 5.20; here the same multisine signal used in the previous two sections is compensated using the RPS method. The figure shows three magnitude spectra: the uncompensated multisine signal (V), the compensated multisine when all its components have phase angles of 0° (V_{RPS}), and the compensated multisine when all components — except 0.1 Hz — have 0° phase angles (V_{RPS2}). In V_{RPS2} the second component, 0.1 Hz, which is closest in frequency to the fundamental, has a




Figure 5.19: Δv resulting from the slope estimated by the RPS fit when there is a mismatch between fit frequency (f_1) and input frequency (f_i).




Figure 5.20: Magnitude spectra of multisine signal — before (V) and after (V_{RPS} and V_{RPS2}) drift compensation using the ramp plus sine method. In V_{RPS} all of the multisine’s components have 0° phase angles; in V_{RPS2} the component at 0.1 Hz has a phase angle of 90° , whilst all other components have 0° phase angles.

phase angle of 90° .

Figure 5.20 also shows that whilst the RPS method gives much better results than the straight line LLS fit method (cf. Figure 5.13), the relative error in the amplitude estimate can still be nearly as high as 20 %.

One option to reduce the errors observed when compensating multisine signals is to improve the quality of the fit. This could be achieved, for example, by adding more sinusoidal components to the fitting function. However, this would also increase the complexity of the fit, thus increasing the required processing power. Also, this approach could become unsustainable as more components are added to the measurement perturbation; although in practice it may be enough to fit only those components closest in frequency to the fundamental of the multisine.

Another option is to filter the multisine signal before feeding it to the RPS fit. From the preceding analysis it is clear that to maximise the effectiveness of this method, those components of the multisine that are closest to the fundamental should be filtered out, as this will result in the greatest error reduction. Since the composition of the measurement perturbation is generally known beforehand, it is possible to use a simple moving mean filter with a window size equal to N_2 , i.e. the number of samples in one cycle of the second multisine component. This will filter out all of the multisine’s components, except for the fundamental, thus producing a signal that can be accurately fitted using the RPS method.

Consider again the multisine from the previous example except with all components' phase angles set to 45° , an added linear drift of -1 mV/s and a dc offset voltage of 3.3 V. Figure 5.21 shows the multisine signal before and after being processed by the moving mean filter (v and v_f , respectively), as well as the ramp fitted by the RPS method in both cases, i.e. r_f and r_{vf} , respectively.

Whilst in the time domain the difference between the two ramps does not appear too significant, the frequency domain view in Figure 5.22 shows that considerable errors in both the amplitude and phase estimates are avoided by filtering the multisine before fitting it using the RPS method.

The main drawback of the RPS method compared to the other methods considered here (the simple linear approximation in particular) is that it involves more steps and is computationally more complex.

A common drawback of all of the methods considered is that they assume that the drift in the voltage is linear. When this is not the case, the compensation will fail to properly remove the drift and by incorrectly fitting a ramp to the signal the compensated signal will contain a linear component that was not present in the original signal. Even if the drift is not perfectly linear, however, considering the scale of the error that occurs when no compensation is used, compensating the signal is still likely to lead to more accurate results than not compensating it. To prevent the scenario in which the act of compensating leads to less accurate results, it is possible for the signal processing stage to include a check in which, for example,




Figure 5.21: Time-domain plot of multisine signal before and after being processed by the moving mean filter (v and v_f , respectively), as well as the ramp fitted by the RPS method in both cases, i.e. r_f and r_{vf} , respectively.

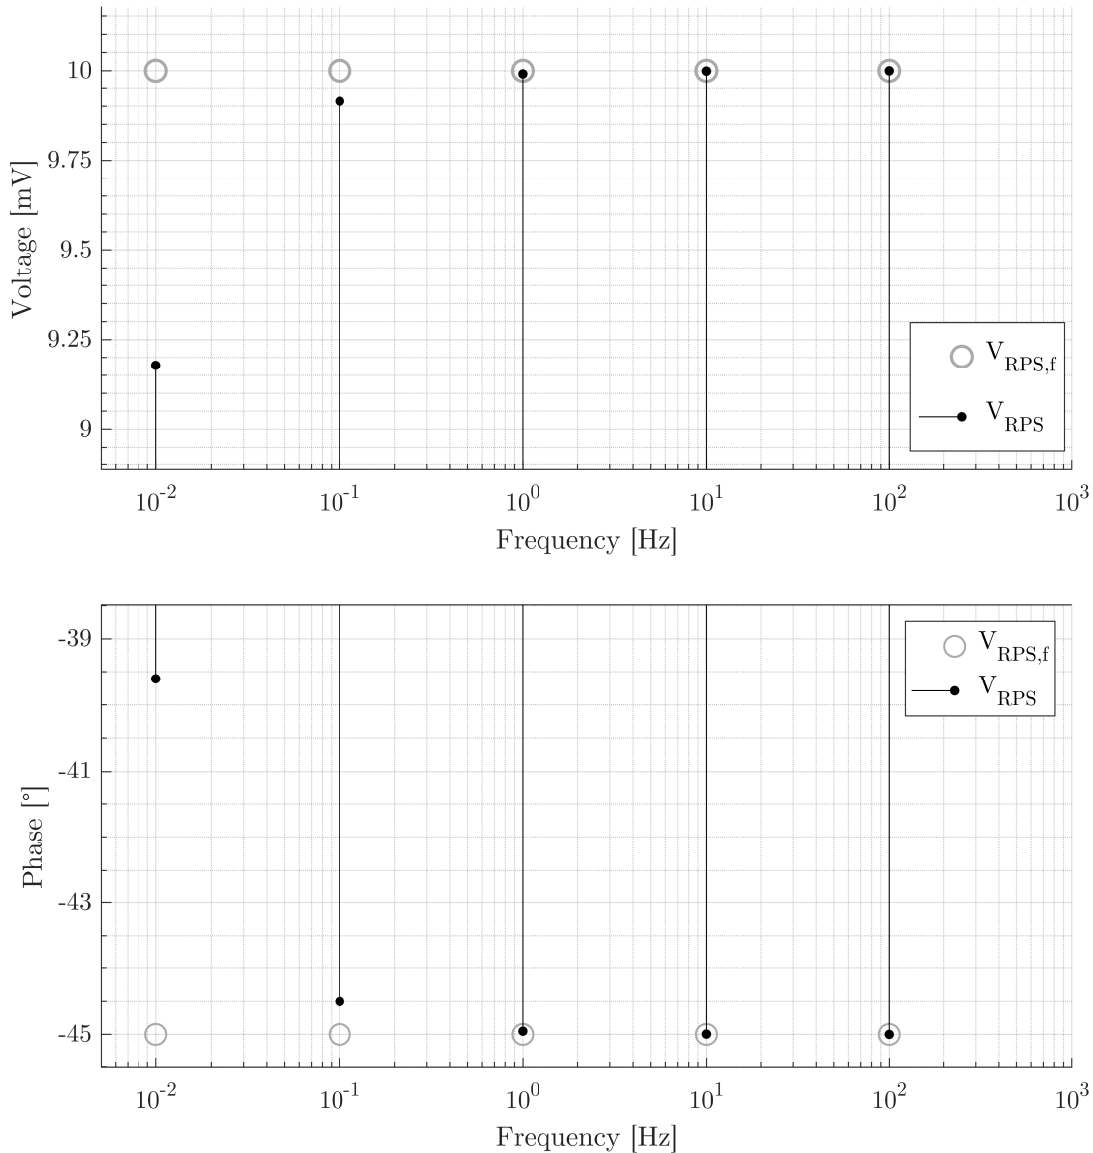


Figure 5.22: Frequency spectra of multisine signal after drift compensation using the RPS method — with ($V_{RPS,f}$) and without (V_{RPS}) filtering the signal before fitting.

the total harmonic distortion (THD) of the multisine components before and after compensation could be compared.

In summary, this section has shown that the drift compensation method based on the RPS fit is capable of correctly identifying linear drifts in both sine and multisine signals. It has also been shown that in order to minimise errors multisine signals need to be filtered before being fitted.

5.3.6.4 Comparison of compensation methods

The purpose of this section is to compare the effectiveness of the three compensation methods discussed and to conclude the analysis of the drift compensation methods.

For the first example, consider again the example given in Section 5.3.3 and illustrated in Figure 5.4 on p. 69 in which a multisine with five components (the lowest of which is at 10 mHz) has a superimposed linear drift of -1 mV/s. All components have identical amplitudes and phase angles of 10 mV and 0° . Figure 5.23 shows the signal both before being compensated (V), as well as after being compensated using the simple linear approximation (V_s), the straight line LLS method (V_{LLS}), and the RPS method (V_{RPS}). Note that in the time domain the compensated signals are difficult to distinguish from each other; for clarity only one of them is shown.

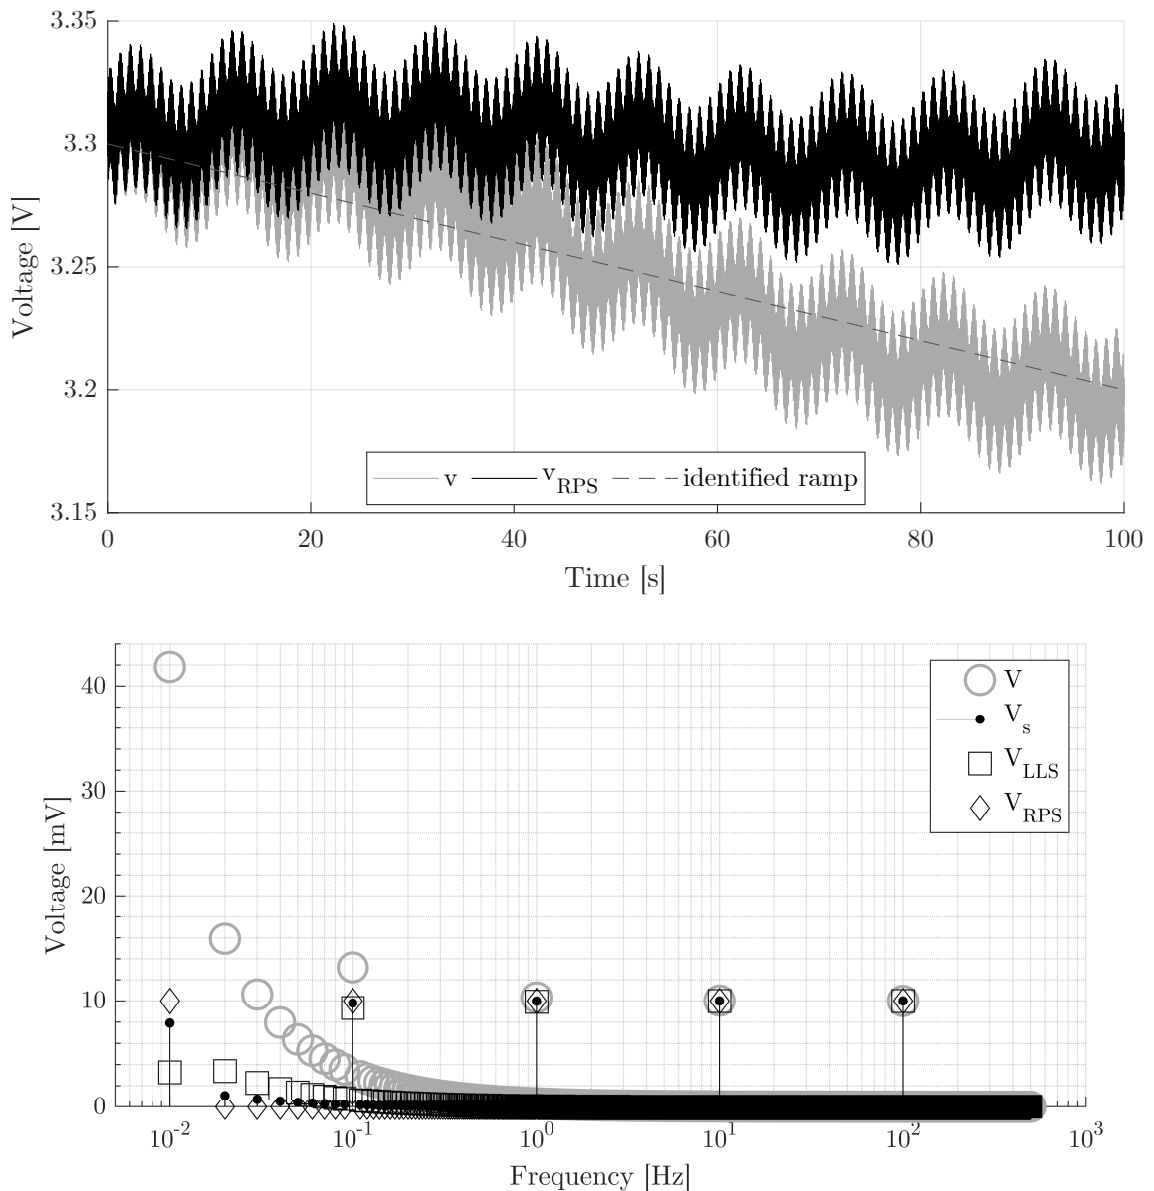


Figure 5.23: Compensation of linear drift in multisine signal (Example 1): the signal is shown before being compensated (V), as well as after being compensated using the simple linear approximation (V_s), the straight line LLS method (V_{LLS}), and the RPS method (V_{RPS}).

The magnitude spectrum in Figure 5.23 shows that the uncompensated signal, V , contains a large amount of spectral leakage due to the superimposed ramp. As a consequence, the amplitude of the fundamental component contains a relative error of over 300 %. The error decreases at higher frequencies, but even the third component at 1 Hz still has an error of over 3 %. The output of the simple linear drift compensation method, V_s , is much improved relative to the uncompensated signal, but significant errors remain. Similarly, the signal compensated using the straight line LLS method, V_{LLS} , is improved relative to the uncompensated signal, but still contains large errors. Finally, it can be seen that the output of the RPS compensator, V_{RPS} , no longer contains any meaningful traces of the drift, and residual errors are negligible. Table 5.1 summarises the results.

Table 5.1: Results of compensation example 1: magnitude and relative error in magnitude, $e_{r,m}$. Results given to two significant figures. Any $e_{r,m} < 0.01$ % has been rounded down to zero.

f [Hz]	0.01	0.1	1	10	100
No compensation					
$ V(f) $ [mV]	42	13	10	10	10
$e_{r,m}$ [%]	320	32	3.2	0.32	0.031
Simple linear approximation					
$ V_s(f) $ [mV]	7.9	9.8	10	10	10
$e_{r,m}$ [%]	21	2.1	0.21	0.021	0
Straight line LLS					
$ V_{LLS}(f) $ [mV]	3.2	9.3	9.9	10	10
$e_{r,m}$ [%]	68	6.8	0.68	0.068	0
Ramp plus sine compensation					
$ V_{RPS}(f) $ [mV]	10	10	10	10	10
$e_{r,m}$ [%]	0	0	0	0	0

In the previous example there was no error in the phase estimation; as previously explained, this is the case when the multisine components are in phase with the components of the ramp. Consider next an example in which the phase angles of the multisine components are all 45° , i.e. they are *not* the same as the phase angles of the ramp components. All other characteristics of the multisine and the superimposed ramp are the same as in the previous example. Magnitude and phase spectra are

shown in Figure 5.24 and the associated results are listed in Table 5.2. Note that in order to improve the clarity of the phase plot any components with a magnitude below $1e-12$ V have been set to zero. As can be seen, failing to compensate for the drift in the voltage can also lead to substantial errors in the estimation of the phase angles. Again, the compensation methods based on the simple linear approximation and the straight line LLS fit do reduce the errors in amplitude and phase estimation compared to the uncompensated case, but residual errors remain significant at low frequencies. The RPS method, on the other hand, correctly compensates for the drift in the signal and residual errors are negligible.

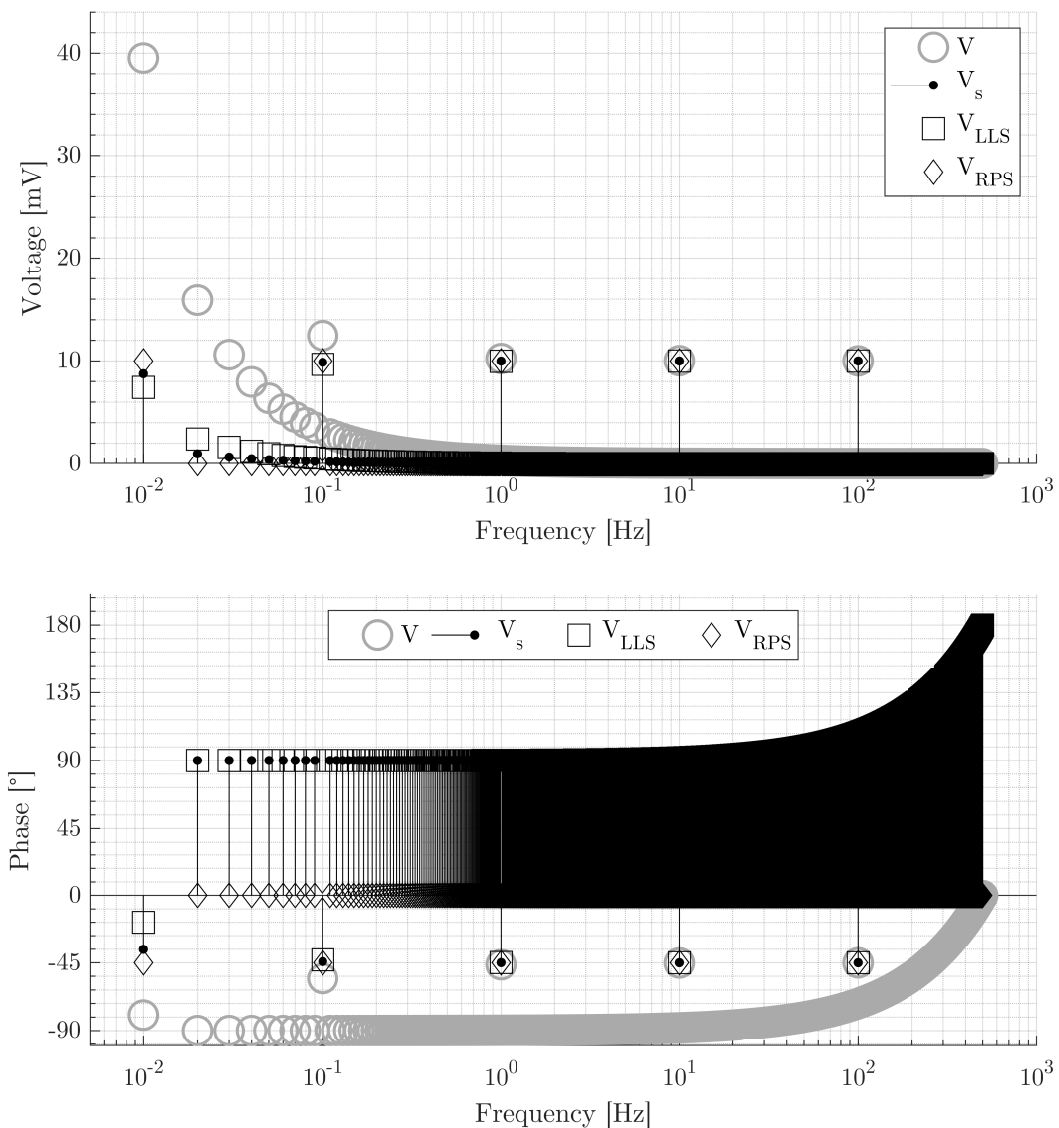


Figure 5.24: Magnitude and phase spectra of compensated and uncompensated multisine signal (Example 2): the signal is shown before being compensated (V), as well as after being compensated using the simple linear approximation (V_s), the straight line LLS method (V_{LLS}), and the RPS method (V_{RPS}).

Table 5.2: Results of compensation example 2: magnitude, phase angle, relative magnitude error, $e_{r,m}$, and absolute phase error, $e_{a,p}$. Results given to two significant figures. Any $e_{r,m} < 0.01\%$ and $e_{a,p} < 0.1^\circ$ rounded to zero.

f [Hz]	0.01	0.1	1	10	100
No compensation					
$ V(f) $ [mV]	40	12	10	10	10
$\angle V(f)$ [°]	-80	-55	-46	-45	-45
$e_{r,m}$ [%]	300	25	2.3	0.23	0.029
$e_{a,p}$ [°]	35	10	1.3	0.12	0
Simple linear approximation					
$ V_s(f) $ [mV]	8.8	9.9	10	10	10
$\angle V_s(f)$ [°]	-36	-44	-45	-45	-45
$e_{r,m}$ [%]	12	1.3	0.14	0.014	0
$e_{a,p}$ [°]	8.9	0.79	0	0	0
Straight line LLS					
$ V_{LLS}(f) $ [mV]	7.4	9.7	10	10	10
$\angle V_{LLS}(f)$ [°]	-18	-43	-45	-45	-45
$e_{r,m}$ [%]	26	3.3	0.34	0.035	0
$e_{a,p}$ [°]	27	2	0.19	0	0
Ramp plus sine compensation					
$ V_{RPS}(f) $ [mV]	10	10	10	10	10
$\angle V_{RPS}(f)$ [°]	-45	-45	-45	-45	-45
$e_{r,m}$ [%]	0	0	0	0	0
$e_{a,p}$ [°]	0	0	0	0	0

The final example is a repeat of the first example, except that a zero-mean white noise signal with a standard deviation of 10 mV has been added to the multisine. The results are summarised in Table 5.3. Comparing the results to those of Table 5.1 shows that whilst adding noise to the signal has somewhat increased the amplitude and phase estimation errors in the straight line LLS and RPS compensation methods, these methods are considerably more robust than the simple linear approximation method, whose relative amplitude estimation errors have trebled compared to the noise-free fit.

Table 5.3: Results of compensation example 3: magnitude, phase angle, relative magnitude error, $e_{r,m}$, and absolute phase error, $e_{a,p}$. Results given to two significant figures. Any $e_{r,m} < 0.01\%$ and $e_{a,p} < 0.1^\circ$ rounded to zero.

f [Hz]	0.01	0.1	1	10	100
No compensation					
$ V(f) $ [mV]	42	13	10	10	9.9
$\angle V(f)$ [°]	-90	-90	-90	-90	-90
$e_{r,m}$ [%]	300	31	2.4	0.32	0.6
$e_{a,p}$ [°]	0	0.25	0.21	0.17	0.1
Simple linear approximation					
$ V_s(f) $ [mV]	3.1	9.2	9.9	10	9.9
$\angle V_s(f)$ [°]	-90	-90	-90	-90	-90
$e_{r,m}$ [%]	69	7.5	1.5	0.069	0.64
$e_{a,p}$ [°]	8.9	0.79	0	0	0
Straight line LLS					
$ V_{LLS}(f) $ [mV]	3.2	9.3	9.9	10	9.9
$\angle V_{LLS}(f)$ [°]	-90	-90	-90	-90	-90
$e_{r,m}$ [%]	68	7.4	1.5	0.068	0.64
$e_{a,p}$ [°]	0.17	0.35	0.23	0.16	0.11
Ramp plus sine compensation					
$ V_{RPS}(f) $ [mV]	9.9	9.9	9.9	10	9.9
$\angle V_{RPS}(f)$ [°]	-90	-90	-90	-90	-90
$e_{r,m}$ [%]	0.88	0.73	0.79	0	0.63
$e_{a,p}$ [°]	0	0.33	0.23	0.16	0.11

To summarise, it has been shown that a drift in the measured voltage can lead to considerable errors in both the amplitude and phase of the measurement components, which will directly propagate into the estimated impedance values. Two existing compensation methods, one based on a simple linear approximation and the other on a straight line LLS fit, were analysed and it was shown that they both have significant limitations related to their ability to accurately fit sinusoidal and multisine signals. A new compensation method was described that uses the LLS algorithm to fit a signal consisting of the sum of a ramp and a sinusoid. It was shown that the RPS fit can be used to accurately extract the characteristic slope of a linear

drift present in both sine and multisine signals — irrespective of sinusoidal phase angles. Finally, it was shown that the identified slope can be used to remove the linear drift from the voltage signal; this leads to more accurate amplitude and phase estimates and thus ultimately enables higher-accuracy impedance measurements.

5.3.7 Sampling below Nyquist rate

Battery packs used in EVs and battery energy storage systems typically consist of hundreds of cells connected in different series and parallel arrangements, which creates the requirement for a large number of ADC input channels. From an economical point of view, monitoring battery packs of such scale may therefore require the use of multiplexed ADCs. This limits the maximum effective sampling rate per channel, and consequently also the highest achievable measurement frequency. A solution to this, which allows the upper end of the measurement bandwidth to be extended under these circumstances, was proposed by the author in [59]. A short summary of the method's fundamental principle follows.

If f_s is the effective channel sampling rate, then it is known from the sampling theorem that it is possible to reconstruct any signal whose bandwidth is limited to $< f_s/2$. From the signal's point of view, if its bandwidth is limited to f_B , then in order to be able to properly reconstruct the signal, it should be sampled at a rate $> 2f_B$, i.e. the Nyquist rate. If the signal's bandwidth is *not* limited and/or it contains components at frequencies $\geq f_s/2$, *aliasing* will occur: those components $\geq f_s/2$ will be folded back into the frequency range below $f_s/2$ and distort the original signal.

In EIS measurements we have control over the frequency content of the signals, so with a careful choice of the frequency components of the measurement perturbation, the observation window length, and the sampling frequency, it is possible to sample frequencies above $f_s/2$ and avoid overlapping frequency content and the associated loss of information.

The proposed method was verified experimentally using the power converter-based EIS setup developed in this work. The associated results — first reported in [59] — are reproduced in Chapter 7.

It is also possible to use the same principle explained above, to reduce the volume of data that requires both processing and storing; however, this was outside the scope of this work.

5.4 Power converter design and control strategy

The principle of power converter-based EIS is to introduce the measurement perturbation required for the impedance determination through control of the power converter. The *control strategy*, or detailed operation of the controller, is therefore expected to be different to a regular power converter implementation in some respects. As will be shown in Section 5.4.1, power converter-based EIS also requires different considerations to be made with regards to the hardware design of the power converter compared to a more conventional converter application.

5.4.1 Power converter design

In regular power converter applications low-pass filters are normally used to smooth out the switching ripple and remove the switching harmonics from the converter output waveforms. In some cases, filters are also used on the converter's input side, for example, if the input source cannot handle switching ripple, or to reduce conducted electromagnetic interference [60]. Similarly to the filters on the output side, input filters in regular power converter applications thus tend to be used to filter out very high frequencies, i.e. the switching frequency and its harmonics, as well as high frequency transients associated with parasitic elements in the switches. Since switching converters typically use frequencies in the order of 10^4 to 10^6 Hz, it is possible to use low-order filters, as the filters' bandwidth can be made much smaller than the switching frequency. For example, a passive second order filter can attenuate the switching ripple by 80 dB when its corner frequency is placed two decades below the switching frequency. Employing the same filter design in a converter used for EIS could attenuate the measurement perturbation, which would mean that the amplitude of the measurement perturbation *before* being filtered, i.e. in the converter waveforms, would need to increase accordingly in order to make up for the attenuation in the filter. To avoid this problem, whilst still allowing the switching ripple to be filtered, thus requires a little more care in the design of the input/output filters. In order to achieve a given attenuation of the switching ripple, the filters may require a much steeper roll-off, particularly if the measurement perturbation contains high frequency components. The filter design needs to take into account the wider system's setup and specific use cases. In general, the filter between the battery and the converter needs to be designed in such a way that the measurement perturbation components fall into the filter's *passband*, whereas the switching frequency falls into its *stopband*. Therefore, if the system should allow the flexibility of connecting the battery on either side of the converter, both input

and output filters should meet this requirement. If, on the other hand, the battery always remains on the same side of the converter (e.g. if a bi-directional converter is used), then the filter on the *opposite* side of the battery can be designed such that both the switching ripple *and* the measurement perturbation are within the stopband. In this latter case, reducing the effect of the excitation signal on the rest of the system can therefore be achieved through proper filter design.

Another consideration in the design of the power converter is the relationship between the switching frequency and the highest measurement frequency. In order to obtain a reasonably clean sinusoidal waveform (i.e. with little distortion due to non-linearity), the switching frequency should be at least ten times greater than the highest frequency component to be contained in the measurement perturbation. Due to the high switching frequencies of typical power converters, there are likely to be other constraints in the converter-based EIS system that limit the upper boundary of the measurement bandwidth before the switching frequency becomes a limiting factor, e.g. transducer bandwidth, ADC conversion rates, or filter specifications.

5.4.2 Control strategy

The development of a control strategy for a power converter-based EIS system depends on a number of fundamental decisions: the perturbation injection method (open or closed-loop); the variable that is to be controlled; the type of measurement perturbation; and the configuration and required flexibility of the system (charge, discharge, or both).

Open and closed-loop perturbation injection methods were compared in 4.3.2. Given its advantages, it was decided to focus on the closed-loop method in this work. As stated before, perturbing the battery current is generally not as sensitive to errors as perturbing the voltage (due to the low impedance of LIBs); thus it was decided to control the battery current. Further, due to its adaptability and its relative ease of implementation, a multisine measurement perturbation was opted for. Section 4.3.4 described the differences between the charge and discharge setups and highlighted that two of the outstanding challenges in power converter-based EIS are associated with the discharge setup, namely: 1) allowing simultaneous control of the converter input and output, in order to both introduce the measurement perturbation and to create a regulated output, and 2) reducing the effect of the measurement perturbation on the load. For this reason, the focus in this work was on the discharge setup.

A unidirectional power converter was used in this work. It was also desired to create a flexible experimental setup that would allow impedance measurements to

be carried out at any point during both charge and discharge cycles. As a result, both the input and output filters of the converter had to be designed so as to allow the measurement perturbation through unattenuated. Filter design could therefore not be explored as a solution to 2). Instead, two different ideas are proposed here to reduce the effect of the measurement perturbation on the load when filtering is not an option. Both involve the use of *two* battery packs with identical specifications.

The first method is centred around a control strategy in which the two battery packs assume different roles: one is controlled to provide a constant output power, whilst the second one uses a much faster output control in order to compensate for the perturbations created by the first. One downside of this approach is that only one of the battery packs can be monitored at a time, as the pack with the dynamic output can not be guaranteed to be in steady state; however, the roles of the two battery packs could be alternated regularly, so as to allow impedance measurements to be obtained from both battery packs on a periodic basis.

The second method proposed to reduce perturbations in the load has the advantage that it would allow both battery packs to be monitored simultaneously. This could be achieved by perturbing the current of both packs with identical excitation signals, except for inverting the relative phases of the measurement components, which would lead to their effective cancellation at the output.

Due to the requirements of more complex control strategies and the doubling in battery pack size, the experimental validation of the proposed solutions fell outside the scope of this work. Instead, the implementation focused on the first challenge described above, i.e. allowing simultaneous control of the converter's input and output.

To summarise, the control strategy in this work centers around the closed-loop measurement perturbation injection of a multisine signal in the battery (converter input) current with simultaneous control of the output voltage in order to provide a regulated output.

The control architecture chosen to realise the proposed strategy is that of a double-loop structure. The control of the battery current has to be relatively fast in order to allow high frequency perturbations to be introduced. The control of the output voltage, on the other hand, can be a lot slower when the load is in steady state, since the controller only needs to compensate for the slowly drifting input (battery) voltage. Based on the differing dynamics of the two controllers, it should therefore be possible to nest the two feedback loops, i.e. to place the current loop *inside* of the voltage loop, and to tune them individually. One downside of this approach is that if the output voltage should experience a relatively fast disturbance,

e.g. due to a change in the load impedance, the controller would be unable to react to it very quickly. A relatively simple (although not ideal) solution to this problem is to temporarily switch from the double-loop structure with slow output voltage control to a single-loop controller with fast voltage control. Whilst this would cause any active impedance measurement acquisition to be stopped, load regulation would normally have higher priority than impedance measurement acquisition. Also, for a valid impedance measurement the battery must be in steady state — the occurrence of large transients would therefore most likely invalidate any impedance acquired under such conditions anyway.

The development of the double-loop control architecture was carried out primarily on the basis of computer simulation. There were two main reasons for this: firstly, this allowed an analytical approach, i.e. classical control engineering techniques could be employed to design the controller; secondly, simulation offers a safe and non-destructive iterative design process. The downside of using this approach is that it neglects many of the non-idealities of the system, e.g. parasitics, and is therefore generally unlikely to result in a controller design that is accurate enough to work in a practical implementation without further modifications; however, it does allow a general understanding of the dynamics of the system to be gained and can thus simplify a practical implementation.

To allow simulation, a model of the plant, i.e. the converter, had to be developed. Since power converters are inherently non-linear and many of the classic control theory techniques only work for linear time-invariant (LTI) systems, the plant model was linearised around an operating point. A small-signal model of the converter was derived using the state-space averaging method [61]. From this, the relevant transfer functions could be obtained which were subsequently used to simulate the system response and design the control architecture.

Chapter 6

Implementation

6.1 Batteries

The selection and design of several parts of the experimental setup, such as the power converter, transducers, and signal conditioning circuits, depended — to some extent — on the size and specification of the battery pack (mainly its output voltage and current). The selection of a suitable battery was therefore one of the first tasks carried out.

One of the novel contributions of this project is the use of power converter-based EIS on a much greater scale than any previous work. As explained in Section 4.3.4, previous works have been limited to using one or two cells; the objective in this work was thus to create a setup involving a cell-arrangement consisting of double figures. Taking into account the available financial resources, a 16S1P battery pack configuration was decided upon, i.e. 16 series, 1 parallel. Another consideration in this choice was that — together with the boost converter topology — this configuration would result in a sufficiently high voltage to drive an inverter stage and an induction motor rated at several kilowatts.

Lithium-iron-phosphate (LFP) was chosen as the type of LIB for this project. The main reason behind this choice was the inherent safety advantage of LFP batteries over other types of LIBs, such as lithium-cobalt-oxide (LCO).

The module that was chosen as the building block for the battery pack was the K2B3V90E: a LFP module manufactured by K2 Energy [62]. Each module consists of twenty-eight 3.2 Ah cells connected in parallel, resulting in an effective module capacity of approximately 90 Ah. With sixteen modules connected in series, this resulted in a battery pack with an effective nominal output voltage of 51.2 V and a nominal energy capacity of approximately 4.6 kWh. The battery pack is shown in Figure 6.1.



Figure 6.1: Battery pack composed of sixteen series-connected K2B3V90E modules. Approximate dimensions (L×W×H): 650×425×290 mm.

To put the pack's energy capacity into context, Table 6.1 lists the energy capacity of the battery packs found in a selection of BEVs available in the UK at the time of writing.

In addition to this relatively large battery pack, a smaller pack consisting of individual cells was set up. The reason for this was twofold: initially, whilst the experimental system was still under development, it enabled the use of lower charge and discharge currents; later, it would be possible to use this smaller battery pack as a reference to establish how well the proposed system worked on larger modules composed of multiple cells in parallel. The cell chosen for this smaller pack was the K226650E02 [32]: a cylindrical 26650 LFP cell¹ that is also manufactured by K2 Energy and is in fact the same cell that is used in the K2B3V90E module. Again, the same 16S1P configuration was used as in the larger battery pack, resulting in the same effective nominal output voltage (51.2 V), but a lower nominal energy capacity of approximately 164 Wh.

A printed circuit board (PCB) was designed to securely hold the cells in place and to provide a fuse-protected output connection, as well as voltage sense paths for all sixteen cells via a DB-25 D-sub connector. This is shown in Figure 6.2.

¹i.e. 26 mm in diameter and 65 mm in length

Table 6.1: Sample of battery electric vehicles available in UK in 2022.

Manufacturer	Model	Battery Capacity (kWh)
BMW	i3	37.9
Ford	Mustang Mach-E	75.7 / 98.7
Honda	e	35.5
Hyundai	Ioniq	40.4
Hyundai	Kona	39 / 64
Kia	Niro EV	64.8
Kia	Soul EV	64
Nissan	Leaf	39 / 59
Peugeot	e-208	50
Renault	Zoe E-Tech electric	52
Renault	Twizy	6.1
Skoda	Enyaq iV	77
Smart	Fortwo	17.6
Tesla	Model 3	60 / 75
Tesla	Model S	75 / 100
Vauxhall	Corsa-e	50
Volkswagen	ID.3	58
Volkswagen	ID.4	77

**Figure 6.2:** Battery pack composed of sixteen series-connected K226650E02 cells on be-spoke PCB. Approximate dimensions (L×W): 286×219 mm.

The voltage sense connections were placed as close to the battery terminals as possible in order to reduce the loop inductance and to eliminate from the measurement the voltage drop appearing across the battery connectors. On the small battery pack, despite placing the voltage sense connections as close to the cell terminals as possible, the voltage drop across the cell holder terminals cannot be completely excluded from the measurement due to the type of cell holder used. Compensation for the contact resistance is thus required.

A protection circuit module (PCM) was used to further enhance the safety of the experimental setup. This provided protection against several conditions: over charge and over discharge of any individual cell/module, over current, over temperature, as well as short-circuit detection.

6.2 Power converter design and control strategy

6.2.1 Power converter design

The power converter used in this work is a dual-phase interleaved boost converter (IBC). The converter's topology is depicted in the schematic diagram of Figure 6.3 and a photograph is shown in Figure 6.4.

The converter is rated for a maximum output power of 2 kW and is operated at a switching frequency of 100 kHz. Its key specifications are listed in Table 6.2.

Whilst power converter-based EIS in principle works with any type of converter, a step-up dc-dc converter was chosen in this work in order to boost the relatively low voltage inherent in batteries. In the dual-phase interleaved topology, the two PWM signals used to drive the switches, Q_1 and Q_2 , are 180° out-of-phase with each other. This can lead to a reduction in the input current ripple. In the particular case of a 50 % duty cycle — depicted in Figure 6.5 — there is (in theory) *no* input current ripple.

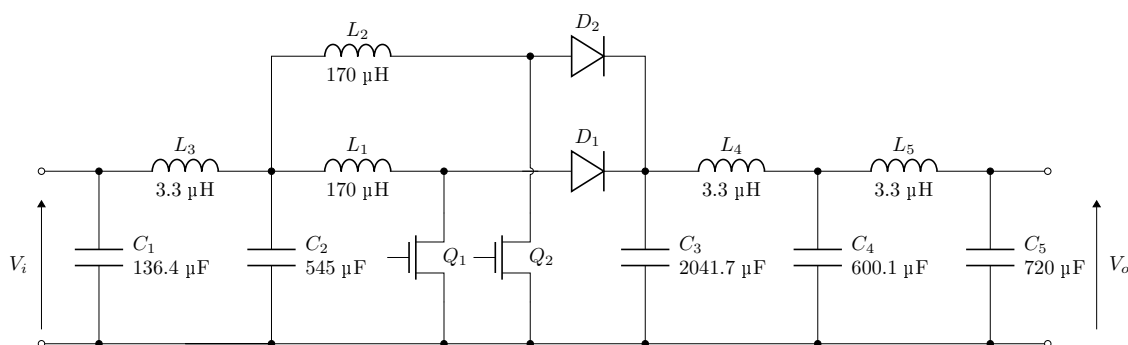




Figure 6.3: Schematic diagram of the dual-phase interleaved boost converter.



(a) Top side.



(b) Bottom side.

Figure 6.4: The dual-phase IBC. Approximate dimensions (L×W×H): 440×200×100 mm.**Table 6.2:** Main specifications of dual-phase interleaved boost converter.

Maximum input voltage	100 V
Maximum output voltage	350 V
Maximum output power	2 kW
Operating switching frequency	100 kHz
Maximum switching frequency	150 kHz




Figure 6.5: Input current ripple in the case of a 50 % duty cycle. d_1 and d_2 represent the switching waveforms for Q_1 and Q_2 , respectively.

The interleaved topology also has the benefit that the input power splits between the two parallel paths; the power rating of the components on each path, $L_{1/2}$, $Q_{1/2}$, and $D_{1/2}$, can thus be halved compared to the regular boost converter topology.

6.2.1.1 Input and output filters

As discussed in Section 5.4.1, in power converter-based EIS the design of the converter’s input and output filters depends not only on the bandwidth of the measurement perturbation but also on the particular setup used. In this work the flexibility was desired to connect the batteries on either side of the converter and therefore both filters were designed so as *not* to attenuate the measurement perturbation.

Figure 6.3 shows that both the input and output filters of the converter are passive filters using only capacitors and inductors. The input filter consists of C_1 , L_3 , and C_2 . To understand the effect of the input filter on the measurement perturbation its frequency response is considered. This is approximated by ignoring component non-idealities, modelling the converter as an independent current source, i_{ref} , and the battery impedance, Z_B , as the “load” — as shown in Figure 6.6.

A small-signal ac analysis using the indicated component values is performed in order to obtain the frequency response (Bode) plot of Figure 6.7. At very low frequencies the capacitors act as open circuits and the inductor acts as a short circuit, consequently a unity dc gain is observed. Two responses are shown in Figure 6.7, since the value of Z_B changes significantly between the two battery packs. A value




Figure 6.6: Model used to analyse the current transfer function of the input filter. Note that the component values are those from the IBC design shown in Figure 6.3.




Figure 6.7: Bode plot of input filter current transfer function.

of 1Ω is used for the small battery pack (SBP), as this corresponds approximately to its impedance at 1 kHz. It can be seen that the transfer function has a real pole at around 230 Hz after which the magnitude response decays at a rate of 20 dB per decade; this is followed by a complex pole pair at approximately 8.3 kHz after which the magnitude decays at 60 dB per decade. It should be noted that the 3 dB cut-off point is inversely proportional to the battery impedance, i.e. the lower the battery's impedance, the higher the cut-off frequency. The modules used in the large battery pack (LBP) consist of 28 cells connected in parallel, their nominal impedance magnitude is therefore 1/28th of the cells'. As the dashed line in the frequency response plot shows, when the LBP is used the filter's stopband starts at around 5.7 kHz.

The preceding analysis also shows that in certain circumstances the input filter can be left out of the converter model without loss of accuracy. The particular

condition that must be fulfilled for this to be the case can be stated (approximately) by saying that the battery impedance at the highest measurement perturbation frequency should be much lower than the impedance of C_2 . If this condition is met, the current through the battery impedance, i_b , will be approximately equal to i_{ref} .

The output filter consists of C_3 , L_4 , C_4 , L_5 , and C_5 . Again, a small-signal ac analysis is used to obtain the frequency response of the output filter. In this case the converter is more accurately modelled as an independent voltage source due to its low output impedance. A resistive load is considered, whose range is determined based on the converter's minimum and maximum power and voltage ratings: the load resistance *minimum* is found by assuming minimum output voltage (approximately 50 V) and maximum output power, i.e. $R_{L(\text{min})} = \frac{V_{o(\text{min})}^2}{P_{o(\text{max})}} = \frac{50^2}{2000} = 1.25 \Omega$; the load resistance *maximum*, on the other hand, is determined by assuming maximum output voltage and minimum output power (assumed to be 1 % of maximum): $R_{L(\text{max})} = \frac{V_{o(\text{max})}^2}{P_{o(\text{min})}} = \frac{180^2}{20} = 1620 \Omega$. The equivalent circuit used in this analysis is shown in Figure 6.8. The frequency response of the voltage transfer function is illustrated in Figure 6.9, showing both responses — with minimum and maximum load.

Figure 6.9 confirms that, for the given range of load resistances ($[1.25, 1260] \Omega$), the output filter does not attenuate frequencies below approximately 1 kHz. Above the corner frequency, the characteristic fourth order filter roll-off of 80 dB per decade can be observed.

The analysis was repeated by replacing the inductors L_4 and L_5 with short circuits and the capacitors C_3 , C_4 , and C_5 with a single capacitor with a capacitance equal to the sum of the three capacitors. The frequency response of this simplified arrangement is shown in Figure 6.10 along with the response of the actual circuit (including the five separate reactive elements); in both cases component non-idealities have been ignored and a load resistance of 1.25Ω was used.

Figure 6.10 shows that there is good agreement between the two circuits. This




Figure 6.8: Model used to analyse the voltage transfer function of the output filter.




Figure 6.9: Bode plot of output filter voltage transfer function.




Figure 6.10: Bode plot of output filter impedance transfer function.

enables the derivation of the converter model to be simplified by replacing the output filter with a single equivalent output capacitor equal in value to the sum of the three capacitors C_3 , C_4 , and C_5 .

6.2.2 Control strategy

As explained in Section 5.4.2, the control strategy in this work is centred around the closed-loop measurement perturbation injection of a multisine signal in the input (battery) current with simultaneous control of the output voltage in order to provide a regulated output. It was also explained that the design of the controllers was undertaken using computer simulation and that to allow this to take place, a model of the power converter (i.e. the dual-phase IBC) had to be developed first. This is covered in Section 6.2.2.1.

6.2.2.1 Derivation of converter model

A note regarding the symbols used in this section. Lowercase symbols with hat notation, e.g. \hat{i} , are used to denote small signals, i.e. small perturbations around the dc operating point or steady-state value around which the model is linearised. Uppercase symbols, e.g. I , are used to denote the steady-state values. Lowercase symbols, e.g. i , denote *large* signals, i.e. $i = I + \hat{i}$. Dot notation, e.g. \dot{v} , is used to indicate (time) derivatives. A bold font is used to denote matrices, e.g. \mathbf{x} .

The small-signal model of the dual-phase IBC is derived using the state-space averaging method [61]. The derivation is loosely based on the work found in [63].

The process is split into four sequential parts:

1. Large-signal model
2. Averaged large-signal model
3. Steady-state model
4. Small-signal model

The parts are named after their output, i.e. each part results in a different model of the converter. The output of the last stage is the small-signal model, i.e. a *linear* model that describes the converter's dynamics around a particular operating point, or steady-state value. From this it is possible to determine the transfer functions that allow the controller design to be carried out, e.g. the control to input current transfer function, $\hat{i}_i(s)/\hat{d}(s)$.

Using the simplifications associated with the input and output filters described in Section 6.2.1.1, the circuit diagram of the dual-phase IBC is redrawn as shown in Figure 6.11 in order to illustrate the quantities and components that are relevant to the model derivation. Note that $C = C_3 + C_4 + C_5$ (cf. Figure 6.3).

Since the inductors in each path are identical, L and r_L will be used to refer, respectively, to the inductance and the parasitic resistance of a single inductor, i.e. $L = L_1 = L_2$ and $r_L = r_{L1} = r_{L2}$.

Aside from the parasitic resistance of the inductor, component non-idealities, such as the parasitic resistances of the capacitors and switches, or the diode's forward voltage, have not been included in the model. The reasoning behind this is that whilst their inclusion would have led to a slight improvement in model accuracy, it would still not have led to perfect agreement between simulated and measured results, as it is difficult to account for all non-idealities and often they are not fully known in the first place. The purpose of deriving the converter model is not to

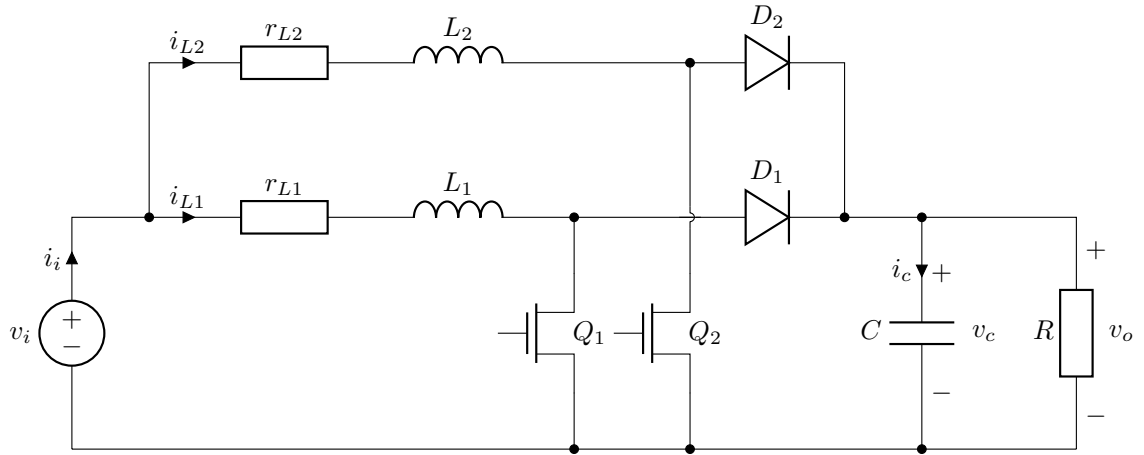


Figure 6.11: Circuit diagram of dual-phase interleaved boost converter used in model derivation.

obtain a perfectly accurate model, but instead to obtain a model that is accurate enough to guide the general design of the controllers, and to inform the controller tuning process. The omission of most component non-idealities still allows this, whilst considerably simplifying the model derivation.

The state-space averaging method uses *state-space representation* to describe the dynamics of the system. The standard form of the state-space representation is given by the two expressions:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \quad (6.1)$$

$$\mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{D}\mathbf{u} \quad (6.2)$$

where (6.1) is known as the *state equation*, \mathbf{x} are the state variables, \mathbf{u} are the input variables, \mathbf{A} is the state matrix, and \mathbf{B} is the input matrix.

(6.2) is known as the *output equation*: \mathbf{y} are the system's outputs, \mathbf{C} is the output matrix, and \mathbf{D} is the feedforward matrix.

Figure 6.11 shows that there are three energy-storing elements in the model (i.e. it is a third order system) and therefore there are three state variables:

$$\mathbf{x} = \begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix} \quad (6.3)$$

1. Large-signal model The assumption is made that the converter operates in continuous conduction mode, i.e. the inductor currents never go to zero. With this assumption in place, and given that the converter uses two switches, the circuit can

assume a total of four modes, or configurations, depending on the state (open/closed) of each switch. The large-signal model therefore consists of four sets of state and output equations — one for each mode.

The different modes will be referred to using a two-bit notation: the first bit representing the state of switch Q_1 , and the second the state of switch Q_2 ; a ‘1’ indicates that the switch is closed, whereas a ‘0’ means the switch is open. For example: mode 10 refers to the circuit configuration in which Q_1 is closed and Q_2 is open.

The state equations for the four different modes are given in (6.4) to (6.7).

$$\dot{\mathbf{x}}_{00} = \begin{bmatrix} \dot{i}_{L1} \\ \dot{i}_{L2} \\ \dot{v}_c \end{bmatrix} = \overbrace{\begin{bmatrix} -\frac{r_L}{L} & 0 & -\frac{1}{L} \\ 0 & -\frac{r_L}{L} & -\frac{1}{L} \\ \frac{1}{C} & \frac{1}{C} & -\frac{1}{RC} \end{bmatrix}}^{\mathbf{A}_{00}} \overbrace{\begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix}}^{\mathbf{x}} + \overbrace{\begin{bmatrix} \frac{1}{L} \\ \frac{1}{L} \\ 0 \end{bmatrix}}^{\mathbf{B}_{00}} \overbrace{v_i}^{\mathbf{u}} \quad (6.4)$$

$$\dot{\mathbf{x}}_{10} = \begin{bmatrix} \dot{i}_{L1} \\ \dot{i}_{L2} \\ \dot{v}_c \end{bmatrix} = \begin{bmatrix} -\frac{r_L}{L} & 0 & 0 \\ 0 & -\frac{r_L}{L} & -\frac{1}{L} \\ 0 & \frac{1}{C} & -\frac{1}{RC} \end{bmatrix} \begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix} + \begin{bmatrix} \frac{1}{L} \\ \frac{1}{L} \\ 0 \end{bmatrix} v_i \quad (6.5)$$

$$\dot{\mathbf{x}}_{01} = \begin{bmatrix} \dot{i}_{L1} \\ \dot{i}_{L2} \\ \dot{v}_c \end{bmatrix} = \begin{bmatrix} -\frac{r_L}{L} & 0 & -\frac{1}{L} \\ 0 & -\frac{r_L}{L} & 0 \\ \frac{1}{C} & 0 & -\frac{1}{RC} \end{bmatrix} \begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix} + \begin{bmatrix} \frac{1}{L} \\ \frac{1}{L} \\ 0 \end{bmatrix} v_i \quad (6.6)$$

$$\dot{\mathbf{x}}_{11} = \begin{bmatrix} \dot{i}_{L1} \\ \dot{i}_{L2} \\ \dot{v}_c \end{bmatrix} = \begin{bmatrix} -\frac{r_L}{L} & 0 & 0 \\ 0 & -\frac{r_L}{L} & 0 \\ 0 & 0 & -\frac{1}{RC} \end{bmatrix} \begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix} + \begin{bmatrix} \frac{1}{L} \\ \frac{1}{L} \\ 0 \end{bmatrix} v_i \quad (6.7)$$

The output equation is the same for all modes:

$$\underbrace{\begin{bmatrix} v_o \\ i_i \end{bmatrix}}_{\mathbf{y}} = \underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}}_{\mathbf{C}} \underbrace{\begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix}}_{\mathbf{x}} \quad (6.8)$$

2. Averaged large-signal model The next step is to combine the four models to create an averaged large-signal model. This is done by finding a weighted average of the matrices, \mathbf{A} , \mathbf{B} , and \mathbf{C} . Since both \mathbf{B} and \mathbf{C} do not change between the various circuit configurations, only \mathbf{A} needs to be averaged. The average is determined by using the duty cycle, d , as the weight, since this represents the proportion of the switching period that each mode is active for. For example, as Figure 6.5 showed, when $d = 0.5$, only two modes are active during any switching cycle: modes 10

and 01; this is because only one out of the two switches is closed at any time, thus $d_{10} = d_{01} = 0.5$.

The average of matrix \mathbf{A} is given by:

$$\mathbf{A} = d_{00}\mathbf{A}_{00} + d_{11}\mathbf{A}_{11} + d_{01}\mathbf{A}_{01} + d_{10}\mathbf{A}_{10} \quad (6.9)$$

Assuming the converter is operating in steady state, there is no value of d for which all four modes are active during one switching period. When $d < 0.5$, there are two intervals in each switching period during which *both* switches are *open*, but at no point are both switches closed simultaneously, i.e. $d_{11} = 0$. When $d > 0.5$, on the other hand, there are two intervals in each switching period during which both switches are *closed*, but they are never open at the same time, i.e. $d_{00} = 0$. Thus, if the $d < 0.5$ case is considered, the $d_{11}\mathbf{A}_{11}$ term will drop out of (6.9) and if $d > 0.5$ is used, $d_{00}\mathbf{A}_{00}$ will drop out instead. However, regardless of which one is used, the result of (6.9) is the same in both cases, namely:

$$\mathbf{A} = \begin{bmatrix} -\frac{r_L}{L} & 0 & -\frac{d'}{L} \\ 0 & -\frac{r_L}{L} & -\frac{d'}{L} \\ \frac{d'}{C} & \frac{d'}{C} & -\frac{1}{RC} \end{bmatrix} \quad (6.10)$$

where $d' = 1 - d$.

The averaged large-signal model is thus given by:

$$\underbrace{\begin{bmatrix} \dot{i}_{L1} \\ \dot{i}_{L2} \\ \dot{v}_c \end{bmatrix}}_{\dot{\mathbf{x}}} = \underbrace{\begin{bmatrix} -\frac{r_L}{L} & 0 & -\frac{d'}{L} \\ 0 & -\frac{r_L}{L} & -\frac{d'}{L} \\ \frac{d'}{C} & \frac{d'}{C} & -\frac{1}{RC} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix}}_{\mathbf{x}} + \underbrace{\begin{bmatrix} \frac{1}{L} \\ \frac{1}{L} \\ 0 \end{bmatrix}}_{\mathbf{B}} \underbrace{v_i}_{\mathbf{u}} \quad (6.11)$$

$$\underbrace{\begin{bmatrix} v_o \\ i_i \end{bmatrix}}_{\mathbf{y}} = \underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}}_{\mathbf{C}} \underbrace{\begin{bmatrix} i_{L1} \\ i_{L2} \\ v_c \end{bmatrix}}_{\mathbf{x}} \quad (6.12)$$

3. Steady-state model The steady-state model can be obtained from the averaged large-signal model by rewriting *large* signal values as *steady-state* values and setting derivative terms to zero. The state and output equations of the steady-state model thus become:

$$\underbrace{\dot{\mathbf{x}}}_{0} = \underbrace{\mathbf{A}}_{\begin{bmatrix} -\frac{r_L}{L} & 0 & -\frac{D'}{L} \\ 0 & -\frac{r_L}{L} & -\frac{D'}{L} \\ \frac{D'}{C} & \frac{D'}{C} & -\frac{1}{RC} \end{bmatrix}} \underbrace{\mathbf{X}}_{\begin{bmatrix} I_{L1} \\ I_{L2} \\ V_c \end{bmatrix}} + \underbrace{\mathbf{B}}_{\begin{bmatrix} \frac{1}{L} \\ \frac{1}{L} \\ 0 \end{bmatrix}} \underbrace{\mathbf{U}}_{V_i} \quad (6.13)$$

$$\underbrace{\begin{bmatrix} V_o \\ I_i \end{bmatrix}}_{\mathbf{Y}} = \underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}}_{\mathbf{C}} \underbrace{\begin{bmatrix} I_{L1} \\ I_{L2} \\ V_c \end{bmatrix}}_{\mathbf{X}} \quad (6.14)$$

4. Small-signal model The final step is to determine the small-signal model. This is achieved by expanding all large signals in the averaged large-signal model (i.e. rewriting, e.g., i_{L1} as $I_{L1} + \hat{i}_{L1}$), and eliminating from it all terms that appear in the steady-state model, as well as any product of two small signal terms. Finally, the state equation is rewritten so as to include the duty cycle in the input vector, $\hat{\mathbf{u}}$.

The small-signal model of the dual-phase IBC is given by:

$$\underbrace{\begin{bmatrix} \dot{\hat{x}} \\ \dot{\hat{i}}_{L1} \\ \dot{\hat{i}}_{L2} \\ \dot{\hat{v}}_c \end{bmatrix}}_{\dot{\hat{\mathbf{x}}}} = \underbrace{\begin{bmatrix} -\frac{r_L}{L} & 0 & -\frac{D'}{L} \\ 0 & -\frac{r_L}{L} & -\frac{D'}{L} \\ \frac{D'}{C} & \frac{D'}{C} & -\frac{1}{RC} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} \hat{x} \\ \hat{i}_{L1} \\ \hat{i}_{L2} \\ \hat{v}_c \end{bmatrix}}_{\hat{\mathbf{x}}} + \underbrace{\begin{bmatrix} \frac{1}{L} & \frac{V_o}{L} \\ \frac{1}{L} & \frac{V_o}{L} \\ 0 & -\frac{(I_{L1}+I_{L2})}{C} \end{bmatrix}}_{\mathbf{B}} \underbrace{\begin{bmatrix} \hat{v}_i \\ \hat{d} \end{bmatrix}}_{\hat{\mathbf{u}}} \quad (6.15)$$

$$\underbrace{\begin{bmatrix} \hat{v}_o \\ \hat{i}_i \end{bmatrix}}_{\hat{\mathbf{y}}} = \underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}}_{\mathbf{C}} \underbrace{\begin{bmatrix} \hat{i}_{L1} \\ \hat{i}_{L2} \\ \hat{v}_c \end{bmatrix}}_{\hat{\mathbf{x}}} \quad (6.16)$$

Note that the small-signal model includes three steady-state values: I_{L1} , I_{L2} and V_o . These can be obtained from the steady-state model by rewriting its state equation in terms of \mathbf{X} and substituting the result into the output equation:

$$0 = \mathbf{A}\mathbf{X} + \mathbf{B}\mathbf{U} \quad (6.17)$$

$$\mathbf{X} = -\mathbf{A}^{-1}\mathbf{B}\mathbf{U} \quad (6.18)$$

$$\mathbf{Y} = \mathbf{C}\mathbf{X} = -\mathbf{C}\mathbf{A}^{-1}\mathbf{B}\mathbf{U} \quad (6.19)$$

The same steps — preceded by the Laplace transformation of state and output equations — are applied to the small-signal model in order to bring it into a form that directly relates the input and output vectors and can thus be used to determine any transfer function in the system.

6.2.2.2 Current controller

The purpose of the current controller is to allow the multisine measurement perturbation to be injected in the converter input current via closed-loop control. Both the small-signal model and the steady-state model of the dual-phase IBC (derived in the previous section) are entered into MATLAB in order to find the control to input current transfer function, i.e. $G(s) = \hat{i}_i(s)/\hat{d}(s)$. The component values used in the model are summarised in Table 6.3. (The relevant schematic diagram is shown in Figure 6.11.)

Table 6.3: Component values used in the small-signal model of the converter.

r_L	1 m Ω
L	170 μ H
C	3362 μ F

To begin with, the small-signal model was based around an operating point defined by a dc input voltage $V_i = 50$ V (corresponding approximately to the average output voltage of the battery packs), and a steady-state duty cycle $D = 0.5$ (i.e. a middle-of-the-range value). The load resistance was chosen as $R = 100$ Ω , as this presents a suitable load to both battery packs. Figure 6.12 shows the Bode plot of $G(s)$ at the stated operating point.

To ensure a) the stability of the system, b) the reduction of the steady-state error, and c) a suitable bandwidth that will enable the desired measurement fre-




Figure 6.12: Frequency response of control to input current transfer function at operating point: $V_i = 50$ V, $D = 0.5$, and $R = 100$ Ω .

quency range of 0.01–100 Hz to be achieved, a compensator, $H(s)$, is designed by following the process outlined in Section 10-5-3 of [55]. $H(s)$ takes the form of a lead compensator with an additional pole at the origin, i.e. an integrator. Its frequency response is shown in Figure 6.13 and its transfer function is given by:

$$H(s) = \frac{12000(s + 16840)}{s(s + 234500)} \quad (6.20)$$

The lead compensator introduces a phase boost that creates a phase margin of approximately 60° which reduces unwanted peaks in the closed-loop magnitude response. The integrator eliminates the steady-state error, and by placing the crossover frequency approximately one decade below the switching frequency, the usable bandwidth is maximised.

Figure 6.14 shows the closed-loop response using the above controller with the converter at the operating point described before. It can be seen that the compensator fulfills the performance criteria stated above in terms of: stability, steady-state error, and bandwidth.

The frequency response of the current controller, shown in Figure 6.13, allows the individual responses of the lead compensator and the integrator to be distinguished quite easily. It can be seen that if the controller were to be simplified by dropping the lead compensator and retaining only the integrator, this would result in a magnitude response closely resembling that of the more complex controller. Its phase response, on the other hand, would be a constant -90° . The phase boost provided by the lead compensator would thus be lost, which could lead to undesired peaks in the magnitude response of the closed-loop system; however, as Figure 6.15 shows,




Figure 6.13: Frequency response of lead compensator plus integrator.




Figure 6.14: Closed-loop frequency response with lead compensator plus integrator with $G(s)$ linearised around operating point defined by: $V_i = 50$ V, $D = 0.5$, and $R = 100$ Ω .

this peak lies above the highest frequency to be measured in this work (100 Hz) — even with the converter at its most unfavourable operating point. In order to simplify the implementation of the controller, the decision was thus made to use an integral-only controller. Should it be desired to further increase the upper measurement frequency limit in the future, the more sophisticated lead compensator plus integrator controller could be adopted.




Figure 6.15: Comparison of closed-loop frequency response obtained from two different controllers: integral controller with $K_i = 100$, and lead compensator plus integrator. The operating point of $G(s)$ was chosen so as to minimise the frequency of the resonant peak and is given by: $V_i = 40$ V, $D = 0.1$, and $R = 450$ Ω .

An additional potential future improvement is briefly discussed next. As explained before, the small-signal model is a linear approximation of the non-linear large-signal model around the chosen operating point. As such, it is only accurate for small deviations around the operating point. Once a controller with the desired characteristics has been determined at a particular operating point, its operation should therefore be validated at other points in the operating range to ensure satisfactory performance across the entire range of operation. Due to the non-linear nature of the converter, the controller is unlikely to work optimally at all operating points using the same parameters. To improve the controller's performance across the operating range, it is possible to use a technique known as *gain scheduling*. In this approach, the small-signal model is evaluated at various operating points, and for each one, the controller parameters are adjusted, or *tuned*, to achieve the desired performance. The parameter values can then be stored and applied dynamically to the controller based on the current operating point.

Whilst some of the preliminary work required in order to implement the dual-loop control of both input current and output voltage was completed, unfortunately the implementation was not finished in time. The practical implementation of the current controller is discussed in more detail in Section 6.3.2.

6.3 Real-time control and data acquisition system

The implementation of both the power converter controller and the data acquisition is centred around the National Instruments CompactRIO 9035 (cRIO). The CompactRIO is an embedded controller featuring both a dual-core Intel Atom processor running a real-time operating system, as well as a Xilinx Kintex-7 field-programmable gate array (FPGA) [64]. Both the real-time processor and the FPGA can be programmed in the LabVIEW graphical programming environment. As shown in Figure 6.16, the controller is housed in a chassis with eight slots that accept a variety of analogue and digital input/output (I/O) modules, known as C-series modules.

The CompactRIO serves a number of functions in the power converter-based EIS system: firstly, it is used to generate the measurement perturbation signal that serves as the reference waveform for the current controller; secondly, it performs the real-time control of the power converter; thirdly, it handles the acquisition of the battery current and voltage signals; and finally, it stores the digitised signals to disk for offline post-processing. In addition to the above, a control panel created in LabVIEW is used to allow: 1) the configuration of the current controller, i.e.



Figure 6.16: National Instruments CompactRIO 9035 embedded controller.

set values and gain terms; and 2) real-time monitoring of the current and voltage signals being acquired.

It is important to note that performing both signal generation and signal acquisition functions on the same system using the same clock source has the important benefit of avoiding any spectral leakage in the acquired signal due to timing errors, e.g. clock jitter, as such errors will equally affect both the generated and acquired signals. Here, that shared clock source is the CompactRIO's 40 MHz clock.

In the following sections the various tasks carried out by the CompactRIO are described in more detail.

6.3.1 Signal generation

This section describes how the current controller reference waveform, i.e. the measurement perturbation, is generated by the CompactRIO.

As stated before, a multisine, i.e. the sum of multiple sine waves, is used as the measurement perturbation in this work. The purpose of the perturbation is to excite the battery in a suitable manner to allow a useful and meaningful EIS measurement. The amplitude, frequency, and phase of the multisine's components must be considered.

The perturbation amplitude should be within a range where it is low enough to evoke a quasi-linear response from the battery and high enough to give a good SNR. For the SBP, an amplitude of 5 mA was determined empirically to meet these requirements. As the modules in the LBP consist of 28 parallel-connected cells (whose impedance is expected to be very closely matched to each other), the perturbation amplitude used with the LBP was increased by a factor 28 relative to that used with the SBP, i.e. 140 mA.

The frequency content of the measurement perturbation should match the impedance spectrum bandwidth requirements. One of the objectives of this work was

to achieve a wide measurement bandwidth, the frequencies of the multisine components have thus been chosen accordingly: 0.01, 0.1, 1, 10, and 100 Hz. This choice of values also meets the requirement explained in Section 5.3.2 that all frequencies in the measurement perturbation should be integer multiples of the lowest frequency component in order to avoid spectral leakage when the observation window length is set to one cycle of the lowest frequency component.

Finally, the phase angle of the multisine components should be considered. If all components are given the same phase angle, the peaks and troughs of the sinusoids will line up and thus increase the crest factor of the resulting waveform. Large peak-to-peak amplitudes can lead to non-linear responses and should thus be avoided. In order to reduce the crest factor of multisine waveforms, it is possible to randomise the components' phase angles. In this work, phase randomisation has not been implemented, however, for the following reasons: the number of components is limited and their amplitude is kept relatively low; also, due to the large difference in frequency between the components, phase randomisation is less effective at reducing the waveform's crest factor.

The multisine reference waveform is generated on the FPGA from 10,000 samples covering one cycle of the $\sin(x)$ function. Each sample is stored as a 16-bit fixed-point number, i.e. 20 kB of nonvolatile memory are required for this purpose.

The signal is generated at a rate of 100 kSa/s. The individual sinusoids that make up the multisine are generated by reading the samples from memory according to their frequency: reading each value once (at the stated rate) gives a frequency of 10 Hz; reading fewer values increases the frequency, e.g. for 100 Hz every tenth value is read; reading values repeatedly, on the other hand, decreases the frequency, e.g. for 1 Hz every value is read ten times. The components are then added together and the resulting multisine waveform is scaled by the set ac amplitude and added to the set dc current. The result forms the reference current, I_{ref} , that is used as part of the closed-loop current control.

6.3.2 Real-time current control

The closed-loop current control is implemented using a proportional and integral (PI) controller. The control loop runs on the CompactRIO's FPGA at the same rate as the desired converter switching frequency, i.e. 100 kHz.

The reference current, I_{ref} , and the battery current, I , are compared to each other, giving the error signal, e_I :

$$e_I = I_{\text{ref}} - I \quad (6.21)$$

The error signal is then fed into the PI controller, where it goes through the proportional and integral branches to produce the duty cycle signal, d :

$$d = K_P \cdot e_I + K_I \int e_I dt \quad (6.22)$$

where K_P and K_I are the proportional and integral gain terms, respectively.

Suitable values for K_P and K_I were determined empirically. As the controller showed satisfactory performance using only the integral term, the proportional branch was left unused, i.e. $K_P = 0$. The value of the integral gain term depended on which battery pack was in use, due to their differing impedances and the associated difference in the measurement perturbation amplitude used with each pack. In testing, the values listed below provided satisfactory performance in terms of stability, steady-state error, and dynamic response.

$$K_{I,\text{SBP}} = 100 \quad (6.23)$$

$$K_{I,\text{LBP}} = 10 \quad (6.24)$$

The output of the controller — the duty cycle — is then utilised to define the two waveforms that are used to switch the converter. The two switching signals, which, as explained in 6.2.1, are 180° out-of-phase with each other, are generated using a NI 9401 digital I/O module. This is connected to the CompactRIO via one of the eight slots in its chassis. In order to protect the dc-dc converter, the duty cycle is software-limited to 95 %.

6.3.3 Signal acquisition

As described in Section 5.2, the 16S1P battery pack connection topology used in this work requires the acquisition of 17 signals (16 voltages and 1 current) to allow the impedance determination of the batteries in the pack. Once the signals have undergone the required conditioning steps — described in Section 6.4 — they are ready to be acquired. The signal acquisition is accomplished using a number of C-series analogue input modules that connect to the CompactRIO chassis.

The 16 battery voltages are acquired using a NI 9205 analogue input module. Its key specifications are summarised in Table 6.4.

The input channels are set up to acquire single-ended signals, with an input voltage range of ± 5 V, and a sampling rate of 200 kSa/s. The NI 9205 is a scanned module, i.e. its input channels are multiplexed to a single ADC; an aggregate sampling rate of 200 kSa/s therefore corresponds to a channel rate of 12.5 kSa/s. By

Table 6.4: Specifications of NI 9205 analogue input module [65].

Input voltage range	± 0.2 V to ± 10 V (programmable)
Resolution	16 bits
Maximum sampling rate	250 kSa/s (aggregate)
Number of input channels	16 differential / 32 single-ended
ADC type	Successive Approximation Register (SAR)

choosing an integer divisor of 100 kHz for the channel sampling rate, any ripple from the converter — which will be at 100 kHz — will be aliased as dc.

A consequence of using a multiplexed ADC to acquire the voltages is that the input channels will not be sampled at exactly the same time. This time delay will affect the phase of the calculated impedance (since the current will not experience the same delay), and must therefore be considered. Here, the highest measurement frequency is 100 Hz; assuming a worst-case maximum delay of 80 μ s (1/12500), this would correspond to a maximum phase angle error of less than 3°. Although it would be possible to correct for this error by characterising and compensating for the delay between channels, this compensation has not been implemented, as the error is systematic and therefore does not influence the system's ability to accurately detect *variations* in the impedance.

The battery current signal is acquired using a NI 9215 analogue input module. The module's key specifications are summarised in Table 6.5.

Table 6.5: Specifications of NI 9215 analogue input module [66].

Input voltage range	± 10 V
Resolution	16 bits
Maximum sampling rate	100 kSa/s/ch
Number of input channels	4 (differential)
ADC type	Successive Approximation Register (SAR)

The current is also sampled at a rate of 12.5 kSa/s.

Once the conversion of the 17 channels is complete, the FPGA writes the samples to a direct memory access (DMA) first-in-first-out (FIFO) memory buffer. When the buffer contains 1700 elements, i.e. 100 samples of each channel, the data is transferred to the real-time processor. From there it is both written to file and displayed to a control panel for monitoring purposes.

As well as the 17 channels for the impedance determination, the CompactRIO

must acquire additional analogue inputs in order to enable the closed-loop control. Here, that means acquiring the battery current. Although the battery current was already acquired above for the purposes of the impedance determination, it must be acquired again using a separate module. This is due to two reasons: firstly, in order to enable the closed-loop control, the current must be sampled at the same rate as the control loop rate, i.e. 100 kHz; and secondly, there is a constraint in the programming environment which prevents I/O channels of the same module from being used in more than one process. A second NI 9215 module is thus used to sample the battery current at a rate of 100 kSa/s. The acquired signal, I , is used in the negative feedback loop of the current controller.

6.3.4 Data storage

The battery current and voltage signals acquired at a rate of 12.5 kSa/s are written to the CompactRIO's nonvolatile memory in a binary file format known as Technical Data Management Streaming (TDMS). Once the acquisition is complete, the files are transferred to a personal computer for post-processing. The signal processing steps — described in Section 5.3 — are carried out in MATLAB. Although falling out of the scope of this work, the processing of the raw data could of course also take place on the CompactRIO. This could provide battery impedance information in real-time, whilst also enabling the volume of data that requires storing to be reduced, as there may be no need to store the raw voltage and current signals.

6.4 Signal conditioning


Before the battery current and voltage signals can be acquired, they must undergo a series of steps in preparation. This section describes those steps in detail.

6.4.1 Current

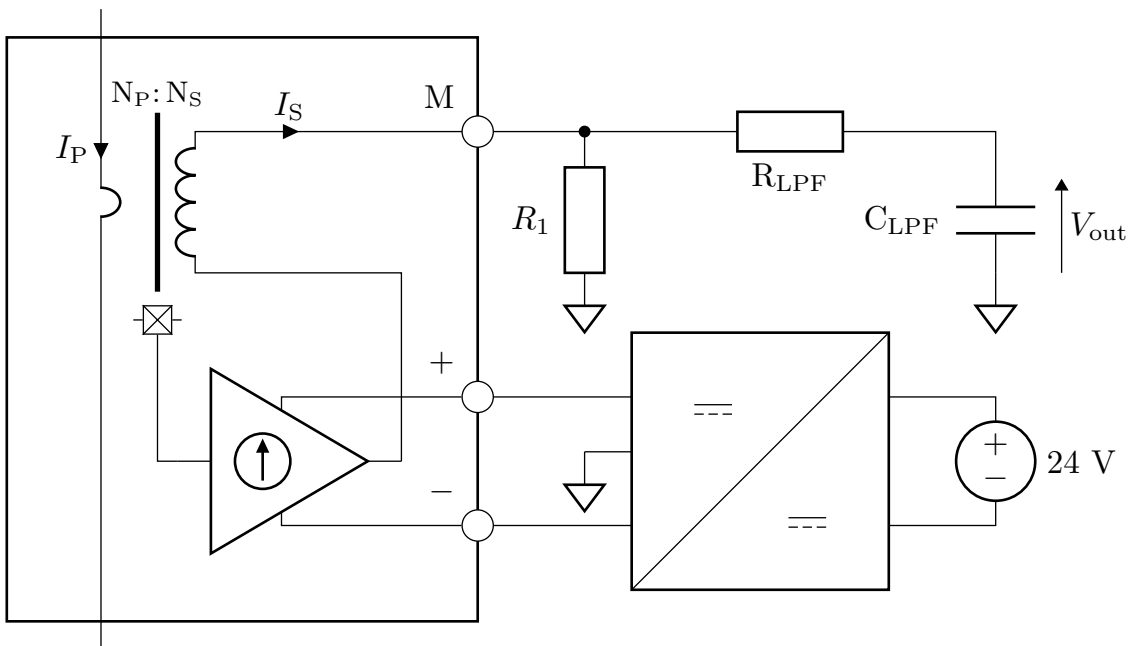
The current measurement is performed using a LEM LA 25-P closed-loop Hall effect transducer [67].

A separate PCB was designed to accommodate the current transducer. The PCB and its schematic diagram are shown in Figure 6.17.

To generate the ± 15 V split power supply required by the active circuitry of the current transducer, the TEN-5-2423 dc-dc converter [68] is used (the +24 V input coming from an auxiliary supply).



(a) Approximate dimensions (L×W): 96×88 mm.



(b) Schematic diagram.

Figure 6.17: The current measurement board.

The conductor carrying the current to be measured is passed through a gap in the plastic case of the LA 25-P; this forms the primary winding around the magnetic core. The current in the secondary winding (located internally to the transducer), I_S , is related to that in the primary winding, I_P , by the turns ratio, i.e. $\frac{I_S}{I_P} = \frac{N_P}{N_S}$, where N_P and N_S are the number of turns in the primary and secondary windings, respectively. The LA 25-P is specified as having a secondary winding with one thousand turns, thus $\frac{I_S}{I_P} = \frac{N_P}{1000}$. A resistor, R_1 , converts the secondary winding current into a voltage. The final step before sampling the current transducer output is to remove any high-frequency noise by passing it through a single-pole RC low-pass filter with a cut-off frequency of roughly 10 kHz. The filter's frequency response allows measurement components up to about 1 kHz to be used before the phase starts to become affected. The following analysis ignores the filter since it does not influence the measurement perturbation components.

We can obtain an expression relating the current and number of turns in the primary winding to the output voltage, by using Ohm's law in combination with the turns ratio:

$$V_{\text{out}} = \frac{N_P \cdot R_1}{1000} \cdot I_P \quad (6.25)$$

To understand what values to choose for N_P and R_1 , it is necessary to consider the specifications of the current transducer, the ADC, and the batteries used in each pack.

The maximum value of N_P is limited by the number of turns that can physically fit through the gap in the current transducer's body, i.e. it depends on the cross-sectional area (CSA) of the cable used. Two different cables are available: the first, rated for 48 A, has a CSA of 6 mm², which allows for only one primary winding; the second, rated for 30 A, has a CSA of 2.5 mm², allowing $N_P \leq 3$.

The output voltage, V_{out} , is measured using the NI 9215 analogue input module. To maximise the resolution of the current measurement, it is important to use the ADC's full input voltage range. In other words, the range of currents to be measured must be mapped to the input voltage range. This means that ideally the maximum current to be measured, $I_{P(\text{max})}$, should lead to $V_{\text{out}} = V_{i,\text{ADC}(\text{max})}$, where $V_{i,\text{ADC}(\text{max})}$ is the ADC's maximum input voltage², i.e. 10 V. To find the value of R_1 that will achieve this, we can rearrange (6.25):

$$R_1 = \frac{V_{i,\text{ADC}(\text{max})}}{I_{P(\text{max})}} \cdot \frac{1000}{N_P} = \frac{V_{i,\text{ADC}(\text{max})}}{I_{S(\text{max})}} \quad (6.26)$$

²The range of currents to be measured is assumed to be symmetrical around 0 A, i.e. $-I_{P(\text{max})} \leq I_P \leq I_{P(\text{max})}$, since the batteries will be both charged and discharged.

There is an additional constraint on the value of R_1 , however. This is due to the current transducer's output stage, whose maximum output voltage, $V_{A(\max)}$, is implied to be approximately 12 V [67]. Part of the transducer's output voltage drops across the secondary winding resistance, R_S , thus reducing the voltage available across R_1 . Since $V_{i,ADC(\max)} = 10$ V, any voltage drop across R_S greater than 2 V will cause a reduction in the achievable measurement resolution. Assuming the worst-case scenario, i.e. the highest rated ambient operating temperature, the maximum value of R_S is given as 85 Ω . The maximum allowable current in the secondary winding that will still allow the maximum resolution to be achieved is therefore found to be:

$$I_{S(\max)} = \frac{V_{A(\max)} - V_{i,ADC(\max)}}{R_{S(\max)}} = \frac{12 - 10}{85} \approx 23.5 \text{ mA} \quad (6.27)$$

This corresponds to maximum primary currents of 23.5 A, 11.75 A, and 7.83 A, when one, two, or three turns are used in the primary winding. The cells used in the SBP have a recommended continuous discharge current of 3.2 A or less; their operating current range is therefore unaffected by this limitation. The modules used in the LBP, on the other hand, have a recommended continuous discharge current of ≤ 90 A. In the case of the LBP, there is therefore a trade-off between higher measurement resolution and maximum measurable current. The decision is made to opt for higher measurement resolution, based on a number of reasons: a current of 23.5 A is deemed high enough for the purposes of this work, as it still enables the modules to be discharged at a moderate rate; it also allows the use of available cables and connectors with current ratings of 30 A and 32 A, respectively.

Another aspect that has not been considered thus far, is that a single current measurement board is required to be used with both battery packs. To ensure the ease of use of the board, the value of R_1 should not change, i.e. R_1 must be chosen so that both current ranges can be measured with adequate resolution. For this reason, the board was designed to offer two separate configurations: one for use with the LBP that can measure up to 23.5 A using a single turn in the primary winding, and the other one for use with the SBP using three turns in the primary winding and with a maximum measurable current of 7.83 A.

Substituting $I_{S(\max)}$ from (6.27) into (6.26) gives the required value of R_1 as 425 Ω . The next preferred value from the E12 series (10 % tolerance) is chosen as 390 Ω . The resistor value is verified using a digital multimeter with a resolution of 0.1 Ω and a rated accuracy of $\pm(0.5\% + 3)$ resulting in the following value for R_1 :

$$R_1 = 426 \pm 1 \Omega \quad (6.28)$$

where the standard uncertainty is based on the stated accuracy specification and the assumption of a uniform error distribution.

The value of R_1 given in (6.28) is used in subsequent data processing steps, in which the current, I_P , is derived from the measured voltage, V_{out} , using the relationship defined in (6.25).

The selected values yield a measurement resolution of approximately 717 μA for the LBP configuration ($N_P = 1$), and 239 μA for the SBP configuration ($N_P = 3$). To put this into context, the amplitude of the measurement perturbations are 140 mA and 5 mA for the LBP and SBP, respectively, showing that the resolution is high enough for the intended measurement.

6.4.1.1 Compensation for offset current

When there is no current flowing in the primary winding of the LA-25P current transducer, a small voltage is still measurable at its output due to the transducer's offset current. In the frequency domain this offset appears only in the dc component and therefore is of no consequence to the accuracy of the impedance determination. If the dc component of the battery current is to be controlled, however, then the offset current should be compensated for in order to improve the accuracy of the control. This can be achieved by obtaining the offset value during periods where there is known to be no current in the primary winding and then simply subtracting it from the measured value.

It should be noted that this step was neglected during the implementation of the controller, mainly because the offset current is relatively low³ and so went largely unnoticed.

6.4.2 Voltage

The battery packs consist of 16 series-connected cells/modules with an average operating voltage of 3.2 V each. As explained in Chapter 5, the common-mode voltage for the cell/module at the high-potential side of the battery pack will thus be almost 50 V. This is too high for most ADCs and is indeed too high for the NI 9205 analogue input module that has been chosen to acquire the voltage signals — its maximum input voltage (sum of signal and common-mode voltage) being 10.4 V [65]. The cell/module voltages are thus obtained differentially using isolation amplifiers.

³Its maximum value is given as 0.2 mA, but in practice it was commonly found to be around half of that.

The cell used in the battery pack [32] has recommended discharge and charge cut-off voltages of 2.50 V and 3.65 V, respectively. This is the range of voltages to be measured.

The input voltage range of the NI 9205 is set to ± 5 V. To maximise the effective resolution of the measured voltage, the range of the signal to be acquired and the ADC input range should overlap as much as possible. This requires a transformation of the cell voltage: the minimum cell voltage of 2.5 V should be mapped to the ADC's minimum input voltage of -5 V, and likewise, the maximum cell voltage of 3.65 V should be mapped to the ADC's maximum input voltage of $+5$ V. This transformation can be achieved in two steps: first the cell voltage range is centered around zero and then it is multiplied by the required gain factor. This is illustrated in Figure 6.18.




Figure 6.18: Mapping of cell voltage range to ADC input voltage range.

To center the cell voltage range around zero, subtract from it the value that lies at the middle of the range, $V_{\text{cell}(\text{mid})}$:

$$V_{\text{cell}(\text{mid})} = \frac{3.65 + 2.5}{2} = 3.075 \text{ V} \quad (6.29)$$

The required gain factor, A_v , is given by the ratio of the ADC input voltage range, $\Delta V_{i,\text{ADC}}$, and the cell voltage range, ΔV_{cell} :

$$A_v = \frac{\Delta V_{i,\text{ADC}}}{\Delta V_{\text{cell}}} = \frac{5 - (-5)}{3.65 - 2.5} = \frac{10}{1.15} \approx 8.7 \quad (6.30)$$

The transformation from the measured cell voltage, V_{cell} , to the ADC input voltage, $V_{i,\text{ADC}}$, can thus be summarised as follows:

$$V_{i,\text{ADC}} = A_v [V_{\text{cell}} - V_{\text{cell}(\text{mid})}] \quad (6.31)$$

A signal conditioning circuit is designed to implement (6.31) in hardware. The circuit is based around the AD204 — a transformer-coupled isolation amplifier [69]. As well as having a $\pm 5\text{ V}$ output range that matches $\Delta V_{i,\text{ADC}}$, the AD204 comes with two additional features that make it particularly useful in this application: 1) an uncommitted operational amplifier at its input, and 2) an isolated $\pm 7.5\text{ V}$ voltage source. The function described by (6.31) can be implemented by using the AD204's op amp in a summing amplifier configuration with two inputs: the cell voltage, V_{cell} , and the middle value of the cell voltage range — only negated, i.e. $-V_{\text{cell}(\text{mid})}$. The latter can be generated using the AD204's isolated voltage source. The schematic of the signal conditioning circuit is shown in Figure 6.19 and describes how this was accomplished in practice.

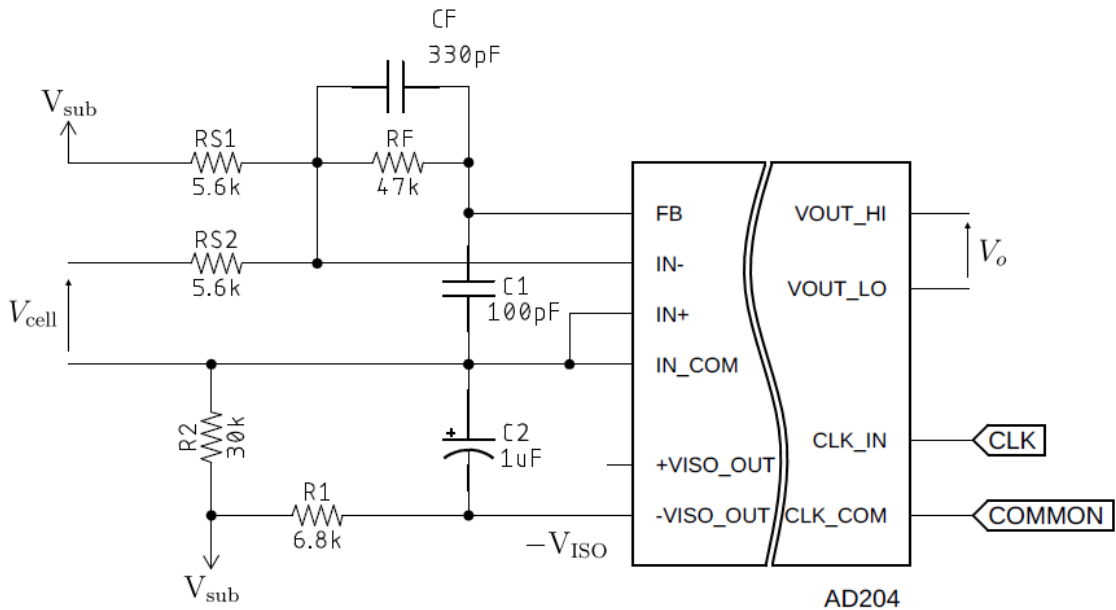


Figure 6.19: Schematic of signal conditioning circuit used in voltage acquisition.

As shown in Figure 6.19, V_{sub} and V_{cell} are the two inputs to the summing amplifier. V_{sub} is generated from $-V_{\text{ISO}}$ (-7.5 V relative to IN_COM) via the potential divider consisting of R_1 and R_2 . The potential divider output is loaded down by the effective resistance seen looking into the summing amplifier input, i.e. R_{S1} ; V_{sub} can therefore be determined as follows:

$$V_{\text{sub}} = -V_{\text{ISO}} \cdot \frac{R_2}{R_1 + R_2} \cdot \frac{R_{\text{S1}}}{R_{\text{S1}} + (R_1 \parallel R_2)} \approx -3.073\text{ V} \quad (6.32)$$

Resistor tolerances and the accuracy of the isolated supply are additional sources of uncertainty that can cause the absolute value of V_{sub} to further deviate from $V_{\text{cell}(\text{mid})}$; this is dealt with in a calibration procedure described in Section 6.4.2.1.

As both input resistors, R_{S1} and R_{S2} , are given the same value, R_S , the relationship between the differential cell voltage, V_{cell} , and the AD204 output voltage, V_o , is described by a well-known result from the analysis of inverting summing amplifiers:

$$V_o = -\frac{R_F}{R_S} (V_{\text{cell}} + V_{\text{sub}}) \quad (6.33)$$

Comparing (6.31) and (6.33), we see that:

$$A_v = -\frac{R_F}{R_S} \quad (6.34)$$

Based on the values chosen for the gain setting resistors (shown in Figure 6.19), it can be seen that the gain has been set to approximately -8.4 . The minus sign reflects the fact that the *inverting* summing amplifier configuration is used. This means that the maximum cell voltage of 3.65 V is mapped to -5 V, and the minimum cell voltage of 2.5 V is mapped to $+5$ V, rather than the other way around.

Using this nominal gain figure (i.e. ignoring resistor tolerance), the resolution at which the cell voltage is measured can be calculated as approximately 18 μV . The voltage perturbations to be measured lie in the region of 0.5 mV; the stated measurement resolution is thus considered to be high enough.

As with the current signal, the voltage is also low-pass filtered with a first-order RC filter with a corner frequency of around 10 kHz.

In order to measure the voltages of all the batteries in the pack, the circuit of Figure 6.19 is repeated sixteen times. The resulting 16-channel voltage conditioning circuit was built on a double-sided PCB — this is shown in Figure 6.20.

The voltage conditioning board is powered from an auxiliary 24 V supply. The TEN5-2413 dc-dc converter [68] is used to convert this into the 15 V supply required by the AD246 — the dedicated power supply/clock driver used in conjunction with the AD204 isolation amplifier [70].

The input connector is a DB-25 D-sub connector; this matches the output connector on the smaller battery pack's PCB (shown in Figure 6.2), as well as the connector soldered onto the shielded multicore cable, which carries the voltage sense signals from the larger battery pack (shown in Figure 6.1).

On the output a DB-37 D-sub connector is used to connect to the input of the NI 9205 ADC module.




Figure 6.20: 16-channel voltage conditioning PCB (top layer).

6.4.2.1 Calibration

From (6.33) it is known that a plot of the output voltage of the signal conditioning circuit, V_o , against the cell voltage, V_{cell} , should result in a straight line with a slope of $-R_F/R_S$ and a root at $V_{\text{cell}} = -V_{\text{sub}}$. However, both the resistor tolerance and the accuracy of the AD204's isolated supply (of which V_{sub} is a function — see (6.32)) will cause the slope and the root to diverge from their nominal values, i.e. they introduce errors. If these errors are not taken into account, they will propagate into the cell voltage estimate and thus the impedance estimate. Since these are systematic errors, i.e. they are constant for a given set of components, a calibration procedure can be used to compensate for them.

The basic concept behind the calibration procedure is to measure the output of the signal conditioning circuit to a series of known inputs⁴, and from the resulting plot of output voltage against input voltage, estimate the value of the slope and the root. These values can then be stored in software and be used to determine V_{cell} from the acquired V_o using the rearranged form of (6.33):

$$V_{\text{cell}} = -\frac{R_S}{R_F}V_o - V_{\text{sub}} \quad (6.35)$$

⁴In theory, as we are dealing with a straight line relationship, only two points are required.

The calibration procedure was carried out by connecting a variable voltage source to the input of the signal conditioning circuit in place of V_{cell} . Its voltage was swept through the range of V_{cell} , i.e. $[2.5, 3.65]$ V, and the output of the signal conditioning circuit, V_o , was acquired by the NI 9205 and written to disk. This was repeated for all 16 channels and the data was then loaded into MATLAB for processing. As an example, Figure 6.21 shows the relationship between V_o and the voltage from the variable voltage source, V_i , for channel 1. Although the exact response differed slightly between channels, the general behaviour was the same across all channels.




Figure 6.21: Output voltage of signal conditioning circuit during calibration (channel 1).

Two things are apparent from Figure 6.21: 1) the relationship between the input and output of the signal conditioning circuit is not exactly a straight line — instead, it appears to consist of *two* straight lines that share a common root; and 2) the range of V_o is not symmetrical around 0 V — rather than being the desired $[-5, +5]$ V, it is closer to $[-6, +4]$ V. Whilst the second issue could easily be caused by the $\pm 10\%$ accuracy of V_{ISO} , an unfavourable combination of R_1 , R_2 and R_{S1} , or indeed a mixture of the two, the first issue is not as easily explained. There appear to be two different gain values at play — evidenced by the differing slopes of the two lines. This appears to be linked to the sign of the output voltage: the range of inputs that leads to negative output voltages is consistently (across all channels) amplified by a larger gain factor than the range of inputs resulting in positive output voltages. Through calibration this gain error is compensated, leaving only a small residual error due to the imperfect estimation of the calibration coefficients.

In order to achieve accurate calibration, the two straight lines were characterised by determining their respective slopes and their common root, i.e. the value of V_i for which $V_o = 0$. This process was repeated for all sixteen channels on the signal conditioning board resulting in 48 coefficients (32 slopes and 16 roots) that are used in conjunction with (6.35) to reverse the acquisition related conversions.

Chapter 7

Results

7.1 Verification of system operation

In order to verify the basic operation of the system — including: control, signal acquisition, and signal processing stages — the SBP (i.e. single cells) was connected to the input of the dual-phase interleaved boost converter (IBC) and a $220\ \Omega$ resistive load was connected to the converter's output. The controller reference was configured for a 1 A dc current and a multisine measurement perturbation with components at 0.01, 0.1, 1, 10, and 100 Hz, each with an amplitude of 5 mA. This resulted in the converter operating at a steady-state duty cycle of around 32 %. At the time of the experiment the SoC of the cells was approximately¹ 30 % and the cells had been rested for about one day prior to the beginning of the experiment. Unless stated otherwise, all results can be assumed to have been obtained at room temperature (approximately 20 °C).

Figure 7.1 shows the raw voltage and current signals covering the length of five cycles of the lowest frequency component — both before and after the reversal of acquisition-related conversions. Prior to the data acquisition, the measurement perturbation had been running for several minutes in order to satisfy the ac steady state requirement. All 16 cell voltages were acquired; however, for clarity, only one (cell 5) is shown in Figure 7.1.

It should be noted that whilst many of the results presented in this section (and those that follow) are for a single, randomly chosen cell in the battery pack, very similar results were obtained for all of the batteries.

From Figure 7.1b it can be noted that the dc component of the current is just below the set value of 1 A. The difference is mainly caused by the lack of compensation for the current transducer's offset current in the controller logic, which was

¹Coulomb counting from full charge was used as the SoC estimation method.

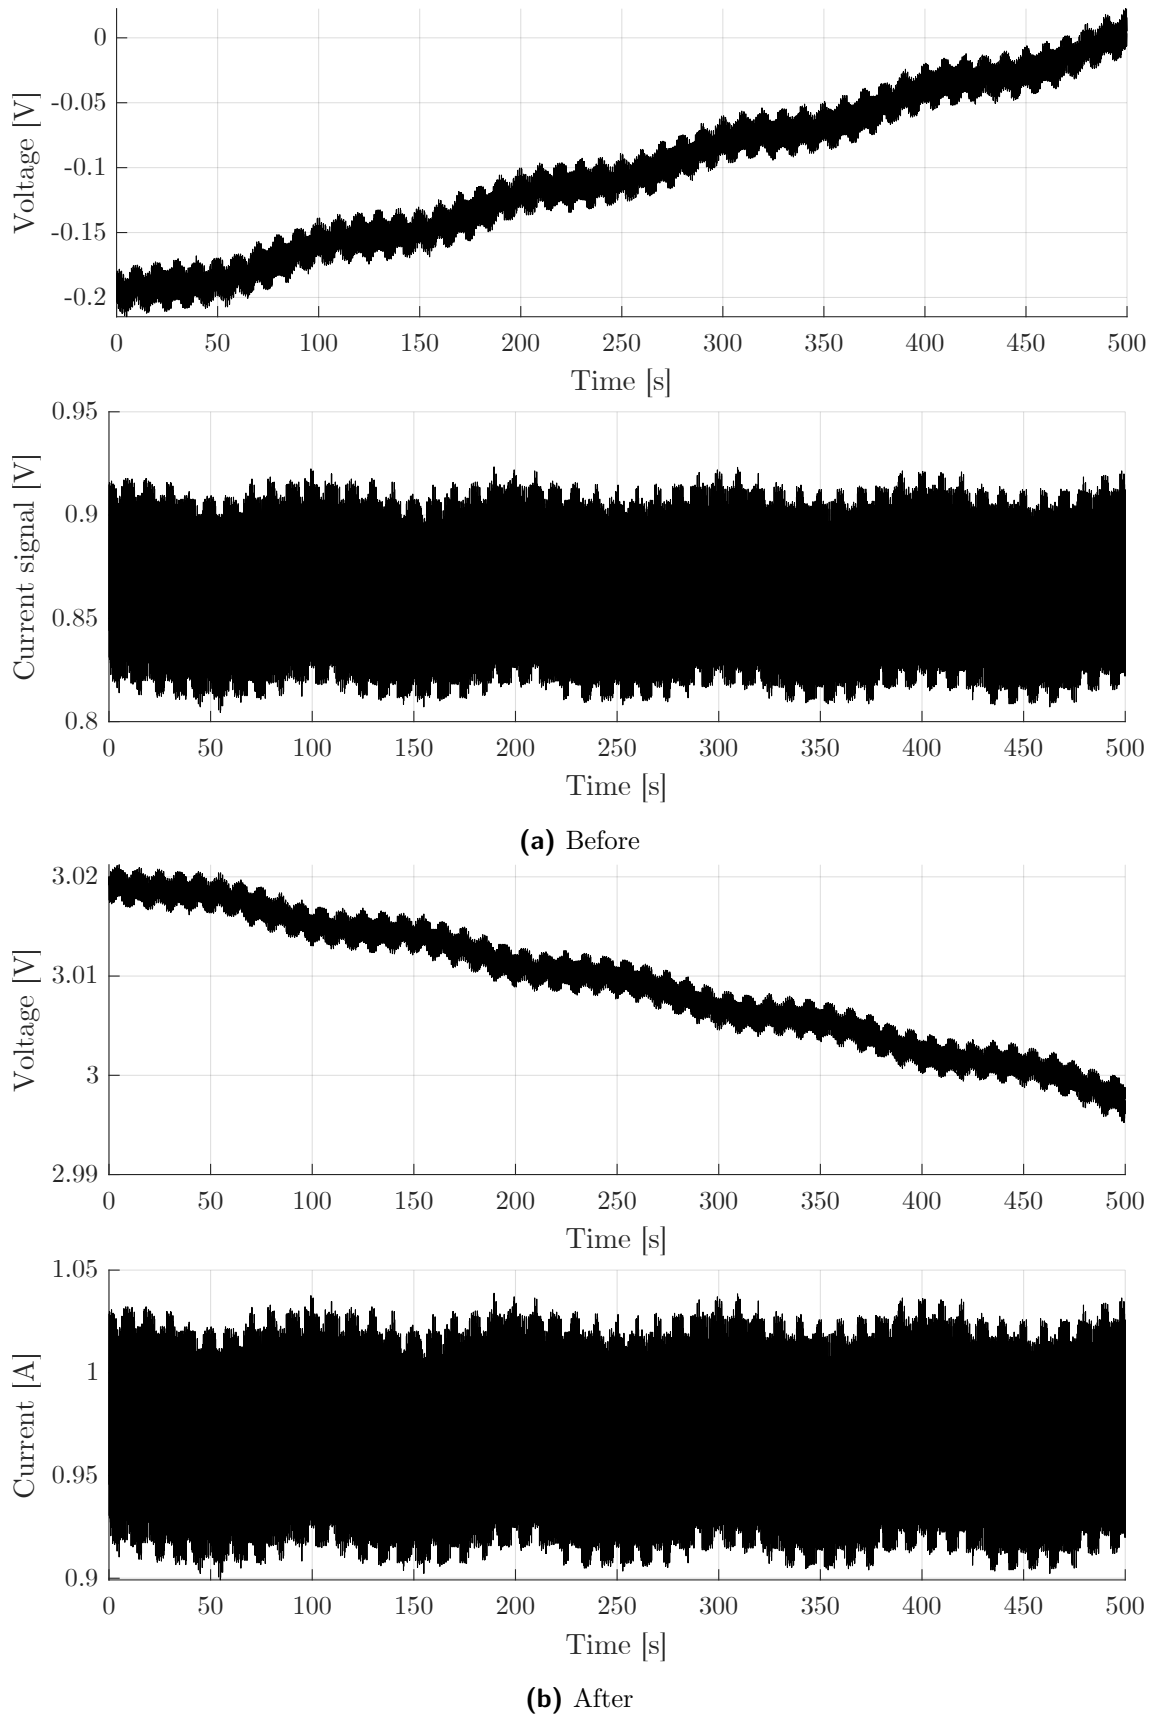


Figure 7.1: Acquired current and voltage signals before and after reversal of acquisition-related conversions.

measured as approximately 45 mA in this experiment.

The next step in the signal processing chain is the signal length adjustment of both current and voltage to one cycle of the lowest measurement perturbation frequency, i.e. 100 s. Figure 7.2 shows the corresponding waveforms.

Figure 7.2 shows that the shape of the current waveform does not exactly match the expected (ideal) shape of a multisine of similar characteristics (cf. Figure 5.5); however, the frequency domain representation of the current — shown in Figure 7.3 — reveals that the controller is operating correctly with all five measurement perturbation components appearing at the expected frequencies. With the exception of the 100 Hz component, whose magnitude is slightly attenuated (-1.3 dB), all frequency components also have the expected (set) amplitude.

From 7.1 and Figure 7.2 it is clear that the voltage signal contains a drift. As previously explained, this is caused by the dc component of the current and the associated change in SoC during the measurement. Before the FFT of the voltage is obtained, the drift in the voltage is compensated using the RPS method. Figure 7.4 shows the uncompensated signal, the ramp estimated by the RPS fit, and the compensated signal in the time domain.




Figure 7.2: Current (top) and voltage (bottom) signals after signal length adjustment.

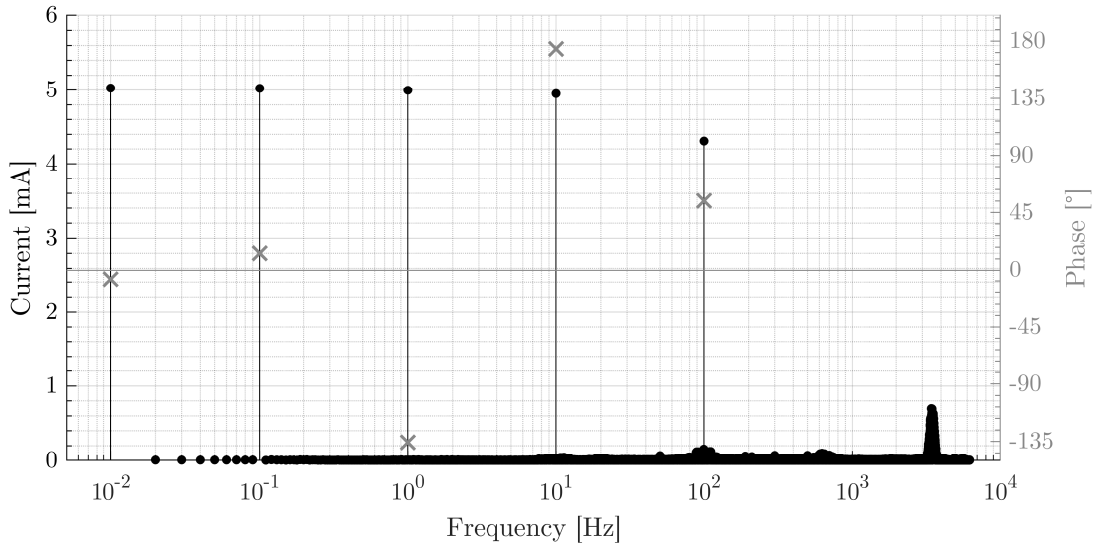


Figure 7.3: Current spectrum. For clarity the phase plot only includes those frequency components that are part of the measurement perturbation.

Figure 7.5 shows the magnitude and phase spectra of the compensated and uncompensated signals. It can be seen that the uncompensated voltage contains considerable distortion, particularly at the low-frequency end; the phase of the low-frequency components is -90° , confirming that the straight line drift approximation can be relatively accurate under the condition of a constant current discharge.

By compensating for the drift in the voltage, magnitude estimation errors of approximately 135 % and 8 % are avoided at 0.01 and 0.1 Hz, respectively. The respective phase errors that are avoided are 81° and 14° .




Figure 7.4: Time domain plot of voltage signal before and after drift compensation using the RPS method.

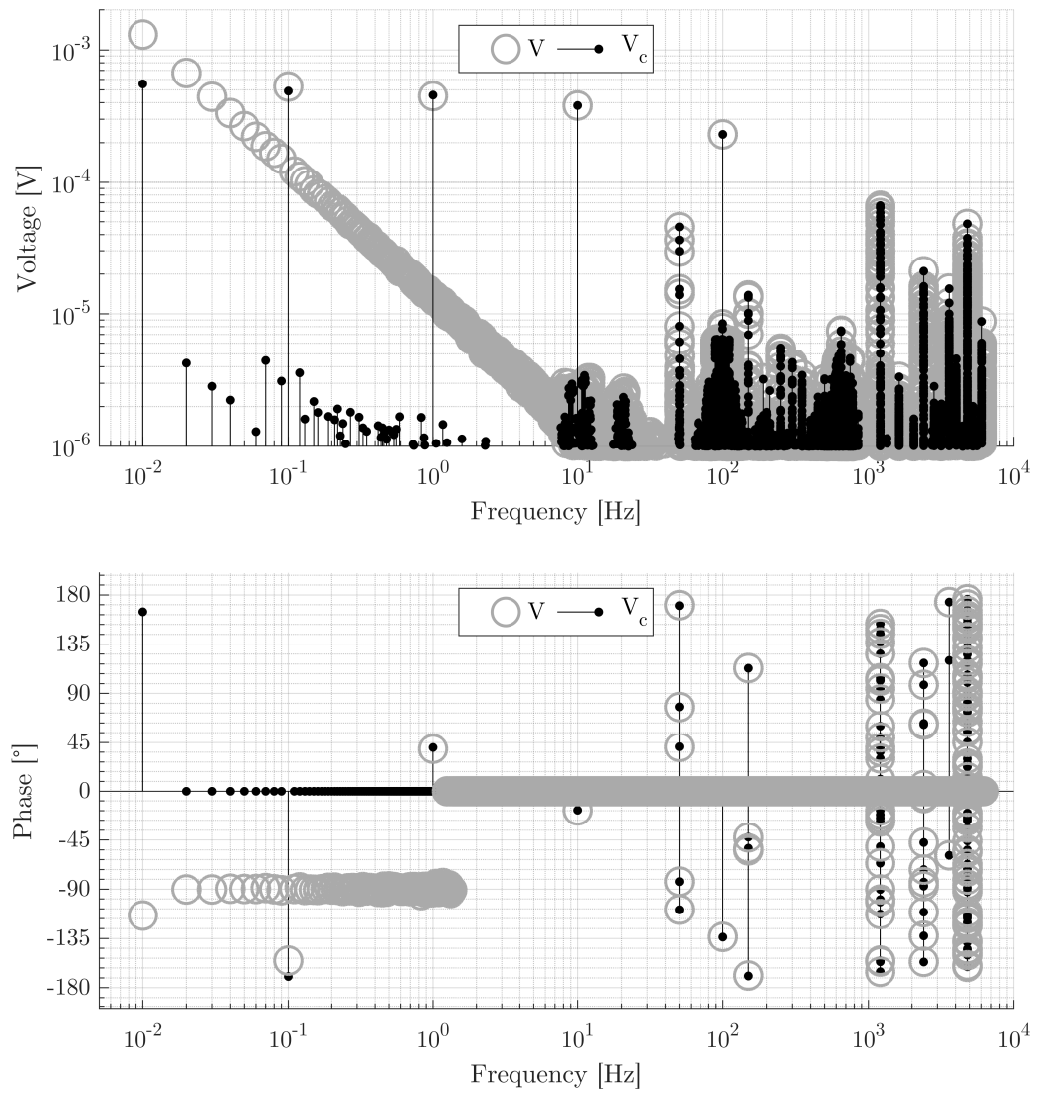


Figure 7.5: The voltage signal in the frequency domain both before and after drift compensation using the RPS method. To improve the clarity of the phase spectrum, only components with magnitude greater than $10 \mu\text{V}$ are shown.




Figure 7.6: Nyquist plots of cell impedance based on compensated and uncompensated voltage signals.

The final step in the signal processing chain is to determine the cell impedance as the ratio of the complex voltage and current signals. Figure 7.6 shows the impedance spectrum obtained from the preceding voltage and current spectra. Plotted alongside the impedance obtained using the compensated voltage are the impedance values obtained using the uncompensated voltage; this further demonstrates the scale of the error introduced by the drift in the voltage, and highlights the importance of using effective drift compensation techniques in online EIS applications.

7.2 Monitoring the high-power battery pack

The impedance of the modules in the LBP was determined using the same multi-sine measurement perturbation employed in the previous section, i.e. with frequency components located at 0.01, 0.1, 1, 10, and 100 Hz. The amplitude of each component was set to 140 mA in order to achieve a similar voltage response as in the test with the SBP, and the dc current was set to 20 A. With a 5 Ω electronic load (Elektro Automatik EL 9200-70 B) connected to the IBC, the converter operated at a steady-state duty cycle of around 30 %. Figure 7.7 shows an overview of the experimental setup used in this section.




Figure 7.7: The experimental setup used to monitor the high-power battery pack.

The SoC of the modules at the time of the experiment was also approximately 30 %. Prior to the beginning of the experiment the modules had been rested for about five days. Similarly to the test involving the SBP, in order to satisfy the ac steady state requirement, the measurement perturbation had been running for several minutes prior to starting the data acquisition.

Figure 7.8 shows the battery pack current in both the time and frequency domains. It can be seen that the controller is operating correctly with the dc component of the current very close to the set value of 20 A (the mean of the current is 19.97 A), and the measurement perturbation components appearing at the expected frequencies and, for the most part, at the expected amplitude; the 10 Hz and 100 Hz components are attenuated by approximately 0.8 dB and 2.4 dB, respectively. This is most likely associated with the bandwidth of the current controller, and the error could thus be reduced through a more robust tuning procedure. It should be noted, however, that the exact amplitude of the perturbation is irrelevant for the purposes of the impedance determination; what matters is that the amplitude is large enough to cause a sufficiently large response in the voltage.

Figure 7.9 shows a time domain plot of both the compensated and uncompensated voltage (of module 5), as well as the ramp estimated through the RPS fit.

Figure 7.10 shows the corresponding frequency domain view (magnitude and phase spectra) of the compensated and uncompensated voltage signals.

Similarly to the results obtained for the SBP in Section 7.1, it can be seen that the uncompensated voltage shows signs of containing a ramp component. By compensating for the drift in the voltage, magnitude estimation errors of 82 % and




Figure 7.8: Time and frequency domain plots of current in high-power acquisition. For clarity the phase plot only includes those frequency components that are part of the measurement perturbation.




Figure 7.9: Time domain plot of voltage in high-power acquisition, both before and after drift compensation using the RPS method.

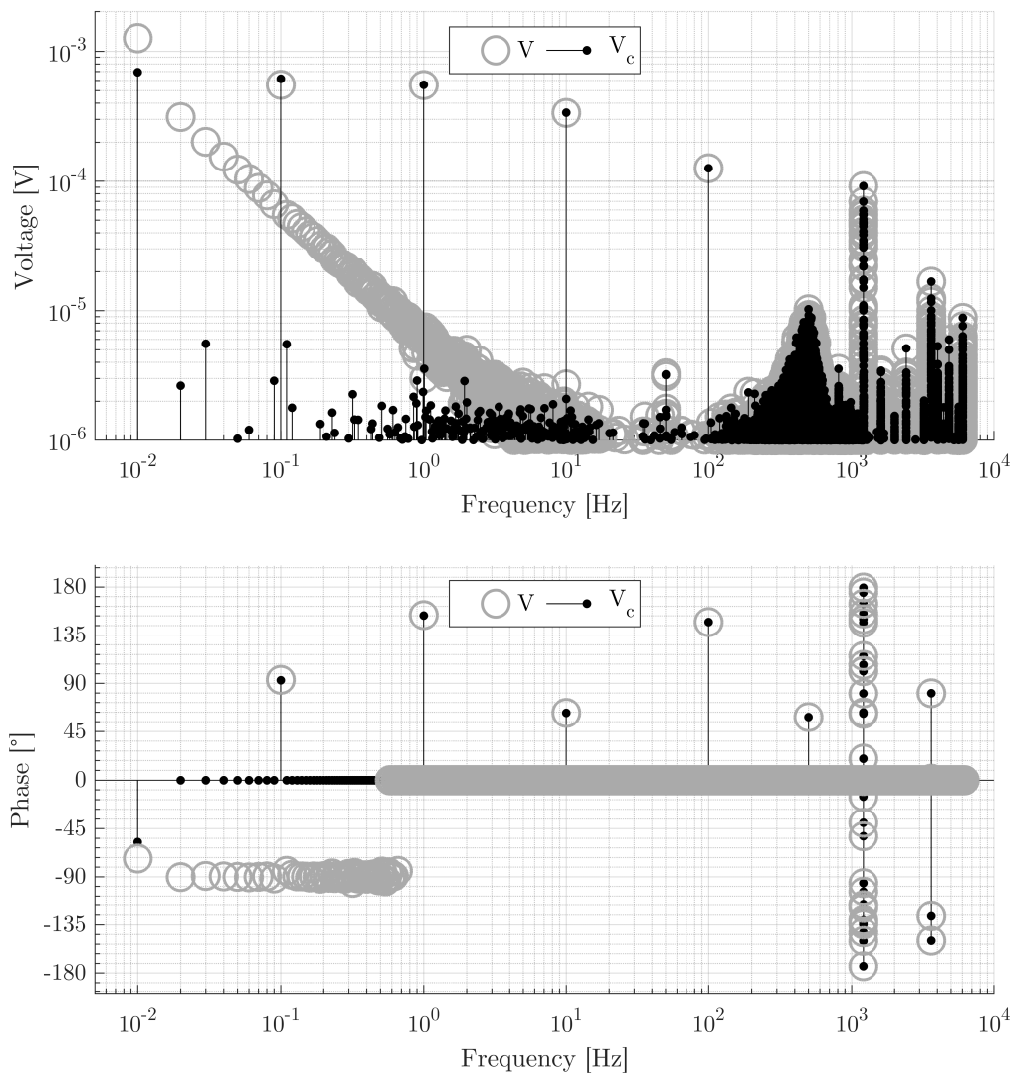


Figure 7.10: Frequency domain plot of voltage in high-power acquisition, both before and after drift compensation using the RPS method. To improve the clarity of the phase spectrum, only components with magnitude greater than $10 \mu\text{V}$ are shown.




Figure 7.11: Nyquist plots of module impedance based on compensated and uncompensated voltage signals.

10 % are avoided at 0.01 and 0.1 Hz, respectively. The phase error avoided at 0.01 Hz is approximately 15° .

From the magnitude spectrum of the voltage it can also be observed that the amplitude of the measurement perturbation components is in the same range as in the results from the SBP. This confirms that the approximation made during the selection of the measurement perturbation amplitude is relatively accurate, i.e. the impedance magnitude of a module consisting of 28 parallel-connected cells is approximately 1/28th of that of a cell.

Figure 7.11 shows the impedance spectra for module 5 obtained using both the compensated and uncompensated voltage signals.

7.3 Measurement repeatability

It is generally desirable for measurements to have high *repeatability* as this reduces the uncertainty associated with the measurement and allows the detection of effects on the measurand that require a high measurement resolution. Repeatability is closely linked to the concept of *precision* and describes the agreement between results obtained from measurements repeated under the same conditions, e.g. same procedure, measurement system, and operating conditions [71].

One way to quantify the repeatability of the power converter-based impedance measurements is to carry out one measurement and then attempt to bring the cell(s) back to the state they were in at the beginning of the measurement. In practice this approach is quite complex, as factors such as cell rest/relaxation time must be

considered; it is also time-intensive, especially if many repetitions are needed. An alternative approach is to acquire a number of measurements in direct succession of each other. This has the advantage of being much simpler, but the downside is that the maximum number of repeated measurements is limited, particularly if very low-frequency perturbations are used, or the dc current is high, as both of these factors will cause the change in SoC between successive measurements to increase and therefore the operating conditions to change. As long as the change in SoC during the length of the acquisition is kept relatively low, however, the state of the cell(s) can be considered as being approximately constant.

The same experimental setup and controller configuration as in Section 7.1 was used to obtain five successive impedance spectra for the cells in the SBP. The lowest measurement perturbation frequency was set to 0.01 Hz and the dc current to 1 A, i.e. the change in SoC during the acquisition was approximately 4.2 %². The cells were rested for approximately one day before the experiment and their SoC was around 40 %.

The spread of the data in the set of five measurements is summarised in Table 7.1 for cell 5. The variability in the real and imaginary parts of the cell's impedance is stated in both absolute and relative terms. The sample standard deviation, σ , is used to quantify the *absolute* variability, whereas the *relative* variability is quantified using the coefficient of variation (CV), defined as the ratio of sample standard deviation to sample mean, i.e. $CV = \sigma/\bar{x}$. Note that the CV is given as a percentage.

Table 7.1: Repeatability results.

f [Hz]	0.01	0.1	1	10	100
Re{Z}					
\bar{x} [mΩ]	110	100	99	81	56
σ [mΩ]	0.48	0.78	0.63	0.42	0.25
CV [%]	0.4	0.7	0.6	0.5	0.4
Im{Z}					
\bar{x} [mΩ]	-11	-5.0	-6.9	-17	-8.7
σ [mΩ]	1.1	0.19	0.061	0.16	0.11
CV [%]	10	4	1	0.9	1.2

Figure 7.12 shows the Nyquist plot of cell 5 obtained from the first measurement along with error bars whose length is determined by the standard deviation of the

²Takes into account that the current was not exactly 1 A due to the transducer's offset current.




Figure 7.12: Nyquist plot with error bars showing variability in measurement set.

real and imaginary parts of the five successive measurements.

The results presented above show that the agreement between repeated measurements is good and the scale of the random error is relatively small. Whilst the results above are for a single cell, very similar data spreads were observed for all of the cells in the battery pack, and also under different conditions, as will be shown in the next section.

In order to determine whether the measurement precision is high enough, the scale of variations in impedance due to factors that may be of interest must be considered. This is analysed in the following section.

7.4 Observable differences in cell impedance

This section will determine the scale of the variations in impedance due to factors that may be of interest, e.g. impedance variations between different cells in a battery pack, or impedance variations due to SoC. The purpose of this is to compare the scale of these variations with the scale of the random error (described in the previous section) in order to determine whether the measurement precision is high enough to detect meaningful changes in cell impedance.

Figure 7.13 is a Nyquist plot depicting the impedance of a number of cells in the SBP. All measurements were obtained at the same time and under the same condi-




Figure 7.13: Nyquist plot showing impedance differences between cells in SBP.




Figure 7.14: Impedance spectra of one cell at different SoC values.

tions as described in the previous section. The SoC of the cells was approximately 40 %. It should also be noted that the cells comprising the SBP are all of the same type, acquired at the same time from a single distributor, and have only ever been used as a pack, i.e. the spread in the results can be attributed mainly to differences in manufacturing and measurement precision.

The plot demonstrates that the precision of the measurements is high enough to detect small impedance differences between cells in a battery pack.

As explained in Chapter 4, the impedance of LIBs is sensitive to changes in SoC. In order to establish whether changes in SoC can be detected using the proposed system, the method used above to obtain the impedance spectra at 40 % SoC is repeated at various SoC values: beginning at 90 %, measurements are acquired in steps of 10 % SoC down to approximately 20 %. Figure 7.14 shows the results for cell 5. As before, the error bars are based on five successive measurements.

Figure 7.14 demonstrates that changes in SoC lead to variations in the cell impedance that are large enough to be observable using the proposed method.

Based on the results presented in this section it can be concluded that the measurement precision of the proposed power converter-based EIS system is high enough to detect meaningful changes in cell impedance.

7.5 Measurements under aliasing conditions

This section describes the experimental validation of the method proposed in [59] (and summarised in Section 5.3.7) as a possible solution to extend the high-frequency end of the measurement bandwidth when the available sampling rate is limited. The results reported herein are reproduced from [59].

A single cylindrical 26650 LFP cell (as used in the SBP) was connected to the

IBC input and a resistive load was connected to the converter's output. The current controller was configured for a 1 A dc current with a superimposed multisine perturbation containing two frequency components: 100.5 Hz and 1000.5 Hz. The signals were initially sampled at a rate of 100 kSa/s over a 10 s window. Figure 7.15 shows the resulting magnitude spectra of both current and voltage, confirming that the components are at the expected frequencies.

The signals were then downsampled to a sub-Nyquist rate of 8 Sa/s. As Figure 7.16 shows, the 100.5 Hz and 1000.5 Hz components are folded back into the $[0, f_s/2)$ range, appearing as aliases at 3.5 Hz and 0.5 Hz, respectively.




Figure 7.15: Voltage and current signals sampled at 100 kSa/s.




Figure 7.16: Voltage and current signals sampled at 8 Sa/s.

Figure 7.16 also shows that, although the components' frequencies have changed, their magnitudes have not changed significantly.

The impedance values obtained from both data sets are listed in Table 7.2.

Table 7.2: Impedance estimation results from measurement under aliasing conditions.

Sample rate	100 kSa/s	8 Sa/s
	Impedance at 100.5 Hz	
$ Z $ [m Ω]	53.7	54.9
$\angle Z$ [°]	-4.7	5.7
	Impedance at 1000.5 Hz	
$ Z $ [m Ω]	51.9	47.8
$\angle Z$ [°]	14.6	8.6

The results in Table 7.2 confirm that, in principle, the proposed method works correctly. The magnitude and phase angle of the impedance values based on the downsampled data are generally of the same order as those based on the data that is unaffected by aliasing. The observed differences in the impedance estimates from the downsampled data can most likely be attributed to an increase in the noise spectral density due to the substantial downsampling factor of 12500, since all of the noise components are condensed down to a very small number of spectral bins. Reducing the downsampling factor and averaging repeated measurements are two approaches that could be used to further increase the accuracy of the impedance estimates.

Chapter 8

Conclusion

The aim of this work was to contribute to the development of the power converter-based EIS technique for diagnostic battery monitoring applications. A survey of the existing literature highlighted three aspects with scope for improvement: 1) limited experimental validation, restricted to monitoring a single cell, or two cells in series; 2) a narrow measurement bandwidth relative to conventional, offline EIS instrumentation; and 3) insufficient consideration of the effects of introducing a measurement perturbation on the overall system, in particular the load. The objectives of this project were thus chosen to address these shortcomings. Specifically, this work set out to: experimentally validate the feasibility of using the power converter-based EIS technique to monitor a battery pack consisting of 16 series-connected LIBs; develop solutions to the technique's bandwidth restrictions; and finally, to propose strategies on how to deal with the effects of the measurement perturbation on the system load.

An experimental setup was developed around a battery pack with a nominal voltage of 51.2 V and nominal energy capacity of approximately 4.6 kWh, consisting of 16 series-connected lithium iron phosphate modules. This was connected to the input side of a dual-phase interleaved boost converter with a maximum rated output power of 2 kW. A proportional plus integral controller was used to control the input (battery) current and inject a multisine measurement perturbation with 5 logarithmically-spaced components located between 0.01 Hz and 100 Hz. A signal conditioning board was developed to allow the 16 battery voltages to be acquired under the high prevailing common-mode voltages. An embedded controller was employed to implement the real-time control and signal acquisition functions.

Experimental results obtained using the developed setup validated the correct operation of the system, thus allowing for the first time a 16S1P 4.6 kWh battery pack to be monitored using the power converter-based EIS method. In addition,

the lower limit of 0.01 Hz represents an order of magnitude decrease in measurement frequency compared to existing power converter-based EIS implementations. With the batteries discharging at a dc current of 20 A, the controller was able to introduce the ac current perturbations at the desired frequencies. The measurement precision was shown to be high enough to allow meaningful impedance variations to be detected: both due to small differences between cells in a battery pack, and also due to changes in SoC.

At low frequencies, SoC-induced variations in the battery voltage were identified as a factor limiting the measurement bandwidth. An analysis of existing voltage drift compensation methods showed that they were significantly limited in their ability to accurately fit sinusoidal and multisine signals: residual errors in the amplitude and phase angle estimation of the voltage after compensation using existing techniques were shown to be substantial. A new drift compensation technique was presented that was shown to overcome some of the limitations of existing methods: the ramp plus sine (RPS) method was shown to be capable of accurately estimating the slope of a linear drift in sine and multisine signals, thus enabling linear drifts in the voltage to be compensated, and allowing the measurement bandwidth to be extended to frequencies in the millihertz range. The main limitation of the RPS method is that — just like existing methods — it assumes that the drift in the voltage is linear.

At high frequencies, maximum ADC sampling rates were considered as the main factor restricting the measurement bandwidth, particularly in applications involving large battery packs requiring potentially hundreds of measurement channels. A solution was proposed that allows sample rates below the Nyquist rate (i.e. twice the signal's bandwidth) to be used, by exploiting the fact that the frequency content of signals in EIS measurements is known in advance. It was shown that by correctly choosing the sampling rate, observation window length, and the frequency components of the measurement perturbation, a loss of information due to aliasing can be avoided, thus enabling the measurement bandwidth to be extended to higher frequencies, which would otherwise be out of range. The main limitation of this method is associated with the restrictions that it places on the allowed frequency content of the measurement perturbation.

The proposed solutions to extend the bandwidth of power converter-based EIS were validated experimentally using the system developed in this work.

With the battery pack discharging at a constant dc current of 20 A, a drift caused by the change in SoC during the acquisition was observed in the battery voltage. It was shown that if left uncompensated, this drift led to errors in the estimation of the voltage amplitude in the range of 80 % and 10 % at 0.01 Hz and 0.1 Hz, respectively.

Further, a phase angle estimation error of approximately 15° could be avoided at 0.01 Hz by compensating for the drift in the voltage using the RPS method.

A multisine measurement perturbation with components at frequencies of 100.5 and 1000.5 Hz was used to validate the sub-Nyquist sampling method proposed to extend the high-frequency end of the measurement bandwidth when the available sampling rate is limited. Current and voltage signals were downsampled to a rate of 8 Sa/s. The spectra of the signals and the resulting impedance estimates confirmed the feasibility of the method.

Finally, several strategies were proposed on how to deal with the effects of the measurement perturbation on the load. In the simplest case, when measurements are restricted to one side of the power converter, appropriate design of the filter on the opposite side of the converter was proposed as a feasible solution.

Power converter-based EIS is an effective method to determine the impedance of batteries and other electrochemical energy storage devices. It has the potential to improve the diagnostic monitoring capability of battery management systems, and thus enhance the safety, efficiency, and reliability of the battery-powered applications that have become such important parts of the ongoing energy transition. As a result, it can be expected to remain the subject of future research. Future work could involve: continuing the development of strategies to deal with the effects of the measurement perturbation, for which some of the ground work was laid in this project; the development of non-linear drift compensation techniques; and the integration of power converter-based EIS in existing commercial battery management systems.

References

- [1] United Nations Framework Convention on Climate Change, “The Paris Agreement.” [Online]. Available: <https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement>. Accessed on: 2021-10-16.
- [2] Department for Transport, “Decarbonising Transport – A Better, Greener Britain,” 2021. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1009448/decarbonising-transport-a-better-greener-britain.pdf. Accessed on: 2021-10-16.
- [3] H. Ritchie, “Sector by sector: Where do global greenhouse gas emissions come from?” Sep. 2020. [Online]. Available: <https://ourworldindata.org/ghg-emissions-by-sector>. Accessed on: 2021-10-16.
- [4] E. Barsoukov and J. R. Macdonald, *Impedance Spectroscopy. Theory, Experiment, and Applications.*, 2nd ed. Wiley-Interscience, 2005.
- [5] A. Jossen, “Fundamentals of battery dynamics,” *Journal of Power Sources*, vol. 154, no. 2, pp. 530–538, Mar. 2006.
- [6] Y. Zhang, C.-Y. Wang, and X. Tang, “Cycling degradation of an automotive LiFePO₄ lithium-ion battery,” *Journal of Power Sources*, vol. 196, no. 3, pp. 1513–1520, Feb. 2011.
- [7] 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,” *Applied Energy*, vol. 102, pp. 885–897, Feb. 2013.
- [8] Y. Olofsson, J. Groot, T. Katrašnik, and G. Tavčar, “Impedance spectroscopy characterisation of automotive NMC/graphite Li-ion cells aged with realistic PHEV load profile,” in *2014 IEEE International Electric Vehicle Conference (IEVC)*, Dec. 2014, pp. 1–6.

- [9] E. Samadani, S. Farhad, W. Scott, M. Mastali, L. E. Gimenez, M. Fowler, and R. A. Fraser, “Empirical Modeling of Lithium-ion Batteries Based on Electrochemical Impedance Spectroscopy Tests,” *Electrochimica Acta*, vol. 160, pp. 169–177, Apr. 2015.
- [10] J. Lee and W. Choi, “Novel State-of-Charge Estimation Method for Lithium Polymer Batteries Using Electrochemical Impedance Spectroscopy,” *Journal of Power Electronics*, vol. 11, no. 2, pp. 237–243, Mar. 2011.
- [11] W. Huang and J. A. A. Qahouq, “An Online Battery Impedance Measurement Method Using DC–DC Power Converter Control,” *IEEE Transactions on Industrial Electronics*, vol. 61, no. 11, pp. 5987–5995, Nov. 2014.
- [12] R. Ferrero, C. Wu, A. Carboni, S. Toscani, M. De Angelis, H. George-Williams, E. Patelli, and P. A. Pegoraro, “Low-Cost Battery Monitoring by Converter-Based Electrochemical Impedance Spectroscopy,” in *2017 IEEE International Workshop on Applied Measurements for Power Systems (AMPS)*. Liverpool, United Kingdom: IEEE, Sep. 2017, pp. 1–6.
- [13] M. A. Varnosfaderani and D. Strickland, “Online impedance spectroscopy estimation of a battery,” in *2016 18th European Conference on Power Electronics and Applications (EPE'16 ECCE Europe)*, Sep. 2016, pp. 1–10.
- [14] J. A. A. Qahouq and Z. Xia, “Single-Perturbation-Cycle Online Battery Impedance Spectrum Measurement Method With Closed-Loop Control of Power Converter,” *IEEE Transactions on Industrial Electronics*, vol. 64, no. 9, pp. 7019–7029, Sep. 2017.
- [15] S. M. Rakiul Islam and S.-Y. Park, “High Precision On-Line Impedance Measurement for a Li-Ion battery,” in *2019 10th International Conference on Power Electronics and ECCE Asia (ICPE 2019 - ECCE Asia)*, May 2019, pp. 2694–2699.
- [16] Z. Xia and J. A. Abu Qahouq, “High Frequency Online Battery Impedance Measurement Method Using Voltage and Current Ripples Generated by DC-DC Converter,” in *2020 IEEE Applied Power Electronics Conference and Exposition (APEC)*, Mar. 2020, pp. 1333–1338.
- [17] C. G. Moral, D. Fernandez, J. M. Guerrero, D. Reigosa, C. R. Pereda, and F. Briz, “Thermal Monitoring of LiFePO₄ Batteries Using Switching Harmonics,” *IEEE Transactions on Industry Applications*, vol. 56, no. 4, pp. 4134–4145, Jul. 2020.

- [18] S. M. R. Islam and S.-Y. Park, "Precise Online Electrochemical Impedance Spectroscopy Strategies for Li-Ion Batteries," *IEEE Transactions on Industry Applications*, vol. 56, no. 2, pp. 1661–1669, Mar. 2020.
- [19] E. Sadeghi, M. H. Zand, M. Hamzeh, M. Saif, and S. M. M. Alavi, "Controllable Electrochemical Impedance Spectroscopy: From Circuit Design to Control and Data Analysis," *IEEE Transactions on Power Electronics*, vol. 35, no. 9, pp. 9933–9942, Sep. 2020.
- [20] A. J. Bard, G. Inzelt, and F. Scholz, Eds., *Electrochemical Dictionary*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012.
- [21] A. J. Bard and L. R. Faulkner, *Electrochemical Methods: Fundamentals and Applications*, 2nd ed. Wiley, 2001.
- [22] T. Christen and M. W. Carlen, "Theory of Ragone plots," *Journal of Power Sources*, vol. 91, no. 2, pp. 210–216, Dec. 2000.
- [23] J. B. Goodenough and K.-S. Park, "The Li-Ion Rechargeable Battery: A Perspective," *Journal of the American Chemical Society*, vol. 135, no. 4, pp. 1167–1176, Jan. 2013.
- [24] K. Persson, V. A. Sethuraman, L. J. Hardwick, Y. Hinuma, Y. S. Meng, A. van der Ven, V. Srinivasan, R. Kostecki, and G. Ceder, "Lithium Diffusion in Graphitic Carbon," *The Journal of Physical Chemistry Letters*, vol. 1, no. 8, pp. 1176–1180, Apr. 2010.
- [25] J. Asenbauer, T. Eisenmann, M. Kuenzel, A. Kazzazi, Z. Chen, and D. Bresser, "The success story of graphite as a lithium-ion anode material – fundamentals, remaining challenges, and recent developments including silicon (oxide) composites," *Sustainable Energy & Fuels*, vol. 4, no. 11, pp. 5387–5416, 2020.
- [26] A. Mahmoudzadeh Andwari, A. Pesiridis, S. Rajoo, R. Martinez-Botas, and V. Esfahanian, "A review of Battery Electric Vehicle technology and readiness levels," *Renewable and Sustainable Energy Reviews*, vol. 78, pp. 414–430, Oct. 2017.
- [27] K. O. Papailiou, Ed., *Springer Handbook of Power Systems*, ser. Springer Handbooks. Singapore: Springer Singapore, 2021.
- [28] Q. Wang, B. Mao, S. I. Stoliarov, and J. Sun, "A review of lithium ion battery failure mechanisms and fire prevention strategies," *Progress in Energy and Combustion Science*, vol. 73, pp. 95–131, Jul. 2019.

- [29] S. McCluer, “Battery Technology for Data Centers and Network Rooms: Ventilation of Lead-Acid Batteries,” 2012.
- [30] L. Energizer Brands, “NiMH Handbook and Application Manual,” 2018. [Online]. Available: https://data.energizer.com/pdfs/nickelmetalhydride_appman.pdf. Accessed on: 2022-03-07.
- [31] J. Vetter, P. Novák, M. R. Wagner, C. Veit, K. C. Möller, J. O. Besenhard, M. Winter, M. Wohlfahrt-Mehrens, C. Vogler, and A. Hammouche, “Ageing mechanisms in lithium-ion batteries,” *Journal of Power Sources*, vol. 147, no. 1, pp. 269–281, Sep. 2005.
- [32] K2 Energy Solutions Inc., “High Capacity K226650E02 Energy Cell Data Sheet,” C004.2 REV-B.
- [33] A123 Systems Inc., “Nanophosphate High Power Lithium Ion Cell ANR26650M1-B,” 2012.
- [34] K. Kutluay, Y. Cadirci, Y. Ozkazanc, and I. Cadirci, “A new online state-of-charge estimation and monitoring system for sealed lead-acid batteries in Telecommunication power supplies,” *IEEE Transactions on Industrial Electronics*, vol. 52, no. 5, pp. 1315–1327, Oct. 2005.
- [35] W. Waag, C. Fleischer, and D. U. Sauer, “Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles,” *Journal of Power Sources*, vol. 258, pp. 321–339, Jul. 2014.
- [36] V. Pop, Ed., *Battery Management Systems: Accurate State-of-Charge Indication for Battery Powered Applications*, ser. Philips Research Book Series. Dordrecht: Springer, 2008, no. v. 9.
- [37] W. Waag and D. U. Sauer, “Adaptive estimation of the electromotive force of the lithium-ion battery after current interruption for an accurate state-of-charge and capacity determination,” *Applied Energy*, vol. 111, pp. 416–427, Nov. 2013.
- [38] J. Aylor, A. Thieme, and B. Johnso, “A battery state-of-charge indicator for electric wheelchairs,” *IEEE Transactions on Industrial Electronics*, vol. 39, no. 5, pp. 398–409, Oct. 1992.
- [39] W. Dreyer, J. Jamnik, C. Guhlke, R. Huth, J. Moskon, and M. Gaberscek, “The thermodynamic origin of hysteresis in insertion batteries,” *Nature Materials*, no. 5, p. 448, 2010.

- [40] T. Huria, G. Ludovici, and G. Lutzemberger, “State of charge estimation of high power lithium iron phosphate cells,” *Journal of Power Sources*, vol. 249, pp. 92–102, Mar. 2014.
- [41] S. Ebbesen, P. Elbert, and L. Guzzella, “Battery State-of-Health Perceptive Energy Management for Hybrid Electric Vehicles,” *IEEE Transactions on Vehicular Technology*, vol. 61, no. 7, pp. 2893–2900, Sep. 2012.
- [42] S. Grolleau, A. Delaille, H. Gualous, P. Gyan, R. Revel, J. Bernard, E. Redondo-Iglesias, and J. Peter, “Calendar aging of commercial graphite/LiFePO₄ cell – Predicting capacity fade under time dependent storage conditions,” *Journal of Power Sources*, vol. 255, pp. 450–458, Jun. 2014.
- [43] H. Yuan and L. Dung, “Offline State-of-Health Estimation for High-Power Lithium-Ion Batteries Using Three-Point Impedance Extraction Method,” *IEEE Transactions on Vehicular Technology*, vol. 66, no. 3, pp. 2019–2032, Mar. 2017.
- [44] M. Bercibar, I. Gandiaga, I. Villarreal, N. Omar, J. Van Mierlo, and P. Van den Bossche, “Critical review of state of health estimation methods of Li-ion batteries for real applications,” *Renewable and Sustainable Energy Reviews*, vol. 56, pp. 572–587, Apr. 2016.
- [45] C. Pastor-Fernández, K. Uddin, G. H. Chouchelamane, W. D. Widanage, and J. Marco, “A Comparison between Electrochemical Impedance Spectroscopy and Incremental Capacity-Differential Voltage as Li-ion Diagnostic Techniques to Identify and Quantify the Effects of Degradation Modes within Battery Management Systems,” *Journal of Power Sources*, vol. 360, pp. 301–318, Aug. 2017.
- [46] A. Eddahech, O. Briat, N. Bertrand, J.-Y. Delétage, and J.-M. Vinassa, “Behavior and state-of-health monitoring of Li-ion batteries using impedance spectroscopy and recurrent neural networks,” *International Journal of Electrical Power & Energy Systems*, vol. 42, no. 1, pp. 487–494, Nov. 2012.
- [47] ISO, “Electrically propelled road vehicles — Test specification for lithium-ion traction battery packs and systems (ISO 12405:4-2018),” Jul. 2018.
- [48] O. Kanoun, A. Y. Kallel, H. Nouri, B. B. Atitallah, D. Haddad, Z. Hu, M. Talbi, A. Al-Hamry, R. Munja, F. Wendler, R. Barioul, T. Keutel, and A. Mangler, “Impedance Spectroscopy: Applications, Advances and Future Trends,” *IEEE Instrumentation Measurement Magazine*, vol. 25, no. 3, pp. 11–21, May 2022.

- [49] L. Ljung, *System Identification: Theory for the User*, 2nd ed., ser. Prentice Hall Information and System Sciences Series. Upper Saddle River, NJ: Prentice Hall PTR, 1999.
- [50] R. Pintelon and J. Schoukens, *System Identification: A Frequency Domain Approach*, 2nd ed. Hoboken, N.J. Piscataway, NJ: Wiley IEEE Press, 2012.
- [51] S. Rodrigues, N. Munichandraiah, and A. K. Shukla, “AC impedance and state-of-charge analysis of a sealed lithium-ion rechargeable battery,” *Journal of Solid State Electrochemistry*, vol. 3, no. 7-8, pp. 397–405, Sep. 1999.
- [52] V. Sauvant-Moynot, J. Bernard, R. Mingant, A. Delaille, F. Mattera, S. Mailley, J.-L. Hognon, and F. Huet, “ALIDISSI, a Research Program to Evaluate Electrochemical Impedance Spectroscopy as a SoC and SoH Diagnosis Tool for Li-ion Batteries,” *Oil & Gas Science and Technology – Revue de l’Institut Français du Pétrole*, vol. 65, no. 1, pp. 79–89, Jan. 2010.
- [53] T. Stockley, M. Bowkett, J. Williams, M. Hathway, and K. Thanapalan, “Advanced EIS Techniques for Performance Evaluation of Li-ion Cells,” *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 8610–8615, 2014.
- [54] J. Xu, C. C. Mi, B. Cao, and J. Cao, “A new method to estimate the state of charge of lithium-ion batteries based on the battery impedance model,” *Journal of Power Sources*, vol. 233, pp. 277–284, Jul. 2013.
- [55] N. Mohan, T. M. Undeland, and W. P. Robbins, *Power Electronics: Converters, Applications, and Design*, 3rd ed. Hoboken, NJ: John Wiley & Sons, 2003.
- [56] H. Wang, A. Gaillard, and D. Hissel, “Online electrochemical impedance spectroscopy detection integrated with step-up converter for fuel cell electric vehicle,” *International Journal of Hydrogen Energy*, vol. 44, no. 2, pp. 1110–1121, Jan. 2019.
- [57] LEM, “Isolated current and voltage transducers. Characteristics — Applications — Calculations,” May 2004, 3rd Edition.
- [58] BioLogic Science Instruments, “Drift correction in electrochemical impedance measurements (EC-Lab — Application Note #17),” 2010.
- [59] A. Sandschulte, R. Ferrero, L. Hardwick, and E. Patelli, “Approach to Wide-Frequency Battery Impedance Measurements in Commercial Applications,” in *2019 IEEE 10th International Workshop on Applied Measurements for Power Systems (AMPS)*. Aachen, Germany: IEEE, Sep. 2019, pp. 1–6.

- [60] R. W. Erickson and D. Maksimović, *Fundamentals of Power Electronics*, 2nd ed. Springer Science+Business Media, 2001.
- [61] R. D. Middlebrook, “Small-signal modeling of pulse-width modulated switched-mode power converters,” *Proceedings of the IEEE*, vol. 76, no. 4, pp. 343–354, Apr. 1988.
- [62] K2 Energy Solutions Inc., “High Capacity K2B3V90E Energy Module Data Sheet,” P0004 REV-A.
- [63] N. Jantharamin and L. Zhang, “Analysis of multiphase interleaved converter by using state-space averaging technique,” in *2009 6th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*, vol. 01, May 2009, pp. 288–291.
- [64] National Instruments, “cRIO-9035.” [Online]. Available: <https://www.ni.com/en-gb/support/model.crio-9035.html>. Accessed on: 2020-11-12.
- [65] National Instruments, “NI 9205 Datasheet,” Jul. 2017.
- [66] National Instruments, “NI 9215 Datasheet,” Mar. 2016.
- [67] LEM, “Current Transducer LA 25-P Data Sheet,” Nov. 2014, Version 3.
- [68] Traco Power, “DC/DC Converters TEN 5 Series, 6 Watt,” Jun. 2016, Rev. January 06. 2015.
- [69] Analog Devices, “Low Cost, Miniature Isolation Amplifiers AD202/AD204,” 2012, REV. D.
- [70] Analog Devices, “Clock Driver AD246,” 2012, REV. 0.
- [71] Joint Committee for Guides in Metrology (JCGM), *JCGM 200:2012 - International Vocabulary of Metrology - Basic and General Concepts and Associated Terms (VIM)*, 3rd ed., 2012.