



The  
University  
Of  
Sheffield.

## **Multilevel Converters for Battery Energy Storage: How Many Levels and Why?**

**By:**

Alexander Petersen

A thesis submitted in partial fulfilment of the requirements for the degree of  
Doctor of Philosophy

The University of Sheffield  
Faculty of Engineering  
Department of Electrical and Electronic Engineering

Submission Date  
12<sup>th</sup> August 2019



# Table of Contents

Abstract	1
List of publications	2
<b>Chapter 1: Introduction and Literature Review</b>	<b>3</b>
1.1 The Evolving Energy Grid	3
1.2 Converters for Grid-Attached Battery Energy Storage	6
1.3 Multilevel Converters	10
1.4 Conclusions	13
<b>Chapter 2: A Method for Predicting Converter Power Loss</b>	<b>19</b>
2.1 The Reference Converter	20
2.2 Loss Metrics	21
2.2.1 On-State Resistance	22
2.2.2 Transient Loss in the Gate	22
2.2.3 Transient Output Loss	24
2.2.4 Gate Drive Loss	24
2.2.4.1 Transient Gate Drive Loss	25
2.2.4.2 Quiescent Gate Drive Loss	26
2.2.5 Diode Loss	28
2.3 Alternative Semiconductor Technologies	30
2.3.1 Insulated Gate Bipolar Transistors (IGBTs)	30
2.3.2 Silicon Carbide (SiC) MOSFETs	30
2.3.3 Gallium Nitride (GaN) High Electron Mobility Transistors (HEMTs)	31
2.4 Conclusions	32
<b>Chapter 3: Practical Parameter Estimation</b>	<b>34</b>
3.1 Slightly Ambiguous Parameters	35
3.2 Pin-to-Pin Device Capacitances	36
3.3 Experimental Pin-to-Pin Capacitance Evaluations	39
3.4 Plateau Voltage	47
3.5 Body Diode Performance	48
3.6 Conclusions	56
<b>Chapter 4: Optimal Power Loss Trends</b>	<b>58</b>
4.1 Generating Results	58

4.2 Results of Power Loss Prediction Method for Silicon MOSFETs	59
4.3 Validation of the Method with SPICE Modelling	63
4.4 Alternative Semiconductor Technologies	66
4.4.1 Insulated Gate Bipolar Transistors (IGBTs)	66
4.4.2 Silicon Carbide (SiC) MOSFETs	67
4.4.3 Gallium Nitride (GaN) High Electron Mobility Transistors (HEMTs)	69
4.5 Experimental Investigation of GaN HEMT Gate Turn-On Energy	72
4.6 Conclusions	77
<b>Chapter 5: Thermal Performance in Multilevel Converters</b>	<b>78</b>
5.1 Finite Element Thermal Modelling	78
5.1.1 FEM Thermal Trends with More Devices	79
5.1.2 FEM Thermal Trends with Pad Size	83
5.2 Modelling with Resistor Network Equivalent	84
5.3 Experimental Validation of Thermal Modelling	90
5.4 Analytical Thermal Optimisation	92
5.5 Transient Thermal Performance	95
5.6 Conclusions	98
<b>Chapter 6: Reliability in Multilevel Converters</b>	<b>99</b>
6.1 The Arrhenius Equation	99
6.2 Results Using Arrhenius Model	101
6.3 Multilevel Converters and Battery Failure	103
6.4 Conclusions	106
<b>Chapter 7: Conclusions and Further Work</b>	<b>109</b>
7.1 Conclusions	109
7.1 Further Work	111
<b>Appendices</b>	<b>112</b>
Appendix A	112
Appendix B	114
Appendix C	117
Appendix D	118
Appendix E	119

## Table of Figures

Figure	Page
<b>1.1 Breakdown of UK energy production 2006 and 2018</b>	<b>3</b>
1.2 UK wind power generation over 2018	4
1.3 Schematics for typical BESS three phase converter and DC-DC converters	7
1.4 Two converter topologies for single phase BESSs	8
1.5 Performance of a crude half bridge single phase BESS	10
1.6 Simplified schematic of an N <sup>th</sup> -order cascaded H-bridge multilevel converter	11
<b>2.1 Simplified schematic of an N<sup>th</sup>-order cascaded H-bridge multilevel converter</b>	<b>20</b>
2.2 Typical MOSFET gate charge-voltage curve	23
2.3 Circuit diagram of a typical isolated gate drive circuit	24
2.4 Gate turn-on and turn-off waveforms for a Fairchild FCH47N60 MOSFET	26
2.5 An H-bridge with added diodes to avoid using body diodes	29
<b>3.1 Curves for variation in MOSFET performance with respect to temperature</b>	<b>35</b>
3.2 A range of MOSFET capacitance-voltage curves from manufacturer datasheets	37
3.3 Two examples of a bilinear approximation for the capacitance-voltage curves	38
3.4 Three identical MOSFET capacitance-voltage curves for different devices	39
3.5 Manufacturer test circuit for measuring MOSFET capacitance	40
3.6 Experimental test circuit for measuring MOSFET capacitance	40
3.7 Result waveforms from experimental MOSFET capacitance test rig	42
3.8 Full test results for Fairchild FCH47N60	44
3.9 Full test results for Texas Instruments CSD17573Q5B	45
3.10 Full test results for Fairchild FQA44N30	46
3.11 Experimental and datasheet calculated Miller charge over a range of $V_{DS}$	47
3.12 The voltage and current ratings MOSFETs that had their body diodes characterised	49
3.13 Plot showing test pulse train for diode characterisation	49
3.14 Body diode test data, along with linearised fit, for Fairchild FDMS86255	49
3.15 Raw data from body diode recovery charge test rig	51
3.16 Error distributions for all body diode testing	51
3.17 Body diode forward resistance against MOSFET voltage rating with best fit	54
3.18 Body diode threshold voltage against MOSFET input capacitance with best fit	54
3.19 Body diode reverse recovery charge against MOSFET output capacitance with best fit	55

<b>4.1 Power loss trends with converter order for optimal devices at 10kHz</b>	<b>59</b>
4.2 Power loss trends with converter order for optimal devices at 80kHz	60
4.3 Power loss trends with converter order for optimal devices at 600kHz	61
4.4 How optimal device parameters change with switching frequency	62
4.5 LTspice model used in validation of power loss in Infineon BSZ042N04NS	63
4.6 Waveform of instantaneous power dissipation for the model shown in figure 4.5	64
4.7 Comparison of on-state conduction loss for silicon MOSFETs and IGBTs	66
4.8 Power loss with increasing switching frequency for SiC and conventional MOSFETs	68
4.9 Power loss trends with the inclusion of GaN power devices	69
4.10 Heat map of the improvement enabled by GaN devices over converter order and frequency	71
4.11 Raw data from an early prototype test rig with excessive ringing	73
4.12 Pictures of the improved test rig	74
4.13 Renders of the daughter board PCB in the improved test rig	74
4.14 Raw test data from improved test rig comparing a GaN device with a silicon one	75
4.15 Gate turn on energy for a range of both GaN and silicon devices	76
<b>5.1 FEM results comparing heatsink performance as number of devices vary</b>	<b>79</b>
5.2 Trend in peak heatsink temperature with an increasing number of devices	79
5.3 Renders of the heatsinks used in the finite element modelling	80
5.4 A series of trend lines of peak heatsink temperatures varying with number of devices	81
5.5 FEM results comparing heatsink performance as thermal pad size varies	82
5.6 A series of trend lines of peak heatsink temperatures varying with thermal pad size	83
5.7 Diagram of an electrical equivalent model of a lossy plate at steady state	84
5.8 Results from electrical equivalent model showing heatsink temperature distributions	86
5.9 Trend derived from electrical equivalent model	86
5.10 Surface plots showing variance in error between the fit and the model results	88
5.11 Scatter plot of error values with respect to a “squareness” coefficient	89
5.12 IR imagery showing results of experimental validation	91
5.13 Fit trend line from experimental validation	91
5.14 Trend in thermal headroom with respect to number of devices in one case	94
5.15 Trend in thermal headroom with respect to number of devices in another case	94
5.16 A comparison of power switching device loss in different modulation schemes	95
5.17 A Foster style thermal model of the internals of a power transistor	96

5.18 The frequency response of the thermal model in figure 5.17	97
<b>6.1 Trend in MTTF with respect to converter order in one case</b>	<b>102</b>
6.2 Trend in MTTF with respect to converter order in another case	102
6.3 Simplified schematic of an $N^{\text{th}}$ -order cascaded H-bridge multilevel converter	104
<b>8.1 Performance of a crude half bridge single phase BESS</b>	<b>112</b>
<b>9.1 Error in the bilinear approximation of <math>C_{\text{iss}}</math></b>	<b>114</b>
9.2 Error in the bilinear approximation of $C_{\text{oss}}$	115
9.3 Error in the bilinear approximation of $C_{\text{rss}}$	115

## Table of Tables

<b>Table</b>	<b>Page</b>
1.1 Representative data for the performance of a wide range of grid energy storage systems	6
3.1 Strength of correlation between body diode forward resistance and MOSFET data	53
3.2 Strength of correlation between body diode threshold voltage and MOSFET data	53
3.3 Strength of correlation between body diode recovery charge and MOSFET data	54
4.1 Comparison between SPICE model results with the new method	65
4.2 List of SiC devices used in performance comparison	68
4.3 List of devices used in experimental validation of GaN turn on energy	75
5.1 Data for heatsinks used in thermal FEM while considering the effect of number of devices	80
5.2 Data for heatsinks used in thermal FEM while considering the effect of thermal pad size	83
5.3 Data for Foster thermal model in figure 5.17	97



## Nomenclature and Glossary

- BESS - Battery Energy Storage System
- BGA - Ball Grid Array, a type of surface mount component package with balls on the underside
- BJT - Bipolar Junction Transistor, a class of transistor
- BMS - Battery Management System, responsible for cell balancing amongst other things
- $C_{DS}$  - Capacitance between the drain and source of a MOSFET
- $C_{GD}$  - Capacitance between the gate and drain of a MOSFET
- $C_{GS}$  - Capacitance between the gate and source of a MOSFET
- $C_{iss}$  - Input capacitance, the sum of the MOSFET gate-source and gate-drain capacitance
- $C_{oss}$  - Output capacitance, the sum of the MOSFET drain-source and the gate-drain capacitance
- $C_{rss}$  - Feedback capacitance, the MOSFET gate-drain capacitance
- DUT - Device Under Test
- FEM - Finite Element Modelling, a type of computer simulation
- $f_{grid}$  - The frequency of the mains utility supply, expressed in Hertz (e.g. 50Hz in UK and Europe)
- $f_{switching}$  - The frequency at which the power electronics in the converter is switching, in Hertz
- GaN - Gallium Nitride, a chemical component of a new class of very high-speed power transistors
- HEMT - High Electron Mobility Transistor, a new set of transistor technologies, including GaN
- HVDC - High Voltage Direct Current, sometimes used for high level grid power transfer
- $I_{D,MAX}$  - The maximum drain current to which a MOSFET is rated
- IR - Infrared, used in this case to refer to thermal imaging equipment that can “see” temperature
- I-V curve - A figure where current through a device is plotted against voltage across the device
- IGBT - Insulated Gate Bipolar Transistor, a class of power transistor
- MMC - Modular Multilevel Converter, a multilevel converter designed in regular, identical blocks
- MOSFET - Metal Oxide Semiconductor Field Effect Transistor, a class of power transistor
- MTTF - Mean Time To Failure, the mean time a device can be expected operate before breaking
- $N_{PWMres}$  - The bit depth of the resolution of the PWM output of the controller, e.g. 8 bit = 256 levels
- OCR - Optical Character Recognition, a technology for computer to read documents
- PCB - Printed Circuit Board
- $Q_{MILLER}$  - Miller Charge, the charge transferred in to the gate during the Miller shelf
- $Q_{RR}$  - The charge transferred back out of a diode during its reverse recovery
- $R_{DS,ON}$  - The on-state resistance between the drain and source in a MOSFET

Si - Silicon, used to refer to typical power switching device technologies

SiC - Silicon Carbide, a chemical component of a new class of both power transistors and diodes

SoC - State of Charge, the charge state of a cell or battery, often expressed as a percentage

SoH - State of Health, the remaining maximum charge a battery will hold compared to new

$T_{AV}$  - Mean temperature over the entire body, or in this case heatsink

PV - Photovoltaic, used to refer to solar electrical (as opposed to thermal) panels

RMS - Root Mean Square, a mathematical transform for AC signals in power calculations

SPICE - A circuit simulation software package

$V_{DS}$  - Drain-source voltage, the voltage at the drain relative to the source in a MOSFET

$V_{DS,MAX}$  - The maximum drain-source voltage to which a MOSFET is rated

$V_{GS}$  - Gate-source voltage, the voltage at the gate relative to the source in a MOSFET

$V_{LINK}$  - The DC link voltage, usually the total voltage across the battery string in this case

$V_{MAINS,RMS}$  - The RMS voltage of the utility mains supply

VRLA - Valve Regulated Lead Acid, a type of electrochemical battery

## Abstract

This work explores the potential benefits of cascaded H-bridge multilevel converters in low-voltage applications, particularly grid-attached battery energy storage systems (BESS). While some benefits of these are discussed in literature, this work seeks to create practical, quantitative models for system performance in terms of a number of key performance parameters. These models are then used to find the trends in these performance parameters with an increasingly high order converter, starting to answer the question of how many levels are best. The system performance parameters modelled are power loss, thermal performance and reliability. Wherever practical models and assumptions are validated, be that experimentally or through comparison with existing methods – this work includes a number of experimental series. The resulting trends explored highlight a number of interesting trends, principally: total power loss can be much lower, particularly at high switching frequencies; system thermal performance can be much improved owing to more efficient heatsink utilisation; and due to these thermal benefits, the system reliability based on switching device failure does not suffer as one might expect, and can in fact be higher under some conditions. The investigation also considers the use of cutting-edge switching device technology, such as gallium nitride power transistors, which a multilevel converter enables the use of, and in turn can significantly reduce power dissipation and increase switching frequency. Overall, the work adds new arguments in favour of multilevel converters in such applications and lowers the barrier to practical implementation by answering a number of questions a designer would likely ask.

The key novel contributions of this work are the results of the trends that were found in terms of converter power loss, system thermal performance and switching device reliability with respect to multilevel converter order – with the methodologies created for these being somewhat novel in their own right. Along the way, however, other novel work was conducted including: an experimental investigation in to the accuracy of voltage-capacitance curves provided by manufacturers; experimental derivation of relationships for predicting MOSFET body diode performance from readily available device parameters; analysis showing the potential impact of GaN devices on converter efficiency; an experimental validation of GaN device gate turn-on energy; creation and validation of empirical relationship for predicting how heatsink performance varies with more devices of a smaller size; as well as an exploration of whether the extreme small size of some modern power transistors could lead to unexpected thermal cycling issues.

## List of Publications

### **Quantitative Power Loss Analysis and Optimisation in Nth-order Low Voltage Multilevel Converters**

A Petersen, D Stone, M Foster, J Davidson

IET Power Electronics 20 Feb 2019

### **An Investigation into the Thermal Benefits of Multilevel Converters**

A Petersen, D Stone, M Foster, D Gladwin

IECON 2018

### **An Experimental Investigation of MOSFET Intrinsic Body Diode Performance**

A Petersen, D Stone, M Foster

ISIE 2018

### **On the impact of current generation commercial gallium nitride power transistors on power converter loss**

A Petersen, D Stone, M Foster

ELECTRONICS LETTERS 53(22):1487-1488 26 Oct 2017

### **Switching loss optimisation of cascaded H-bridge converters for bidirectional grid-tie battery energy storage systems**

A Petersen, D Stone, M Foster, D Gladwin

IECON 2017

# Chapter 1:

## Introduction and Literature Review

This chapter will outline the aims and objectives of this work and set out to explain the journey that was taken to arrive at this. Initially this will consist of a discussion of the field in a broad sense, before homing in on pertinent research and discussing some shortcomings in it that can be addressed.

### 1.1 The Evolving Energy Grid

In recent years, in the UK and around the world, a major shift in electrical energy generation has begun, away from large, centralised, hot power stations, generally powered by fossil fuels, and towards renewable energy sources such as wind and solar. Taking the UK as an example, between 2006 and 2018 the proportion of electricity derived from renewable energy sources has increased from 5% to 30% [1], as shown in figure 1.1. This is due in part to a worldwide effort to reduce emissions of greenhouse gases contributing to man-made climate change, enshrined in a variety of international political agreements such as the 2009 EU Renewable Energy Directive or the various articles of the 2015 Paris Agreement [2, 3].

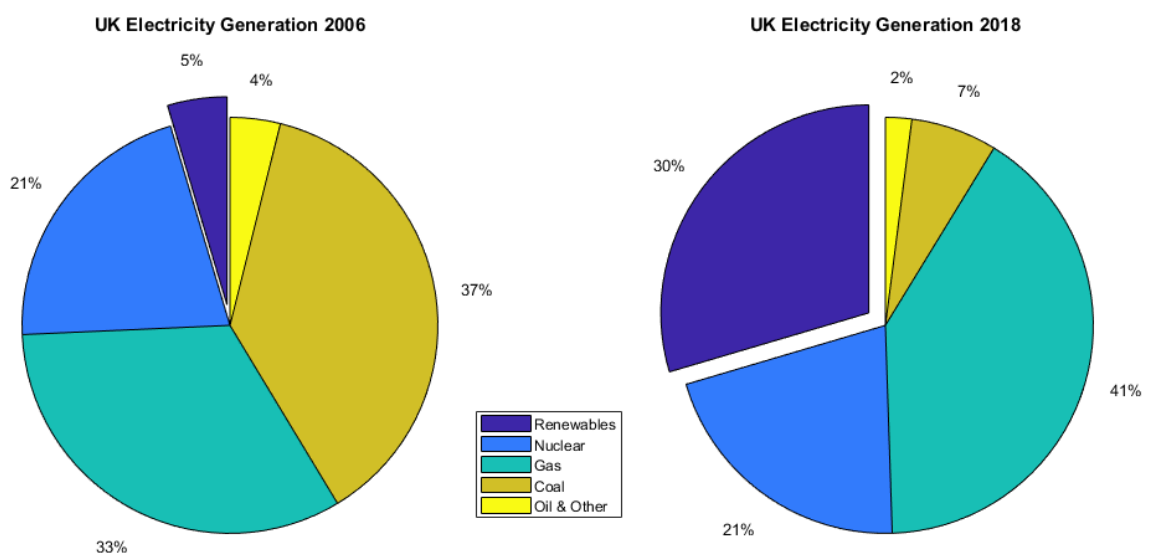


Fig. 1.1 A breakdown of UK electrical energy generation sources in 2006 and 2018, respectively. [4, 5]




Fig. 1.2 UK wind electricity generation over the year of 2018 [6].

Renewable energy sources typically include hydroelectric, solar photovoltaic (solar PV), wind and bioenergy (in the UK mostly burning wood and waste). The geography of the UK limits the amount of practical hydroelectric energy generation and, as such, there has been little growth in this area in recent years and is unlikely to be in the immediate future, but in other areas there has been significant growth: namely solar PV and wind generation. In fact, from 2017 to 2018 there has been a 30.1% increase in installed offshore wind generating capacity, with total installed wind generation capacity having surpassed 20GW [5]. Solar PV has also grown significantly, though perhaps as a result of current government policy, growth from 2017 to 2018 was a mere 4.2%, though total installed capacity still sits at 13GW [5]. This level of generation when compared with peak UK electricity demand of approximately 50GW [6], shows renewable generation capacity has become substantial.

The increasing influence of power sources such as wind and solar, while advantageous from the perspective of lowering overall CO<sub>2</sub> emissions, does have the significant drawback of unpredictable generation capacity. Figure 1.2 shows the variation of wind power generation over the year of 2018 in the UK, with output ranging from over 12GW to almost zero. While this is predictable to some extent, very precise medium- and long-term prediction is very challenging, only being as accurate as the weather forecast it is based on. Solar may seem that it would be more predictable, with cycles of the day being trivial to evaluate, but the extent of any overcast weather can be predicted with a similar accuracy as that of wind. While this illustrates the importance of diversity in an energy grid, as well as the value of hot generation that can rapidly ramp up and down such as with combined cycle gas turbine plants, the impact of this can be alleviated using large-scale energy storage.

Today in the UK, the only large-scale energy storage technology that has been implemented is hydroelectric pumped storage, whereby water is pumped from a lower reservoir to an upper reservoir to store energy during a surplus and allowed to flow back down through turbines to generate electricity in times of need. The largest of these (and, indeed, the largest in Europe) is Dinorwig power station in North Wales, with a peak generation capacity of 1872MW, a peak power input capacity of 1650MW and a total energy storage capacity of approximately 12GWh [7], making up approximately half of pumped storage capacity in the UK. Pumped storage has its drawbacks however, not only in land use and requiring specific local geography but also in large capital outlay, with Dinorwig having cost the UK government £425million in 1974 (£4.35billion equivalent in 2018) [7] making it the most expensive civil engineering project ever undertaken by the UK government at the time. Furthermore, the response time is somewhat limited, with a 0-80% output change with fore planning being possible in 12 seconds in the best possible case [7].

Pumped storage is well suited to bulk storage of energy, but less so to high speed, dynamic response. Other storage options can work much better for this, including flywheels [8, 9], supercapacitor banks [10, 11], compressed air [12, 13], cryogenic energy storage [14, 15] and, of course, electrochemical batteries. Electrochemical batteries have themselves many variants, including nickel, lead and lithium-based chemistries. Nickel has largely gone out of favour in recent years, as lithium batteries have greater energy density, higher power capacity, fewer technical limitations in charging and are no longer significantly more expensive [16]. Lead chemistries, be it flooded lead-acid batteries or gel type valve-regulated lead acid batteries (VRLAs), are a tempting prospect for grid-attached energy storage as their increased weight is not a barrier, as is the case of electric vehicles, and they are generally cheaper and easier to manage and even recycle at end of life [17]. However, lithium battery solutions have superior lifespan and lower maintenance requirements than a flooded lead-acid battery, while enabling higher charge/discharge rates relative to capacity (known as C-rate) – all this while new lithium chemistries are being actively developed and prices are continuously decreasing.

Table 1.1 shows a range of important performance parameters for the aforementioned range of potential grid energy storage solutions, while they can be difficult to directly compare this table draws together figures from a range of sources. Within the realm of faster responding systems, i.e. those with shorter discharge durations, batteries and flywheels are the principal mature technologies [18], with cryogenic storage being particularly far from practical, commercial realisation.

When considering a battery energy storage system (BESS), the principal question is what battery chemistry to use, with the two main options being lead-based chemistries and lithium-based chemistries. While there are a range of chemistries in both of these categories with different specific properties, there are general trends between the two. Specifically, all lithium chemistries have higher

specific energy, high specific power and longer lifespan than their lead-based counterparts [19-22]. While lead does have advantages in terms of reduced capital outlay, this is offset significantly by the shorter lifespan. Furthermore, as previously mentioned, lithium battery costs are reducing, owing to the significant global increase in demand over recent years [23]. All things considered, it seems probable that lithium battery technology will form the backbone of BESSs for the foreseeable future.

Storage System	Power (MW)	Discharge Duration	Efficiency (%)	Approx. Lifespan
Pumped Hydro	250 - 2000	6 - 18 hours	75 - 87	>30 years
Flywheel	0.5 - 5	15 sec to 15 min	93	20 years
Supercapacitor	10	Up to 30 seconds	90	>50,000 cycles
Compressed Air	15 - 400	2 - 24 hours	54 - 88	35 years
Cryogenic	5 - 200	5 hours	50 - 60	25 years
Li-Ion battery	5	25 min to hours	90	15 years
VRLA battery	1-20	40 min to hours	75 - 80	4 - 8 years
VRB flow battery	0.04 - 2	4 - 8 hours	75 - 80	10 years

Table 1.1 Typical or representative performance of a range of grid-scale energy storage solutions, data derived from [24], with supplementary cryogenic data from [14].

## 1.2 Converters for Grid-Attached Battery Energy Storage

So, far it has been discussed that there is demand for methods of storing energy from the national grid, that a BESS is a viable method for achieving this, and that lithium batteries are probably the preferred chemistry for such a system. The question now is: what type of converter should be used to interface a battery with the electrical grid? In this research we will focus on converters that form a bidirectional AC-DC interface. However, when closely integrating batteries with the DC side of a solar PV or DC microgrid systems, a DC-DC converter would normally be employed [25-27].

For a grid-tie BESS there is an apparent de facto standard for general converter topology, shown in figure 1.3a and featuring in numerous publications including [28-30], as well as being seen in almost all commercial converters built for such an application as it is extremely straightforward while achieving adequate performance. Some, rather than directly using the battery voltage as the DC link utilise an intermediate bidirectional DC-DC converter to transform the battery voltage (usually up to avoid exceeding the voltage rating of commonly available converters). Figure 1.3b shows a typical bidirectional DC-DC converter as might be used in such an application (such as in [31, 32]), or even an isolated bidirectional DC-DC converter [28], which naturally grants the additional benefit of the battery no longer being electrically referenced to the mains.



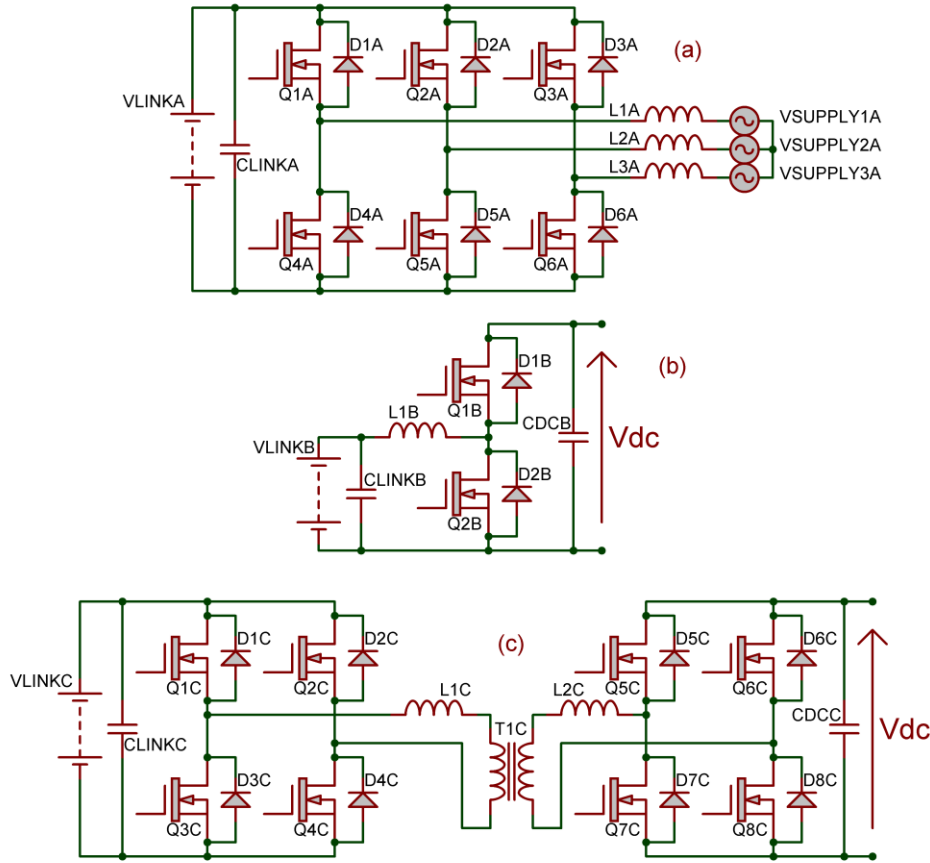


Fig. 1.3 Schematics for (a) a typical bidirectional AC-DC interface for three-phase BESSs, (b) a bidirectional DC-DC converter, and (c) an isolated bidirectional DC-DC converter.

The topology shown in figure 1.3a is an adequate converter to interface a battery (with or without intermediate DC-DC converter) with a three-phase utility mains supply, but this does not translate to a single-phase supply. Operation on a single-phase supply is clearly important and will be essential, as the use of energy storage in a domestic setting continues to increase, as few homes have a three-phase connection.

Figure 1.4a shows how a typical full bridge, also known as an H-bridge, can readily interface a battery with a single-phase mains supply, as in [33-35], and is certainly the most common topology seen within literature. Figure 1.4b shows that with a split DC supply one can operate with only a half bridge [36, 37], therefore half as many switching devices, though with the disadvantage of requiring double the total DC link voltage for the same peak voltage at the AC side as compared with the full bridge solution. However, the main disadvantage of the converter in figure 1.4b is usually the need for two separate DC supplies, or more often splitting a single DC supply using two large capacitors, but balancing these capacitors becomes a significant technical challenge by itself – the use of a battery solves this outright, as by simply dividing one large battery in to two smaller batteries the DC link is split in half with no additional difficulty. These converters, as with the three phase topologies, are sometimes paired with an intermediate DC-DC converter that may or may not be isolated [33, 34].




Fig. 1.4 Two converter topologies for single phase grid tie BESS: (a) a full bridge converter and (b) a half bridge converter with a split DC link.

Any of these converter topologies can bidirectionally pass power between a DC side and an AC side, as would be desired within a BESS, but the use of a battery has some additional concerns, a key one amongst those is balancing, particularly when using lithium chemistries [38, 39], where Ni and Pb chemistry cells can self-regulate under overcharge. Battery imbalance is where the cells that make up the battery are not all at the same voltage. For example, imbalance would be an issue as if not all cells are at the same state, as charging the full string up to its maximum voltage could result in cells within the battery exceeding their individual maximum voltages – the same is true for an under-voltage condition. A system that monitors for imbalance and attempts to remedy it is broadly referred to as a battery management systems (BMS), with most commercial lithium battery packs including a BMS from the manufacturer that will usually also monitor other factors that are important to safety, such as cell temperature [40].

Upon a BMS noting an imbalance, however, what can be done to address it? Naturally, the BMS must be able to communicate to the converter itself and reduce or even stop operation of the system to avoid an imbalance getting out of control, with few practical implementations of balancing circuitry capable of keeping up with the rate of imbalance in the worst case. The most common form of cell balancing is a switched shunt resistor network, where each cell has a resistor across it that the BMS can turn on or off to dissipate any excess energy thermally and permit the whole string to charge up fully without exceeding the maximum voltage of any single element – this is clearly a very simple and reliable method, presumably contributing to its apparent popularity. Switched shunt resistor networks are common even in very large batteries, for instance, this is the balancing mechanism in a 1MWhr lithium-titanate battery used in host research group as a grid- attached BESS [41].

There are numerous more sophisticated balancing methods, many of which enable recovery of excess energy rather than wasting it. These include: PWM current shunting that can change how two series cells pass current; resonant converter shunting which is much like PWM current shunting but more efficient and complex; boost shunting whereby excess energy is boosted and passed to some external energy storage; multi-winding transformers where power can be driven in or out at every cell while maintaining electrical isolation; multiple transformers with the benefit over a multi-tapped transformer that multiple cells can be doing different things simultaneously; a switched transformer where a single transformer is routed to the cell that needs balancing; and the list goes on [42, 43]. These various implementations all have their specific advantages and drawbacks [44], and which is most appropriate depends on the application and the priorities of the designer.

Clearly, cell balancing is a significant part of the design process for a BESS. Any BMS is omitted in all of the converters shown in figure 1.3 and 1.4, but it is important to not overlook this when considering a system as a whole.

Balancing is not the only concern a designer must consider when utilising a large battery pack, another being degradation. Over their lifetime, the capacity of the battery will decrease, which will also decrease the maximum power the battery can safely deliver. A designer may have accounted for this from the outset and planned for the inevitable degradation to some extent, which the data in table 1.1 shows can be in just a few years, but eventually the batteries will no longer be able to fulfil the system requirements and will have to be replaced. In practice, usually one cell degrades faster than the others, not only limiting overall system performance, but also degrading faster still. As time passes, the most degraded string elements in the system will start to seriously limit overall system performance.

This phenomenon was observed during some preliminary work where the converter shown in figure 1.4b was constructed (see Appendix A), and some data from its operation is shown in figure 1.5. Figure 1.5a shows typical current and voltage waveforms while energy flows from the mains into the battery in a satisfactory manner. However, in figure 1.5b under discharge back into the mains there are regular phenomena (highlighted by the yellow circles) where the system fails to meet the target output current. This occurred owing to one among eight 6V VRLA batteries being used having been damaged during some earlier, unrelated long-term testing, thus showing that one bad element can cripple overall system performance – in this example, a converter otherwise shown to be capable of 20A peak not being able to attain 10A despite most of the battery still operating perfectly.

Even the relatively exotic forms of BMS cannot fully address this issue, which is unfortunate as this is a key barrier to the use of second-life electric vehicle (EV) traction batteries in BESS applications [45, 46] – a use case than not only promises to lower the cost of grid energy storage, but

also reduces the environmental impact of the increasing numbers of large, end of life EV battery packs.

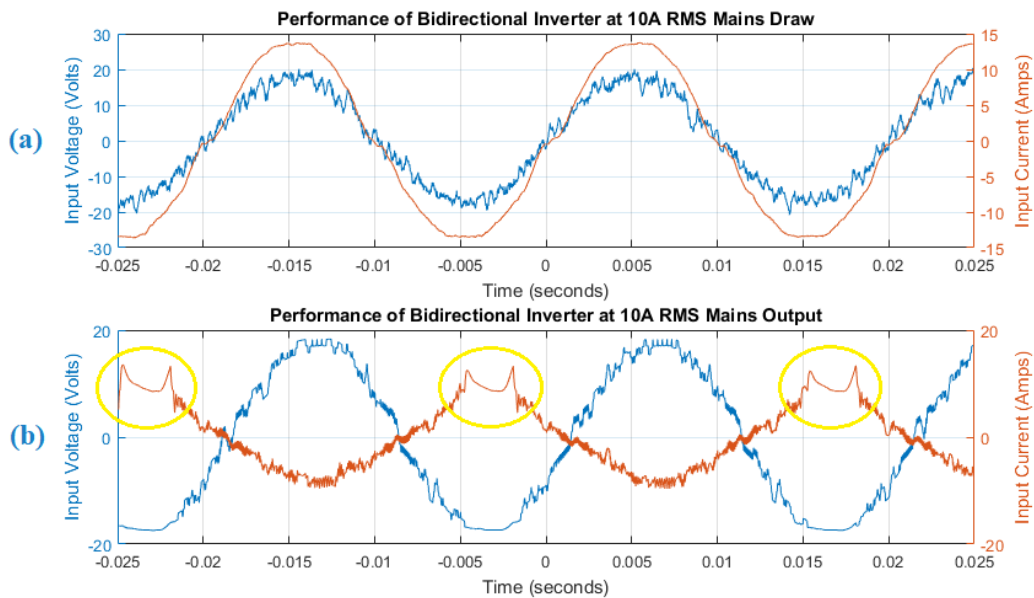


Fig. 1.5 Behaviour of a (poorly designed) converter of the design shown in figure 1.4b

### 1.3 Multilevel Converters

There is a potential solution to both the problem of a large, complex BMS and limiting system performance by partially degraded batteries, and that is through the use of a multilevel converter. Multilevel converters currently only see widespread implementation in situations where there is no practical alternative, such as in power converters for high voltage DC (HVDC) power transmission systems [47-49], where there are no semiconductor devices rated for the very high voltages involved. Here, a multilevel converter allows the high voltage to be shared over a large number of series connected devices. There are a range of multilevel converter topologies, the most common being the flying capacitor multilevel converter, the diode clamped multilevel converter and the cascaded H-bridge converter [50, 51]. The cascaded H-bridge topology, unique among these by its requirement for a large number of high power isolated DC supplies, is ideally suited to BESS applications as a subdivided battery easily provides this.

First proposed in 2009 [52], using a cascaded H-bridge multilevel converter for a grid-attached BESS permits an almost entirely arbitrary usage of the battery attached at each bridge within the converter. For reference, an N-level cascaded H-bridge multilevel converter is shown in figure 1.6. This allows for energy conservative string balancing, which significantly reduces the burden on a BMS, and furthermore allows for the avoidance of severely degraded string elements. This can help in overcoming the barriers to the use of, for example, second life EV traction batteries, as well as coping

much better with the inevitable degradation that occurs with long-term use of electrochemical batteries.




Fig. 1.6 A simplified circuit diagram of an  $N^{\text{th}}$  order cascaded H-bridge multilevel converter.

The work presented in [52] specifically considers a system where nickel-based batteries are utilised on a three-phase grid connection at 6.6kV. Furthermore, the system is not tested under conditions where batteries have significantly different initial state of charge (SoC) or very different capacity, instead demonstrating a system where the various battery strings have differing voltages over a relatively narrow range. The system demonstrates that the strings not only trend to equilibrium, but that that equilibrium is maintained by the control methodology. Whilst this system is not highly optimised, for instance 600V transistors are used across a DC link that does not exceed 300V, the principle had been practically demonstrated.

Further work has also demonstrated this, for example in [53] a similar converter is considered, though in this case with lithium batteries and a more typical 415V three-phase connection. Similar to the work in [52], this does not demonstrate the ability of this topology to overcome gross mismatches in state of health and state of charge, and instead focusses on overcoming subtle imbalances and maintaining of equilibrium across the different strings. Gross mismatches in string performance, and the ability of the inverter topology to overcome this, is explored in [54], with a small demonstration of three cascaded bridges where two of the strings consist of 7Ah batteries while the other is only 4Ah. It is shown that with suitable control the SoC of these three different batteries can be maintained while

the system is under load, even with such a heavy mismatch that could be representative of a severely degraded battery element.

The work presented in [55] is a good representative of the current state of the art in research into cascaded H-bridge multilevel converters for BESS applications with string balancing. In this case rather than a single bridge having a small battery element (relative to a non-multilevel solution), each bridge has a single cell thereby eliminating the need for any external BMS – though this has the key disadvantage of failing to fully utilise the blocking voltage of even the lowest voltage rated transistors. This paper does stretch the capabilities of the system in terms of balancing severely mismatched cells, with a 55% difference in initial states of charge and a 45% difference in nominal capacity between the 12 cascaded bridges, showing the nominal voltage of the cells quickly converging to within 5mV and maintaining this state. This paper also discusses superior SoC estimation by accounting for the impact on the measured cell voltage owing to current transients and goes in to great detail regarding how best to implement the balancing algorithm. This shows that a multilevel cascaded H-bridge BESS can handle batteries found in a wide array of conditions in a way that other converter topologies cannot.

Multilevel converters, such as the cascaded H-bridge multilevel converter, can have benefits when applied to other applications. For instance, in motor drives a multilevel converter can significantly reduce total harmonic distortion. This may be desirable in an environment where a large number of drives is creating significant EMC issues [56], or perhaps in a proposed aircraft electric actuator where a minimum of ripple torque is desired[57]. In either case, it goes to further reinforce the point that lower voltage multilevel converters, i.e. multilevel converters outside their traditional ‘HVDC’ use cases, have great potential.

An obvious first question a designer might ask upon exploring the possibility of using a multilevel converter is how many levels should be used. In traditional multilevel converter applications this is easily answered by finding the highest voltage rated part that fulfils other design requirements and comparing that with the total voltage seen across the converter. In these new and novel applications, however, there is no obvious single constraint that can inform a decision in this way, so perhaps the decision process for the aforementioned papers could be explored.

In the case of [52], the battery voltage is selected to be in the range of 200-300V, quoted as being similar to that in an electric vehicle (circa 2009), which given a 6.6kV connection then defines the number of cascaded bridges required – this seems arbitrary as there is no specified need to have a similar battery voltage to that of an electric vehicle. In [53], 8 cascaded bridges are utilised with no discussion as to why. This is similar to [55] in the sense that there is no discussion as to how the decision was made, but in this case 12 cascaded bridges are used. Meanwhile, in [54] there are only three cascaded bridges used, again with no discussion as to why. It would appear that in all of these

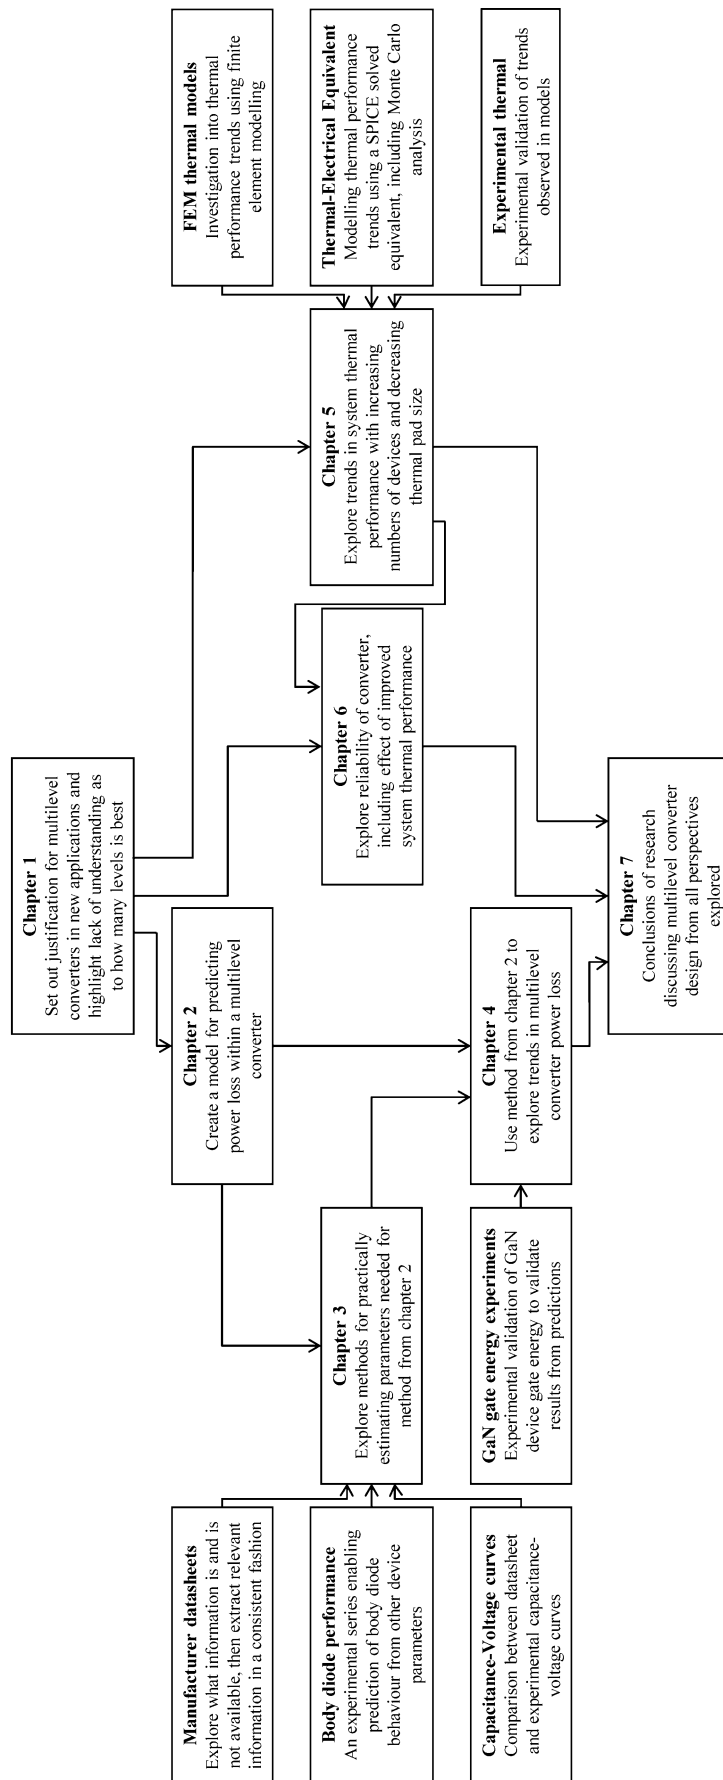
cases, a sufficient number of levels has been selected such that the principle being discussed can be demonstrated, but there appears to be little consideration beyond that. While this is an adequate method for presenting research, it does still leave a potential designer with no way of answering that most basic of questions: how many levels? To address this question, this thesis examines the losses within the converter against the number of levels in an attempt to provide one optimum method by which the number of levels may be selected.

## 1.5 Conclusions

Following this discussion, there are a number of key conclusions that can be drawn. Firstly, electrical energy storage in batteries is going to be an important part of energy grids worldwide in the years to come. Furthermore, the cascaded H-bridge topology shows great hope in addressing some of the big potential issues with grid-attached BESSs, both large and small. That along with the other potentially beneficial applications of multilevel converters (such as low noise motor drives) should adequately demonstrate that a better understanding of multilevel converters will be necessary in the future. Finally, while it may have been proven in an academic context that multilevel converters can be practically implemented and give benefits, the path for a potential designer of a system of this nature is unclear as there are no clear guidelines for choosing the appropriate number of levels. Therefore, it will be the goal of this research to create methods for assessing an optimum number of levels in a multilevel converter, focussing in their use in grid attached BESSs.

To achieve this goal, the work begins in chapter 2, where a method for prediction of power loss in an  $N^{\text{th}}$ -order multilevel converter is described, with chapter 3 going on to look at the practical considerations of finding the data needed for this method. Chapter 4 will use this method to explore the trends in multilevel converter power loss over both switching frequency and multilevel converter order. Chapters 2-4 consider a wide range of switching devices, including state of the art wide bandgap devices such as silicon carbide and gallium nitride. Chapter 5 will explore trends in system thermal performance as the order of a multilevel converter increases, which is incorporated into the work in chapter 6 where the overall reliability of the converter is considered. Chapter 7 will draw together the overall conclusions, highlight the main novel contributions and suggest further work. A flowchart of how the primary bodies of work and the chapters of this thesis are interconnected is included below.

# Thesis Overview





## References for Chapter 1

- [1] Ofgem. (2019, 08/02/2018). *Electricity generation mix by quarter and fuel source (GB)*.
- [2] Office for National Statistics, "UK Energy Statistics, Q1 2018, Statistica Press Release," E. I. S. Department for Business, Ed., ed, 2018.
- [3] United Nations, "Paris Agreement," ed, 2015.
- [4] S. Mercer, "Digest of United Kingdom Energy Statistics (DUKES) 2008," Department of Trade and Industry, Ed., ed, 2007, pp. 129-130.
- [5] Office for National Statistics, "Digest of UK Energy Statistics (DUKES) 2018 Chapter 6: Renewable sources of energy," E. I. S. Department for Business, Ed., ed, 2018.
- [6] Templar. Gridwatch [Online].
- [7] First Hydro Company. (2019, 11/02/19). *First Hydro Company - Dinorwig Power Station*. Available: <https://www.fhc.co.uk/en/power-stations/dinorwig-power-station/>
- [8] D. Peralta, C. Cañizares, and K. Bhattacharya, "Practical Modeling of Flywheel Energy Storage for Primary Frequency Control in Power Grids," in *2018 IEEE Power & Energy Society General Meeting (PESGM)*, 2018, pp. 1-5.
- [9] T. A. Taj, H. M. Hasanien, A. I. Alolah, and S. M. Muyeen, "Dynamic performance enhancement of a grid-connected wind farm using doubly fed induction machine-based flywheel energy storage system," in *IREC2015 The Sixth International Renewable Energy Congress*, 2015, pp. 1-5.
- [10] Q. Kang, X. Xiao, H. Yi, and Z. Nie, "Energy management in grid connected PV systems with supercapacitor energy storage," in *2011 International Conference on Electrical Machines and Systems*, 2011, pp. 1-5.
- [11] X. Tang, W. Deng, and Z. Qi, "Research on Micro-grid voltage stability control based on supercapacitor energy storage," in *2009 International Conference on Electrical Machines and Systems*, 2009, pp. 1-6.
- [12] V. C. Patil, P. I. Ro, and R. K. Ranganath, "End-to-end efficiency of liquid piston based ocean compressed air energy storage," in *OCEANS 2016 MTS/IEEE Monterey*, 2016, pp. 1-5.
- [13] L. Xing, "Integrated market-fit affordable grid-scale energy storage & compressed air energy storage," in *IET Conference on Power in Unity: a Whole System Approach*, 2013, pp. 1-31.
- [14] G. Brett, "Lessons learned developing liquid air energy storage," in *Energy Storage: What's Next for the Grid?*, 2015, pp. 1-24.
- [15] G. Brett and M. Barnett, "Utility-scale energy storage: Liquid air a pioneering solution to the problem of energy storage," in *IET Conference on Power in Unity: a Whole System Approach*, 2013, pp. 1-18.
- [16] S. M. A. S. Bukhari, J. Maqsood, M. Q. Baig, S. Ashraf, and T. A. Khan, "Comparison of Characteristics -- Lead Acid, Nickel Based, Lead Crystal and Lithium Based Batteries," in *2015 17th UKSim-AMSS International Conference on Modelling and Simulation (UKSim)*, 2015, pp. 444-450.
- [17] J. A. McDowall, "Substation battery options: present and future," *IEEE Power Engineering Review*, vol. 20, pp. 4-7, 2000.
- [18] V. A. Boicea, "Energy Storage Technologies: The Past and the Present," *Proceedings of the IEEE*, vol. 102, pp. 1777-1794, 2014.
- [19] C. F. Brasil and C. L. S. Melo, "A comparative study of lead-acid batteries and lithium iron phosphate batteries used in microgrid systems," in *2017 Brazilian Power Electronics Conference (COBEP)*, 2017, pp. 1-7.
- [20] Y. Yang, N. Yang, and H. Li, "Cost-benefit study of dispersed battery storage to increase penetration of photovoltaic systems on distribution feeders," in *2014 IEEE PES General Meeting | Conference & Exposition*, 2014, pp. 1-5.
- [21] A. Dietz, S. Hörlin, and N. Graß, "High voltage battery storage system for multiuse," in *2016 Eleventh International Conference on Ecological Vehicles and Renewable Energies (EVER)*, 2016, pp. 1-6.
- [22] M. Nko and S. D. Chowdhury, "Storage Batteries the Future for Energy Storage," in *2018 IEEE PES/IAS PowerAfrica*, 2018, pp. 705-709.

- [23] C. Curry, "Lithium-Ion Battery Costs and Market," Bloomberg New Energy Finance 5th July 2017.
- [24] Y. S. Wong, L. L. Lai, S. Gao, and K. T. Chau, "Stationary and mobile battery energy storage systems for smart grids," in *2011 4th International Conference on Electric Utility Deregulation and Restructuring and Power Technologies (DRPT)*, 2011, pp. 1-6.
- [25] C. L. Espinosa, "Asynchronous non-inverter buck-boost DC to DC converter for battery charging in a solar MPPT system," in *2017 IEEE URUCON*, 2017, pp. 1-4.
- [26] T. Wu and J. Hu, "Dual-input DC-DC power converter for solar battery charger," in *2016 IEEE 11th Conference on Industrial Electronics and Applications (ICIEA)*, 2016, pp. 1201-1206.
- [27] N. K. Lujara, J. D. v. Wyk, and P. N. Materu, "Power electronic loss models of DC-DC converters in photovoltaic applications," in *IEEE International Symposium on Industrial Electronics. Proceedings. ISIE'98 (Cat. No. 98TH8357)*, 1998, pp. 35-39 vol.1.
- [28] R. Singh, S. F. Taghizadeh, and N. M. L. Tan, "Battery energy storage system for PV output leveling based on bidirectional isolated DC-DC converter," in *2013 IEEE Student Conference on Research and Development*, 2013, pp. 84-89.
- [29] S. Ponnaluri, G. O. Linhofer, J. K. Steinke, and P. K. Steimer, "Comparison of single and two stage topologies for interface of BESS or fuel cell system using the ABB standard power electronics building blocks," in *2005 European Conference on Power Electronics and Applications*, 2005, pp. 9 pp.-P.9.
- [30] M. D. Vijay, I. Hussain, B. Singh, and G. Bhuvaneshwari, "Energy management and control of SECS and BESS integrated AC microgrid," in *2017 IEEE 26th International Symposium on Industrial Electronics (ISIE)*, 2017, pp. 975-980.
- [31] J. G. Pinto, V. Monteiro, H. Gonçalves, B. Exposto, D. Pedrosa, C. Couto, *et al.*, "Bidirectional battery charger with Grid-to-Vehicle, Vehicle-to-Grid and Vehicle-to-Home technologies," in *IECON 2013 - 39th Annual Conference of the IEEE Industrial Electronics Society*, 2013, pp. 5934-5939.
- [32] Y. Wang, F. Qin, and Y. Kim, "Bidirectional DC-DC converter design and implementation for lithium-ion battery application," in *2014 IEEE PES Asia-Pacific Power and Energy Engineering Conference (APPEEC)*, 2014, pp. 1-5.
- [33] S. A. Singh, G. Carli, N. A. Azeez, and S. S. Williamson, "A modified Z-source converter based single phase PV/grid inter-connected DC charging converter for future transportation electrification," in *2016 IEEE Energy Conversion Congress and Exposition (ECCE)*, 2016, pp. 1-6.
- [34] N. Takaoka and J. Itoh, "Battery energy storage system with isolated single-phase matrix converter using center-tapped transformer for power decoupling capability," in *2015 International Conference on Renewable Energy Research and Applications (ICRERA)*, 2015, pp. 1021-1026.
- [35] Z. Guo, C. Zheng, L. Zhao, and X. Qi, "Control strategy study of Bi-directional energy storage converter," in *2017 12th IEEE Conference on Industrial Electronics and Applications (ICIEA)*, 2017, pp. 567-571.
- [36] W.-L. Ming, "Active Control of Voltage Ripple in Power Electronic Converters," PhD thesis, University of Sheffield, 2015.
- [37] B. Lin and D. Chen, "Single-phase neutral point clamped AC/DC converter with the function of power factor corrector and active filter," *IEE Proceedings - Electric Power Applications*, vol. 149, pp. 19-30, 2002.
- [38] Aeroflex, "AN8635-1 An Introduction Cell Balancing and Monitoring Lithium Ion Batteries," 2008.
- [39] Y. Barsukov, "Battery Cell Balancing: What to Balance and How " Texas Instruments 2004.
- [40] F. Altaf, L. Johannesson, and B. Egardt, "Simultaneous Thermal and State-of-Charge Balancing of Batteries: A Review," in *2014 IEEE Vehicle Power and Propulsion Conference (VPPC)*, 2014, pp. 1-7.
- [41] S. Nejad, D. T. Gladwin, M. P. Foster, and D. A. Stone, "Parameterisation and online states estimation of high-energy lithium-titanate cells," in *IECON 2017 - 43rd Annual Conference of the IEEE Industrial Electronics Society*, 2017, pp. 7660-7665.

- [42] J. Cao, N. Schofield, and A. Emadi, "Battery balancing methods: A comprehensive review," in *2008 IEEE Vehicle Power and Propulsion Conference*, 2008, pp. 1-6.
- [43] J. Qi and D. D.-C. Lu, "Review of battery cell balancing techniques," in *2014 Australasian Universities Power Engineering Conference (AUPEC)*, 2014, pp. 1-6.
- [44] M. Daowd, N. Omar, P. V. D. Bossche, and J. V. Mierlo, "Passive and active battery balancing comparison based on MATLAB simulation," in *2011 IEEE Vehicle Power and Propulsion Conference*, 2011, pp. 1-7.
- [45] C. Koch-Ciobotaru, A. Saez-de-Ibarra, E. Martinez-Laserna, D. Stroe, M. Swierczynski, and P. Rodriguez, "Second life battery energy storage system for enhancing renewable energy grid integration," in *2015 IEEE Energy Conversion Congress and Exposition (ECCE)*, 2015, pp. 78-84.
- [46] E. Martinez-Laserna, E. Sarasketa-Zabala, I. V. Sarria, D. Stroe, M. Swierczynski, A. Warnecke, *et al.*, "Technical Viability of Battery Second Life: A Study From the Ageing Perspective," *IEEE Transactions on Industry Applications*, vol. 54, pp. 2703-2713, 2018.
- [47] M. Haeusler. HVDC Solutions for Integration of the Renewable Energy Resources [Online].
- [48] J. Rodriguez, S. Bernet, B. Wu, J. O. Pontt, and S. Kouro, "Multilevel Voltage-Source-Converter Topologies for Industrial Medium-Voltage Drives," *IEEE Transactions on Industrial Electronics*, vol. 54, pp. 2930-2945, 2007.
- [49] F. Hahn, M. Andresen, G. Buticchi, and M. Liserre, "Thermal Analysis and Balancing for Modular Multilevel Converters in HVDC Applications," *IEEE Transactions on Power Electronics*, vol. PP, pp. 1-1, 2017.
- [50] A. K. Koshti and M. N. Rao, "A brief review on multilevel inverter topologies," in *2017 International Conference on Data Management, Analytics and Innovation (ICDMAI)*, 2017, pp. 187-193.
- [51] I. D. Pharne and Y. N. Bhosale, "A review on multilevel inverter topology," in *2013 International Conference on Power, Energy and Control (ICPEC)*, 2013, pp. 700-703.
- [52] L. Maharjan, S. Inoue, H. Akagi, and J. Asakura, "State-of-Charge (SOC)-Balancing Control of a Battery Energy Storage System Based on a Cascade PWM Converter," *IEEE Transactions on Power Electronics*, vol. 24, pp. 1628-1636, 2009.
- [53] M. Vasiladiotis and A. Rufer, "Balancing control actions for cascaded H-bridge converters with integrated battery energy storage," in *2013 15th European Conference on Power Electronics and Applications (EPE)*, 2013, pp. 1-10.
- [54] C. M. Young, N. Y. Chu, L. R. Chen, Y. C. Hsiao, and C. Z. Li, "A Single-Phase Multilevel Inverter With Battery Balancing," *IEEE Transactions on Industrial Electronics*, vol. 60, pp. 1972-1978, 2013.
- [55] E. Chatzinikolaou and D. J. Rogers, "Cell SoC Balancing Using a Cascaded Full-Bridge Multilevel Converter in Battery Energy Storage Systems," *IEEE Transactions on Industrial Electronics*, vol. 63, pp. 5394-5402, 2016.
- [56] S. Sane, S. Sharma, and S. K. Prasad, "Implementation of modular multilevel inverter in local trains of Mumbai," in *2017 2nd IEEE International Conference on Recent Trends in Electronics, Information & Communication Technology (RTEICT)*, 2017, pp. 132-136.
- [57] K. W. Sun, T. J. Summers, C. E. Coates, and C. J. Starkey, "Implementation of a multilevel cascaded H-Bridge drive for an Ironless Axial Flux Permanent Magnet motor for electric vehicle applications," in *2017 19th European Conference on Power Electronics and Applications (EPE'17 ECCE Europe)*, 2017, pp. P.1-P.9.

# Chapter 2:

## A Method for Predicting Converter Power Loss

In the previous chapter it was concluded that there are clear potential benefits to be gained through the use of multilevel converters in settings where one might not otherwise realise; be it string balancing in grid-attached battery energy storage systems (BESS), or lower distortion and noise in motor drives. It was also concluded that there is no clear guidance on how to answer one of the first questions to be asked in the design of multilevel converters: how many levels? Furthermore, it was shown that there is little discussion as to the effect moving to a multilevel solution might have on more general system performance metrics, such as power loss, cost, thermal performance, etc. Therefore, it is the goal of this research to try and fill in this apparent void in knowledge, which might facilitate the practical implementation of multilevel converters in more areas.

There are a number of factors that must be considered in selecting the optimal design of a converter, most with a high degree of interdependence. The cost, mass and size of the system are obvious examples of this, with it being very difficult to determine these in isolation as such parameters are clearly dependant on just about every design decision made throughout the complete design. As a start to this multidimensional analysis, the process began with one low-level parameter that is mostly independent from others: power loss in the converter related to the switching device characteristics.

This study will exclude power loss in the main system inductor, as the inductor design would again vary dependant on a number of factors with a non-trivial relationship. It also excludes power loss in ancillary systems such as cooling and control, as they are similarly dependant on numerous factors that will not be known until more work has been done. This initial method only includes loss in the switching devices and the systems immediately and exclusively related to the switching devices, such as the gate drivers.

Total system power loss will be evaluated by considering contributions from every individual source of loss in the system and then adding them together, under specific design constraints and operating conditions that are described in the forthcoming chapter. While there are many factors in calculating the power loss, the optimisation is trivial, being univariate in this case. That is to say, the optimum under any given set of conditions is simply the one with the lowest total power loss.

There is a strong focus on maintaining a practical approach, so the method must only use data that is readily available to anyone, such that it could be readily implemented in practice. Performance data on a large set of commercially available devices will be compiled, and then the power loss can be calculated under the chosen operating conditions. The emphasis on maintaining a practical approach means the only viable data source will be manufacturer datasheets – parameters that can only be found experimentally should not be used, as this places an unreasonable burden on a potential system designer. It was found that any required parameters that are absent from datasheets can be inferred from information that is available, though the details on this are discussed in chapter 3.

## 2.1 The Reference Converter

While multilevel converters can come in all shapes and sizes, to provide focus for this work, a single topology is chosen and a specific initial rating. This therefore narrows the switching device dataset, reducing the range of devices to the point of being manageable, by limiting the range of acceptable drain current ratings. This aims to provide an example of the approach, which would be applicable to other topologies and ratings, with variations to the dataset of devices chosen.

As this research was spawned from research into grid-tie BESS, a grid-tie BESS was used as the reference converter. Of the numerous types of multilevel converter that exist, such as diode clamped and flying capacitor, only the cascaded H-bridge multilevel converter easily translates to a BESS application, as this converter requires individual DC sources which may be derived from battery modules which make up the BESS, and as a result the reference converter will be of this configuration. To illustrate the topology, an  $N^{\text{th}}$ -order cascaded H-bridge multilevel converter is shown in figure 2.1.




Fig. 2.1 A simplified circuit diagram of an  $N^{\text{th}}$ -order cascaded H-bridge multilevel converter for grid-tie battery energy storage applications.

To carry out the topology optimisation a specification is required to inform the required voltage and current rating of the switching devices. Therefore, AC and DC side voltages must be defined, as well as the rated system power which will allow for inference of system rated current. The selected values are as follows for the reference application:

- 230V 50Hz single phase grid connection,
- 500V nominal DC link (i.e. total battery voltage),
- 6kW power capacity.

The grid connection specification is due to this being the UK and EU standard grid connection voltage at single phase. A single-phase connection is suitable as a three-phase converter is (beyond control) little more than three identical converters in the same unit, i.e. there is no additional complexity to be considered in the three-phase case, so there is no need to consider it immediately. The 500V nominal DC link is derived from considering the worst-case peak grid voltage (~350V), and assuming a typical lithium cell technology a 500V DC link corresponds to a minimum voltage of ~400V (for instance a LiFePO<sub>4</sub> cell has maximum voltage of 4.1V and a minimum of 2.9V [1]). This overhead of 50V is enough to allow for dynamic string avoidance (as discussed in chapter 1) under worst case conditions in all but the lowest order multilevel converters. The 6kW power rating seems reasonable for a BESS appropriate for the domestic setting, being approximately equivalent to a medium-power electric vehicle charger, therefore potentially complementing vehicle-to-grid technologies that are likely to form part of the future home energy system as the smart grid becomes more mature. In addition, this rating is above the maximum domestic solar installation rating of 4kW and covers the power rating of most of the appliances in a domestic setting (except, for example, a 10kW shower).

## 2.2 Loss Metrics

To create a figure for system power loss, we shall consider a variety of sources of system loss in turn, the sum of these being the figure for total loss for a given set of conditions. Specifically, the sources of loss considered for the semiconductor devices within the converter are:

- switching device on-state resistance,
- transient loss in the gate,
- transient drain-source or ‘output’ loss,
- gate driver losses, both transient and quiescent,
- power losses in the diodes.

These will be considered in turn. For the purposes of this chapter, it is assumed that all quantities required are readily available from datasheets, though chapter 3 will discuss in great detail the variety of ways in which this is not necessarily true. The analysis will initially only consider MOSFETs, but other power electronic device technologies will be considered later.

### 2.2.1 On-State Resistance

All MOSFETs have a finite on-state resistance between drain and source,  $R_{DS,ON}$ . This is a dominant source of loss in many applications, particularly at low switching frequencies. Figures for  $R_{DS,ON}$  are available on any and all manufacturer datasheets.

Figure 2.1 shows the configuration of a cascaded H-bridge converter, and in operation either the switching device pair Qx-1 and Qx-3 are conducting or the pair Qx-2 and Qx-4 are conducting, where  $x$  in the range 1-N. This is omitting dead time, but as dead time forms a small part of the overall switching period this was deemed reasonable. Therefore, it can be said that there are  $2N$  devices in the conduction path, where  $N$  is the number of cascaded bridges in the converter, also referred to as its ‘order’. Converter order,  $N$ , is distinct from the number of output levels the converter is capable of output, that we shall call  $n$  – output levels being the number of distinct output voltage levels the converter can achieve. They are related by the expression:

$$n = 2N + 1 \quad (2.1)$$

Knowing the number of devices in the conduction path and their individual resistance, the only other value required to calculate loss is the RMS current. This is found from the grid voltage connection and the RMS power rating, as given in the converter specification. Therefore, the power loss due to on-state conduction is:

$$P_{R_{DS,ON}} = 2NI_{RMS}^2 R_{DS,ON} , \quad I_{RMS} = \frac{P_{RATED}}{V_{MAINS,RMS}} \quad (2.2)$$

### 2.2.2 Transient Loss in the Gate

This is a measure of the power dissipated in the gate of switching devices in the power converter, and is found by estimating the charge-voltage curve during turn-on, which is related to the energy lost at the gate during a single cycle by:

$$E_{GATE} = \int_0^{\sum Q_G} V_{GS}(Q_G) dQ_G \quad (2.3)$$

Figure 2.2 shows a typical gate charge-voltage curve for a MOSFET (or IGBT for that matter). The gradients of the curve in figure 2.2 from 0 to  $Q_1$  and from  $Q_1+Q_{MILLER}$  to  $\sum Q$  are easily derived from the device datasheet, as they can be determined from the input capacitance,  $C_{iss}$ . However, this is not the common terminology used on manufacturer datasheets, instead using the terminology: input capacitance,  $C_{iss} = C_{GS} + C_{GD}$  (measured with drain shorted to source); output capacitance,  $C_{oss} = C_{DS} + C_{GD}$  and the feedback capacitance,  $C_{rss} = C_{GD}$ . It should be noted that  $C_{iss}$  varies with respect to the drain-source voltage, as do the other pin-to-pin capacitances  $C_{oss}$  and  $C_{rss}$ .

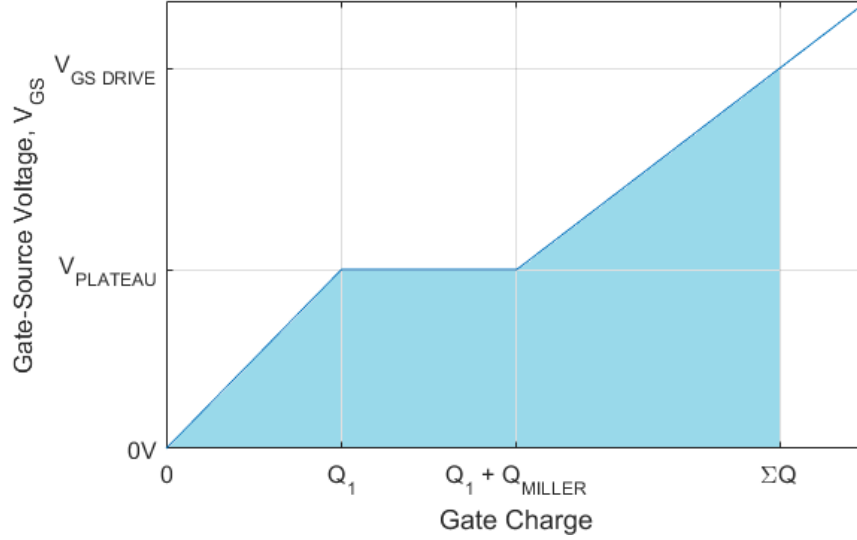


Fig. 2.2 A typical MOSFET gate charge-voltage curve, with some key values annotated.

The flat region in the centre of the plot in figure 2.2 (as well as figure 2.4) is known as the Miller shelf [2]. Here, the gate-source voltage remains constant while the transistor turns on, during which time the gate-drain capacitance charges through the gate. The Miller charge is calculated from the integral of the feedback capacitance  $C_{rss}$  and the output capacitance  $C_{oss}$ , with respect to the drain-source voltage. As such, the Miller Charge is (non-linearly) correlated to the maximum drain-source voltage across which the device is switching.

The calculations are further complicated as the maximum drain-source voltage,  $V_{DS,MAX}$ , varies sinusoidally with time as a result of the AC grid connection. This is accounted for below:

$$Q_{MILLER,AV}(V_{DS}) = \frac{\int_0^{\pi/2\omega} Q(V_{DS,MAX} \sin(\omega t)) dt}{\pi/2\omega}, \omega = 2\pi f_{grid} \quad (2.4)$$

The lack of an RMS calculation in equation 2.4 is acceptable as it evaluates over the first quarter cycle of the utility supply waveform, from  $t=0$  to  $t=\pi/2\omega$ . During this period of the output, as the signal is strictly positive and increasing, this integral and the standard RMS calculation yield identical results, and equation 2.4 is significantly easier to work with than a true RMS calculation.

With the charge-voltage curve calculated, and, by extension, the energy dissipated in the gate in a single switching cycle as in equation 2.3, it is trivial to extend the energy dissipation calculation to the total power dissipated into the gates of the MOSFETs throughout the converter. Since only one bridge switches at any one time and each of the four MOSFETs in the H-bridge goes through a ‘turn on’ once per  $f_{switching}$  cycle, the total power dissipated in the gates is:

$$P_{GATE} = 4E_{GATE}f_{switching} \quad (2.5)$$



### 2.2.3 Transient Output Loss

Output loss is the power that is expended in charging the capacitance between the drain and the source,  $C_{oss}$ , during ‘turn on’ and ‘turn off’ of the switching device. The integral of capacitance with respect to voltage yields charge, and the integral of charge with respect to voltage yields energy, therefore the energy dissipated during a single switching event is given by:

$$E_{OUT} = \iint_0^{V_{DS,MAX}} C_{oss}(V_{DS})d^2V_{DS} = \int_0^{V_{DS,MAX}} Q_{oss}(V_{DS})dV_{DS} \quad (2.6)$$

However, as  $V_{DS}$  is sinusoidally time varying, equation 2.6 needs to be adjusted accordingly in a similar fashion to equation 2.4:

$$E_{OUT,AV} = \frac{\int_0^{\pi/2\omega} E_{OUT}(V_{DS,MAX} \sin(\omega t))dt}{\pi/2\omega}, \quad \omega=2\pi f_{grid} \quad (2.7)$$

The total power can be calculated from the average energy dissipated in a single cycle, as shown in equation 2.7, multiplied by the switching frequency and the number of output loss events per switching period. As previously mentioned, only one bridge switches during a single  $f_{switching}$  cycle, but each device in the bridge goes through both a charge and discharge cycle, yielding the expression:

$$P_{OUT} = 8E_{OUT,AV}f_{switching} \quad (2.8)$$

### 2.2.4 Gate Drive Loss

A generic gate drive was considered to derive an expression for loss in the gate drive. Figure 2.5 shows the gate drive topology used, with M1 being the main power MOSFET. The labels ‘SYSTEM A’ and ‘SYSTEM B’ shown in figure 2.5 denote where the main power MOSFET connects to the rest of the converter. The control signal labelled ‘CONTROL’ can turn on or off the phototransistor within the optocoupler labelled ‘OPTO’ as required. This in turn controls the power stage Q1 and Q2 to either pull the gate high, to  $V_G$ , or to ground (relative to the source of M1), through the gate resistor  $R_G$  for current limiting.




Fig. 2.3 Circuit diagram of the gate drive considered for the method.

For simplicity, an optocoupled solution with an isolated DC-DC converter was used instead of a transformer isolated gate drive. The losses in a transformer isolated gate drive would be different, but as the overall effect on converter loss is unlikely to be large this is not considered a significant issue (this is shown to be true later in chapter 4). This decision does not affect the validity of the comparative results of this research, but it serves as an example of a way in which the absolute results may not have a great degree of accuracy, while the relative results maintain their validity.

To compute the losses in the gate drive, transient driver losses and quiescent losses will be calculated separately.

#### 2.2.4.1 Transient Gate Drive Loss

Transient gate drive loss is calculated with reference to figure 2.2. The highlighted area under the curve is the energy dissipated in the gate ( $E_{GATE}$ ), while the product of the drive voltage and the total charge is the total energy being put in by the gate drive. Therefore, the transient energy lost in the gate drive, i.e. the area above the curve bounded by the gate drive voltage, is the difference between the two:

$$E_{DRIVE,T} = (C_{ISS}V_{GS,DRIVE} + Q_{MILLER,AV})V_{GS,DRIVE} - E_{GATE} \quad (2.9)$$

This is correct in principle, however it omits the inherent gate resistance in the physical device. This is a figure that datasheets often provide, but it is little more than a ‘ballpark’ figure. In reality it is impossible to know precisely what the gate resistance is in advance, and it is very difficult to measure. As a result, the separation between transient gate drive loss and transient loss in the gate is not strictly as described here, with a greater, but unspecified, amount being dissipated in the gate and less in the driver. This is a moot point however, as in a practical calculation the total gate charge ( $\Sigma Q$  in figure 2.2) is multiplied by the gate drive voltage ( $V_{GS,DRIVE}$  in figure 2.2), and that is considered as the sum of the transient energy lost in both the gate and the driver in a single switching cycle. The disadvantage is that predicting exact power loss in the device, needed for thermal calculations for instance, has this limitation. Thankfully, it is later shown that gate dissipation losses are almost negligible compared to other sources further limiting the impact of this concession in the method.

To calculate power loss per module, energy is then multiplied by four, for each switching device operating during a switching cycle, and again multiplied by the switching frequency. Equation 2.10 also incorporates  $\eta$ , which is a measure of efficiency of the isolated DC supply for the gate drive,  $V_G$  in figure 2.3, which, while load dependent, is typically 75%, as found in datasheets [3].

$$P_{DRIVE,T} = 4E_{DRIVE,T}f_{switching}/\eta \quad (2.10)$$

### 2.2.4.2 Quiescent Gate Drive Loss

To calculate the quiescent power dissipation in the gate drive circuit, the component values must be found, and therefore the peak current requirement of the driver must be determined. To find the peak current requirement, the maximum permissible time to perform a single switching operation must be found, as faster switching for a given gate charge profile would require higher current – moving a given charge in a smaller time period means higher current.

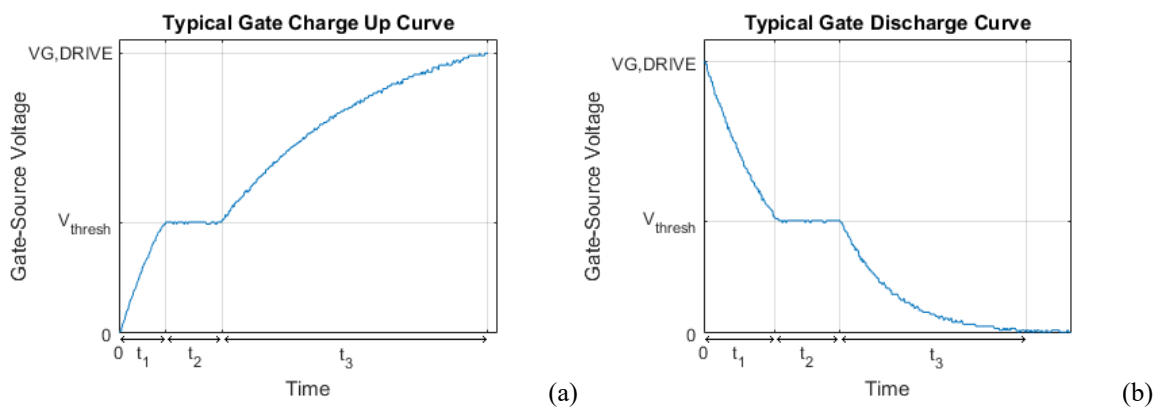


Fig. 2.4 A plot of the turn-on (a) and turn-off (b) behavior of a typical MOSFET with respect to time, labelled with key values. This is experimentally derived data for the Fairchild FCH47N60.

The turn-on behaviour of a MOSFET is shown in figure 2.4a. The device has turned on fully at time  $t_2$ . The time before  $t_1$  is an exponential relationship dictated by the relationship between the input capacitance  $C_{iss}$  and the gate resistance  $R_g$ . Between  $t_1$  and  $t_2$  is the Miller Shelf. As part of the gate dissipation calculations, the Miller Charge has already been calculated, and with a given plateau voltage and gate resistance the time can be found.

During the period 0 to  $t_1$ , there is a classic resistor-capacitor exponential charge between  $C_{iss}$  and  $R_g$ , and in the period  $t_1$  to  $t_2$  the voltage is fixed. The expressions in these two periods being:

$$t_{ON} = t_1 + t_2, \quad t_1 = -\ln\left(1 - \frac{V_{PLATEAU}}{V_{GS,DRIVE}}\right) C_{iss} R_g, \quad t_2 = \frac{Q_{MILLER} R_g}{V_{PLATEAU}} \quad (2.11)$$

Maximum switching time is calculated from the resolution of the PWM occurring and the fundamental switching frequency of the converter. For example, if the converter is running with a 100kHz switching frequency and 8-bit PWM, the minimum time base is  $T/2^8 \approx 40\text{ns}$ , where  $T$  is the minimum time increment possible for an 8-bit PWM clock. An estimation of gate drive peak current to attain this speed is then easily derived from information readily available about the device and the application. However, this method proved too restrictive, as many devices that should have been capable were deemed to be too slow – therefore the switching period constraint was relaxed. This results in a more reasonable design constraint but at the cost of slightly higher harmonic distortion.

An effort was made to find a more concrete method for deciding dead time. However, it appears to be an area where a certain degree of estimation is expected. As previously mentioned, it stands to reason that as dead time increases it is likely to increase distortion in the power converter output, and literature was explored in an attempt to quantify this. The research presented in [4] does show the impact of dead time when compared with having none in one case, even showing how the impact on total harmonic distortion (THD) varies with modulation depth, but the dead time selected is not justified and nor is there any trend presented with regards to the impact of dead time. Meanwhile, [5] does show a trend of the impact of dead time, but only on the effectiveness of a specific control algorithm, rather than the converter's performance at large, or converter losses. In [6], potential methods for compensating for the impact of dead time in popular control algorithms are discussed, but with no useful quantifying of the impact it has on loss in the first place. There is also significant research on the impact of dead time on THD in class D (switching) audio power amplifiers [7, 8], but these results are not easily translated into the context of a grid-connected power converter.

Current literature yields no feasible method for predicting the harmonic distortion in a generalised case, and dead time's impact upon it. Therefore, increasing the switching period constraint by a factor of three, as described previously, was found to yield credible results with respect to devices being capable of high switching frequencies or not, so this estimation was used for the calculations. This is a worthwhile compromise for the generation of this generalised, comparative metric as it will be a consistent approximation across the device comparison.

Equation 2.11 shows how to calculate the turn-on time. The turn-off time is computed similarly, though as can be seen in the discharge curve in figure 2.6,  $t_1$  will be different due to a larger voltage swing occurring. The edge condition for this being:

$$(t_{ON} + t_{OFF})(1.2t_{dead}) = \frac{3}{f_{switching} 2^{N_{PWM\ res}}} \quad (2.12)$$

The sum of turn-on and turn-off time,  $t_{ON}$  and  $t_{OFF}$ , with the addition of the dead time,  $t_{dead}$  (inflated by 20% as a safety margin), permits the calculation of the required gate resistor,  $R_g$  (see figure 2.5). The peak current requirement of the gate drive is then calculated from the gate resistor  $R_g$  and the peak drive voltage  $V_{GS,DRIVE}$ , i.e.  $I_{MAX} = V_{GS,DRIVE}/R_g$ .

To calculate the quiescent loss of the gate driver, i.e. power loss that does not occur as a direct result of the switching transient, from the calculated peak current requirement requires inspection of the circuit diagram in figure 2.3. A key source of quiescent loss is  $R_1$ , its value being related to the gate resistor by the gain of the main drive transistors  $Q_1$  and  $Q_2$ . For example, if  $Q_1$  and  $Q_2$  were to have a nominal current gain of 100,  $R_1$  would be 100 times the size of  $R_g$ . The quiescent power loss in that resistor would then be the  $P_{R1} = V_{GS,DRIVE}^2/R_1$ .

There is also quiescent loss in the isolated DC-DC converter. While every device is a little different the loss tends to be approximately 15% of rated output [3]- this is included in the loss estimation also. Furthermore, losses in the opto-isolator are due to the infrared LED inside, with drive current as high as 30mA [9]. Knowing that on average across the converter the LEDs in the gate drives are on half of the time, and the drive voltage (taken as 5V here), this source of quiescent gate drive loss can be quantified.

While it may seem unnecessary to evaluate and include 120mW in an LED in a 6kW power converter and would often be omitted (perhaps quite reasonably), it was included in this case owing to the sheer number of gate drives that are present in high order power converters. For instance, in the case of 20 cascaded bridges, there are 80 gate drives – making it not insignificant as the number of levels increases.

### 2.2.5 Diode Loss

The power switching devices are not necessarily the only semiconductors in the converter, with diodes also forming an important part of converter design. During the dead time of a switching period, the diodes commutate the ‘free-wheel’ current – the current driving into the load (in this case the grid). MOSFETs have an intrinsic diode from the source to the drain, due to the p-n silicon junction inherent in their design, but these are often not used instead choosing to use higher performance external switching diodes. One of the reasons for this is that the way the body diode is formed on the die makes it unsuitable for sustained current, but in this application it would only be carrying brief pulses of current so that is not a serious problem. The key problem is the lack of data characterising the performance of the body diodes.

Figure 2.5 shows an example of how external switching diodes can be used to disable the intrinsic body diode (labelled e.g. D1-Q1), while still allowing free-wheel conduction. D1-1 forces the body diode D1-Q1 to never conduct by being anti-series, while D1-2 ensures there is still a path for free-wheel. The forward voltage drop of MOSFET devices is low enough that the diode D1-1 is needed, if the body diode had a high forward voltage this diode would be unnecessary as diode D1-2 would clamp the voltage low enough that the body diode could not become forward biased, but that is not the case here.




Fig. 2.5 A single H-bridge with high performance diodes used to avoid intrinsic body diode conduction.

While the addition of these diodes does indeed immobilise the MOSFET body diode, and therefore remove a potential unknown as body diodes are infrequently characterised by the manufacturer, there is a potentially very significant disadvantage, especially in the case of multi-level converters. The diodes enumerated as DN-1 are in the main conduction path along with the main switching devices QN, meaning there will be significant additional loss, with two diodes (one high-side and one low-side) in the conduction path for every bridge in the converter at all times. Even in the case of a single H-bridge, this would have an impact on overall power loss. The question is whether the reduction in power loss attained by removing these additional diodes from the conduction path is overshadowed by the increased power dissipation due to the reverse recovery of the intrinsic diode, as well as their presumably inferior forward conduction performance during dead time.

There are two cases where the body diode has power loss: forward conduction and reverse recovery. Forward conduction is relatively easy to calculate, as current through the diodes during dead time will be the same as the main system current, i.e. the mains current. Knowing the RMS system current, we then need to know the voltage across the diode for a given forward current. Power dissipation during reverse recovery depends on the reverse recovery charge and the voltage across which that charge will transfer. Reverse recovery charge would need to be derived somehow, but the voltage across which that charge must transfer is simply the DC link across the bridge in question. Therefore, to calculate power loss while using the body diode we must know the reverse recovery charge and the diode current-voltage curve – a method for this will be discussed in the next chapter.

## 2.3 Alternative Semiconductor Technologies

### 2.3.1 Insulated Gate Bipolar Transistors (IGBTs)

A relatively modern technology, with reliable IGBT technology only really being realised in the 1990s, they are supposedly a best of both worlds between a MOSFET and a BJT by combining the bipolar output stage of a BJT with the insulated gate, and hence low gate current requirements, of a MOSFET [10]. They are currently extremely popular in many larger power converters: from motor drives to grid-tie battery energy storage systems. They are generally not suited to higher switching frequencies, normally operated below 20kHz, with some devices even stating they are not capable of hard switching at their rated current above 5kHz [11].

The key difference in predicting the power loss of an IGBT and a MOSFET is the on-state conduction loss. While the gate of an IGBT is very similar to a MOSFET, the output of the device behaves in much the same way as a BJT (while maintaining a body diode). As a result, datasheets provide curves for the collector-emitter saturation voltage ( $V_{CE,SAT}$ ) over a range of collector currents. This curve is similar to that of a diode and can be adequately approximated by a voltage drop and a resistance.

### 2.3.2 Silicon Carbide (SiC) MOSFETs

Silicon carbide MOSFETs are a relatively new class of power device, though have comfortably entered industrial applications. Extending the method to include them is very straightforward, being essentially identical to silicon MOSFETs in terms of parameters available and application of the power loss prediction method [12]. One of the key benefits of SiC MOSFETs over conventional silicon devices is their access to significantly higher voltage ratings. However, as increasing the number of levels in multi-level converters lowers the device voltage requirement, they are unlikely to be competitive in the application under consideration.

For a different application the process presented here may be directly applied to these devices. SiC MOSFETs, at least anecdotally, tend to have lower on-state resistance but at the expense of a gate that needs to be driven harder [13]. The body diode will perform differently to conventional silicon equivalents, with much higher forward voltage typically. SiC MOSFETs do not typically exist in voltage ratings lower than 600V, and such are very unlikely to compete in a multilevel case against much lower voltage rated silicon devices, however, they may prove optimal in the single bridge comparison case. To this end, data will be extracted from a number of datasheets and compared with that found for conventional silicon devices.

### 2.3.3 Gallium Nitride (GaN) High Electron Mobility Transistors (HEMTs)

GaN HEMTs are a very new class of power device, with some devices just coming to market that are relevant to this power range. They promise much faster switching and lower on-state resistance, with few apparent disadvantages other than their monetary cost [14]. Despite being very different to a silicon MOSFET in construction, the method for estimating power loss is very similar, with datasheet layouts and available information being almost identical.

One key difference is that the gate drive voltage is much lower, no more than 5V in the case of GaN devices, rather than at least 10V in the case of silicon MOSFETs. The method for predicting gate associated power dissipation remains otherwise identical. The other key difference is in the body diode, GaN devices do not have a body diode in quite the same way as a silicon MOSFET, lacking a P-N junction in their physical construction, but negative bias will raise the gate voltage up to the point the device will conduct [15], which behaves much like a body diode. During ‘body diode forward conduction’ the voltage drop is the threshold voltage of the device along with the product of the drain-source on-state resistance and the current – all known quantities used in other parts of the method and therefore easy to calculate. Better still, the GaN device ‘body diode’ exhibits zero reverse recovery, not only removing the need to include it in the method, but also clearly being of benefit to overall system power loss.



## 2.4 Conclusions

This chapter has outlined a method for calculating power loss in the reference converter, illustrated in figure 2.1, for a given device, for a given number of converter levels. The method also allows for determining if any given device is suitably rated under a given set of conditions (converter order, switching frequency, etc.). The method focusses on silicon MOSFETs, but there is discussion of how this method is applicable to other switching device technologies, including state of the art wide bandgap devices, with varying degrees of modification being required. A simple sum of the various contributors gives total system power loss, and this can be applied to an optimisation in terms of device selection under a given set of conditions.

While the individual elements of this model may not be novel in their own right, the combination of these component parts forms a new method that enables a practical, large scale evaluation of loss in order to explore trends in a way not done before now. The concessions made to practical implementation of this method (which will be expanded on significantly in the next chapter), make this particularly useful in helping solve the core, immediate questions of a potential designer of a system of this type.

There is a focus in this research on creating a practical method, and the following chapter will discuss in great detail how the various quantities this method requires can all be found, inferred or estimated from information available on manufacturer datasheets.

## References for Chapter 2

- [1] Enix Energies, "18650 LIFEPO4 BATTERY 1500mAh 3.2V Part No: ACL9011," ed, 2006.
- [2] P. Horowitz and W. Hill, *The Art of Electronics*, 3rd ed.: Cambridge University Press, 2016.
- [3] Traco Power. DC/DC Converters TMA Series, 1 Watt [Online]. Available: [tinyurl.com/yceokz63](http://tinyurl.com/yceokz63)
- [4] M. I. Jahmeerbacus and M. Sunassee, "Evaluation of Selective Harmonic Elimination and sinusoidal PWM for single-phase DC to AC inverters under dead-time distortion," in *2014 IEEE 23rd International Symposium on Industrial Electronics (ISIE)*, 2014, pp. 465-470.
- [5] A. Khaligh, J. R. Wells, P. L. Chapman, and P. T. Krein, "Dead-Time Distortion in Generalized Selective Harmonic Control," *IEEE Transactions on Power Electronics*, vol. 23, pp. 1511-1517, 2008.
- [6] S. Wang, W. Song, J. Ma, J. Zhao, and X. Feng, "Study on Comprehensive Analysis and Compensation for the Line Current Distortion in Single-Phase Three-Level NPC Converters," *IEEE Transactions on Industrial Electronics*, vol. PP, pp. 1-1, 2017.
- [7] F. Koeslag, H. d. T. Mouton, H. J. Beukes, and P. Midya, "A detailed analysis of the effect of dead time on harmonic distortion in a class D audio amplifier," in *AFRICON 2007*, 2007, pp. 1-7.
- [8] I. D. Mosely, P. H. Mellor, and C. M. Bingham. (1999, 10th June) Effect of dead time on harmonic distortion in class-D audio power amplifiers. *IEEE Electronics Letters*. 950-952.
- [9] Silicon Laboratories, "Silicon Labs Si826x 4A Isolated Gate Drivers," ed, 2013.
- [10] B. K. Bose, "Recent advances in power electronics," *IEEE Transactions on Power Electronics*, vol. 7, pp. 2-16, 1992.
- [11] Infineon. IRG4BC20FPbF Fast Speed IGBT [Online]. Available: [bit.ly/33q6UEM](http://bit.ly/33q6UEM)
- [12] K. M. Speer and S. Banerjee, "State of the SiC MOSFET: Device evolution, technology merit and commercial prospects," Littlefuse2016.
- [13] ROHM Semiconductor, "The Next Generation of Power Conversion Systems Enabled by SiC Power Devices," 2013.
- [14] A. Lidow, "GaN Makes a Frontal Attack on Silicon Power MOSFETs," ed. Bodo's Power Systems, 2019.
- [15] Efficient Power Conversion. (2009, Application Note 002: Fundamentals of Gallium Nitride Power Transistors.

# Chapter 3:

## Practical Parameter Estimation

In the previous chapter, a series of expressions were derived to model losses in an  $N^{\text{th}}$ -order cascaded H-bridge multilevel converter for a medium power grid-tie battery energy storage system. A sum of these loss components permits the evaluation of total power converter loss based on a wide range of parameters, and therefore a broad design optimisation can be performed. However, prior to the optimisation, a dataset of real-world devices and their relevant performance characteristics is required. This chapter will discuss the ways in which this is non-trivial, and outline the methods used to generate this dataset.

Core to the estimation and optimisation is the creation of a dataset of real devices with all the parameters the optimisation model requires. Unfortunately, there are a number of challenges to the extraction of useful parameters from device datasheets. The vast majority of devices considered in this thesis are silicon MOSFETs, and while there are loose conventions on what information manufacturers provide, even among the relatively narrow subset of switching devices considered, the layout of these documents varies significantly, preventing automatic data extraction and resulting in a time-consuming, human driven data extraction process. While many manufacturers provide online datasets with all of their devices and their key performance metrics, not only is this not universal or consistent, but the expressions derived in the previous chapter requires more data than manufacturers commonly provide in this format.

In the derivation of the model there are a number of performance parameters that were extracted, some of which are straightforward, some of which less so. The datasheets parameters extracted are:

- Maximum drain-source voltage rating,  $V_{\text{DS,MAX}}$
- Maximum drain current rating,  $I_{\text{D,MAX}}$
- On-state drain-source resistance,  $R_{\text{DS,ON}}$
- Gate-source plateau voltage,  $V_{\text{PLATEAU}}$
- Pin-to-pin device capacitances
- Unit cost
- Technology (e.g. Si, SiC, GaN)
- Body diode performance characteristics

While some of these parameters are unambiguous, such as the maximum drain-source voltage rating, and what semiconductor technology is used, others are much less so, with the ambiguity of the available parameters discussed below:

### 3.1 Slightly Ambiguous Parameters

Maximum drain current rating is a good example of a slightly ambiguous parameter. On any MOSFET supplier website, a user can sort by  $I_{D,MAX}$  and find a value taken from the first page of a device datasheet. However, maximum drain current rating varies inversely with temperature – the higher the temperature, the lower the rating. Figure 3.1a illustrates this trend for a specific device. The headline current rating figures presented by manufacturers follow no fixed standard, with different manufacturers choosing different temperatures at which to quote this figure. While most quote at room temperature (20-25°C), it is not universal and must therefore be accounted for in the extraction of data. All datasheets provide information on how  $I_{D,MAX}$  varies with temperature, so a temperature at which to extract drain current rating was chosen as 80°C for this work, and all devices will have  $I_{D,MAX}$  extracted from the datasheet under these conditions.

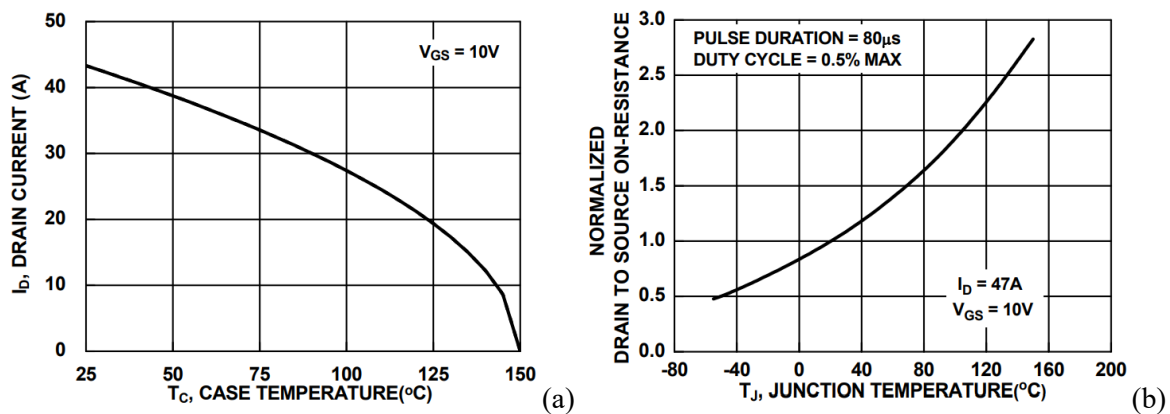


Fig. 3.1 Curves from the datasheet for the ON Semiconductor/Fairchild FCH47N60 [1] - (a) showing maximum drain current and (b) showing on-state resistance, both varying with device temperature.

On-state drain-source resistance is also related to junction temperature, though this time they are positively correlated, with a higher temperature leading to a higher on-state resistance, as shown in figure 3.1b. Once again, the on-state resistance at 80°C was chosen here as the representative value.  $R_{DS,ON}$  also varies with the gate source voltage used in the driver. This was selected at the outset to be 10V, a figure whose selection has an impact on other parts of the model also. This value was selected as for all MOSFETs this is a sufficiently high voltage for the device to be fully on, and at or very near minimum  $R_{DS,ON}$ .

The junction temperature under operation could be any value, over a wide range, with 80°C representing the low end of what that might be common. Although on-state resistance varies with junction temperature, junction temperatures varies with power dissipation (among other factors), and power dissipation varies with on-state resistance. Parameter co-dependence is therefore a significant obstacle, and to allow the development of this model, without a complete model for the entire system being a pre-requisite, an initial fixed operating temperature was taken as a first approximation.

Clearly, as the total system model nears completion other relevant system properties, such as thermal solution performance, could be integrated into this part of the model.

Unit cost is also not entirely unambiguous, as cost varies based on volume, supplier, geographic location, etc. Again, as this method focusses on relative accuracy rather than absolute accuracy, a consistent method being used for all devices is adequate provision. Costs were taken from large UK suppliers (Mouser/Farnell/RS), for quantities in the order of 500-2000 units (different products have different breakpoints, hence the range), with cost data compiled in November 2016 in GBP.

### 3.2 Pin-to-pin device capacitances

The pin-to pin capacitances of a MOSFET vary substantially with the drain-source voltage, with this relationship documented in every MOSFET datasheet. As the integrals of these capacitances with respect to the drain-source voltage forms an important part of the method developed here, extraction of these curves is essential. This has some practical challenges, as discussed below.

The pin-to-pin capacitances are: the drain-source capacitance,  $C_{DS}$ , the gate-drain capacitance,  $C_{GD}$ , and the gate-source capacitance,  $C_{GS}$ . These are not commonly available on manufacturer datasheets, instead: input capacitance,  $C_{iss} = C_{GS} + C_{GD}$  (measured with drain shorted to source); output capacitance,  $C_{oss} = C_{DS} + C_{GD}$  and the feedback capacitance,  $C_{rss} = C_{GD}$ , are more usually found. This terminology conveniently collects the pin-to-pin capacitances in to the groups that they would usually be used in for any sort of calculation, and the method used here is not an exception. This convention is completely universal across all datasheets seen during this work.

The practical difficulty comes from the extreme non-linearity found in the relationship between these quantities and the drain-source voltage,  $V_{DS}$ . Figure 3.2 shows the curves relating these capacitances with the drain-source voltage for four different but typical devices, extracted from a range of datasheets. Not only is the relationship non-linear, but they are all non-linear in different ways which makes for a challenge in precisely extracting these curves in large numbers to form a data set.

A technical approach was explored at first, using optimal character recognition (OCR) and edge detection to try and extract the data. This proved prohibitively difficult, as there is so much inconsistent clutter in these various graphs, such as the labels over plotted lines (as can be seen in figure 3.2), that this method was quickly abandoned. Therefore, and regrettably, the data must be extracted from these graphs manually.

These curves are so irregular that a precise manual extraction process would be time prohibitive when hoping to extract data from tens of devices. A simplification was required, and while not necessarily the same for each type of capacitance, it must be consistent across devices. A bilinear approximation was selected as a balance between quality of estimation and ease of data entry process

in the case of  $C_{oss}$  and  $C_{rss}$ . An example of this bilinear simplification can be seen in figure 3.3. This simplification means that only three figures need extracting for  $C_{oss}$  and  $C_{iss}$ : maximum capacitance, minimum capacitance and ‘corner voltage’. In the case of  $C_{iss}$ , a simple constant value seemed an adequate approximation. Therefore, just a total of four numbers are inferred from the manufacturer datasheets by a human operator for each device considered and entered into the dataset.

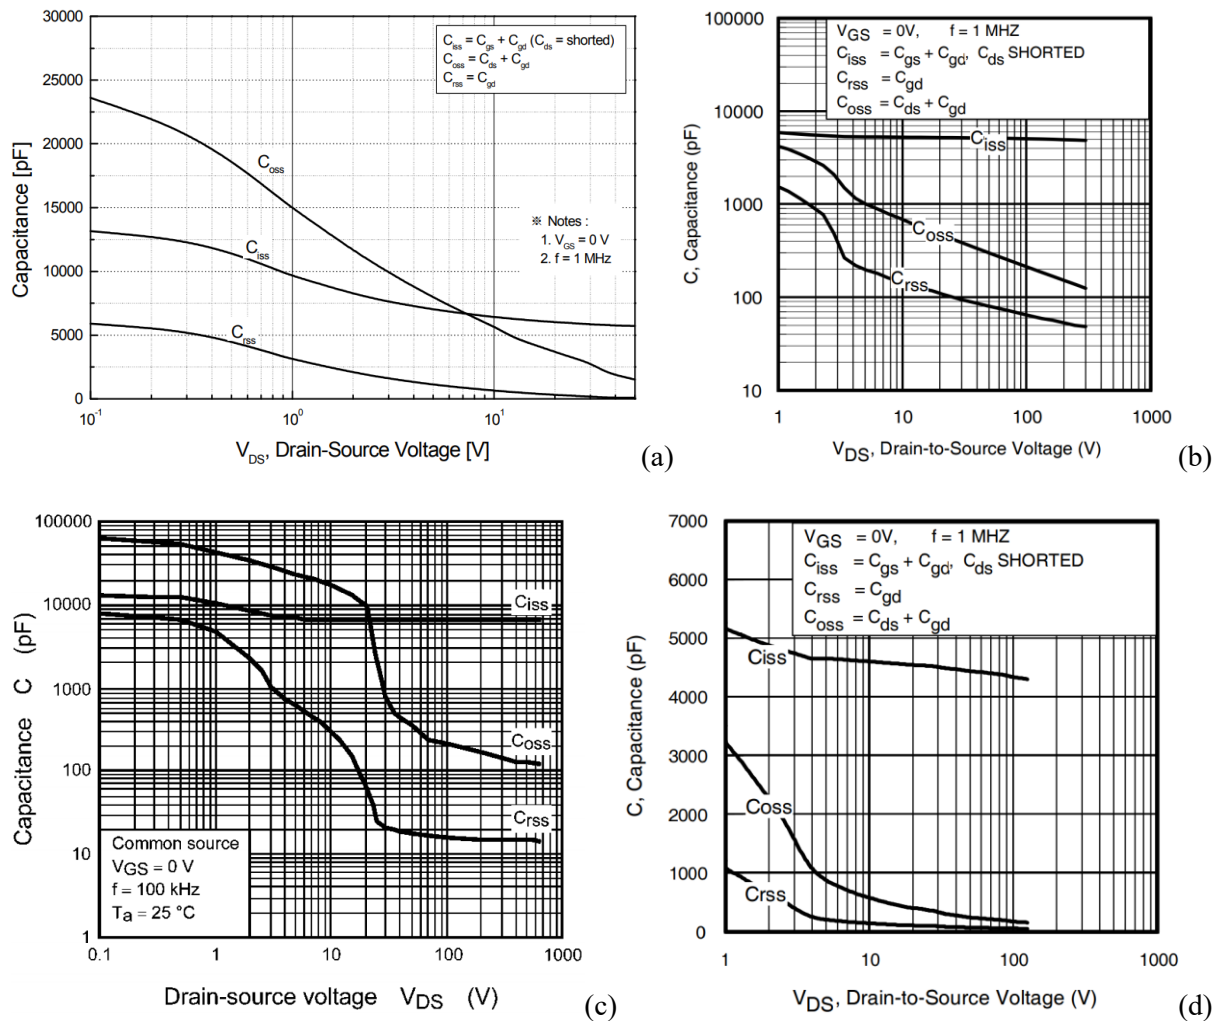


Fig. 3.2 Curves from manufacturer datasheets showing variation in device capacitances with drain-source voltage for (a) Fairchild/On Semiconductor FCH47N60 [1], (b) Infineon AUIRFP4409 [2], (c) Toshiba TK49N65W [3], and (d) International Rectifier IRFP4229 [4].

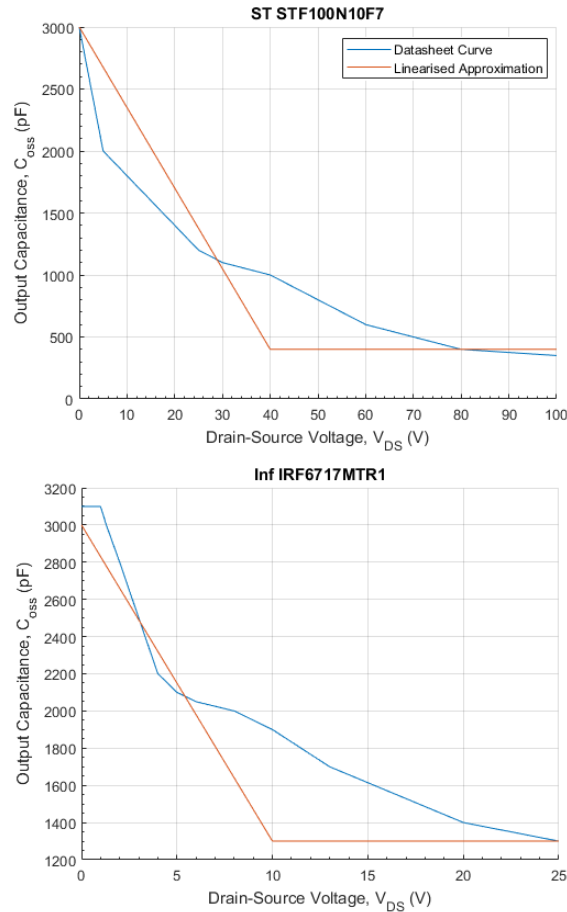


Fig. 3.3 Curves for output capacitance varying with drain-source voltage, along with bilinear simplifying approximation, devices under inspection are the STMicroelectronics STF100N10F7 [5] and the Infineon IRF6717MTR1 [6]. Axes are all linear, unlike those in figure 3.2.

While extracting exact curves for every device was concluded to be time-prohibitive, after the complete dataset had been assembled, a cross section of the devices had the curves precisely mapped, so that a comparison could be performed between the precise datasheet information and the bilinear/linear approximations.

The detailed results of this analysis can be found in Appendix B, but to summarise, there can be significant error induced by this linearising approximation. In some (in fact most) cases, the error between the two is not more than 25%, but in the most extreme case the error was found to be almost 140%. This simplification clearly induces a large amount of error in some cases, though it is worth noting that the approximations with the largest error do have the smallest impact on the loss figure overall, as parts of the model dependant on  $C_{iss}$  contribute much less to total system loss than those dependent on  $C_{rss}$ . Even then, this led to some reasonable questions being asked of the validity of this method.

During the data entry process, however, some suspicions were raised as to the precision of the information made available by the manufacturers. Figure 3.4 shows three sets of capacitance curves

for three different devices, with different current ratings, on-state resistances, voltage ratings, etc. The three curves are identical (other than the horizontal axis being longer for the device with higher voltage rating in figure 3.4c), and this was found to be far from an uncommon occurrence during the data entry process. Some even appear to be low-quality, compressed bitmap copies of their vector counterparts in other, often older datasheets. This suggests a potential lack of rigor in process by which these curves are derived. Or perhaps it is simply difficult to reliably predict over a large manufacturing run, and so a whole series of broadly similar devices are provided with the same capacitance data as it is likely to still be within error bounds, but these error bounds are not known to the user of said datasheets. It was concluded that before the impact of the error in the proposed estimation could be considered, the accuracy of the original data provided should be explored.




Fig. 3.4 Capacitance curves varying drain-source voltage for three different devices:

(a) Infineon BSB165N15NZ3[7], (b) Infineon IPD200N15N3[8] and (c) Infineon IPP320N20N3[9].

### 3.3 Experimental Pin-to-Pin Capacitance Evaluation

An experimental rig was developed to permit the evaluation of capacitance with respect to the drain-source voltage. Some data sheets document their testing procedure for some or all parameters provided on the datasheet, and capacitance is no exception. Figure 3.5 is representative of the best sort of test circuit diagram one is likely to find provided in a datasheet.

An almost exact replica of this circuit was attempted, but it was found difficult to maintain stability of the high-side MOSFET, believed to be operating as a constant load. An attempt was also made to produce a clean step current source for driving the gate, but issues were encountered with stray capacitance resulting in unacceptable ringing on the leading edge of any drive signal. Instead, an alternative circuit was devised which maintains the principle of the documented test procedure, while being easier to implement. This new test circuit is shown in figure 3.6.

Figure 3.6a shows a simplified circuit diagram of the testing apparatus. There are some significant differences between the datasheet test circuit in figure 3.5 and that used in experimentation. The



changes in testing procedure should not be so significant as to affect the results, however. The key elements of the test apparatus will be outlined herein.

The high side load is now a constant current driver, rather than the linearly biased matched MOSFET arrangement shown in figure 3.5. Figure 3.6b shows the basic configuration of the constant current driver, using an LM317 linear regulator. The LM317 is widely used as a robust, simple adjustable linear regulator that operates as a 1.25V fixed regulator, which with a potential divider on the output can achieve any voltage above that (though not above the supply voltage), or with a series shunt resistor as shown in figure 3.6b can regulate for constant current. Figure 3.6b shows, as an example, a 100Ω resistor which would correspond to a constant current output of 12.5mA – this can be adjusted by simply adjusting the series resistor but did not exceed the tens of milliamps range during testing for thermal reasons.

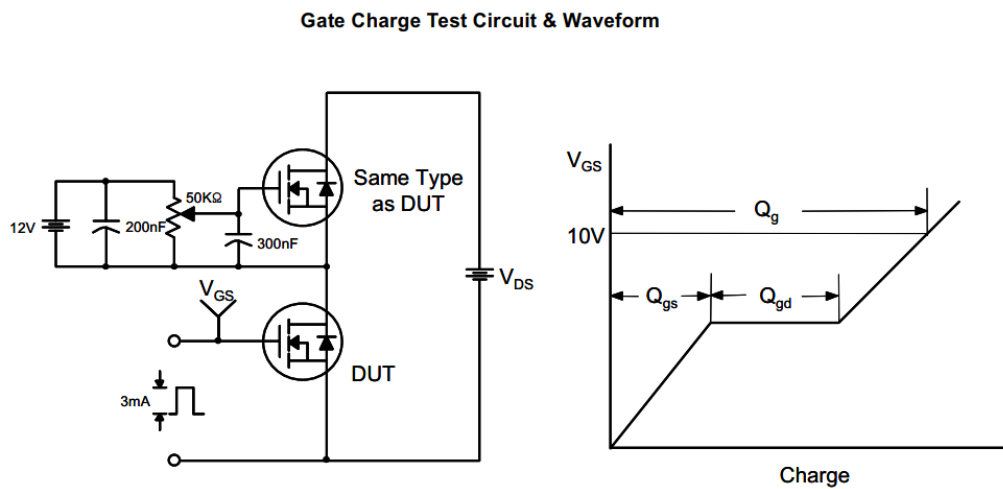


Fig. 3.5 Test circuit for estimating gate charge from the datasheet for the Fairchild Semiconductor FQA44N30 [10].

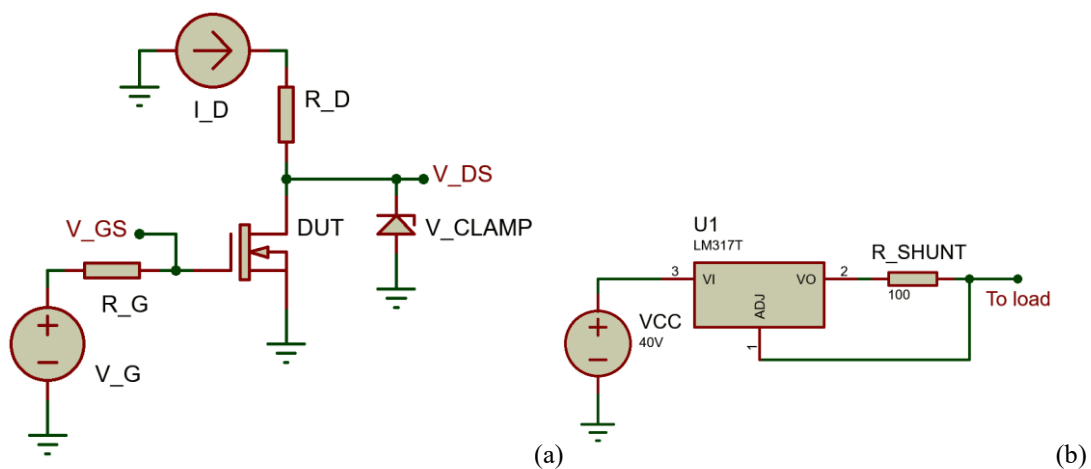


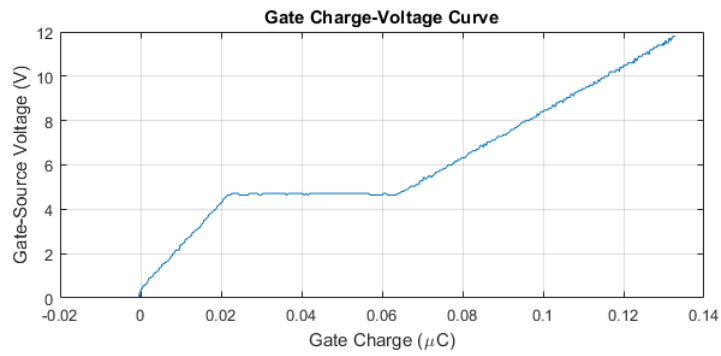
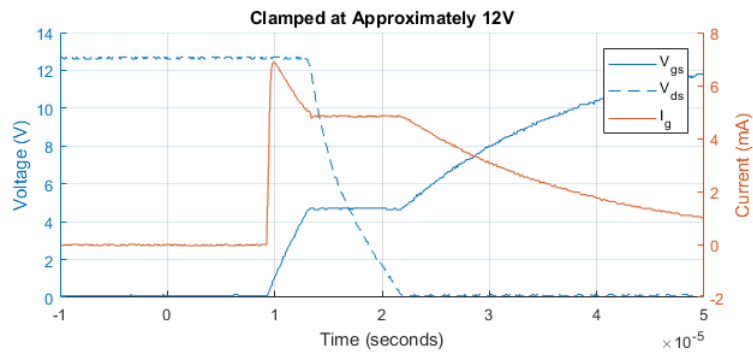
Fig. 3.6 Test circuit used in experimental exploration of device capacitances: (a) showing the circuit as a whole and (b) focussing on the constant drain current driver.

Between the drain and the source of the device under test (DUT) was a clamping diode, shown in figure 3.6a as  $V_{CLAMP}$ . A range of Zener diodes allowed for charge curves to be plotted for a variety of drain-source voltages. The use of a constant current driver and a clamping diode allows for a very low noise and reliable (assuming thermal stability) switching voltage, along with few undesirable transients during the DUT turning on. The linear regulator did not stabilise its output in a satisfactory way after a step load was applied from open circuit. The maximum clamping voltage was chosen as 30V, as most datasheets show capacitance curves largely levelling out above 30V, even in 500V and 600V class devices. A maximum clamping voltage of 30V resulted in choosing a supply voltage of 40V, as shown in figure 3.6b, so even accounting for shunt resistor voltage drop, the regulator will not approach dropout.


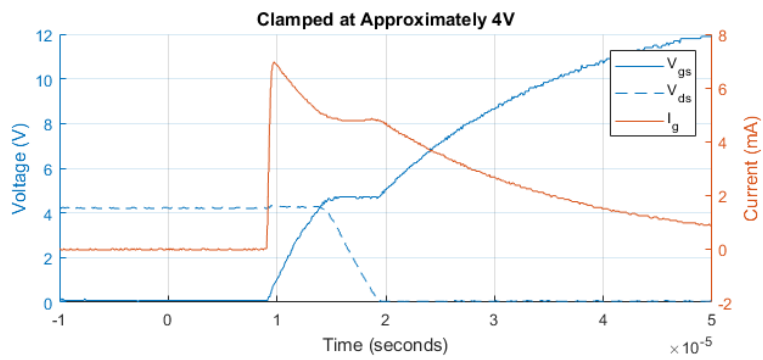
The resistor  $R_D$  in figure 3.6a (or  $R_D$ ) was placed on the high side of the DUT to permit drain current measurement, and to ensure that the constant current driver was maintaining regulation. Measurement of drain current along with drain-source voltage would allow characterisation of the capacitance. While this did indeed allow for validation of the current regulation, it could not perform accurate drain current measurements. This is due to the significant capacitance of the clamping diode, which is given on manufacturer datasheets, though it was not explored how accurate the diode capacitance data is. This Zener diode capacitance discharges through the drain of the DUT, but not through the shunt resistor, and so not all drain current is measured. While this colours results for measuring  $C_{oss}$ , it has no adverse impact on measurements of  $C_{rss}$  and  $C_{iss}$ .

The gate drive is performed by the Avago Technologies HCPL-3120 [11]. While a constant current driver was explored, satisfactory results could not be attained, so instead a voltage source gate drive with a relatively high gate resistance was utilised. The gate resistance  $R_G$  (or  $R_G$ ) was chosen as 2k $\Omega$  initially, to permit neglecting the DUT's inherent gate resistance, which is normally quoted as 10-20 $\Omega$  nominal and is very difficult to precisely measure and therefore account for. Voltage was measured at both the driver,  $V_G$  in figure 3.6a (or  $V_G$ ), and the gate,  $V_{GS}$  in figure 3.6a (or  $V_{GS}$ ). Measuring the voltage at both sides of the resistor  $R_G$  also allows for inferring current through it and therefore the current into the gate. Therefore, the gate-charge voltage curve can be measured, allowing for experimental validation of both  $C_{iss}$  and  $C_{rss}$ .

A range of clamping Zener diodes were used, with clamping voltages at (approximately) 30V, 20V, 12V, 8V, 6V, 4V, 0.5V and 0V. The 0.5V clamp was provided by a forward biased 1N4148 standard silicon diode, rather than a reverse biased Zener diode, and the 0V clamp was a jumper wire. The results of two example test runs are shown in figure 3.7, showing both the raw test data and the inferred gate charge-voltage curves for two different clamping voltages. Note the longer plateau or 'Miller shelf' in the gate-charge curve in figure 3.7a as compared to 3.7b, due to higher  $V_{DS,MAX}$ .



(a)



(b)

Fig. 3.7 Two sets of results for the experimental rig in figure 3.6 at clamping voltages ( $V_{DS,MAX}$ ) of (a) approximately 12V and (b) approximately 4V for the Fairchild Semiconductor FQA44N30 [10].

Three devices were comprehensively tested representing a typical high voltage device, a typical low voltage device and the device for which the estimation were worst: the ON Semiconductor/Fairchild FCH47N60 [1], the Texas Instruments CSD17573Q5B [12] and the Fairchild Semiconductor FQA44N30 [10], respectively. A comparison of precise datasheet information, estimated datasheet information and experimental results for these devices is shown in figures 3.8-3.10. Figure 3.9 has one fewer plot, as the device in question only has a maximum voltage rating of 30V, so testing at approximately 30V was omitted.

Figures 3.8 and 3.9 both show how the error induced by the estimation is relatively small when compared with the error between datasheet information and real, experimental results. It is worth noting that the results at higher clamping voltages are of greater importance as devices operating far below their rated voltage are unlikely to be optimal in the final optimisation. Figure 3.10 shows the device for which the error in the linearising estimation was found to be highest, made worse by the fact that in this case, the precise datasheet information appears quite close to the experimental data.

This returns us to the question of whether the error induced owing to the bilinear approximation made before is acceptable. This testing suggests that the data that manufacturers make available regarding the pin-to-pin capacitances is only accurate to within approximately half an order of magnitude. Perhaps this is due to difficulties in maintaining this value over production runs, or perhaps manufacturers don't feel the accuracy of these figures is important enough to their customers to spend the money doing accurate testing for every device. Reasoning aside, this large error puts the error of up to 140% from the linearising simplification in context. As the error bounds in the original data are so large, then the additional error from the bilinear simplification, while significant, is acceptable.

A further analysis was conducted to investigate whether there was any impact found by varying the drain current or gate resistor. If the test apparatus was functioning correctly, there should be no difference, and figure 3.11 shows the Miller charge (the charge during the plateau of the gate charge-voltage curve) for a range of clamping voltages. The figures from the precise datasheet information, the linearised simplification and experimental results under four permutations of gate resistance and drain current values are shown and labelled in figure 3.11. The consistency between the experimental tests shows no unexpected behaviour and therefore consistent test results. Figure 3.11 also highlights that while there is error between the precise datasheet information and its bilinear simplification, that error is small when compared with the error between either of them and real-world test results.

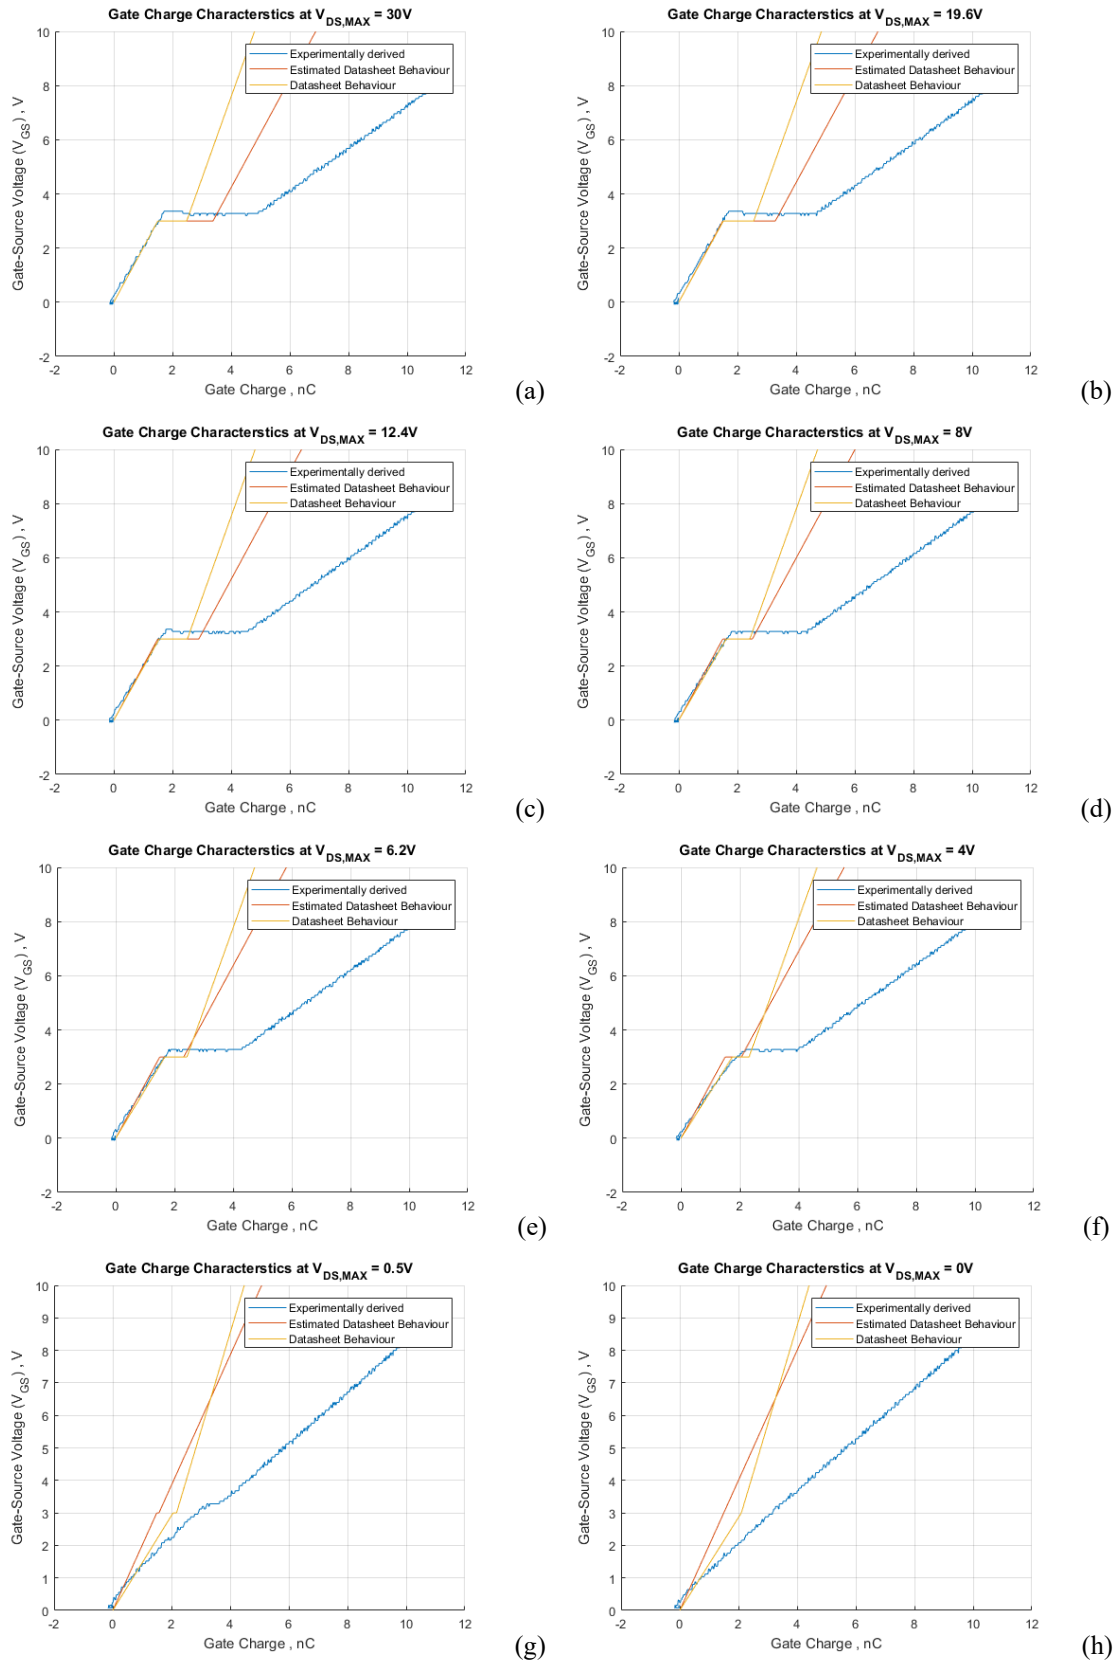


Fig. 3.8 Full results at a range of clamping voltages for the ON Semiconductor/Fairchild FCH47N60 [1].

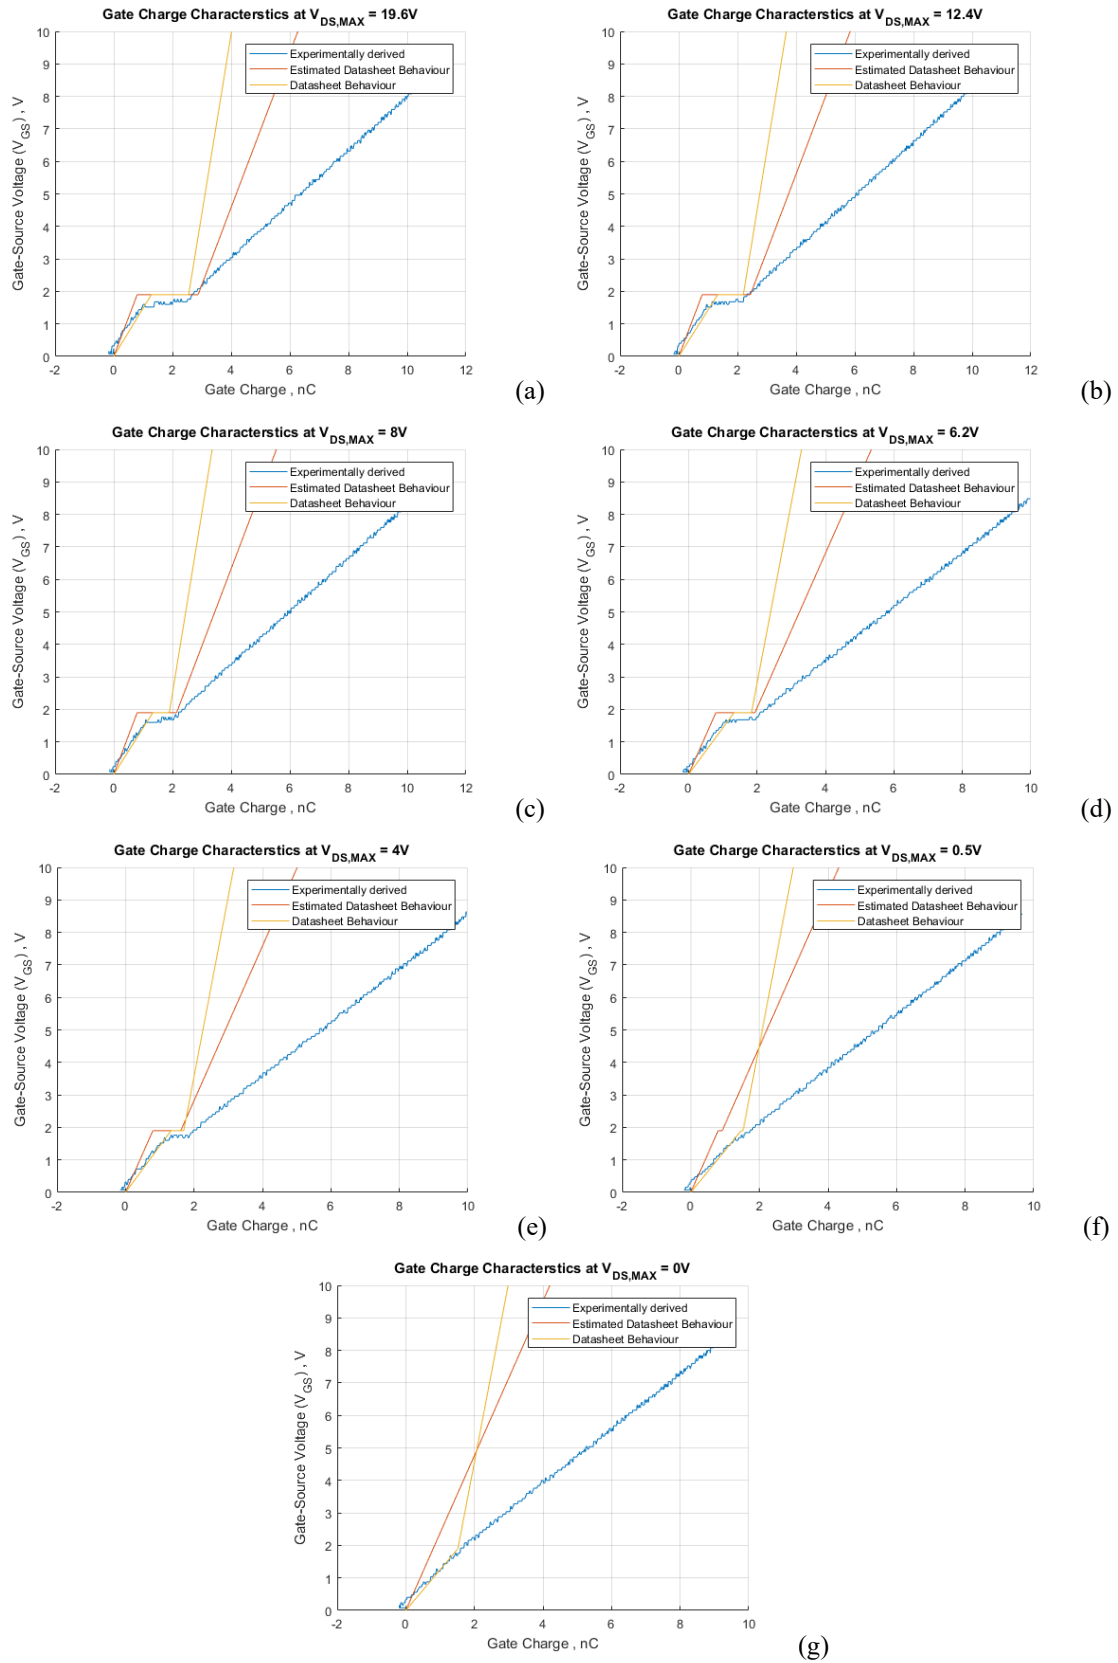


Fig. 3.9 Full results at a range of clamping voltages for the Texas Instruments CSD17573Q5B [12].

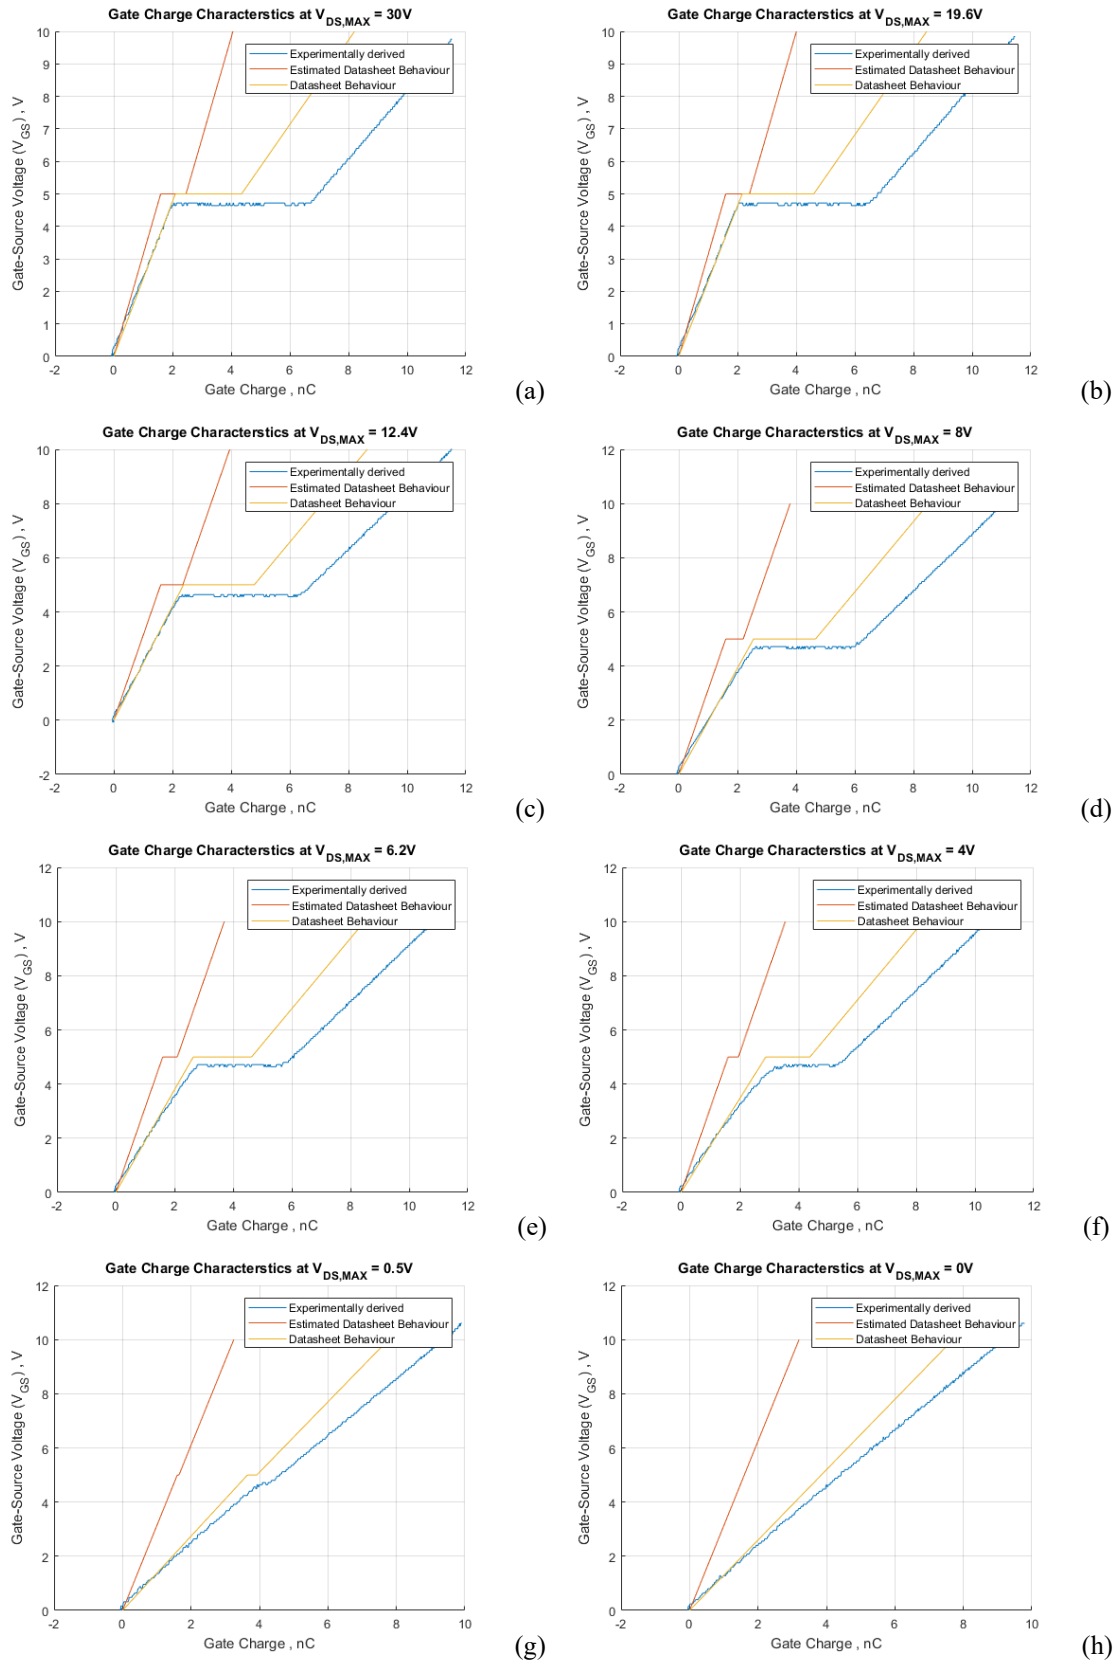


Fig. 3.10 Full results at a range of clamping voltages for the Fairchild Semiconductor FQA44N30 [10].

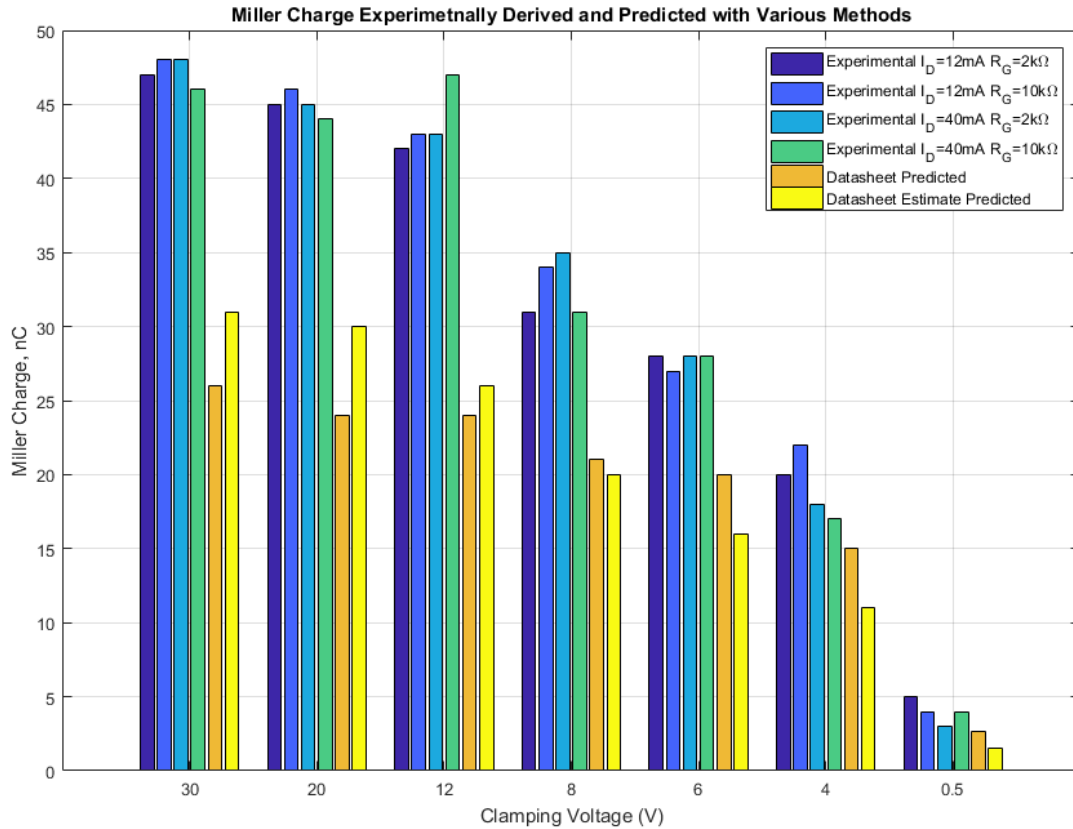


Fig. 3.11 Miller charge for a range of clamping voltages, comparing analytical predictions with multiple experimental configurations.

### 3.4 Plateau Voltage

The ‘plateau voltage’ is the gate-source voltage at which the Miller shelf occurs, for instance in figure 3.7 it can be seen that the plateau voltage is slightly less than five volts. While plateau voltage is available on some datasheets it is not universal, however threshold voltage is. These two quantities are not identical, and threshold voltage is usually quoted as a wide range with a maximum and a minimum but no explicit typical value. For example, the Infineon BSB165N15NZ3 [7] quotes a gate threshold voltage in the range 2-4V, but a gate plateau voltage of 5.2V.

The experimental rig used in the exploration of pin-to-pin capacitances was used to explore this with a handful of devices, and datasheets that quoted threshold voltage tend to be overestimates, with average gate threshold voltages being underestimates. As a compromise, typical gate threshold voltage multiplied by a factor of 1.25 was found to make a good estimate of real-world performance, with error of less than 20% over the sample of six devices tested. As small variations in this parameter have a relatively small impact of the loss metrics overall, this was considered acceptable. Particularly when, as mentioned previously, one recalls this method prioritises relative precision over absolute accuracy.



### 3.5 Body Diode Performance

As outlined in chapter 2, this method aims to be able compare the use of internal MOSFET body diodes rather than using additional, external switching diodes. While conventional wisdom suggests avoiding using body diodes for commutation during dead time, in a multilevel converter clearly the forward voltage drop of external diodes, with two diodes in the conduction path for every bridge, could be an issue. As a result, we want to evaluate the loss when using the intrinsic body diode – not only in forward conduction, but also in the reverse conduction that will occur due to the reverse recovery of these diodes, something negligible in external, high-performance switching diodes.

While there are a small number of devices designed with the body diode in mind, such as IXYS's HiPerFET series, most manufacturers appear to assume the body diode will not be utilised, and as such generally there is no information characterising its performance on manufacturer datasheets. This information is required to calculate power loss, so an investigation was performed to find whether it is possible to infer body diode performance from some information that is widely available.

A review of the literature in this field found little of use. Existing discussion in literature focuses instead on the experimental exploration of a specific device (or very narrow range of devices) [13-16], or discusses a die-level model to predict some parameters [17, 18]. This is of little help in the derivation of a practical prediction from readily available information over a broad range of devices. Initially, the derivation of a model from die level and up may seem sensible, but this is impractical as manufacturers seldom provide even the most basic information about the device geometry, and any sort of additional testing required on the devices was not deemed to be within the bounds of a practical method – a key emphasis of this work.

The goal is, therefore, to explore the possibility of a correlation between relevant body diode performance metrics and readily available device parameters. The use of readily available parameters enables the comparison of many devices over a range of conditions at once without costly and time-consuming testing of large sets of devices.

This will be achieved through large scale experimental characterisation of body diodes over a wide array of ratings, then exploring the correlation with datasheet parameters. Once these relationships are found, it will allow extrapolation to all silicon MOSFETs. Rather than limiting the analysis to only devices relevant to the reference converter specification outlined in chapter 2, the decision was made to explore MOSFETs as a whole. Figure 3.12 shows the range of voltage and current ratings (the two main parameters for a MOSFET), with this intended to be representative of a full range of voltage and current ratings for mainstream, commercial silicon MOSFETs. A full list of the devices used can be found in Appendix C.




Fig. 3.12 The voltage and current ratings of the MOSFET devices that had their body diodes characterised.




Fig. 3.13 A plot demonstrating the pulsed current used in tracing the I-V curves of the DUTs.

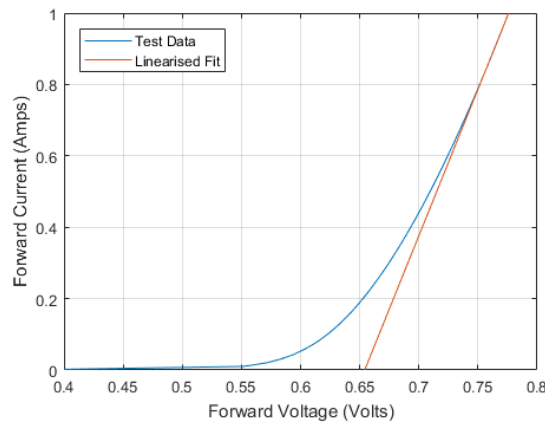


Fig. 3.14 I-V curve test data, along with linearised fit, for Fairchild FDMS86255.

As explained in chapter 2, there are two key performance metrics required for computation of power loss in body diodes during operation of the converter: the current-voltage (or I-V) curve of the diode for calculating loss during forward conduction of the diode, and the reverse recovery charge ( $Q_{RR}$ ) for calculating the impact of reverse conduction on converter power loss.

Tracing the I-V curve of the body diode was a straightforward process thanks to the use of an instrument designed for this application, a Keithley 2612A Sourcemeter. The Sourcemeter is an extremely precise four quadrant programmable power supply with current and voltage measurement, with the ability to program custom automatic testing procedures and easily return the data for later processing. It will run to relatively low current, 1A in this case, but this should be high enough to reach the linear part of the I-V curve.

The testing profile consists of a number of current pulses of increasing magnitude, the pulses being just long enough for the voltage to stabilise (1ms), and then a sample is taken. The duty cycle of the pulse train is kept low, with off periods of 100ms, to ensure that there is negligible thermal deviation during the testing – this is illustrated in figure 3.13. This does mean that this testing is only representative of body diode performance at room temperature, a significant compromise but addressing this would require a more time than was available. The gate was shorted to the source throughout this experiment to ensure that the MOSFET stayed off.

The data derived must be reduced to some representative parameters. A diode I-V curve near its conduction threshold can be accurately modelled using the Shockley equation [19], but performance near the conduction threshold of the body diode is irrelevant as we are attempting to model power loss in a pulsed current case, so this is not an appropriate model. At higher current, resistance dominates performance, and as such the body diode I-V curve is simplified to a linear model that is represented by an on-state resistance and a ‘simplified threshold voltage’. The comparison between actual test data and this approximation can be seen in figure 3.14. The 1A maximum current during testing assures the linear region of the diode I-V curve is reached.

Measurement of the reverse recovery current was performed with a custom test platform which rapidly transitions a diode from forward conduction to reverse voltage biased and measures the voltage across the device and current that flows out of it. An example of the raw data from this testing apparatus is shown in figure 3.15.

The orange line in figure 3.15 shows the source-drain voltage of the MOSFET under test, i.e. the voltage across the body diode. Before the drive voltage,  $V_{IN}$ , swings negative, the forward voltage of the diode can be clearly seen. The time from which the drive voltage drops below that of the device under test,  $V_{DUT}$ , to the point at which  $V_{DUT}$  reaches zero, the body diode is said to be in reverse recovery. As there is a known resistor (in this case 10k $\Omega$ ) between the drive voltage and the DUT, the current during this time period can be inferred, and by extension the charge that passes out of the device before reverse recovery completes – this is the reverse recovery charge,  $Q_{RR}$ . While there is some non-ideal behaviour, such as the imperfect transition of  $V_{IN}$  from positive to negative, this was concluded to be real behaviour rather than an instrumentation issue as it did not vary with different probes in different configurations, and so the data is accurate, and the behaviour can be accounted for.




Fig. 3.15 An example of the raw data extracted from the body diode reverse recovery test rig, in this case for the Infineon BSZ042N04.

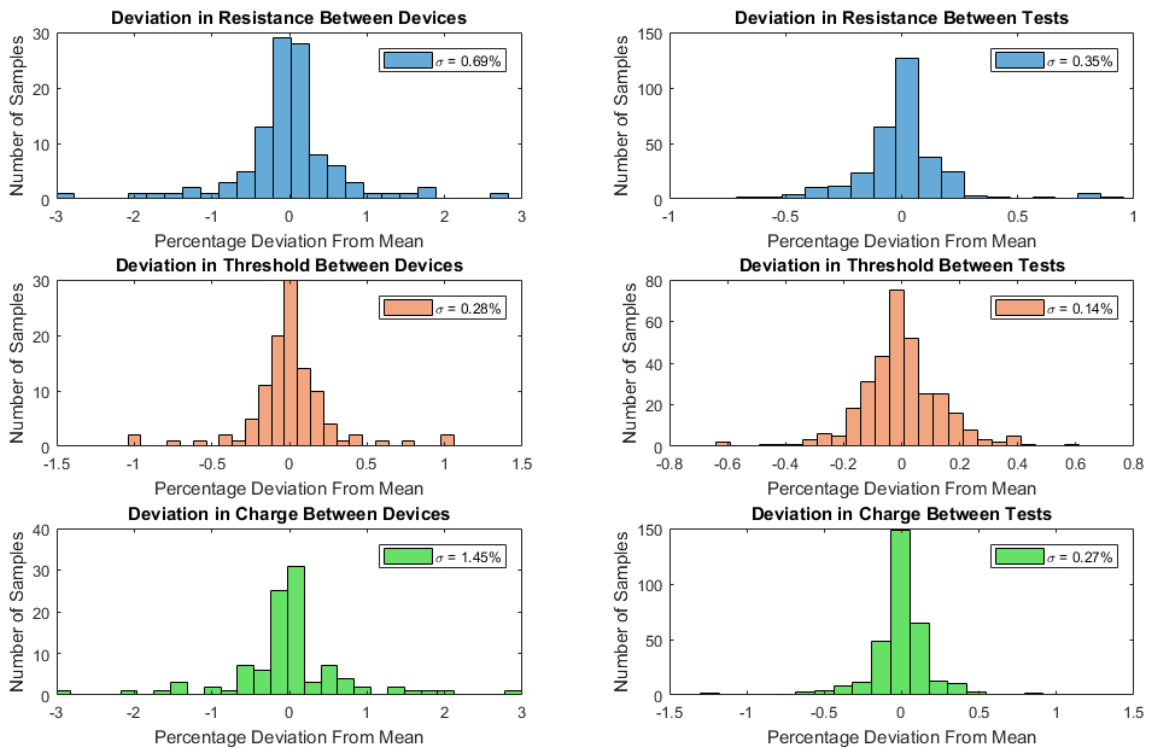


Fig. 3.16 Distributions representing error in experimentation (right column) and apparent manufacturing tolerances (left column) for the three experimentally derived parameters along with standard deviation,  $\sigma$ .

These tests were performed for all of the devices listed in Appendix D. Furthermore, to investigate the variation from device to device and the noise in the testing procedure, each device had three tests performed on two supposedly identical units. This will permit for investigation as to the noise of measurement as compared with the apparent manufacturing tolerances, with the hope being that the noise in the testing procedure would be lower than the inherent spread from device to device. The

results of this are shown in figure 3.16, with the standard deviations in each case shown in the legend. The results show that for all three body diode performance parameters being measured, manufacturing tolerances contribute more error than the measurement noise, therefore permitting the claim that the experimental rig is not contributing additional error. Furthermore, the distributions are approximately normal, as one might expect, further supporting the validity of testing as a skewed distribution might suggest a systematic source of error.

Now that the three key performance parameters have been experimentally derived for a large set of devices, the statistical analysis can be performed. These measured values are to be investigated for correlation with device parameters universally available through manufacturer datasheets. Voltage and current rating are two obvious parameters, but others were selected also. The datasheet parameters to be investigated for correlation with experimentally derived performance metrics are:

- maximum drain-source voltage rating,
- maximum drain current rating,
- nominal threshold voltage,
- device capacitances ( $C_{oss}$ ,  $C_{iss}$  and  $C_{rss}$ )
- and maximum power dissipation.

The device capacitances and nominal threshold voltage are closely linked to die geometry, and therefore possibly correlated with the performance metrics of interest. The maximum power dissipation seemed worth investigating as this is representative of the ‘bulk’ of the device, though this is affected significantly by the package rather than just the properties of the die. The capacitance values, which vary with respect to drain-source voltage, are all evaluated at 1V. The maximum power dissipation, which varies with temperature, is evaluated at the only temperature consistently available, 25°C.

The strength of correlation between the performance metrics and device parameters will be evaluated using the Pearson correlation coefficient [20]. Upon inspection of test data, first order linear and logarithmic fits were explored. While the Pearson correlation coefficient is normally expressed in the range -1 to 1, in this case only the strength of correlation is needed, so the modulus of this coefficient is used. Therefore, the coefficient that is normally in the range 1 to -1 is now in the range 0 to 1, with 1 representing perfect correlation and 0 representing no correlation.

Tables 3.1-3.3 show the modulus of the Pearson correlation coefficient between each of the selected device parameters for all four types of fit considered, for each body diode performance metric respectively. A ‘log-x fit’ (as seen in tables 3.1-3.3) describes a linear fit between the performance metric and the logarithm of device parameter. A ‘log-y fit’ is the opposite of this, and

linear and log-log fits are self-explanatory. The optimal fits for each case, namely that with the highest correlation coefficient, is shown in figures 3.17-3.19.

MOSFET device parameter	linear fit	log-x fit	log-y fit	log-log fit
Drain-Source Voltage Rating	0.61	0.67	0.59	0.67
Drain Current Rating	0.08	0.08	0.12	0.14
Nominal Threshold Voltage	0.55	0.51	0.55	0.51
$C_{OSS}$ (@1V)	0.16	0.41	0.17	0.45
$C_{ISS}$ (@1V)	0.27	0.29	0.30	0.34
$C_{RSS}$ (@1V)	0.28	0.43	0.29	0.47
Max Power (@25°C)	0.24	0.60	0.25	0.65

Table 3.1 Modulus of Pearson's coefficient of correlation between body diode forward resistance and MOSFET device parameters.

MOSFET device parameter	linear fit	log-x fit	log-y fit	log-log fit
Drain-Source Voltage Rating	0.20	0.23	0.21	0.24
Drain Current Rating	0.64	0.74	0.63	0.73
Nominal Threshold Voltage	0.28	0.25	0.28	0.26
$C_{OSS}$ (@1V)	0.32	0.70	0.32	0.69
$C_{ISS}$ (@1V)	0.77	0.83	0.78	0.81
$C_{RSS}$ (@1V)	0.56	0.71	0.56	0.70
Max Power (@25°C)	0.46	0.61	0.47	0.61

Table 3.2 Modulus of Pearson's coefficient of correlation between body diode simplified threshold voltage and MOSFET device parameters.

MOSFET device parameter	linear fit	log-x fit	log-y fit	log-log fit
Drain-Source Voltage Rating	0.68	0.56	0.75	0.7
Drain Current Rating	0.19	0.26	0.36	0.5
Nominal Threshold Voltage	0.32	0.3	0.51	0.51
$C_{OSS}$ (@1V)	0.91	0.77	0.6	0.92
$C_{ISS}$ (@1V)	0.67	0.51	0.74	0.79
$C_{RSS}$ (@1V)	0.73	0.52	0.67	0.75
Max Power (@25°C)	0.44	0.45	0.51	0.79

Table 3.3 Modulus of Pearson's coefficient of correlation between body diode reverse recovery charge and MOSFET device parameters.




Fig. 3.17 Body diode forward resistance plotted against and fitted to MOSFET device voltage rating.




Fig. 3.18 Body diode simplified threshold voltage plotted against and fitted to MOSFET input capacitance,  $C_{ISS}$ .




Fig. 3.19 Body diode reverse recovery charge plotted against and fitted to MOSFET output capacitance,  $C_{OSS}$ .

Figure 3.17 shows the correlation between MOSFET drain-source voltage rating and body diode resistance, which has a correlation coefficient of 0.67. While this is the best correlation found for the diode resistance, it is notably weaker than the optimal for voltage drop and reverse recovery charge. This is still a very worthwhile estimation, and even with its shallow gradient it reduces the standard deviation of the error compared with taking a mean value from  $21\text{m}\Omega$  to  $11\text{m}\Omega$ .

Figure 3.18 shows the correlation between MOSFET input capacitance,  $C_{ISS}$ , and body diode simplified threshold voltage. This is the strongest correlation found for diode voltage with a correlation coefficient of 0.83, or more precisely -0.83 as it is a negative correlation. As with figure 3.17, note that the x-axis is labelled logarithmically to more appropriately display the type of fit.

Figure 3.19 shows the correlation between MOSFET output capacitance,  $C_{OSS}$ , and body diode reverse recovery charge. This is a very good fit, with a correlation of 0.92 and few outliers. The fit in this case is not a log-x fit but rather a log-log fit, and as such the figure is plotted with both axes presented logarithmically.

These three best fits are expressed numerically in the three expressions shown in equation 3.1-3.3. These three expressions permit predictions for the three key body diode performance metrics required for the calculations derived in chapter 2 from readily available datasheet parameters.

$$R_D = 0.0303 \log_{10}(V_{DS,MAX}) + 0.0746 \quad (3.1)$$

$$V_D = -0.0734 \log_{10}(C_{ISS}) + 0.914 \quad (3.2)$$

$$\log_{10}(Q_D) = 0.793 \log_{10}(C_{OSS}) - 11.21 \quad (3.3)$$



### 3.6 Conclusions

This chapter has highlighted how difficult it is to obtain quantitative information on MOSFET devices to allow a detailed comparison to be made between devices. Although manufacturers provide tabulated and graphical data on their datasheets there are inconsistencies in the format in which this data is presented. Experimental investigations have also shown there is marked difference between datasheet values and those obtained from measurements. This chapter has explored the numerous ways in which there can be more to consider, ranging from the ensuring consistent interpretation to inferring whole new sets of information. While some of these data extraction methods might make some concessions that one might rather avoid, this is an exercise in making the best of what there is. Much academic research chooses not to engage with such often frustrating practical limitations on what information is available, but the core of this research is the idea that this could be almost immediately deployed in industry – so rather than being frustrated by practical limitations on available information, embracing them. The shortcomings of the manufacturer datasheets aside, after spending the time to ensure the extraction of consistent and reasonable data and predictions, we have successfully formed a firm foundation for the work moving forward.

The key novel work is the relationships that have been experimentally derived which enable first order approximations of body diode performance from readily available information – something not before possible. Furthermore, this chapter has documented the experimental investigation of the accuracy of manufacturer datasheet information in terms of pin-to-pin capacitances, which were found to be quite poor. Overall, the work described in this chapter ensures that the method described in chapter 2 is practically applicable in a consistent manner over a large set of devices using data available to anyone. The next chapter will proceed to use this methodology to optimise device selection with respect to system power loss, something that would not have been possible without the work done in these two chapters.

## References for Chapter 3

- [1] ON Semiconductor. FCH47N60 [Online]. Available: [tinyurl.com/ya4lh4w3](http://tinyurl.com/ya4lh4w3)
- [2] Infineon. AUIRFP4409 [Online]. Available: [tinyurl.com/y7nyvgyc](http://tinyurl.com/y7nyvgyc)
- [3] Toshiba. TK49N65W [Online]. Available: [tinyurl.com/ycjnmq2](http://tinyurl.com/ycjnmq2)
- [4] International Rectifier. IRFP4229 [Online]. Available: [tinyurl.com/ycdur53y](http://tinyurl.com/ycdur53y)
- [5] STMicroelectronics. STF100N10F7 [Online]. Available: [tinyurl.com/y7bvvagh](http://tinyurl.com/y7bvvagh)
- [6] International Rectifier. IRF6717MTR [Online]. Available: [tinyurl.com/ydfnsvmn](http://tinyurl.com/ydfnsvmn)
- [7] Infineon. BSB165N15NZ3 [Online]. Available: [tinyurl.com/ya8kony5](http://tinyurl.com/ya8kony5)
- [8] Infineon. IPP200N15N3 [Online]. Available: [tinyurl.com/y9bn3mda](http://tinyurl.com/y9bn3mda)
- [9] Infineon. IPB320N20N3 [Online]. Available: [tinyurl.com/ydf7asvj](http://tinyurl.com/ydf7asvj)
- [10] Fairchild Semiconductor. FQA44N30 [Online]. Available: [tinyurl.com/y9xqrrt5](http://tinyurl.com/y9xqrrt5)
- [11] Avago Technologies. HCPL-3120 [Online]. Available: [tinyurl.com/y8dwg6lr](http://tinyurl.com/y8dwg6lr)
- [12] Texas Instruments. CSD17537Q5B [Online]. Available: [tinyurl.com/ycaxe8ls](http://tinyurl.com/ycaxe8ls)
- [13] T. Appel and H. G. Eckel, "Bipolar reverse recovery behavior of the body-diode of a SiC JFET," in *2013 15th European Conference on Power Electronics and Applications (EPE)*, 2013, pp. 1-6.
- [14] C. D. Brown and B. Sarlioglu, "Reducing Switching Losses in BLDC Motor Drives by Reducing Body Diode Conduction of MOSFETs," *IEEE Transactions on Industry Applications*, vol. 51, pp. 1864-1871, 2015.
- [15] T. Funaki, "A study on performance degradation of SiC MOSFET for burn-in test of body diode," in *2013 4th IEEE International Symposium on Power Electronics for Distributed Generation Systems (PEDG)*, 2013, pp. 1-5.
- [16] J. Roig, S. Mouhoubi, F. D. Pestel, N. Martens, F. Bauwens, H. Massie, *et al.*, "Body-diode related losses in Shield-Plate FETs for SiP 12V-input DC/DC buck converters operating at high-frequency (4MHz)," in *2012 24th International Symposium on Power Semiconductor Devices and ICs*, 2012, pp. 291-294.
- [17] C. Schmidt and M. Roebnitz, "A Performance Comparison of SiC Power Modules with Schottky and Body Diodes," in *PCIM Europe 2017; International Exhibition and Conference for Power Electronics, Intelligent Motion, Renewable Energy and Energy Management*, 2017, pp. 1-8.
- [18] R. Siemieniec, O. Blank, M. Hutzler, L. J. Yip, and J. Sanchez, "Robustness of MOSFET devices under hard commutation of the body diode," in *2013 15th European Conference on Power Electronics and Applications (EPE)*, 2013, pp. 1-10.
- [19] W. Shockley, "The Theory of p-n Junctions in Semiconductors and p-n Junction Transistors" in *The Bell System Technical Journal*, 1949, 28 (3): 435-489
- [20] K. Pearson, "Notes on regression and inheritance in the case of two parents," in *Proceedings of the Royal Society of London*, 1895, 58 : 240-242.

# Chapter 4:

## Optimal Power Loss Trends

In chapter 2 a method was derived, based on a set of equations, for predicting power loss in a 6kW, single phase, cascaded H-bridge grid attached battery energy storage system, for an arbitrary number of cascaded bridges. In chapter 3 a discussion was had on how to best estimate the parameters required for the aforementioned method. Therefore, it is now possible to evaluate power loss for this reference converter specification, with any number of cascaded bridges and for any commercial device from just the datasheet. This chapter now explores the trends in device selection for minimising losses over a range of converter levels

### 4.1 Generating Results

To begin, a significant dataset of approximately 80 currently commercially available (as of November 2017) silicon MOSFETs was compiled (listed in Appendix B). The voltage and power specification of the reference converter, as described at the beginning of chapter 2, shows that any devices need a current rating of at least 26A ( $6\text{kW}/230\text{V}\approx 26\text{A}$ ), so all devices had a maximum continuous drain current rating in the range 28-50A. In all other regards, however, the dataset contains a very diverse selection of conventional silicon MOSFETs, particularly in terms of voltage rating, with voltage ratings from 710V to just 25V. The dataset includes very contemporary devices, as well as those which have been on the market much longer.

Over a wide range of converter order, the decision was made for every device in the dataset as to whether it was suitable for use under those conditions. This decision was based upon the voltage rating of the devices – as the converter order increases, the switching devices are only exposed to a fraction of the total DC link and therefore do not require as high a maximum drain-source voltage rating. The equation describing this relationship is shown in equation 4.1, where  $V_{\text{DS,MAX,REQ}}$  is the minimum required drain-source voltage rating,  $V_{\text{LINK}}$  is the total DC link,  $N$  is the converter order, and  $f_{\text{safety}}$  is a safety margin factor, set at 1.1 in this case for a safety margin of 10%.

$$V_{\text{DS,MAX,REQ}} = \frac{V_{\text{LINK}}}{N} f_{\text{safety}} \quad (4.1)$$

For any and all devices considered capable of being used in a given converter order, the power dissipation, as in chapter 2, was calculated. In each case of converter order, the device with the lowest total power loss was considered to be the optimal device. As the main goal was to explore the implication of converter order on total power loss, the best-case power loss is plotted against the converter order, as shown in figure 4.1.

## 4.2 Results of Power Loss Prediction Method for Silicon MOSFETs

Figure 4.1 shows the trend in power loss in the optimal case with increasing converter order, remembering that ‘optimal’ is defined as selecting the device that results in the minimum total converter power loss associated with the switching devices and their drivers. In this case at a switching frequency of 10kHz, representing a low switching frequency that is still credible in a modern power converter. Also displayed in figure 4.1 is a simple cost (£) figure, derived exclusively from summing the cost of the switching devices (with the cost figures found using the method described in the previous chapter). This is clearly not a complete cost model, but it serves to give some indication while not taking part in the optimisation. Furthermore, the bars representing total power loss are subdivided into the separate sources of loss as outlined in chapter 3.

The main conclusion to be drawn from figure 4.1 is that converter power loss can be lower with a higher order converter. In this case, solely from the perspective of converter power loss, the optimal number of cascaded bridges would be ten. Noting the dominance of the dark blue colour in the bars shows that on-state conduction loss is the greatest by far, as is perhaps unsurprising at a low switching frequency. The trend in on-state conductance loss shows that the lower on-state resistance of lower voltage rated devices is enough to offset the larger number of devices in the conduction path.

Notable in figure 4.1 are the ‘steps’ that the power loss seems to take down between certain values of converter order, for example between nine and ten cascaded bridges there appears to be a sharp step down. This shows the point at which a new, lower voltage rated device becomes available and becomes the optimal device (60V devices in this example), with that mostly maintaining superiority until the next voltage rating threshold. Similar ‘steps’ are seen in the cost curve, just not always down.




Fig. 4.1 Variation in minimum total power loss in cascaded H-bridge converters switching at 10kHz.

Figure 4.2 shows the result of the analysis evaluated with a switching frequency is 80kHz. The first difference to notice between figures 4.1 and 4.2 is that in the case of figure 4.2 the power loss is generally higher, not only in worst case but also the best case. This is to be expected as higher switching frequencies (all else being equal) do lead to higher power loss – though of course higher switching frequency has a number of other benefits that may outweigh the cost in power dissipation in the switching devices, such as requiring smaller passive components.

Furthermore, in this case the dominance of on-state conduction losses is not as acute. Naturally, at a higher switching frequency the on-state resistance does not get higher and neither does the energy required to turn a device on, but the number of turn-on events over a given period has increased and so too will the total power loss.

If this is the case however, why is it that the on-state conduction loss for a converter order of one is so much higher in the case of figure 4.2 than in figure 4.1? This is because the method has found that the device selected at 10kHz (as in figure 4.1) as optimal in the case of a single H-bridge, is not appropriate at the higher frequency of 80kHz (as in figure 4.2). It is notable that this does not appear to be happening at higher converter order, hinting at another significant potential benefit of moving to multilevel converters. The lower voltage rated devices that the use of a multilevel converter grant access to, tend to have significantly lower capacitances, and as a result are more able to switch at higher frequencies than their higher voltage rated counterparts. As higher switching frequency tends to enable miniaturisation of power converters, it is not difficult to see how this alone could be a serious argument for multilevel converters in such applications.




Fig. 4.2 Variation in minimum total power loss in cascaded H-bridge converters switching at 80kHz.

Figure 4.3, once again, displays similar information to that in figure 4.1 and 4.2, but in this case at a switching frequency of 600kHz, meant to represent something of a boundary case where increasing switching frequency is a priority, perhaps in the interest of reducing the size of filter passives. Once again, the power loss is greater across the board in the case of figure 4.3, than in figures 4.1 and 4.2, due to the higher switching frequency, particularly in the case of a single H-bridge.

In fact, the outlier case in figure 4.3 shows a total power loss across four switching devices in excess of 200W, which is far beyond the package thermal limits of the device in question, not to mention dissipating over 200W across four gate drivers! Clearly this is not a realistic design scenario, and serves to reinforce the point made previously that multilevel converters can enable easier access to higher switching frequencies.

The cost lines in figures 4.1-4.3 serve to highlight how considering only one aspect of system design is unlikely to yield a truly optimal result. For instance, in figure 4.3 it may seem that a solution with twelve cascaded bridges is very nearly optimal, and while it is from the perspective of power loss the orange line clearly shows that the cost of the switching devices is more than five times higher than in the case of five cascaded bridges. As discussed in previous chapters, this analysis, while informative, only forms one part of a larger discussion and cannot confidently select an optimal converter specification in isolation.




Fig. 4.3 Variation in minimum total power loss in cascaded H-bridge converters switching at 600kHz.

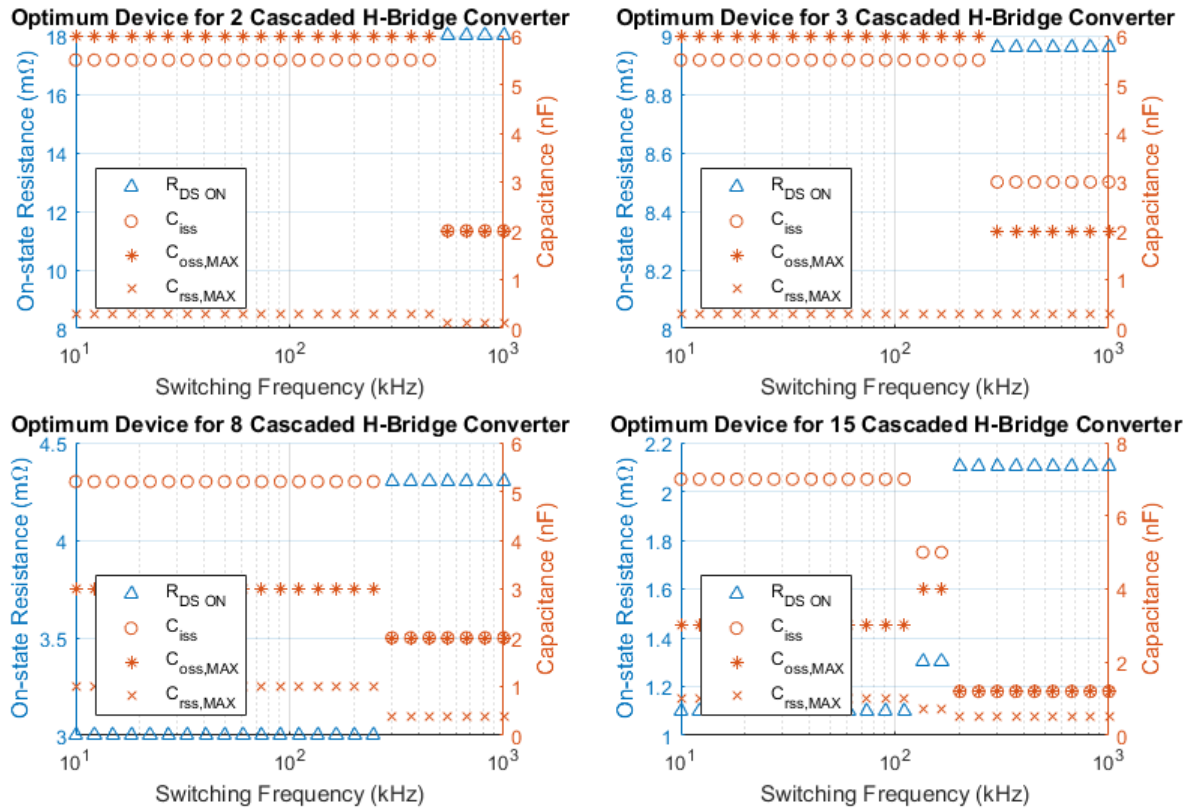


Fig. 4.4 How optimal devices parameters change with increasing switching frequency, over a range of converter orders.

Less obvious in figures 4.1-4.3 is how the parameters of the optimal devices tend to change with frequency. This was discussed somewhat in regards the change at one H-bridge between the results in figures 4.1 and 4.2, where at higher frequency a higher on-state resistance ( $R_{DS,ON}$ ) was accepted in the interest of a device capable of operating at the higher frequency. This trend can be seen over the full range that the method operates.

Figure 4.4 shows how the ratings of the optimal device at a given converter order vary with the switching frequency of the converter, each graph showing this for a different, fixed value of converter order. In all cases, on-state resistance of the optimal device increases as switching frequency increases and the pin-to-pin capacitances reduce. This shows the algorithm behaving as we would expect: at very low frequencies such as the results shown in figure 4.1, where on-state resistance dominates, on-state resistance must be minimised even if that means greater capacitance values; meanwhile at higher frequencies such as in figure 4.3 the transient losses start to be much more significant and a lower capacitances will be selected even if that comes at the cost of higher on-state resistance.

### 4.3 Validation of the Method with SPICE Modelling

While any assumptions the method are built on have been justified and as such the data it produces should be accurate, it is clearly desirable to perform some level of validation in these figures. Unfortunately, experimental validation is impractical as the large sets of devices would lead to expensive and time-consuming testing. However, comparing results with some other established analytical approach will increase confidence in the results given. Specifically, simulation in LTspice IV was performed, and power loss figures were attained for the exact same conditions as used here, and the results compared.

Figure 4.5 shows an example model used in validation of power loss prediction in one device, in this case the Infineon BSZ042N04NS. In fact, all of the devices considered in this validation are Infineon devices within their ‘OptiMOS’ series. A representative sample of these high-quality SPICE models enables validation over a wide range of conditions.

There are four parameters to be set in the SPICE simulation shown in figure 4.5: the supply voltage V2, the load resistance R2, the gate drive voltage V1 and the gate resistance R1. The supply voltage and load resistance serve to represent the voltage that the switching devices would experience in the system, for example 50V would represent a 500V DC link split over ten cascaded bridges, while the resistance of  $1.916\Omega$  limits system current so when M1 is completely on the current is identical in this case to the RMS system current in the method. R1 is already calculated in the method derived in chapter 2, so that is trivial to find, with the waveform parameters in V1 coming from the peak gate voltage and the switching frequency of the converter, with a representative duty cycle of 50%.

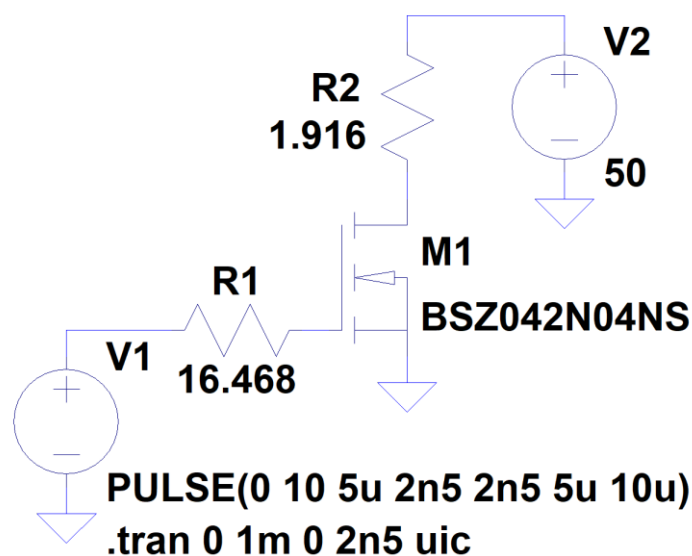


Fig. 4.5 LTspice model used in validation of power loss in Infineon BSZ042N04NS.






Fig. 4.6 Waveform of instantaneous power dissipation with time for the model shown in figure 4.5.

To calculate instantaneous power loss, the power loss between the drain and the source and the power supplied by the gate drive power supply, V1, are summed. The power loss between the drain and the source is the product of the drain source voltage and the drain current, while the gate drive power is the product of the voltage of V1 and the current flowing from it.

After performing a transient simulation on the model in figure 4.5, the resulting power trace is shown in blue in figure 4.6. One can clearly see the off period where power loss is zero, the on period where there is modest power loss, and the period during which the switching occurs where power loss is much higher for a short period of time. Also shown in figure 4.6 is the average power over this simulation period of 100 $\mu$ s, for reference.

For a limited range of devices, namely those for which a reliable SPICE model could be found, the average power loss was computed for a wide range of potential converter conditions, under these identical conditions the power loss was calculated using the new method, and the power loss figures were compared. Unfortunately, SPICE modelling does not allow us to separate the power loss in to separate parts in the way that the new method does, so we may only compare the total power loss. The results of these analyses are shown in table 4.1.

Device Name	Switching Frequency (kHz)	No of Cascaded Bridges	New Method Prediction (W)	SPICE Model Prediction (W)	Difference (%)
Infineon BSC076N06NS3	250	6	71	69	2.7
Infineon BSC320N20NS3	400	3	168	146	13
Infineon BSZ042N04NS3	100	10	63	64	1.6
Infineon BSZ0904NS1	20	14	81	91	12
Infineon IPD053N06N	500	5	48	37	22
Infineon IPP200N15N3	40	2	61	61	0.3
Infineon IPP320N20N3	800	4	229	184	20
Infineon IRFZ44N	250	8	268	234	13

Table 4.1 A comparison of results from the new method alongside SPICE simulation results.

Table 4.1 shows that the error between the SPICE simulation results and the new method is small, with a maximum difference of 22%, and a mean difference of just 11%. The simulations tend to skew towards higher switching frequencies, a deliberate decision as power loss at lower switching frequencies is dominated by on-state resistance. It is of note that in the case of the Infineon BSZ0904NS1, even with a low switching frequency of 20kHz, there is a 12% discrepancy between the two – a surprisingly large error in conditions where on-state resistance is dominating, as predicting a resistive power loss should be relatively simple. Further investigation shows that the figures given for  $R_{DS,ON}$  within the SPICE model simply do not align closely with the datasheet values in the case of the BSZ0904NS1, unlike the others – why this might be is unknown.

There are a number of reasons why the results of the method may reasonably differ from the results produced by the SPICE model, for instance the evaluation of on-state drain-source resistance ( $R_{DS,ON}$ ). As explained in chapter 3, the method evaluates  $R_{DS,ON}$  at a junction temperature of 80°C to better represent  $R_{DS,ON}$  under actual operating conditions – these SPICE models, however, do not. Furthermore, the SPICE models presumably include a nominal gate resistance where the method (as discussed in chapter 3) does not as this is poorly defined. As a result, the value for gate resistance, external or otherwise, is not identical in the two cases.

Overall, the small error (as compared to other unavoidable sources of error as discussed in chapter 3) between the SPICE model predictions for power loss and the new method developed in this thesis grants great confidence in the new method. Furthermore, it is notable that there is no particular bias towards one being higher than the other and no tendency for a greater discrepancy at higher switching frequencies, both of which would suggest some sort of systematic bias. This SPICE validation along with the rigorous validation of assumptions made in the method demonstrate the reliability of the method and bolsters confidence in the accuracy of the predictions made.

## 4.4 Alternative Semiconductor Technologies

### 4.4.1 Insulated Gate Bipolar Transistors (IGBTs)

As outlined in chapter 2, the fact that IGBTs do not exist at lower voltage ratings in the same way as silicon MOSFETs, with voltage ratings seldom below 600V, means that IGBTs presumably cannot compete at higher converter order for the application studied here. Furthermore, IGBTs are not able to operate at switching frequencies as high as MOSFETs, with many datasheets stating they are unsuitable for switching frequencies above as little as 5kHz, so are unable to compete with MOSFETs at higher switching frequencies. The question remains, how do IGBTs compare with silicon MOSFETs in a single bridge at a low switching frequency. After all, if IGBTs have significantly lower loss than silicon MOSFETs under these conditions then the benefits of moving to a multilevel converter with low voltage MOSFETs may be moot.

As we are only comparing the two classes of devices at low switching frequency it is fair to compare just the on-state conduction power loss of these two classes of devices. This is fortunate as the data that would be needed to predict power loss due to driving the gate of the IGBT in the same level of detail as the case of the MOSFETs is not available, due to extremely patchy and inconsistent datasheets (in details relevant to switching related losses, at least). However, those datasheets that do contain data for IGBT gate capacitance suggest that it is slightly higher but comfortably within an order of magnitude when compared with similarly rated MOSFETs, further making it reasonable to neglect in their comparison.




Fig. 4.7 Comparison of on-state conduction loss for silicon MOSFETs and IGBTs

Figure 4.7 shows the on-state conduction loss of a single transistor carrying a current of 26A, equivalent to 6kW at 230V, IGBT output power loss being calculated differently to MOSFET power loss as outlined in chapter 2. Figure 4.7 shows the power loss under these conditions for all of the MOSFETs used in the results earlier in this chapter (blue), as well as a sample of IGBTs of relevant voltage rating (orange) – all devices have similar current rating (30-50A) as with the previous analysis. There are only 600V IGBTs shown as lower voltage rated devices in the relevant power range could not be found, and while IGBTs are available at higher voltage ratings, these ratings are not necessary in this application and obviously have inferior performance, so have been omitted.

As we already know, lower voltage devices have significantly lower per-device power loss due to their much lower on-state resistance, but IGBTs can be seen to have similar power loss to their silicon MOSFET contemporaries, but not the least. The fact that figure 4.7 shows IGBTs as ‘middle of the pack’ as compared to similarly rated silicon MOSFETs under the only conditions they can compete shows that it is reasonable to omit IGBTs from the overall optimisation as they have no real impact on the conclusions. However, it does show that at low frequency with a single bridge an IGBT is a reasonable device to use as compared with a MOSFET, as common wisdom would attest.

#### 4.4.2 Silicon Carbide (SiC) MOSFETs

Silicon carbide MOSFETs may be a relatively new power switching devices class, but already see use in numerous industrial applications. They have a lower theoretical limit on on-state resistance and can be made to withstand higher voltage than silicon MOSFET equivalents. However, they do require a higher gate driver voltage, and have an inferior body diode (at least in terms of forward voltage drop). With the exception of the body diode and the higher gate drive voltage, the power loss prediction method can be applied to a SiC MOSFET in the same way as with a normal silicon MOSFET. While silicon devices were driven with 10V peak gate drive voltage in generating results for the method, the SiC devices were driven with 18V, in accordance with manufacturer advice given in application notes [1]. The body diode performance cannot be predicted by the method derived in the previous chapter, but the devices considered have relevant performance figures given in the datasheets – they are not widely better documented but rather a small number of devices were selected with body diode data.

As SiC MOSFETs are not available in voltage ratings lower than 600V (though they are available much higher), they are unlikely to be competitive with significantly lower voltage rated silicon devices, and so are not of the utmost relevance to the multilevel scenario, but the question remains as to whether they are optimal in the case of a single H-bridge. To this end, the power loss for a small range of SiC MOSFETs in this application was calculated over a range of switching frequencies, alongside a selection of silicon devices of similar voltage rating. The results of this analysis are shown in figure 4.8.




Fig. 4.8 Power loss with increasing switching frequency for a range of MOSFETs, both silicon and SiC.

Figure 4.8 shows a range of switching devices, including silicon MOSFETs (in blue) and SiC devices (in red). They all increase their power loss with increasing switching frequency, as one would expect, with many devices reaching a frequency above which the devices are incapable of operating. There are two main observations: the SiC devices are comfortable operating at switching frequencies that these relatively high voltage silicon MOSFETs are not capable, and secondly that the SiC devices do not have the lowest loss of all devices, though it is very close and SiC is better on average. For completeness, the devices included for the results in figure 4.8 are listed in table 4.2.

The primary conclusion of this analysis is that in the case of a single bridge solution, a SiC device is certainly a reasonable solution, and may well be optimal in the case of a high frequency device. This does not extend to the multilevel case, however, as unlike conventional silicon MOSFET technology, there are no lower voltage rated devices to benefit from a multilevel use case. In higher voltage and higher current applications SiC devices would likely outshine the conventional silicon competition under almost all conditions, but it seems in the 500-700V, 30-60A range the difference is not that significant.

Device Name	Technology	Voltage Rating
IXYS MKE38RK600DFELB	Silicon	600
ST STW56N60DM2	Silicon	600
Fairchild FCH47N60N	Silicon	600
Infineon IPW65R045C7	Silicon	650
Toshiba TK49N65W	Silicon	650
ST STW56N65M2	Silicon	650
Infineon IPB65R045C7	Silicon	700
ST STW62N65M5	Silicon	700
ST STI57N65M5	Silicon	700
ROHM SCT 3060AL	Silicon Carbide	650
USCi UJ3C065080K3S	Silicon Carbide	650
USCi UF3C065040K4S	Silicon Carbide	650

Table 4.2 A list of devices used in analytical exploration of SiC device performance

#### 4.4.3 Gallium Nitride (GaN) High Electron Mobility Transistors (HEMTs)

GaN HEMTs that can operate at significant power are a very new tool in the figurative toolbox of a power electronics designer, with great promise for sparking a new generation of extremely efficient power converters capable of operating at switching frequencies in excess of 1MHz [2]. The specific series of devices considered in these analyses are Efficient Power Conversion (EPC)'s range of sixth generation eGaN devices. These devices were chosen as one of the only that could be found rated in the current range of interest with comprehensive datasheets and easy availability of parts, at the point that this was investigated (June 2017).

In the same way as with the silicon MOSFETs, a dataset was compiled from manufacturer datasheets for a small but representative sample of devices – nearly half of EPC's full range. The voltage ratings of the GaN devices included range from 160V to 30V. There were no devices available of suitably high current rating with a higher voltage rating at the time of writing (Nov 2018) - though devices such as the Texas Instruments LMG3410 are rated to 600V it only has a current rating of 12A.

As outlined in chapter 2, the method for predicting power loss of GaN devices is almost identical to silicon MOSFETs (with the exception of modelling body diode loss), so direct comparison with results for silicon devices is possible. The results of the optimisation under a given set of conditions, both including and excluding the GaN devices, are shown in figure 4.9.

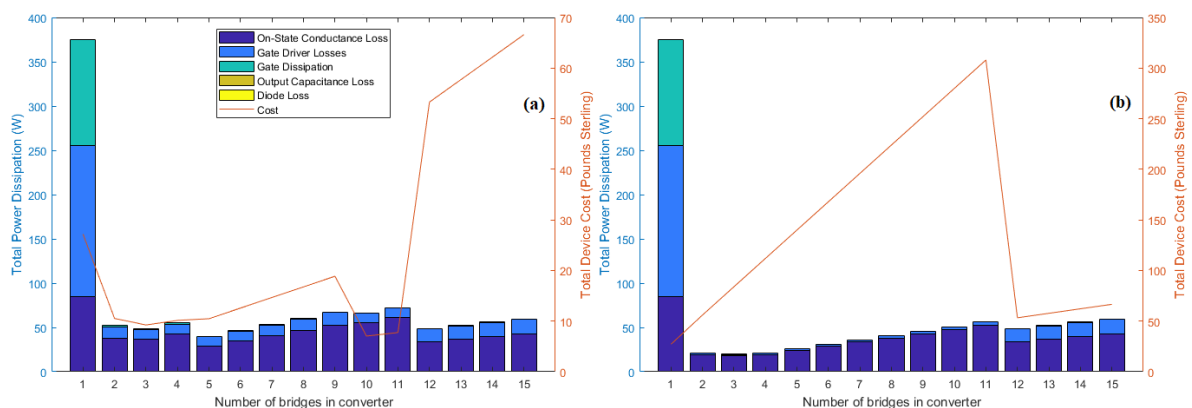


Fig. 4.9 Variation in minimum total power loss in cascaded H-bridge converters switching at 500kHz (a) excluding GaN devices and (b) including GaN devices.

Figure 4.9a shows the trend in power loss for the optimal devices selection with increasing converter order, the same results as in figures 4.1-4.3 but now at a switching frequency of 500kHz. Figure 4.9b shows the same results only now including EPC's sixth generation eGaN devices in the dataset alongside the existing silicon MOSFETs.

There are clearly a series of devices ranging from two cascaded bridges to eleven where there is a difference in the selected devices in the two cases – this is where the EPC eGaN devices proved to be optimal over any of the silicon MOSFET devices in the dataset. Particularly at lower converter order, there is a significant reduction in total power dissipation of the optimum case, as much as 62%, in fact. Notably, at very high converter order, in excess of 11 cascaded bridges, the silicon devices maintain their position as optimal due to their extremely low on-state resistance which the GaN devices cannot compete at the very low voltage rated devices (25-30V). Also, there is no difference in the case of a single bridge converter as none of the GaN devices considered are rated to sufficiently high voltage.

Also of note is the difference in the orange line, representing raw combined cost of the transistors in the converter, between figures 4.9a and 4.9b. The GaN devices may be optimal in terms of the very low power loss they enable, but they clearly significantly inflate the system cost with these GaN transistors costing approximately £8 per unit while a comparable silicon MOSFET might cost less than £1 per unit. This is particularly true at high converter order where the low voltage MOSFETs see a significant unit cost reduction, whereas the GaN devices do not.

Clearly under some conditions the use of GaN devices can enable significantly lower total system power loss. To explore how these benefits change with switching frequency and converter order, figure 4.10 was created. To create figure 4.10, the optimal (i.e. minimum) power loss scenario was found over a range of switching frequencies and converter order, both including and excluding GaN devices in the device dataset. In the cases where GaN devices enabled lower power loss, the extent to which this was case was evaluated and plotted as a heatmap.




Fig. 4.10 A plot showing the improvement in total power loss enable by GaN devices over silicon MOSFETs.

Figure 4.10 shows that the greatest benefits are to be gained through the use of GaN power transistors at high switching frequencies and a relatively small number of cascaded bridges. More yellow colours show greater benefits by GaN device use, up to 73% in fact, while in the deep blue regions silicon MOSFET devices were still optimal. The white region in the top left denotes an area in which no devices in the dataset are capable – GaN devices are not rated to sufficiently high voltage while MOSFETs rated to high enough voltage cannot operate at such a switching frequency. The greatest gains are to be found in the range of two to five cascaded bridges and at higher switching frequencies. At the higher numbers of cascaded bridges, the low voltage rated (<60V) GaN devices fail to improve on the low on-state resistance of equivalent silicon devices. However, the turn-on gate energy remains remarkably low, as can be seen in the case of ten and eleven cascaded bridges where the GaN devices are competing with 30V silicon devices and still prove optimal at very high switching frequency.

The results shown in figure 4.10 does not exceed a switching frequency of 1MHz, despite the GaN devices advertising that they can operate at full current rating up to 5MHz. This is not included as many (if not most) of the silicon devices would not be capable of this due to the stray inductances and capacitances in the package design, and these high-frequency parasitic effects are not included in any part of the method due to the difficulty with evaluating this. The EPC eGaN devices have taken these issues so seriously that one of the main revisions from the fifth generation to the sixth was a complete redesign of their proprietary power BGA package to reduce capacitance between the pads. As a result, it can be said that not only are there significant benefits in terms of reduced power loss to be gained



through the use of GaN power devices, but also access to extremely high switching frequencies, comfortably in excess of 1MHz.

At the time of this research being performed (July 2017), GaN power device are not capable of operating at simultaneously high enough current and voltage for this reference converter. As a result, a multilevel converter in this application is being used for similar reasons to a traditional multilevel converter application, a method of overcoming insufficient voltage rating of the semiconductor technology. However, in this case we are electing not to use devices that are suitably rated to access the raft of benefits enabled by this emerging power semiconductor technology. It is worth note that during writing of this (January 2019) Texas Instruments just released a GaN-based half bridge module rated for 600V at 40A, which would be adequate for this application without a multilevel implementation.

This analysis shows that one of the key benefits of GaN power devices appears to be the extremely low gate turn-on energy. The difference appears so striking that some validation of this was deemed wise, in case of some unanticipated phenomenon occurring in the use of this new class of power semiconductor device.

#### 4.5 Experimental Investigation of GaN HEMT Gate Turn-On Energy

The testing procedure was not dissimilar to that used in the investigation of pin-to-pin MOSFET capacitances discussed in the previous chapter. In fact, the same testing apparatus was originally employed, but was found to operate in an unsatisfactory manner when testing the GaN devices. After extensive investigation it was concluded that this was almost certainly predominantly an instrumentation issue. It appeared that the gate of the GaN devices is so high impedance that even the inductance from the loop of the oscilloscope ground lead was causing ringing on the signal rendering any results useless, even with a large gate resistance to damp out oscillation. An example of this ringing can be seen in figure 4.11, showing the case after ringing had already been significantly attenuated.




Fig. 4.11 Data from initial GaN/MOSFET gate energy test rig showing ringing on signals.

Clearly, the instrumentation error (the ringing) in figure 4.11 would spoil any results gathered. So, a new test bed was developed, designed to minimise the distance, and therefore loop inductance, between the pads of the device under test (DUT) and the oscilloscope probe as much as possible. In the interest of time the previous test apparatuses were built by hand on traditional prototyping board, but as the GaN devices to be tested are only available in a ball grid array (BGA) package a printed circuit board (PCB) would be required. As multiple devices were to be tested – with different physical packages in many cases – a motherboard would be built for the test circuit with each device being built on an interchangeable daughterboard.

The daughterboard was the focus of the effort to minimise the loop inductance in the measurement. All the devices tested were in surface mount packages, so the shortest path from the pads to the test probes would be through the PCB. The test daughterboards have via stitching from one side of the board where the device is mounted through to large pads outside the soldermask on the other side of the boards that can be readily probed. This, in conjunction with a change in the application of the oscilloscope test probes, yielded successful test results. Figure 4.12 shows the motherboard and daughterboards together, with figure 4.12b showing a small piece of tinned copper wire used on the scope probe to significantly reduce the loop in the ground bond.




Fig. 4.12 Pictures of (a) the mated test PCBs, with (b) showing probing method.




Fig. 4.13 Renders of daughterboard PCB for testing of EPC eGaN power transistors.

An example of the daughterboards used in the successful test rig is shown in figure 4.13, shown in this case prepared for the proprietary BGA package used by the sixth generation EPC eGaN devices. Figure 4.13 shows that this board has two components in addition to the main transistor: they comprise a Zener diode clamp and series resistor that were added due to concerns about gate overvoltage if the ringing was not only in measurement, but this was found to be unnecessary and so remained unpopulated – these were absent in the daughterboards designed for the silicon devices that were tested. The large, upright black blocks in the render in figure 4.13 are surface mount standard 0.1inch single in-line headers for connecting to the motherboard.

The test rig depicted in figure 4.12 also differs from that used in the capacitance validation experiment performed in chapter 3 in that the load in the circuit was no longer a voltage clamped constant current circuit, but instead a large inductance. The DUT is pulsed on for a sufficiently short time that the drain current cannot rise to the point of the potentially damaging the device. This large inductance ( $17\mu\text{H}$ ) was paired with a high performance Schottky diode, for free wheel current flow. This also allowed the test rig to operate to significantly higher maximum drain-source voltage.

To perform the test, a relatively short (approx.  $5\mu\text{s}$ ) on pulse was driven into the gate of the DUT through a known resistor. The drive signal, gate-source voltage and drain-source voltage were all synchronously measured through an oscilloscope. The gate current is found and measured over the period from the drive voltage raising about  $0.1\text{V}$ , to the time that the drain-source voltage falls below  $0.1\text{V}$ . Also knowing the gate-source voltage makes it easy to calculate the gate turn-on energy for any DUT. Note that more energy flows into the gate after the device is fully on, but we are measuring the energy to turn on the device and nothing more.

The devices used in the testing are listed in table 4.3. They span a range of voltage ratings for both silicon and GaN technologies. They are all from the main dataset used for the results given at the beginning of this chapter, and so have similar current rating to one another in the range of 30-50A. The silicon devices have a maximum gate drive voltage of  $10\text{V}$  in this test, as low as is normally reasonable and the value used in main analytical method, whereas the GaN devices only have a maximum gate drive voltage of  $5\text{V}$ , also reflective of the analytical method.

Device Name	Technology	Voltage Rating
Infineon BZS0904NSI	Silicon	30
Fairchild FDMC86340	Silicon	80
Fairchild FDMS86103	Silicon	100
Infineon BSB165N15NZ3	Silicon	150
EPC EPC2029	Gallium Nitride	80
EPC EPC2034	Gallium Nitride	160

Table 4.3 A list of devices used in experimental validation of GaN gate turn-on energy




Fig. 4.14 Curves for the turn-on transient for the Infineon BSB165N15NZ3 and the EPC EPC2034.




Fig. 4.15 Trends in gate turn-on energy for the range of devices listed in table 4.3.

Figure 4.14 shows the turn-on transients for a silicon device alongside an equivalent GaN device, captured using the experimental apparatus as used to find gate energy. There are a number of key differences, firstly the time between the drive signal starting (at time =  $0\mu\text{s}$ ) and the device turning on, shown by the drain-source voltage falling to zero, is significantly longer in the case of the silicon device. Not only does the GaN device turn on faster, but the gate-source voltage is significantly lower during this turn-on period, reducing the turn-on energy even further. There also appears to be a much-reduced Miller shelf in the case of the GaN device, in line with the much-reduced feedback capacitance documented in the manufacturer datasheet.

Figure 4.15 shows the trends in power loss with increasing maximum drain-source voltage for a range of silicon and GaN devices. It shows that GaN devices do indeed have much lower gate turn-on energy than similarly rated conventional silicon devices, with the 150V silicon device having almost ten times the gate turn-on energy of the 160V GaN device. In fact, both the 80V and 160V GaN devices manage to achieve lower energy than the 30V silicon device - a device explicitly added to the test as a result of its very low gate turn-on energy. Furthermore, the results show the power loss increasing with greater  $V_{DS,MAX}$  owing to an increase in the Miller shelf as the feedback capacitance must charge over a large voltage. The GaN devices show their very low feedback capacitance by increasing power loss with respect to  $V_{DS,MAX}$  at a more gradual rate than the silicon devices with which they are compared.

The experiment confirms that the GaN devices considered do, indeed, have a much lower gate turn-on energy than silicon equivalents, going some way to explaining the large improvement in converter power loss, particularly at high frequencies.

## 4.6 Conclusion

This chapter has shown, using the methodology outlined in chapter 2 and 3, that a multilevel converter can have significantly lower total power loss than a conventional single bridge solution, and even higher order converters with a great many devices can have competitively low power loss. Furthermore, moving to a multilevel solution grants particularly great benefits in a converter operating at higher switching frequency due to the access to smaller, lower voltage rated devices that tend to more easily operate at higher switching frequencies. The results of this method were validated over a limited selection of devices using a SPICE derived power loss model, which showed strong correlation.

In the discussion of power devices other than silicon MOSFETs it was concluded that at this power level IGBTs were outperformed by MOSFETs, though SiC MOSFETs can outperform high voltage silicon MOSFETs at higher switching frequencies. Neither IGBTs or SiC MOSFETs are available in voltage ratings lower than 500-600V, so do not benefit from a multilevel converter topology in this case study. On the other hand, GaN devices do not currently exist in this current range with sufficient voltage rating to be used without use of a multilevel topology. GaN devices excel at very high frequencies where even low voltage silicon devices struggle, and under some conditions can outperform silicon devices at lower switching frequencies. Multilevel converter technology could enable the use of these devices today, enabling massive potential reductions in system power loss and size.

One of the key novel contributions shown in this chapter are the trends in multilevel converter performance described toward the beginning of the chapter (particularly in figures 4.1 to 4.3), which have not been done before and make an insightful companion to a significant and growing body of research within the academic community. The other primary novel contributions are in the area of Gallium Nitride power devices: not only does the extensions of the method to include GaN devices give numerical insight into the benefits that GaN devices will enable as they become mainstream, but the experimental investigation of GaN devices gate energy as compared traditional silicon devices validates some very striking numbers only before seen claimed on manufacturer datasheets. Between chapters 2 to 4 the trends on multilevel converter performance with respect to converter power loss have been thoroughly explored, the next chapter will now explore this from the perspective of system thermal performance.

## References

- [1] CREE. 2008, Application Considerations for SiC MOSFETs.
- [2] Efficient Power Conversion. 2009, Application Note 002: Fundamentals of Gallium Nitride Power Transistors.

# Chapter 5:

## Thermal Performance in Multilevel Converters

Over the course of the previous three chapters a method for prediction power loss of switching devices in a multilevel converter was created, the means by which we can practically derive the relevant parameters was investigated, and finally the system power loss trends over converter order and switching frequency was explored. It was concluded that a cascaded H-bridge multilevel converter can lead to lower total power loss and ease the transition to higher switching frequencies, especially so if one considers the new classes of devices that a multilevel converter might grant a designer access to.

In this chapter we shall explore the implications on system thermal performance when moving to a multilevel solution of increasing order, with a view to investigate the generalised trends. Again, there is a focus on maintaining a practical method that could be easily applied to a large number of devices over a wide range of conditions without prohibitively time-consuming analyses. This will, at least initially, revolve around steady-state finite element modelling (FEM).

### 5.1 Finite Element Thermal Modelling

It stands to reason that a well distributed thermal load means a heatsink can be used more efficiently. In the limit of one device, the resulting local hotspot will drive up junction temperature as the dissipation becomes a point source, as compared to identical thermal power dissipation over a large number of thermal contact points from multiple devices. The results in the previous chapter have shown that total power dissipation in a multilevel converter need not be higher (and can in fact be lower), so looking at trends over varying converter order, for a given power level, is somewhat informative.

To investigate the relationship between system thermal performance and converter order, a number of simulations using steady-state thermal finite element modelling (FEM) were performed to see how the peak temperature under the device, and by extension junction temperature, varies with number of devices. To focus the analysis, a reference converter specification was selected. The reference converter specifications are identical to those decided upon at the beginning of chapter 2, namely:

- Nominal 500V DC link
- Maximum average power capacity of 6kW
- Cascaded H-bridge configuration



### 5.1.1 FEM Thermal Trends with More Devices

FEM simulations were performed for a range of generic heatsink designs with the number of devices ranging from one, up to twenty-five. Figure 5.1 shows the temperature distribution in the case of one device compared to the case of twenty-five. The total power dissipation is identical, but the peak temperature has been reduced: from 159°C to 75°C, given an ambient temperature of 25°C. Peak temperature is of interest as this is the temperature of the heatsink where the thermal pad is in contact. The thermal model includes radiative, conductive and convective thermal models for greatest completeness, all with model parameters offset to typical values (later denoted in table 5.1, described as Heatsink 2) used to represent an aluminium heatsink in air without forced air cooling.




Fig. 5.1 FEM results in the case of a single device dissipating as compared to many devices.




Fig. 5.2 The resulting trend line from a series of FEM simulations as the number of devices mounted on the heatsink increases, for identical power dissipation and cooling.



The simulation results shown in figure 5.1 highlight the significant gains to be made, but in order to explore the trends in greater depth, the simulation was run multiple times for varying numbers of devices and a curve of best fit was applied, as seen in figure 5.2. Under all circumstances, the devices are placed on a grid evenly spaced across the heatsink, hence considering integer multiples, e.g. 3x3 4x5, etc. The curve of best fit was found, by inspection, to be of the form, where  $x$  is number of devices:

$$y = ax^{-b} + c \quad (5.1)$$

This trend holds for one type of device with one set of thermal conditions, which is indicative, but the goal of this analysis is to create a more generalised expression. To this end, further validation was performed. Similar analysis was conducted with different heatsinks and different thermal parameters in an attempt to validate a general expression applicable to any thermal scenario with a minimum of initial information. Table 5.1 shows the conditions of each test, while figure 5.3 shows renders of the four heatsinks used in this extended FEM analysis.

Heatsink Parameters	Heatsink Number			
	1	2	3	4
Convection Coeff. (W/m <sup>2</sup> K)	10	12	9	10
Convection Temperature (°C)	20	30	25	28
Emissivity	0.2	0.25	0.2	0.18
Radiation Temperature (°C)	200	180	210	200
Total Power Dissipated (W)	100	100	30	40
Length (mm)	600	500	150	240
Width (mm)	600	300	120	120
Depth (mm)	25	85	50	97

Table 5.1 A table of the parameters used in the analyses with results shown in figure 5.4.




Fig. 5.3 Renders of the 3D heatsink models used in the analyses to follow. Numbers correspond to those used in table 5.1.

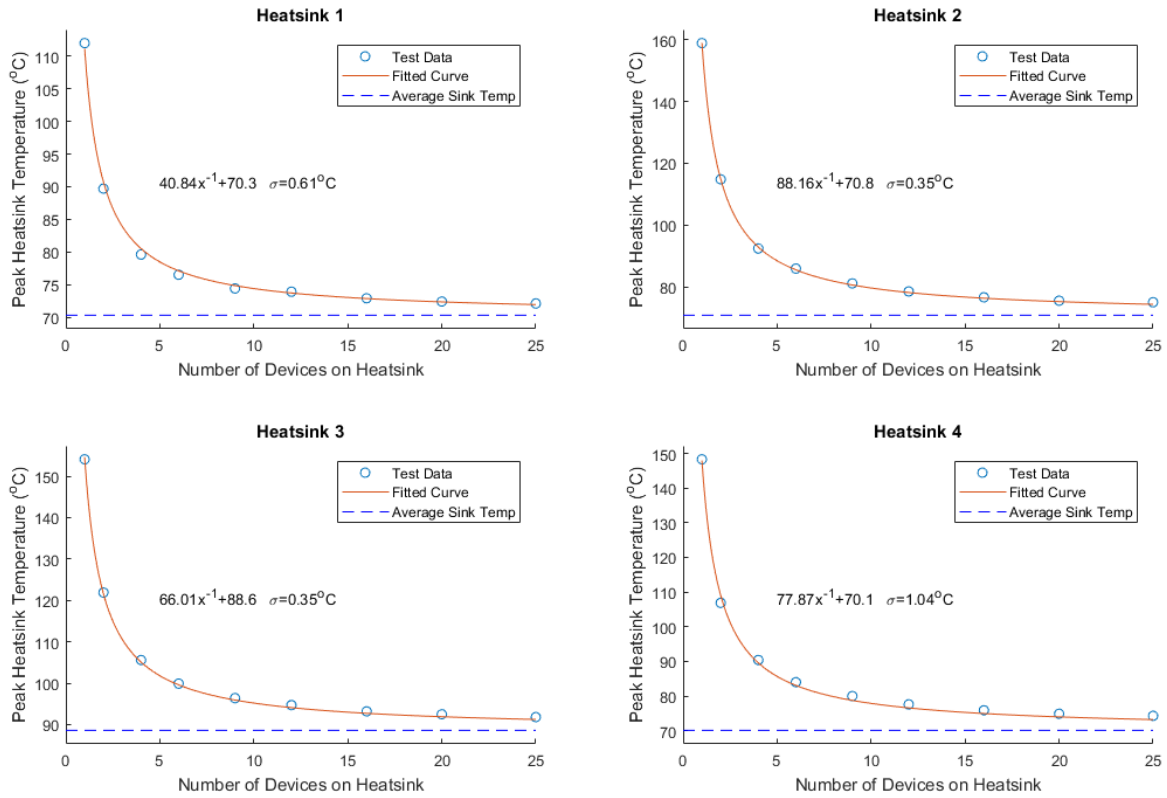


Fig. 5.4 The resulting trend line from a series of FEM simulations as the number of devices mounted on the heatsink increases, all else being equal, under the four configurations described in table 1. Standard deviation of fitted curve is shown each plot as  $\sigma$ .

Inspection of the further analysis, shown in figure 5.4, permitted some simplification to equation 5.1, yielding a generalise expression:

$$y = ax^{-1} + T_{AV} \quad (5.2)$$

Equation 5.2 has two unknowns, where  $T_{AV}$  is the average temperature of the heatsink and does not change dependant on number of devices, and  $a$  must be solved for. This simplified expression means that the trends in heatsink performance can be predicted for any heatsink and thermal parameters by running a singular FEM analysis to find  $T_{AV}$  and  $a$ .

The standard deviation in the lines of best fit are annotated as  $\sigma$  in figure 5.4 and show that use of the generalised, when compared to equation 5.1, expression in equation 5.2 still maintains very low error over the wide range of conditions considered. Therefore, this more generalised fit is shown to be good.

### 5.1.2 FEM Thermal Trends with Device Pad Size

As discussed in previous chapters, one of the key benefits afforded by the use of multilevel converters is that as number of levels in the converter increases, the designer gains access to lower voltage rated devices, which can have higher performance than their higher voltage counterparts in numerous ways. One key, relevant, exception to this is that these devices tend to be physically smaller. To investigate whether this has a significant effect, much as before, a pair of FEM runs with a given heatsink under two extremes of pad size, with all other parameters identical, were performed. The results are shown in figure 5.5, with a 4mm<sup>2</sup> pad resulting in a 122°C peak while a 192mm<sup>2</sup> pad results in a peak temperature of only 94°C.

Figure 5.5 shows, much as before with number of devices, that pad size does indeed have a significant impact of peak heatsink temperature. To find an expression that describes the relationship between the impact on the peak heatsink temperature and the pad size, the same method was utilised as in the case where the impact of the number of devices on the heatsink was considered. This was performed with a view to generating an expression similar to equation 5.2, but now for the impact of pad size on peak heatsink temperature, so that their relative impact can be easily considered. A range of pad sizes, from packages all capable of current in the order of fifty amps, were considered:

- TO-247 at 192mm<sup>2</sup>
- TO-220 at 120mm<sup>2</sup>
- D<sup>2</sup>PAK at 48mm<sup>2</sup>
- TDSON-8 at 16mm<sup>2</sup>
- TSDSON-8 at 4mm<sup>2</sup>




Fig. 5.5 FEM results in the case of four small devices dissipating heat, compared to four devices with large pads.

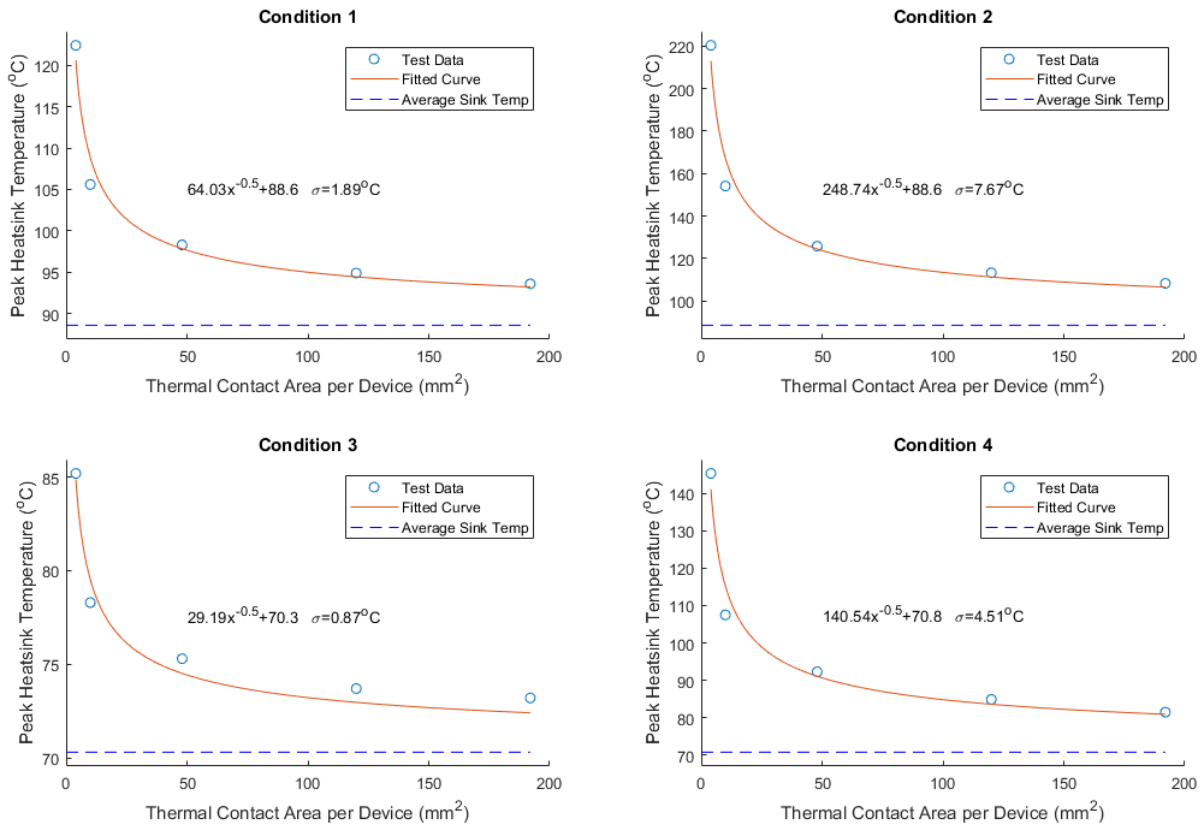


Fig. 5.6 The resulting trend line from a series of FEM simulations as the device thermal pad size increases, all else being equal, under the four test conditions described in table 2. Standard deviation of fitted curve is shown each plot as  $\sigma$ .

Heatsink Parameters	Test Condition Number			
	1	2	3	4
Convection Coeff. (W/m <sup>2</sup> K)	9	9	10	12
Convection Temperature (°C)	25	25	20	30
Emissivity	0.2	0.2	0.2	0.25
Radiation Temperature (°C)	210	210	200	180
Total Power Dissipated (W)	30	30	100	100
Length (mm)	150	150	600	500
Width (mm)	120	120	600	300
Depth (mm)	50	50	25	85
Number of Thermal Pads	4	1	12	6

Table 5.2 A table of the parameters used in the analyses with results shown in figure 5.6.

Figure 5.6 shows the trends in peak heatsink temperature with respect to pad size for four different sets of thermal conditions. The conditions of these four cases are defined in table 5.2. The data points in figure 5.6 are accompanied by the line of best fit. The function describing the fit takes the same

general form as that shown in equation 5.1, but it differs from equation 5.2. The generalised function for the case of peak heatsink peak temperature variation with pad size is:

$$y = bx^{-0.5} + T_{AV} \quad (5.3)$$

Equations 5.2 and 5.3 allow for bulk analysis of the thermal performance of a range of converter configurations and devices for a given heatsink and specification from a single FEM simulation. This allows for an informed, practical selection of the optimum device and number of levels for a given multilevel converter based on real world data (i.e. manufacturer datasheets), optimised from the perspective of minimum junction temperature.

## 5.2 Modelling with Resistor Network Equivalent

While FEM is a popular and powerful method for performing thermal modelling, it is of course not the only option. A popular approach, very common in simple first-order thermal predictions, is the use of electrical equivalent networks, where current is an analogue for thermal flux (i.e. power) and voltage is analogous to temperature. Models of this nature can incorporate reactive (i.e. capacitive or inductive) components to model properties such as thermal mass, but in the case a steady state thermal model these are unnecessary. A diagram of the heatsink thermal model with N by M devices is shown in figure 5.7.




Fig. 5.7 A diagram of a thermal equivalent electrical model for a lossy plate, i.e. a simplified heatsink.

The voltage at every node in the equivalent circuit in figure 5.7 represents the temperature at that nodes location on the lossy plate, with each node having: a potential power input represented by a current source, a thermal resistance to ambient represented by  $R_{th,a}$ , a thermal resistance to the next node horizontally represented by  $R_{th,x}$ , and a thermal resistance to the next node vertically represented by  $R_{th,y}$ . While the diagram shows a current source at every node, this is not necessarily the case.  $R_{th,x}$  and  $R_{th,y}$  are affected both by the thermal conductivity of the material from which the heatsink is constructed as well as the thickness.  $R_{th,a}$  is affected by the thermal properties of the heatsink materials but also the surface area available to dissipate into the air, which would depend on the depth of any fins on the heatsink as well as the area that each node would represent – in fact, as there is no explicit radiative, conductive, or convective dissipation model accurately quantifying  $R_{th,a}$  would be quite challenging.

Fortunately, in this case absolute accuracy is not of concern. Rather, this model will be used to validate the claims made based upon FEM, that the fit expressed in equation 5.2 and that this is true regardless of specific heatsink performance parameters. Therefore, the challenges in accurately quantifying  $R_{th,x}$ ,  $R_{th,y}$  and  $R_{th,a}$  need not be addressed. Initially, figures were chosen such that overall system performance appeared similar to that shown in figure 5.1, specifically  $R_{th,x} = 0.25^{\circ}\text{C}/\text{W}$ ,  $R_{th,y} = 0.25^{\circ}\text{C}/\text{W}$ ,  $R_{th,a} = 1000^{\circ}\text{C}/\text{W}$  and total power dissipation was 400W.

This model will be evaluated against the fit derived before and outlined in equation 5.2. Figure 5.8 shows a heatmap of four cases where the total power dissipation is the same, but the number of devices varies – all four cases are displayed with different colour maps so that reading the maximum and minimum on the legend can immediately show the large difference as the number of devices vary. The grid of devices in the resistor network used for these simulations is 120x120, making for a good balance between resolution and reasonable computation time. For simplicity, ambient temperature is considered as zero in this case as it would only result in a DC offset in final figures and we are not interested in absolute results, but rather the trends. The solutions are calculated by invoking Ngspice through MATLAB to solve the system as an operating point analysis, i.e. DC steady-state condition. For the full range of results under these conditions, some seen in figure 5.8, the trend of peak temperatures with an increasing number of devices is shown in figure 5.9 and fitted to equation 5.2 in exactly the same manner as shown in figure 5.4, with a suitably low standard deviation of just  $0.43^{\circ}\text{C}$ . This shows a good fit under one set of conditions, but this can now be rigorously explored over a very wide set of conditions.




Fig. 5.8 Model predictions over a range of numbers of devices using the model shown in figure 5.7.




Fig. 5.9 Trend of peak temperatures against number of devices for results in figure 5.8 fitted to equation 5.2.

To further explore the validity of the claim that the trend described in equation 5.2, a Monte Carlo analysis will be performed on the four main parameters in the resistor network model:  $R_{th,x}$ ,  $R_{th,y}$ ,  $R_{th,a}$  and the total power dissipation. A Monte Carlo is where a spread of values are considered for the variables in concern, and the output is then explored for every possible combination of these with the output in this case being the quality of fit to the trend shown in figure 5.2. Many Monte Carlo analyses use a normal distribution of the variables in question, such as in the case where the impact of tolerance in parts is to be investigated, but this is not necessary in this case as we are trying to explore the trends out to the limits with no particular weighting and so a uniform distribution is used.


In this case, the Monte Carlo analysis will consider the four variables in question evenly and linearly spaced over an order of magnitude, approximately centred around the values used in the simulations shown in figure 5.8, with those having been selected to mimic behaviour similar to that found in the earlier finite element analyses. Specifically, the ranges for the variables are:

- $R_{th,x}$  in the range of 0.05 to 0.5 °C/W
- $R_{th,y}$  in the range of 0.05 to 0.5 °C/W
- $R_{th,a}$  in the range of 100 to 1000 °C/W
- Total power in the range of 100W to 1000W


Under every combination of these variables, a number of these simulations are performed in the same manner as in the case of the FEM-based analysis conducted previously. Specifically, in turn the total power is passed through 1 device, 2 devices, 4, 6, 9, 12, 16, 20 and finally 25 devices, spaced evenly over the heatsink in question, then fitted to equation 5.2. The quality of this fit is then found, expressed as the standard deviation between the fitted curve and the nine points corresponding to each test case. In this case, standard deviation is not expressed in terms of temperature owing to the wide range of temperatures that are experienced over such a wide range of parameters, with peak heatsink temperatures varying from 0.3°C and 534°C. Therefore, the standard deviation is expressed as a percentage of the total temperature range seen on the heatsink, therefore normalising it.

Plotting the data has some challenges also, with it not being possible to display this normalised standard deviation figure with respect to four dimensions simultaneously. Instead, it is plotted with respect to two dimensions, with the other two variables set to a fixed value in the middle of their potential range. Figure 5.10a shows this with  $R_{th,x}$  and  $R_{th,y}$  as the two axes, as these two variables are the principle representatives of geometry (as well as material properties), and considering these two together will allow exploration of the impact of ‘squareness’, i.e. is there an impact in scenarios where  $R_{th,x}$  and  $R_{th,y}$  are not approximately equal. Similarly, figure 5.10b shows the case where total power and  $R_{th,a}$  are the two axes, these being representative of power into the system and capacity for power to leave the system, and may well have a relationship from this perspective.





(a)



(b)

Fig. 5.10 Two plots showing how the error between the model and the fit varies with model parameters.

The key results to take from figure 5.10 is that the maximum standard deviation from the fit is 2.63% of the full temperature range, with the worst case over the full Monte Carlo analysis being only 2.76% of the full temperature range. Figure 5.10a shows that the greatest deviation from the fit is found where  $R_{th,x}$  does not approximately equal  $R_{th,y}$  confirming the previously mentioned suspicion that the quality of the fit is worse in cases where the heatsink is far from square, though it is worth noting that even in the case where the heatsink is ten times wider than it is tall (that is to say  $R_{th,x}$  is ten times that of  $R_{th,y}$ ) the standard deviation from the fit is still only 2.63% of the full temperature range. This trend is expanded upon in figure 5.11, where the deviation for all cases in the Monte Carlo analysis is plotted with respect to  $R_{th,x}/R_{th,y}$ , where 1 represents a perfectly square heatsink, which does indeed correspond to the smallest deviation from the fits. If one assumes that the heatsink will not be more than twice as tall as it is wide (or vice versa) then the maximum deviation is a mere 0.83% of the full temperature range. Figure 5.10b shows that there is a general trend towards greater error in the fit at higher power while  $R_{th,a}$  has little impact, though the total error is still low and has a small impact when compared to that of heatsink squareness.

To conclude, it has been demonstrated that even over a very wide range of permutations of parameters, well beyond those of a practical system, the fit proposed in equation 5.2 is of a good quality over the full range, and therefore should be valid in any actual design cases. The only exception to this is in the case of heatsinks that are significantly larger in one planar direction than the other, where error in the fit does increase, but even in fairly extreme cases, the error remains relatively low. It proved impractical to investigate the impact of pad size using this same model as grids small enough to be computed in a reasonable timeframe proved so coarse as to interfere with results.




Fig. 5.11 A plot showing how error in the fit varies with how square the heatsink is.

### 5.3 Experimental Validation of Thermal Modelling

The analysis so far is based exclusively on steady-state thermal finite element modelling, and while this should have a high degree of accuracy, an experimental validation is necessary to increase confidence in these trends. The key issues in attaining accurate results from FEM is inaccurate material parameters or incomplete models. This does not call into question the validity of equations 5.2 and 5.3, however, as the model includes conductive, convective and radiative thermal mechanisms, and is also shown to be valid with a range of material properties (see table 1). Therefore, there should be no reason for failing to conform the previously defined relationship.

To minimise the impact of the sensors on measurements taken, a large heatsink and large thermal dissipative devices (power resistors) were used. Specifically, a 300mm x 300mm x 40mm heatsink (ABL 165AB3000B) was used, heated by varying numbers of TO-247 100W power resistors (Vishay LTO100F4R700JTE3), TO-247 being a package often used for transistors. A number of 2mm holes were machined from the back (fin side) of the heatsink to within 1mm of the front (device side) and packed with thermal compound to accurately measure temperature at strategic locations, namely under mounted devices, using thermocouples. A schematic of a machined heatsink can be found in Appendix D. The relatively large pad of the TO-247 power device makes the 2mm diameter channel thermally negligible, especially when packed with thermal compound. The apparatus was also photographed in infra-red, with emissivity calibrated to conform with thermocouple measurement - the thermal images for the four configurations tested are shown in figure 5.12. Power dissipation was 180W total in all cases, and the apparatus was given ample time (approximately 2 hours) to reach thermal equilibrium in ambient conditions of  $18^{\circ}\text{C} \pm 1^{\circ}\text{C}$ .

Figure 5.12 shows a good fit of the experimentally derived data to the FEM derived equation 5.2. Therefore, it is shown that the relationship predicted through analysis of computer simulation shows strong agreement with that measured in the real world.

While it would be desirable to undertake a similar experimental validation of the impact of pad size on thermal performance, as the devices get smaller the impact of the measurement equipment on the thermal junction becomes prohibitively large.




Fig. 5.12 Thermal imaging of the experimental apparatus, viewed from the front (device side) with two, four, nine and twenty-five devices mounted. Temperature scale is identical in each case, as is total power dissipation.




Fig. 5.13 A plot of peak heatsink temperature against number of devices, along with best fit conforming with equation 5.2, derived experimentally and presented in the same form as figures 5.2 and 5.4.

## 5.4 Analytical Thermal Optimisation

The two trends shown in equations 5.2 and 5.3 describe how peak heatsink temperature varies with number of devices and thermal pad area. These are the two key variations in the device-heatsink interface that present themselves as a system uses a larger number of lower voltage rated devices, as would be the case in a multilevel converter of increasing order.

To evaluate real-world thermal performance, a dataset of commercial devices was compiled, along with a number of relevant thermal performance metrics extracted from manufacturer datasheets. The full list of the devices considered can be found in appendix A. All devices considered are rated to between thirty and fifty amps, in line with the reference converter specification outlined in section 1. The key device parameters extracted from the datasheets of those listed in appendix A are:

- maximum drain-source voltage rating,
- maximum junction temperature,
- thermal contact area,
- thermal resistance from junction to thermal contact area,
- whether it has an electrically insulated thermal contact.

Evaluating junction temperature of a device under a given set of conditions is now possible. The number of levels in a converter determines whether a device is suitably electrically rated, as a greater number of levels reducing voltage stresses on each device, in same fashion as described in chapter 4. The number of levels in the converter will give the number of thermal pads, and the compiled dataset will yield the pad size – this enables calculation of the peak heatsink temperature for the results of the single FEM analysis conducted at any state. The junction temperature is then the sum of the peak heatsink temperature and the temperature rise due to the power dissipated in each device and the thermal resistance to the heatsink. The thermal resistance from the junction to the case was extracted from the datasheet, a representative thermal conductivity of  $5\text{WK}^{-1}\text{m}^{-1}$  for thermal grease, and the thermal resistance of an insulated medium (if required), in this case using a typical thermal conductivity of a Kapton insulator of  $0.46\text{WK}^{-1}\text{m}^{-1}$ .

Some of these devices are surface mount and are designed to sink their heat through the PCB that they are electrically bonded to. To calculate the junction temperature in these cases, these devices will be considered as mounted to an aluminium substrate PCB. These integrate an insulator and have a typical thermal conductivity of  $1\text{WK}^{-1}\text{m}^{-1}$ , according to a reference page from Epectec [1].

To explore the total impact of multilevel converters on thermal performance - given a specific heatsink, power dissipation and ambient temperature – FEM is used to calculate the peak and average

temperature for a given of pad size and number thereof. From this one set of results, for every number of levels (and, by extension, number of devices), the junction temperature is calculated for every device of a sufficient voltage rating using equations 5.2 and 5.3. The optimal device under any given conditions is the one with the greatest ‘junction temperature margin’, which is the difference between the calculated junction temperature and its rated maximum, as found on the manufacturer datasheet.

The results of this method for a given heatsink are shown in figures 5.14 and 5.15. Figure 5.14 shows the trend for a total power dissipation of 30W, while figure 5.15 shows the trend for a system with the same peak and average heatsink temperature but with a total dissipation of 200W. The points on the plots represent the maximum junction temperature margin that any device is capable of offering under those conditions. Results for aluminium-backed PCB bonded heatsinking and conventional, external heatsinking are shown separately, as noted in the legend. The heatsink used in this case is ‘Heatsink 2’ as defined in table 5.1.

Figures 5.14 and 5.15 show that there are significant benefits to be gained from distributing a thermal load across a larger number of devices on a single heatsink, but that the benefits diminish for increasingly large numbers of devices. Figure 5.14 shows that the benefits suffer from diminishing returns even more so at lower power, because the power dissipated in each device is so low compared to its maximum power carrying capacity that there is very little temperature change from heatsink to junction.

Additionally, the smaller pad area of the lower voltage devices used in high order multilevel converters is not a significant issue, because at higher order the power flowing through each thermal pad is so small that there is no significant increase in peak heatsink temperatures, and by extension no significant increase in junction temperature.

This analysis shows that a multilevel converter can have significant thermal benefits over conventional topology. It also shows that high order multilevel converters grant slim benefits over those with a more reasonable number of levels in a thermal context. Finally, this analysis shows that the greatest benefits are to be gained in the case where a heatsink is working hard, while there are slim benefits to be gained in the case of a heatsink with large operational overhead.




Fig. 5.14 Trend in junction temperature margin with an increasing number of devices for a total power of 30W.




Fig. 5.15 Trend in junction temperature margin with an increasing number of devices for a total power of 200W.

## 5.5 Transient Thermal Performance

As the transition is made to a higher order multilevel converter, the modulation of the switching devices varies also. In a conventional single H-bridge grid-tie battery energy storage system, one pair of switching devices switches during the positive half cycle of the mains, while the other pair operates in the other half of the cycle. While the switching frequency is high enough to avoid of thermal cycling, the thermal mass of even the die being more than enough to filter out frequencies in the order of kilohertz, this may not be true of low frequency mains utility supply at 50/60Hz.

An understanding of the impact of modulation on thermal cycling is important, as depth of thermal cycle can have as much, if not more, of an impact on reliability as average junction temperature [2, 3]. The mechanism of thermal cycling failure is predominantly the thermal gradient in the internal elements of the device, stressing bonding media between these internal elements of a semiconductor device, as they have differing coefficients of thermal expansion.

In a cascaded H-bridge multilevel converter, only one bridge is switching at any time, so any single device will be switching less of the time. Devices in a multilevel converter, unlike a normal single bridge converter, also spend time in the on-state for significant periods of time (multiple milliseconds) without switching. Therefore, as multilevel converter order increases, the time during which a device is left in a given state can vary, as does the time during which switching losses are accrued. While there are numerous modulation strategies that can be used, a comparison between a simple modulation strategy in a single bridge converter and that in a multilevel converter is illustrated in figure 5.16.

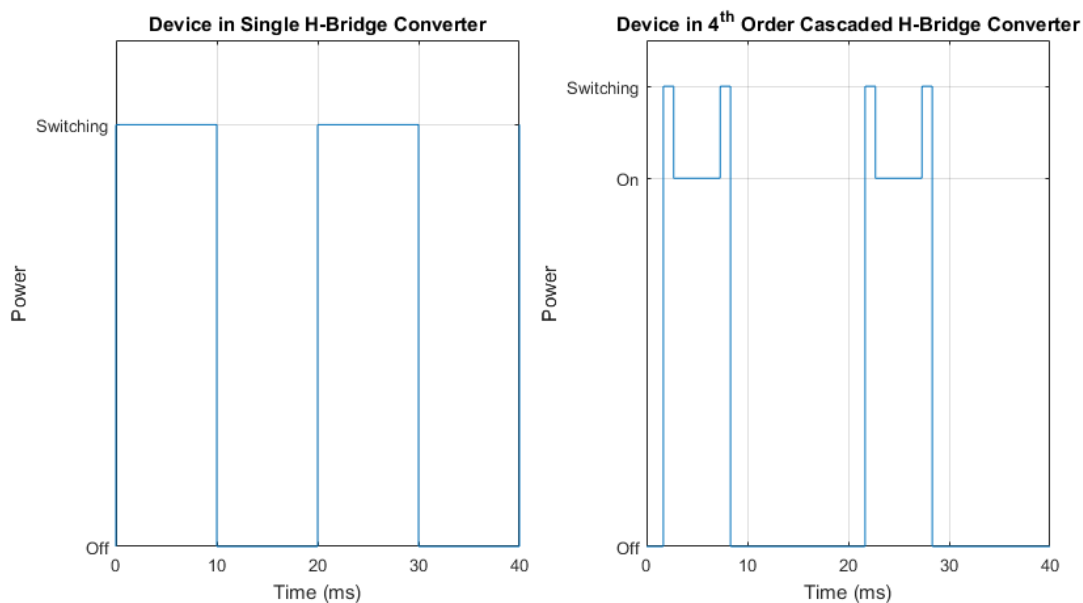


Fig. 5.16 Two plots comparing the modulation of a single device in a conventional converter to that in a typical multilevel counterpart.



Anecdotally, semiconductor thermal systems have enough thermal mass that modulation frequencies in the order of tens of Hertz can be neglected. While it seems reasonable to accept this on the scale of a heatsink, or even a device package, is that necessarily true on the scale of the die? Some of the lower voltage rated devices considered (as low as 30V 40A) have very small die elements which therefore have very small thermal mass compared to, for instance, that of a 600V IGBT.

Obtaining any data regarding die size in commercial products from their manufacturers proved impossible, so a device was dissected, and the die element measured. While the smallest devices proved impractical to dissect, and the measurements were somewhat imprecise, these figures were rounded down to one significant figure to form a lower bound of thermal mass. A standard Foster thermal model was used [4], as shown in figure 5.17, with the parameters used listed in table 3. While many of the parameters in table 5.3 vary with temperature, the model is only meant as a first order approximation, so figures found at a fixed temperature were deemed sufficient. Figure 5.17 shows the heatsink as a DC voltage source which represents a fixed temperature, as even a modest heatsink has sufficient thermal mass for any signals of interest to not significantly affect it. This assumption is easy to check as if there is negligible temperature variation at the package level then it must be smaller still at heatsink level. The die, the bonding medium and the package (or at least the metal thermal contact of the package) are modelled separately.

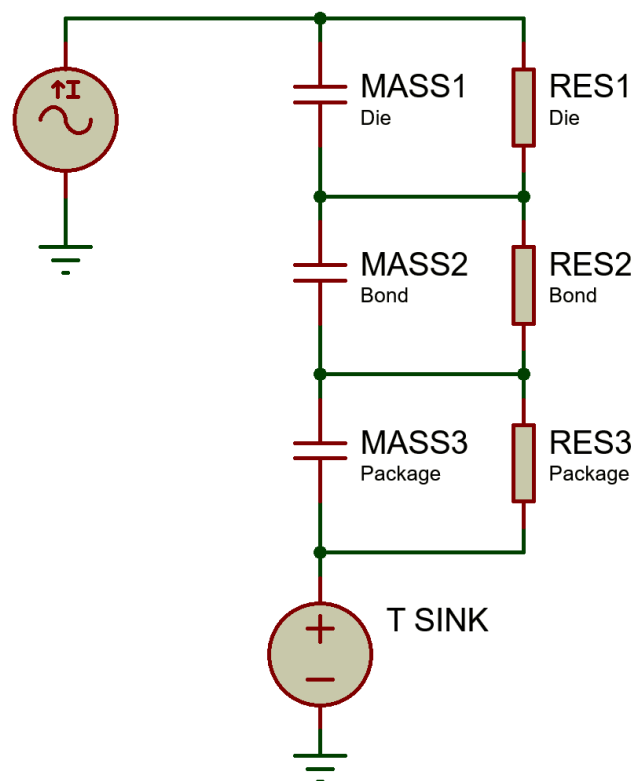


Fig. 5.17 A Foster style thermal model of a transistor, including the die, bonding medium, and package.

Parameter	Value
Silicon Thermal Conductivity	$159\text{Wm}^{-1}\text{K}^{-1}$
Silicon Specific Heat Capacity	$712\text{Jkg}^{-1}\text{K}^{-1}$
Silicon Density	$2.33 \times 10^3\text{kgm}^{-3}$
Die Dimensions	1.2mm x 1.6mm x 0.2mm
Solder Thermal Conductivity	$66.8\text{Wm}^{-1}\text{K}^{-1}$
Solder Specific Heat Capacity	$227\text{Jkg}^{-1}\text{K}^{-1}$
Solder Density	$7.29 \times 10^3\text{kgm}^{-3}$
Bond Dimensions	1.2mm x 1.6mm x 0.05mm
Copper Thermal Conductivity	$388\text{Wm}^{-1}\text{K}^{-1}$
Copper Specific Heat Capacity	$385\text{Jkg}^{-1}\text{K}^{-1}$
Copper Density	$8.96 \times 10^3\text{kgm}^{-3}$
Package Conductive Medium Dimensions	2.2mm x 1.8mm x 1mm

Table 5.3 Parameters used to calculate values in equivalent circuit shown in figure 5.17.

Using the model in figure 5.17 with the parameters in table 5.3 allows the thermal frequency response of the internals of the transistor to be calculated, which is shown in figure 5.18. Material properties in table 5.3 are readily available, and dimensions were found by dissecting an Infineon BSZ0904NS1.

Figure 5.18 shows that at frequencies below 1Hz the admittance is constant at 1.37S, and that as frequency increases the admittance decreases, showing that higher frequency power signals will result in smaller temperature variation. The left vertical-axis in figure 13 shows admittance rather than impedance to maintain the appearance of a bode plot. This plot shows that at 100Hz, the frequency of modulation in a multilevel converter, there is significant attenuation of any thermal cycling, even in the limiting case of an extremely small semiconductor device. Therefore, it is reasonable to say that the transition from a single bridge configuration to multilevel topology, and the use of physically small devices that comes with that, has no significant effect on thermal cycling of devices.




Fig. 5.18 Frequency response of Foster thermal model in figure 5.17 with parameters outlined in table 5.3.

## 5.6 Conclusions

In this chapter we have concluded that there are significant thermal benefits to be gained through the use of multilevel converters from the perspective of system thermal performance. By distributing a thermal load more evenly across a heatsink, the peak temperatures, and by extension device junction temperatures can be reduced. Furthermore, it was shown that the small thermal pad area possessed by some of the devices used in high order multilevel does little to offset the thermal benefits gained through efficient use of the heatsink. The thermal benefits were found to suffer from diminishing returns, but less so when compared to a single bridge case where the system is being operated without a large amount of thermal headroom, as would likely be the case in a commercial product. This analysis was based on trend found through FEM simulation series, but this was successfully validated to an extent experimentally.

Two of the key novel results are equations 5.2 and 5.3, which give an expression for how the peak heatsink temperature (which is linearly related to the device junction temperature) varies with the number of devices and the size of the thermal pad on those devices, for a given heatsink and total power. These relationships were initially derived from finite element modelling but were then validated both by a different type modelling, SPICE in this case, but also experimentally validated. These relationships enabled the trends which were then explored and formed the other novel contribution outlined in this chapter: how system thermal performance changes as multilevel converter gets more levels, something not done before.

This chapter concluded by exploring the possibility that the change in modulation strategy that is a result of moving to a multilevel converter might have an impact on the depth of thermal cycling that a switching device experience. This was a concern as depth of thermal cycle, just like average operating temperature, can have a large impact of reliability. It was concluded that the change in modulation is unlikely to have a significant impact even in the case of the very small devices. Load cycling will, of course, age the semiconductor devices in the converter, but that would be equally true in the case of a conventional or a multilevel converter as therefore is not of interest.

## References for Chapter 5

- [1] Epectech. (11 June 2018). *Aluminium PCBs - High Performance and Low-Cost Materials*. Available: <http://www.epectec.com/pcb/aluminum-circuit-boards.html>
- [2] S. Chang, R. Wang, Y. Xiang, and G. Chang, "Solder cracking mechanism correlation to alloy composition under thermal cycling stress (part 1)," in *2009 International Conference on Electronic Packaging Technology & High Density Packaging*, 2009, pp. 552-557.
- [3] United States Department of Defense, "MIL-HDBK-217F Reliability Prediction of Electronic Equipment," 1991.
- [4] M. Janicki and A. Napieralski, "Considerations on Electronic System Compact Thermal Models in the Form of RC Ladders," in *2019 IEEE 15th International Conference on the Experience of Designing and Application of CAD Systems (CADSM)*, 2019, pp. 1-4.

# Chapter 6:

## Reliability in Multilevel Converters

In chapter 4 the benefits that multilevel converters could give from the perspective of converter power loss were explored, and it was that there are significant potential benefits to be gained by moving to a multilevel converter topology, including high order multilevel converters. In chapter 5 the benefits of multilevel converters from the perspective of thermal performance were explored, and again it was found that there are benefits to be gained by moving to a multilevel solution over a conventional approach, but that these benefits suffer from diminishing returns with high order converters failing to grant significant additional benefits over a lower order multilevel implementation. In this chapter we discuss the impact of the use of multilevel converters in a context of overall system reliability.

### 6.1 The Arrhenius Equation

A common source of failure in any power converter is the failure of the main power switching devices [1]. It would therefore stand to reason that as the number of devices increases that the rate of failure would also increase, which is true. With converter order,  $N$ , the number of devices increases by a factor of  $N$  and therefore the mean time to failure (MTTF) decreases by a factor of  $N$ . This can also be expressed in terms of an alternative metric, the failure rate, denoted by  $\lambda$ , which is equal to  $MTTF^{-1}$ .

However, the previous chapter concluded that the use of a multilevel converter topology can improve overall system thermal performance, thereby reducing switching device junction temperature, and lower temperature operation of switching devices leads to increased reliability [2]. To balance this advantage in terms of reliability against the aforementioned disadvantage, the increase in reliability owing to lower temperatures must be quantified.

Fortunately, the relationship between temperature and reliability is well known, owing to the way in which manufacturers do reliability testing on devices. Rather than testing the devices under their use conditions where they are designed to operate for long periods without failure, the devices are tested at artificially elevated temperatures to accelerate failure in order to complete testing within a reasonable timeframe. These results from this accelerated testing timeframe can then be adjusted to represent reliability under actual operating conditions, using a relationship known as the Arrhenius equation. This same relationship can therefore be used to quantify the impact on reliability enabled by the lower junction temperatures seen in multilevel topologies. This can then be combined with the

simple reliability penalty incurred by having more devices to find the overall impact on system reliability of moving to a multilevel topology.

As reliability estimation is clearly of great interest to manufacturers and customers, there are numerous handbooks and guides issued by industry [3-5] and even rigorous military guidelines [2, 6]. There is apparent consensus on a thermal de-rating methodology, and although it may be old, the Arrhenius model is still the de facto standard method for thermal derating. Even in academia, though there is discussion on the limitations of the method [7, 8], there are still many examples of use in current research [9, 10].

The mean time to failure (MTTF), the inverse of the failure rate, is given by the expression [3]:

$$MTTF = 1/\lambda = \frac{D H A_f}{r} \quad (6.1)$$

- MTTF = mean time to failure
- $\lambda$  = failure rate, devices per unit time
- D = number of devices under test
- H = test time per device
- $A_f$  = Acceleration factor, derived from Arrhenius equation
- r = the number of failures

Where the Arrhenius equation is:

$$A_f = e^{\left[ \frac{E_a}{k} \left( \frac{1}{T_{use}} - \frac{1}{T_{stress}} \right) \right]} \quad (6.2)$$

- $E_a$  = activation energy of failure mode in question,
- k = Boltzmann's constant
- $T_{use}$  = Use temperature, absolute
- $T_{stress}$  = Test temperature, absolute

Not all of the parameters in equations 6.1 and 6.2 are relevant to this analysis, and designed from the perspective of device testing. In testing, a large number of devices are put on test to see how long before they fail, so for example, more devices for the same failure rate means higher MTTF. For our use, the relationship must be inverted, and combining with some simplification and rearrangement shows that:

$$1/MTTF \propto D e^{\frac{1}{k}\left(\frac{1}{T_{use}} - \frac{1}{T_{stress}}\right)} \quad (6.3)$$

The absolute figure for MTTF is not of interest, rather the impact of the use of multilevel converters on MTTF. As such, a new metric is created, the MTTF coefficient  $k_{MTTF}$ , defined as:

$$k_{MTTF} = \frac{1}{D e^{\frac{1}{k}\left(\frac{1}{T_{use}} - \frac{1}{T_{stress}}\right)}} \quad (6.4)$$

## 6.2 Results Using Arrhenius Model

We have now derived an expression in equation 6.4 that estimates the change in reliability as a result of a change in device junction temperature, while also including a factor to account for the total number of devices in the converter. This figure for reliability is representative of relative MTTF owing to failure of the power switching devices within the power converter. As previously mentioned, an increase in number of devices (as in the case of a higher order multilevel converter) will clearly reduce the MTTF, while a lower operational temperature would increase it. In the previous chapter it was shown that a higher order multilevel converter can decrease operational junction temperature, therefore there are factors suggesting a potential for both increased and decreased system reliability through increased converter order. By combining equation 6.4 with the results from the previous chapter, the balance between these two effects can be investigated.

The trend shown in figure 6.1 is based upon the thermal conditions shown in figure 5.9, which represented a limit case of a heatsink with plentiful thermal capacity, that is to say a large heatsink with a small amount of power dissipated in it. The MTTF coefficient is normalised so that it is unity in the case corresponding to the minimum possible number of devices: two, as in the case of a single half bridge converter. In figure 5.9 the minimal utilisation of the heatsink results in minimal thermal benefits through the use of a multilevel converter in any case, and very rapidly diminishing return, and this is reflected in the lack of improvement in reliability. Across the board the reliability is worse the more devices there are, though it is worth noting that the reliability cost has been significantly attenuated, for instance in the case of 48 devices the MTTF has reduced by a factor of 10 when it would be 24 if not accounting for thermal benefits.




Fig. 6.1 An example trend of the normalised MTTF coefficient with respect to number of devices in the converter under low power/strong heatsink conditions.




Fig. 6.2 An example trend of the normalised MTTF coefficient with respect to number of devices in the converter under high power/weak heatsink conditions.

Figure 6.2 shows the trend in reliability for the opposite limit case as in figure 5.10 where a marginal heatsink is used, that is to say a small heatsink with a lot of power dissipated into it where significant thermal benefits were found through the implementation of a multilevel solution. Again, the MTTF coefficient is normalised to unity in the case of two devices, and under these conditions that represents the worst reliability with even a converter containing 48 devices have a marginally superior MTTF. It is worth remembering that this is a limit case, where the comparison is against just

two devices operating at the absolute maximum rated junction temperature, which is not necessarily a realistic design scenario.

Figure 6.1 and 6.2 demonstrate the overall system reliability trends, while factoring in system thermal performance under the two extreme states, from a thermal perspective. In the pessimistic case shown in figure 6.1 there is a partial offsetting of the significant reliability disadvantages, particularly in higher order converters, while the optimistic case in figure 6.2 shows that there can be significant reliability benefits even in high order converters. Most realistic design scenarios would likely fall somewhere between the two, with some potential benefit (or at least minimal penalty) in the case of a multilevel converter of modest order, while a higher order converter would almost always suffer a reliability penalty over a non-multilevel solution other than under extreme conditions. This conclusion whereby a handful of levels in a converter grants benefits that fall away as the number of levels gets large is an interesting result as it appears to reinforce what many might instinctively expect, unlike those previously considered, there is a penalty for there simply being more switching devices as was not necessarily the case from the perspective of power loss.

### 6.3 Modular Multilevel Converters and Battery Failure

All of the analyses so far describe a system that consists of a single, monolithic converter. In a multilevel converter, however, there is scope for the system to be easily implemented in a modular fashion. Considering the structure of a cascaded H-bridge multilevel converter in particular, as shown in figure 6.3, each of the groups of four switching devices with its associated DC link, drive and control could form a standalone module, the combination of a number of these making up the converter as a whole. Similar to considering the use of multilevel converters in low voltage applications in the first place, this may seem like an increase in complexity and technical challenge with a lack of clear overall benefit, but this is not necessarily the case, as there is discussion in the literature regarding the benefits of modular multilevel converters (MMCs).

In [11], for example, the means by which a cascaded H-bridge converter can be used to achieve cell balancing and dynamic damaged/degraded string avoidance are demonstrated, and it also describes how this type of system is inherently modular from the perspective of overall system design. This is true with output voltage scaling being as simple as adding additional submodules, where a submodule in the of the converter shown in figure 6.3 being comprised of one H-bridge and the DC link element consisting of some number of cells. From a design perspective this is clearly desirable as a designer can design a submodule and then readily scale the overall system to any voltage with little additional effort. Exploiting this sort of modularity is the reason that MMCs are already in use in industry in medium and high voltage multilevel converter applications [12].



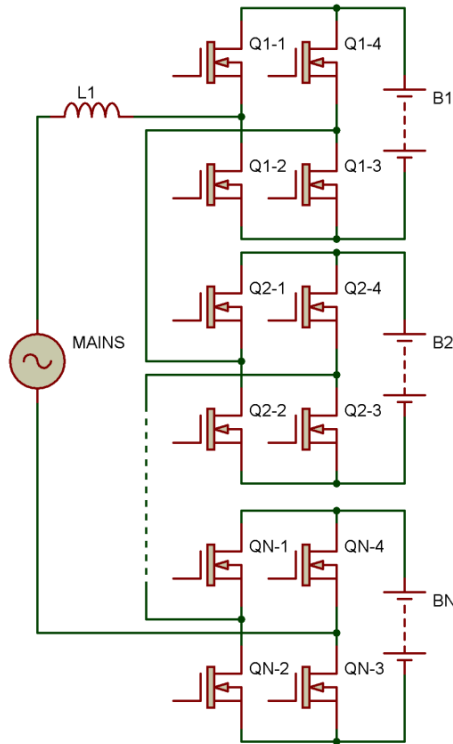


Fig. 6.3 A simplified schematic of a  $N^{\text{th}}$  order multilevel converter.

Further work such as [13-15] consider an MMC where each submodule operates as a distinct thermal entity and looks at attempting to minimise thermal excursions and outliers through control. Both as new devices, but particularly as devices age, there will be an inherent spread in device performance resulting in some devices running hotter than others or going through deeper thermal cycles, this would contribute to lower reliability as has been discussed in this chapter as well as chapter 5. This work may be looking at large MMCs used in power converters for medium and high voltage applications, but the principles also apply to our use case of lower voltage rated devices.

When looking at reliability, it is of course valuable to look at MTTF as in the first portion of this chapter, but it is also worth considering how the system is likely to respond to failure. While a semiconductor failure is an important failure mode, in the case of a battery energy storage system (BESS), another likely source of system failure comes from the cells within the battery pack itself. Catastrophic cell failure is not likely to be overcome by any topology particularly well, as it is likely to result in the physical destruction of the system, but a failure of this type is also relatively unlikely as any large battery pack will have extensive monitoring of cell voltage, temperature, etc. to avoid this case under almost any circumstances. Cell failure through gradual degradation over time, however, is a serious concern. A manufacturer's datasheet for high quality, industry standard lithium-based cells [16] only characterises capacity out to a lifetime of 250 cycles, at which point they already show significant capacity degradation – testing in the research field shows that this loss only gets worse at longer lifespans [17]. The relatively short testing periods available suggests that lifetime is likely to be

single figure years (though exact lifetime depends on a host of factors), so how a converter responds to this failure mode is very important.

Papers such as [18, 19] describe how control of a modular multilevel converter can be configured to respond to faults in BESS applications in particular. Both demonstrate that an MMC can readily have a control scheme implemented that allows overcoming battery fault states within the converter, becoming faulty suddenly or otherwise. This in tandem with papers discussed in chapter 1 showing the ability to balance load between cells that perform differently but would perhaps not be considered faulty [11, 20, 21] show that this topology shows great promise when dealing with aging batteries.

Serviceability is another key consideration that makes a modular approach appealing. Current MMC applications may be on a different scale than the proposed domestic scale BESS [22], but serviceability is still of concern. A user would surely prefer that under the conditions that part of the battery pack of their storage system was degraded, that the faulty part of the system could be readily swapped out to renew the performance of the system. Discussion previously shows that the converter may still be able to operate to some extent with faulty modules, even. As a value proposition, a converter with tolerance of faulty and degraded components with easy repair may be preferable over a converter that will perhaps last longer without intervention but then need complete replacement or refurbishment.

## 6.4 Conclusions

This chapter has gone some way in addressing one of the most obvious counterarguments to the use of a multilevel converter in applications where it may be conventionally considered unnecessary: will that not adversely affect reliability? Using the Arrhenius relationship it has been found that when one considers the improved thermal performance found through the use of the multilevel converters, that in the worst case any reliability ‘costs’ are significantly attenuated, while under best case conditions the reliability can actually improve. As for whether these best-case conditions would be likely in practical system, this work does not draw a conclusion as to that, but it is clear that switching device reliability would not be harmed to the extent that one might expect, and in edge cases could actually be improved.

While insightful, this analysis does only consider the power switching devices as the failure source and despite that being a probable failure mode, other aspects of the system such as the gate drivers are unlikely to find a significant thermal benefit to offset worsening system reliability with increasing number. This reinforces the impression that a very high order converter is unlikely to give a good reliability proposition in terms of MTTF of the power converter.

A discussion was then had surrounding how well a converter would be likely to deal with a fault, with the conclusion drawn that the failure within the converter is just as detrimental to operation in a multilevel converter as in any other case, but owing to the inherent modularity of this topology it should be easier to enable easy maintenance. Also discussed was the fact that in a battery energy storage system (BESS) the cells making up the battery pack, while outright catastrophic failure is quite unlikely in a well-designed system, will degrade to the point of being effectively faulty over a relatively short period of time. Not only is a multilevel converter capable of dealing more gracefully with degradation, but the modularity of the system can deal better with outright faults.

Overall, the reliability of the system as a whole is quite complex, but it is certainly unfair to say that the significant increase in devices brought about by a multilevel topology definitely leads to a decrease reliability. Furthermore, a multilevel converter can give many more options under certain types of system fault, particularly in the case of a BESS, as well as being a boon for easy maintenance. A 20kWh battery pack made up of 10Wh 19650 Li cells has 2000 individual cell elements, and with topologies currently used any one of those being seriously faulty leads to a failure in functionality and challenging maintenance prospects. If batteries have inherently short lifespans then perhaps ease of repair and partial replacement is more important than having extremely long converter life, and if so then a multilevel converter looks a tempting proposition.

## References for Chapter 6

- [1] H. Wang, M. Liserre, and F. Blaabjerg, "Toward Reliable Power Electronics: Challenges, Design Tools, and Opportunities," *IEEE Industrial Electronics Magazine*, vol. 7, pp. 17-26, 2013.
- [2] United States Department of Defense, "MIL-HDBK-217F Reliability Prediction of Electronic Equipment," 1991.
- [3] W. J. Vigrass. (2013, Calculation of Semiconductor Failure Rates.
- [4] S. Speaks. (2009, Reliability and MTBF Overview.
- [5] P. Ellerman. (2012, Calculating Reliability using FIT & MTTF: Arrhenius HTOL Model *MicroNote 1002*.
- [6] M. Pecht, "The Influence of Temperature on Microelectronic Device Failure Mechanisms," US Army Research Laboratory, Electronics & Power Sources Directorate 1993.
- [7] J. V. Manca, W. Wondrak, K. Croes, W. D. Ceuninck, V. D. Haeger, L. D. Schepper, *et al.*, "The Arrhenius relation for electronics in extreme temperature conditions," in *HITEN 99. Third European Conference on High Temperature Electronics. (IEEE Cat. No. 99EX372)*, 1999, pp. 29-32.
- [8] M. S. Cooper, "Investigation of Arrhenius acceleration factor for integrated circuit early life failure region with several failure mechanisms," *IEEE Transactions on Components and Packaging Technologies*, vol. 28, pp. 561-563, 2005.
- [9] Y. Yangyang and F. Chen, "Worst-case prediction using the Arrhenius model," in *2009 Annual Reliability and Maintainability Symposium*, 2009, pp. 69-71.
- [10] J. A. Mittereder, N. S. Cronk, S. C. Binari, G. D. Via, D. Fanning, H. Tserng, *et al.*, "RF Arrhenius life testing of X-band high voltage GaAs PHEMTs," in *2009 Reliability of Compound Semiconductors Digest (ROCS)*, 2009, pp. 103-108.
- [11] E. Chatzinikolaou and D. J. Rogers, "Cell SoC Balancing Using a Cascaded Full-Bridge Multilevel Converter in Battery Energy Storage Systems," *IEEE Transactions on Industrial Electronics*, vol. 63, pp. 5394-5402, 2016.
- [12] J. Rodriguez, S. Bernet, B. Wu, J. O. Pontt, and S. Kouro, "Multilevel Voltage-Source-Converter Topologies for Industrial Medium-Voltage Drives," *IEEE Transactions on Industrial Electronics*, vol. 54, pp. 2930-2945, 2007.
- [13] J. Gonçalves, D. J. Rogers, and J. Liang, "Submodule Temperature Regulation and Balancing in Modular Multilevel Converters," *IEEE Transactions on Industrial Electronics*, vol. 65, pp. 7085-7094, 2018.
- [14] F. Hahn, M. Andresen, G. Buticchi, and M. Liserre, "Thermal Analysis and Balancing for Modular Multilevel Converters in HVDC Applications," *IEEE Transactions on Power Electronics*, vol. 33, pp. 1985-1996, 2018.
- [15] M. K. Bakhshizadeh, K. Ma, P. C. Loh, and F. Blaabjerg, "Indirect thermal control for improved reliability of Modular Multilevel Converter by utilizing circulating current," in *2015 IEEE Applied Power Electronics Conference and Exposition (APEC)*, 2015, pp. 2167-2173.
- [16] Samsung. Introduction of INR18650-25R [Online].
- [17] M. J. Smith, D. T. Gladwin, and D. A. Stone, "A Comparison of the Effects of Charging Strategies on Lithium-ion Cell Performance in High Temperature Environments," presented at the IEEE-ICIT, Melbourne, Australia, 2019.
- [18] K. Xu, S. Xie, X. Wang, B. Zhang, and S. Bian, "Model Predictive-based Fault-tolerant Control for Modular Multilevel Converters Without Redundant Submodules," in *2018 IEEE International Power Electronics and Application Conference and Exposition (PEAC)*, 2018, pp. 1-5.
- [19] Y. Ma, H. Lin, X. Wang, Z. Wang, and Z. Ze, "Analysis of Battery Fault Tolerance in Modular Multilevel Converter with Integrated Battery Energy Storage System," in *2018 20th European Conference on Power Electronics and Applications (EPE'18 ECCE Europe)*, 2018, pp. P.1-P.8.
- [20] R. H. Williams, M. P. Foster, and D. A. Stone, "A Model-based Predictive Current Controller for a back-to-back connected multilevel converter aerospace starter-generator," in *IECON*

- 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society, 2016, pp. 2618-2623.
- [21] M. Vasiladiotis and A. Rufer, "Balancing control actions for cascaded H-bridge converters with integrated battery energy storage," in *2013 15th European Conference on Power Electronics and Applications (EPE)*, 2013, pp. 1-10.
- [22] M. Haeusler. HVDC Solutions for Integration of the Renewable Energy Resources [Online].

# Chapter 7:

## Conclusions and Further Work

### 7.1 Conclusions

In chapter 1 it was discussed how a cascaded H-bridge multilevel converter has been widely demonstrated to have significant potential benefits in application where it has not been conventionally considered worthwhile. This work, while valuable, was yet to address one of the first questions a designer might ask upon attempting to implement such a system: how many levels should be used? Furthermore, work in these novel multilevel converter applications focusses on specific benefits without discussion of the impact of wider system performance. The goal of this research was therefore to create practical, quantitative methods for answering the question of how many levels is best – along the way discovering the impact of wider system performance.

There are numerous aspects that should be considered when analysing system performance as a whole, the first of which considered in this work being the power loss in the converter. Chapters 2 and 3 created a method for evaluating power loss under a given set of conditions based on data that can be easily found for any devices. Chapter 3 makes particular contributions in both experimental analysis showing apparent accuracy of some datasheet information as well as enabling prediction of MOSFET body diode performance from other readily available device parameters through analysis of the results of another experimental series, leading to the publication of “An Experimental Investigation of MOSFET Intrinsic Body Diode Performance”.

Chapter 4 used the methodology outlined in chapters 2 and 3, along with data extracted from datasheets for a wide range of MOSFETs and found that there are potentially significant benefits to be found through the use of a multilevel converter in terms of system power loss. This holds true even in a high order converter and is predominantly due to the lower voltage rated devices that a higher order multilevel converter grants access to having a much lower on-state resistance – so much lower an on-state resistance that even with so many more devices in the conduction path the total resistance is still reduced. The benefits of a multilevel solution to system power loss are particularly pronounced at higher switching frequencies, where the physically larger devices associated with higher voltage ratings tend to consume a lot of power during switching transients. This shows that the implementation of a multilevel converter eases access to higher switching frequencies, which could result in smaller, lighter and cheaper converters. These results form the major part of the conference paper “Switching loss optimisation of cascaded H-bridge converters for bidirectional grid-tie battery

energy storage systems”, and later the journal paper “Quantitative Power Loss Analysis and Optimisation in Nth-order Low Voltage Multilevel Converters”.

In chapter 4 the use of alternative semiconductor technologies was discussed. While discussions regarding IGBTs and SiC MOSFETs largely concluded that in this application there is little benefit to be gained if any, there was promise in the use on GaN power devices. GaN devices capable of significant power applications are a relatively new class of device and are not yet widely available in both high voltage and current ratings, therefore the use of a multilevel converter can permit the use of these devices rated at a lower voltage. These devices were modelled in a very similar way as the MOSFETs and had a profound effect on the potential for reduced power loss, leading to much reduced loss under many conditions. This was particularly true at higher switching frequencies, owing to the GaN devices miniscule gate turn-on energy, with system power loss savings up to 73%. The benefits were less so with very high order converters as the lowest voltage rating GaN devices considered with 60V, which cannot compete with the low on-state resistance of 30V and 25V Si MOSFETs. As these results were so striking, an experimental validation of gate-turn on energy for a small sample of GaN and Si MOSFET devices was performed to ensure no error in the model – it was shown that there was no error in the model. This was published in IEEE Electronics Letters as “On the impact of current generation commercial gallium nitride power transistors on power converter loss”.

Chapter 5 continued to explore system performance with multilevel converters of increasing order, but this time from the perspective of system thermal performance. Specifically, the work proposes that as a given thermal power dissipation is spread over more devices mounted on a heatsink, that heatsink is more efficiently utilised and leads to less acute hotspots and therefore lower junction temperatures. This proved to be true, and also was found that this trend can be easily predicted. An equation was proposed where all coefficients can be easily found from a single instance of, for instance, a finite element model, therefore enabling analysis of trends without a prohibitive computational workload. This trend was found using finite element modelling and validated using a thermal-electrical equivalent circuit model as well as experimentally. This trend was shown to hold true independent of the conditions of the thermal system, this was shown particularly rigorously using the thermal-electrical equivalent model to perform a Monte Carlo analysis. This work was published as “An Investigation into the Thermal Benefits of Multilevel Converters”.

The work concluded in chapter 6 where a discussion was had regarding reliability, this being one of the most obvious counterarguments to the use of a multilevel converter where it is not necessary – conventional wisdom suggesting more devices, more problems. An analysis of mean time to failure of the main power devices showed that the thermal benefits that were found and quantified in chapter 5 at the very least offset the reduction in reliability that might be expected, and in some conditions a multilevel implementation can actually have improved reliability of the power switching devices over

a conventional converter, though with increasing converter order any such benefits diminish. Chapter 6 then went on to discuss how the inherent modularity can help the system as a whole deal with failure, both unexpected as in power devices and inevitable as in the case of a large battery bank. It concludes by suggesting that if a battery has such a short lifetime then perhaps how gracefully a system can manage battery degradation matters at least as much as how long it is before something else breaks.

Overall, for the first time, this work has provided practical, data-driven analysis of multilevel converter performance in the reference use case in terms of system power loss, system thermal performance and reliability of switching devices. This work can help a potential designer take that first step, as well as showing that multilevel converters in an application such as this can have benefits not widely discussed in literature (such as reduced power loss) and not suffer nearly as badly from issues that might be expected such as worsened reliability. Together, this removes some of the barriers to this technology moving from academic interest to practical, commercial use.

## 7.2 Further Work

This body of work has fallen slightly short of being able to enable a complete answer to the question ‘exactly how many levels should be used in a particular application?’. The work has considered the impact on the system in terms of power loss, thermal performance and reliability, but there are other key system factors that should be considered, including: the design of the main system inductor; quantifying the benefits of a reduced BMS as well as the cost of increased system complexity; quantifying the potential benefits of a scalable, modular system; as well as good models for system cost and physical size.

Upon completion of these additional models, a multivariate optimisation can be performed to find the optimal configuration for a given specification, though this would also require the creation of a cost function that could be modified to reflect a designer’s priorities. This cost function would not be a trivial undertaking as how would one, for instance, balance system size against the additional complexity in design. Furthermore, later models could feed back in to earlier ones, for example the results of a thermal optimisation could feed back in to the estimation of the device parameters (such as  $R_{DS,ON}$ ) in the evaluation of system power loss.

Furthermore, the work that has been performed can be extended. This is particularly true in the case of quantifying reliability as the only quantitative analysis is in the reliability of the switching devices. This would need extending to other system components, such as gate drivers, but also the discussion regarding the benefits of modularity and the ability of multilevel converter to overcome battery degradation need a quantitative element.



# Appendices

## Appendix A

In the interest of gaining experience with building power electronics, a straightforward hardware prototype based on the converter shown in figure 1.4b was constructed. This served the secondary objective of possibly building upon the work of a previous student, though this work never materialised owing to finding a research niche elsewhere.

For convenience, the entire system was transformer isolated from mains and operated at low voltage: with a 12V AC mains supply and a pair of 24V lead acid (VRLA) battery strings, the DC supplies on each side of the link must exceed the peak mains voltage at any time, hence 24V. Lead acid was selected as a chemistry for simplicity and also to avoid risks and complexity associated with charging lithium batteries, such as balancing and inadvertant overvoltage.

Control was maintained at its simplest, in order attain results that show as many as possible of the artefacts associated with a overly simplified control topology, and to enable rapid prototype development. A sinusoidal current reference derived from the mains voltage was fed in to a comparator along with a measurement of actual current, directly driving the half bridge. The secondary half bridge was not utilised.

After numerous design revisions, the system functioned well enough for its performance to be evaluated. Said performance is shown in figure 1.5 and is shown here again in figure 8.1.




Fig. 8.1 Behaviour of the prototype converter of the design shown in figure 1.4b

The upper plot in figure 8.1 shows the converter's performance when acting as a rectifier, i.e. storing energy from the grid in the batteries. The performance is non-ideal, but functional. The noise on the sampled mains voltage (the blue line) is due to the transformer isolation, as the voltage source therefore has the dominant inductance in the system, making it difficult to measure without the switching noise present. However, this with no precise analysis it is clear that a device with performance such as this would be in gross breach of any EMC legislation. Not only due to the magnitude of the switching noise, but also in that it is not fixed frequency, making it difficult to characterise. However, as this device was never intended for actual use, but more as an experimental platform and learning exercise, this is acceptable.

When operating as an inverter, however, the performance is worse. While the broad trend of the curves is as desired, the noise is much worse than when operating as a rectifier. This is due the voltage in the battery strings sagging under load, exacerbated by the fact that the batteries are operating at relatively high C-rate, and some the batteries in the strings utilised were hardly in pristine condition.

There is, however, a more noticeable effect of the degraded battery strings shown in figure 8.1, namely the smooth sag in the waveform seen at the peak of the positive current output. This is due to the fact that one section of one of the battery strings had very poor state of health and could not support the current demanded, resulting in the complete failure to attain the required output, despite the cells running well below their ratings. This demonstrates how the performance of the entire system can be severely limited just one small element of a battery string.

Overall, this hardware prototype was a success in as much as it significantly streamlined the process of building future hardware, and also demonstrating one of the key issues associated with use of single-level converters for grid-battery interface applications: lack of robustness to degraded elements within a battery string and no capacity for cell balancing.

## Appendix B

To explore the error induced by using a linearising simplification while extracting capacitance voltage curves from manufacturer datasheets, an analysis was performed. This analysis compared the charge integral over voltage, with respect to the maximum voltage of the integral, between the precise datasheet curves and the linearising simplification. The nature of the linearising simplification is shown in figure 3.3. Only a small subset of devices was considered, using devices with a representative range of drain-source voltage ratings, and they are listed in table 9.1.

IXYS MKE38RK600DFELB	Fairchild FCH47N60N	IR IRFP260N
Toshiba 2SK2967	Fairchild FQA44N30	Inf IPP200N15N3
ST STF100N10F7	Fairchild FDMS86103L	Inf IPD30N08S2
Inf BSZ042N06NS	Inf IRFZ44NPBF	Inf BSZ042N04NS
TI CSD17573Q5B	Inf IRF6717MTR1	

Table 9.1 Devices used in evaluation of error in approximation, with results shown in figures 9.1-9.3.

Abbreviations: ST = STMicroelectronics, IR = International Rectifier, Inf = Infineon, TI = Texas Instruments.

The error between the linearised estimation and the precisely extracted data is shown for the three different capacitances in figures 9.1 to 9.3.




Fig. 9.1 Range of error incurred in estimation of input capacitance,  $C_{iss}$ , for the devices listed in table 9.1.




Fig. 9.2 Range of error incurred in estimation of output capacitance,  $C_{oss}$ , for the devices listed in table 9.1.




Fig. 9.3 Range of error incurred in estimation of feedback capacitance,  $C_{rss}$ , for the devices listed in table 9.1.

Figure 9.1 shows the error between the integral of the input capacitance,  $C_{iss}$ , with respect to drain-source voltage,  $V_{DS}$ , of the estimation as compared with the precisely extracted datasheet curve, plotted over the range of  $V_{DS}$  up to its maximum rating. The reader will recall that the input capacitance has the extreme simplification of being approximated by a constant value. The error ranges from approximately -50% to +20%. The results also highlight some of the shortcomings of

human data entry – for instance, the approximation with the most extreme deviation at -50% goes on to almost zero error at the top end of its voltage range, while a better balance would clearly have been struck by estimating high and having a lower maximum error.

Figure 9.2 shows similar error of the estimation as compared to precise datasheet information, but in this case the integral of the output capacitance,  $C_{oss}$ , with respect to  $V_{DS}$  is considered. The estimation in this case is bilinear, as shown in figure 3.3. The error in this case is higher than the case for  $C_{iss}$ , with maximum error from approximately -30% to +80%. This is again exacerbated by the shortcomings of human data entry, with this being a relatively difficult estimation to make accurately, particularly when one remembers it is a linear simplification made from figures that usually have one or more axes expressed logarithmically. The peak in error on each of the curves occurs at the corner of the bilinear approximation, as would be expected even if the human element were perfect.

Finally, figure 9.3 shows the estimation error for the integral of the feedback capacitance,  $C_{rss}$ , with respect to  $V_{DS}$ . The error is larger again in this case with maximum error from -30% to +140%. These error values could all be improved by revising the estimation parameters, but that would defeat the point, this is a true and fair representation of the human error and difficulty involved in this process. Exploration of options that slightly increased complexity did not yield significantly better results, for instance an exponential or  $x^N$  fit with a similar corner selection process. While some fits may have been better in some cases, it was worse in others and made the human estimation stage more challenging – the bilinear estimation was concluded to be the best compromise.

## Appendix C

List of all silicon MOSFET devices used in analysis in chapter 2:

Inf BSB165N15NZ3	Inf BSZ040N06LS5	Inf IPB65R045C7	Inf IPW65R045C7	IXYS MKE38RK600
Inf IPP65R045C7	ST STW62N65M5	Toshiba TK49N65W	ST STW56N60DM2	Fairchild FCH47N60N
ST STW56N65M2	ST STI57N65M5	Inf BSC320N20NS3	Inf IPP320N20N3	IR IRFP260N
Fairchild FDP2710	IR IRFP4229	Vishay SUM45N25	Toshiba 2SK2995	Toshiba 2SK2967
Inf AUIRFP4409	Toshiba 2SK3176	ST STB40NF20	Inf IRFB4137	Fairchild FQA44N30
IXYS IXFT50N30Q3	Inf IPA075N15N3	IR IRFI4228	Inf BSB165N15NZ3G	Inf IPP200N15N3
Inf BSC190N15NS3	Fairchild FDMS86255	Fairchild FDMS86200	Inf BSC360N15NS3G	ST STF100N10F7
Toshiba TK40A10N1	Vishay IRFP064PBF	Inf IRFI4410ZPBF	Toshiba TK34A10N1	Fairchild FDMS86103
Toshiba TK46A08N1	Fairchild FDMC86340	Toshiba TK35A08N1	Renesas RJK0852DPB	Inf IPD30N08S2
Inf BSC076N06NS3G	Inf IPD053N06N	Fairchild FDD86540	ST TF100N6F7	Inf BSZ042N06NS
Toshiba TPCA8048-H	Inf IRFI1010N	Inf IPD30N06S2	Inf AUIRFZ46NL	Inf IRFZ44NPBF
Vishay SQD50N05-11	Inf AUIRFZ44N	Toshiba TK50P04M1	Toshiba TPCA8045-H	Inf BSZ042N04NS
Toshiba TPCA8015-H	Vishay SQD50N04-4	Inf IPD50N04S4L-08	ONsemi NVTFS5811	TI CSD17573Q5B
Toshiba TPCA8026	Fairchild FDMC8010	Inf IRFH5301TR2PBF	Inf BSZ0904NSI	Inf IRF6717MTR1
Vishay SIR862DP-T1	Fairchild FDMS3602S	Inf BSZ036NE2LS		

List of all EPC eGan devices used in supplementary analysis:

EPC EPC2034	EPC EPC2023	EPC EPC2015C	EPC EPC2029	EPC EPC2032	EPC EPC2033
-------------	-------------	--------------	-------------	-------------	-------------

URL to download .xlsx of database compiled:

[www.sheffield.ac.uk/eee/research/emd/fetdb](http://www.sheffield.ac.uk/eee/research/emd/fetdb)

## Appendix D

List of all silicon MOSFET devices used in diode characterisation:

AUIRF1010	IPP200N15	IXFL210N30	FDMS86255	IRF840	SQJA86EP
AUIRFP4409	IPP320N20	MDP1921	FDN359AN	IRFB4227	STP16NF06
BSC076N06	IPW60R041	MKE38RK600	FDN8601	IRFML8244	STP55NF06
BSP318S	IRF3315	NVTFS5811	FDU3N40	IRFP064	STY145N65
BSZ036NE2	IRF530N	PSMN1R2	FQA44N30	IRFP4229	TK40A10
BSZ042N04	IRF640	RFD14N05	IPA65R280	IRFP4668	TK49N65
FCH47N60	IRF740	SI4840	IPB034N03	IRFS7734	TK72A12
FDBL86210	IRF7493	SPP20N60	IPB65R045	IRFU224	TPCA8026
FDL100N50	IRL8113	SQD50N05	IPB65R660	IRL2703	TPCA8048

## Appendix E

Schematic of heatsink machine for experimental analysis in chapter 4

