THE UNIVERSITY OF SHEFFIELD

Department of Electronic and Electrical Engineering

Electrical Machines and Drives Research Group

Thesis submitted in partial fulfilment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Advanced thermal modelling and management techniques to improve power density in next generation power electronics

by

Jonathan Neil Davidson

Supervised by Prof David Andrew Stone and Dr Martin Paul Foster

March 2015

to Mam, Dad, Katharine,
Chris and Joshua

Summary

This thesis sets out a series of new techniques to improve the thermal management of power electronics. The work is motivated by the increasing impetus to design smaller, more energy efficient electronic power systems for a range of applications, notably electric vehicles. Thermal management is an increasingly important tool which can facilitate improvements in power density through better monitoring and control of system temperatures. This thesis seeks to deliver improvements in implementing this strategy.

A review of the state of the art in thermal management is reported, focussing on temperature measurement, thermal characterisation and system modelling techniques. In addition, novel techniques for arbitrary dissipation control and die temperature measurements in semiconductor devices are presented. A novel analysis of the limitations of low-order thermal models is also described. Improvements and applications of these techniques form the basis of this thesis.

The pseudorandom binary sequence (PRBS) technique for system identification is applied throughout the thesis to characterise thermal systems. A mathematical analysis is provided, together with a novel technique to determine the minimum gain which can be identified by PRBS techniques in the presence of noise. A novel improvement to the PRBS technique for typically ten times more noise resilient measurements is then developed based on mathematical mixing of different frequency PRBS signals. In parallel, a novel technique is formulated to estimate the temperature throughout a multiple device system using digital IIR filters and PRBS thermal characterisation, which achieves errors of 3-5% when demonstrated practically. By combining these techniques, a comprehensive temperature estimation and control methodology is implemented for a multiple device system under active cooling. Finally, the expansion of the proposed methodologies to steady-state die temperature estimation is presented with comparable accuracy to surface temperature measurements, increasing the usefulness of the developed techniques in a practical setting.

Publications

Portions of the work contained in this thesis have been disseminated in the following internationally published works:-

Journal publications

[1]**JN Davidson**, DA Stone and MP Foster, '*Arbitrary waveform power controller for thermal measurements of semiconductor devices*', Electronics Letters 48, (7), March 2012, pp. 400-402

[2]**JN Davidson**, DA Stone and MP Foster, '*Minimum gain identifiable when pseudorandom binary sequences are used for system identification in noisy conditions*', Electronics Letters, 49, (22), October 2013, pp. 1388-1389

[3]**JN Davidson**, DA Stone and MP Foster, '*Required Cauer network order for modelling of thermal transfer impedance*', Electronics Letters, 50, (4), February 2014, pp. 260-262

[4]**JN Davidson**, DA Stone, MP Foster and DT Gladwin, '*Improved bandwidth and noise resilience in thermal impedance spectroscopy by mixing PRBS signals*', IEEE Transactions on Power Electronics, 29, (9), September 2014, pp. 4817-4828

[5]**JN Davidson**, DA Stone, and MP Foster, '*Real-time prediction of power electronic device temperatures using PRBS-generated frequency-domain thermal cross-coupling characteristics*', IEEE Transactions on Power Electronics, 30, (6), June 2015, pp. 2950-2961

Conference proceedings

[6]**JN Davidson**, DA Stone and MP Foster, '*Mixed PRBS technique for thermal impedance spectroscopy of power electronic systems*', Power Control and Intelligent Motion (PCIM) 2013, Nuremberg, May 2013, pp. 1047-1054

[7]**JN Davidson**, DA Stone, MP Foster and CR Gould, '*Prediction of device temperatures in an electric vehicle battery charger system by analysis of device thermal cross-coupling*', 15th European Conference on Power Electronics and Applications (EPE), Lille, September 2013

[8]**JN Davidson**, DA Stone and MP Foster, '*Real-time temperature monitoring and control for power electronic systems under variable active cooling by characterisation of device thermal transfer impedance*', 7th IET international conference on Power Electronics, Machines and Drives (PEMD), Manchester, April 2014

Manuscripts under consideration for publication in journals

[9]**JN Davidson**, DA Stone, MP Foster and DT Gladwin, (Under Review) '*Real-time temperature prediction in a multiple device power electronics system subject to dynamic cooling*', IEEE Transactions on Power Electronics

[10]**JN Davidson**, DA Stone and MP Foster, (Under Review) '*Measurement and characterisation technique for real-time die temperature prediction of MOSFET-based power electronics*', IEEE Transactions on Power Electronics

Acknowledgements

As I approach the end of my PhD, it gives me great pleasure to reflect on all the invaluable help and support I have received, and to show my appreciation for it. So many people made this PhD the pleasant journey it was. In this small space I wish to offer thanks to those special people who have had a great impact on my work.

I would firstly like to acknowledge the support and encouragement of my academic supervisors, Dave Stone and Martin Foster, who have kept me on track these last three years. Thanks also to Dan Gladwin, whose proofreading of papers has been a great help.

I would like to thank my parents, Keith and Judith, who have offered encouragement and advice throughout, and whose labours have made this thesis readable. Also thanks to Katharine, Chris, Auntie Irene and Auntie Peggy who came to Sheffield to support me.

I will acknowledge those colleagues who have helped me in their own language below.

Special thanks to Chris Davenport, who has always been there and given his friendship freely throughout my degree and PhD. Thanks also to David Anderson and Adam Samuel for their support.

Dan Schofield is acknowledged for encouragement, technical assistance and support.

我要感謝曾志華的幫助. (Chi Tsang)

با تشکر از شهاب نژاد برای حمایت و تشویق (Shahab Nejad)

I'd like to thank Dave Hewitt for the interesting discussions and help in the lab.

Je voudrais remercier le gars du bureau, Dalil Benchebra, pour ses mathématiques étonnantes.

Quisiera expresar mi agradecimiento a Jonathan Gómez por sus consejos, bromas y multímetro!

Finally, thanks to Andy Fairweather, Dan Rogers, James Green, Huw Price and Chris Gould for their help and advice.

Contents

Nomenclature

Symbol	Meaning	Context (if not a general symbol)
a	bipolar power amplitude	power spectral density of PRBS
$\boldsymbol{a}$	a vector of IIR filter coefficients	IIR filters
A	a constant	
A	junction area	diode forward voltage drop
A	an example PRBS	Chapter V
$\boldsymbol{b}$	a vector of IIR filter coefficients	IIR filters
B	an example PRBS	Chapter V
c	constant	
C	electrical capacitance	
C_{ox}	MOSFET oxide capacitance	
C_{th}	thermal capaciity	
d	constant	
D	diffusivity coefficient	diode forward voltage drop
e	elementary charge	diode forward voltage drop
e	exponential constant	
e	PID error signal	PID control analysis
E_{g}	bandgap	diode forward voltage drop
f	frequency	
f_{c}	corner frequency	low pass filters
f_{P}	PRBS frequency	
f_{P}^{-}	low frequency PRBS clock rate	

Symbol	Meaning	Context (if not a general symbol)
f_{P}^{+}	high frequency PRBS clock rate	
f_{s}	sampling frequency	
$\mathcal{F}$	discrete Fourier transform	
F_0	DC value of a PRBS frequency spectrum	
G	length of vector $\boldsymbol{a}$	IIR filters
h	Planck constant	
H	an unknown system (or its frequency response)	
$\widetilde{H}$	measured gain of H under noisy conditions	
H	length of vetor $\boldsymbol{b}$	IIR filters
i	index of discrete time	
i_b	base current	
i_{bias}	bias current	
i_{c}	collector current	
i_{d}	drain current	
i_{e}	emitter current	
I	current	
I_{s}	saturation current of a diode	
j	imaginary unity	
k	Boltzmann constant	
k	index of discrete frequency	
K_{p}	PID controller proportional gain	
K_{u}	ulimate gain of a system	

Symbol	Meaning	Context (if not a general symbol)
L	MOSFET channel length	MOSFET transients
L	minority carrier diffusion length	diode forward voltage drop
L	greatest of G and H	IIR filters
$\mathcal{L}$	Laplace transform	
m^*	effective mass of electron (m_e^*) or hole (m_h^*)	
n	PRBS bit length	
n_av	number of samples averaged	
n_i	intrisnic carrier density	
n_s	number of samples	
N	density of acceptors (N_a) or donors (N_d)	MOSFET transients
N	shorthand for $2^n - 1$	PRBS equations
p	pole of a system	
$\wp$	the PRBS function	
q	power dissipation or heat flux (time domain)	
Q	power dissipation or heat flux (frequency domain)	
Q	power amplitude	power spectral density of PRBS
$r_\mathrm{ds(on)}$	MOSFET channel resistance	
R	resistance	
R_th	thermal resistance	
s	Laplace operator variable	
$\boldsymbol{s}$	state vector	IIR filters
S	Seebeck constant	

Symbol	Meaning	Context (if not a general symbol)
t	time	
t_{conv}	ADC conversion time	
t_{meas}	disconnection time for measurement	
T	absolute temperature (time domain)	
T_{d}	PID controller differential period	PID control analysis
T_{i}	PID controller interation period	PID control analysis
T_{u}	ultimate period of a system	PID control analysis
v_{be}	base-emitter voltage	
v_{bias}	bias voltage	
v_{cc}	positive supply voltage	
v_{ce}	collector-emitter voltage	
$v_{\mathrm{ce(sat)}}$	BJT saturation voltage	
v_{d}	drain voltage (with respect to ground)	
v_{ds}	drain-source voltage	
v_{i}	input voltage	
v_{n}	noise voltage	
V	DFT of voltage	
V_{dd}	positive supply voltage	
V_D	diode forward voltage drop	
V_{gs}	gate-source voltage	
V_{i}	DFT of input voltage	
V_{n}	DFT of input noise	
V_{o}	DFT of output voltage	

Symbol	Meaning	Context (if not a general symbol)
V_{th}	MOSFET threshold voltage	
W	MOSFET channel width	MOSFET transients
x	ordinal of heat source	
$x(t)$	impulse-like contribution to response	Laplace system analysis
y	ordinal of temperature measurement point	
$y(t)$	steady-state contribution to response	Laplace system analysis
Z	electrical impedance	
$\mathcal{Z}$	$\mathcal{Z}$-transform	
z	$\mathcal{Z}$-transform operator variable	
Z_{th}	thermal impedance	
z_{th}	temperature response of a system	
$Z_{\mathrm{th}\cdot x\to y}$	equivalent Norton thermal impedance between x and y	
α	a device constant	diode forward voltage drop
α	thermal diffusivity	first principles thermal modelling
Δ	number of standard deviations	minimum gain/thermal impedance
δ	duty cycle	
δ	blower speed	
ϵ_{Si}	dielectric constant of silicon	MOSFET transients
λ	thermal conductivity	first principles thermal modelling
λ	a delaying parameter	IIR filters
μ	carrier mobility	diode forward voltage drops
π	fundamental constant	

Symbol	Meaning	Context (if not a general symbol)
ϕ	angle of complex number	
Φ_F	Fermi level	diode forward voltage drop
Φ_{xx}	power spectral density	
φ	scaled angular frequency in the range $0 \leq \varphi \leq \pi$	IIR filters
$\varpropto$	of the order of	
θ	relative temperature (time domain)	
θ_n	noise in temperature measurement	
Θ	relative temperature (frequency domain)	
τ	time constant	
ω	angular frequency	
ω_p	PRBS angular frequency	
ω_s	angular sampling frequency	

List of figures

List of tables

Chapter I

Introduction

This Chapter introduces the body of work which forms this thesis. It explains the motivation for the presented research and introduces the issues surrounding thermal modelling in a practical context, describing the challenges facing modern power electronics designers. Finally, it sets out the structure of the thesis and states the contribution made by each chapter.

1.1 Motivation

For the foreseeable future, the development of power electronics will be critical to the success of a wide swathe of industries. Power electronics bridges the gap between traditional analogue small-signal technologies, such as amplifiers, and traditional power technologies, such as transformers, by enabling the voltage-current transformation of all categories of electrical waveform.

Strategic visions by academic and industrial experts in the power electronics field highlight the importance of increasing technology penetration of power electronics in the automotive [1.1], aerospace [1.2], locomotive [1.3] and marine [1.4] industries. Power electronics can supplant existing technologies in these industries due to the improvements in size, weight, efficiency and, consequently, cost compared to traditional technologies. These considerations provide the impetus for power electronics research.

The development of power electronics is necessarily multifaceted, with research on-going into novel topologies [1.5], novel devices [1.6], novel semiconductor materials [1.7], novel control algorithms [1.8] and novel applications, such as local electricity distribution [1.9]. However, one area upon which an increasing number of researchers are focussing is the thermal management of power electronics. This field seeks to develop techniques for easy measurement, estimation, prediction and control of temperature as well as developing new physical arrangements of devices to improve

thermal conductivity and reduce hot spots in designs. This thesis aims to further the state of the art in this field.

1.2 Trends in design constraints for power electronics

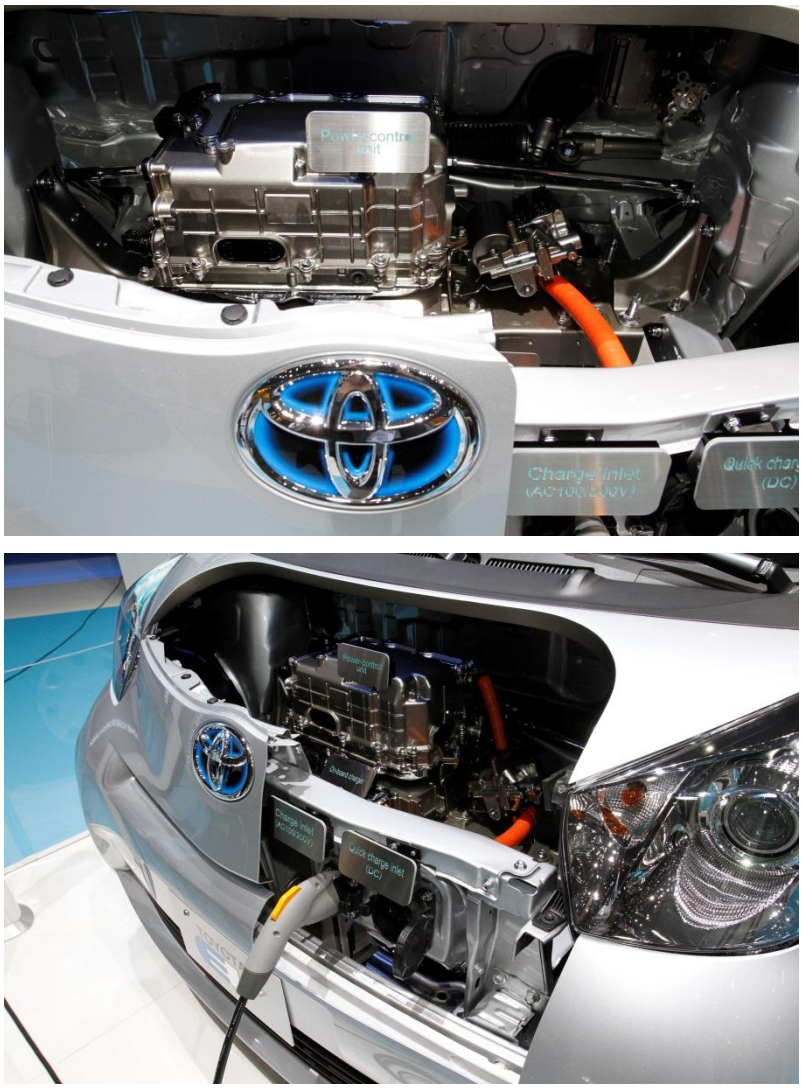

Fig. 1.1 Power electronics under the bonnet of a prototype Toyota electric car [1.10]

The design of power electronics has multiple constraints placed upon it by the application for which it is to be used. In electric vehicles, for example, the electronic design is constrained by the available space under the bonnet, for which power electronic systems, motors, batteries and the mechanical transmission system compete. Fig. 1.1 shows the battery charger power electronics for a prototype electric car, presented by Toyota in 2011. The limited space available and the large size of state-of-the-art power electronics can clearly be seen in these photographs. Additionally, for

transport applications in particular, the mass of the power electronics is an issue. In the case of motorcars, for example, each additional kilogramme of weight represents typically 70 W of additional rated power provision [1.11]. For these reasons, there is an increasingly strong impetus to produce smaller and lighter power electronic systems.

1.8 kg
circa 1990

0.9 kg
circa 2000

Fig. 1.2 Two electrically identical 300 W computer power supplies from different periods showing the size and weight reduction which has occurred over time

Reductions in size and weight are occurring through all branches of power electronics. Fig. 1.2, for example, shows the changes in computer power supplies over a 10 year period. The newer power supply, although it has identical electrical inputs, outputs and power levels, is approximately half the weight and volume of its predecessor. The driver in this case is a demand for smaller, more integrated computers.

This type of improvement can only be achieved with consideration of the thermal aspects of the design. For example, efficiency improvements reduce the losses in power electronics systems, reducing the need for large heatsinks. Novel devices materials, such as silicon carbide (SiC), can operate at higher temperatures than traditional silicon [1.12] and therefore allow increased power dissipation for a given cooling arrangement by increasing the thermal gradient between the device and ambient. However, these improvements are expensive with SiC devices costing three to five times more than silicon [1.13]. Thermal management, on the other hand, can often be incorporated into

existing microcontrollers using information already available. It is therefore a vital area of research, and one to which this thesis intends to offer a contribution.

1.3 Novelty

This thesis deals with power devices whose junctions operate under thermal steady-state conditions. That is to say, the transient effects of device switching on junction temperature are not evaluated. Instead, the effect of changes in load and therefore average dissipation are studied. The novelty presented in this thesis covers three broad areas:-

1. Improvements to and analysis of the pseudorandom binary sequence (PRBS) method of system identification

   The new techniques and analysis apply when system identification is inhibited by noise on the measurement signals. This is a common scenario for the identification of real-world systems. A new calculation of the lowest gain that can be identified by maximum length PRBS techniques in the presence of noise is presented in Chapter IV while a technique to allow the identification of lower gains by modifying the PRBS is presented in Chapter V. For a typical system, the modification results in a ten-fold reduction in the minimum gain compared to a standard PRBS technique of comparable complexity.

2. Real-time temperature estimation techniques based on Fourier analysis of thermal systems under dynamic cooling

   The new techniques convert the system response in the form of its thermal transfer impedances which can be identified using the new PRBS method. By converting these impedances into a digital IIR filter representation, the temperature response of a system to a given power input can be predicted. The technique is extended to systems under dynamic cooling (where the thermal impedances change) and displays estimations with a low error of typically 3-4% when compared to measured temperatures on the same system.

3. Circuit topologies to control power dissipation in semiconductor devices while measuring the junction or die temperature

The topologies allow the power dissipation in a semiconductor device, typically a MOSFET, to be controlled by a demand voltage. When coupled with the temperature estimation technique mentioned above, a die temperature estimator is produced. The topologies presented in this work are for a steady-state estimator with an estimation update rate of 10 Hz or less. However, the estimation techniques could equally be applied to the transient response (update rate of 10 kHz or more) following a suitably high frequency characterisation.

1.4 Thesis structure

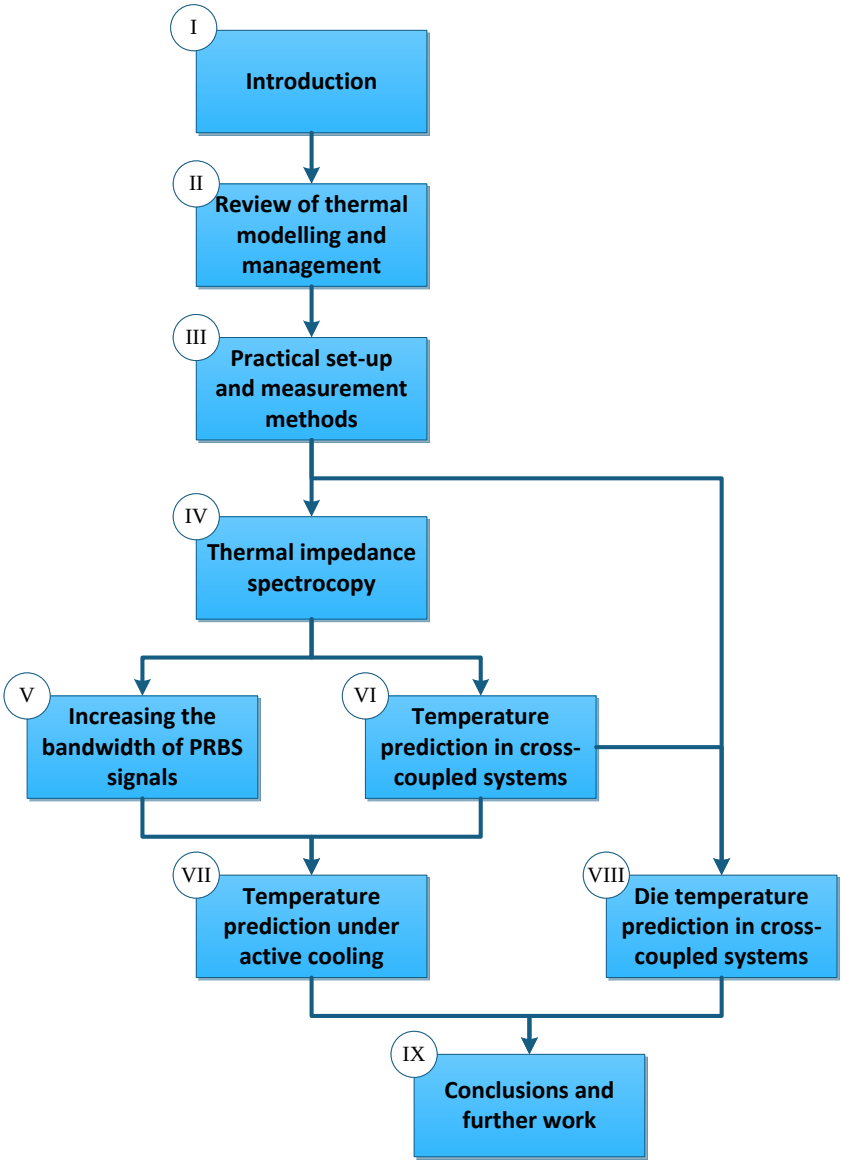

Fig. 1.3 Chapter-by-chapter structure of the thesis showing methodology development

The thesis is organised into several interconnected chapters, each describing a unit of work. The publications corresponding to each chapter can be found in section 1.5. A full list of chapter titles and the fully referenced academic papers can be found earlier in the thesis. The thesis has a predominantly linear structure, with each chapter expanding and improving on the previous. In some cases, the work diverges to explore two parallel strands of research, which are later merged to address a single objective. The thesis structure is shown diagrammatically in Fig. 1.3. A brief description of each chapter is provided below.

Chapter I introduces the thesis and explains its relevance to the modern engineering community. It provides a breakdown of the structure and summarises the thesis.

Chapter II reviews the state of the art with regards to the thermal modelling and management of power electronics, covering equivalent circuit models, temperature measurement, thermal impedance measurement, model parameterisation, temperature prediction and applications of pseudorandom binary sequences. The review identifies some of the challenges posed by current modelling techniques which will be addressed by later chapters.

Chapter III presents novel techniques for the measurement of relevant temperatures in power electronics. In particular, a new method for controlling power dissipation is proposed. A technique to measure the die temperature of semiconductor devices while simultaneously controlling dissipation is also proposed. Additionally, the Chapter sets out the experimental arrangements used in the thesis to measure temperature, characterise thermal systems and process the resulting data.

Chapter IV introduces thermal impedance spectroscopy using the pseudorandom binary sequence (PRBS) as a test signal. The Chapter presents a novel analysis of the order of complexity required of thermal models to accurately represent thermal transfer impedance. The mathematics of the PRBS is reported and a new method for the analysis of its noise resilience is provided. PRBS techniques, which are used in **Chapters V to VIII**, draw heavily on the analysis in this Chapter.

An improvement to the PRBS characterisation technique is proposed in ***Chapter V***. The method is based on mixing two PRBS signals of different frequencies using a variety of mathematical operators to produce a new PRBS-like signal. Characterisation using the new signal is valid over the combined valid frequency ranges of the source signals. The technique increases the bandwidth and improves the noise resilience of characteristics compared to standard PRBS signals. This method is later used in **Chapter VII** for rapid characterisation.

Chapter VI documents the development of a temperature estimation method based on an initial thermal characterisation of a system using the techniques developed in **Chapters III and IV**. By determining the thermal impedance between pairs of devices in the frequency domain, the temperature response of each device in a system can be estimated based on power dissipation data. The technique is implemented using digital filters to describe the thermal cross-coupling between devices. These filters can then be used to estimate temperatures in a computationally efficient manner. The methodology is demonstrated in this Chapter for estimation of surface temperatures in a four-device test system.

Chapter VII brings together the techniques from **Chapters V and VI** to produce a comprehensive temperature estimation and control method for a thermal system under active cooling. Following characterisation of a demonstrator using the mixed PRBS technique from **Chapter V**, a digital filter-based estimation method is produced. In contrast to the technique in **Chapter VI**, the new method can estimate the device surface temperatures throughout a system under variable cooling. The technique is then expanded to control the temperature by the modulation of active cooling.

The characterisation and temperature estimation techniques proposed in **Chapters V to VII** operate on surface temperatures measured by thermocouples. ***Chapter VIII*** applies the technique to estimate the temperatures of the semiconductor dice in power devices. These temperatures are more relevant for reliability considerations. Using the simultaneous die temperature measurement and power control technique from **Chapter III** together with the PRBS characterisation, a cross-coupling thermal model of the system is produced. A real-time temperature estimator is then implemented using a digital filter for high precision sensorless die temperature monitoring. The temperature estimations produced in this work are updated at a rate of 10 Hz or less and therefore do not capture the thermal effects of switching transients although they are sensitive to changes in load. However, the estimation techniques developed in **Chapters V to VII** could be used for this type of estimation if the thermal cross-coupling is characterised at sufficiently high frequency.

Finally, ***Chapter IX*** draws the work together and summarises the improvement to the state of the art which have been developed. It suggests further work which could be

undertaken following from the work presented in this thesis, and which will increase the thesis' impact on the field of thermal modelling and management of power electronics.

1.5 Publication interconnections

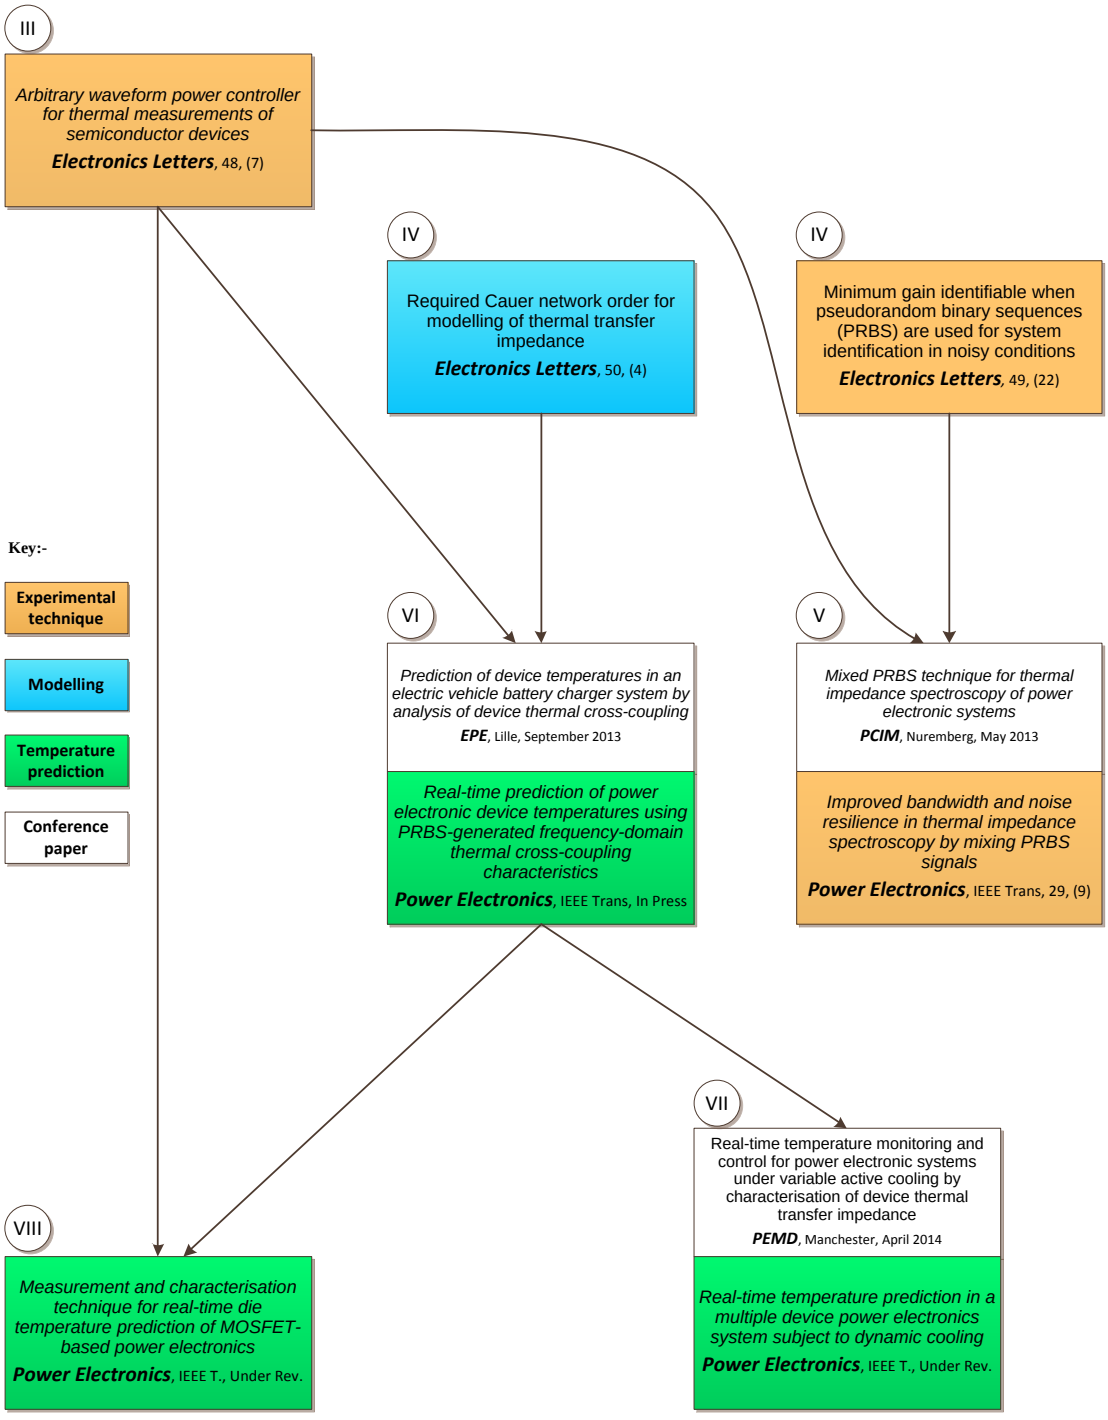

Fig. 1.4 Links between published papers with corresponding chapter numbers

Several of the techniques and analyses documented in this thesis have been published in internationally recognised journals and conference proceedings. Fig. 1.4 shows the links between papers generated by this thesis. In addition, the chapter which documents the work presented in each paper is also indicated. The papers cover three areas: experimental technique, thermal modelling and temperature estimation. Links between the papers show methodology development from experimental technique to modelling to comprehensive temperature estimation. These interconnections are broadly similar to the methodology development shown in chapter structure but with variations because there is not a one-to-one correlation between papers and chapters.

1.6 References

[1.1] S. Dusmez and A. Khaligh, '*A Charge-Nonlinear-Carrier-Controlled Reduced-Part Single-Stage Integrated Power Electronics Interface for Automotive Applications*', IEEE Transactions on Vehicular Technology, vol. 63, pp. 1091-1103, 2014.

[1.2] K. Rajashekara, '*Parallel between More Electric Aircraft and Electric/Hybrid Vehicle Power Conversion Technologies*', IEEE Electrification Magazine, vol. 2, pp. 50-60, 2014.

[1.3] F. Ma, A. Luo, X. Xu, H. Xiao, C. Wu and W. Wang, '*A Simplified Power Conditioner Based on Half-Bridge Converter for High-Speed Railway System*', IEEE Transactions on Industrial Electronics, vol. 60, pp. 728-738, 2013.

[1.4] T. Ericsen, N. Hingorani and Y. Khersonsky, '*Power electronics and future marine electrical systems*', IEEE Transactions on Industry Applications, vol. 42, pp. 155-163, 2006.

[1.5] A. Boglietti, A. El-Refaie, O. Drubel, A. Omekanda, N. Bianchi, E. Agamloh, M. Popescu, A. Di Gerlando and J. Borg Bartolo, '*Electrical Machine Topologies: Hottest Topics in the Electrical Machine Research Community*', IEEE Industrial Electronics Magazine, vol. 8, pp. 18-30, 2014.

[1.6] A. Rodriguez, D. Vega, R. Najar and M. Pina, '*Novel Electronic Devices in Macroporous Silicon: Design of FET Transistors for Power Applications*', IEEE Transactions on Electronic Devices, vol. 58, pp. 3065-3071, 2011.

[1.7] M. Werquin, C. Gaquiere, Y. Guhel, N. Vellas, D. Theron, B. Boudart, V. Hoel, M. Germain, J. C. De Jaeger and S. Delage, '*High power and linearity*

performances of gallium nitride HEMT devices on sapphire substrate', Electronics Letters, vol. 41, pp. 46-47, 2005.

[1.8] H. L. Ginn and C. Guangda, '*Digital Control Method for Grid-Connected Converters Supplied With Nonideal Voltage*', IEEE Transactions on Industrial Informatics, vol. 10, pp. 127-136, 2014.

[1.9] Z. Rui, W. Dan, M. Chengxiong, L. Jiming, Y. Jiawei, Y. Yang, C. Xun and Z. Jun Feng, '*Dual active bridge synchronous chopper control strategy in electronic power transformer*', IET Electric Power Applications, vol. 8, pp. 89-97, 2014.

[1.10] Toyota, "New Toyota EV prototype," ed. Flickr (Sunnyvale, USA), 2011. Licensed under CC BY-NC-ND 2.0.

[1.11] Wikimedia Foundation. (2014, 6 August). *Power-to-weight ratio*. Available: en.wikipedia.org/w/index.php?title=Power-to-weight_ratio&oldid=619378736

[1.12] J. Rabkowski, D. Peftitsis and H. P. Nee, '*Silicon Carbide Power Transistors: A New Era in Power Electronics Is Initiated*', IEEE Industrial Electronics Magazine, vol. 6, pp. 17-26, 2012.

[1.13] S. Allen, '*Silicon Carbide MOSFETS for High Powered Modules*', *in Proc.* Applied Power Electronics Conference (APEC), Long Beach, CA, USA, 2013.

Chapter II

Current state of technologies for thermal modelling and management of power electronic devices

In the introduction, it was seen that the predominant trend in power electronics is for smaller, more thermally managed systems. Improvements in size and thermal management can only be achieved, however, with increased accuracy in temperature estimation, for which high quality thermal modelling is required. This Chapter reports on the state of the art in thermal modelling and the areas upon which it relies, such as temperature measurement, presentation of data and system identification. The Chapter provides the reader with the underlying knowledge to appreciate the improvements proposed in subsequent chapters to thermal modelling, management and temperature control.

2.1 Relevance to the thesis

This thesis is concerned with improving the power density of power electronic systems through the advanced thermal modelling of the active semiconductor devices from which such systems are constructed. To realise this objective, an understanding is required of thermal modelling, with which complex real-life thermal systems can be simplified to allow better control and monitoring. A number of techniques used for thermal modelling are therefore introduced in this Chapter (section 2.2).

Accurate thermal models rely on high quality temperature data for parameter fitting and verification. There is a trade-off, however, between high precision measurements at the surface of devices and lower precision but more relevant temperature measurements on the semiconductor die. This Chapter reports on the available temperature measurement techniques and appraises their usefulness (section 2.3).

Many models are fitted from thermal impedance data. Methods of measuring the thermal impedance (section 2.4) and fitting models to it (section 2.6) are provided. Thermal impedance can be expressed graphically in many different ways. While the underlying data is the same, the apparent closeness of a model fit differs depending on how the data is interpreted. A discussion of the forms in which thermal impedance data is presented is therefore provided (section 2.5).

The derived thermal models of the system can be used to predict the temperature response of a device to a given stimulus. Several methods of estimating and predicting the temperature of power electronics are proposed and developed in this thesis. A discussion of the techniques reported in literature is therefore provided (section 2.7).

Because a significant portion of this work concerns the measurement of thermal impedances using pseudorandom binary sequence-based system identification techniques, the state of the art in this area is also presented (section 2.8).

2.2 Thermal modelling

2.2.1 Types of modelling

Thermal modelling is the mathematical representation of the way in which the temperature of a system responds to changes in its heating and cooling and due to other external and internal effects (such as boundary temperatures and changes in states of matter). Modelling can be broadly divided into two categories:-

1. Parametric modelling
2. Non-parametric modelling

Parametric models involve the determination of a specific finite set of unknown variables. The most frequently used type is the thermal equivalent circuit of resistors and capacitors. For this type of model, a design phase is required to define the model in terms of the particular parameters which are to be determined (for example, the arrangement of components in an equivalent circuit). The range of real-life characteristics which can be well approximated by the model is highly dependent on the design. That is to say, careful parameter setting is likely to produce a good approximation to measured characteristics only for a carefully designed model. Careless

model design, by contrast, limits model effectiveness because no combination of parameter values for the design will accurately match practical characteristics.

Non-parametric models, conversely, are produced directly by the measured characteristics themselves. One technique would be to determine the transfer function between the system's inputs and outputs directly from the measured data. No assumptions are made about the system's operation (that is, there is no design phase) and therefore the model has no parameters to determine. This means that such models are not susceptible to difficulties in model definition and parameter fitting; however, they can be more complicated in order to exactly match imperfect fitting characteristics.

Several types of model are commonly used for thermal analysis. The most relevant for the work in this thesis are described below.

2.2.2 Equivalent circuit thermal analogues

2.2.2.1 Physically derived lumped parameter models

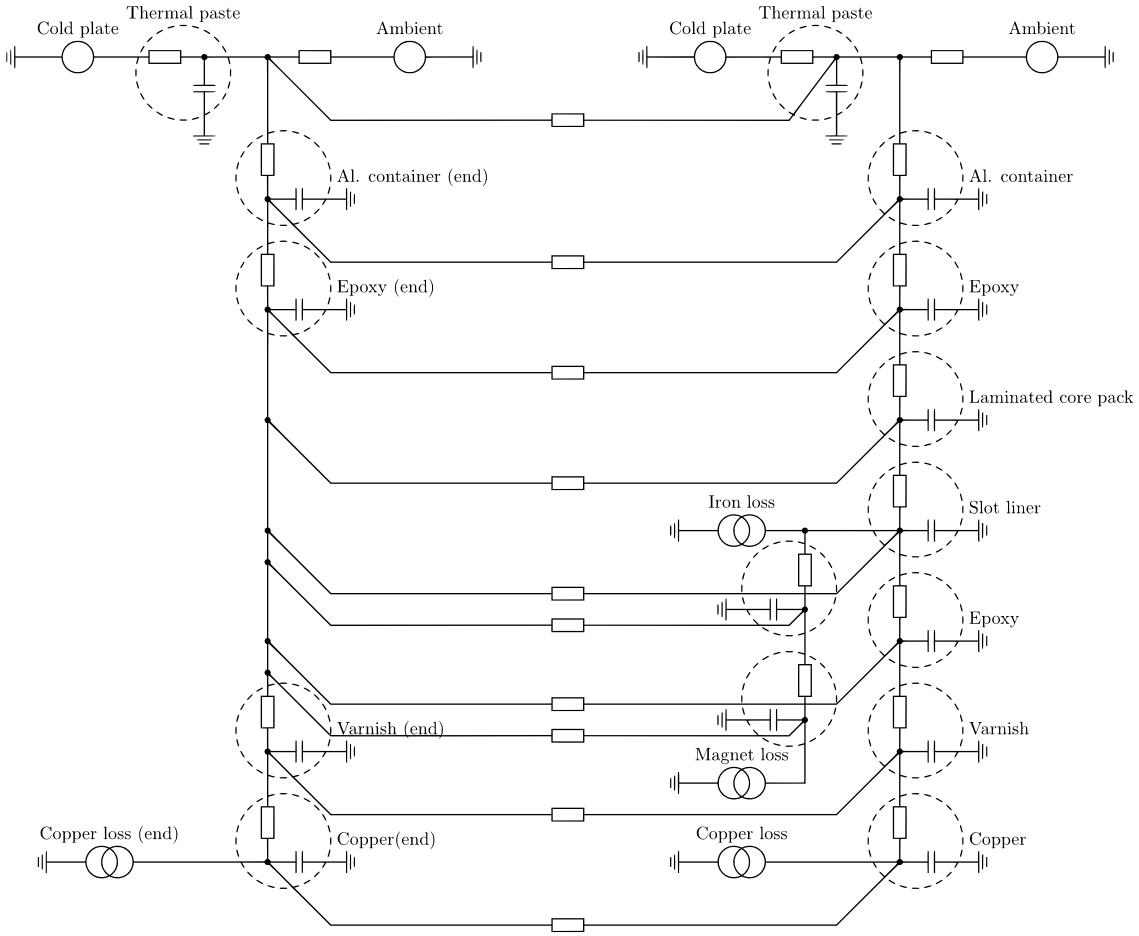

Fig. 2.1 Typical large lumped parameter model. Adapted from *Performance analysis and thermal modelling of a high-energy-density pre-biased inductor* by Wrobel *et al* [2.1].

The most common type of thermal model is the thermal circuit analogue, which is a type of parametric model. These equivalent circuits are designed to match the physical structure of a system. A model is formed by considering the system as a set of thermally homogenous regions or *lumps* and modelling each as a capacitor (analogous to its heat capacity). The lumps are then interconnected with their neighbours using resistors (analogous to thermal resistance). The result is referred to as a lumped parameter model. Heat dissipation, q, is modelled as a current source that results in heat flux (akin to electrical current) flowing through the equivalent circuit and giving rise to a temperature difference, θ, (akin to voltage) across components. A typical large thermal circuit analogue, presented by Wrobel *et al* [2.1] to model an inductor for power electronic

systems, is shown in Fig. 2.1. Other authors model the thermal attributes using a variety of different systems [2.2][2.3][2.4].

2.2.2.2 Compact thermal circuit models

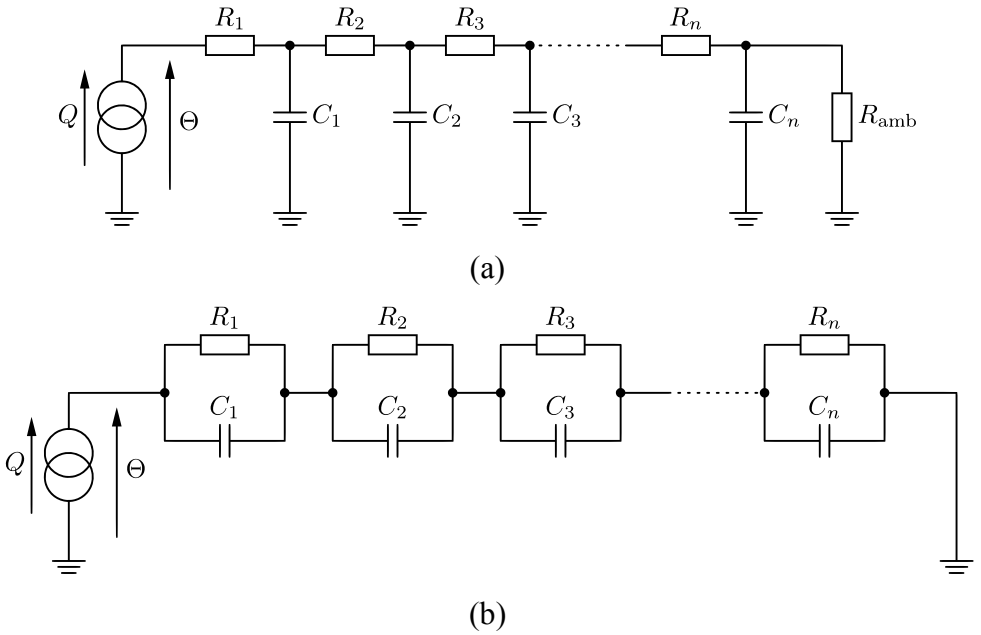

Fig. 2.2 Cauer (a) and Foster (b) thermal equivalent circuit analogues

Alternatively, a standard-form thermal equivalent circuit can be used, with its parameters fitted from the measured response of a real system. Two networks are typically used for this purpose, namely the Cauer and Foster networks, which are shown in Fig. 2.2(a) and (b) respectively.

Cauer networks are one dimensional (1-D) structural models of a material. They are usually considered as one-port networks, although they may be treated in two-port form by moving the temperature measurement point. Mussallam *et al* [2.5] report that for 1-D thermal flows, a resistor-capacitor (RC) T-network represents the thermal characteristics of layers of material. When several of these networks are placed in series to represent several layers, the Cauer network is formed. Each node on the network has physical significance and represents the temperature at a physical location.

Foster networks, however, have no physical meaning and each element merely represents a time constant present in the network. They are a type of non-parametric

model because there is no design phase in the development of the model. Instead, they are fitted to step response data as a way of producing a simple model with the correct response [2.6]. Because the constituent components in a Foster network are not analogous to physical properties, only the left-most node has physical significance. It therefore follows that Foster networks can only be used in one-port form. Several Foster networks can be interconnected, however, to model different parts of a thermal circuit.

Gerstenmaier and Wachutka [2.7] observe that any linear network can be used to describe thermal flows as long as it can fit the transient response of a one-port thermal circuit; hence Cauer and Foster networks, which are both able to fit the response, are equivalent. Lin *et al* [2.8] formalises this argument by reporting the conversion between second order Foster and Cauer networks. This property can be exploited to allow a structural Cauer network to be found from measured Foster networks, using techniques described in [2.9].

2.2.2.3 Processing of thermal models

Equivalent circuit models are ubiquitous in thermal modelling because of their simplicity and the ability to analyse such models using standard electrical analysis tools. One such tool, which is well known to electrical engineers and frequently used in many contexts, is the *simulation programme with integrated circuit emphasis* (SPICE). SPICE is regularly used to analyse thermal circuit analogues to estimate temperature responses to simulated inputs for arbitrary circuits [2.10][2.11][2.12]. The programme has advantages over simple circuit analysis because it can model components whose thermal resistance changes according to other modelled variables, such as temperature. In [2.12], for example, Nuño *et al* propose a technique to reduce complex fluid dynamics models developed for a heater into to equivalent circuits for simulation in SPICE. Their models specifically take into account changes in the emissivity of the exposed surface and changes in the specific heat capacity and density of the surrounding air with the demand on the system. The authors show that these effects can accurately be modelled using voltage-controlled resistors in SPICE.

As an alternative to SPICE, pure mathematical approaches can be developed using traditional circuit analysis techniques such as nodal circuit analysis, together with

matrix algebra, to produce simple equations for rapid analysis. Hsu and Vu-Quoc [2.13], for example, describe the underlying mathematics for equivalent thermal circuit-based mathematics. In this paper, the authors use the equations to form equivalent circuits based on finite element analysis (see below). The thermal equivalent circuit can then be coupled with the electrical circuit for comprehensive electro-thermal simulation.

2.2.3 Finite element analysis

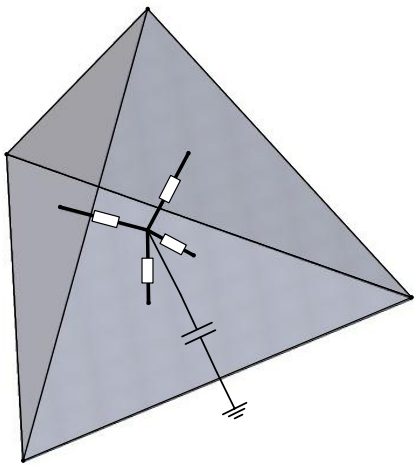

Fig. 2.3 Tetrahedral finite element model

Thermal analysis of three dimensional (3-D) systems can be performed by dividing the system into a finite set of elements and assign a compact thermal model to each 3-D element, a process that is known as meshing. A typical model consists of a general network of resistors and capacitors [2.14]. The basic variety for thermal analysis consists of a thermal point mass at the centre of a polyhedron with a thermal resistance to each surface. The overall model is formed by connecting resistors at adjacent faces. Boundary conditions, such as known temperatures or heat fluxes, can be defined for elements or surfaces, to create a complete model for analysis. A variety of polyhedrons can be used for analysis but tetrahedrons are commonly used because they can be arranged to fit into a variety of physical shapes. A simple tetrahedron element model is shown in Fig. 2.3.

Finite element models tend to be large, however, with many thousands of elements. Numerical methods are therefore used to solve problems or perform simulations. While SPICE can be used, specialist software such as ANSYS or COMSOL Multiphysics is usually employed. Models are produced offline using information on the physical

arrangement of materials and heat-sources in the system. Processing is computer-intensive due to the large model size. This technique is therefore suitable for offline evaluation of thermal properties but not generally suitable for online temperature estimation.

2.2.4 First principles model derivation

Thermal analysis can be performed directly from first-principles mathematics. This approach usually relies on solving the heat equation, given in (2.1).

$$\frac{\partial \theta}{\partial t} - \alpha \nabla^2 \theta = \frac{\alpha}{\lambda} q \tag{2.1}$$

Where θ and q are the time-domain temperature and power fields, respectively; α and λ are the thermal diffusivity and thermal conductivity of the material, respectively, ∇ is the Laplace operator and t is time.

Bryant *et al* [2.15], for example, apply the heat equation to a physical model of a power electronic module to predict temperature response in multiple devices due to cross-coupling between them. To reduce the mathematical complexity, they perform a cosine expansion on the equation so that a discrete Fourier series-based method can be used for analysis. They conclude that this technique is suitable for modelling thermal effects in power converters and verify its effectiveness on an IGBT-based power module.

The mathematical complexity of Fourier- and differential equation-based techniques limits their usefulness for online temperature estimation and control. However, these techniques represent a useful design tool for evaluating the thermal attributes of a system under typical load conditions.

2.3 Temperature measurement

In order to characterise the thermal properties of an electronic system, it is necessary to measure the temperature response to an applied power input. The location where temperature is measured is of particular importance. External measurement techniques, such as thermocouples or infra-red radiation measurement, typically characterise the

temperature at some distance from the point of heating. The temperature measured at this point has a different characteristic to that measured at the junction or heating region of a device, since there is a thermal resistance and thermal mass between the point of heating and the point of measurement in the case of external measurement.

Alternatively, the temperature at the junction or heating region can be determined by measuring temperature-sensitive parameters of the semiconductor die which dissipates the power. This approach is advantageous since the temperature of the semiconductor junction is interesting to a design engineer, owing to the temperature limitations of the materials from which it is constructed. However, measuring the junction temperature is difficult because it typically requires the measurement of characteristics, such as band gap or turn-on voltage, which the device was not designed to facilitate.

A distinction may be drawn between the die and junction temperature of a device. Whereas die temperature refers to the temperature anywhere on the semiconductor die, junction temperature refers specifically to the temperature in the active region of the device, such as a MOSFET channel or BJT junction. Although junction temperature is the most significant temperature, there are limited techniques for its measurement because any technique must measure some property of the semiconductor arrangement at the junction. There are thus more die temperature measurement techniques since properties of nearby semiconductor arrangements, such as the temperature of the intrinsic diode in a MOSFET, can be used in addition to junction temperature measurement techniques. The junction temperature can be estimated from the temperature elsewhere on the die by adapting the techniques described in 2.7.

In this section, several techniques for the measurement of temperature are reported. They are discussed below.

2.3.1 External temperature measurement

Measurement of infra-red radiation is a widely used optical technique for determining the temperature at the surface of device packaging [2.15][2.16][2.17]. Its usefulness is limited by the variation in emissivity between materials and by detector cost. It can,

however, be used on packaged chips in a test environment to measure the external temperature.

Alternatively, direct measurement of the outside of a device can be achieved with a temperature sensor. Each of these sensors makes use of some physical effect which varies with temperature (a number of which are described below). In each case, the sensor is specially designed to make use of the physical effect while achieving accuracy appropriate to the sensor cost. Thermocouples are one widely used type of external sensor due to their small size, low cost and fast response [2.18]. They rely on the Seebeck effect, where a voltage is developed across two junctions between two different metals at different temperatures. However, they have a very low output voltage which must be conditioned prior to sampling [2.19]. Thermocouples measurements are considered in detail in Chapter III.

Temperature measurements based on attached sensors suffer from a number of difficulties. The sensor is normally placed away from the semiconductor die (the temperature of interest). In addition, the sensor has its own thermal capacity and resistance which affect the results by introducing an additional time constant. They also increase the overall thermal mass and affect cooling at the site of the sensor. These limitations are, however, offset by the increased accuracy and measurement bandwidth that can normally be achieved by a connected sensor because it is possible to make a direct electrical connection from a sensor to a data logger.

2.3.2 Junction temperature estimation from p-n junction I-V curve

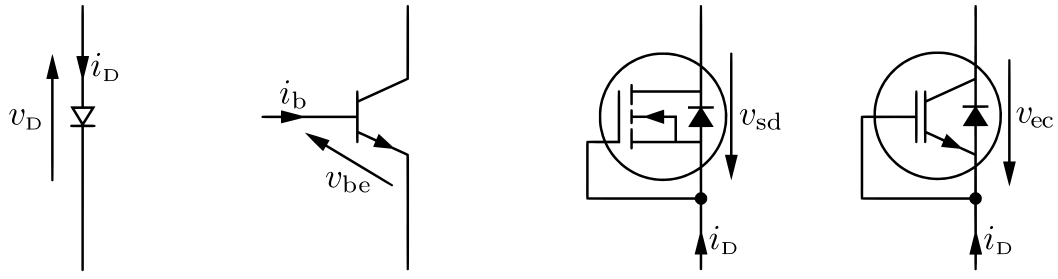

Fig. 2.4 Relevant currents and voltages for measuring temperature from the forward voltage drop in (a) a diode; (b) a bipolar transistor; (c) the intrinsic diode in a MOSFET; and (d) the intrinsic diode in an IGBT

The diode equation, stated in (2.2), is the relationship between the forward voltage drop and the conduction current of a p-n junction [2.20].

$$V_{\mathrm{D}} = \frac{kT}{\mathrm{e}} \ln\left(\frac{I}{I_{\mathrm{s}}} - 1\right) \tag{2.2}$$

Where V_{D} is the forward voltage drop, I is the conduction current, T is junction temperature, I_{s} is the reverse leaking current, k is the Boltzmann constant and e is elementary charge. Iskenovic and Mitic [2.20] use this relationship to directly measure the junction temperature of a diode. By applying very short bursts of constant current to minimise self-heating effects from the measurement system, the forward voltage drop can be measured. Because short pulses cause difficultly in accurate sampling of the voltage, the output voltage is low-pass filtered and a bespoke technique applied for measurement. A quartz crystal oscillator is used for timing and the number of pulses observed between two points along the voltage decay is measured. This gives an estimate of V_{D} which does not require sampling with an analogue-to-digital converter.

The diode equation can be used to measure MOSFET temperatures by using the intrinsic antiparallel diode present in those devices [2.21]. During the non-conducting period, a small bias current is driven through the intrinsic diode and the resultant voltage is measured. Similar techniques can be applied to other devices which contain a p-n junction, such as seen in the work of Yong-Seok and Seung-Ki [2.22]. As well as an antiparallel intrinsic diode, IGBTs feature an on-state voltage drop, $v_{\mathrm{ce(sat)}}$ which varies with both collector current and junction temperature. The relationship is almost linear with a typical 5 K maximum error generated by this assumption. A device is first characterised off-line to populate physical parameters. These parameters are then used online to predict the temperature from the collector current and voltage drop.

In all these cases, care must be taken because the voltage drop depends not only on temperature but also on current. If the current is not constant, the effect of both temperature and current must be characterised to calibrate the sensor, increasing the complexity of the technique.

2.3.3 Junction temperature estimation from resistance

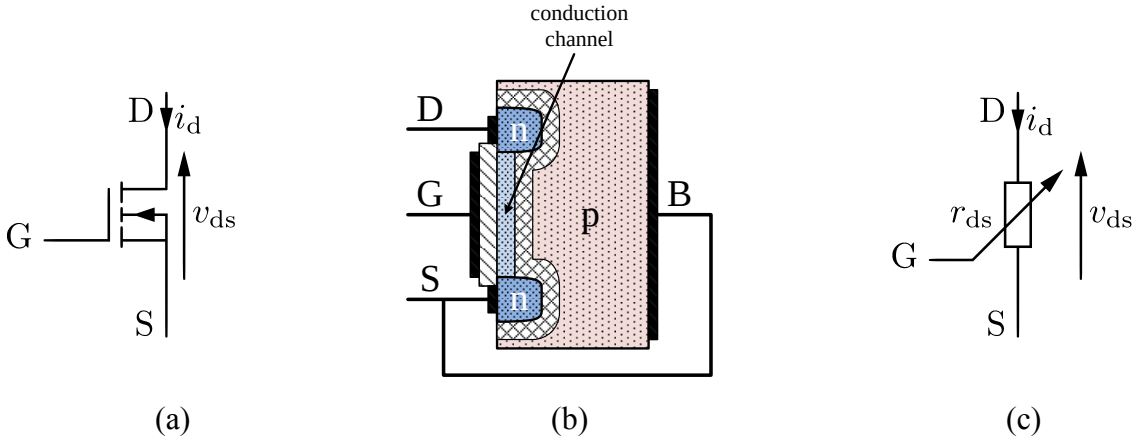

Fig. 2.5 Conduction channel resistance in an on-state MOSFET. (a) Schematic of MOSFET with relevant voltages and currents indicated. (b) Structure of MOSFET with conduction channel shown. (c) Equivalent circuit of conduction channel.

Majority-carrier semiconductor devices, including MOSFETs, exhibit temperature-dependence on their resistance similar to most other conductors. Pyo and Sheng [2.23] uses this relationship to estimate the junction temperature of a MOSFET and diode. The relationships between forward voltage and current are characterised using short pulses to limit self-heating. In normal operation, current and voltage can be measured and the junction temperature can be calculated from the relationship obtained for that device.

Koenig *et al* [2.24], however, directly characterise the temperature dependency of on-state resistance $r_{\mathrm{ds(on)}}$. This, they suggest, can be entered into a look-up table for the device, allowing junction temperature to be calculated online by measuring current and voltage and hence calculating $r_{\mathrm{ds(on)}}$.

2.3.4 Junction temperature estimation from MOSFET transients

Since MOSFETs are often used in switching, the dependency between temperature and measureable transients may be used to infer device temperature. Barlini *et al* [2.25] observe that because electron mobility and therefore resistance is related to temperature, and because MOSFET turn-on time is governed by oxide capacitance, the turn-on waveform can used to infer junction temperature. It is shown that the rate of change in drain current during MOSFET turn-on [2.25] is

$$\frac{\mathrm{d}I_d}{\mathrm{d}t} = \mu(T) C_{\mathrm{ox}} \frac{W}{L} \left(V_{\mathrm{gs}} - V_{\mathrm{th}}(T) \right) \cdot \frac{\mathrm{d}V_{\mathrm{gs}}}{\mathrm{d}t} \tag{2.3}$$

Where I_d is the drain current, C_{ox} is the gate oxide capacitance, W and L are the width and length of the channel respectively, V_{gs} is the gate voltage and t is time. The electron mobility, μ, and threshold voltage, V_{th}, are functions of absolute temperature, T. From measuring $\mathrm{d}I_d/\mathrm{d}t$, the temperature can therefore be inferred.

A similar technique is applied by Chen *et al* [2.26] and Wang *et al* [2.27]. In these cases, the temperature dependency of the threshold voltage is used to infer temperature. It can be shown [2.28] that

$$\frac{\mathrm{d}V_{\mathrm{th}}}{\mathrm{d}T} = \frac{\mathrm{d}\Phi_{\mathrm{F}}}{\mathrm{d}T} \cdot \left(1 + \frac{e}{C_{\mathrm{ox}}} \sqrt{\frac{\epsilon_{\mathrm{Si}} N_{\mathrm{a}}}{kT \ln(N_{\mathrm{a}}/n_{\mathrm{i}})}} \right) \tag{2.4}$$

Where V_{th} is threshold voltage, Φ_{F} is the Fermi level, e is elementary charge, N_{a} is the density of acceptor atoms, k is the Boltzmann constant, ϵ_{Si} is the dielectric constant of silicon and n_{i} is the intrinsic carrier density. Φ_{F} is itself temperature-dependent [2.29], as shown in (2.5).

$$\frac{\mathrm{d}\Phi_{\mathrm{F}}}{\mathrm{d}T} \approx \frac{1}{T} \left(\frac{E_{\mathrm{g}}(T=0)}{2e} - \left| \Phi_{\mathrm{F}} - \Phi_{\mathrm{F(i)}} \right| \right) \tag{2.5}$$

Where E_{g} is the band gap and $\Phi_{\mathrm{F(i)}}$ is the Fermi level of an intrinsic semiconductor (that is, the midpoint of the bandgap). Chen *et al* use this dependency to measure junction temperature by adding a large resistor to the gate to delay turn-on time and increase the sensitivity to temperature. Results verify that the turn-on delay can be used to estimate junction temperature. In this case, the results are accurate to 3 K.

2.3.5 Embedded temperature sensors

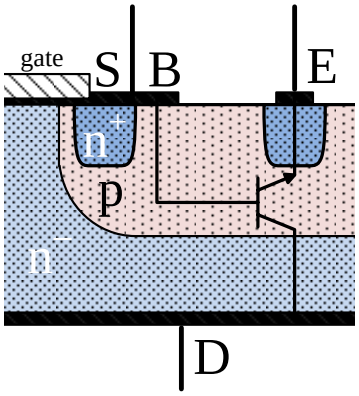

Fig. 2.6 Modified DMOS cell in a power MOSFET exposing microscopic BJT for temperature measurement

Several methods of integrating temperature sensors into devices are reported in the literature. Pfost *et al* [2.30], for example, propose the integration of a number of temperature sensors into a power MOSFET. Since many power MOSFETs are implemented as an array of many vertically diffused MOS (DMOS) cells, Pfost *et al* argue that re-purposing a small proportion of cells for temperature measurement will have little electrical impact while providing several temperature readings in proximity to the active region. A microscopic NPN transistor is embedded into the MOSFET structure, as show in Fig. 2.6. By applying a bias current between base and emitter, the forward voltage drop of the p-n junction can be measured and used to determined temperature as described in section 2.3.2. The proposed sensors were designed for characterisation of test devices; however, the authors note their potential suitability for normal use applications as well.

An alternative approach is taken by Dibra *et al* [2.31] who present a design for an on-chip temperature difference sensor based on the Seebeck effect. This sensor measures the difference between two regions of the semiconductor die and is therefore suitable for power devices with co-integrated logic—so-called smart power technology. These temperature sensors target pulsed power applications where there is likely to be a significant temperature difference between the integrated power MOSFET and the logic circuitry under transient operation. Through close thermal coupling with the active area of the device, it is possible to rapidly measure the thermal transient without the measurement noise associated with p-n junction-based temperature sensors.

Many commercial solutions for embedded temperature sensors exist, particularly for sensors embedded into power electronic modules. Fuji Electric [2.32], for example, have developed a motor-drive IGBT module with a p-n junction-based temperature sensor thermally coupled to each of the six IGBTs. These sensors allow device temperature changes during various use scenarios, such as regenerative braking, to be monitored. In one example, a 5°C temperature increase observed on the case of the module masks the true 65°C internal temperature increase which can be measured using the embedded sensors.

2.3.6 Other die temperature measurement techniques

Direct band gap semiconductors emit radiation when excited electrically (electroluminescence) or by an external light source (photoluminescence) [2.33]. The frequency of the radiation is dependent on the band-gap of the semiconductor, which is temperature sensitive [2.34]. This effect has therefore been used to measure exposed die temperature [2.35].

Photons incident on a semiconductor die are scattered inelastically (that is, their energy changes) due to interactions with phonons in the semiconductor [2.36]. Since phonon energy is temperature-dependent, the spectrum of the scattered photons is also temperature dependent, and can be used to measure temperature [2.37]. This effect is known as Raman scattering.

In his review, Blackburn [2.29] summarises a number of additional external measurement techniques. Several of these techniques, and the techniques above, require the semiconductor die to be exposed and are therefore only useful within the laboratory on test devices.

2.4 Measurement of the thermal impedance in practical circumstances

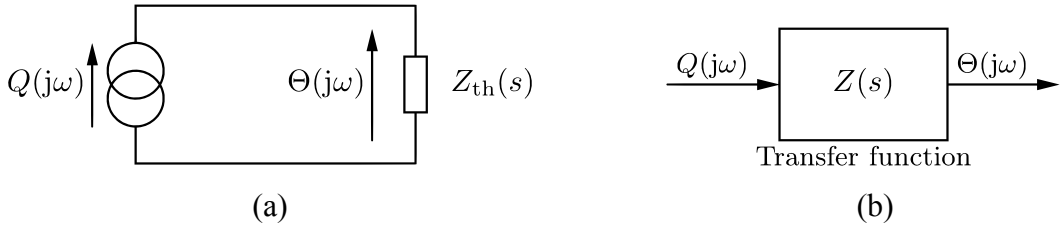

Fig. 2.7 Thermal impedance between heat source and temperature measurement. (a) Circuit diagram. (b) Equivalent transfer function.

In order to parameterise models or to develop non-parametric models, some characteristics of the real system are required. Thermal impedance, which is the frequency-domain relationship between the heat dissipated and the temperature response, is one useful characteristic. For thermal systems, the impedance is the transfer function between heat dissipation and temperature, as shown in Fig. 2.7.

Musallam *et al* [2.38] have reported using a Cauer thermal analogue to estimate and control the temperature in a power converter to improve reliability. The authors use known thermal characteristics to control temperature but, in situations where the thermal characteristics are not known, techniques to derive the characteristics can be used to improve thermal management. Accurate and efficient methods to derive these characteristics in the frequency domain are therefore required. Where the power dissipation in a circuit is changeable, the thermal resistance, which is merely a measure of the expected temperature rise due to a static dissipation, is insufficient for accurate temperature estimations. Identifying the frequency-dependent complex thermal impedance is therefore required.

2.4.1 Measuring thermal impedance using step functions

A common approach to measuring the thermal impedance is to apply a step function of power to the thermal system, measure the transient response and convert it into a Foster Network [2.39]. This technique is limited in several ways. First, the Foster networks characterised can only model self-impedance. Where transfer impedance is relevant, such as in the case of cross-coupling [2.40], a more general network such as a two-port

Cauer network is required. Second, the conversion from a thermal transient to a frequency domain representation (and therefore to a more general model) is difficult because the conversion process is very sensitive to noise, particularly at higher frequencies, as shown by Salleras *et al* [2.41]. In addition, to extract the response for a long time constant, the response must be obtained for the whole period as no technique is possible which can extract a low frequency response without a long characterisation measurement period.

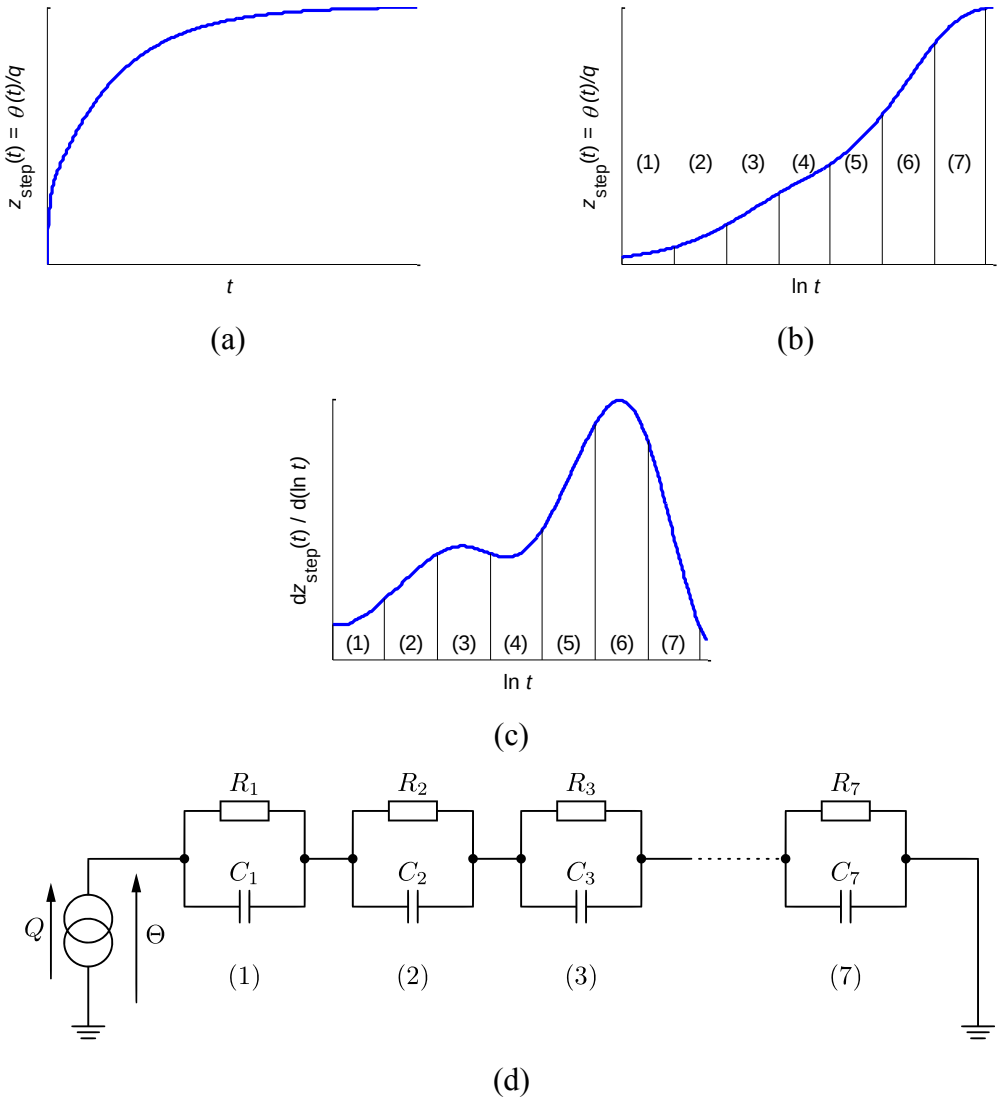

Fig. 2.8 Parameterisation of a Foster network from step response. (a) Step response. (b) Step response against logarithmic time. (c). Derivate of step response against logarithmic time. (d) Foster network with corresponding step response regions indicated.

Székely and Bien [2.6] report a method to identify the Foster thermal analogue of a semiconductor device by analysing the transient temperature response to a step of power. Their method is shown in Fig. 2.8. By mapping the time-domain derivative of the step response into the logarithmic time domain and dividing the response into sections, the time-constants of each Foster element can be calculated. A convolution-based process is developed to perform the conversion. This process results in a Foster network whose step response is an approximation of the initially measured response. A good explanation with reference to thermal measurements of electronic packaging is provided by Yang *et al* [2.42].

2.5 Presentation of thermal impedance

The format in which results are presented is important. While the same data is presented for each format, it is important to note that some methods of presentation lend themselves more easily to the visual extraction of information.

2.5.1 Impedance Bode plots

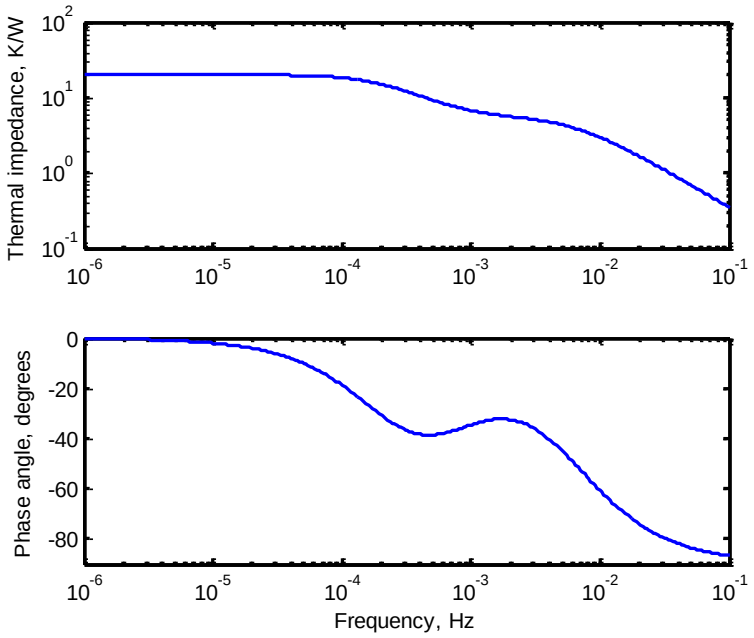

Fig. 2.9 Bode plot of a typical thermal system

Bode plots are the typical method of presenting frequency domain data. Amplitude and phase are separated and plotted against logarithmic frequency. In the case of power gain, for which Bode plots were developed, it is common to plot the amplitude data in

decibels, a logarithmic unit. For impedances, however, no such unit exists and therefore the amplitude is plotted on a logarithmic scale. Bode plots allow changes in system response with frequency to be viewed easily and, since it is usually amplitude that is important, irrelevant data can easily be ignored. A typical Bode plot of the self-impedance of a system is shown in Fig. 2.9.

Bode plots are widely used to present frequency-dependent data. Christiaens and Beyne [2.43], for example, extract frequency domain thermal impedance data from a microchip attached to a large ceramic substrate. These allow comparison of the peak thermal impedance and time constants between different microchip materials. David *et al* [2.44] similarly present thermal data in Bode plot format, in this case for microwave semiconductor devices. By analysing the shape of the Bode plot, they are able to conclude that the thermal circuit in question is equivalent to a low-pass filter.

2.5.2 Nyquist impedance plots

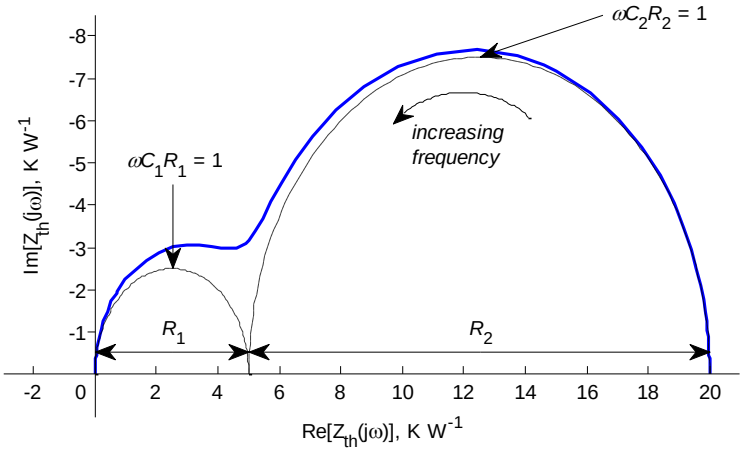

Fig. 2.10 Nyquist plot of thermal self-impedance. The thermal system appears like a two-element Foster network. Parameters R_1, R_2, C_1 and C_2 can be found by inspection as indicated.

An alternative method of presentation is offered by the Nyquist diagram. In this case, the real and imaginary parts of the thermal impedance are plotted against each other on a single graph. Frequency data is not plotted but, in typical systems, it increases nonlinearly from right to left. An overview of this technique is provided by Kawka *et al*

[2.45]. When temperature measurement is at the heat source, the Nyquist plot resembles a series of semicircles displaced along the real axis. The radii and centre frequencies of these semicircles can be used to infer properties of the layers through which the heat flux must travel before being dissipated into ambient. This type of plot is useful in fitting a Foster network to the response, as demonstrated in Fig. 2.10. The technique is particularly well suited to characterisation of packages because these are typically formed from layers of different semiconductor materials. Poppe *et al* [2.46], for example, use Nyquist diagrams to present thermal data. In this case, the differing thermal impedances measured along the same path in opposing directions are reported. The Nyquist diagram is used to clearly show the differences.

2.5.3 Structure functions

Poppe *et al* [2.46] further present their results as structure functions. In this representation, the cumulative thermal resistance is plotted against the cumulative thermal capacitance for the infinitesimally divided Cauer network which represents the thermal circuit [2.47]. It is called the structure function because it represents the heat-flux through the structure [2.6]. When plotted in this format, straight horizontal and vertical sections of the graph show areas of lumped resistance and capacitance respectively. In the case of resistance, this will usually refer to an interface between two regions where a substantial boundary resistance exists. The structure function is typically calculated from the transient thermal response where heat dissipation and temperature measurement are applied at the same location. Salleras *et al* [2.47] highlight the difficulties posed by the calculation method, which is based on convolution, and report that high precision calculations are required, consisting of 50-100 significant digits for a typical 10-stage network.

2.6 Parameterisation of models from thermal impedance

To model the thermal characteristics of devices and circuits, they must be excited by an input heat dissipation and have their temperature response measured. By analogy to electrical systems, the thermal impedance is calculated from the quotient of temperature and power. In some cases, only the thermal resistance (the real part of thermal impedance) is required. This is used to predict the temperature rise caused by a static dissipation once all transient effects have passed. Kim and Shin [2.48], for instance,

calculate the thermal resistance of various LED packaging arrangements by applying a static heat dissipation and measuring the temperature change after the transient response has passed. Similarly, Reih *et al* [2.49] use variations in the gain of a heterojunction transistor with temperature to measure the thermal resistance of several device geometries under different power dissipation levels. The authors extend this work to several adjacent devices integrated onto the same die, and develop a model to estimate the thermal interference between devices. The authors use the information they obtain to identify a trade-off between electrical and thermal performance when selecting a device structure and circuit layout.

In both [2.48] and [2.49], the devices in question are small and, particularly in the case of LEDs, unlikely to experience fast-varying power dissipation under normal operation. Because of this, it is reasonable to consider only thermal resistance and neglect time-varying parameters. Where this is not the case, however, thermal impedance must instead be calculated, which allows for accurate modelling under transient and time-varying dissipation conditions.

Thermal impedance can be calculated by measuring the transient response in temperature to a step change in heat dissipation. This approach is taken by Shammas *et al* [2.50] who use the transient data to calculate parameters in a thermal circuit analogues of semiconductor packages. Lin *et al* [2.8] apply this technique to LED packages and, by analysis of the resulting transient, are able to draw conclusions on the thermal structure of the LED's packaging.

2.6.1 Characterisation of thermal cross-coupling

The calculation of cross-coupling characteristics is a well-studied area of research. Cross-coupling is essentially the transfer impedance of a system between a heat dissipation point and a separate temperature measurement point. Typically, the heat dissipation point is an area of high dissipation under typical use. The temperature measurement point is a point of interest: in practical systems it may be a point with a low temperature limit. Many cross-coupling characteristics can be obtained for a system, since there may be many pairs of points which are of interest. Research papers frequently characterise the thermal cross-coupling by fitting an electrical equivalent

circuit of resistors and capacitors to the system [2.40][2.51][2.52]. Other methods are also used. Evans *et al* [2.39], for example, describe the development of a model based on state-space differential equations of the heat equation which is simplified using a computer algorithm.

Cross-coupling and other models can be used to estimate the temperature under a given power demand. Brückner and Bernet [2.51], for example, report a temperature estimation technique for a voltage source converter module. By constructing an electrical analogue for the thermal system and calculating power input from electrical characteristics, the junction temperatures of the semiconductor devices are estimated. In this case, however, temperature is not calculated with respect to time but rather the temperature of each element for each operating condition is evaluated. In a system with variable operating conditions, a method is often required which predicts the temperature at a given time for an arbitrary dissipation profile.

2.7 Real-time estimation and prediction of device temperatures

Once identified, thermal models can be used to estimate and predict the temperature at each system element due to the power dissipations in all elements. These prediction methods are important in power electronics because they allow a design to be evaluated in terms of its compliance with the rated maximum temperature of elements and its likely longevity under a range of operating conditions. Techniques such as active cooling may be used to maintain system parameters within tolerances. This concept has been applied extensively for on-chip temperature management in the VLSI sector. X. Wang *et al* [2.53], for example, use model-predictive control to calculate the temperature and attempt to hold it at a set-point by varying power consumption. This approach has applications in power electronics where the maximum component temperatures are constraints to the safe operating envelope. H. Wang *et al* [2.52] report a method which uses a grid of resistors, capacitors and current sources to model a planar microchip. By constructing a finite difference equation, the temperature response at several points can be estimated in real-time for arbitrary power input. This estimate is improved using error compensation from a limited number of on-board temperature sensors. Similarly, Musallam and Johnson [2.54] use a combination of the practical and

simulated temperature response to a step power input to generate parameters for a Laplace-domain model. By transforming the model into the time domain and converting to a difference equation, the temperature due to an arbitrary power input may be computed in real time. In [2.54], the model is parameterised and the difference equations computed for each pair of system elements, hence the effect of cross-coupling is explicitly taken into account.

In recent years, similar approaches have been applied to power electronics. For example, James *et al* [2.40] have reported using a step function to evaluate the auto- and cross-coupling between elements of multichip power electronic module and the authors fit these characteristics to a Foster network. The resulting model is used to estimate the temperature response for an arbitrary power input. In addition, Holmes *et al* [2.55] present a method using a simple equivalent circuit and a closed-loop Luenberger observer system to estimate device temperatures in real time. The effectiveness of this model is limited, however, because the equivalent circuit must be formulated and its component values populated before it can be used to make estimations.

2.7.1 Estimation of junction or die temperature

Real-time estimation of the junction or die temperature is an important field because it is this temperature which is relevant for consideration of device ratings and reliability. Several techniques have been reported in recent literature. Musallam *et al* [2.56] present a temperature estimator for a single power MOSFET mounted on a heatsink based on a one-dimensional (1-D) thermal model. By analysing the heat transfer path from the active region through the physical structure of the device, the authors populate a 7th order Cauer network using published device characteristics. The paper reports that the network can be reduced to 3rd order without significant impact on results. Verification is achieved using the forward diode voltage drop method (discussed in the 2.3.2) to infer die temperature and shows that a reasonable estimation can be achieved.

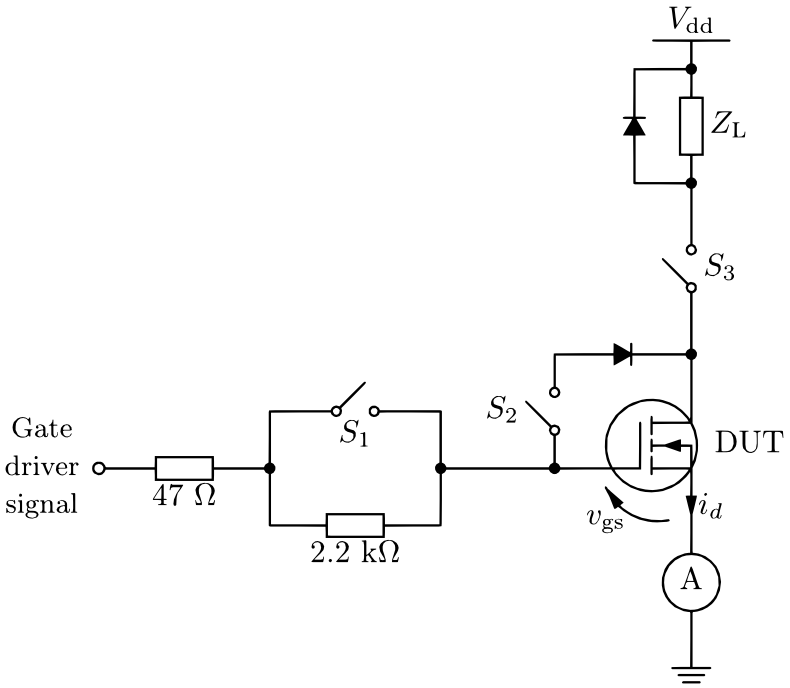

Fig. 2.11 Online threshold voltage measurement circuitry proposed by Chen *et al*

An alternative method of temperature estimation and modelling is reported by Chen *et al* [2.57]. Their method is based on the computation of junction temperature from the MOSFET's threshold voltage using the method described in section 2.3.4. The proposed measurement system is implemented on a MOSFET switching an inductive load, Z_L, and its topology is reproduced in simplified form in Fig. 2.11. Under normal operation S_1 and S_2 are open and S_3 is closed and the measurement system has no impact. For measurements, S_3 is momentarily opened to disconnect the device-under-test (DUT) from the load, while S_1 and S_2 are momentarily closed to increase turn-off time and thereby improve measurement accuracy. The channel current, i_d is monitored following turn-off and the gate voltage v_gs is recorded at the zero-crossing (this is the threshold voltage, V_th). Using this method, the cooling curve calculated from threshold voltage measurements is used to parameterise a 3rd order Foster network, which may then be used as the basis of a real-time temperature estimator. Through repeated online measurement of the cooling curve over the device's lifespan, the technique is able to update the model to account for device ageing. Reported results show good agreement between estimated results and a real cooling curve for a single device thermal system, although further work is needed to address loss of temperature measurement calibration and to overcome difficulties from the momentary load disconnection in operation.

A distinction may be drawn between the transient and steady state estimation of junction temperatures. On the one hand, steady-state estimation of junction temperatures determines the temperature change due to changes in average power dissipation. This data is useful, for example, when analysing the effect of changing load of a power converter. The estimated temperatures are proportional to the junction temperature averaged over a period of 50 milliseconds or more. Transient temperature estimation, on the other hand, is sensitive to temperature changes caused by individual switching events where time constants are of the order of 100 microseconds. This high frequency transient analysis is useful when considering the effect of switching on junction temperature within the semiconductor and high frequency repetitive thermal stress.

2.7.2 Implementations using difference equations

Estimations of temperature response using thermal models can be a complex matter involving computationally intensive numerical techniques [2.39][2.58][2.59]. In practice, if the problem can be reduced to a small set of equations to calculate, a potential implementation becomes practical. Converting a model into a collection of finite difference equations is a useful method which yields a simple estimation and prediction technique.

James *et al* [2.40] model the cross-coupling in a thermal system by fitting a Foster network to the step response between each pair of devices, using techniques discussed earlier in this Chapter. To embed a temperature predictor based on such a model into a microcontroller, the equivalent circuit must undergo transient analysis. For a given model, a set of finite difference equations are produced by applying the trapezium rule to differential equations for each capacitor [2.60]. These equations are then used iteratively to generate a temperature prediction for an arbitrary input.

2.8 Applications of pseudorandom binary sequences

Pseudorandom binary sequences (PRBS) can be used to excite a system so that the impedance response can be determined. They are often used to generate band-limited white noise. A PRBS is a series of variable-length input pulses which, due to the many switching events, has improved content at higher frequency. System identification using PRBS techniques has been in use for many decades. Davies [2.61] presents an in-depth

analysis of the technique and its application. Recently, PRBS techniques have been applied to a variety of applications including parameter estimation in electrical generators [2.62], testing of digital-to-analogue converters [2.63] and modelling of batteries [2.64].

Wellstead [2.65] describes the techniques which can be used for system identification using PRBS signals. These are based on computing direct spectral estimates of the input PRBS signal and the corresponding output. At its most basic, the reported technique calculates the transfer function (or thermal impedance) as the quotient of the discrete Fourier transforms (DFTs) of the input and output signals. Wellstead, however, describes techniques to improve data quality, such as spectral smoothing and windowing of the DFTs.

The applications of PRBS are broad. In early work on the subject, Hampton [2.66] discusses the use of PRBS as an alternative to digitised white noise generated through natural processes. In his work, time and frequency-domain characteristics of the PRBS are described while the derivation of the power spectral density is reported. Hampton presents a practical implementation of a noise generator based on PRBS which he supposes can replace traditional noise generation techniques. Furthermore, a 1992 inventors' patent [2.67] proposes the production of very long pseudorandom noise by taking the exclusive-or of two PRBS outputs clocked at different frequencies for improved white noise characteristics.

Since then, numerous applications of PRBS have been found. Fairweather *et al* [2.64] use PRBS to identify the state of health and charge of lead-acid batteries. By applying a PRBS discharge current to the terminals of a battery and sampling the resulting voltage perturbations, the frequency-domain impedance of the battery can be characterised. This is used to parameterise Randles' model of a battery.

Miao *et al* [2.68] have used PRBS to identify the transfer function of switching power converters. The method used is similar to Fairweather *et al*; however, Miao *et al* additionally use cross-correlation to reduce the effect of noise. In this case, by repeating the PRBS for several periods and averaging the period together by cross-correlation, the

amount of noise present is reduced and hence a cleaner transfer function can be obtained. The result is that a more accurate model of the converter can be obtained.

PRBS has also been used for thermal modelling. Gabano *et al* [2.69] report the analysis of heat transfer through a wall by applying heat dissipation in the form of a PRBS to one side and measuring the temperature at the opposing side. Similarly, Defer *et al* [2.70] reports the thermo-physical characterisation of an arbitrary material by applying PRBS heat flux to its face and measuring the resulting temperature change. In both cases, the thermal impedance of the system is derived.

2.9 Chapter conclusions

This Chapter has reported on many of the available techniques for thermal modelling and temperature measurement. Modelling techniques can be parametric, where a specific model design is produced and parameterised for each system, or non-parametric, where a general model is derived for any system. The parameters of models can be determined using offline analysis which uses material properties and knowledge of the structure, but not actual readings. Although thermal analysis using this data is useful for rapid prototyping, a final design requires practical measurement of thermal impedances in order to increase accuracy and verify the design. Methods of measuring thermal impedance and forming models from the data are therefore reported, together with implementations of temperature estimators and predictors. Techniques to measure relevant temperatures, which are required for thermal impedance measurements, are reported. Applications of pseudorandom binary sequences are also given, since these form the basis of the thermal impedance measurement techniques proposed in this thesis.

2.10 References

[2.1] R. Wrobel, N. McNeill and P. H. Mellor, '*Performance Analysis and Thermal Modeling of a High-Energy-Density Prebiased Inductor*', IEEE Transactions on Industrial Electronics, vol. 57, pp. 201-208, 2010.

[2.2] B.-H. Lee, K.-S. Kim, J.-W. Jung, J.-P. Hong and Y.-K. Kim, '*Temperature Estimation of IPMSM Using Thermal Equivalent Circuit*', IEEE Transactions on Magnetics, vol. 48, pp. 2949-2952, 2012.

[2.3] R. Rashmi, G. A. Armstrong, S. J. Harrington, A. Bousquet and S. Nigrin, '*Equivalent circuit model for thermal resistance of deep trench isolated bipolar transistors*', *in Proc* Bipolar/BiCMOS Circuits and Technology Meeting (BCTM), 2010, pp. 261-264.

[2.4] W. H. Tang, Q. H. Wu and Z. J. Richardson, '*Equivalent heat circuit based power transformer thermal model*', IEE Electric Power Applications, vol. 149, pp. 87-92, 2002.

[2.5] M. Musallam, P. P. Acarnley, C. M. Johnson, L. Pritchard and V. Pickert, '*Estimation and control of power electronic device temperature during operation with variable conducting current*', IET Circuits, Devices & Systems, vol. 1, pp. 111-116, 2007.

[2.6] V. Székely and T. Van Bien, '*Fine structure of heat flow path in semiconductor devices: A measurement and identification method*', Solid-State Electronics, vol. 31, pp. 1363-1368, 1988.

[2.7] Y. C. Gerstenmaier and G. Wachutka, '*Calculation of the temperature development in electronic systems by convolution integrals*', *in Proc* Semiconductor Thermal Measurement and Management Symposium, 2000, pp. 50-59.

[2.8] Y. Lin, Y. Lu, Y. Gao, Y. Chen and Z. Chen, '*Measuring the thermal resistance of LED packages in practical circumstances*', Thermochimica Acta, vol. 520, pp. 105-109, 2011.

[2.9] L. Codecasa, '*Canonical forms of one-port Passive Distributed thermal networks*', IEEE Transactions on Components and Packaging Technologies, vol. 28, pp. 5-13, 2005.

[2.10] A. Chvala, D. Donoval, J. Marek, P. Pribytny, M. Molnar and M. Mikolasek, '*Fast 3-D Electrothermal Device/Circuit Simulation of Power Superjunction MOSFET Based on SDevice and HSPICE Interaction*', IEEE Transactions on Electronic Devices, vol. 61, pp. 1116-1122, 2014.

[2.11] S. Lineykin and S. Ben-Yaakov, '*Modeling and Analysis of Thermoelectric Modules*', IEEE Transactions on Industry Applications, vol. 43, pp. 505-512, 2007.

[2.12] F. Nuno, J. Parrondo, A. M. Pernia, R. Barrio, P. J. Villegas, M. J. Prieto and P. Barquilla, '*Equivalent Electrical Model and Software Tool for SPICE-Compatible Thermal Simulations of High-Power Resistors*', IEEE Transactions

on Components, Packaging and Manufacturing Technology, vol. 4, pp. 859-869, 2014.

[2.13] J. T. Hsu and L. Vu-Quoc, '*A rational formulation of thermal circuit models for electrothermal simulation. Part I. Finite element method*', IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications, vol. 43, pp. 721-732, 1996.

[2.14] T. W. C. Mak and L. Feinstein, '*A simplified modeling technique for finite element thermal analysis of a plastic SOIC package*', IEEE Transactions on Components and Packaging Technologies, vol. 22, pp. 6-10, 1999.

[2.15] A. Bryant, N.-A. Parker-Allotey, D. Hamilton, I. Swan, P. A. Mawby, T. Ueta, T. Nishijima and K. Hamada, '*A Fast Loss and Temperature Simulation Method for Power Converters, Part I: Electrothermal Modeling and Validation*', IEEE Transactions on Power Electronics, vol. 27, pp. 248-257, 2012.

[2.16] D. Gautam, D. Wager, F. Musavi, M. Edington, W. Eberle and W. G. Dunford, '*A review of thermal management in power converters with thermal vias*', *in Proc* Applied Power Electronics Conference and Exposition (APEC), 2013, pp. 627-632.

[2.17] C. Hsien-Chie, C. Wen-Hwa and I. C. Chung, '*Integration of Simulation and response surface methods for thermal design of multichip modules*', IEEE Transactions on Components and Packaging Technologies, vol. 27, pp. 359-372, 2004.

[2.18] R. J. Redding, '*The electrical measurement of temperature*', Students' Quarterly Journal, vol. 17, pp. 61-65, 1947.

[2.19] J. D. Nash, D. R. Caldwell, M. J. Zelman and J. N. Moum, '*A Thermocouple Probe for High-Speed Temperature Measurement in the Ocean*', Journal of Atmospheric and Oceanic Technology, vol. 16, pp. 1474-1482, 1999.

[2.20] P. S. Iskrenovic and D. B. Mitic, '*Temperature measurement by means of semiconductor diode in pulse mode*', Review of Scientific Instruments, vol. 63, pp. 3182-3184, 1992.

[2.21] Z. Jakopovic, Z. Bencic and F. Kolonic, '*Important properties of transient thermal impedance for MOS-gated power semiconductors*', *in Proc* Industrial Electronics (ISIE), 1999, pp. 574-578 vol.2.

[2.22] K. Yong-Seok and S. Seung-Ki, '*On-line estimation of IGBT junction temperature using on-state voltage drop*', *in Proc* Industry Applications Conference (IAS), 1998, pp. 853-859 vol.2.

[2.23] S. Pyo and K. Sheng, '*Junction temperature dynamics of power MOSFET and SiC diode*', *in Proc* Power Electronics and Motion Control Conference (IPEMC), 2009, pp. 269-273.

[2.24] A. Koenig, T. Plum, P. Fidler and R. W. De Doncker, '*On-line Junction Temperature Measurement of CoolMOS Devices*', *in Proc* Power Electronics and Drive Systems (PEDS), 2007, pp. 90-95.

[2.25] D. Barlini, M. Ciappa, M. Mermet-Guyennet and W. Fichtner, '*Measurement of the transient junction temperature in MOSFET devices under operating conditions*', Microelectronics Reliability, vol. 47, pp. 1707-1712, 2007.

[2.26] H. Chen, V. Pickert, D. J. Atkinson and L. S. Pritchard, '*On-line Monitoring of the MOSFET Device Junction Temperature by Computation of the Threshold Voltage*', *in Proc* Power Electronics, Machines and Drives (PEMD), 2006, pp. 440-444.

[2.27] R. Wang, J. Dunkley, T. A. DeMassa and L. F. Jelsma, '*Threshold voltage variations with temperature in MOS transistors*', IEEE Transactions on Electronic Devices, vol. 18, pp. 386-388, 1971.

[2.28] G. Groeseneken, J. Colinge, H. E. Maes, J. C. Alderman and S. Holt, '*Temperature dependence of threshold voltage in thin-film SOI MOSFETs*', IEEE Electronic Device Letters, vol. 11, pp. 329-331, 1990.

[2.29] D. L. Blackburn, '*Temperature measurements of semiconductor devices - a review*', *in Proc* Semiconductor Thermal Measurement and Management Symposium, 2004, pp. 70-80.

[2.30] M. Pfost, D. Costachescu, A. Mayerhofer, M. Stecher, S. Bychikhin, D. Pogany and E. Gornik, '*Accurate Temperature Measurements of DMOS Power Transistors up to Thermal Runaway by Small Embedded Sensors*', IEEE Transactions on Semiconductor Manufacturing, vol. 25, pp. 294-302, 2012.

[2.31] D. Dibra, M. Stecher, A. Lindemann, J. Lutz and C. Kadow, '*Seebeck difference - temperature sensors integrated into smart power technologies*', *in Proc* Power Semiconductor Devices and ICs (ISPSD), 2009, pp. 216-219.

[2.32] H. Ichikawa, T. Ichimura and S. Soyano, '*IGBT Modules for Hybrid Vehicle Motor Driving*', Fuji Electric Review, vol. 55, pp. 46-50, 2009.

[2.33] B. G. Streetman and S. K. Banerjee, *Solid state electronic devices*, 6 ed. New Jersey, USA: Prentice Hall, 2006.

[2.34] Y. P. Varshni, '*Temperature dependence of the energy gap in semiconductors*', Physica, vol. 34, pp. 149-154, 1967.

[2.35] M. Naci Inci and T. Yoshino, '*Use of super luminescence diodes in absolute temperature measurement*', *in Proc* Lasers and Electro-Optics Europe, 2000.

[2.36] Y. H. Kao, M. N. Islam, J. M. Saylor, R. E. Slusher and W. S. Hobson, '*Raman effect in AlGaAs waveguides for subpicosecond pulses*', Journal of Applied Physics, vol. 78, p. 2198, 1995.

[2.37] W. Zong Liang, C. Jun, Z. Sa Sa, L. Guang Ping, W. Wei Jie, J. Shuo, L. Xiang Zhi, L. Xiao Hui, L. Sha, S. Bo Ning and L. Yong Ning, '*Spatial Resolution Improvement of Distributed Raman Temperature Measurement System*', IEEE Sensors Journal, vol. 13, pp. 4271-4278, 2013.

[2.38] M. Musallam, P. P. Acarnley, C. M. Johnson, L. Pritchard and V. Pickert, '*Power electronic device temperature estimation and control in pulsed power and converter applications*', Control Engineering Practice, vol. 16, pp. 1438-1442, 2008.

[2.39] P. L. Evans, A. Castellazzi and C. M. Johnson, '*Automated Fast Extraction of Compact Thermal Models for Power Electronic Modules*', IEEE Transactions on Power Electronics, vol. 28, pp. 4791-4802, 2013.

[2.40] G. C. James, V. Pickert and M. Cade, '*A thermal model for a multichip device with changing cooling conditions*', *in Proc* Power Electronics, Machines and Drives (PEMD), 2008, pp. 310-314.

[2.41] M. Salleras, M. Carmona and S. Marco, '*Issues in the Use of Thermal Transients to Achieve Accurate Time-Constant Spectrums and Differential Structure Functions*', IEEE Transactions on Advanced Packaging, vol. 33, pp. 918-923, 2010.

[2.42] Y. Yang, R. Master, G. Refai-Ahmed and M. Touzelbaev, '*Transient Frequency-Domain Thermal Measurements With Applications to Electronic Packaging*', IEEE Transactions on Components, Packaging and Manufacturing Technology, vol. 2, pp. 448-456, 2012.

[2.43] F. Christiaens and E. Beyne, '*Transient thermal modeling and characterization of a hybrid component*', *in Proc* Electronic Components and Technology Conference, 1996, pp. 154-164.

[2.44] S. David, W. Batty, A. J. Panks, R. G. Johnson and C. M. Snowden, '*Thermal transients in microwave active devices and their influence on intermodulation distortion*', *in Proc* Microwave Symposium Digest (MTT-S), 2001, pp. 431-434 vol.1.

[2.45] P. Kawka, G. De Mey and B. Vermeersch, '*Thermal Characterization of Electronic Packages Using the Nyquist Plot of the Thermal Impedance*', IEEE Transactions on Components and Packaging Technologies, vol. 30, pp. 660-665, 2007.

[2.46] A. Poppe, Y. Zhang, J. Wilson, G. Farkas, P. Szabo, J. Parry, M. Rencz and V. Szekely, '*Thermal Measurement and Modeling of Multi-Die Packages*', IEEE Transactions on Components and Packaging Technologies, vol. 32, pp. 484-492, 2009.

[2.47] M. Salleras, J. Palacin, G. Carles and S. Marco, '*Difficulties on the estimation of the thermal structure function from noisy thermal impedance transients*', *in Proc* Thermal, Mechanical and Multiphysics Simulation and Experiments in Micro-Electronics and Micro-Systems (EuroSime), 2006.

[2.48] L. Kim and M. W. Shin, '*Thermal Resistance Measurement of LED Package with Multichips*', IEEE Transactions on Components and Packaging Technologies, vol. 30, pp. 632-636, 2007.

[2.49] J. S. Rieh, D. Greenberg, B. Jagannathan, G. Freeman and S. Subbanna, '*Measurement and modeling of thermal resistance of high speed SiGe heterojunction bipolar transistors*', *in Proc* Silicon Monolithic Integrated Circuits in RF Systems, 2001, pp. 110-113.

[2.50] N. Y. A. Shammas, M. P. Rodriguez and F. Masana, '*A simple method for evaluating the transient thermal response of semiconductor devices*', Microelectronics Reliability, vol. 42, pp. 109-117, 2002.

[2.51] T. Bruckner and S. Bernet, '*Estimation and Measurement of Junction Temperatures in a Three-Level Voltage Source Converter*', IEEE Transactions on Power Electronics, vol. 22, pp. 3-12, 2007.

[2.52] H. Wang, S. X. D. Tan, L. Guangdeng, R. Quintanilla and A. Gupta, '*Full-chip runtime error-tolerant thermal estimation and prediction for practical thermal management*', *in Proc* Computer-Aided Design (ICCAD), 2011, pp. 716-723.

[2.53] X. Wang, K. Ma and Y. Wang, '*Adaptive Power Control with Online Model Estimation for Chip Multiprocessors*', IEEE Transactions on Parallel and Distributed Systems, vol. 22, pp. 1681-1696, 2011.

[2.54] M. Musallam and C. M. Johnson, '*Real-Time Compact Thermal Models for Health Management of Power Electronics*', IEEE Transactions on Power Electronics, vol. 25, pp. 1416-1425, 2010.

[2.55] J. D. Holmes, M. P. Foster and D. A. Stone, '*System-wide Temperature Estimation for IMS based Power Electronics Circuits*', *in Proc* International Conference on Power Electronics and Drive Systems, 2009, pp. 1081-1084.

[2.56] M. Musallam, P. P. Acarnley, C. M. Johnson, L. Pritchard and V. Pickert, '*Open loop real-time power electronic device junction temperature estimation*', *in Proc* IEEE International Symposium on Industrial Electronics, 2004, pp. 1041-1046 vol. 2.

[2.57] H. Chen, B. Ji, V. Pickert and W. Cao, '*Real-Time Temperature Estimation for Power MOSFETs Considering Thermal Aging Effects*', IEEE Transactions on Device and Materials Reliability, vol. 14, pp. 220-228, 2014.

[2.58] I. Swan, A. Bryant, P. A. Mawby, T. Ueta, T. Nishijima and K. Hamada, '*A Fast Loss and Temperature Simulation Method for Power Converters, Part II: 3-D Thermal Model of Power Module*', IEEE Transactions on Power Electronics, vol. 27, pp. 258-268, 2012.

[2.59] Y. Yonghong, G. Zhenyu, Z. Changyun, R. P. Dick and S. Li, '*ISAC: Integrated Space-and-Time-Adaptive Chip-Package Thermal Analysis*', IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 26, pp. 86-99, 2007.

[2.60] R. C. Dorf, *The Electrical Engineering Handbook*. Florida: CRC Press, 1997.

[2.61] W. D. T. Davies, *System Identification for Self-Adaptive Control*. London: Wiley-Interscience, 1970.

[2.62] A. J. Saavedra-Montes, J. M. Ramirez-Scarpetta, C. A. Ramos-Paja and P. Malik, '*Identification of excitation systems with the generator online*', Electric Power Systems Research, vol. 87, Jun 2012.

[2.63] M. Singh, M. Sakare and S. Gupta, '*Testing of high-speed DACs using PRBS generation with Alternate-Bit-Tapping*', *in Proc* Design, Automation & Test in Europe Conference & Exhibition, 2011.

[2.64] A. J. Fairweather, M. P. Foster and D. A. Stone, '*Modelling of VRLA batteries over operational temperature range using Pseudo Random Binary Sequences*', Journal of Power Sources, vol. 207, pp. 56-59, 2012.

[2.65] P. E. Wellstead, '*Non-parametric methods of system identification*', Automatica, vol. 17, pp. 55-69, 1981.

[2.66] R. L. T. Hampton, '*A hybrid analog-digital pseudo-random noise generator*', *in Proc* Proceedings of the spring joint computer conference, 1964, pp. 287-301.

[2.67] S. C. Albers, T. J. Callaghan, W. L. Rasmussen and R. A. Pajak, "Noise generator using combined outputs of two pseudo-random sequence generators," US patent 5153532, 1992.

[2.68] B. Miao, R. Zane and D. Maksimovic, '*System identification of power converters with digital control through cross-correlation methods*', IEEE Transactions on Power Electronics, vol. 20, pp. 1093-1099, 2005.

[2.69] J. D. Gabano, T. Poinot and H. Kanoun, '*Identification of a thermal system using continuous linear parameter-varying fractional modelling*', IET Control Theory & Applications, vol. 5, pp. 889-899, 2011.

[2.70] D. Defer, E. Antczak and B. Duthoit, '*The characterization of thermophysical properties by thermal impedance measurements taken under random stimuli taking sensor-induced disturbance into account*', Measurement Science & Technology, vol. 9, pp. 496-504, 1998.

Chapter III

Practical set-up of experimental rigs and methods of measurement for thermal analysis of power electronics

This Chapter proposes a number of the techniques which are used in this thesis for measurement and control of temperature and power. These methods are presented in detail and in general terms so that they may be used later in this thesis without full explanation. First, a simple and easy-to-implement circuit for arbitrarily controlling the instantaneous power dissipation in a device is proposed. By measuring the heating effect on an electronic product, an engineer can evaluate the effectiveness of its design. The controller is demonstrated practically with a MOSFET using an arbitrary waveform, and power dissipated in the device agrees well with intended dissipation. Typical experimental arrangements and measurement techniques used in the thesis are then reported. The properties and attributes of thermocouple-based measurements of device surface temperatures are considered and shown to be appropriate for high bandwidth measurements. Finally, die temperature measurements of a MOSFET are proposed as an alternative to measuring surface temperature. By applying a small bias current to the intrinsic diode in a MOSFET and measuring the voltage, the temperature can be calculated. A circuit is therefore designed to allow measurement to take place simultaneously with device power dissipation. Results show improved die temperature estimation compared to the simpler surface temperature measurement techniques.

3.1 Instantaneous power dissipation control in an active semiconductor device

There is an increasing drive to reduce the size of electronic products. For high power products, the thermal ratings of components present a significant constraint on the design. An engineer must therefore ensure that components are adequately spaced and cooled to ensure that ratings are observed. Improvements in efficiency notwithstanding,

reductions in volume can be achieved by more effective heat sinking, forced cooling and allowing components to operate closer to their ratings for longer [3.1][3.2][3.3]. To make these improvements, it is important to understand how the power dissipation in one device affects the temperature and characteristics of other components [3.3]. By applying an arbitrary power dissipation waveform to a device, mimicking operating conditions, it is possible to observe the transient and steady-state responses of its packaging, circuit board and heat sink arrangement. It is also possible to observe how the collective power dissipation and ambient conditions around a group of devices bring each device close to its individual rating and affect its reliability. As a result, an engineer can evaluate the effectiveness of the thermal design during product development.

In order that a device accurately mimics its operating dissipation, a circuit is required to drive it such that its instantaneous power dissipation matches the desired operating waveform. Previous literature [3.4] has presented a method for this which involves pulse-width modulating a constant current and relying on the almost constant on-state resistance of a MOSFET, $r_{\mathrm{ds(on)}}$, to dissipate power corresponding to an arbitrary waveform. In addition to complex switching logic and dummy loads, this set-up also requires feedback on the true value of $r_{\mathrm{ds(on)}}$ to account for differences between devices. It will not work with bipolar transistors, for example, whose on-state characteristics are not defined by a resistance. A more general purpose and easier to implement design is therefore presented in this section.

3.1.1 Design methodology

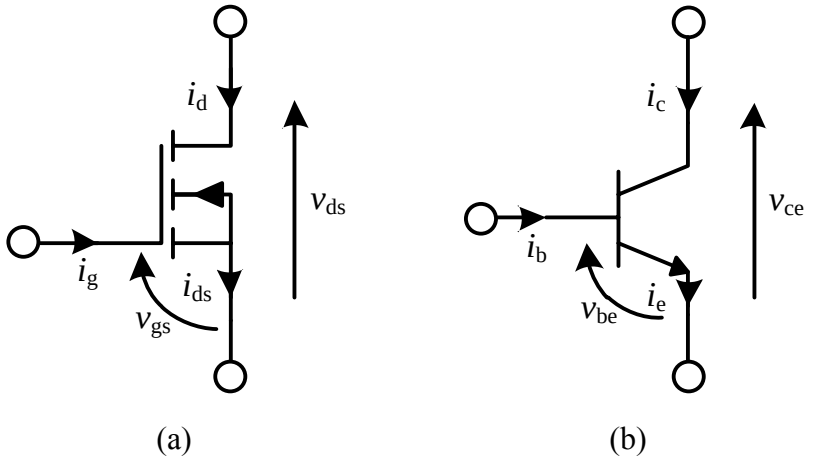

Fig. 3.1 Circuit diagrams of MOSFET (a) and BJT (b) with relevant voltages and currents indicated

For a bipolar junction transistor (BJT) or field-effect transistor (FET), such as those shown in Fig. 3.1, the instantaneous power dissipation (q) can be expressed as in (3.1) and (3.2) (see [3.5]). Power dissipation due to gate or base current can be neglected since gate current and the internal resistance of the gate terminal are negligible and h_{FE}, the BJT current gain, is large for most devices.

$$q_{\mathrm{FET}} = v_{\mathrm{ds}} i_{\mathrm{d}} \tag{3.1}$$

$$q_{\mathrm{BJT}} = v_{\mathrm{ce}} i_{\mathrm{c}} + v_{\mathrm{be}} i_{\mathrm{b}} = \left(\frac{v_{\mathrm{ce}}}{1 + 1/h_{\mathrm{FE}}} + \frac{v_{\mathrm{be}}}{\mathrm{h_{FE}} + 1} \right) i_{\mathrm{e}} \approx v_{\mathrm{ce}} i_{\mathrm{e}} \tag{3.2}$$

where q_{FET} and q_{BJT} are the instantaneous power dissipation in the FET and BJT respectively; v_{ds} and v_{ce} are the instantaneous drain-to-source (for a FET) and collector-to-emitter (for a BJT) voltage; i_{d} and i_{e} are the instantaneous drain (for a FET) and emitter (for a BJT) current; i_{g} and i_{b} are instantaneous gate (for a FET) and base (for a BJT) current; and i_{c} and h_{FE} are the collector current and current gain of a BJT respectively.

Considering the case of a MOSFET, the power dissipated can be controlled by a variation in v_{ds}, i_{d} or both. Hypothetically possible approaches include varying v_{ds} while keeping i_{d} constant, but this prevents q dropping to zero because of the on-state resistance; and varying i_{d} while keeping v_{ds} constant, but this requires complex circuitry to ensure stability since the gate signals for each transistor in the circuit are in anti-phase. In Fig. 3.2, a design which varies v_{ds} and i_{d} in proportion to each other is presented.

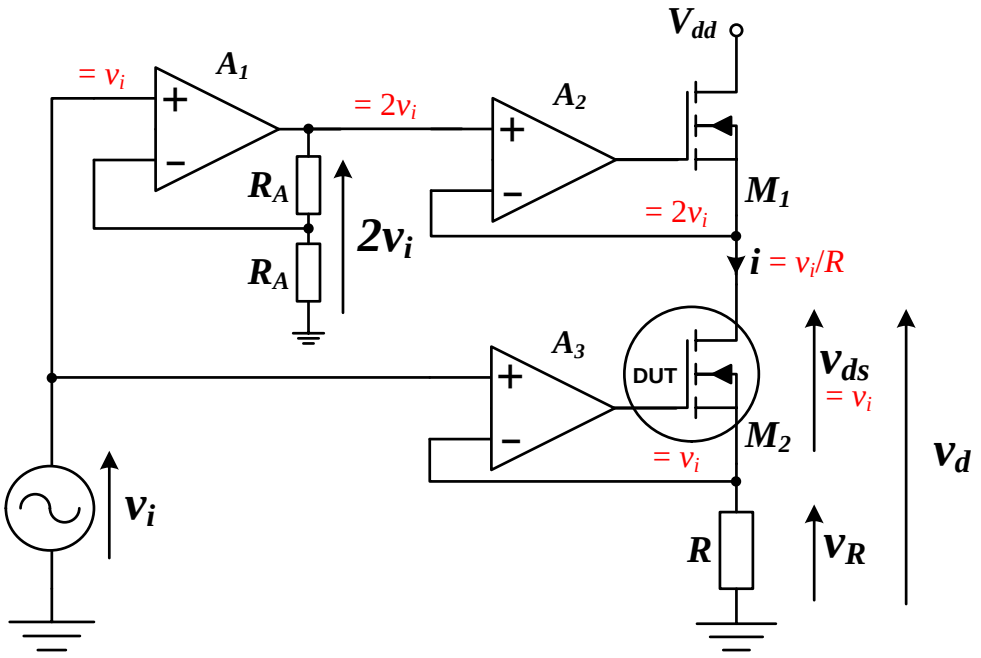

Fig. 3.2 Proposed design of the arbitrary waveform power controller

The combination of A_3, M_2 and R form a constant current sink. The op-amp feedback controls the drain-source resistance of M_2 such that $v_{\mathrm{R}} = v_{\mathrm{i}}$, hence $i = v_{\mathrm{i}}/R$. This is valid as long as the voltage v_{d} is sufficiently greater than v_{R} such that $v_{\mathrm{ds}} \geq i \times r_{\mathrm{ds(on)}}$ (i.e. $R \geq r_{\mathrm{ds(on)}}$) or $v_{\mathrm{ce}} \geq v_{\mathrm{ce(sat)}}$ for a BJT, the reason for which will be given shortly.

The combination of A_1 and resistors R_A provide an amplifier of gain 2. A_2 and M_1 are arranged in voltage-follower configuration and hence they keep v_{d} at $2v_{\mathrm{i}}$. The drain voltage (v_{d}) is therefore twice the source voltage (v_{R}) and it follows that $v_{\mathrm{ds}} = v_{\mathrm{R}}$. This means the voltages across R and M_2 are identical and, since the same current passes through both, they have identical instantaneous power dissipation as given in (3.3). R can be adjusted to set the appropriate level of power dissipation and to ensure v_{ds} and i stay within device ratings for the required dissipation.

$$q(t) = \frac{v_{\mathrm{i}}(t)^2}{R} \tag{3.3}$$

Because the same current passes through the MOSFET and resistor, the power dissipated is determined by the effective resistance of each component. For equal

dissipation, these two resistances must be equal. The channel resistance of a MOSFET may be increased to any value above $r_{\mathrm{ds(on)}}$, and therefore this is the minimum working resistance for R, which is fixed. For BJTs, R must be large enough such that the peak input voltage, $\hat{v}_{\mathrm{i}}$, is greater than $v_{\mathrm{ce(sat)}}$ at $i = \hat{v}_{\mathrm{i}}/R$ otherwise the base of M_2 will saturate and the required voltage drop will not be possible.

The input waveform is set to the square root of the required arbitrary power waveform and scaled according to the required average power. The waveform would be supplied by an arbitrary function generator or from the output of a digital-to-analogue converter (DAC) from a microcontroller.

3.1.2 Accuracy

Limitations to the accuracy of the power controller result from the component tolerances and op-amp accuracy. The true power dissipation can be calculated from $i_{\mathrm{d}}v_{\mathrm{ds}}$ where i_{d} and v_{ds} are accurately measured. Since i_{d} depends on the accuracy of R and voltage follower A_3, and v_{ds} depends on the tolerance of resistors R_A and accuracy of op-amps A_1 and A_2, high performance and low tolerance components must be selected in a practical implementation. The tolerance of resistors R_A will typically be around 0.1% for accurate voltage doubling. Power resistor R must also be low tolerance with reduced temperature sensitivity as it directly affects power as shown in (3.3). The effect of temperature sensitivity can be reduced by effective cooling of R. The gain, input offset voltage and slew rate of the op-amps introduce an offset voltage between the desired and actual output voltage and therefore high performance op-amps are required. Notwithstanding these design considerations, at the low frequencies required for the thermal experiments described in this thesis, op-amp proportional gain produces good results.

3.1.3 Results

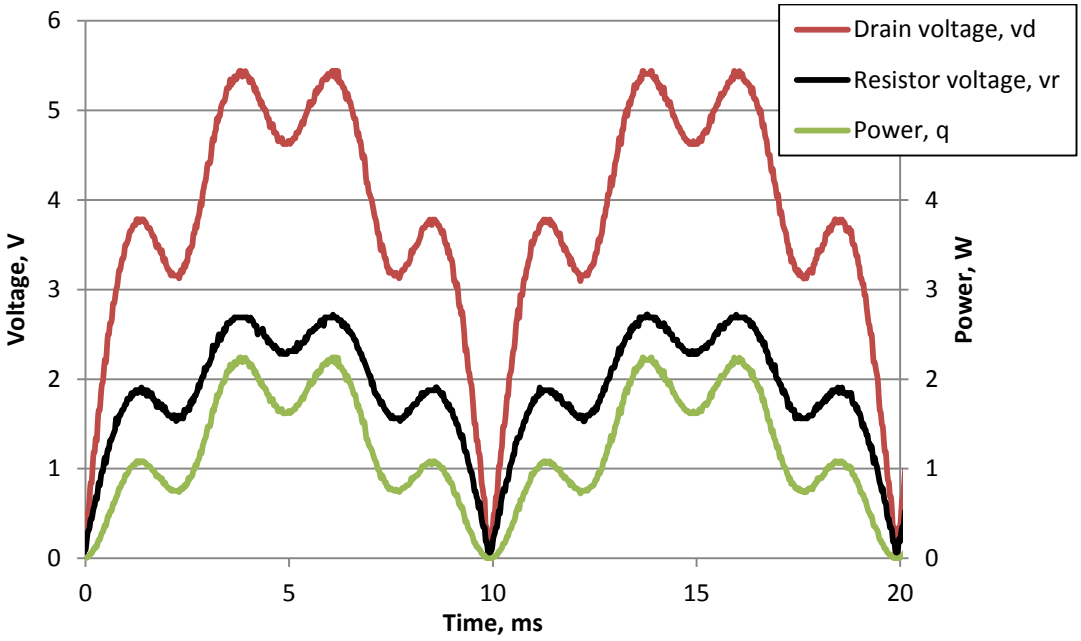

Fig. 3.3 Voltage and power waveforms observed on implemented system

Fig. 3.3 shows the measured voltage (v_{R}, v_{d}) and power (q) characteristics for the power controller circuit in Fig. 3.2, as constructed with op-amp LM324, MOSFET IRF510, $R_A = 1\,\mathrm{k\Omega}$ and $R = 3.3\,\Omega$. V_{dd} is regulated at 12 V. An arbitrary power dissipation function, $q(t)$, of $A(\sin\omega t + \sin^2 2\omega t + 1)$ is selected, where A is a constant (fixed at 2.65 W), t is time and ω is angular frequency. From (3.3), the input voltage waveform, $v_{\mathrm{i}}(t)$, must be $\sqrt{AR(\sin\omega t + \sin^2 2\omega t + 1)}$. An 8-bit R-2R ladder DAC driven by a microcontroller is used to produce this waveform at 100 Hz.

The power output calculated from v_{i}^2/R shows good agreement with the intended function; they have a Pearson product-moment correlation coefficient of 0.998, an average power difference of 0.8% and the noise level is low. The direct control of both voltage and current ensures that power control is independent of semiconductor device nonlinearities.

Fig. 3.4 Practical implementation of the power controller

This simple circuit, which can control the power dissipation in an active semiconductor device by controlling the device's voltage and current, has applications throughout this thesis. Its practical construction is shown in Fig. 3.4. By applying an arbitrary power waveform which mimics operating or test power dissipation, it is possible to evaluate the effect heating has on a particular system design, and therefore modify the system accordingly. This circuit is therefore repeatedly used in the thesis where the power dissipation in an active semiconductor device is to be controlled.

3.2 Experimental arrangements and the measurement of external temperatures

3.2.1 Operation of thermocouples

Throughout this thesis, proposed experimental techniques are verified by applying power waveforms to discrete power electronic components and measuring the resulting temperature. The primary way by which temperatures are measured in the presented work is by using K-type thermocouples. Thermocouples are passive temperature-voltage transducers whose operation is based on the Seebeck effect wherein a voltage is generated across a temperature gradient. Thermocouples utilise a junction at unknown temperature between two different metals. The voltage generated at a distant point of

known temperature, called the cold junction, is measured to calculate the difference in temperature between the two junctions. Fig. 3.5 shows a schematic of the arrangement.

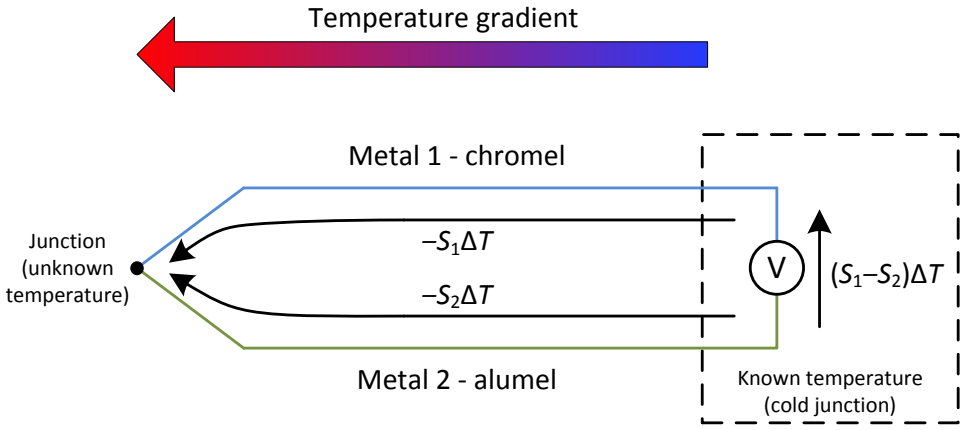

Fig. 3.5 Measurement of temperatures using K-type thermocouples and the Seebeck effect

Since the two junctions are at different temperatures, there must be a temperature gradient across the wires. Due to the Seebeck effect, there will also be a proportional voltage gradient. However, for thermocouples, only the voltage and temperature difference between the terminals is relevant. The overall voltage generated across wire 1 is

$$\Delta V_1 = -S_1 \Delta T \tag{3.4}$$

Where S_1 is the Seebeck constant for the particular metal from which the wire is made, and ΔT and ΔV_1 are the differences in temperature and voltage, respectively, across the wire. In a thermocouple, two different metals are used and the temperature at the cold junction is the same for both metals, therefore a voltage, V, is apparent between the metals at the cold junction. Current is not allowed to flow in the wires and therefore, from Kirchoff's voltage law, the measured voltage can be written as

$$V = \Delta V_1 - \Delta V_2 = (S_1 - S_2)\Delta T \tag{3.5}$$

$S_1 - S_2$ is the temperature coefficient: it is the difference between the Seebeck coefficients for the particular metals used. In this thesis, the K-type (chromel-alumel)

thermocouple is used because its operating temperature range (0 to 292°C) covers the operational temperature range of the power electronic devices considered. For these metals, the temperature coefficient is 392 μV/K. The absolute temperature is calculated by adding the calculated ΔT to the cold junction temperature, which is measured using a different type of temperature sensor, in this case capable of measuring absolute temperature. This process is known as cold junction compensation.

Commercial thermocouples may be produced with time constants of a few hundred milliseconds [3.6]. In practice, thermal time constants of systems to be characterised are much longer than this value, usually at least several seconds and often many orders of magnitude greater. The rapid response time of a thermocouple therefore allows frequency-domain temperature measurements to be made at the higher frequencies required. The low thermal mass and low cost of thermocouples are also attractive features. The cold junction sensor is usually slower than the thermocouple. As long as the cold junction remains at a steady temperature, however, only occasional recalculation of the cold junction temperature is required.

3.2.2 Attachment of thermocouples to devices

The thermocouple is attached to the area whose temperature is to be measured. In most cases, the temperature of interest is the active junction on the semiconductor die inside a power electronic device. The thermocouple is therefore attached as close as practicable to the die, typically on the metal part of the packaging adjacent to the die. Fig. 3.6 shows a diagram of a thermocouple attached to the tab of a TO-220 device.

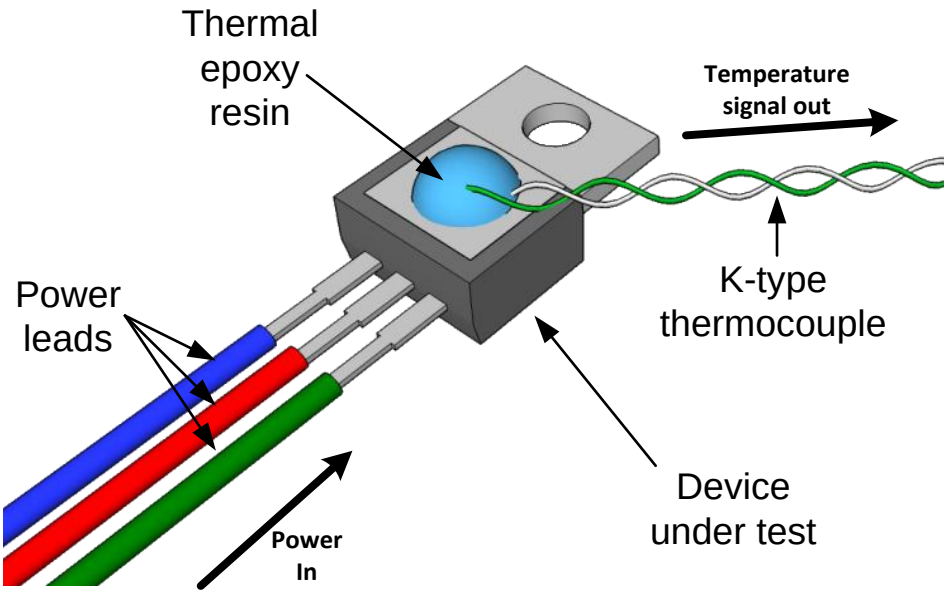

Fig. 3.6 Thermocouple-based measurement of temperatures

The thermocouple is secured using a type of epoxy resin adhesive specifically designed for high thermal conductivity and low electrical conductivity. For the work documented in this thesis, *Electrolube thermal bonding compound (TBS)* was used, which has a thermal conductivity of 1.1 W/mK and an electrical resistivity of 1 TΩm. As small a volume of epoxy as is practicable is used to attach the junction of the thermocouple to the device in order to minimise the effect the thermal mass and resistance of the adhesive has on temperature measurements. Care is also taken to ensure the junction does not touch the metal tab of power devices, since both are commonly attached to external circuits on the same power supply and the electrical coupling would therefore affect the voltage (and hence temperature) measurement.

3.2.3 Experimental arrangement

An experimental set-up is produced for each experiment based on the general experimental arrangement shown in Fig. 3.8 (page 57). Power dissipation in devices is controlled by an external circuit. For example, for semiconductor devices, the arbitrary waveform power controller presented in section 3.1 is used. The controller receives a demand signal from a waveform generator.

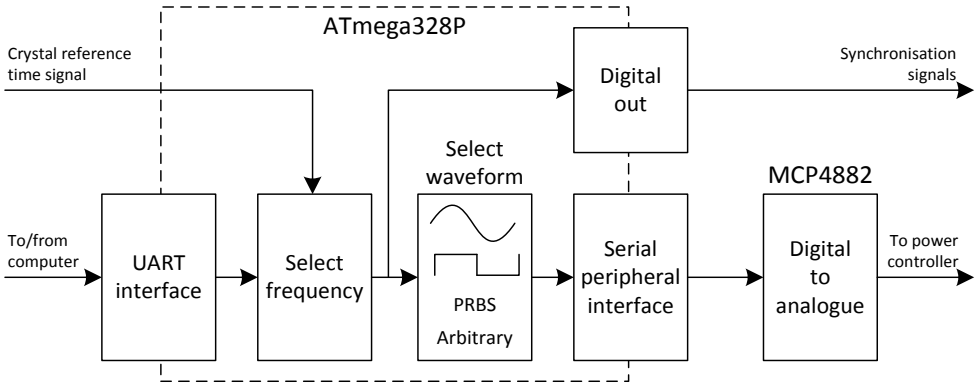

Fig. 3.7 Waveform generation and experimental control

For the purposes of this thesis, a bespoke microcontroller-based waveform generator was produced. The topology of the waveform generator is shown in Fig. 3.7. The generator is capable of producing arbitrary waveforms at frequencies of up to 100 Hz and as low as 1 μHz, which covers all necessary signals in this work. Good timing and low jitter is ensured by the use of a calibrated high precision (10 ppm tolerance) quartz crystal. The voltage output is generated by an MCP4882 digital-to-analogue converter, attached to the microcontroller. Synchronisation signals are also generated to provide a time reference for the signal processing stage.

Temperature, power and synchronisation signals are recorded using a LabJack U6 datalogger. The LabJack has a sixteen channel bipolar non-differential 24-bit analogue-to-digital converter with programmable gain. For high precision measurements (range ±0.1 V), it has an effective resolution of 2.4 μV at 25 Hz, allowing a temperature resolution of 6 mK to be achieved when K-type thermocouples are used. At low precision (range ±10V), its effective resolution is 160 μV, more than sufficient for measurements of power dissipation [3.7].

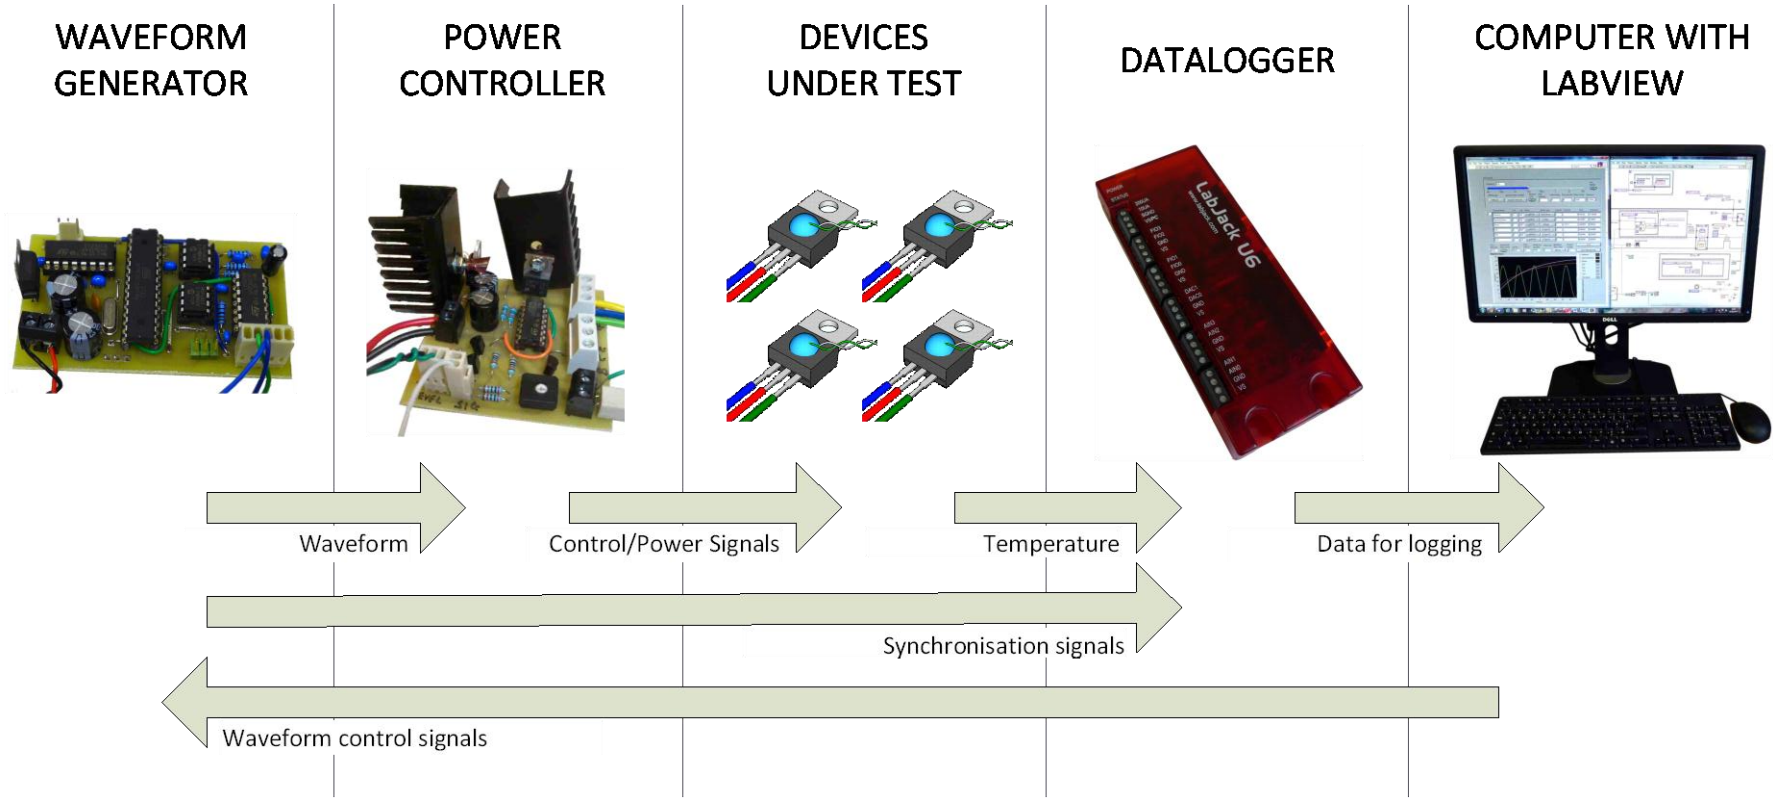

Fig. 3.8 Practical implementation of temperature logging and experiment control

The LabJack is interfaced with a computer running LabView, which provides experimental control and management. The LabView programmes are modified for each experiment, but they all operate by controlling the output waveform and frequency of the waveform generator (communicated with over serial port) and recording all temperature, power and synchronisation data from the LabJack. Once the experiment is complete, the data is transferred to MATLAB for post-processing and analysis in line with the requirements of each experiment.

3.2.4 Thermal imaging

In addition to thermocouple measurement, surface temperatures can be recorded using a 2-dimensional thermal imaging camera, as shown in Fig. 3.9.

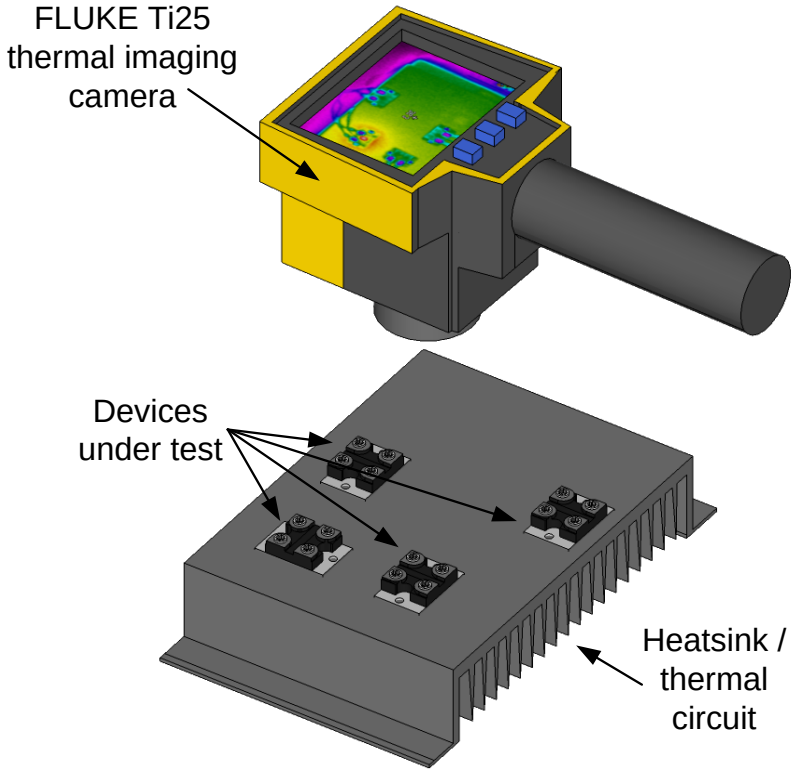

Fig. 3.9 Thermal imaging of devices and heatsink

This type of measurement allows the temperature over an entire heatsink to be recorded simultaneously without affecting the system by invasively adding sensors. Images can be used to determine a temperature profile across a system or to observe which devices are operating at the highest temperatures under given operating conditions. While these features offer advantages, there are a number of drawbacks. Thermal images are affected

by the colour of the exposed surfaces (blacks appear hotter than silvers, for instance). In addition, all interesting temperatures must not be visually obscured. This contrasts with thermocouple measurements which can be taken in difficult-to-observe places. In addition, thermal imaging cameras are expensive with basic models usually limited to low sampling rates but still costing thousands of pounds. These cameras are therefore useful for supplementary analysis, particularly where the overall temperature profile is interesting, whereas thermocouple measurements provide fast accurate single point temperature readings.

3.3 Measurement of die temperature

Thermocouple-based temperature measurements are effective at measuring the surface temperatures of devices with high rapidity and sensitivity. However, in most circumstances it is the temperatures of the semiconductor dice which are of interest. While a temperature measurement of the metallic tab on a power device is a close approximation of die temperature, the thermal resistance between the tab and the thermal mass of the device itself introduce an additional time constant and temperature attenuation, which limits the usefulness of the data.

To avoid this problem, direct temperature measurement of the die is possible in many cases. A number of techniques are available, several of which are described in Chapter II. In this thesis, characterisation of thermal properties is achieved by applying special power dissipation waveforms while simultaneously measuring temperature. In most cases, thermocouple-based surface temperature measurements are used due to their simplicity of implementation. However, in some case it would be preferable to know die temperature since it is the active junction on the semiconductor die which is subject to thermal limits. Although it would be preferable to measure junction temperature, such measurements usually require the exclusive use of the semiconductor component which normally provides power dissipation. Therefore, unless the die is specifically redesigned to incorporate a separate on-die temperature sensor, power dissipation and temperature measurement cannot occur simultaneously.

3.3.1 Temperature measurement using the diode forward voltage drop

Die temperature estimation is achieved by measuring the forward voltage drop across the intrinsic diode in a MOSFET when it is forward biased by a small fixed current. This normally parasitic device operates in the same way as a normal PN junction diode and its electrical characteristics are therefore described by the Shockley diode equation, which relates the bias current through the diode to its forward voltage drop [3.8], as shown in (3.6).

$$I = I_{\mathrm{s}}\left(\exp\left(\frac{eV}{kT}\right) - 1\right) \tag{3.6}$$

Where V is the forward voltage drop, e is elementary charge, k is the Boltzmann constant and T is absolute temperature. I_{s} is the diode's saturation current, which can be defined in terms of the physical characteristics of the device [3.8] as

$$I_{\mathrm{s}} = Aen_{\mathrm{i}}^2\left(\frac{D_{\mathrm{h}}}{L_{\mathrm{h}}N_{\mathrm{d}}} + \frac{D_{\mathrm{e}}}{L_{\mathrm{e}}N_{\mathrm{a}}}\right) \tag{3.7}$$

Where A is the junction area, D_{h} and D_{e} are the diffusivity coefficients of holes and electrons, respectively; L_{h} and L_{e} are the minority carrier diffusion length of holes and electronics; and N_{d} and N_{a} are the donor and acceptor atom concentrations, respectively. n_{i} is the carrier concentration in an intrinsic semiconductor of the same material. It can be written as [3.8]

$$n_{\mathrm{i}}^2 = 4\left(\frac{2\pi kT}{h^2}\right)^3 (m_{\mathrm{e}}^* m_{\mathrm{h}}^*)^{3/2} \exp\left(-\frac{E_{\mathrm{g}}}{kT}\right) \tag{3.8}$$

Where h is the Planck constant, m_{e}^* and m_{h}^* are the effective electron and hole masses, respectively; and E_{g} is the semiconductor bandgap. Because most of the terms in equations (3.7) and (3.8) are physical, material or structural constants, these constants can be grouped together into a single device constant, α, to simplify the expression. A simplified expression for I_{s} is therefore given in (3.9).

$$I_s = \alpha T^3 \exp\left(-\frac{E_g}{kT}\right) \tag{3.9}$$

$$\alpha = 4eA\left(\frac{D_h}{L_h N_d} + \frac{D_e}{L_e N_a}\right)\left(\frac{2\pi k}{h^2}\right)^3 (m_e^* m_h^*)^{3/2} \tag{3.10}$$

The Shockley diode equation in (3.6) can be rearranged to give the forward voltage drop as a function of other parameters. In practical circumstances, $I \gg I_s$, therefore the Shockley diode equation can be rewritten as

$$V = \frac{kT}{e} \cdot \ln\left(\frac{I}{I_s}\right) \tag{3.11}$$

By substituting in I_s, the following expression is obtained.

$$V = \frac{kT}{e} \cdot \ln\left(\frac{I}{\alpha T^3 \exp\left(-\frac{E_g}{kT}\right)}\right) \tag{3.12}$$

Since it is a mathematical identity that $\ln(a/b) \equiv \ln a - \ln b$, the expression for V can be rewritten as

$$\begin{aligned} V &= \frac{kT}{e} \cdot \left[\ln\left(\frac{I}{\alpha T^3}\right) + \frac{E_g}{kT}\right] \\ &= \frac{kT}{e} \ln\left(\frac{I}{\alpha T^3}\right) + \frac{E_g}{e} \end{aligned} \tag{3.13}$$

Equation (3.13) gives the forward voltage drop, V, as a function of the temperature, T, and is the basis for temperature estimation using the intrinsic diode. If the bias current, I, is kept constant then the change in voltage due to change in temperature can be determined. An expression for this gradient is given in (3.14).

$$\frac{\partial V}{\partial T} = \frac{k}{e} \cdot \left[\ln\left(\frac{I}{\alpha T^3}\right) - 3\right] \tag{3.14}$$

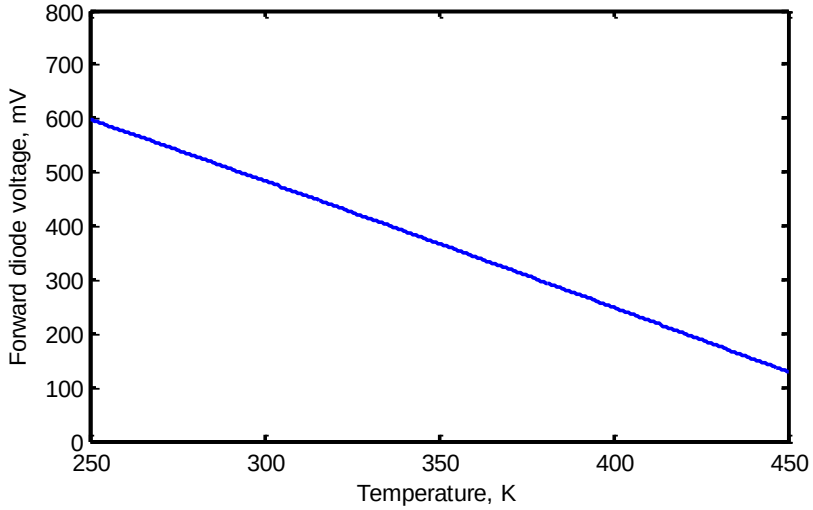

(a)

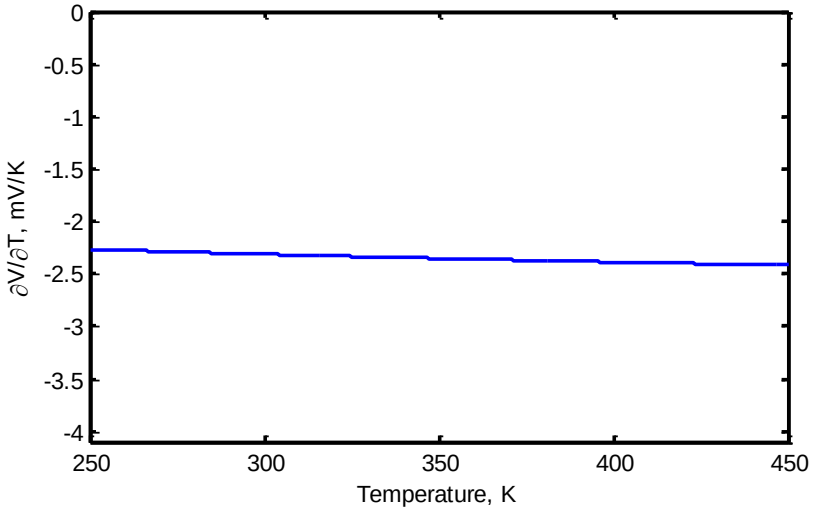

(b)

Fig. 3.10 Theoretical relationship between forward voltage drop and junction temperature of the intrinsic diode**.** (a) Voltage against current from (3.13). (b) Gradient of curve from (3.14).

For a practical power diode, α is of the order of 1 A/K^3 (where A/K^3 are the SI units of α; α is calculated from I_s for diode MUR840 at 150°C [3.9]). Fig. 3.10(a) shows the relationship between forward voltage drop of the intrinsic diode and die temperature at a constant biasing current of 10 mA, calculated from (3.13). The gradient of the line is shown in Fig. 3.10(b). It can be seen that for practical temperatures (250 to 450 K, i.e. −20 to 180°C), the relationship between voltage and temperature is near-linear. This is

clearly shown by the gradient plot which shows little variation in gradient over the entire range.

3.3.2 Calibration of the diode temperature sensor

It is therefore possible to produce a die temperature sensor using the intrinsic diode in a MOSFET. The device must first be calibrated to determine the gradient and intercept of the relationship. To achieve this calibration, the device under test (DUT; a MOSFET) is placed in an insulated box, as shown in Fig. 3.9. The inside of the box is heated by a separate heating element (a power resistor), so that the temperature inside the box is uniform after the passing of transient effects. A thermocouple is attached to the DUT to provide an independent temperature reading for the calibration.

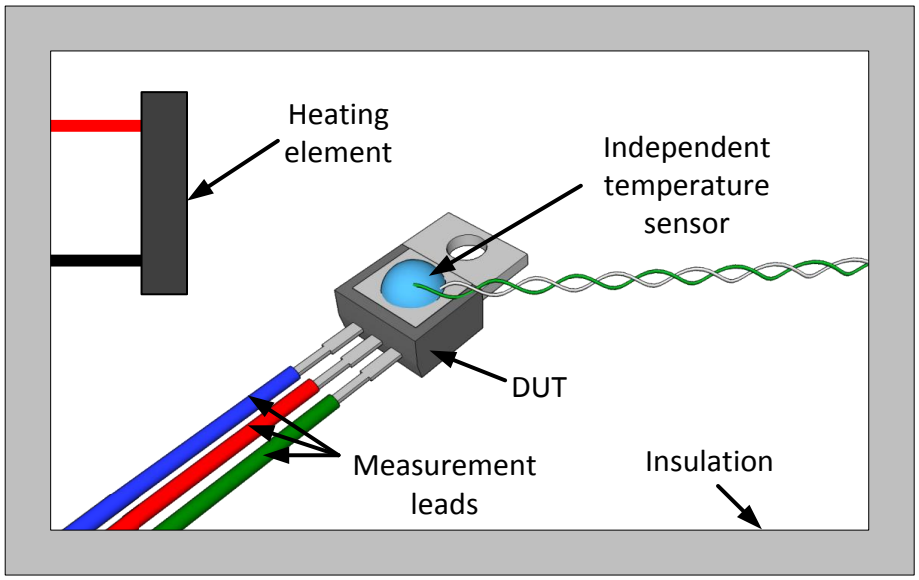

Fig. 3.11 Experimental arrangement for intrinsic diode temperature measurement calibration

A 10 mA bias current is applied between source and drain (with the gate shorted to the source to prevent the formation of a conduction channel) so that the current passes through the intrinsic diode. The interior of the box is heated to a number of temperatures in turn and the forward voltage drop of the intrinsic diode and the actual box temperature is measured at each. A settling time is permitted at each new temperature to ensure the die and case temperatures of the MOSFET are equal. Fig. 3.12 shows the measured relationship between the case temperature measured by the thermocouple and the diode voltage drop. A line of best fit is shown which demonstrates

the linearity of the result. The correlation of the results to the line is excellent demonstrating that good die temperature determination can be achieved.

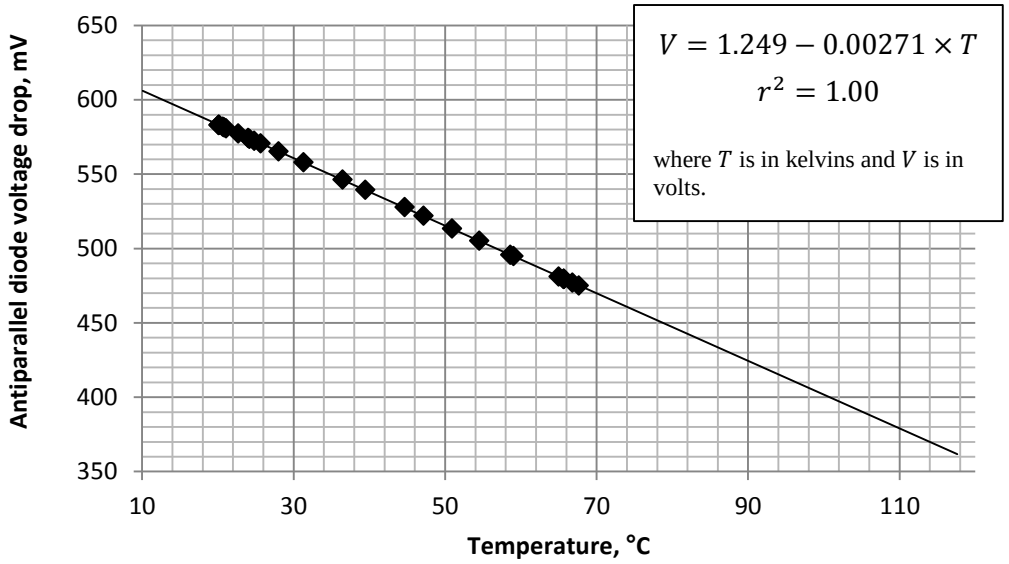

Fig. 3.12 Practical results of intrinsic diode measurement with line of best fit shown. The equation of the line and the Pearson product-moment correlation coefficient (r) are shown.

3.3.3 Near-simultaneous controlled power dissipation and die temperature measurement

For thermal characterisation, it is necessary to measure the die temperature while simultaneously dissipating power in the die. This presents a significant difficulty because it is not practical to measure temperature by using the device in a high dissipation mode. To overcome this, a circuit is developed to allow a MOSFET device to be switched between dissipating and measurement modes so that both abilities may be used. By performing the switch rapidly and adjusting power dissipation to account for the non-dissipating periods, power can be controlled with a sufficient bandwidth for practical purposes while near-simultaneous temperature measurement can also be achieved. The concept of the circuit is shown in Fig. 3.13.

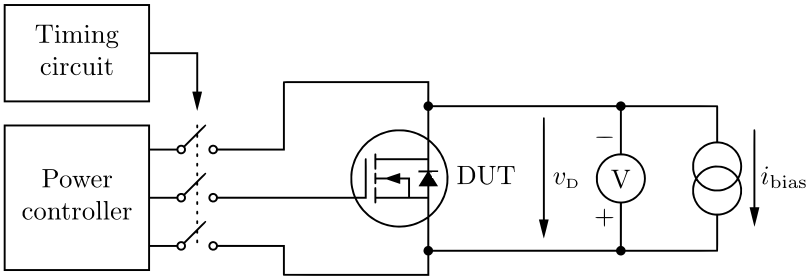

Fig. 3.13 Simplified circuit diagram of MOSFET with heat dissipation and intrinsic diode temperature measurement

A current source, connected from source to drain of the device under test (DUT), provides a small bias current, i_{bias}. This current is typically 10 mA for a power MOSFET, meaning that it does not cause significant power dissipation as a result of the measurement (i.e., dissipation is typically less than 10 mW for the duration of the measurement). A voltmeter is connected across the intrinsic diode to measure the forward voltage drop from which the die temperature is calculated.

The arbitrary waveform power controller presented in section 3.1 is also connected to the DUT. Under normal operation, this power controller determines the dissipation in the device. However, in this arrangement the current source allows a small amount of current to bypass the device, resulting in a small reduction in power dissipated. If accuracy is particularly important, the input waveform to the controller can be adjusted to correct for the error; however, this error is normally negligible. The power controller can be disconnected to allow measurement to take place.

A measurement interval is selected to be sufficiently long such that the proportion of time required for voltage measurement is small, whilst the measurement interval is also sufficiently short such that the bandwidth requirements are met. An interval of between 0.1 and 1 s is a practical range since most thermal time constants are significantly longer and enough time is available for voltage measurement. At the beginning of each measurement interval, the power controller is disconnected and the dissipation ceases. The drain-source voltage is therefore determined by the bias current and intrinsic diode. After a short settling time, the voltage is sampled and the power controller is reconnected for the remainder of the measurement interval. By keeping the disconnect time small (for example, less than 20% of the measurement interval) and increasing

power dissipation by the same proportion, near-simultaneous measurement can be achieved.

The two parts of this circuit, the power controller and current source, operate in intrinsic and therefore special control and isolation are required. The current source is isolated from the power controller circuit because it requires a power source of the reverse polarity. The schematic of the current source and voltage measurement is shown in Fig. 3.14.

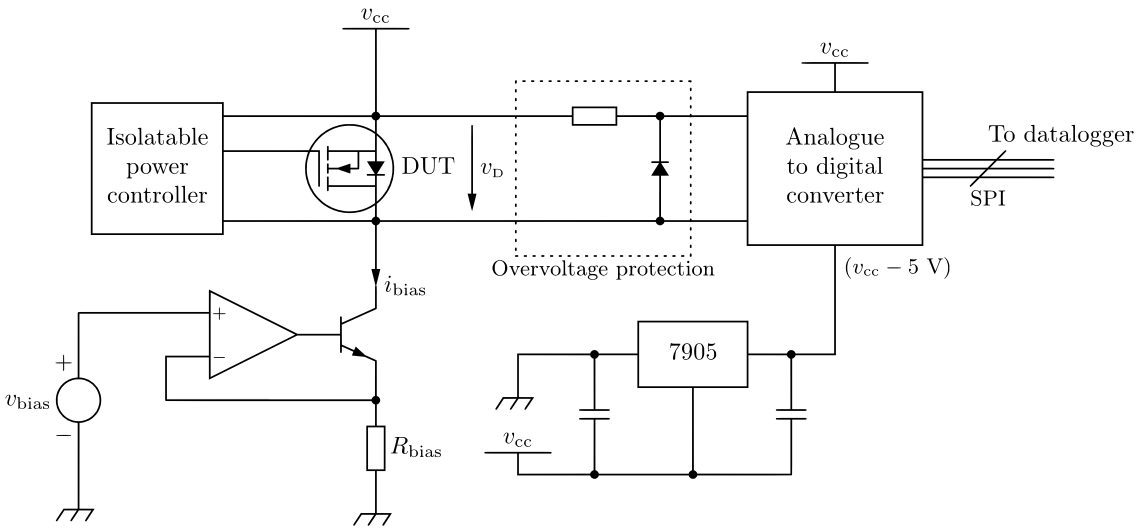

Fig. 3.14 Isolated intrinsic diode forward voltage measurement (note the upside-down MOSFET, labelled *DUT*)

The bias current is provided a constant current sink, controlled by v_{bias} and R_{bias}. An analogue-to-digital converter (ADC) performs the voltage measurement. Because the absolute voltage at the cathode of the diode is ill-defined, the ADC is instead referenced to the supply voltage, v_{cc}. The negative supply voltage for the ADC is achieved using a 7905 negative voltage regulator which provides a −5 V supply with respect to v_{cc}. The ADC performs differential voltage measurement and sends this data to a datalogger using the serial peripheral interface (SPI) protocol. The ADC is given overvoltage protection for the case when the DUT is being used to dissipate and therefore a large negative voltage would otherwise be applied to the ADC's inputs.

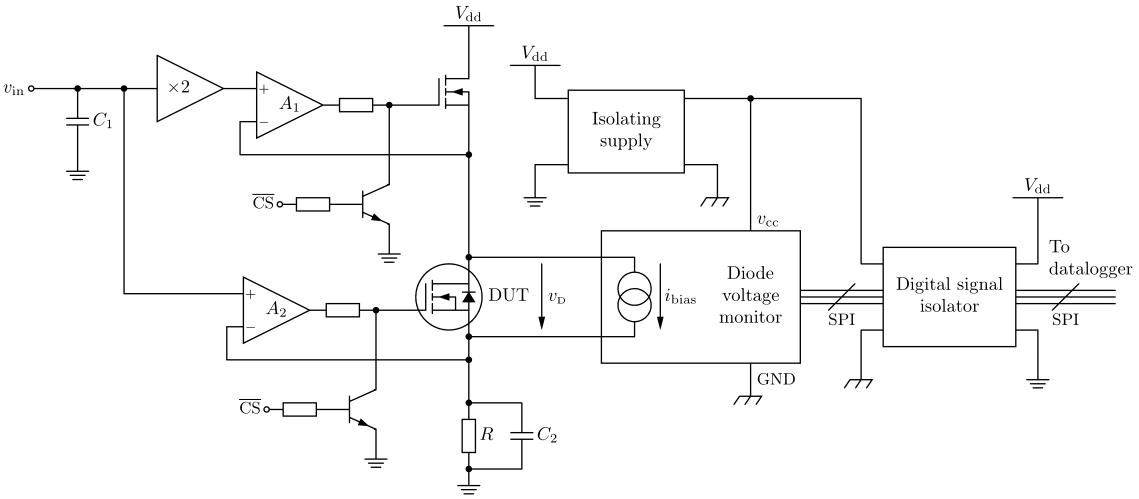

Fig. 3.15 Power controller circuit for the *DUT*, with the position of the intrinsic diode voltage monitor labelled

The power controller, shown in Fig. 3.15, is a modified version of the power controller proposed in section 3.1. In this case, the gates of both the DUT and the other MOSFET are grounded when a signal—called circuit select (CS)—is not applied (i.e. when $\overline{\mathrm{CS}} = 1$). The grounding performs the task of the switch shown in Fig. 3.13. Decoupling capacitors C_1 and C_2 are added to reduce noise caused by the rapid current change. The power dissipated in the circuit is, from (3.3),

$$q = \frac{v_{\mathrm{in}}^2}{R}\delta \tag{3.15}$$

Where δ, the duty cycle, is the proportion of time for which the power controller is connected. The diode voltage monitor is powered from an isolated supply to allow the DUT voltage to float freely with operating conditions. A digital signal isolator is used to couple the SPI output of the voltage monitor ADC to the unisolated supply so that it can be logged. The circuit is operated using a microcontroller to time the operation of $\overline{\mathrm{CS}}$ and the sampling of the voltage and convert voltage measurements to temperature. For precise temperature measurement, a high precision ADC is required. At this level of precision, economically viable ADCs can perform only a few conversions per second, so accuracy is important. Sigma-delta conversion, which uses oversampling to reduce the complexity of quantisation is the preferable form of conversion because it is less sensitive to aliasing noise [3.10].

A simplified timing diagram is given in Fig. 3.16.

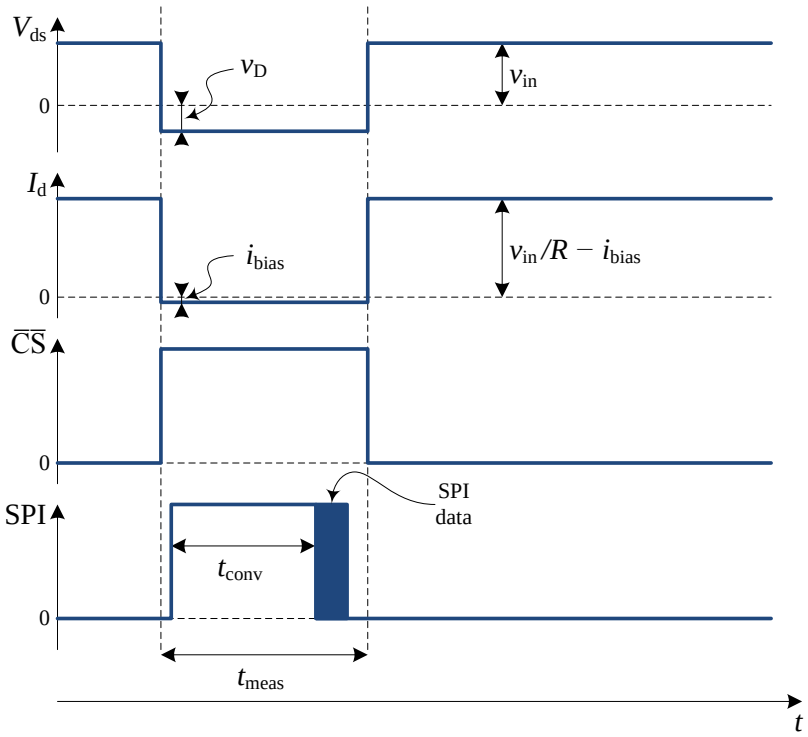

Fig. 3.16 Timing diagram for MOSFET heat dissipation and die temperature measurement

For most of the time, the circuit operates in power control mode with $V_{ds} = v_{in}$ and $I_d \approx v_{in}/R$. At the beginning of the measurement interval, the power controller is disconnected and the intrinsic diode determines the voltage. The diode voltage, v_D, is measured after a short settling time by activating the ADC. Once conversion is complete, the SPI data is read and the DUT returns to power dissipation mode.

3.3.4 Results

To demonstrate practical implementation, an MCP3553 22-bit sigma-delta ADC was used. This ADC has a sampling rate of 60 Hz, therefore an ADC conversion period (t_{conv}) of 17 ms was required. Buffers of 0.5 and 2.5 ms were added at the beginning and end of the conversion, respectively, to prevent errors arising from settling transients, leading to a total measurement period (t_{meas}) of 20 ms. The time between successive samples (the overall sampling rate of the circuit) was chosen as 200 ms, giving a duty cycle of $\delta = 0.9$.

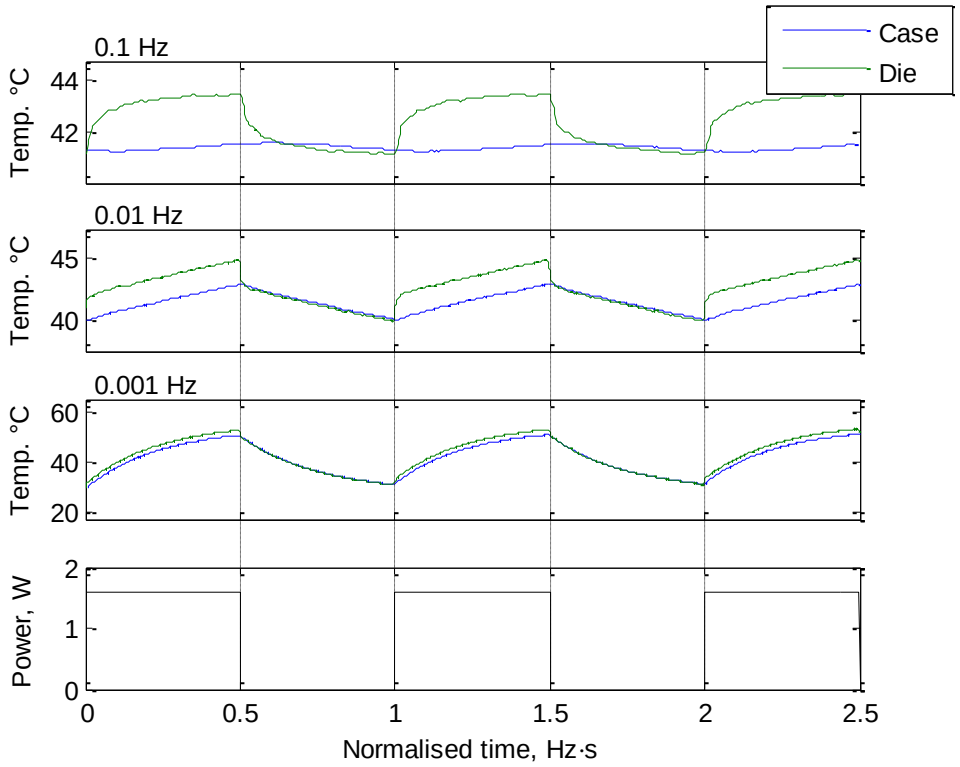

Fig. 3.17 Response of case and die temperature to square wave power input at multiple frequencies

A 1.6 W amplitude square wave power input was applied to an IRF510, packaged in a TO-220 case, attached to a 5 K/W heatsink at three frequencies in turn: 0.1 Hz, 0.01 Hz and 0.001 Hz. Fig. 3.17 shows the case and die temperature responses for each frequency. For ease of comparison, the time base of each graph is normalised to the frequency of the power input. In each case, the die temperature exceeds the case temperature by approximately 2°C while the device is dissipating; the gap in temperature closes when the dissipation stops.

The die temperature approximates a first order low pass temperature response added to the case temperature. This is because the thermal capacity of the die is small compared to the case and therefore the power entering the die is almost the same as that entering the case. The temperature difference is caused by the thermal resistance between the two. A small response time can be observed at high frequency, but may be neglected at low frequency.

At high frequencies, the difference between die and case temperatures is pronounced. For the 0.1 Hz signal, the case temperature is barely affected while the die temperature shows pronounced changes with dissipation. However, for the low frequency case, the additional 2°C on the die temperature is small compared to the overall change in case temperature. In addition, the short time constant of the die-to-case relationship means that for all but rapid changes to load, the additional temperature from dissipation at the die is proportional to power (that is, the die appears as a purely resistive element).

It may be noted that the die temperature appears to fall below the case temperature occasionally, an observation most clearly made in the high frequency case. In reality the die temperature cannot fall below the case temperature; this observation is the result of measurement error (either an underestimation of the die temperature or an overestimation of the case temperature). The error is less that 0.1°C, an uncertainty which may be tolerated in most practical systems.

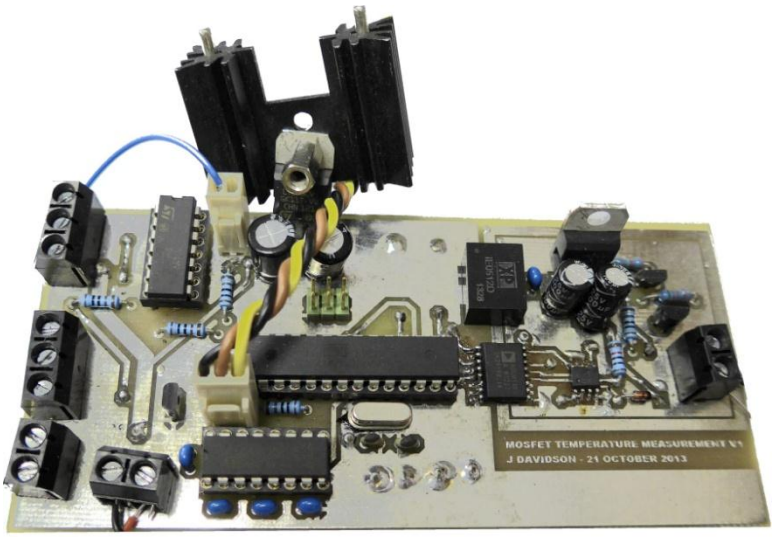

Fig. 3.18 Practical implementation of the die temperature measurement and dissipation controller

The measurement technique proposed in this section can perform die temperature measurement whilst almost simultaneously controlling power dissipation in a MOSFET, and has a number of applications. The practical implementation of the system is shown in Fig. 3.18. By operating as an arbitrary waveform power controller more than 80% of the time, then switching to a mode where the intrinsic diode is biased and its voltage

drop is measured, the die temperature can be calculated while the device mimics steady-state operating condition. Using this information, thermal impedance characteristics can be derived at the semiconductor die of a device, allowing the estimation and control techniques presented in this thesis to be applied with reference to die temperature, rather than purely case temperature.

3.4 Chapter conclusions

This Chapter has introduced a number of techniques that will be used in later chapters to control experiments and measure the thermal characteristics of power electronic devices. An arbitrary waveform power controller is proposed which is used to control the power dissipation in semiconductor devices. Power dissipation waveforms can therefore be applied to mimic normal operation or as test signals to evaluate and characterise devices, systems and heatsink arrangements. Temperature measurements are required to perform these characterisations, and the practical implementation of thermocouple-based measurements is described, along with details of typical experimental arrangements. These techniques are used throughout the thesis for characterisation of devices and evaluation of techniques. Finally, a circuit for the measurement of MOSFET die temperatures is proposed, which offers near-simultaneous controlled power dissipation in the same device. This circuit allows characterisation of devices from the semiconductor die, allowing for thermal management of devices closer to their operating limits.

3.5 References

[3.1] J. W. Kolar, U. Drofenik, J. Biela, M. L. Heldwein, H. Ertl, T. Friedli and S. D. Round, '*PWM converter power density barriers*', *in Proc* Power Conversion Conference, Nagoya, 2007, pp. 1656-1676.

[3.2] M. Musallam, P. P. Acarnley, C. M. Johnson, L. Pritchard and V. Pickert, '*Estimation and control of power electronic device temperature during operation with variable conducting current*', IET Circuits, Devices & Systems, vol. 1, pp. 111-116, 2007.

[3.3] J. D. Holmes, M. P. Foster and D. A. Stone, '*System-wide Temperature Estimation for IMS based Power Electronics Circuits*', *in Proc* International Conference on Power Electronics and Drive Systems, 2009, pp. 1081-1084.

[3.4] Z. Jakopovic, '*Computer controlled measurement of power MOSFET transient thermal response*', *in Proc* Power Electronics Specialists Conference (PESC), 1992, pp. 1025-1032 vol.2.

[3.5] A. S. Sedra and K. C. Smith, *Microelectronic Circuits*, 5th ed. Oxford: Oxford University Press, 2004.

[3.6] Omega Engineering, *Comparison of time constant vs. overall outside diameter of bare thermocouple wires or grounded junction thermocouples in air*. Stamford, CT, USA, 2014.

[3.7] LabJack Corporation. (2014). *LabJack U6 User's Guide*. Available: labjack.com/support/u6/users-guide/appendix-b

[3.8] J. Allison, *Electronic engineering semiconductors and devices*. London: McGraw-Hill, 1990.

[3.9] Fairchild Semiconductor, *MUR840, MUR860, RURP840, RURP860 Datasheet*, 2013.

[3.10] B. E. Boser and B. A. Wooley, '*The design of sigma-delta modulation analog-to-digital converters*', IEEE Journal of Solid-State Circuits, vol. 23, pp. 1298-1308, 1988.

Chapter IV

Techniques for the thermal impedance spectroscopy of power electronic systems

In this Chapter, techniques are presented for thermal impedance spectroscopy. Initially, pure sine wave spectroscopy, a technique commonly used for the analysis of electroceramics, is reported. By applying a sinusoidal input and measuring the gain and phase of the resulting output, the impedance of a system can be determined at a single frequency. The process is repeated for several frequencies over the frequency band of interest. Since this is an arduous process, with separate measurement and lead-in time required for each reading, an alternative technique is proposed which takes several readings but requires only one measurement and lead-in time. This pseudorandom binary sequence (PRBS) technique is explored in terms of signal generation, spectral content and practical application. Finally, noise and computational intensity considerations are taken into account. An expression is derived for the minimum gain that can be determined using a particular PRBS. The reader can therefore select an appropriate technique to characterise a system accurately and rapidly with noise resilience and with computational efficiency.

4.1 Introduction to impedance spectroscopy

Impedance spectroscopy is a technique widely used for a number of applications, including electrochemical analysis [4.1], battery characterisation [4.2] and electrical analysis [4.3]. It is a form of frequency domain system identification whereby the impedance or gain and phase of a system is determined at individual frequencies over a bandwidth of interest. In this thesis, the same concepts are applied to thermal systems, and the results are used as an empirical system model for temperature estimation and control.

At its simplest, impedance spectroscopy involves the application of a sine wave of voltage or current to a system while the relative amplitude and phase of the resulting

current or voltage are measured. This is known as the pure sine wave method. Systems under test can be either one-port, where the voltage and current are measured at the same port and hence the self-impedance spectrum is calculated, or two-port, where voltage and current are measured at different ports and hence the transfer impedance, or *cross-coupling*, is measured. The operation of the pure sine wave method for these two system types are shown diagrammatically in Fig. 4.1.

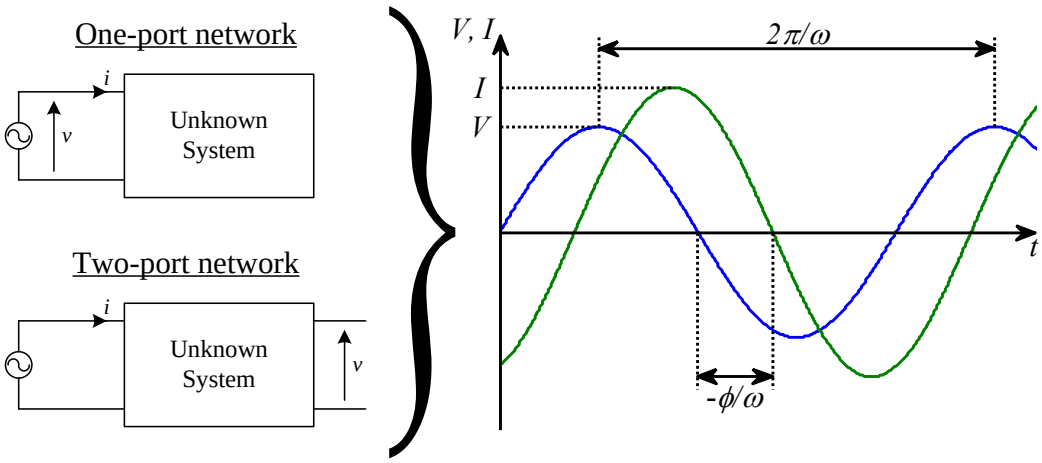

Fig. 4.1 Pure sine wave impedance spectroscopy of a system

In this example, a sinusoidal current is applied to one port and the corresponding voltage measured. The impedance is then calculated as the ratio between voltage and current, as shown below.

$$Z(\omega) = \frac{V(\omega)}{I} \angle \phi(\omega) \tag{4.1}$$

where V is the measured amplitude of the resulting voltage waveform following current excitation at angular frequency, ω, and ϕ is the phase shift. $Z(\omega)$ is the frequency-dependent impedance. To characterise the system, a sine wave of current is injected and the impedance is measured at a number of frequencies over a frequency band of interest. A typical characterisation consists of 5 to 10 measurements per decade of frequency.

The time taken for each measurement depends on the frequency being measured and the prevailing time constant of the system. After the sinusoid is applied to the system, the

output must be allowed to settle into its final sinusoidal output. Initially, the imposition of the sinusoid onto an unexcited system causes an impulse to be superimposed. This is demonstrated in Fig. 4.2, shows the initial response of a typical system.

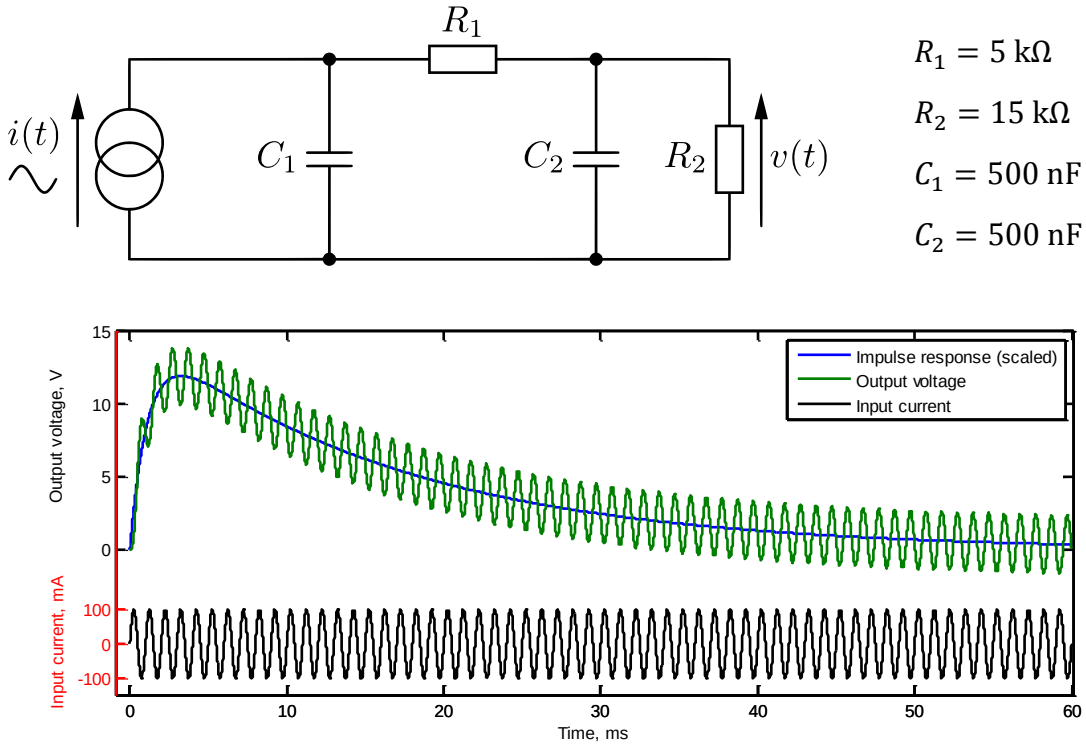

Fig. 4.2 Practical measurement and setting time

The cause of the initial transient can be explained using Laplace analysis [4.4]. For a capacitor-resistor ladder, such as that in Fig. 4.2, the impedance spectrum can be written in the Laplace domain as an equation of the form

$$Z(s) = \frac{A}{(s - p_1)(s - p_2) \cdots (s - p_3)} \tag{4.2}$$

Where A is a constant, $\boldsymbol{p}$ is a vector of the system poles and s is the Laplace operator variable. The response to a sine wave current excitation is found by multiplying $Z(s)$ by the Laplace transform of a sine wave, $I(s)$, as shown in (4.3).

$$\mathcal{L}\{v(t)\} = Z(s) \cdot \mathcal{L}\{i(t)\} = Z(s) \cdot \mathcal{L}\{I \sin \omega t\} \tag{4.3}$$

Where $\mathcal{L}$ is the Laplace transform, I is the amplitude of the current sine wave and ω is the angular frequency. $v(t)$ and $i(t)$ are the time (t) domain representations of voltage and current respectively. Substituting in $Z(s)$ and performing the Laplace transform yields

$$V(s) = Z(s) \cdot I(s) = \frac{A}{(s-p_1)(s-p_2)\cdots(s-p_3)} \cdot \frac{I\omega}{s^2+\omega^2} \tag{4.4}$$

This can be expressed in partial fraction form, noting that the $s^2+\omega^2$ term can be separated into $(s-\mathrm{j}\omega)(s+\mathrm{j}\omega)$, as shown in (4.5).

$$V(s) = \frac{c_1}{s-p_1} + \frac{c_2}{s-p_2} + \cdots + \frac{c_n}{s-p_3} + \frac{d_0}{s-\mathrm{j}\omega} + \frac{d_1}{s+\mathrm{j}\omega} \tag{4.5}$$

Where $\boldsymbol{c}$ is a vector of constants related to the partial fractions from $Z(s)$ and $\boldsymbol{d}$ is a pair of constants relating to the partial fractions from $I(s)$. Using the cover-up rule, d_0 and d_1 can be calculated, as given in (4.6) and (4.7).

$$d_0 = \left.\frac{V(s)}{s-\mathrm{j}\omega}\right|_{s=\mathrm{j}\omega} = \frac{Z(\mathrm{j}\omega)I}{2\mathrm{j}} \tag{4.6}$$

$$d_1 = \left.\frac{V(s)}{s+\mathrm{j}\omega}\right|_{s=-\mathrm{j}\omega} = -\frac{Z(-\mathrm{j}\omega)I}{2\mathrm{j}} \tag{4.7}$$

$Z(\mathrm{j}\omega)$ is a complex variable. For simplicity, it can be expressed in polar form as given in (4.8).

$$Z(\pm\mathrm{j}\omega) = |Z(\mathrm{j}\omega)| \exp(\pm\mathrm{j}\phi) \tag{4.8}$$

Substituting (4.6) to (4.8) into (4.5), and taking the inverse Laplace, yields

$$v(t) = x(t) + y(t) \tag{4.9}$$

where $x(t)$ and $y(t)$ are the functions given in (4.10) and (4.11) respectively.

$$x(t) = c_1 \exp(p_1 t) + c_2 \exp(p_2 t) + \cdots + c_n \exp(p_n t) \tag{4.10}$$

$$\begin{aligned} y(t) &= d_0 \exp(\mathrm{j}\omega t) + d_1 \exp(-\mathrm{j}\omega t) \\ &= \frac{|Z(\mathrm{j}\omega)| \exp(\mathrm{j}\phi)\, I}{2\mathrm{j}} \exp(\mathrm{j}\omega t) - \frac{|Z(-\mathrm{j}\omega)| \exp(-\mathrm{j}\phi)\, I}{2\mathrm{j}} \exp(-\mathrm{j}\omega t) \\ &= I|Z(\mathrm{j}\omega)| \frac{\exp(\mathrm{j}\omega t + \mathrm{j}\phi) - \exp(-\mathrm{j}\omega t - \mathrm{j}\phi)}{2\mathrm{j}} \\ &= I|Z(\mathrm{j}\omega)| \sin(\omega t - \phi) \end{aligned} \tag{4.11}$$

For stable systems, all poles, $\boldsymbol{p}$, have negative real parts and therefore $x(t)$ decays as $t \to \infty$, leaving the long term response as sinusoid $y(t)$. $x(t)$ is closely related to the impulse response of the system since the same equation (with different coefficients, $\boldsymbol{c}$) is obtained when $Z(s)$ is decomposed into partial fractions an transformed into the time domain. The response of the system to a sinusoid is therefore an amplified (or attenuated) phase-shifted sinusoid ($y(t)$) superimposed on a scaled impulse response-like perturbation ($x(t)$). This is the effect observed in Fig. 4.2. To accurately measure the phase and gain, it is therefore necessary to delay measurement until the impulse response of the system no longer has an effect, which is typically after a wait of 3-5 times greater than the dominant time constant. This requirement significantly increases measurement time for systems with long time constants.

4.1.1 Application of thermal systems

The demand for increased power density of power electronics in recent years is providing an ever greater impetus for the advanced thermal modelling techniques where real-time processing adapts the cooling behaviour to the changing thermal state of a system. In this case, the thermal modelling of the devices in a system is necessary to predict a thermal cycling failure, the likelihood of which can be reduced by integrated thermal management [4.5]. Such analyses commonly use electrical circuit analogues, where the thermal mass is modelled by a capacitor, C_{th}, and the thermal resistance is modelled by a resistor, R_{th}. A current source in such analogues represents the heat flux or generation, Q, and the voltages represent the temperature, T.

The pure sine wave spectroscopy measurement technique can be applied to a thermal system. In this case, the input current is equivalent to applied power dissipation and the output voltage is equivalent to measured temperature. A significant difference between electrical and thermal systems is that practical measurements performed on thermal systems can only apply positive power dissipation (that is, it is nontrivial to remove thermal energy controllably from a single point). Therefore pure sine wave thermal impedance spectroscopy must use an offset sinusoid as its input power waveform. This additionally adds an offset to the resultant temperature waveform; however, no additional settling time is required compared to the zero-centred waveform since the effect, akin to adding the system's step response to the output, requires the same settling time as the impulse response.

In the remainder of this Chapter, pure sine wave characterisation of thermal systems will be used to infer the level of complexity required of equivalent circuit models of thermal systems. Thermal systems typically have increased time constants compared to electrical systems, leading to significant measurement time requirements. Thus an alternative approach which can be used to simultaneously characterise the thermal system over a band of frequencies while also reducing experimental time will subsequently be described.

4.2 Order of the Cauer network required to model transfer impedance

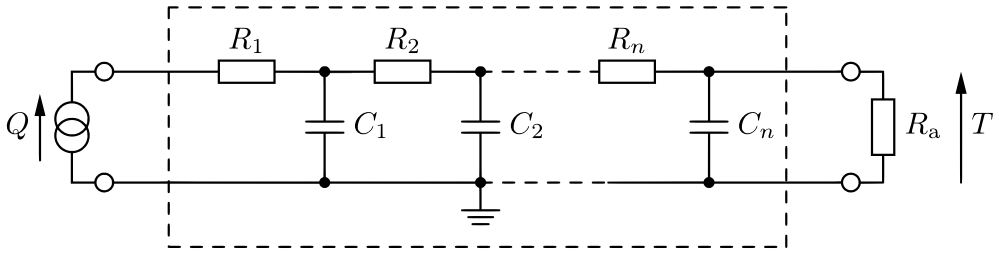

Fig. 4.3 Two-port Cauer network with heat source, Q, and ambient connection, R_a

Cauer- and Foster-based networks are typically used in modelling thermal systems as electrical analogues. In their one-port form, they are mathematically identical [4.6]; however, the Cauer network is a closer representation of the physical form of the thermal circuit because each node represents a real temperature. In the Foster network,

by contrast, only the leftmost node has physical significance. In addition, thermal cross-coupling, which is the transfer impedance between the two system components, is used to analyse the thermal response of the system with multiple heating elements, as will be seen in Chapter VI. In this case, a two-port network must be used and since the Foster network does not have an appropriate two-port form, a two-port Cauer network is used, such as the one shown in Fig. 4.3.

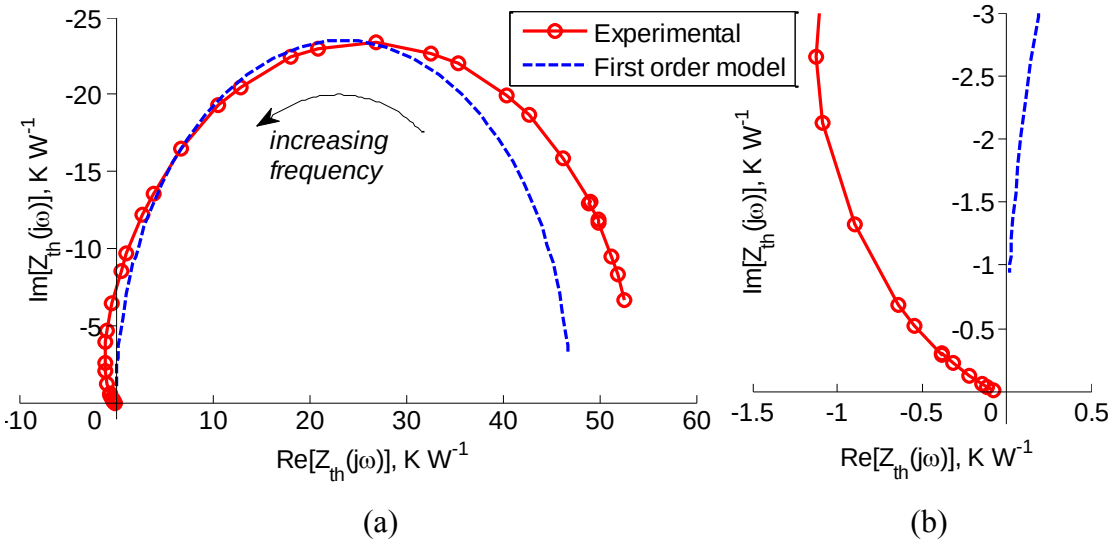

(a) (b)

Fig. 4.4 Nyquist plot of experimental transfer impedance spectrum for a TO220 device with fitted first order model. (a) Full plot. (b) Zoomed in around the origin.

In this section, pure sine wave thermal impedance spectroscopy is used to determine the appropriate order of the Cauer network required in a thermal analysis. Impedance spectroscopy is a technique for analysing the constituent components of ceramics by virtue of their differing electrical time constants [4.7]. The data are plotted as Nyquist diagrams and the characteristics are found by inspection. In recent years, this technique has been applied to thermal analysis, particularly for packaging and microelectronics [4.8][4.9]. To illustrate the limitations of the present thermal modelling, the transfer impedance between the active region of a TO220-packaged transistor and a thermocouple glued to its tab was measured experimentally by using impedance spectroscopy, as shown in Fig. 4.4. Typical thermal modelling would consider this as a first-order problem and therefore use a single element Cauer network. However, such a model occupies only a single quadrant in the Nyquist diagram, whereas the characteristic of the real device occupies multiple quadrants. For even the simplest of

thermal arrangements, the present models are inaccurate. An analysis of the network length required for an accurate thermal modelling is therefore presented.

4.2.1 Required network order

A device may be modelled as a series of adjoined discrete materials, each with a thermal mass and a thermal conductivity. A Cauer network model can be developed by taking these properties and assigning the corresponding capacitance, C_{th}, and resistance, R_{th}, to each Cauer element. In reality, R_{th} and C_{th} are the distributed components; however, to maintain ease of analysis they are lumped together so that their effects can be modelled as discrete components. To build a model of a complete thermal system, the system is first divided into single material blocks each of which is assigned a thermal resistance and a capacitance. The complete system model is formed by interconnecting these blocks and inserting the heat sources and the ambient connections wherever appropriate. Typically, each single material is modelled as a single $R_{\mathrm{th}}C_{\mathrm{th}}$ element. However, to approximate the distributed components, the bulk $R_{\mathrm{th}}C_{\mathrm{th}}$ element for each material block can be divided into n smaller $R_{x(\mathrm{th})}C_{x(\mathrm{th})}$ elements according to (4.12), which results in an nth-order Cauer network for each block.

$$R_{x(\mathrm{th})} = \frac{R_{\mathrm{th}}}{n}; \qquad C_{x(\mathrm{th})} = \frac{C_{\mathrm{th}}}{n} \tag{4.12}$$

Although the additional elements improve the approximation to a distributed component, they also add to the computational complexity of the model. The order of a Cauer network model is the number of RC elements, n, from which it is formed. To find the optimal order for frequency response modelling, the impedance spectra of a single block of material modelled as a Cauer network of varying order is considered. A block defined only by its thermal properties, in this case R_{th} = 10 K/W and C_{th} = 0.1 Ws/K, is modelled as an nth-order Cauer network with an ambient connection of resistance $R_{\mathrm{a(th)}}$= 10 K/W to a fixed temperature. Nonlinear heat flows, such as convection and radiation, are neglected in this section. The thermal impedance spectrum, $Z_{\mathrm{th}}(\mathrm{j}\omega)$, is calculated according to (4.13).

$$Z_{\mathrm{th}}(\mathrm{j}\omega) = \frac{T(\omega)}{Q}\angle\phi(\omega) \tag{4.13}$$

where j is the imaginary unit, ω is the angular frequency, Q is the heat flux into the material, T is the temperature at the boundary with the ambient and ϕ is the phase difference between T and Q.

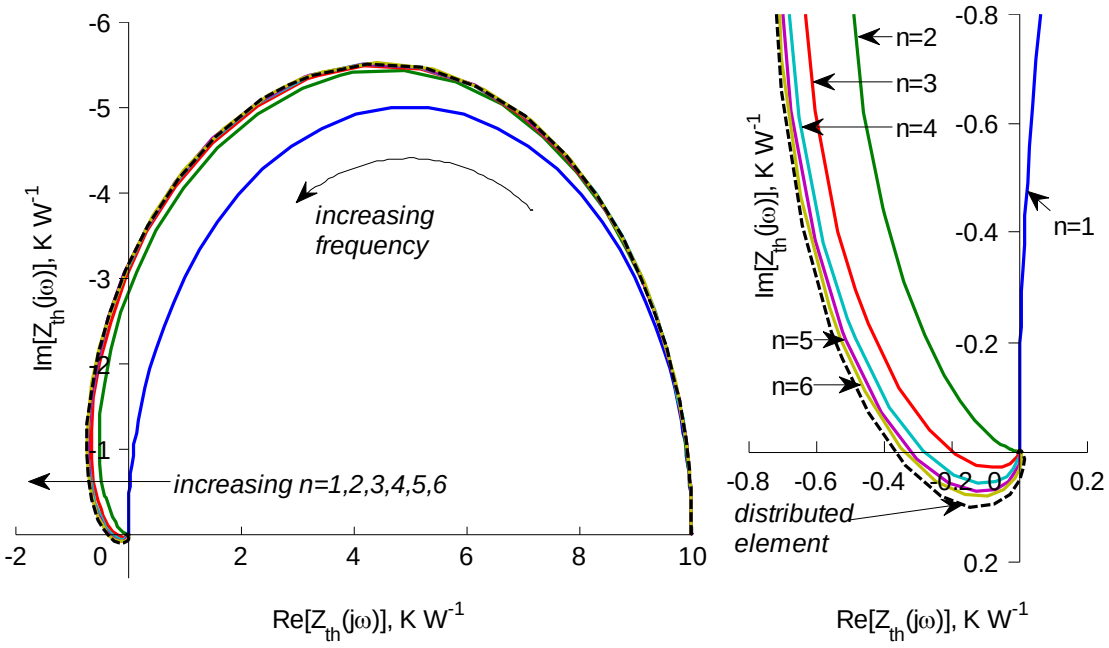

Fig. 4.5 Nyquist plot of a single material Cauer network approximation of different orders, n. (a) Full plot. (b) Zoomed in around the origin.

Fig. 4.5 shows the transfer impedance Nyquist plots of the material for different Cauer network orders, n. As n increases, the curves converge as they become more representative of a distributed element. It is interesting to note that the shapes of the higher-order curves enter all the quadrants of the Nyquist diagram at a high frequency, unlike those commonly seen for the electroceramics impedance spectra (such as those in [4.7]). This can be explained by the frequency-domain flux-to-temperature transfer function, Z_{th}, of the two-port Cauer network, which, for an nth-order network, takes the form below. This equation can be derived by using the leapfrog ladder technique noting that the input is a current source [4.10].

$$Z_{\mathrm{th}}(\mathrm{j}\omega) = \frac{1}{A_n(\mathrm{j}\omega)^n + A_{n-1}(\mathrm{j}\omega)^{n-1} + \cdots + A_1(\mathrm{j}\omega)^1 + A_0} \tag{4.14}$$

where $\boldsymbol{A}$ is a set of constants. Since each of $(\mathrm{j}\omega)^1$ to $(\mathrm{j}\omega)^4$ occupies a different quadrant in the Nyquist plot, a curve in all the quadrants can be approximated by using a Cauer network of $n \geq 4$. A model consisting of four elements is therefore the minimum complexity required to accurately model a single material because $n = 4$ is the smallest order where all quadrants of the Nyquist plot are entered. The $n = 1$ curve is especially inaccurate because it represents a first-order network and therefore occupies only one quadrant. It shows a significant difference from the distributed element over all the frequencies. This demonstrates the limitations inherent in modelling a single material block or device by using a single element network.

4.2.2 Measurement position

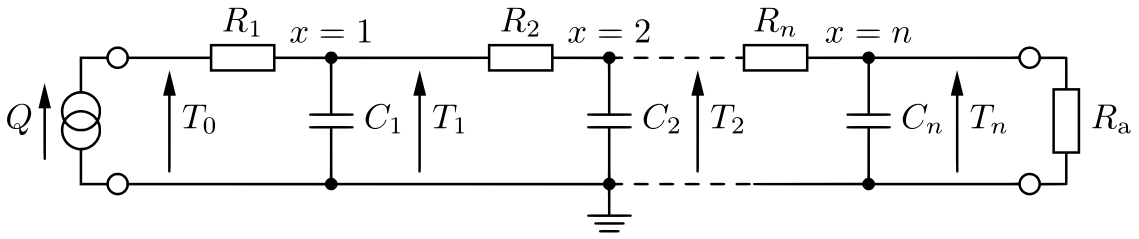

Fig. 4.6 Cauer two-port network with temperature measurements taken at various points along the material

For a high-order Cauer network, the position in the network at which the temperature is measured also has an effect. A two-port Cauer network with several measurement points is shown in Fig. 4.6. By plotting the Nyquist curves of the transfer impedance between the heat flux source, Q, and each temperature, T, the effect of the several adjoined $R_{x(\mathrm{th})}C_{x(\mathrm{th})}$ elements can be analysed, as shown in Fig. 4.7.

At the heat source, the curve is of the first order and therefore occupies only a single quadrant. As the measurement is taken further away from the heat source, the curves take a form similar to the distributed element seen in Fig. 4.5. With each additional element between the source and the measurement, the curves occupy an extra quadrant of the Nyquist diagram. Therefore, for an accurate approximation of the transfer

impedance partway through a material, at least four Cauer elements are required between the heat flux source and the temperature measurement.

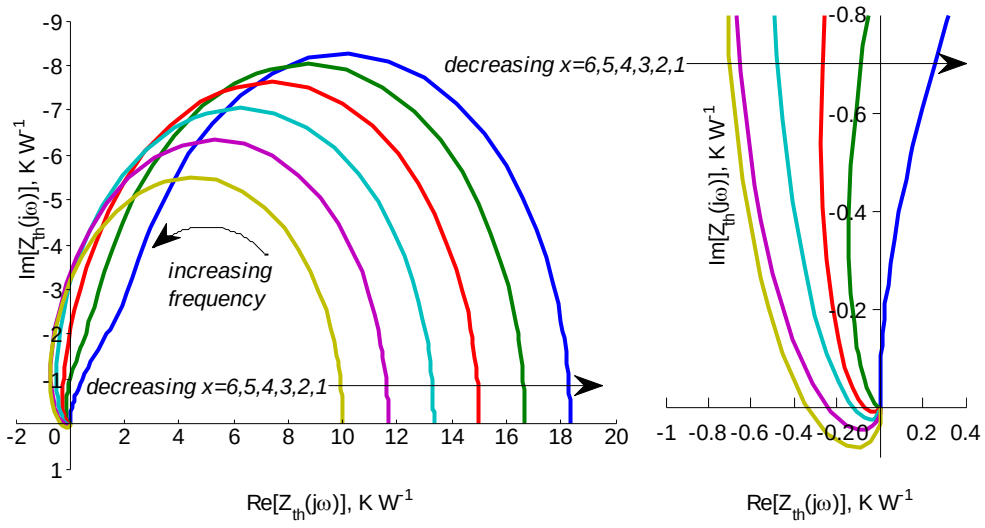

(a) (b)

Fig. 4.7 Nyquist plot of a single material Cauer network approximation with transfer impedance taken at different nodes, x. (a) Full plot. (b) Zoomed in around the origin.

It has been seen that it is important to model a thermal circuit with a sufficient order to capture the behaviour in all the quadrants, particularly in the situations where the high-frequency response is important. One such situation is the thermal management of the power devices under a rapidly changing load, for example, in high pulsed power applications, where the resulting large changes in temperature cause a wire bond failure and shorten the device's lifespan [4.5]. In this case, the first-order model does not represent a good approximation of a distributed element at a high frequency although it shows a better agreement at a low frequency. This is because, unlike the distributed element result, it cannot cross the imaginary axis and therefore represents an inaccurate model for the applications with rapid power transients.

4.2.3 Applicability to thick material layers

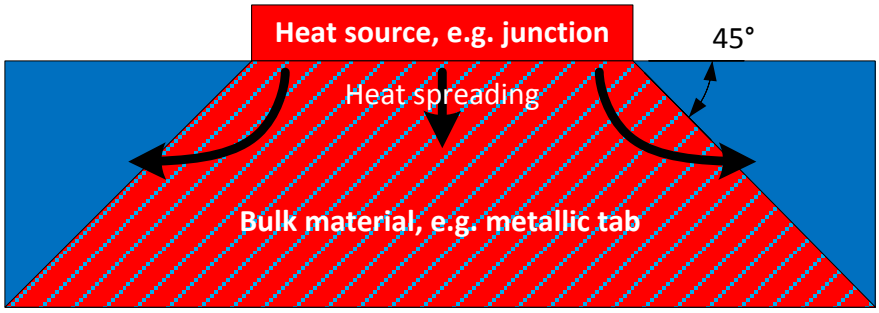

Fig. 4.8 Heat spreading in bulk material

In practical devices, a lower material layer may be larger in cross-sectional area than the preceding layer as shown in Fig. 4.8. For thicker materials, this means that the effective thermal resistance per unit length through the material may decrease and the effective thermal capacity per unit length through the material may increase due to heat spreading. Heat spreading is the effect where heat flux entering a bulk material from a small heat source will typically flow outwards at a angle of typically 45° to the horizontal [4.11]. This effect is significant for thicker bulk materials where the contact area with a heat source is small compared to total area. In a packaged device, this may occur between the junction and the bulky metallic tab. In these cases, it is therefore invalid to assume uniform thermal resistance and capacitance throughout the single material layer; instead, the values must be adjusted to take account of the effect.

4.2.4 Summary of Cauer network order analysis

Thermal impedance spectroscopy is used to assess the required number of the Cauer network elements to model a single material block. Although such blocks are commonly modelled as the first-order Cauer networks, their thermal transfer impedance occupies only a single quadrant of a Nyquist diagram in contrast to the experimental measurements. A Cauer network consisting of at least four elements between the heat flux input and the temperature measurement is therefore required for the model to correctly occupy all of the quadrants of the Nyquist diagram at a high frequency. The results demonstrate that the Cauer networks with orders higher than four do not substantially improve the approximation to the distributed components, indicating that the four elements are a sensible compromise between accuracy and complexity.

4.3 Generation of thermal impedance spectra using pseudorandom binary sequences

Pure sine wave spectroscopy is an arduous process since the impedance and phase at each frequency must be measured and evaluated separately. The process is particularly time-consuming for thermal systems where time constants are long and consequently the frequency band of interest may be in the microhertz range. Even neglecting settling time requirements, 6 logarithmically spaced readings between 100 μHz and 1 mHz inclusive require over 7 hours. If a single reading could replace that range, the time would be reduced to less than 3 hours. When settling time considerations are taken into account, the difference is more stark, since the pure sine wave method requires six separate settling periods as opposed to the one required for a single reading.

Using a step-response to characterise the thermal system is a possible alternative. In this case, a step change in power is applied to the input and the resulting frequency response measured. However, the accuracy limitations discussed in Chapter II limit its usefulness, especially in systems with repeated rapid demand changes. If a modest increase in complexity would reduce this limitation without increasing measurement time, that would be a more accurate alternative. Therefore, in this section, the use of a pseudorandom binary sequence-based technique is introduced. Pseudorandom binary sequences (PRBS) can be used to excite a system so that the impedance response can be determined. PRBS techniques represent a modest increase in complexity over step-response testing. System identification using PRBS techniques has been in use for many decades. These techniques have improved spectral content at higher frequency due to their many switching events and so find other applications in the generation of band-limited white noise. PRBSs resemble a random sequence of ‘0’s and ‘1’s but are finite-length and predetermined [4.12]. They are produced using linear-shift feedback registers of length n by taking the exclusive-NOR from several taps. In this thesis, only maximum-length PRBSs are considered; these are sequences of length $N = 2^n - 1$ where every state of the feedback register, except the trivial all-‘1’s state, is used in a sequence.

The transfer function of a system can be calculated over a range of frequencies by applying a PRBS input and measuring the resulting output. In thermal systems, the

transfer function desired is typically the self impedance at one point (where the transfer function is between power and temperature at that point) or the cross-coupling between two points (where the transfer function is between power input at one point and temperature output at the other). To analyse thermal systems, an electrical analogue is again used, where temperature is akin to voltage and power is akin to current. In this way a thermal system can be considered using the familiar terms used for an electrical system. In the thermal system, power (which in practice must always be positive since it is controlled by power dissipation) is applied and the resulting temperature variation is measured. The complex thermal impedance, Z_{th}, is the transfer function relating these two. It is found by calculating the quotient of the power and temperature in the frequency domain as shown in (4.15).

$$Z_{\mathrm{th}}(\omega) = \frac{T(\omega)}{Q(\omega)} \tag{4.15}$$

where T is the output temperature, Q is the input power and ω is the angular frequency.

4.3.1 Production of PRBS

PRBS are generated by linear-shift feedback registers of length n by taking the exclusive-NOR from several taps. The shift register necessary for a 4-bit PRBS is shown in Fig. 4.9.

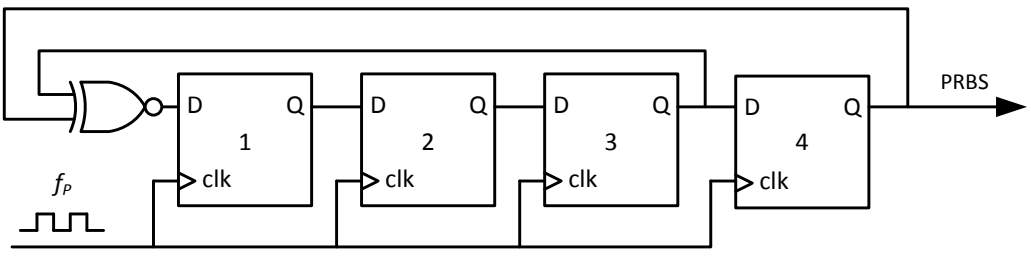

Fig. 4.9 Implementation of a 4-bit PRBS using a linear-shift feedback register

The length of the PRBS required is determined by a combination of the desired bandwidth of the signal, the acceptable computational complexity of the analysis and the time available for the experiment. The bandwidth and consequently the time

required for an experiment increase with PRBS-length as shown in (4.18) below. The appearance of a 4-bit PRBS in the time domain is shown in Fig. 4.10.

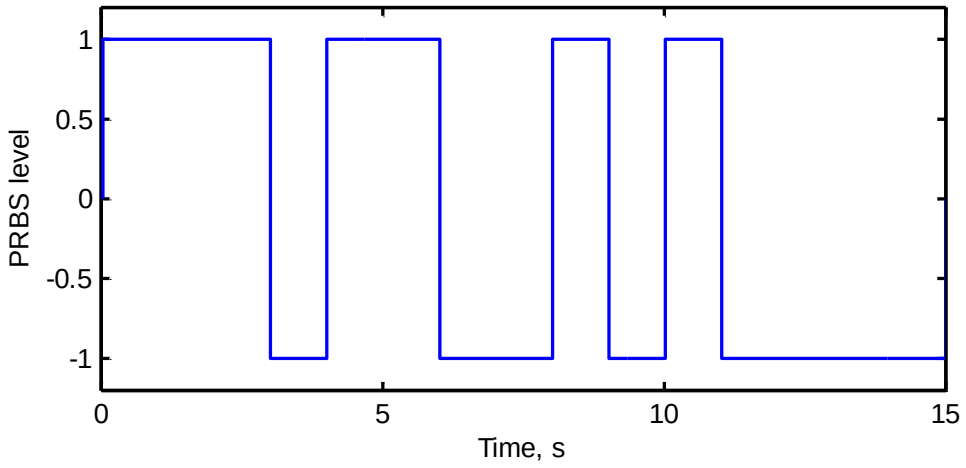

Fig. 4.10 Time domain representation on the 4-bit PRBS realised as a bipolar signal

4.3.2 Frequency domain analysis of PRBS

Conversion of time-domain signals to frequency-domain representations is most efficiently performed using the discrete Fourier transform (DFT) for which there exist several computationally efficient algorithms. Algorithms requiring of the order of $n_s \log_2 n_s$ complex calculations for an n_s-length sequence are available [4.13]. The equation of a PRBS' DFT can be calculated from its power density spectrum, Φ_{xx}, which is shown by Davies [4.12] and Hampton [4.14] to be a line spectrum of discrete points bounded by (4.16). The discrete points on the power spectrum are separated by f_P/N. Fig. 4.11 shows the power spectral density plot of the four bit PRBS from Fig. 4.9.

$$\Phi_{xx}(f) = \frac{a^2}{f_P}\frac{(N+1)}{N}\left[\frac{\sin(f\pi/f_P)}{f\pi/f_P}\right]^2 \tag{4.16}$$

Where the PRBS is a bipolar signal and shifts between $+a$ and $-a$ watts, the sequence is generated by n shift registers giving it a length $N = 2^n - 1$, and is controlled by a clock at frequency f_P Hz.

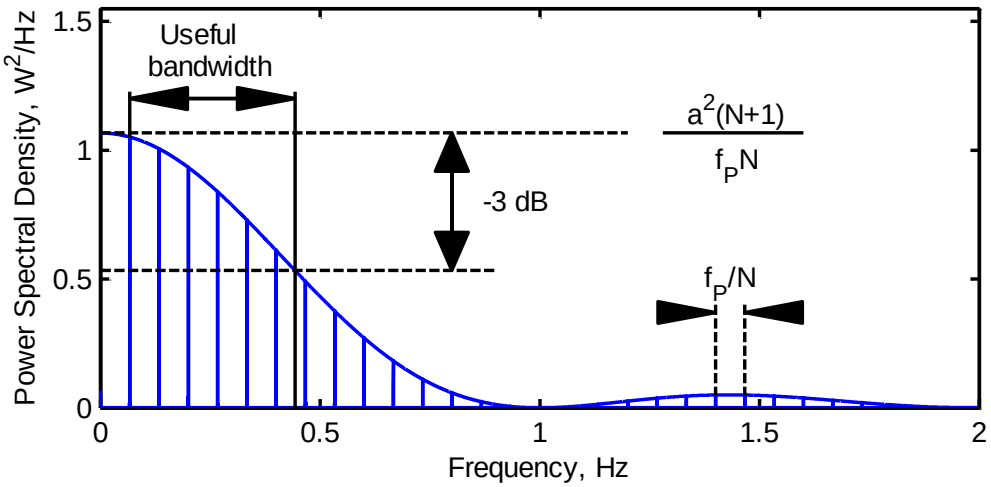

Fig. 4.11 Power spectral density of the 4-bit PRBS generated by the shift register above

The band-limit of a PRBS is the frequency at which its power drops by half. This occurs when

$$\left[\frac{\sin(f\pi/f_{\mathrm{P}})}{f\pi/f_{\mathrm{P}}}\right]^2 = 0.5 \quad \Rightarrow \quad f \approx \frac{f_{\mathrm{P}}}{2.3} \tag{4.17}$$

The discrete points on the power spectrum are separated by f_{P}/N. From (4.17) it can be seen that the frequencies over which the PRBS is valid are those in equation (4.18) for positive integers k.

$$f = \frac{kf_{\mathrm{P}}}{N}; \quad \frac{f_{\mathrm{P}}}{N} \leq f \leq \frac{f_{\mathrm{P}}}{2.3} \tag{4.18}$$

Some existing literature has used $f_{\mathrm{P}}/3$ as the band limit [4.12]. In this thesis, $f_{\mathrm{P}}/2.3$ is used, which is in accordance with (4.17) and allows analysis in a wider band. For simplicity, discrete analysis is usually performed, typically discrete Fourier analysis using the DFT. The effect of signal length and sampling rate on the result of a DFT must therefore be accounted for.

Since the Fourier transform of the signal's autocorrelation function is equal to both the squared modulus of the signal's Fourier transform [4.15] and the power spectral density

of the signal [4.12], the Fourier transform of a PRBS can be derived. This is expressed as a DFT in (4.19).

$$\Phi_{\mathrm{xx}}[k] = \frac{|\mathcal{F}(p[i])|^2}{n_{\mathrm{s}} f_{\mathrm{s}}}; \quad f = \frac{k f_{\mathrm{s}}}{n_{\mathrm{s}}}; \quad \forall \quad 0 \leq i, k \leq n_{\mathrm{s}} - 1 \tag{4.19}$$

where n_s is the number of samples, p is the PRBS function, f_{s} is the sampling rate and $\mathcal{F}$ is the DFT. The divisor $n_{\mathrm{s}} f_{\mathrm{s}}$ is a scaling factor to account for the discrete nature of the DFT. Substituting (4.16) into (4.19) yields the equation of the DFT: a line spectrum bounded by a *sinc* function expressed below (since sequence duration $N/f_{\mathrm{p}} = n_{\mathrm{s}}/f_{\mathrm{s}}$)

$$|\mathcal{F}(p[i])| = a\sqrt{N+1}\frac{f_{\mathrm{s}}}{f_{\mathrm{P}}} \cdot \frac{\sin(f\pi/f_{\mathrm{P}})}{f\pi/f_{\mathrm{P}}} \tag{4.20}$$

Fig. 4.12 shows the DFT of a 4-bit ($N = 15$) 1 Hz sequence. The use of a logarithmic scale is typical in the representation of Bode plots, which are an easy to understand representation of transfer functions. The Bode plot of the PRBS shows a smooth transition from the useful band at low frequency to the high frequency region where identification cannot take place.

The DFT has a near-flat spectrum below the band limit, whose height can be found from (4.20) as $f \to 0$. This value is shown on the Bode plot as F_0 and is given in (4.21).

$$F_0 = |\mathcal{F}(p[0])| = a\sqrt{N+1}\frac{f_{\mathrm{s}}}{f_{\mathrm{P}}} \tag{4.21}$$

The PRBS is almost balanced. It has one more low state than high states when generated using a linear shift feedback register with an exclusive-*NOR* gate such as in Fig. 4.9 above. If a precisely balanced signal is desired, an extra low output can be added to the sequence. When an extra '0' is added to the end of the sequence, it results in the DFT given in Fig. 4.13 below.

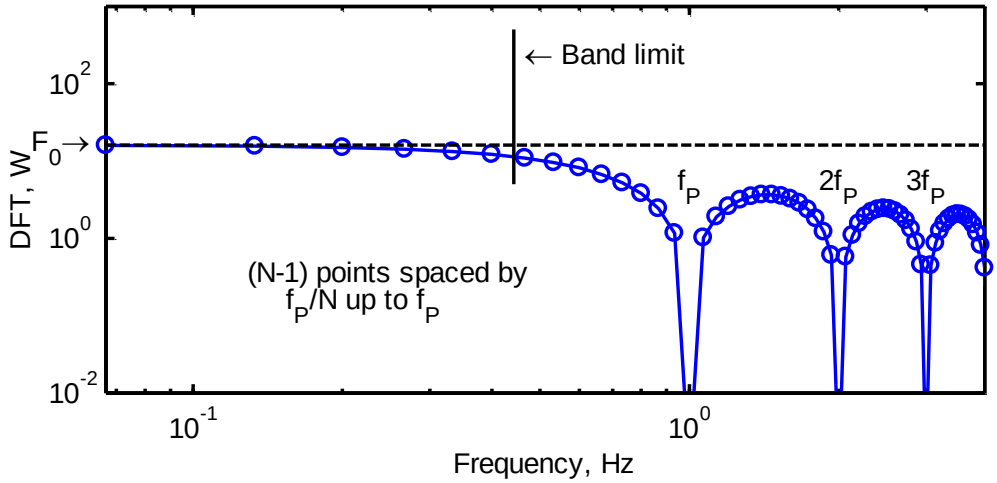

Fig. 4.12 Discrete Fourier transform of a 4-bit 1 Hz PRBS

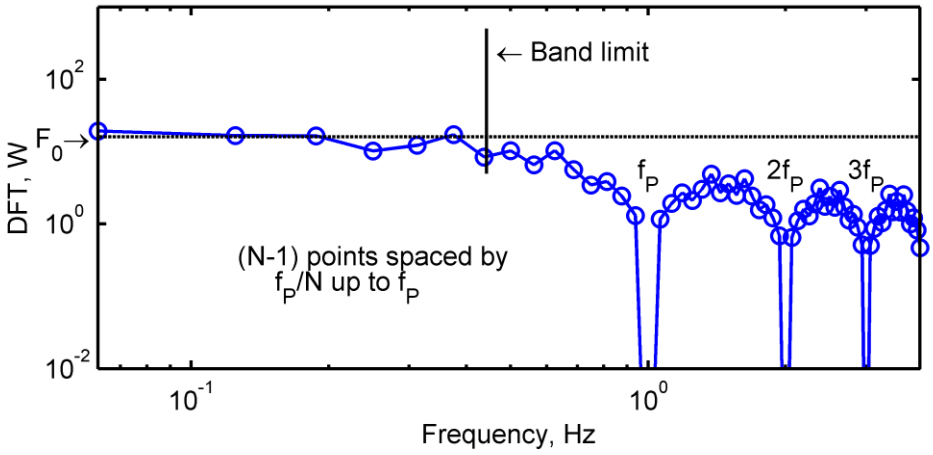

Fig. 4.13 Discrete Fourier transform of the 4-bit 1 Hz PRBS with a balancing '0'

The Bode plot of the extended PRBS has a similar but not identical envelope to the equivalent standard PRBS. It has a similar bandwidth and can therefore also be used for system identification. However, its distorted spectral content across frequencies means that some higher frequencies have more content than lower frequencies. Thermal systems are most sensitive to lower frequency signals and there is no compelling reason to balance the PRBS, therefore this extended signal is not used.

For either PRBS type, the thermal impedance of the system is calculated using (4.22) by dividing the DFT of output temperature by the DFT of input power across the frequency domain. This approach is explored fully in [4.13], where techniques to improve data quality such as windowing and spectral smoothing are also described.

$$Z_{\mathrm{th}}[k] = \frac{\mathcal{F}(T[i])\,[k]}{\mathcal{F}(Q[i])[k]}; \quad f = k\frac{f_{\mathrm{s}}}{n_{\mathrm{s}}}; \quad t = \frac{i}{f_{\mathrm{s}}} \quad \forall \quad 0 \leq k, i \leq n_{\mathrm{s}} - 1 \qquad (4.22)$$

where Z_{th} is the complex thermal impedance, $\mathcal{F}$ is the discrete Fourier transform, i is the time-domain sample index, k is the frequency-domain sample index, n_{s} is the number of samples, f_{s} is the sampling frequency, f is frequency and t is time. This technique can be therefore be used as an alternative to pure sine wave spectroscopy.

4.4 Minimum gain identifiable when PRBS techniques are used for system identification in noisy conditions

4.4.1 Introduction

The transfer function of a system, H, is most simply identified by taking the quotient of the Fourier transforms of its output and input when excited by PRBS [4.13]. In practical systems, noise will be present in the measured output voltage and this is especially problematic in systems with low-gain transfer functions where the input signal is of limited amplitude, as the signal-to-noise ratio at frequencies of interest is insufficient for accurate system identification. One example of this is thermal systems whose gains fall off rapidly at modest frequencies.

Existing literature reports methods of reducing noise; for example, using cross-correlation [4.16]. In this section, however, the minimum gain identifiable is calculated in a system with additive white Gaussian noise (AWGN) present in its output, allowing an engineer to discount noise-affected frequencies in the identified response and determine whether noise-reduction techniques are required.

4.4.2 Theoretical

For an unknown system, H, with AWGN present in the output, the experimental process to identify H can be modelled as in Fig. 4.14. In a noiseless system, $H(\omega)$ can be exactly identified over the useful frequency band of the PRBS. However, when noise, v_{n}, is present, only an approximation to $H(\omega)$ is found, which is denoted $\widetilde{H}(\omega)$. Although in reality the noise is present before the DFT is taken, it is shown being added afterwards in Fig. 4.14. The linearity of the DFT means this has no effect on the result,

but the mathematics is simplified. The transfer function calculated is therefore the sum of the actual transfer function plus a noise-dependent term. This is shown for angular frequency, ω, in (4.23).

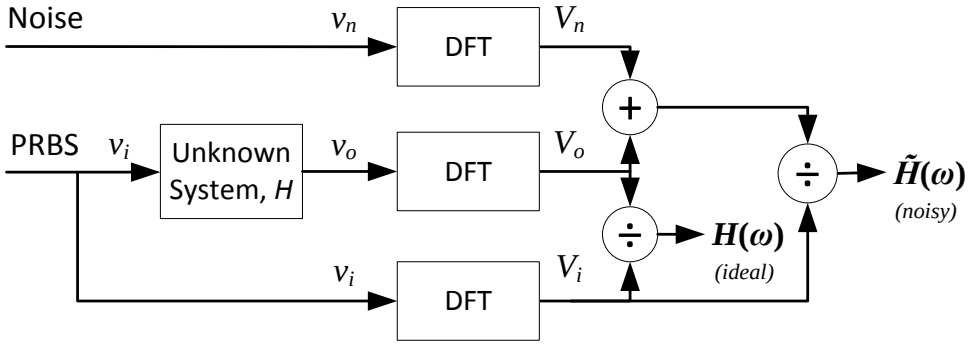

Fig. 4.14 Process for system identification using PRBS showing noiseless output, $H(\omega)$, and noisy output, $\widetilde{H}(\omega)$

$$\widetilde{H}(\omega) = \frac{V_{\mathrm{o}}(\omega) + V_{\mathrm{n}}(\omega)}{V_{\mathrm{i}}(\omega)} = H(\omega) + \frac{V_{\mathrm{n}}(\omega)}{V_{\mathrm{i}}(\omega)} \tag{4.23}$$

where V_{i}, V_{n} and V_{o} are the DFTs of the input, noise and ideal output signals respectively. It was shown in (4.16) that the power spectral density (PSD) of a PRBS is as given (4.24). Over the useful frequency range, the spectrum is approximately flat with height derived from (2) when $f \to 0$.

$$\Phi_{\mathrm{xx}}(f) = \frac{V^2}{f_{\mathrm{P}}} \frac{N+1}{N} \left[\frac{\sin(f\pi/f_{\mathrm{P}})}{f\pi/f_{\mathrm{P}}}\right]^2 \tag{4.24}$$

where $\Phi_{\mathrm{xx}}(f)$ is the PSD at frequency f of a bipolar PRBS, amplitude V, length $N = 2^n - 1$, where n is the shift-register length, clocked at frequency f_{P}. To calculate the DFT, it must be noted that the PSD of a PRBS is the Fourier transform of its autocorrelation function [4.14] and also that the Fourier transform of its autocorrelation function is the squared modulus of its Fourier transform [4.15]. Converting the Fourier transform to the DFT (and scaling accordingly), the PSD can therefore be written

$$\Phi_{\mathrm{xx}}[k] = \frac{|V_{\mathrm{i}}|^2}{n_{\mathrm{s}} f_{\mathrm{s}}} \tag{4.25}$$

where V_i is the DFT of the PRBS, f_s is the sampling rate, n_s is the total number of samples and k is the discrete frequency index where $f = kf_{\mathrm{P}}/N$. The $n_{\mathrm{s}} f_{\mathrm{s}}$ divisor accounts for the scaling difference between the continuous Fourier transform and the DFT. Substituting (4.24) into (4.25) when $f \to 0$ (and noting that sequence duration $N/f_{\mathrm{P}} = n_{\mathrm{s}}/f_{\mathrm{s}}$ for a PRBS) gives an approximation of the DFT of a PRBS over its useful range, given in (4.26).

$$|V_i| \cong V \frac{f_{\mathrm{s}}}{f_{\mathrm{P}}} \sqrt{N+1} \tag{4.26}$$

The DFT of AWGN, V_{n}, consists of real and imaginary components which are independently and normally distributed [4.17]. $|V_{\mathrm{n}}|$ is therefore Rayleigh distributed [4.18]. From Parseval's relation [4.15] it is possible to state the following.

$$\frac{1}{n_{\mathrm{s}}} \sum_{i=0}^{n_{\mathrm{s}}-1} |V_{\mathrm{n}}[i]|^2 = n_{\mathrm{s}} \overline{v_{\mathrm{n}}^2} \tag{4.27}$$

where $\overline{v_{\mathrm{n}}^2}$ is the noise power, a measurement of the noise present in the signal. $|V_{\mathrm{n}}|$ takes the form of a mean value with random perturbations. The mean value is the mean of the corresponding Rayleigh distribution and the amplitude of the random perturbations is a function of the standard deviation. The mean and standard deviation are derived in literature [4.18] and are given in (4.28) and (4.29) respectively. Using (4.27), they are also expressed in terms of noise power, $\overline{v_{\mathrm{n}}^2}$.

$$\overline{|V_{\mathrm{n}}|} = \sqrt{\frac{\pi}{2} \frac{1}{2n_{\mathrm{s}}} \sum_{i=0}^{n_{\mathrm{s}}-1} |V_{\mathrm{n}}[i]|^2} = \sqrt{\frac{\pi}{4} n_{\mathrm{s}} \overline{v_{\mathrm{n}}^2}} \tag{4.28}$$

$$\text{st. dev.}\,|V_\text{n}| = \sqrt{\frac{4-\pi}{2}\frac{1}{2n_\text{s}}\sum_{i=0}^{n_\text{s}-1}|V_\text{n}[i]|^2} = \sqrt{\left(1-\frac{\pi}{4}\right)n_\text{s}\overline{v_\text{n}^2}} \tag{4.29}$$

$|V_\text{n}|$ consists of these two components: the mean value and random perturbations which are in effect bounded by a multiple of the standard deviation. From (4.23) and (4.24), the minimum gain identifiable in the presence of noise, H_min, can be found when

$$|H(\omega)| \approx \frac{\max|V_\text{n}|}{V(f_\text{s}/f_\text{P})\sqrt{N+1}} \tag{4.30}$$

The numerator is a combination of the mean of the noise DFT plus a multiple, Δ, of the standard deviation. To calculate H_min, (4.28) and (4.29) are substituted into (4.30) noting that the number of samples, n_s is Nf_s/f_P. This value is given in (4.31).

$$H_\text{min} = \frac{1}{2}\sqrt{\frac{\overline{v_\text{n}^2}}{V^2}\frac{f_\text{P}}{f_\text{s}}\frac{N}{N+1}}\left(\sqrt{\pi}+\Delta\sqrt{4-\pi}\right) \tag{4.31}$$

where Δ is the number of standard deviations of noise perturbations which is acceptable for the application. H_min at $\Delta = 0$ is refered to as noise floor because it is the gain in a bode plot of $\widetilde{H}$ around which the noise is centred.

4.4.3 Results

To illustrate the effect of noise, a 1 V 8-bit ($N = 255$) bipolar PRBS is applied to the system, H, in (4.32) under varying noise power conditions. This system is a three element 1 kΩ–1 F RC ladder. The PRBS clock frequency is 30 kHz and the output is sampled synchronously at 30 kHz.

$$H(s) = \frac{1}{(10^{-9})s^3 + (6\times10^{-6})s^2 + (5\times10^{-3})s + 1} \tag{4.32}$$

where s is the Laplace operator variable.

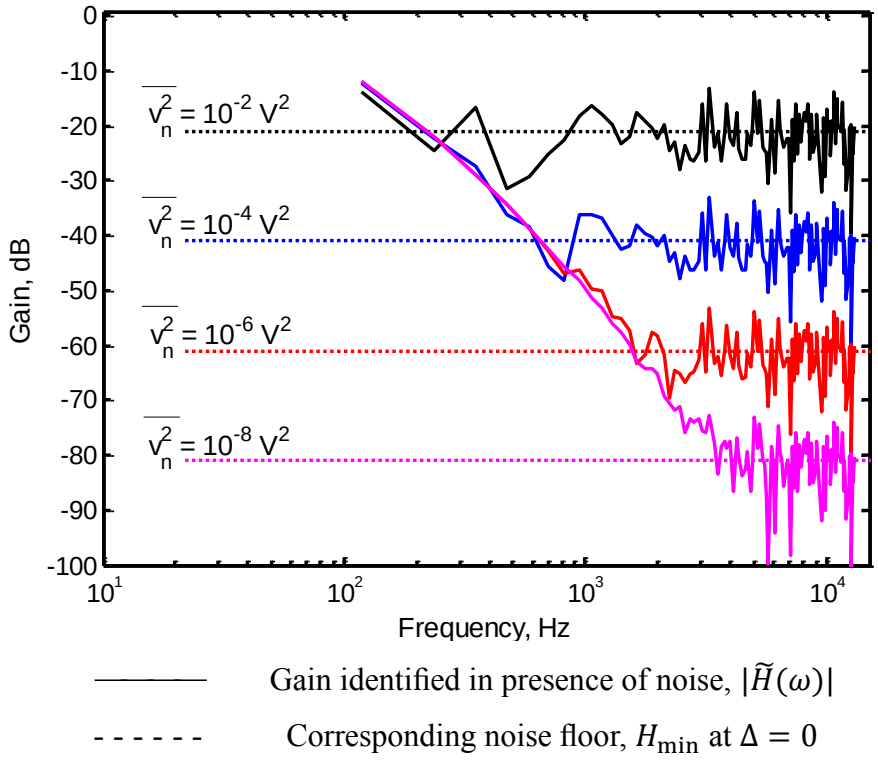

Fig. 4.15 Noise floor at varying noise power

Fig. 4.15 shows the relationship between the noise power, $\overline{v_n^2}$, and the bode plot of the identified transfer function, $\widetilde{H}(\omega)$. $\widetilde{H}(\omega)$ is presented for several noise floors calculated from (4.31) at $\Delta = 0$, each of which is indicated by a dashed line. There is good agreement between the noise floor and the gain at which noise dominates the identified transfer function. This confirms that for AWGN the gain where noise dominates is calculable from the Rayleigh distribution.

Fig. 4.16 shows the relationship between the number of standard deviations of noise, Δ, above the mean which may be discarded for an acceptably noise-free result. As Δ increases, the noise remaining in the signal decreases but so too does the useful bandwidth, thus a compromise is required. The level of noise acceptable is dependent on the application but, by inspection, truncating the signal when $|H|$ falls under the $\Delta = 2$ line removes the majority of the noise while leaving a large part of the transfer function intact.

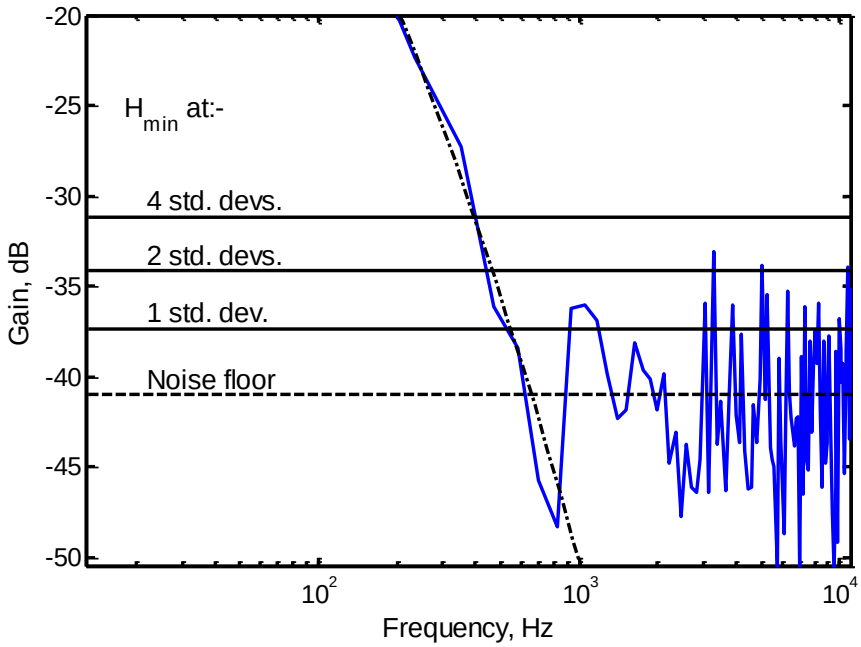

Fig. 4.16 Minimum gain identifiable at varying standard deviations relative to the noise floor for $\overline{v_{\mathrm{n}}^2} = 10^{-4}\ \mathrm{V}^2$

It is therefore possible to estimate the minimum gain of a system which can be identified using PRBS techniques. With this information, an engineer can recognise when the transfer function has become unacceptably affected by noise and hence confidently discard the affected portion of the result.

4.4.4 Discussion

The result in (4.31), reiterated below for convienence, illustrates some important considerations in the choice of PRBS to use in a partiular situation.

$$H_{\min} = \frac{1}{2}\sqrt{\frac{\overline{v_{\mathrm{n}}^2}}{V^2}\frac{f_{\mathrm{P}}}{f_{\mathrm{s}}}\frac{N}{N+1}}\left(\sqrt{\pi} + \Delta\sqrt{4-\pi}\right) \tag{4.31}$$

From the equation it can be seen that noise resilience may be improved by

- Decreasing the external noise, $\overline{v_{\mathrm{n}}^2}$
- Increasing the PRBS voltage, V
- Increasing the sampling to PRBS clock frequency ratio, $f_{\mathrm{s}}/f_{\mathrm{P}}$

Adjusting the PRBS clock frequency itself is not an option because that would affect the available bandwidth as shown by (4.18). Likewise, it is notable that adjusting the sequence length, N, has no significant effect on noise resilience (although it does affect bandwidth). Based on this information, it would seem sensible to use a high PRBS voltage and high sampling rate for greatest noise reslience.

A high PRBS voltage might pose a number of issues depending on the application. In the thermal characterisation of power electronics considered in this thesis, the input is a varying positive power dissipation. An increase in level would therefore result in increased temperature and, since power electronics have upper temperature ratings, a limit on the practical input power is imposed.

	Calculation	Order of mangitude (f_{P} and N constant)
Variables to store	$2n_{\mathrm{s}}\Big\vert_{n_{\mathrm{s}}=(f_{\mathrm{s}}/f_{\mathrm{P}})N}$	$\mathcal{O}(f_{\mathrm{s}})$
Complex calculations required	$\mathcal{O}(n_{\mathrm{s}}\log_2 n_{\mathrm{s}})\Big\vert_{n_{\mathrm{s}}=(f_{\mathrm{s}}/f_{\mathrm{P}})N}$	$\mathcal{O}(f_s \log_2 f_s)$

Table 4.1 Variation in PRBS computational complexity and storage requirements with switching frequency, f_{s}

Increasing the sampling rate is a less significant issue since frequencies of interest for power electronics tend to be in the sub-hertz range, and faster sampling is therefore easy to implement. A potential difficulty with using a high simpling rate, however, relates to the computational complexity required to calculate the impedance. To identify the system using PRBS, the complete set of readings of the input and output must be stored, and the DFTs of each taken and the quotient evaluated. Since a DFT requires of the

order of $n_s \log_2 n_s$ complex multiplications [4.13], the computational requirements are as given in Table 4.1.

It can therefore be seen that PRBS techniques are most computationally intensive for high sampling rates, or for long high-frequency sequences. However, greatest noise resilience is achieved, for a given bandwidth, when the greatest sampling rate is used. It is therefore necessary to find a compromise between these two competing objectives.

4.4.5 Summary of noise issues in PRBS spectroscopy

It has been shown that in the frequency domain, additive white Gaussian noise has a Rayleigh-distributed DFT with non-zero mean. Using this property, the minimum gain of the system identifiable by PRBS in a noisy environment has been calculated. The relationship between the noise floor and the identified transfer function has been presented with good agreement shown between the estimated and the actual gain at which noise dominates. An acceptable noise level is discussed in terms of the number of standard deviations above the noise floor in a signal. By choosing a sensible number of standard deviations of noise acceptable, an engineer can determine the gain under which the effect of noise is too great and hence discard the affected portions of the identified transfer function. Likewise, the engineer can select a PRBS sequence and sampling rate to deliver the required noise immunity for a particular application.

4.5 Chapter conclusions

This Chapter has explored methods of thermal impedance spectroscopy, a technique for system identification. Initially, pure sine wave identification was presented as a possible method because it is used extensively in other fields of study. Using this technique, the order of the Cauer network required to accurately model thermal transfer impedance, or *cross-coupling*, was determined. Because the pure sine wave technique is arduously long for thermal systems, a pseudorandom binary sequence based technique was reported as an alternative. This technique involves the application of a special signal with uniform frequency content over the frequency band of interest that can be used to measure the thermal impedance over the entire band simultaneously. This technique was then further analysed in terms of its noise resilience. Through first-principles mathematical derivation, the minimum gain which can be successfully determined from

the technique in a practical noisy system is determined. In addition, the increased computational complexity required for noise-resilient fast-sampling is discussed. This value allows a suitable PRBS to be selected for rapid, accurate, noise-resilient and computationally-efficient thermal impedance spectroscopy.

4.6 References

[4.1] R. G. M. Oliveira, M. C. Romeu, M. M. Costa, P. M. O. Silva, J. M. S. Filho, C. C. M. Junqueira*, et al.*, '*Impedance spectroscopy study of Na2Nb4O11 ceramic matrix by the addition of Bi2O3*', *Journal of Alloys and Compounds,* vol. 584, pp. 295-302, 2014.

[4.2] D. Andre, M. Meiler, K. Steiner, C. Wimmer, T. Soczka-Guth and D. U. Sauer, '*Characterization of high-power lithium-ion batteries by electrochemical impedance spectroscopy. I. Experimental investigation*', *Journal of Power Sources,* vol. 196, pp. 5334-5341, 2011.

[4.3] I. Shitanda, K. Inoue, Y. Hoshi and M. Itagaki, '*Analysis of dye-sensitized solar cells with current collecting electrodes using electrochemical impedance spectroscopy, with a finite element method*', *Journal of Power Sources,* vol. 247, pp. 475-480, 2014.

[4.4] M. M. Aziz. (2013). *ECM2105 - Control Engineering: System Response Lecture, Exeter University*. Available: people.exeter.ac.uk/mmaziz/ecm2105/ecm2105_n4.pdf

[4.5] P. Cova and F. Fantini, '*On the effect of power cycling stress on IGBT modules*', *Microelectronics and Reliability,* vol. 38, pp. 1347-1352, 1998.

[4.6] K. V. V. Murthy and R. E. Bedford, '*Transformation between Foster and Cauer Equivalent Networks*', *IEEE Transactions on Circuits and Systems,* vol. 25, pp. 238-239, 1978.

[4.7] J. T. S. Irvine, D. C. Sinclair and A. R. West, '*Electroceramics: characteristation by impedance spectroscopy*', *Advanced Materials,* vol. 2, pp. 132-138, 1990.

[4.8] A. Poppe, Y. Zhang, J. Wilson, G. Farkas, P. Szabo, J. Parry*, et al.*, '*Thermal Measurement and Modeling of Multi-Die Packages*', *IEEE Transactions on Components and Packaging Technologies,* vol. 32, pp. 484-492, 2009.

[4.9] S. Vintrou, N. Laraqi and A. Bairi, '*Calculation and analysis of thermal impedance of microelectronic structures from analytical models*', *Solid-State Electronics,* vol. 67, pp. 45-52, 2012.

[4.10] M. E. Van Valkenburg, "Analog filter design," ed New York: CBS College Publishing, 1982, pp. 459-461.

[4.11] Y. Koito, S. Okamoto and T. Tomimura, '*Two-dimensional numerical investigation on applicability of 45 degrees heat spreading angle*', *Journal of Electronics Cooling and Themal Control,* vol. 4, pp. 1-11, March 2014.

[4.12] W. D. T. Davies, *System Identification for Self-Adaptive Control*. London: Wiley-Interscience, 1970.

[4.13] P. E. Wellstead, '*Non-parametric methods of system identification*', *Automatica,* vol. 17, pp. 55-69, 1981.

[4.14] R. L. T. Hampton, '*A hybrid analog-digital pseudo-random noise generator*', *in Proc* Proceedings of the spring joint computer conference, 1964, pp. 287-301.

[4.15] A. V. S. Oppenheim, R.W., *Discrete-Time Signal Processing*. New Jersey: Prentice Hall, 1999.

[4.16] B. Miao, R. Zane and D. Maksimovic, '*System identification of power converters with digital control through cross-correlation methods*', *IEEE Transactions on Power Electronics,* vol. 20, pp. 1093-1099, 2005.

[4.17] J. Sijbers, A. J. den Dekker, E. Raman and D. Van Dyck, '*Parameter estimation from magnitude MR images*', *International Journal of Imaging Systems and Technology,* vol. 10, pp. 109-114, 1999.

[4.18] N. L. Johnson, S. Kotz and N. Balakrishnan, *Continuous Univariate Distributions, vol. 1*. New York, Wiley: Wiley, 1994.

Chapter V

Increasing the bandwidth and noise resilience of PRBS-based system identification for thermal impedance spectroscopy by mixing sequences

In this Chapter, a new technique, which characterises thermal circuits with improved bandwidth and noise resilience by using a modified pseudorandom binary sequence (PRBS) technique, is presented and compared with the state of the art. By mixing PRBS signals, a new signal is generated which can be used to obtain the thermal impedance spectrum of power electronic systems. The proposed technique increases the useful frequency range of a PRBS by mixing two identical sequences at different frequencies. The new signal incorporates the frequency responses of both contributions. Mixing can be performed using a number of mathematical operators and analysis reveals that AND is the operator of choice since it has the lowest average input power for the same effectiveness. The bandwidth, frequency domain representation and noise resilience of PRBS signals are also reported. It is shown that the noise floor is significantly reduced under the mixed technique, which allows lower impedances to be measured under noisy measurement conditions. For a typical 8-bit PRBS, mixing reduces the noise floor by a factor of 10.5. Simulated and experimental validation are performed and results show the mixed scheme offers increased bandwidth, reduced computation and improved noise resilience compared to single PRBS techniques.

5.1 Introduction

Chapter IV introduced the pseudorandom binary sequence (PRBS) method of system identification, which represents an improvement on the step response method. PRBS, however, increases the computational complexity required to generate a thermal model since it involves frequency domain analysis, typically using the discrete Fourier transform (DFT). It has been seen that PRBS-based methods are most computationally intensive for long sequences or where sample rates are high. It is, however, under these circumstances that the noise resilience and bandwidth of the PRBS is greatest.

In addition, to extract the response for a system containing long time constants, the response must be obtained for the whole period of the longest time constant as no technique is possible which can extract a low frequency response without a long characterisation measurement period. Since the experimental time is already arduously long, it is prudent to apply an excitation signal which has more favourable frequency content, subject to practical limitations such as complexity of the measurement system design and computational complexity for processing.

In this Chapter, a variant of the PRBS is proposed which extends the bandwidth of a standard PRBS into higher frequencies without increasing the experimental time and with only a modest increase in complexity. Additionally, the resilience of the bandwidth extension to noise is greater than that of the standard PRBS with the same bandwidth. This improved noise-resilience counteracts problems due to inaccuracy in system identification at high frequency, which is usually the result of noise in the system.

There are a number of applications in which accurate modelling of high frequency effects can provide significant improvement in performance. Active thermal management, such as the system described in [5.1], is a good example of such an application where an observer system is used to provide estimates of component temperature in order to allow active control of cooling and other forms of temperature management. If a model with inaccurate high frequency response is employed, a greater margin of safety is required by the controller to prevent overheating, resulting in a less efficient design.

It is also possible to apply other techniques to reduce noise. For example, several repeats of the same sequence averaged together reduce uncorrelated noise, as will be shown later in this Chapter. An improved averaging technique is reported by Miao, *et al* [5.2], for example, which uses a cross-correlation method to directly extract the response from several repeats. Both the repeat-and-average method and the cross-correlation method, however, require several repeats of the excitation PRBS, which increases experimental time.

Instead, in this Chapter a technique which reduces noise for the high frequency signal, compared to running the two PRBS sequences separately is presented. In addition, the technique increases the useable frequency range of a PRBS-like signal by mixing two PRBSs at different clock frequencies. The resulting signal and its response can be processed to extract the impedance spectrum over the range of both PRBS sequences with only a modest increase in complexity. Since the proposed technique requires fewer experiments to cover the same frequency range (reducing the need for experimental lead-in time) and fewer PRBS repeats to achieve the same noise level, the experimental time required is significantly reduced.

5.2 Mixing pseudorandom binary sequences

5.2.1 Production

It was shown in Chapter IV that a standard PRBS has a maximum frequency determined by its clock frequency, f_{P}, and a bandwidth determined by its length, N, as shown in (5.1).

$$\frac{f_{\mathrm{P}}}{N} \leq f \leq \frac{f_{\mathrm{P}}}{2.3} \tag{5.1}$$

To extend the frequency range using a standard PRBS, either N must be increased or two separate sequences must be run consecutively. Increasing N increases the number of samples required and therefore the computational complexity required to calculate the DFTs. Each sequence requires a separate settling time to overcome transient effects from the system. It is important to note that measurements must be taken for at least the longest time constant to be identified and therefore running two sequences consecutively increases the amount of time required for the experiment beyond that expected from sequence length because two settling periods would now be required. For these reasons, another solution is desirable.

In this section, a scheme is proposed whereby two sequences are mixed together to produce a new signal which includes the frequency range of both. By choosing two sequences which are adjacent in the frequency domain, the valid frequency range of the derived response can be extended. Fig. 5.1 shows the mixing scheme proposed. Two

identical PRBS sequences are added together and the response of the thermal system taken. Both the input power (q) and the output temperature (θ) signals are processed (indicated by double line). The high and low frequency responses are calculated separately.

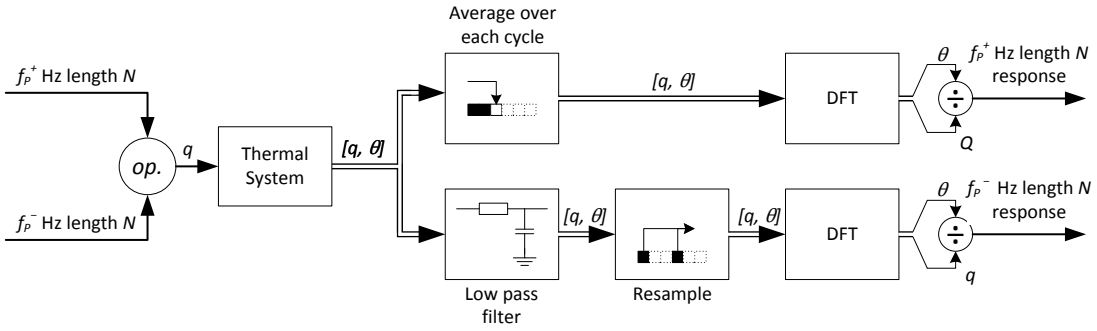

Fig. 5.1 Proposed PRBS mixing scheme

Two identical PRBSs are generated at different frequencies. The higher frequency sequence, which runs at a frequency f_{P}^{+}, is repeated and the duration of the repeated signal exactly coincides with the duration of the lower frequency PRBS, which runs at f_{P}^{-}. This means that the mixed PRBS is periodic with the same period as low frequency source sequence. These two sequences may be mixed together using one of a variety of mathematical operators. Each operator has its own characteristics which will be explored in detail in section 5.6. For now, the *AND* function is considered. To maintain input power, the amplitude of the mixed signal is doubled so that for inputs in the range $[0, Q]$, the mixed output is in the range $[0, 2Q]$. This way, each constituent signal part in the mixed signal has the same power as its corresponding source. The implications of this are explored in section 5.6.

The choice of frequencies is determined by the sequence length and the frequency band desired. From (5.1), choosing the upper limit of the low frequency sequence to be the lower limit of the high frequency sequence, the following relationship is obtained.

$$f_{\mathrm{P}}^{+} = \frac{N}{2.3} f_{\mathrm{P}}^{-} \tag{5.2}$$

Note that the value of f_P^+ will need rounding down such that a whole number of high frequency sequences fits exactly into the low frequency sequence, thus preventing spectral leakage since sampling will be aligned to both source sequences. Therefore the frequency band covered has been extended to

$$\frac{f_{\mathrm{P}}^-}{N} \leq f \leq \frac{f_{\mathrm{P}}^+}{2.3}; \quad \text{or} \quad \frac{1}{N} f_{\mathrm{P}}^- \leq f \leq \frac{N}{5.3} f_{\mathrm{P}}^- \tag{5.3}$$

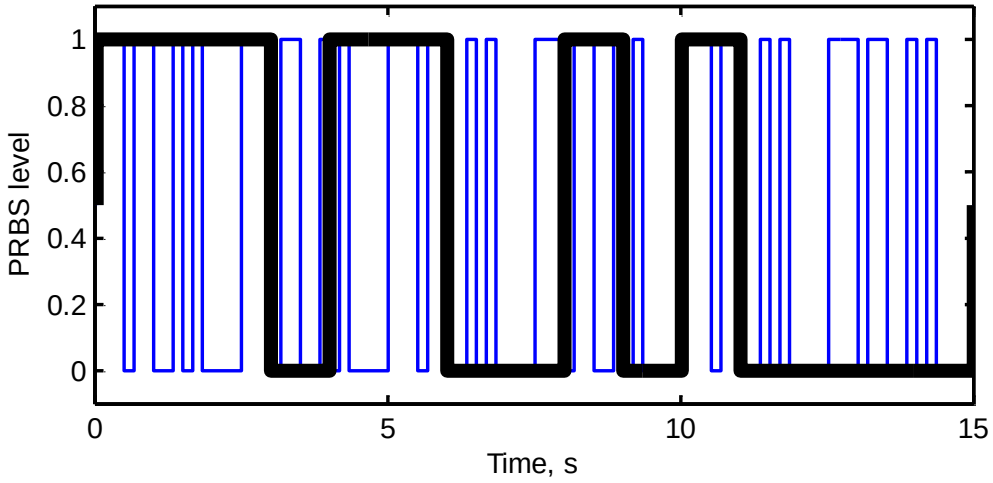

(a)

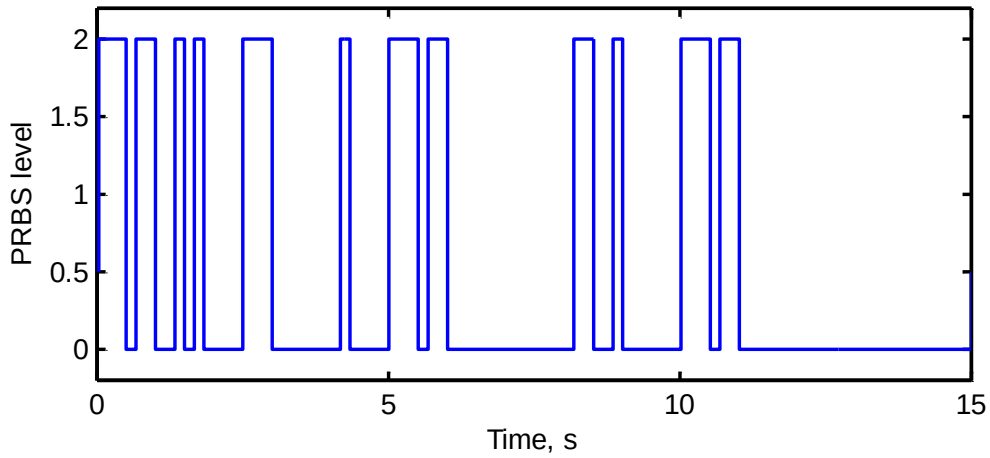

(b)

Fig. 5.2 Source (a) and mixed (b) PRBS signals. Two identical sequences at different frequencies are mixed using the *AND* operator to produce a hybrid mixed sequence with properties of both originals.

By way of example, if a 4-bit 1 Hz sequence (such as that considered in Chapter IV) were to be used as the low frequency signal, the high frequency signal would run at 6 Hz. Thus in each full 1 Hz sequence, six repeats of the 6 Hz sequence are present. To illustrate this effect, Fig. 5.2(a) shows the two original sequences and Fig. 5.2(b) shows the result of mixing them.

The response of a thermal system to the mixed signal is obtained and the high and low frequency responses are extracted separately, as described below.

5.3 Frequency domain analysis and extraction

The concept of mixing signals to shift and replicate frequency content has been extensively applied to analogue signals for many purposes including, for example, modulation and frequency shifting [5.3]. These procedures rely on heterodyning, where two sine waves at different frequencies are multiplied together resulting in a new waveform consisting of two sinusoids which may be calculated according to (5.4).

$$A_1 \sin(\omega_1 t) \cdot A_2 \sin(\omega_2 t) = \frac{A_1 A_2}{2} \left[\cos\big((\omega_1 - \omega_2)t\big) - \cos\big((\omega_1 + \omega_2)t\big)\right] \quad (5.4)$$

The two sinusoids are termed heterodynes and oscillate at the sum and difference of the source frequencies respectively. They are added together to form the resulting signal. When used for modulation, the frequencies are usually well separated with sometimes only one heterodyne retained (the other being removed by filtering). However, if the frequencies are close together and both are retained, the resulting sinusoids will occupy a small frequency band.

It is possible to consider any signal, including a PRBS, as the infinite summation of sinusoids whose frequencies and amplitude are determined from the Fourier transform of the signal. Applying this concept to the PRBS, the effect of the proposed mixing scheme can be explained in terms of heterodyning. If the PRBS is unipolar (i.e. it consists of '0's and '1's only), then the binary *AND* operator is equivalent to multiplication. The effect of mixing can be described in the frequency domain as sum of the heterodynes formed by mixing every frequency in one signal with every frequency

in the other. Fig. 5.3(a) and (b) show linear frequency plots of the DFTs of a 13 Hz and 1 Hz 5-bit PRBS. As both signals are sampled at 52 Hz and are 1 second in length, the high frequency signal contains 13 repeats.

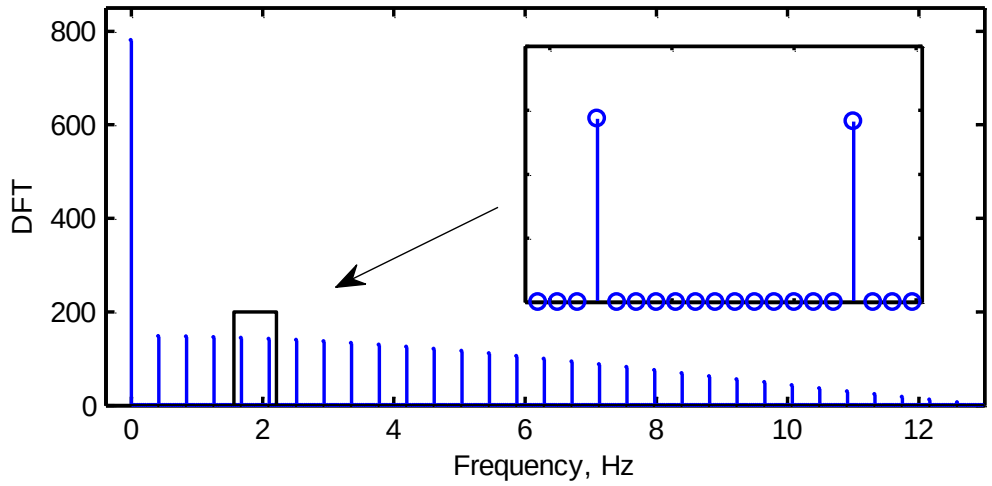

(a)

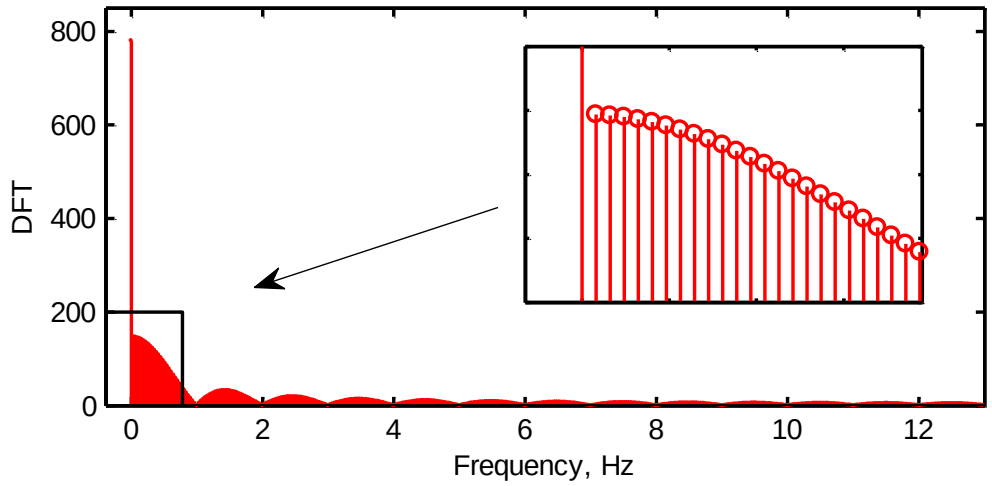

(b)

Fig. 5.3 Linear frequency plots of the high frequency (a) and low frequency (b) individual contributions with areas of interest

One feature of the unipolar PRBS is that it has a large time-domain average compared to the amplitude of its constituent frequencies. This is shown in Fig. 5.3 as the component at 0 Hz, which is of equal height for both plots since their respective sequences are the same length and have the same average.

In Fig. 5.3(a), the effect on the DFT of repeating a sequence can be seen. Because the 13 Hz PRBS is repeated 13 times, only every 13th point is non-zero and the non-zero points follows the normal envelope of a DFT. The DFT is therefore spread out to cover the higher frequencies. However, the low frequency range before the first component at non-zero frequency is not part of the useful frequency band. Instead, this region is covered by the low frequency PRBS in Fig. 5.3(b).

When mixed together, high amplitude frequency components on both plots produce heterodynes with the greatest amplitude while lower amplitude frequency components produce much smaller heterodyne amplitudes. In both plots, the largest component is at 0 Hz and therefore the heterodynes produced from that component have the greatest amplitude. Considering just these resulting high-amplitude heterodynes, the DFT of the resulting waveform would be the two source DFTs superimposed on each other. The other heterodynes are much lower in amplitude and form a low amplitude noise floor in the resulting DFT. Fig. 5.4 shows the DFT of the mixed signal. Colours have been used to highlight which source signal is responsible for each frequency component.

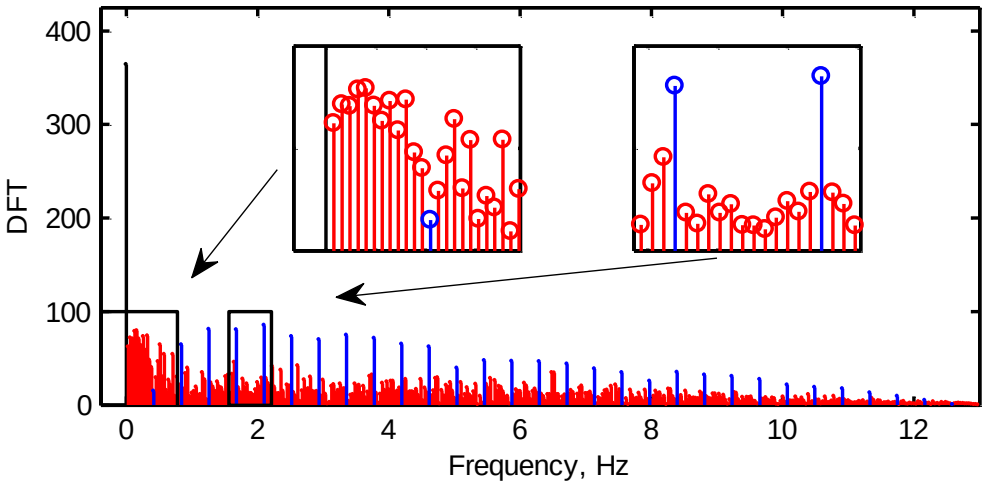

Fig. 5.4 Linear frequency plot of mixed signal DFT with areas of interest inset

The high frequency, high amplitude components remain separated by 12 low amplitude points, a feature which is used to recover the original PRBS as will be seen later. The low frequency components of the low frequency signal remain present in the output. As long as the 0 Hz component in the source signals is sufficiently greater that the other

DFT values, this noise floor will be low enough to be neglected and can be removed when the results are processed, which is described later.

From here on, an 8-bit PRBS is used to illustrate the technique. The use of a longer PRBS provides greater clarity in bode plots of the signal's DFT. From (5.2), the mixing frequency ratio f_P^+/f_P^- is 110. Fig. 5.5 compares the frequency spectrum of the mixed signal for a 0.1/11.0 Hz 8-bit *AND*-mixed sequence with that of the individual contribution PRBS signals. The high frequency individual contribution PRBS DFT has been scaled by a factor of 110 since it is repeated 110 times over the course of the low frequency signal.

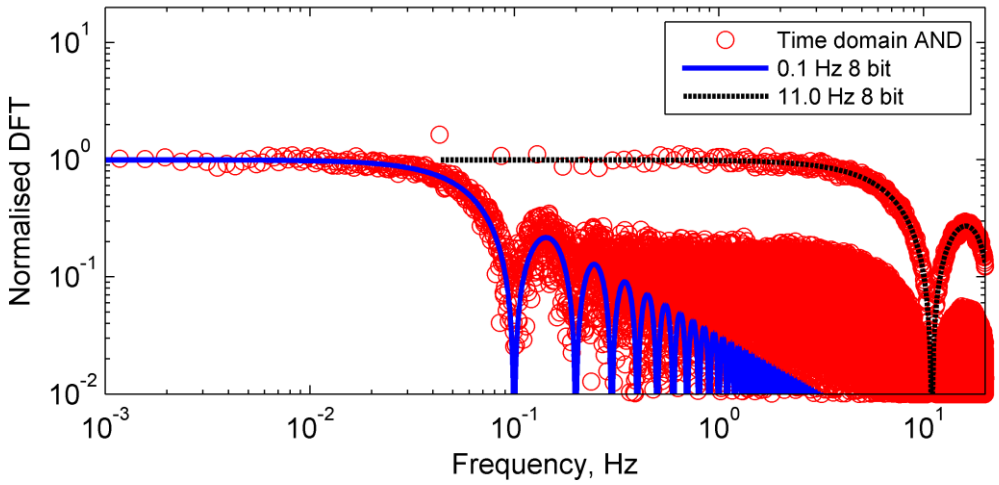

Fig. 5.5 DFT of PRBS sources and mixed signal

The DFT of the mixed PRBS is represented by circles while the lines represent each PRBS alone. Each 110th point of the mixed PRBS agrees well the high frequency PRBS spectrum, especially at higher frequency, in the same way as in Fig. 5.4. The first 109 points of the mixed DFT show good agreement with the low frequency PRBS. In this figure, it is easier to see the separation between the useful frequencies (which show good agreement with the source PRBS spectra) and the noise floor produced by the low amplitude heterodynes. The high amplitude frequency points extend the bandwidth to cover the bandwidths of both individual contributions.

5.3.1 High frequency extraction

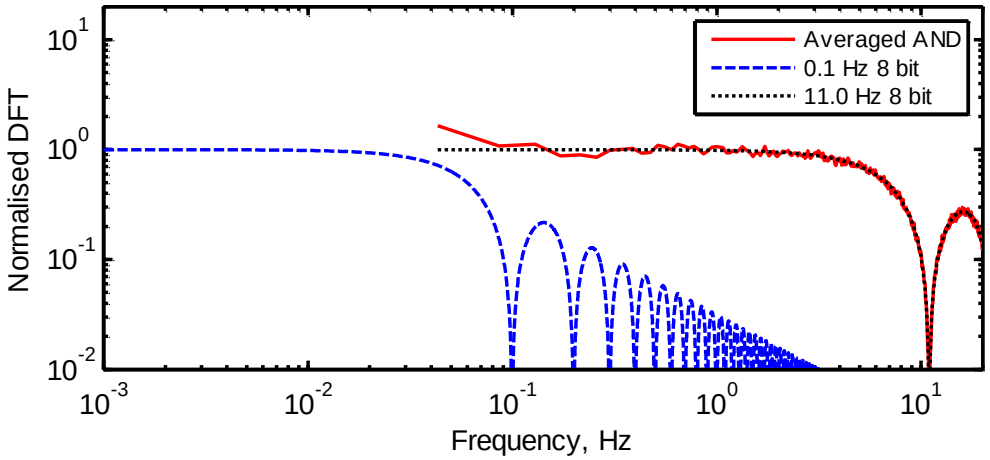

Fig. 5.6 DFT of per sequence averaged AND-mixed sequence compared to PRBS sources

Since high frequency data are to be extracted, the response of the system must be sampled at appropriately high frequency. However, because the high frequency sequence (and therefore its response) is repeated many times only one version of it need be stored. To achieve this, the input and response values at each stage along the repeated sequence are averaged together, producing a single short-length input sequence and response signal to be analysed. This operation is performed by a microcontroller during the experiment thus reducing the number of samples which are to be stored and processed by a factor of $f_{\mathrm{P}}^{+}/f_{\mathrm{P}}^{-}$.

After the whole sequence is complete, the averaged high frequency data can be converted into the frequency domain by taking the quotient of the temperature and power DFTs, discussed in Chapter IV. This is performed according to (5.5) below.

$$Z_{\mathrm{th}} = \frac{\mathcal{F}(\theta)}{\mathcal{F}(q)} \tag{5.5}$$

Where Z_{th} is the thermal impedance and $\mathcal{F}$ is the DFT operator. In our example, this has the effect of sampling the DFT at every 110th point and, since the averaging is synchronised to the high frequency PRBS, it therefore removes the effect of the low frequency PRBS leaving only what would be expected for the high frequency PRBS.

The response has the same useful bandwidth as the high frequency PRBS from which it was derived, shown in (5.1) for $f_P = f_P^+$. Fig. 5.6 shows the DFT of this signal compared to the source signal DFTs. The extracted response follows the high frequency PRBS response well, deviating only at the low end of the spectrum.

5.3.2 Low frequency extraction

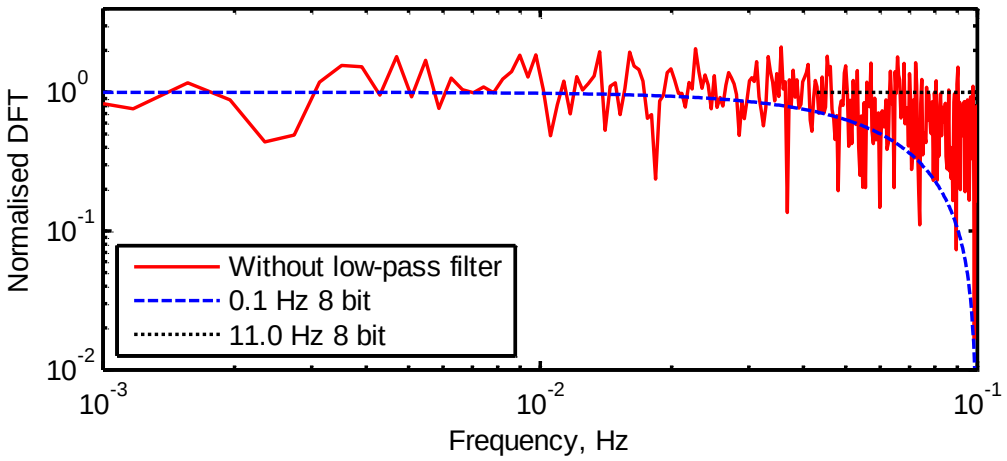

(a)

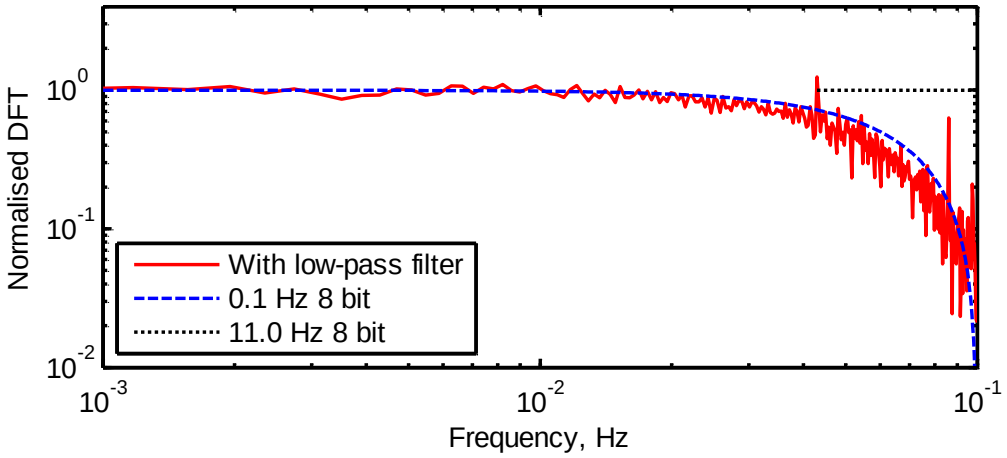

(b)

Fig. 5.7 Comparison of DFT without (a) and with (b) low pass filtering

The DFT at low frequency shows good agreement with the *AND*-mixed sequence (as shown in Fig. 5.5). However, in a practical system, storing and processing all sampled data requires significant memory and high processing capability. For n_s samples, of the order of $n_s \log_2 n_s$ complex multiplications are required to calculate the frequency response [5.4]. Resampling the data would therefore reduce computation time by a

factor of at least the ratio of frequencies (110 in this case). This process therefore allows the low frequency data to be recovered in a less computationally intense way.

However, directly resampling the raw data at a lower frequency causes significant aliasing as, by virtue of the PRBS's pseudorandom nature, the changes in state of the PRBS between samples are not captured. Fig. 5.7(a) shows the DFT of a directly resampled mixed signal (sampled at $f_s = f_P^+/110$). The frequency response differs significantly from the low frequency source, which results in erroneous recovery of the signal.

Low frequency data is therefore extracted after removing the effect of aliasing. Both input and output waveforms are low pass filtered, using corner frequency $f_c = 4f_P^-$. This value is chosen so as not to significantly attenuate the useful data while still significantly reducing the effect of the high frequency PRBS. The filter can be implemented as a digital filter such as that in (5.6), allowing for easy implementation.

$$y[i] = ax[i] + (1-a)y[i-1]; \qquad a = \frac{f_c}{f_s + f_c} \tag{5.6}$$

where y is the filtered output and x is the unfiltered input. The filtering process allows the data to be resampled without aliasing, reducing the amount of data to store and process. The data is therefore re-sampled at a suitable rate; typically f_P^+/f_P^- of its original rate. It will therefore consist of the same number of points as the high frequency data. Both input and response must be filtered and resampled equally to maintain the relationship between them. The re-sampled data can be processed into the frequency response by equation (5.5), the result of which agrees well with the low frequency source as shown in Fig. 5.7(b).

The thermal impedance is calculated by applying these methods to both the input power and output temperature readings.

5.4 Demonstration on a simulated single material

The proposed technique is demonstrated by taking a model thermal system as an example and applying mixed PRBS signals to it. The response is simulated and the mixed PRBS technique applied to it to calculate the thermal impedance. This can then be compared to the ideal thermal impedance taken directly from the model to validate the technique.

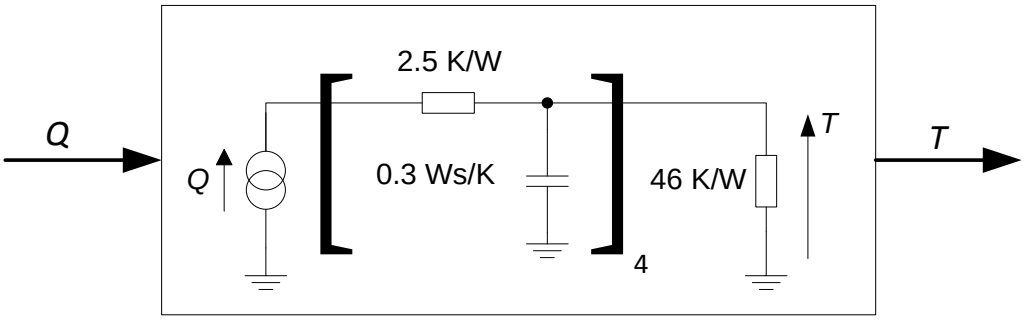

Fig. 5.8 Simulated equivalent circuit of a single material thermal system

An example Cauer network consisting of four identical RC elements with a resistive ambient connection was chosen as the thermal system to be analysed and is as shown in Fig. 5.8. In this Chapter, all temperatures are considered relative to ambient, which is denoted symbolically by a ground connection. The thermal network has a fourth-order transfer impedance shown in (5.7), which is used in the simulations. The thermal resistance and capacitance values were chosen to produce a simple one-material model with a similar thermal response to a TO220-packaged device. Practical results on such a device are evaluated later in this Chapter.

$$Z_{\mathrm{th}}(s)|_{s=\mathrm{j}\omega} = \frac{1}{0.127s^4 + 1.02s^3 + 2.31s^2 + 1.30s + 0.00217} \tag{5.7}$$

Two 8-bit PRBSs are *AND*-mixed to create a new two-level signal. $f_{\mathrm{P}}^{-} = 0.1\,\mathrm{Hz}$ is chosen as the lower frequency and therefore, from (5.2), $f_{\mathrm{P}}^{+} = 11.0\,\mathrm{Hz}$. The mixed PRBS is generated and processed according to the schematic in Fig. 5.1 using the digital low pass filter in (5.6). Fig. 5.9 shows the thermal impedance frequency response of the example thermal system as found using the mixed technique (solid lines) compared to

the ideal response of the system (circles). The vertical line differentiates the two responses, which are extracted separately.

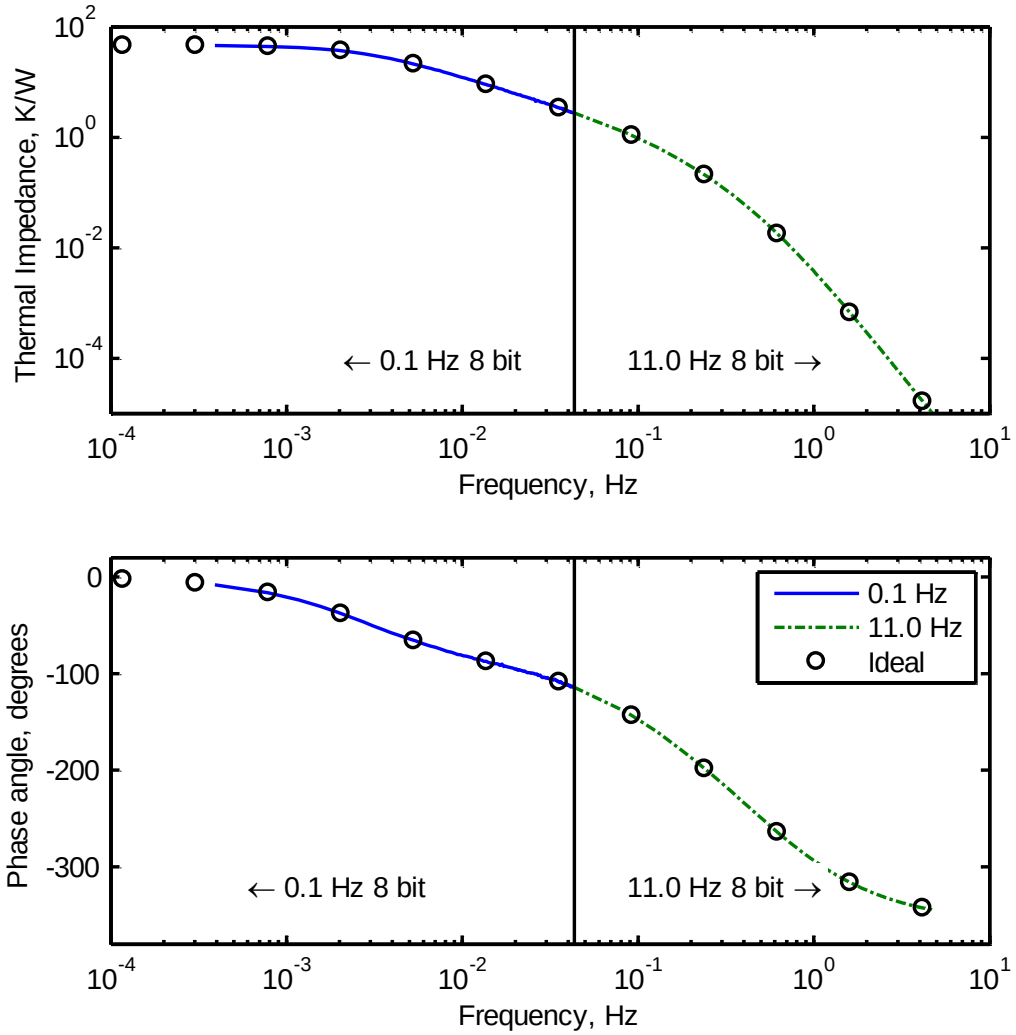

Fig. 5.9 Bode plot of the thermal system derived from mixed PRBS scheme (lines) and ideal (circles)

Results show excellent agreement between both the high and low frequency PRBS results and the ideal filter response over all frequencies of interest. This validates the technique for a linear system such as a thermal equivalent circuit.

5.5 Effect of noise

Results above are for a noiseless system; in reality, noise is present in all measurement systems. When measuring temperature, noise can be particularly intense given the low voltage output of many thermal sensors. To simulate typical measured noise, additive

white Gaussian noise (AWGN) is added to the response of the thermal system. This type of noise has a normal distribution of instantaneous amplitude and an approximately flat power spectrum.

5.5.1 Minimum thermal impedance recoverable

Assuming that the noise measured is AWGN, the effect on a single PRBS can be modelled. An analysis of the minimum gain that can be identified in the presence of noise was reported in Chapter IV. That analysis referred to a bipolar PRBS; however, for thermal analysis the unipolar PRBS is used instead.

Applying the mathematical derivation from Chapter IV, the minimum impedance that can be identified in the presence of noise is

$$Z_{\min} = \sqrt{\frac{\overline{\theta_{\mathrm{n}}^2}}{Q^2}\frac{f_{\mathrm{p}}}{f_{\mathrm{s}}}\frac{N}{N+1}}\left(\sqrt{\pi} + \Delta\sqrt{4-\pi}\right) \tag{5.8}$$

where Δ is the number of standard deviations of noise power perturbations which is acceptable for the application and $\overline{\theta_{\mathrm{n}}^2}$ is the noise on the temperature signal. This equation differs from the equivalent result in Chapter IV because Q is the amplitude of the unipolar power input (whereas the equivalent V was half the peak-to-peak voltage in Chapter IV). The minimum impedance at zero standard deviations ($Z_{\min}$ when $\Delta = 0$) is referred to as the noise floor in this Chapter.

5.5.2 Minimum thermal impedance under the mixed PRBS scheme

Where the noise is uncorrelated to the PRBS frequency, it can be reduced by averaging several repeated sequences [5.2]. Assuming the noise is AWGN, it can be shown (from [5.5]) that averaging n_{av} repeats reduces the initial noise power, $\overline{\theta_{\mathrm{n,0}}^2}$, to

$$\overline{\theta_{\mathrm{n}}^2} = \frac{\overline{\theta_{\mathrm{n,0}}^2}}{n_{\mathrm{av}}} \tag{5.9}$$

where $\overline{\theta_{\mathrm{n}}^2}$ is the new noise power. A significant advantage of the mixing scheme described in this Chapter over using two separate PRBS schemes is that the high

frequency data are obtained from averaging several sequence responses. This causes a significant reduction in noise for the high frequency response where, since thermal systems tend to resemble low-pass filters, the impedance is likely to be lower than the low frequency response.

In the mixed scheme therefore, the noise floor at high frequency is reduced by a factor of $\sqrt{f_P^+/f_P^-}$ over that of the high frequency source PRBS. This is a significant advantage over the using single PRBSs because the increased noise floor is not significant at low frequency but the reduced floor at high frequency, a reduction of 10.5 times for an 8 bit signal, allows recovery of lower thermal impedance in less time and with minimal added complexity.

More noise-resilient methods of processing the high frequency data are also available. Miao *et al* [5.2] proposes a technique where the cross-correlation of the input and output is calculated. This is averaged to reduce its length to that of an unrepeated cross-correlation and converted into the frequency domain using the DFT. This technique can be used to further reduce the noise; however, it requires the entire sequence to be stored before processing. Because an objective of this Chapter is to only modestly increase computational complexity, the proposed technique has been presented using DFTs only.

5.6 Choice of operator

There are several options for the mixing of PRBS sequences. The Boolean operators, *AND* ($\cdot$), *OR* ($\vee$) and *XOR* ($\oplus$) (or their inverted equivalents, *NAND*, *NOR* and *XNOR*) each produce a two-level sequence. The additive operator generates a three level sequence.

Many of the binary operators are equivalent. Since a PRBS is an almost-balanced almost-random sequence, inverting it has no effect. Therefore, from de Morgan's theorem [5.6], for two PRBSs A and B, the equivalencies in equations (5.10) to (5.12) can be written. Hence by comparing *AND*, *OR*, *XOR* and the additive operator, all sensible operators are considered.

$$A \cdot B \equiv \bar{A} \cdot \bar{B} = \overline{A \vee B} \tag{5.10}$$

$$A \vee B \equiv \bar{A} \vee \bar{B} = \overline{A \cdot B} \tag{5.11}$$

$$A \oplus B \equiv \overline{A \oplus B} = \bar{A} \oplus \bar{B} \tag{5.12}$$

AND and *OR* operators also show a degree of equivalency. The similarity is due to the linearity of the system. Neglecting the steady state components of the PRBS and temperature response, the response of the system for an *OR* gate is

$$\theta(t) = z_{\mathrm{th}}([\mathscr{p}(f_{\mathrm{P}}^{+}t) \vee \mathscr{p}(f_{\mathrm{P}}^{-}t)]_2) \tag{5.13}$$

where z_{th} is the time-domain thermal response of the system and $\mathscr{p}$ is the PRBS function. Multiplying this through by (-1) is equivalent to replacing *OR* with *NOR* and since *NOR* is equivalent to *AND* (equation (5.10)), result in

$$-\theta(t) = z_{\mathrm{th}}\left(\overline{[\mathscr{p}(f_P^{+}t) \vee \mathscr{p}(f_P^{-}t)]_2}\right) = z_{\mathrm{th}}([\mathscr{p}(f_P^{+}t) \cdot \mathscr{p}(f_P^{-}t)]_2) \tag{5.14}$$

Hence interchanging *AND* and *OR* is equivalent to multiplying the a.c. component of the input power and therefore the temperature response by (-1), which does not affect the DFT amplitudes or the result. However, the average power in the mixed signal differs between operators, as shown in Table 5.1. This affects the steady-state operating temperature during the experiment. For instance, *AND* has a low average power of 0.4315 compared to 1.4352 for *OR*. This allows a higher input amplitude and therefore higher signal-to-noise ratio while keeping temperature within thermal limitations when *AND* is used.

	AND	*OR*	*XOR*	Additive operator
Operation	$2 \times [A \cdot B]_2$	$2 \times [A \vee B]_2$	$2 \times [A \oplus B]_2$	$A + B$
Average ($f_P^+ / f_P^- = 110$)	0.4315	1.4352	1.0036	0.9333

Table 5.1 Comparison of the average power of mixing operator outputs for an 8-bit PRBS

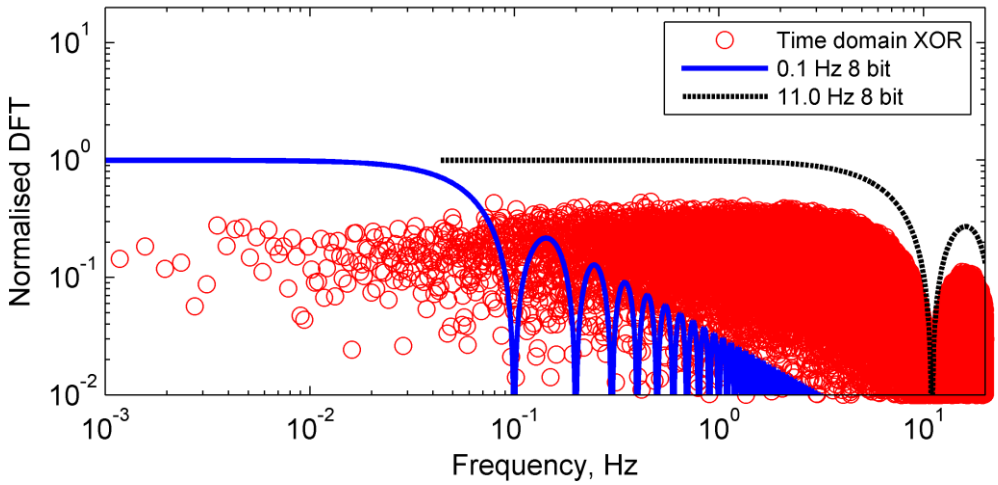

Fig. 5.10 DFT of the mixing scheme using the XOR operator

The *XOR* operator shows significant difference with *AND*, *OR* and the additive operator. Essentially, *XOR* produces a longer random binary sequence, and has established applications for this purpose [5.7]. However, *XOR*ed sequences have characteristics which do not correspond to that expected of a normal PRBS of that length and do not follow the characteristics of the source signals well before or after processing. In particular, *XOR*ed sequences display significant attenuation over the source sequences. Fig. 5.10 shows the DFT of an *XOR*ed mix of 0.1 and 11.0 Hz 8-bit PRBSs. The low quality of the DFT limits *XOR*'s usefulness as a mixing operator.

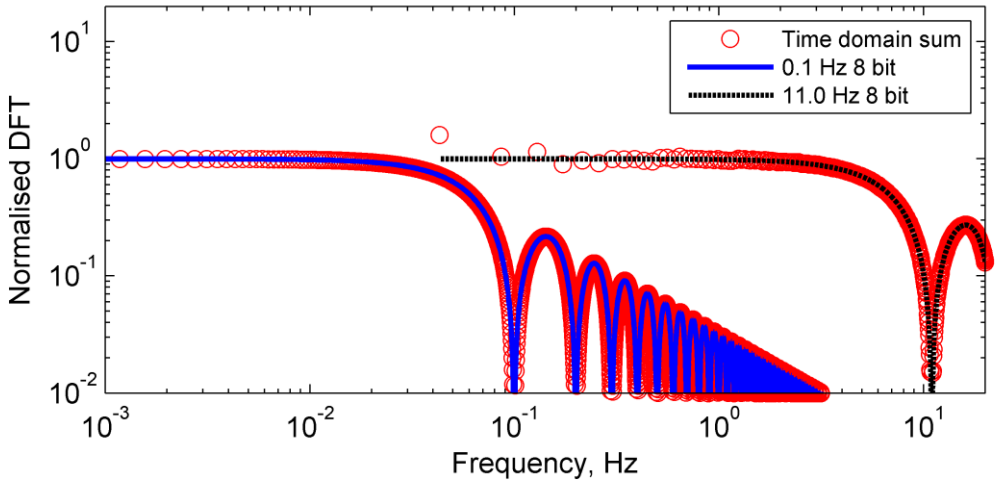

Fig. 5.11 DFT of the mixing scheme using the additive operator

The additive operator produces a three-level PRBS and has an elegant DFT, as shown in Fig. 5.11. This is due to the DFT's linearity which means that the source DFTs themselves can be added together, allowing for very good agreement between source and mixed DFTs. This advantage is reduced, however, by the resulting average power being more than twice that of the *AND* operator. In addition, the additive operator results in a three level output requiring more complex control circuitry. After the low and high frequency extraction techniques have been applied, the frequency response is very similar to that of *AND* and *OR*, hence the additive operator is otherwise equivalent to these operators.

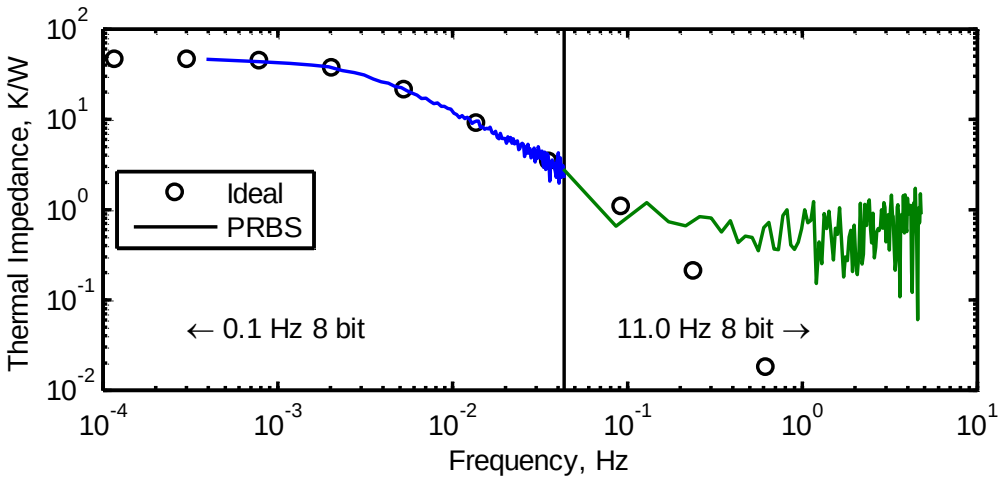

(a)

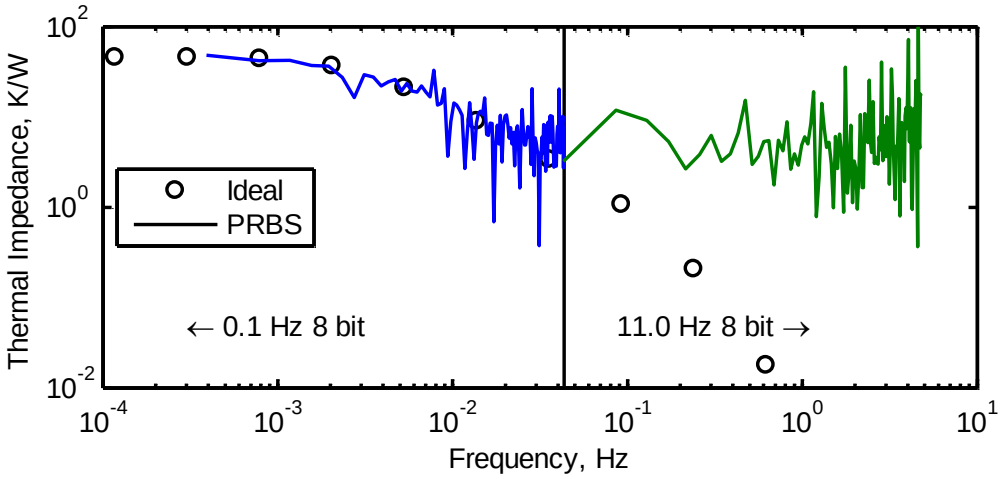

(b)

Fig. 5.12 Effect of AND, OR and additive (a) and XOR (b) operators on noisy mixed PRBS scheme

Fig. 5.12 shows the impedance spectra generated for the thermal circuit in Fig. 5.8. To demonstrate the differences between operators, AWGN of amplitude $\overline{\theta_{\mathrm{n}}^2} = 100\ \mathrm{K}^2$ is added. This is a larger quantity than might normally be expected; however, its use here allows differences between operators to be clearly apparent. Identical white noise signals are added in both cases so that differences are due to the operators alone.

Spectroscopy using *AND*, *OR* and additive operators have identical responses and are shown in Fig. 5.12(a). As shown in Fig. 5.12(b), spectroscopy using the XOR operator is less resilient to the noise on the system. A vertical line separates the low and high frequency responses.

The most appropriate mixing operator is therefore *AND* because it produces a sequence with a low mean power whose resulting spectrum is noise resilient. In addition the *AND* operator's mean power, in the range $[0,\ 2Q]$, is similar to that of the individual contribution PRBS in the range $[0,\ Q]$. This means that, for the same thermal limitations, mixing two PRBS with *AND* does alter the noise floor at low frequency and reduces it by $\sqrt{f_{\mathrm{P}}^{+}/f_{\mathrm{P}}^{-}}$ at high frequency, a 10.5 times reduction for 8-bit sequences. For the additive or *OR* operators, the increased average power will increase the average temperature potentially above the thermal limits of the system.

5.7 Comparison of PRBS techniques

Table 5.2 compares the attributes of the tests necessary to characterise the thermal impedance of a system between 392 µHz and 4.8 Hz. The sampling rate for each example is $f_s = 4f_{\mathrm{P}}$.

Comparing the separate and mixed schemes, the mixed scheme is clearly superior. While covering the same frequency range and requiring the same number of complex calculations (to calculate the input and output DFTs), the mixed scheme allows the measurement of a much lower Z_{th} than the separate scheme. It is also marginally faster. This demonstrates its effectiveness as an alternative to separate PRBSs.

Criterion	0.1/11.0 Hz 8-bit (separate)	0.1/11.0 Hz 8-bit (mixed)	11.0 Hz 15-bit (single)
Frequency range	392 μ - 4.8 Hz	392 μ - 4.8 Hz	336 μ - 4.8 Hz
Sample rate, f_s	0.4/44 Hz	44 Hz	44 Hz
No. of complex calculations (from $n_s \log_2 n_s$)	40 777	40 777	4 456 300
Values to be stored	4 080	4 083	262 136
Lowest Z_{th} at $Q = [0, 1]$ W, $\Delta = 2$ $\overline{\theta_n^2} = 10$ mK^2	181 mK/W	17 mK/W	181 mK/W
Time required excluding settling time	2 573 s	2 550 s	2 978 s
Time required with settling time	7 673 s	5 100 s	5 958 s

Table 5.2 Comparison of PRBS methods for calculating the complex thermal impedance of a thermal system

An alternative is to use a single longer PRBS to cover the entire range by itself. In this case, a 15-bit PRBS is necessary to cover the same frequency range. It is important to note, however, that this PRBS covers a slightly wider frequency range (since f_P^+/f_P^- for the two PRBSs is not a power of two, the selected single PRBS must go lower in frequency if clocked at f_P^+). The time required for the experiment is therefore comparable to that required for the mixed scheme, differing only due to its increased bandwidth. Its failing, however, is the huge number of complex calculations required to determine the DFT. This severely restricts its application in situations where computational power is limited. There is also the requirement to store more data, as the whole sequence must be stored prior to the DFTs being calculated.

Before readings can be taken, the input signal must be applied for a settling time so that transient effects pass and do not affect results. The reasons for this requirement were explored in Chapter IV. For the mixed (and long single) PRBS methods, only one

settling time is required for the entire frequency range since only one experiment is required. By contrast, each separate PRBS requires its own settling time. The settling time is related to the longest time constant in a system and not to the type of PRBS used. An advantage of the proposed technique is that only one settling time is required to cover the same bandwidth as two separate PRBSs. The long PRBS can be discounted as a viable technique since it requires significantly more data storage and processing, contrary to the objectives of the Chapter.

5.8 Experimental comparison between mixed and unmixed techniques

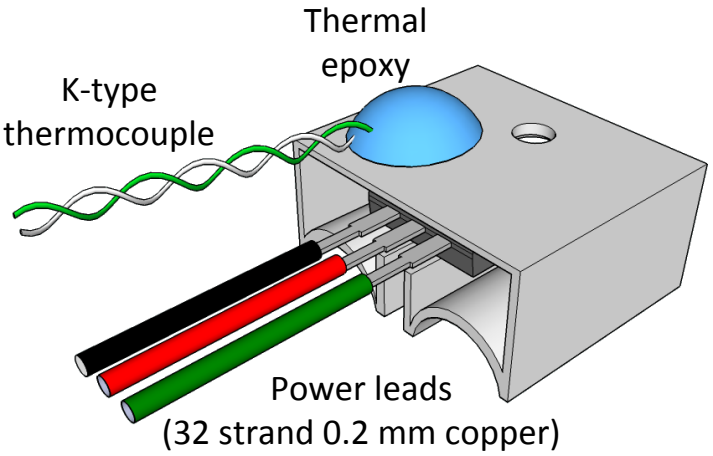

Fig. 5.13 Experimental set up with TO220 package mounted on a 20 K/W heatsink

An IRF510 MOSFET on a 20 K/W heatsink was suspended by its leads in an enclosed but uninsulated environment as shown in Fig. 5.13. A K-type thermocouple was attached to the back of the heatsink using thermal bonding epoxy. Two 8-bit ($N = 255$) PRBS sequences of 0.1 Hz and 11.0 Hz clock frequency (f_P) respectively were mixed using the *AND* operator and applied as a power waveform to the MOSFET. The power input was controlled using the technique reported in Chapter III. The waveform had a power range of [0, 4.88] W. The temperature response sampling frequency was 50 Hz and the results were processed using the mixed PRBS technique described above. The same set-up was then excited with separate source signals with input in the range [0, 2.44] W so that the average input powers were comparable to the mixed PRBS input.

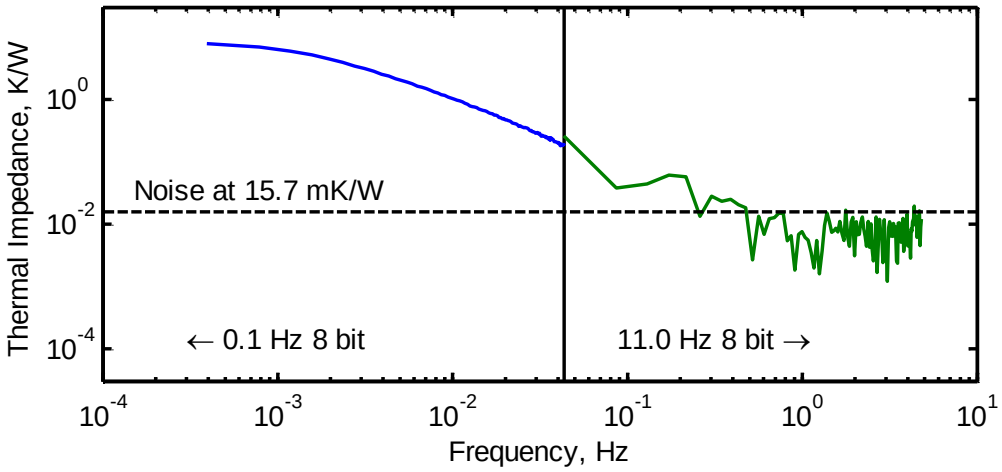

(a)

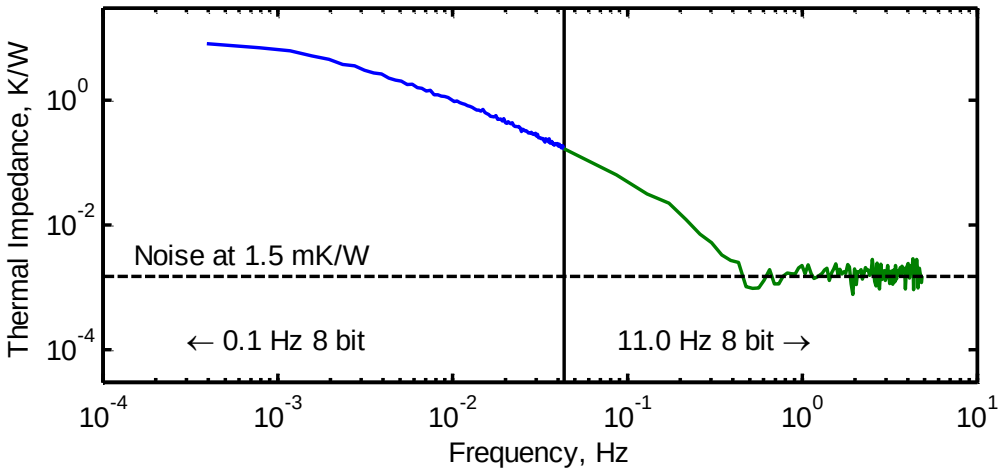

(b)

Fig. 5.14 Bode plot of experimental results from mixed PRBS scheme on a power transistor using (a) the individual PRBS contributions and (b) the *AND*-mixed PRBS. The thermal impedance below which noise is present is indicated

The impedance spectra derived from the source PRBS and the mixed PRBS are presented in Fig. 5.14(a) and (b) respectively. A vertical line separates the low and high frequency responses for each technique. The two methods show excellent agreement at low frequency where measurements of the higher thermal impedance are not affected by noise. At high frequency, however, it is evident that noise has limited the range of impedances over which the spectrum could be accurately calculated. The noise present on the mixed result is significantly reduced compared to that in the source PRBS signals. In the mixed PRBS result, the noise floor is at 1.5 mK/W. From (5.9), noting that 110 repetitions of the high frequency sequence are present in this case, the noise

floor expected of the separate results is at 15.7 mK/W as indicated in the figures. Despite the difference in the input power range, the average temperatures of the mixed and separate experiments are comparable at 37.5°C and 37.1°C respectively.

These results validate the mixed PRBS technique as a method of increasing the signal-to-noise ratio and therefore increasing the useful bandwidth of a PRBS in the noisy conditions typical of thermal tests. It is recognised that this example is of a simple thermal circuit where more basic methods may suffice. The main applications for this technique are envisioned to be those where the thermal model is high order, such as in [5.8].

5.9 Chapter conclusions

A technique to increase the usable frequency range of a PRBS signal when used for thermal impedance measurements has been presented. By mixing two identical PRBS signals at different frequencies, a new signal with properties of both source signals is created. The *AND*, *NAND*, *OR*, *NOR* and additive operators produce input signals which result in identical responses, while the response generated by the *XOR* operator is less noise resilient. The *AND* operator has the lowest average power for the same effectiveness and is therefore the operator of choice. The additive operator, however, has a cleaner and therefore more easily calculable discrete Fourier transform. This advantage is offset by its average power level which is almost twice that of the *AND* operator, and its requirement for a three-level input signal. When the mixed PRBS is applied to a thermal system, the power and temperature waveforms can be processed to determine the complex impedance response over the combined frequency range of both source PRBSs. Mathematical validation has shown that in ideal conditions, the thermal impedance spectrum derived exhibits excellent agreement with the correct result. In noisy conditions, the mixed technique produces a high frequency response which is more noise resilient than the corresponding source PRBS due to the calculation method involving the averaging of several repeats of a signal. For a typical 8-bit PRBS, a 10.5-fold reduction in noise can be expected. Experimental validation on a heatsinked transistor demonstrates the increased noise resilience and bandwidth of the mixed technique using the *AND* operator, while the average power and average temperature, do not differ significantly between *AND*-mixed and separate PRBS techniques.

5.10 References

[5.1] J. D. Holmes, M. P. Foster and D. A. Stone, '*System-wide Temperature Estimation for IMS based Power Electronics Circuits*', *in Proc* International Conference on Power Electronics and Drive Systems, 2009, pp. 1081-1084.

[5.2] B. Miao, R. Zane and D. Maksimovic, '*System identification of power converters with digital control through cross-correlation methods*', IEEE Transactions on Power Electronics, vol. 20, pp. 1093-1099, 2005.

[5.3] B. Sklar, *Digital Communications: Fundementals and Applications*, 2nd ed. Englewood Cliffs, New Jersey: Prentice Hall, 2001.

[5.4] A. V. S. Oppenheim, R.W., *Discrete-Time Signal Processing*. New Jersey: Prentice Hall, 1999.

[5.5] N. L. Johnson, S. Kotz and N. Balakrishnan, *Continuous Univariate Distributions, vol. 1*. New York, Wiley: Wiley, 1994.

[5.6] P. Horowitz and W. Hill, *The Art of Electronics*. New York: Cambridge Universtiy Press, 1989.

[5.7] S. C. Albers, T. J. Callaghan, W. L. Rasmussen and R. A. Pajak, '*Noise generator using combined outputs of two pseudo-random sequence generators*', US patent 5153532, 1992.

[5.8] O. S. Senturk, L. Helle, S. Munk-Nielsen, P. Rodriguez and R. Teodorescu, '*Power Capability Investigation Based on Electrothermal Models of Press-Pack IGBT Three-Level NPC and ANPC VSCs for Multimegawatt Wind Turbines*', IEEE Transactions on Power Electronics, vol. 27, pp. 3195-3206, 2012.

Chapter VI

Estimating the temperature of power electronic devices in real time using PRBS-generated thermal cross-coupling characteristics

This Chapter presents a technique to estimate the temperature response of a multi-element thermal system based on the thermal cross-coupling between elements. The complex frequency-domain cross-coupling of devices is first characterised using a pseudorandom binary sequence (PRBS) technique. The characteristics are then used to estimate device temperatures for a known input power waveform using a discrete Fourier transform-based technique. The resulting estimation shows good agreement with an example practical system used for evaluation. To reduce the computational complexity of the initial method, a digital infinite impulse response (IIR) filter is fitted to each cross-coupling characteristic. A high correlation fit is demonstrated which produces a near-identical temperature response compared to the initial procedure whilst requiring fewer mathematical operations. Experimental validation on the practical system shows good agreement between IIR filter estimations and practical results. It is further demonstrated that this agreement can be substantially improved by taking feedback from an internal reference temperature. Additionally, the proposed IIR filter technique allows the efficient calculation of future device temperatures based on simulated input, facilitating future temperature predictions.

6.1 Introduction

Chapters IV and V presented the pseudorandom binary sequence technique as a method for determining the thermal self impedance of devices in a power electronics system. However, practical systems typically consist of several power devices mounted in close proximity, often on a shared cooling system. When considered in isolation, the temperature of each of these devices is a function of only its own power dissipation waveform and the cooling system employed. In a compact system, however, its

temperature is also heavily dependent on dissipation waveforms in other components. This dependency is termed the cross-coupling between the components.

The importance of cross-coupling is demonstrated by Fig. 6.1 which shows two thermal images of four power devices (1, 2, 3, 4) mounted on a heatsink. If device 2 dissipates 95 W then its temperature rises to 113°C, which is acceptable for a silicon device. However, if both devices 2 and 4 each dissipate 95 W simultaneously, the temperature at each device rises beyond 134°C, which is potentially damaging to many silicon devices. The 21°C difference in temperature results from cross-coupling which would not be predicted by simple thermal modelling. This demonstrates the need to take cross-coupling effects into account for accurate thermal modelling.

Chapter II reported on the extensive previous literature pertaining to cross-coupling characterisation. Methods used include the fitting of electrical circuit analogues [6.1][6.2][6.3] and the development of state-space differential equations [6.4]. Once identified, cross-coupling characteristics can be used to calculate the temperature at each system element due to the power dissipations in all elements. These estimation methods are important in power electronics because they allow a design to be evaluated in terms of its compliance with the rated maximum temperature of elements under a range of operating conditions.

Some techniques, such as that proposed by Brückner *et al* [6.1] consider the peak temperature which each device will attain under various operating conditions. Although such works consider the cross-coupling between power devices, the temperature is not calculated with respect to time. In a system with variable operating conditions, a method which estimates the temperature at a given time for an arbitrary dissipation profile is required. Techniques which do calculate the real-time temperature response have been applied to the very large scale integration (VLSI) technology area [6.3][6.5]. Applications have also been found in power electronics packaging [6.6].

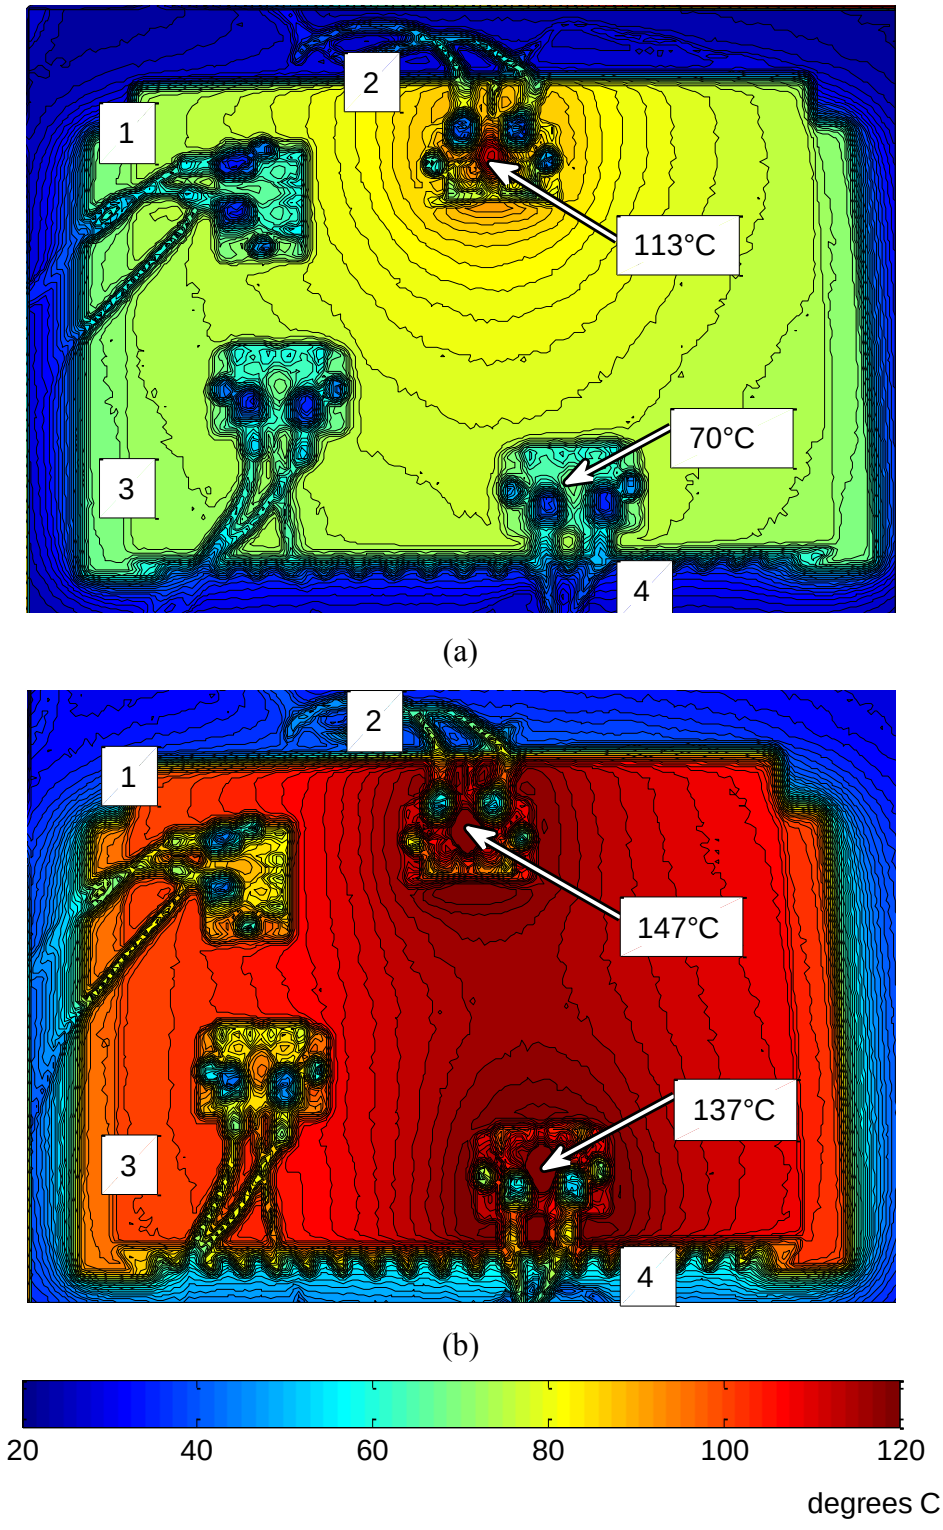

Fig. 6.1 Thermal images of four devices mounted on a heatsink with the temperatures of devices 2 and 4 indicated. (a) Device 2 dissipating alone (b) Both device 2 and device 4 dissipating.

In recent years, similar approaches have been applied to discrete component power electronics. For example, James *et al* [6.2] have reported using a step function to evaluate the auto- and cross-coupling between elements of multichip power electronic module and fit these characteristics to a Foster network. The resulting model is used to estimate the temperature response in real-time for an arbitrary power input. This method has the disadvantage that the difference equations are not generated directly from the characterisation. They must instead be generated from an equivalent circuit model using transient circuit analysis.

In this Chapter, the cross-coupling is directly measured for use in temperature estimation using a PRBS technique. To provide real-world focus, the thermal design of a typical H-bridge converter, an application where there are significant thermal constraints, is used for evaluation [6.7]. The techniques in this Chapter are particularly applicable to power converters for electric vehicles (EVs) where thermal constraints are more severe than in static applications [6.8].

The approach adopted in this Chapter assumes the positions of devices in a thermally coupled system are fixed. It is therefore necessary to characterise the thermal properties of the system for each device arrangement, making the approach most suited for production systems, rather than prototype systems. In addition, a superposition-based approach is used to evaluate the thermal response with multiple devices dissipating, which can cause difficulty where the thermal properties of the system are highly temperature-dependent. The criteria to ensure the accuracy of this superposition-based approach are therefore evaluated in 6.2.2. This Chapter also assumes that a technique to estimate the dissipation in each device is available. Development of such a technique is a nontrivial matter since the power dissipation is often highly temperature-dependent. This problem is out of the scope of this Chapter; however, the temperature predictions produced could be used as an input to whatever model is used.

Notwithstanding the above assumptions and limitations, the proposed techniques have a number of advantages. Foremost, the documented procedure can be used to produce very computationally efficient temperature estimators and predictors while still retaining good accuracy. More accurate approaches such as finite element modelling, by contrast, require considerable computer resources. Additionally, the need to construct and

parameterise a descriptive electrical equivalent circuit is removed by direct measurement of the cross-coupling characteristics, meaning the described procedure requires limited engineer time to implement.

6.2 Definition and generation of cross-coupling characteristics

6.2.1 Definition

A thermal system can be modelled as an electrical equivalent circuit, consisting of resistors and capacitors. Fig. 6.2 shows a typical thermal equivalent circuit between multiple heat sources (represented as current sources) and temperature readings (represented as voltages).

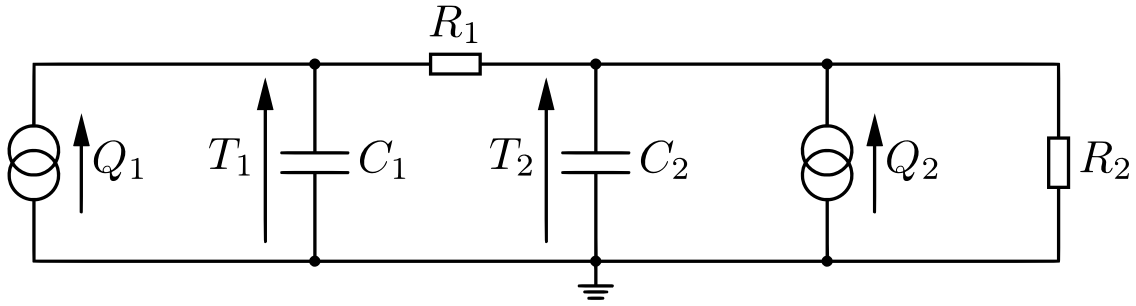

Fig. 6.2 Schematic of multiple heat sources and temperature measurements

The topology of the network model differs both between systems and with the type of modelling used. As in all linear two-port systems, however, the relationship between a single power source and single temperature reading can be simplified into a Norton equivalent circuit with impedance dependent on frequency. Fig. 6.3 shows the Norton circuit between arbitrary heat source *x* and temperature measurement point *y*. In this work, Q, Z_{th} and Θ refer to the frequency domain representations of heat flux, thermal impedance and temperature respectively. In addition, all temperatures reported are relative to ambient which is represented symbolically as an earth connection.

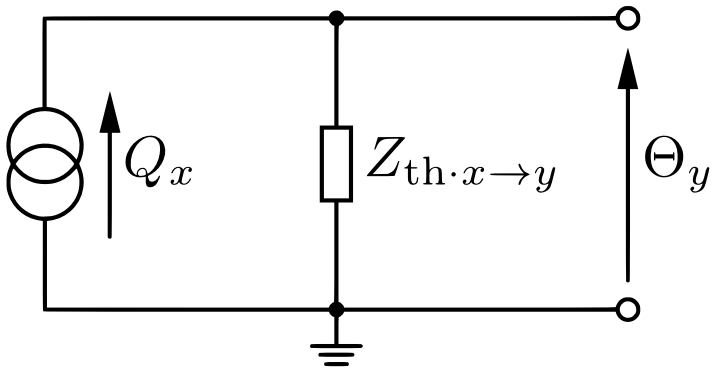

Fig. 6.3 Norton equivalent circuit of the thermal circuit between heat source *x* and measurement point *y*

The cross-coupling, $Z_{\mathrm{th}\cdot x\to y}$, is the complex frequency-domain transfer impedance between heat source *x* and measurement *y*. By applying Ohm's law and the theory of superposition, the temperature at all measurement points due to all device dissipations can be expressed in matrix form as shown in (6.1). The $n \times m$ thermal impedance matrix is the combined characteristics of the entire system.

$$\begin{pmatrix} \Theta_1 \\ \vdots \\ \Theta_m \end{pmatrix} = \begin{pmatrix} Z_{\mathrm{th}\cdot 1\to 1} & \cdots & Z_{\mathrm{th}\cdot n\to 1} \\ \vdots & \ddots & \vdots \\ Z_{\mathrm{th}\cdot 1\to m} & \cdots & Z_{\mathrm{th}\cdot n\to m} \end{pmatrix} \begin{pmatrix} \mathrm{Q}_1 \\ \vdots \\ \mathrm{Q}_n \end{pmatrix} \tag{6.1}$$

Where n and m are the number of devices and temperature measurement points in the system respectively. In this work, both temperature measurement and power input are taken for each of the power devices, therefore $n = m$. Equation (6.1) must be evaluated for each frequency to generate the complete temperature response.

6.2.2 Superposition in thermal analysis

The heat equation, introduced in Chapter II and restated below, may be used to calculate the temperature response of a system to a given power input.

$$\frac{\partial \theta(t)}{\partial t} - \alpha \nabla^2 \theta(t) = \frac{\alpha}{\lambda} q(t) \tag{6.2}$$

Where $\theta(t)$ and $q(t)$ are the time-domain temperature and power fields, respectively; α and λ are the thermal diffusivity and thermal conductivity of the material, respectively; ∇ is the Laplace operator and t is time. For a system with multiple heat sources $q_1 \cdots q_n$, the temperature field can be calculated by solving (6.3).

$$\frac{\partial\theta(t)}{\partial t} - \alpha\nabla^2\theta(t) = \frac{\alpha}{\lambda}\sum_{i=1}^{n} q_i(t) \tag{6.3}$$

If an equation of the form (6.2) is generated for each of $q_1 \cdots q_n$, and the generated equations are summed, the following expression can be written. This equation represents the superposition of the expected temperature responses to each of the heat sources individually.

$$\sum_{i=1}^{n}\left[\frac{\partial\theta_i(t)}{\partial t} - \alpha\nabla^2\theta_i(t)\right] = \frac{\alpha}{\lambda}\sum_{i=1}^{n} q_i(t) \tag{6.4}$$

This expression can be decomposed into (6.5).

$$\sum_{i=1}^{n}\frac{\partial\theta_i(t)}{\partial t} - \alpha\sum_{i=1}^{n}\nabla^2\theta_i(t) = \frac{\alpha}{\lambda}\sum_{i=1}^{n} q_i(t) \tag{6.5}$$

By inspection, it can therefore be seen that

$$\frac{\partial\theta(t)}{\partial t} - \alpha\nabla^2\theta(t) = \sum_{i=1}^{n}\frac{\partial\theta_i(t)}{\partial t} - \alpha\sum_{i=1}^{n}\theta_i(t) \tag{6.6}$$

This implies that the individual temperature responses can be superimposed, as seen in (6.7).

$$\theta(t) = \sum_{i=1}^{n}\theta_i(t) \tag{6.7}$$

If the above analysis is valid, it follows that superposition can be used to analyse a thermal system with multiple heat sources. The analysis is valid if α and λ are constants: if they are not, the generation of (6.5) is invalid because the values of α and λ from each source equation cannot be factored out of the summation.

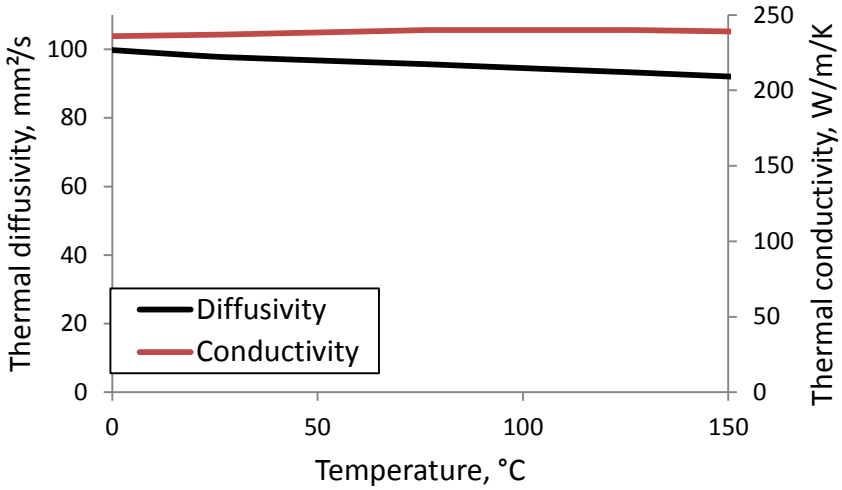

Fig. 6.4 Thermal diffusivity and conductivity of aluminium [6.9]

In this Chapter, the thermal properties of an aluminium heat sink are studied. Fig. 6.4 shows the thermal diffusivity and conductivity of aluminium over the normal operating temperature range. It can be seen that there is little variation in either property under practical temperatures and therefore the assumption that α and λ are constants is reasonable. It is therefore valid to use superposition in the thermal analysis.

6.2.2.1 Impact of temperature-dependent device properties on superposition

The superposition theorem considered above assumes that the heat dissipation in each device can be known accurately. This calculation is limited if only based on a simple electrical characterisation of the device, such as from an estimation of device resistance and capacitance, because many device properties have temperature dependency. Instead, feedback on the estimated device temperatures is required to adjust the dissipation model in order to take account of these secondary effects.

This Chapter, however, does not consider the difficulties in determining device temperature. The work assumes that a viable application-specific method of estimation exists and therefore the input power dissipation is accurate. Instead, the Chapter

focusses on methods of temperature estimation from known power dissipation. Assuming that the power dissipation can be accurately known and the medium of heat transfer fulfils the requirements set out in 6.2.2 (as is the case for an aluminium heatsink), superposition can validly be used. In cases where the requirements of 6.2.2 are not fulfilled, such as in the modelling of phase-change materials, another more complex approach must be adopted.

6.2.3 Characterisation of cross-coupling using pseudorandom binary sequences

To characterise $Z_{\mathrm{th}\cdot x\to y}$ between components *x* and *y* over a frequency range, the single sequence PRBS is used. In brief, this PRBS technique applies a specially generated excitation signal which has an almost flat power spectrum over the band of frequencies given in (6.8). In this Chapter, angular frequencies are used in equations to simplify the mathematics. A thorough description of the technique is provided in Chapter IV.

$$\frac{\omega_{\mathrm{p}}}{N} \leq \omega \leq \frac{\omega_{\mathrm{p}}}{2.3} \tag{6.8}$$

Where ω is the valid frequency band, ω_{p} is the clock frequency of the PRBS and N is the number of periods of ω_{p} that are present in a single sequence. By applying this signal as an input power waveform to a thermal system and measuring the resulting temperature waveform, the system can be characterised over the frequency band. In this way, the response of the system can be determined over a range of frequencies simultaneously. As demonstrated in Chapter IV, the cross-coupling spectrum between two points in a system under test is generated by applying a PRBS input power waveform to device *x*, measuring the temperature response at point *y* and then taking the quotients of the discrete Fourier transforms of power and temperature as shown in (6.9).

$$Z_{\mathrm{th}\cdot x\to y}(\omega) = \frac{\mathcal{F}\left(\theta_y(t)\right)}{\mathcal{F}\left(q_x(t)\right)} = \frac{\Theta_y(\omega)}{Q_x(\omega)} \tag{6.9}$$

Where $\mathcal{F}$ is the discrete Fourier transform operator and ω is angular frequency. θ_y and q_x are the time domain temperatures and power readings at points *y* and *x* respectively.

6.3 Experimental arrangement

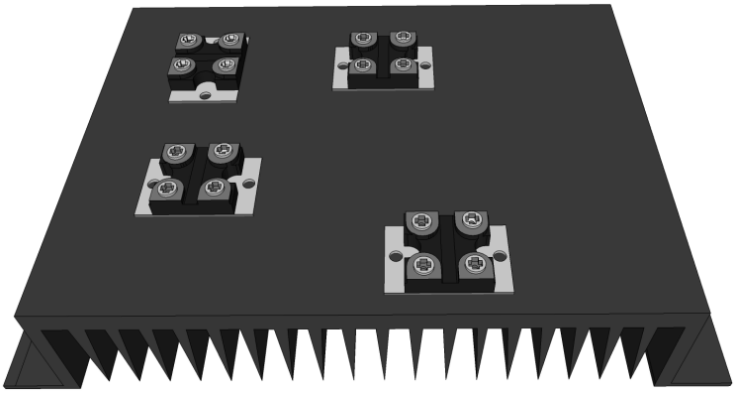

(a)

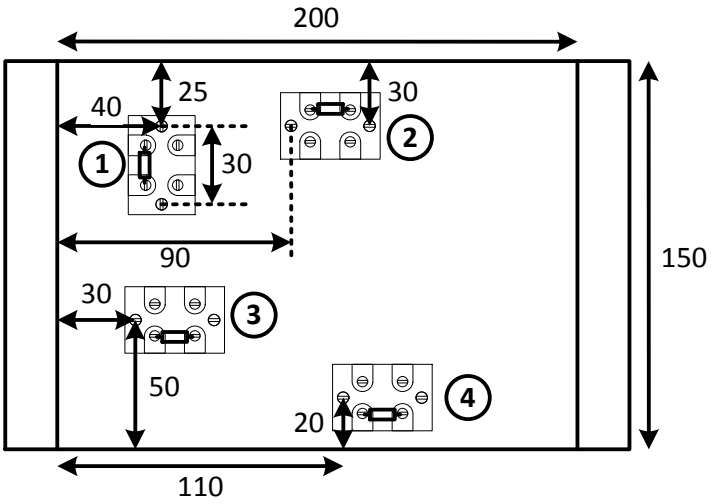

(b)

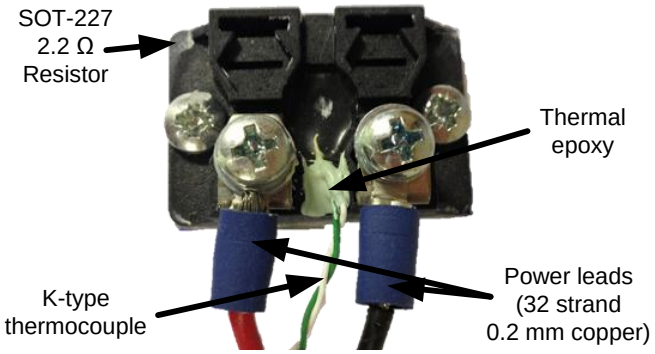

(c)

Fig. 6.5 Heatsink arrangement. (a) 3D view. (b) Plan view with measurements in millimetres and resistor number circled. (c) Close up of individual device.

A typical arrangement for power electronics is used to demonstrate the techniques outlined in this Chapter. An H-bridge based multi-kilowatt converter, with four

SOT-227 packaged power devices mounted on a passively cooled 0.47 K/W heatsink, is selected. The power devices are placed irregularly as shown in Fig. 6.5 to reflect the fact that the placement of devices is a compromise between thermal and electrical design [6.10]. The MOSFETs in the design are substituted for power resistors since this allows the power dissipated to be controlled more easily and removes the need to construct a complex control circuit. A power controller determines the instantaneous dissipation in the devices, which can be set to match the operational losses in a converter so that the thermal characteristics of the system are unchanged.

A thermocouple is attached with thermal epoxy resin to the top of each package in the recess directly above each resistor to measure the device temperature as shown in Fig. 6.5(c). In this Chapter, Θ_1 to Θ_4 and Q_1 to Q_4 refer to the frequency domain temperature and power dissipation in devices 1 to 4, as identified in Fig. 6.5(b), respectively.

Temperature measurement was performed using a LabJack U6 measuring to an effective resolution of 0.06 K at 25 samples per second, giving the system a maximum accessible bandwidth of 12.5 Hz. The rise time of the thermocouples is 0.1 s while the bandwidth of the waveform generator and power dissipation controller are in excess of 1 kHz. Overall, the system bandwidth is therefore at least 10 Hz. In this work, the maximum frequency used is 0.1 Hz, a value sufficiently below the system bandwidth as to make results independent of the measurement system. The power spectral density of the noise from the thermocouple measurement was 1.4×10^{-4} K^2/Hz.

6.4 Measurement of thermal cross-coupling

To measure the thermal cross-coupling between the devices, a pseudorandom binary sequence (PRBS) was applied to each resistor using the technique described in Chapter IV. Each device in turn was excited with a 0.004 Hz 8-bit PRBS with amplitude 95 W and the resulting temperature responses at all devices were logged. Each experiment was repeated with three cycles of a 0.1 Hz PRBS to extend the valid frequency range. From (6.8), the valid frequency band of the PRBS is therefore 15.6 μHz to 43 mHz.

Fig. 6.6 shows the cross-coupling calculated between each pair of devices on the heatsink. In addition, the thermal auto-coupling, $Z_{\mathrm{th}\cdot x\to x}$, is measured and indicated.

This is the temperature rise in a device due to its own dissipation and is greater in value than the cross-coupling between devices, with the difference being particularly large at high frequency.

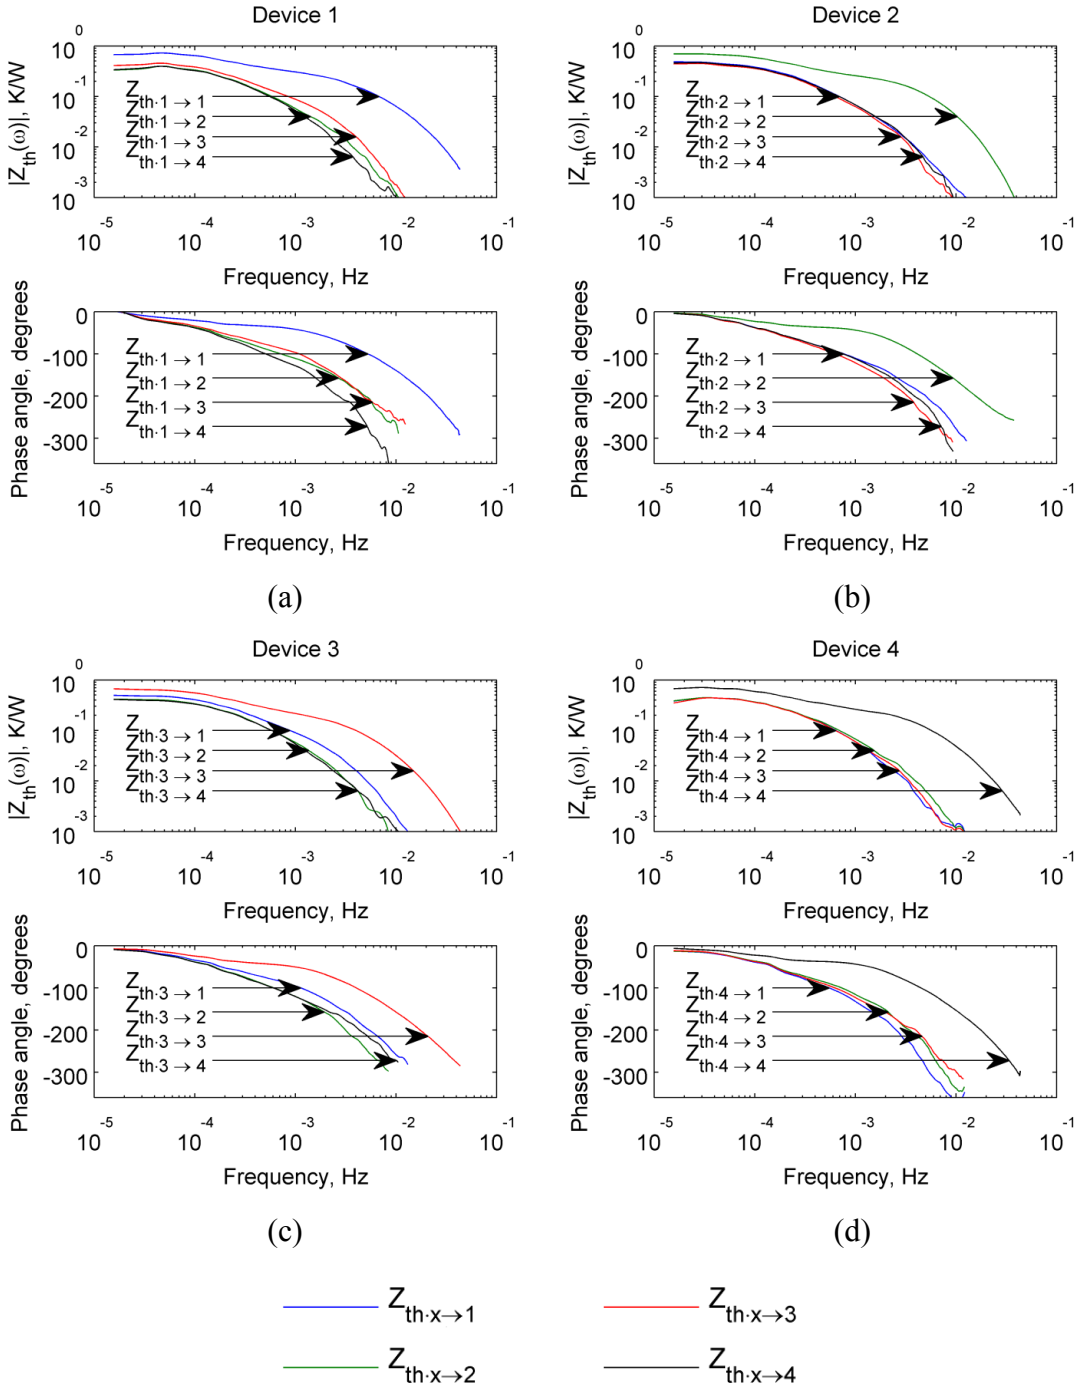

Fig. 6.6 Cross-coupling between components. (a) From device 1. (b) From device 2. (c) From device 3. (d) From device 4.

The cross-coupling characteristics between different devices are broadly similar in shape because they are mounted on the same heatsink under similar conditions. However, there are notable differences in the coupling between different devices. The coupling between devices 1 and 3, for example, is significantly elevated compared to other devices (evident in Fig. 6.6(a) and (b)). This is due to the close proximity of the devices. Similarly, the coupling between devices 1 and 2 is counter-intuitively lower than that of devices 1 and 3 even though they are physically separated by the same distance. This is due to the heatsink's fin arrangement lowering the thermal resistance between devices 1 and 3, but not devices 1 and 2.

6.5 Real-time estimation of temperatures

6.5.1 Estimation using frequency domain data

Having obtained the cross-coupling characteristics between each pair of devices, the temperature response of each device due to arbitrary input power waveforms in any combination of the devices can be estimated. In this section, how this estimation can be made using frequency-domain analysis of the input powers and cross-coupling is examined. Fig. 6.7 shows the procedure for estimating the temperature using the discrete Fourier transform (DFT). First, for each pair of devices *x* and *y*, the DFT of the temperature at device *y* due to the power dissipation in device *x* is calculated and termed $\Theta_{y(x)}$. Subsequently, the temperature at each device is calculated by summing the contributions from each power source and evaluating the inverse DFT. This procedure is described in detail below.

The input power, $q_x(t)$, which is to be analysed using this procedure will be defined over the finite time period $0 \leq t \leq t_{\max}$. However, it is not possible to calculate the temperature response over the same period using this power waveform directly because the discrete Fourier transform which is used in this procedure assumes that its input is periodic. Assuming the temperature difference from each device to ambient is initially zero, the output of the inverse DFT used in processing must also begin and end with the value zero. To achieve this, the simulation period is extended and the input power waveform $q_x(t)$ is padded with zeros such that any stored thermal energy at time $t_{\max}$ (due to the input power) is dissipated to ambient before the end of the simulation period, as shown in (6.10), thereby fulfilling the DFT's periodicity criterion. It is normally

reasonable to wait three to five times the longest time constant in the system for it to settle to steady state, which results in a simulated final temperature of 0.7% to 5% of the present value. In the following example, a delay of three time constants is used.

$$q'_x(t) = \begin{cases} q_x(t), & t \le t_{\max} \\ 0, & t_{\max} < t \le t_{\max} + 3\tau_{\max} \end{cases} \tag{6.10}$$

Where $t_{\max}$ is the maximum time for which $q_x(t)$ is defined and $\tau_{\max}$ is an estimate of the longest time constant of the system.

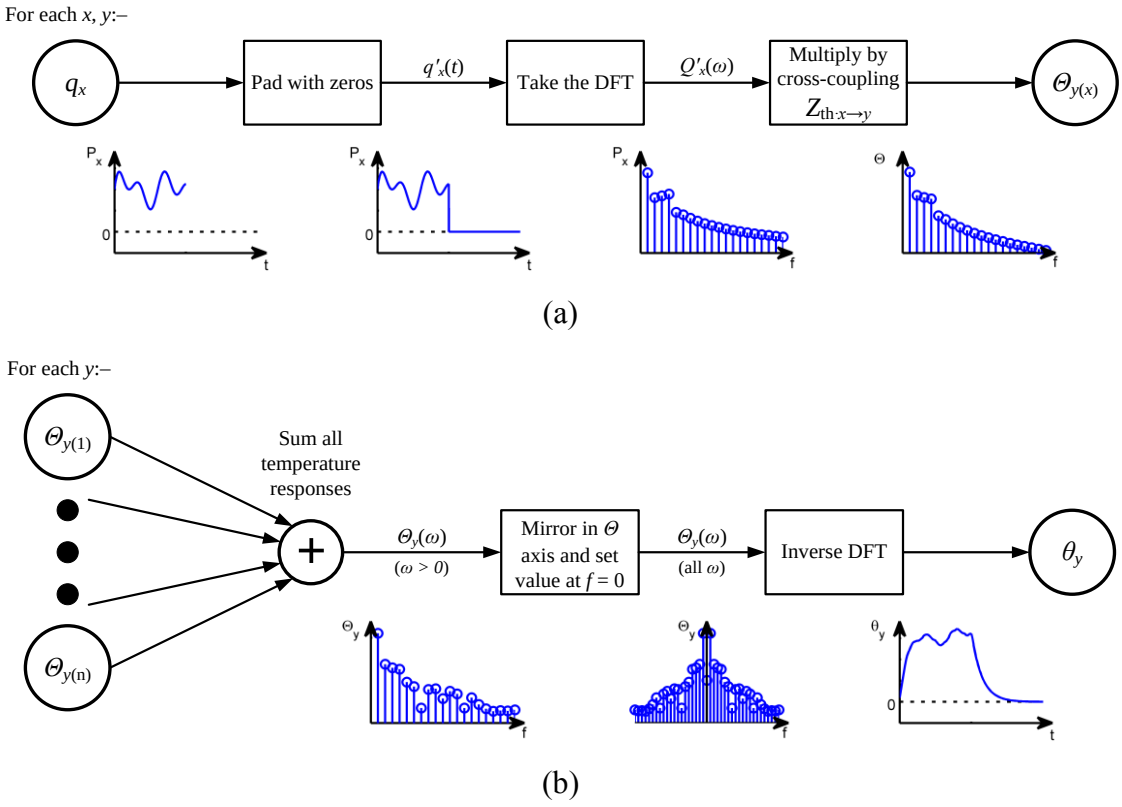

Fig. 6.7 Procedure for calculating the temperature response from cross-coupling data with example waveforms shown. (a) For calculating the frequency domain response at *y* due to individual heat source *x*. (b) For calculating the time domain response at *y* due to all heat sources.

To calculate the temperature at device *y*, the DFT ($\mathcal{F}$) of the padded input power is first taken and multiplied by the cross-coupling between device *y* and each of the other devices at every frequency in the DFT to find $\Theta_{y(x)}$ for each pair. The temperature response is the sum of these temperature DFTs as shown in (6.11). This procedure is

equivalent to a single line of the thermal impedance matrix equation given in (6.1). Because the valid frequency range of the cross-coupling, $Z_{\mathrm{th}\cdot x\to y}$, may be smaller than the frequency range of the input power, $Z_{\mathrm{th}\cdot x\to y}$ is extended to ensure it covers the same range. $Z_{\mathrm{th}\cdot x\to y}(\omega_{\min})$ is therefore used for frequencies in the input power lower than the minimum, $\omega_{\min}$, for which $Z_{\mathrm{th}\cdot x\to y}$ is defined. Similarly, the cross-coupling is assumed to be zero for frequencies higher than the maximum defined frequency, $\omega_{\max}$,. This is expressed mathematically in (6.12).

$$\Theta_{\mathrm{y}}[k] = \sum_{x} \mathcal{F}\big(q'_x(t)\big)[k] \cdot Z'_{\mathrm{th}\cdot x\to y}(\omega) \tag{6.11}$$

$$Z'_{\mathrm{th}\cdot x\to y}(\omega) = \begin{cases} Z_{\mathrm{th}\cdot x\to y}(\omega_{\min}), & 0 < \omega < \omega_{\min} \\ Z_{\mathrm{th}\cdot x\to y}(\omega), & \omega_{\min} \le \omega \le \omega_{\max} \\ 0, & \omega > \omega_{\max} \end{cases} \tag{6.12}$$

Where $k = \omega(n_{\mathrm{s}} - 1)/\omega_{\mathrm{s}}$ is the discrete frequency index, ω_{s} is sampling rate and n_{s} is the number of samples. The above description is for the positive frequencies present in the DFT. Because the temperature waveform produced is real, its DFT at negative frequencies is the complex conjugate mirror of its DFT at positive frequencies [6.11]. The values at negative frequencies in the DFT are therefore calculated according to (6.13).

$$\Theta_y[-k] = \Theta_y^*[k] \tag{6.13}$$

The value at frequency zero ($\Theta_y[0]$) is proportional to the average value of θ_y and therefore the average temperature. This can be seen from the synthesis and analysis equations of the DFT [6.11]. In this thesis, the temperature difference is calculated from ambient. Because the system is initially at ambient, $\Theta_y[0]$ is set such that the initial temperature difference is zero, according to (6.14).

$$\Theta_y[0] = -\sum_{\substack{k=-(n_{\mathrm{s}}-1)/2 \\ k\neq 0}}^{(n_{\mathrm{s}}-1)/2} \Theta_y[k] \tag{6.14}$$

The inverse DFT ($\mathcal{F}^{-1}$) is then taken and the time domain temperature signal obtained, producing the final expression for the temperature response, θ_y, as shown in (6.15).

$$\theta_y(t) = \mathcal{F}^{-1}\left(\Theta_y[k]\right) \tag{6.15}$$

To validate the technique, devices 2 and 4 were set to simultaneously dissipate 95 W square waves at 2 mHz and 0.2 mHz respectively. The temperature responses of all devices were estimated using the procedure described above and are shown in Fig. 6.8(a). In this case, it is assumed that the power input is periodic and that all transients have passed and therefore padding has not been added according to (6.10). The waveforms were also applied experimentally to the example heatsink and the temperature response recorded over one cycle once steady state was reached as shown in Fig. 6.8(b).

The two plots show good agreement with each other. The correlation coefficient, r^2, is greater than 0.99 and the root-mean-square error is less than 2.4 K or 3% of the total range for each waveform, indicating that the cross-coupling method proposed is effective for temperature estimation.

6.5.2 Improved computational efficiency by fitting a digital IIR filter

The direct use of frequency domain characteristics discussed above is computationally intensive since the calculation of forward and inverse discrete Fourier transforms is required. In order to provide a real time prediction of temperature, the whole input power waveform must be analysed and the entire temperature response generated. For a long power profile, the computational power required approaches that of a desktop PC. Ideally, the proposed technique could be implemented on an embedded system installed alongside a power converter to monitor input power and estimate the temperature accordingly. A microcontroller-based system would be particularly suitable as modern power electronics converters and drives typically already contain such a device [6.12].

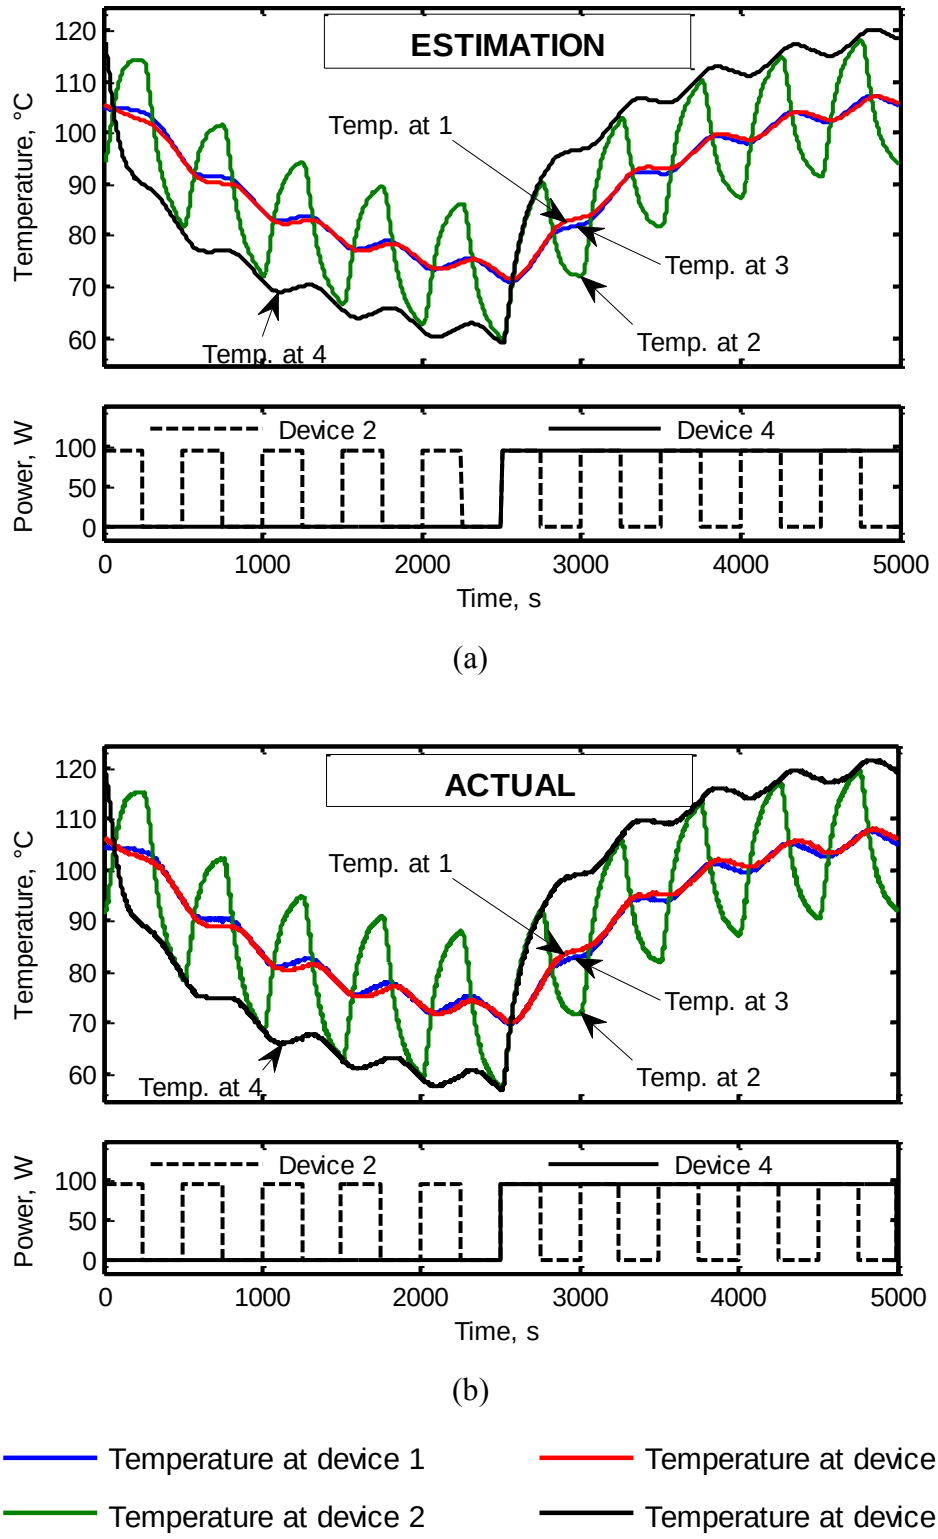

Fig. 6.8 Estimated (a) and measured experimental (b) responses to 0.2 mHz and 2 mHz square wave dissipation

The usual approach to simplifying a thermal system is to model it as an equivalent circuit of resistors, capacitors and current sources, known as a lumped parameter model [6.13]. These models can be constructed directly by calculation based on the geometry of devices and their heatsinks without reference to experimental data. As discussed in Chapter II, lumped parameter models are commonly used for real-time temperature predictions in a number of applications including simulation of buildings and their heating systems [6.14] and electrical machines [6.15]. Their simplicity means they are easy to implement on embedded systems. Lumped parameter models have also been applied to temperature predictions in power electronic systems. Previous literature has discussed the use of Luenberger observers [6.16]; however they are limited because an equivalent circuit must be formulated and its component values populated before it can be used to make predictions.

Alternatively, the thermal impedance of a system can be modelled using a simplified electrical equivalent circuit; for example, by fitting a Foster network to a step response [6.2]. To embed a temperature predictor based on such a model into a microcontroller, the equivalent circuit must undergo transient analysis. For a given model, a set of finite difference equations is produced by applying the trapezium rule to differential equations for each capacitor [6.17]. These equations are then used iteratively to generate a temperature estimation for an arbitrary input. If a lumped parameter model is used, the equations must be generated from the fitted equivalent circuit. In the method proposed in this Chapter, the difference equations are directly generated from cross-coupling techniques, removing inaccuracy from step-response characterisation and avoiding the need for transient circuit analysis.

To achieve this, a digital filter is fitted directly to the cross-coupling characteristics. A digital filter is a series of additions and multiplications performed on the current and previous samples of the input power waveform. Filters are described by two vectors, $\boldsymbol{a}$ and $\boldsymbol{b}$. From [6.11], the output of the filter may be calculated as shown in (6.16). From this equation, simple difference equations can be generated directly [6.11].

$$\theta[i] = \sum_{\lambda=1}^{G} a_\lambda \theta[i-\lambda] + \sum_{\lambda=0}^{H} b_\lambda q[i-\lambda] \tag{6.16}$$

where i is the sample number and G and H are the lengths of vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ respectively. The filter is said to be causal because no information from future values is required to calculate the present output value. This is necessary for real-time estimation. The frequency response of the system is

$$Z_{\mathrm{th}}(z) = \frac{\sum_{\lambda=0}^{G} b_\lambda z^{-\lambda}}{1 + \sum_{\lambda=1}^{H} a_\lambda z^{-\lambda}}; \quad z = e^{-\mathrm{j}2\pi\left(\frac{\omega}{\omega_\mathrm{s}}\right)} \tag{6.17}$$

where z is the $\mathcal{Z}$-transform operator variable, j is the imaginary unit, ω_s is the sampling frequency and ω is angular frequency.

There are two types of filter available. The finite impulse response (FIR) filter, where $a_\lambda = 0$ for $\lambda \geq 1$, is the simpler but the response to a power input is of finite length since after H zero-input samples the output falls to zero. This type of filter is very easy to implement as only a memory of the previous H samples is required. Its usefulness is limited, however, by the length of the filter commonly necessary to model a real system at a sensible sampling rate and by the group delay that results from a long filter. FIR design techniques also tend to have linear phase [6.11], which is not generally true of cross-coupling.

Alternatively, an infinite impulse response (IIR) filter may be fitted. This filter has arbitrary length $\boldsymbol{a}$ and $\boldsymbol{b}$ vectors which may take any real value. Because the output is fed back, an IIR filter's response to an input waveform is typically infinite in length. In addition, IIR filters are more computationally efficient since they typically require shorter vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ than FIR filters, meaning fewer previous samples have to be stored [6.11].

In this Chapter, IIR filters are therefore used to represent the cross-coupling frequency responses identified. Fitting is based on the MATLAB `invfreqz` function [6.18]. This uses the Gauss-Newton iterative method to find vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ which minimise the mean-square-error between the desired frequency response and the curve fit in (6.18) to form a stable IIR filter.

$$\sum_{k=0}^{n_s-1}\left|Z_{\text{th}}(\omega(k))-\frac{B\big(\varphi(k)\big)}{A\big(\varphi(k)\big)}\right|^2 \tag{6.18}$$

$$B(\varphi)=\sum_{\lambda=0}^{G} b_\lambda e^{-\mathrm{j}\lambda\varphi}\,; \qquad A(\varphi)=1-\sum_{\lambda=1}^{H} a_\lambda e^{-\mathrm{j}\lambda\varphi} \tag{6.19}$$

$$\varphi(k)=\frac{\pi}{\omega_{\text{max}}}\omega(k) \tag{6.20}$$

$$\omega(k)=\omega_{\text{min}}\left(\frac{\omega_{\text{max}}}{\omega_{\text{min}}}\right)^{\frac{k}{n_s-1}} \tag{6.21}$$

where A and B are functions of $\boldsymbol{a}$ and $\boldsymbol{b}$, respectively, given in (6.19), k is the frequency index and n_s is the number of samples over which Z_{th} is calculated. For correct weighting of the different frequencies in the minimisation of function (6.18), a logarithmic spread of frequencies is used between ω_{min} and ω_{max}, which are the minimum and maximum frequencies over which Z_{th} is defined, as shown in (6.21). $\omega(k)$ is the real angular frequency for which $Z_{\text{th}}(\omega)$ is defined. Since the Fourier transform is periodic around the unit circle, $\varphi(\omega)$ is the scaled angular frequency in the range $0 \leq \varphi \leq \pi$. For both ω and φ, negative frequencies are omitted due to symmetry.

This method of minimisation was selected since it considers the complex values of the cross-coupling, as opposed to the magnitude alone. By doing this, both phase and magnitude are fitted closely. Vectors of length 4 and 7 were chosen for $\boldsymbol{a}$ and $\boldsymbol{b}$ respectively as these values were the smallest that yielded good results. Although the characterisation is a computationally intensive procedure, it is performed offline only once.

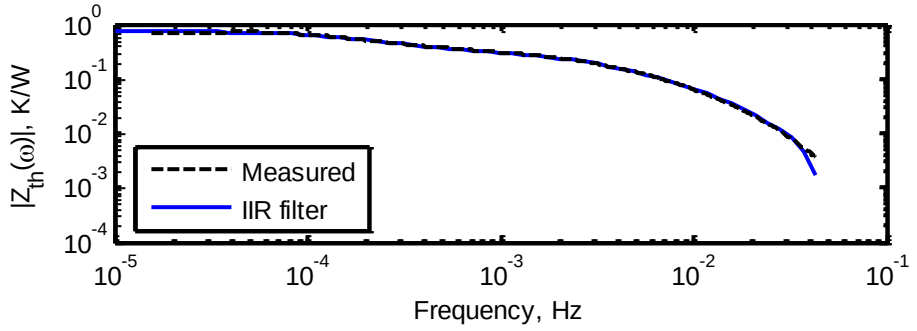

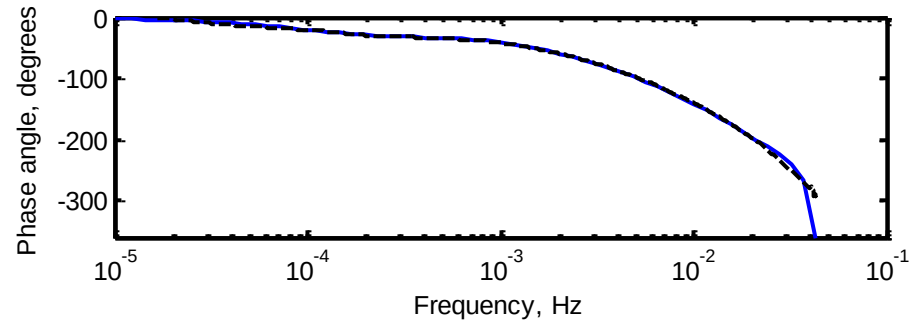

(a)

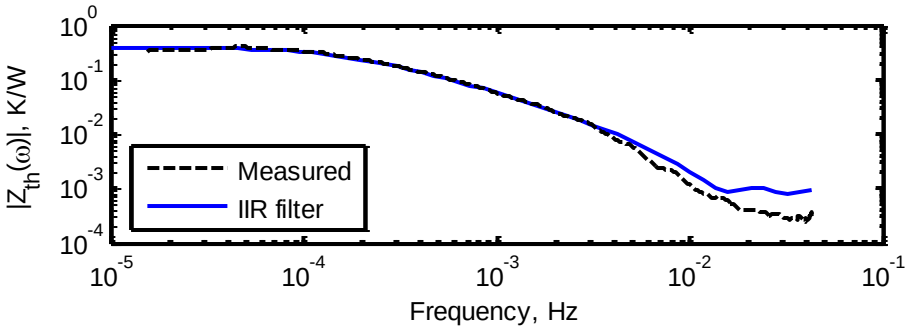

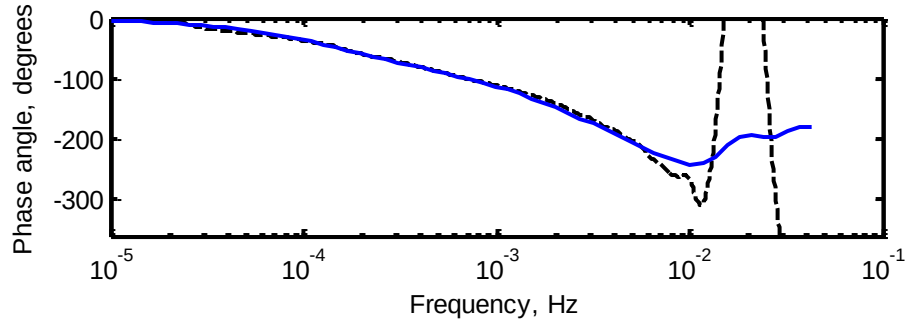

(b)

Fig. 6.9 Typical Bode plots for measured and fitted IIR filter spectra for (a) auto-coupling $Z_{\mathrm{th}\cdot 1\to 1}$ and (b) cross-coupling $Z_{\mathrm{th}\cdot 1\to 2}$.

Fig. 6.9(a) and (b) show the typical Bode plots of measured and fitted IIR filter spectra for typical auto-coupling ($Z_{\mathrm{th}\cdot x\to x}$) and cross-coupling ($Z_{\mathrm{th}\cdot x\to y}$) characteristics respectively. The agreement between the IIR filter and the measured spectrum is very

good in the case of the auto-coupling spectrum due to its higher impedance and therefore greater noise immunity. The cross-coupling spectrum shows reduced agreement at high frequency, particularly in phase, caused by the effect of noise in the original measurement.

Noise in the system results in a noise floor, $|Z_{\text{noise}}|$, in the measurement of thermal impedance under which neither the magnitude nor phase can be measured accurately. For the experimental set-up in this work, the spectral density of the noise was 4.3×10^{-4} K^2/Hz. Since three repeated sequences of the high frequency PRBS are averaged, the effective noise is reduced by a factor of three to 1.4×10^{-4} K^2/Hz. It was shown in Chapters IV and V that the spectral density for a unipolar PRBS power waveform with respect to frequency, f, is

$$\Phi_{\text{PRBS}}(f) = \frac{Q^2}{4f_{\text{p}}} \cdot \frac{N+1}{N}\left[\frac{\sin f\pi/f_{\text{p}}}{f\pi/f_{\text{p}}}\right] \tag{6.22}$$

Where Q is the PRBS amplitude and f_{p} is the PRBS clock frequency. At the highest frequency in the useable bandwidth, $f = 43$ mHz, the PRBS frequency is $f_p = 0.1$ Hz and the spectral density is therefore $\Phi_{\text{PRBS}}(43 \times 10^{-3}\text{ Hz}) = 1.6 \times 10^{5}$ W^2/Hz. Because the Fourier transform is proportional to the square root of the power spectrum, the impedance measured due to noise alone is, from (6.9),

$$|Z_{\text{noise}}| = \frac{\sqrt{\Phi_{\text{noise}}}}{\sqrt{\Phi_{\text{PRBS}}}} = 9.4 \times 10^{-5}\text{ K/W} \tag{6.23}$$

For measured impedances approaching this value, noise adversely affects results. For measurement of phase, the effect is more pronounced since the phase angle of noise is random whereas the amplitude of noise is limited to around $|Z_{\text{noise}}|$. This is the reason for the error seen at low values of thermal impedance in Fig. 6.9(b). However, since the impedance is low at noise-affected frequencies, the effect of the error on the resulting estimations is negligible. A full explanation of system identification using PRBS in noisy conditions was given in Chapter IV.

The temperature response may be calculated using (6.16) or, alternatively, with the more efficient direct form II implementation given in [6.11]. The sampling frequency of the filter is set by the Nyquist limit of the cross-coupling according to (6.24).

$$\omega_s = 2\omega_{max} \tag{6.24}$$

The agreement between temperature estimations made using the direct cross-coupling procedure and estimations made from a fitted IIR filter is excellent. Fig. 6.10(a) and (b) show estimated temperature responses for an arbitrary input power waveform (Fig. 6.10(e)) using each of these techniques respectively. For clarity, each of these figures consists of a view of the entire profile on the left and the magnified peak region (where differences are greatest) on the right. The input power waveform is the EU driving cycle [6.19] where the power demand is variable and becomes larger over time, before dropping to zero. This waveform was selected because it gives a representative cycling power dissipation in a converter which may be intended for electric vehicle use. The temperature response is therefore a steady but perturbed increase followed by an exponential decay. The root-mean-square difference between the direct cross-coupling estimations and the IIR filter estimations is 0.13 K, which is negligible and demonstrates their equivalency. These results are discussed in more detail in section 6.6.

6.5.3 Real temperature feedback

The accuracy of the temperature response can be improved with feedback from a single temperature sensor. By measuring the temperature at one point for which the cross-coupling characteristics are known, the difference between the calculated and practical temperature can be calculated. Assuming that this difference is consistent between all the points for which cross-coupling is known, the temperature at these points can be accurately calculated according to (6.25).

$$\Theta'_y = \Theta_{estimation,y} + \left(\Theta_{real,r} - \Theta_{estimation,r}\right) \tag{6.25}$$

Where $\Theta_{estimation,r}$ and $\Theta_{real,r}$ are estimation and feedback temperatures at point r respectively. Θ'_y is the feedback-corrected temperature at point y. This technique can be applied to the IIR filter method described above, and for the example heatsink

arrangement presented in this Chapter. For this experiment, the reading is corrected using the temperature data for device 3, which is selected because of its central location. The effectiveness of this technique is discussed in section 6.6 and shown in Fig. 6.10(c).

6.5.4 Temperature response extrapolation

Predicting the temperature response to future power input is a useful feature of a temperature control system [6.16]. For example, if it appears likely that the thermal limits are to be exceeded in the future, actions can be taken to avoid this, such as limiting power, increasing cooling or warning the user. A significant advantage of the IIR filter estimation technique compared to the use of the DFT on cross-coupling characteristics directly is the ease with which prediction can be made for simulated input. By duplicating the filter's internal state and inputting simulated input power (perhaps by duplicating recent power dissipation), the temperature at any point in the future can be predicted efficiently and accurately. Since sample rates of less than 10 Hz are typically required for the IIR filter, predicting one minute ahead will require less than 600 filter iterations. Calculations of the necessary complexity could be performed by a typical microcontroller in a few milliseconds.

6.6 Comparison of proposed techniques

To demonstrate the techniques, an arbitrary power waveform is used. A power cycle, based on a vehicle running the European driving cycle waveform [6.19], is selected as it is well known, is of an appropriate length for the time constants of the heatsink arrangement used in this work and is sufficiently arbitrary to permit the easy evaluation of the techniques. For simplicity, the power handling requirement is assumed to be proportional to the vehicle velocity described in the driving cycle.

Fig. 6.10(a-c) shows the estimation temperature response of the example heatsink arrangement to the input power waveform in Fig. 6.10(e). To validate the results, this input power waveform was applied practically to the example system and the resulting temperature response in shown in Fig. 6.10(d). In each case, the same power was applied to each device. The left-hand side of each figure shows the response over an extended period while the right-hand side shows detail around the peak temperature point.

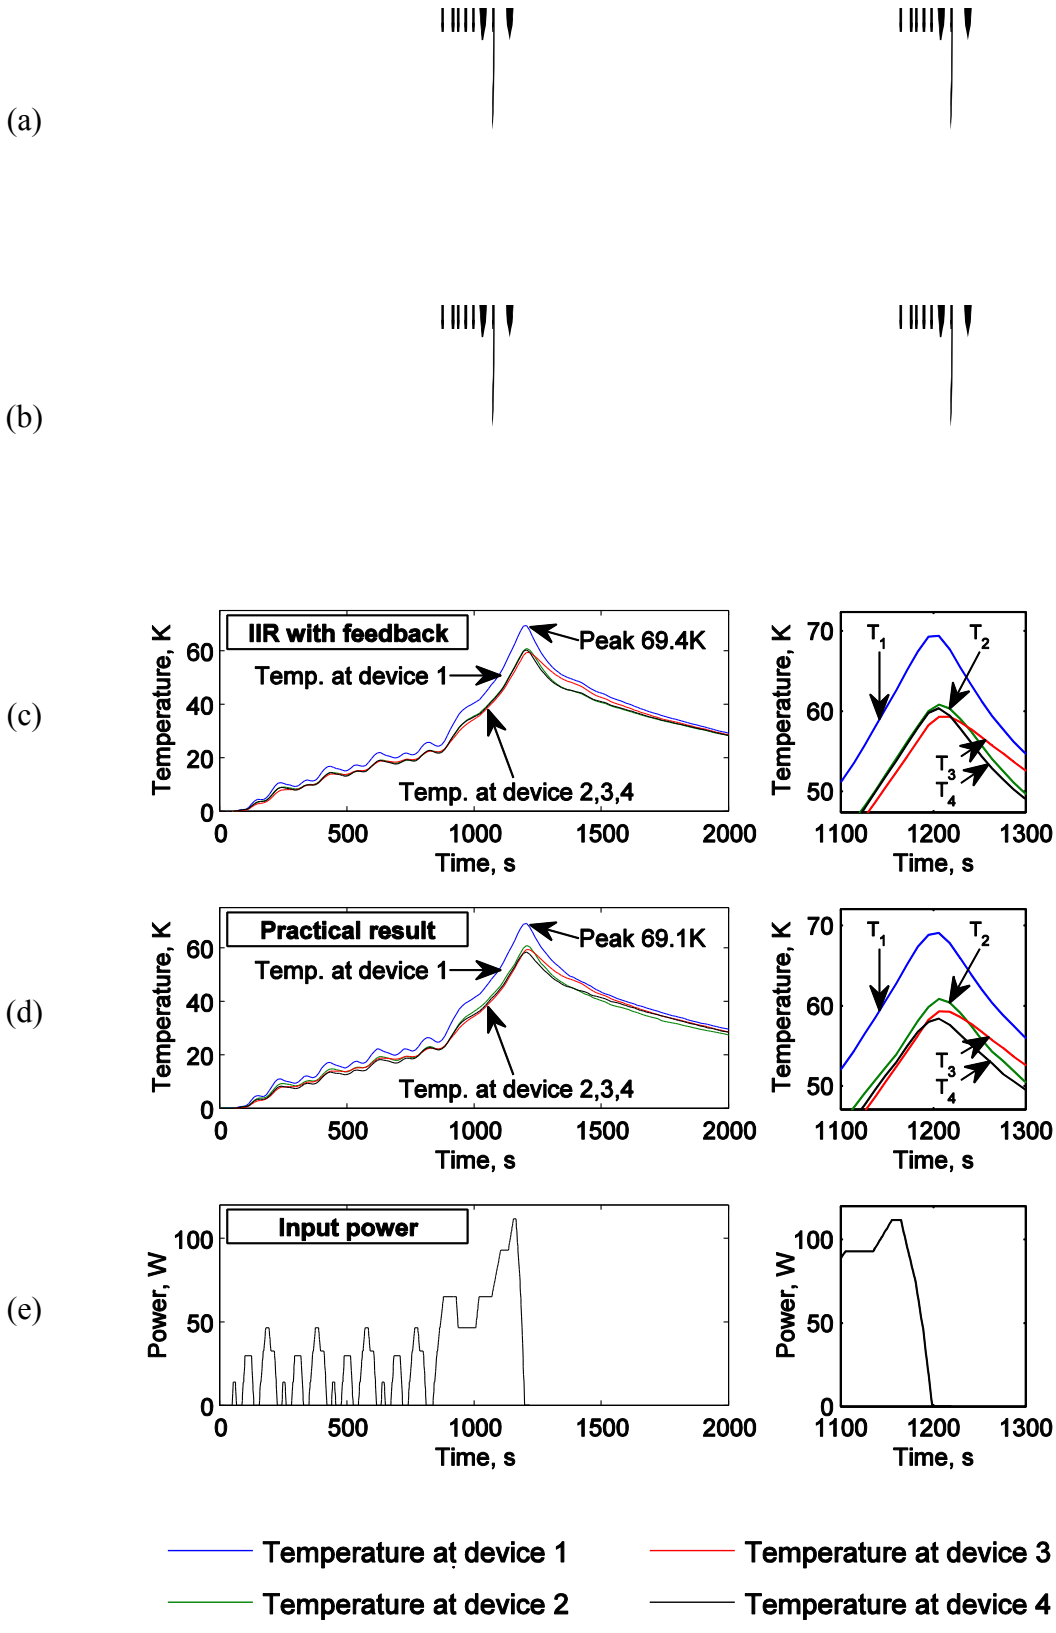

Fig. 6.10 Temperature responses relative to ambient for an arbitrary power waveform (ambient temperature 24°C). (a) Response from frequency domain cross-coupling. (b) Response from fitted IIR filter. (c) Response from feedback-corrected IIR filter (with feedback from θ_3). (d) Practical result. (e) Input power.

There is good correlation between all estimated responses and the practical response. The shapes of the curves show excellent agreement with the practical result. This is most evident around the temperature peak where only a minor difference in the relative positions of T_2 and T_4 can be seen. This error is likely to relate to a minor change in the ambient conditions between the characterisation of the system and the practical result measurement. The error is apparent in all three estimated responses because each result is derived from the same cross-coupling characteristics and input power waveforms. There is little difference in waveform shape produced between the three estimation methods, demonstrating their equivalency.

	Root-mean-square error from practical result (K)			
	At T_1	At T_2	At T_3	At T_4
Cross-coupling	4.0	3.7	3.7	3.8
Fitted IIR filter	3.8	3.4	3.7	3.7
Feedback-corrected IIR	0.47	0.54	0[a]	0.51

[a] by definition

Table 6.1 Accuracy of estimations using different methods compared to practical result

The root-mean-square error between each estimation method and the practical result is given in Table 6.1. Cross-coupling and fitted IIR filter results both show around a 4 K error in each temperature reading, which corresponds to 6% of the peak temperature. This error is significantly reduced by feedback from the temperature of device 3 to around 0.5 K, which corresponds to 1% of the peak temperature.

The computational requirements of the estimation techniques are compared in Table 6.2. These figures are for a sequence of 174 time steps where a further 334 are required for settling time padding in the direct cross-coupling calculation, making the sequence length 512 for efficient Fourier transform calculation. The direct cross-coupling procedure is significantly more complex than using a fitted IIR filter since the required number of multiplication-additions to calculate the entire temperature response is an order of magnitude higher. In addition complex number arithmetic is required since

most of these operations are part of a discrete Fourier transform, further increasing computational complexity.

	To calculate one time step		To calculate entire sequence	
	Multiplication-additions	Variables stored	Multiplication-additions	Variables stored
Cross-coupling	10 750[b]	10 441 024	10 750[b]	10 441 024
Fitted IIR filter[c]	11	7	1 914	7

[b] Assuming a DFT requires $n_s \log_2 n_s$ multiplication-additions for an n_s-length sequence

[c] with or without feedback correction

Table 6.2 Microprocessor resources required per cross-coupling characteristic

The difference is more significant if only a single time step is to be calculated. In this case, the direct cross-coupling procedure still requires that the whole sequence is processed which means there is no reduction in required resources. However, for the IIR filter, only a single iteration is required. The result is three orders of magnitude difference in calculation complexity.

The difference in required storage also demonstrates an advantage of the IIR filter, which requires only the 7 element filter state vector to be stored per cross-coupling characteristic for the direct form II implementation [6.11], whereas the direct cross-coupling method requires the length of both the padded sequence and its DFT to be stored simultaneously during the DFT. This comparison neglects the need to store the cross-coupling characteristics, which are also smaller in the case of the IIR filter, in that case consisting only of 10 variables in vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ combined.

To relate these resource requirements to a practical situation, the timing requirements for a standard microcontroller were calculated. For the direct cross-coupling calculation, processor time is dominated by the calculation of the discrete Fourier transforms which requires 178 ms for a forward transform and 196 ms for an inverse transform when implemented using fixed-point arithmetic on an ATmega1284P clocked at 10 MHz.

Each calculation therefore takes 374 ms. By comparison, 174 iterations of the IIR filter requires 416 µs with or without feedback. These figures relate to calculating the response per cross-coupling characteristic. For the arrangement in this Chapter, four temperatures are calculated from four power inputs, meaning that more than four times these resources are required.

From this data, it is clear that the fitted IIR filter is significantly more computationally efficient than using the direct cross-coupling procedure. If the direct cross-coupling procedure is used, at least half of processor time is required for temperature estimation. Using the proposed filter technique releases this processor time for other tasks, such as temperature prediction under several possible future loads or electrical management of the system. With the addition of feedback temperature correction, a very close estimation of the temperature response can be made.

6.7 Chapter conclusions

A technique to estimate the temperature response of a multi-element thermal system has been presented. Initially, a characterisation was performed using a PRBS technique, calculating the complex frequency domain cross-coupling. Subsequently, the characteristics can be used to estimate the device temperature responses for a known input power waveform. The resulting estimation generated by the presented method shows good agreement with the practical system used for evaluation. To reduce the computational complexity of the initial method, a digital IIR filter is fitted to each cross-coupling characteristic. An excellent correlation is demonstrated which produces a near-identical temperature response compared to the direct cross-coupling procedure. Further temperature responses can be extrapolated by inputting expected future values to the filter. Experimental validation on the practical system shows good agreement between IIR filter estimation and practical results. This agreement can be substantially improved by taking feedback from a reference temperature.

6.8 References

[6.1] T. Bruckner and S. Bernet, '*Estimation and Measurement of Junction Temperatures in a Three-Level Voltage Source Converter*', *IEEE Transactions on Power Electronics,* vol. 22, pp. 3-12, 2007.

[6.2] G. C. James, V. Pickert and M. Cade, '*A thermal model for a multichip device with changing cooling conditions*', *in Proc* Power Electronics, Machines and Drives (PEMD), 2008, pp. 310-314.

[6.3] H. Wang, S. X. D. Tan, L. Guangdeng, R. Quintanilla and A. Gupta, '*Full-chip runtime error-tolerant thermal estimation and prediction for practical thermal management*', *in Proc* Computer-Aided Design (ICCAD), 2011, pp. 716-723.

[6.4] P. L. Evans, A. Castellazzi and C. M. Johnson, '*Automated Fast Extraction of Compact Thermal Models for Power Electronic Modules*', *IEEE Transactions on Power Electronics,* vol. 28, pp. 4791-4802, 2013.

[6.5] X. Wang, K. Ma and Y. Wang, '*Adaptive Power Control with Online Model Estimation for Chip Multiprocessors*', *IEEE Transactions on Parallel and Distributed Systems,* vol. 22, pp. 1681-1696, 2011.

[6.6] M. Musallam and C. M. Johnson, '*Real-Time Compact Thermal Models for Health Management of Power Electronics*', *IEEE Transactions on Power Electronics,* vol. 25, pp. 1416-1425, 2010.

[6.7] J. W. Kolar, U. Drofenik, J. Biela, M. L. Heldwein, H. Ertl, T. Friedli*, et al.*, '*PWM converter power density barriers*', *in Proc* Power Conversion Conference, Nagoya, 2007, pp. 1656-1676.

[6.8] M. Yilmaz and P. T. Krein, '*Review of integrated charging methods for plug-in electric and hybrid vehicles*', *in Proc* Vehicular Electronics and Safety (ICVES), 2012, pp. 346-351.

[6.9] eFunda. (2015, 4 March). *General information on element aluminium*. Available: http://www.efunda.com/materials/elements/element_info.cfm?Element_ID=Al

[6.10] N. V. Queipo, J. A. C. Humphrey and A. Ortega, '*Multiobjective optimal placement of convectively cooled electronic components on printed wiring boards*', *IEEE Transactions on Components, Packaging, and Manufacturing Technology, Part A,* vol. 21, pp. 142-153, 1998.

[6.11] A. V. S. Oppenheim, R.W., *Discrete-Time Signal Processing*. New Jersey: Prentice Hall, 1999.

[6.12] S. Sirisukprasert, L. Jih-Sheng and L. Tian-Hua, '*Optimum harmonic reduction with a wide range of modulation indexes for multilevel converters*', *IEEE Transactions on Industrial Electronics,* vol. 49, pp. 875-881, 2002.

[6.13] P. R. Strickland, '*The Thermal Equivalent Circuit of a Transistor*', *IBM Journal of Research and Development,* vol. 3, pp. 35-45, 1959.

[6.14] M. M. Gouda, S. Danaher and C. P. Underwood, '*Quasi-adaptive fuzzy heating control of solar buildings*', *Building and Environment,* vol. 41, pp. 1881-1891, 2006.

[6.15] A. Ridge, R. McMahon and H. P. Kelly, '*Detailed thermal modelling of a tubular linear machine for marine renewable generation*', *in Proc* Industrial Technology (ICIT), 2013, pp. 1886-1891.

[6.16] J. D. Holmes, M. P. Foster and D. A. Stone, '*System-wide Temperature Estimation for IMS based Power Electronics Circuits*', *in Proc* International Conference on Power Electronics and Drive Systems, 2009, pp. 1081-1084.

[6.17] R. C. Dorf, *The Electrical Engineering Handbook*. Florida: CRC Press, 1997.

[6.18] MathWorks. (2013, 8 July). *MATLAB 2013a documentation: invfreqz*.

[6.19] *EEC Journal Officiel No. C81*, 1990.

Chapter VII

Estimating temperatures in a dynamically cooled multiple device system in order to control cooling and thereby reduce thermal cycling

This Chapter extends the estimation technique described in Chapter VI to a thermally-coupled, multiple device system subject to dynamic cooling. Using a demonstrator system, the thermal transfer impedance between pairs of devices is determined in the frequency domain for a quantised range of active cooling levels using a technique based on pseudorandom binary sequences. For each cooling level and pair of devices, a digital IIR filter is produced which can be used to directly estimate temperature from device input power. When the cooling level changes, the filters in use are substituted and the internal states of the old filters are converted for use in the new filter. Two methods for filter state conversion are developed—a computationally efficient method which is suited to temperature estimation in a single device system, and a more general method which requires increased memory and processing capacity. To demonstrate an application of the technique, the single device estimator is used as part of a cooling control system to reduce the thermal cycling without requiring online temperature measurement. Results show that all device temperatures can be estimated with low error using a system which is suitable for integration on an embedded processor, and that thermal cycling can be reduced when controlled cooling is implemented.

7.1 Introduction

Following the development of improved measurement and modelling techniques in Chapters III to V, Chapter VI proposed a temperature estimation technique for a multiple device system. It was proposed that the technique could be used to monitor device temperatures thereby preventing overheating and consequently reducing system size. The technique relied on an offline frequency-domain characterisation procedure to generate a thermal model. The model was then implemented as a set of infinite impulse response (IIR) digital filters, with two fitted for each pair of devices. This solution is

limited to cases of static cooling because changes in cooling level affect the thermal model. In this Chapter, the technique is extended to cover temperature estimation in a dynamically cooled system.

As well as preventing over-temperature conditions, reliability considerations must be taken into account when designing a system because damage to power converters accumulates due to multiple heating and cooling events until system failure results [7.1]. This is termed thermal cycling. It is important to recall that a simple method of thermal management, often used in cheaper systems, is to incorporate large heatsinks and active cooling systems to maintain device temperatures well below their rated maxima under all load conditions [7.2]. This solution suffers from two problems. First, the system and cooling must be sized for the worst-case condition which only occurs when the system is under full load (or fault conditions) in high ambient temperature conditions for a prolonged operating periods. However, this is an unusual situation, and, under practical conditions, the system operates significantly below the worst case. Secondly, where active cooling is not controlled, temperatures rapidly increase and decrease with system load which can lead to significant thermal cycling [7.3]. Controlled and managed active cooling is therefore required.

Before cooling control can be implemented, a method of temperature estimation under dynamic cooling is required. Many authors use an analytical approach to address the problem [7.4][7.5]. This approach is limited by the necessity to determine the physical parameters of the derived equations, which is especially difficult in cases where thermal cross-coupling between devices is to be considered. In these cases, a numerical approach is preferable since complex parameters need not be determined. Possible approaches include the offline modelling of a system using finite element analysis (FEA) and computational fluid dynamics (CFD) [7.6]. Models developed from FEA can later be simplified and incorporated into a control system. The FEA approach is limited by the requirement for the model to precisely match the practical system. In addition, numerical modelling of turbulent flow is difficult [7.7]. It is shown by Rodgers *et al* [7.8] that more practical thermal analysis is preferable in many cases. In this Chapter, an experiment-based characterisation of a power electronic system under controlled-cooling is proposed.

Whereas Chapter VI proposed a method to estimate the temperature over several devices coupled on a statically cooled heatsink by characterising the cross-coupling between them using a simple pseudorandom binary sequence (PRBS) technique, this Chapter uses the improved mixed PRBS technique described in Chapter V. The mixed technique is applied to each power device in turn and the resulting temperature response at all devices is measured. The cross-coupling between each pair of devices over a band of frequencies is determined and used to estimate the temperature response to a given power input and cooling level using a novel technique which results in accurate prediction.

Using the output of the estimator, the temperature rise in electronic devices is controlled to reduce cycling while ensuring devices stay within maximum ratings. Because this means that the over-temperature conditions can be more readily estimated, heatsinking requirements are reduced. This fast-response controlled active cooling prevents temperatures from exceeding ratings when there is a rapid load increase on a system operating close to its thermal limit. This means that device temperatures can be held close to their thermal limit, reducing the temperature safety margin required. The addition of an active cooling element (such as a fan or blower) increases cost and size, but this problem is mitigated by the consequential reduction in heatsink size. Due to the weight of heatsinks, adding active cooling is particularly relevant in transport applications where reduced weight is paramount. Active cooling which is intelligently controlled has other advantages. For example, temperature variations due to varying load can be reduced, lessening fatigue due to thermal cycling.

7.2 Thermal system characterisation

7.2.1 Cross-coupling

In Chapter VI, it was shown that a linear thermal system can be represented by the thermal transfer impedance, or *cross-coupling*, between its heat sources and points where the temperature is to be measured. In a power electronics system, an engineer is normally concerned with only the few points whose temperatures are limited by device physics and construction, and are likely to exceed rated levels. Typically, these are the electronic power devices which dissipate high power and are damaged by overheating. To model the thermal cross-coupling, the problem is considered using Norton's theorem

and superposition theory. The relationship between each dissipating device and each temperature of interest is approximated using the equivalent electrical circuit in Fig. 7.1.

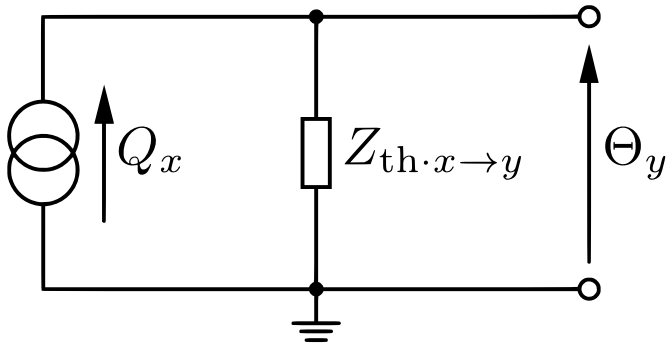

Fig. 7.1 Norton equivalent electrical circuit between heat dissipation point x and temperature measurement point y

Q_x is used to denote the power dissipated at point x, Θ_y for the temperature at point y due to the dissipation at point x and $Z_{\mathrm{th}\cdot x\to y}$ for the cross-coupling between points x and y. To simplify the mathematics, these quantities are used in the frequency domain which is signified using upper case symbols (lower case symbols are used for the time domain). From Ohm's law and superposition theory, the temperature at any point y due to the power dissipations at all points in the system can be expressed in matrix notation as in (7.1) and (7.2). It is possible to use simple point-by-point multiplication in these calculations because the mathematics is carried out in the frequency domain.

$$\boldsymbol{\Theta}(\mathrm{j}\omega) = \boldsymbol{Z}_{\mathbf{th}}(\mathrm{j}\omega) \cdot \boldsymbol{Q}(\mathrm{j}\omega) \tag{7.1}$$

$$\boldsymbol{Z}_{\mathbf{th}} = \begin{pmatrix} Z_{\mathrm{th}\cdot 1\to 1} & \cdots & Z_{\mathrm{th}\cdot n\to 1} \\ \vdots & \ddots & \vdots \\ Z_{\mathrm{th}\cdot 1\to m} & \cdots & Z_{\mathrm{th}\cdot n\to m} \end{pmatrix}; \quad \boldsymbol{\Theta} = \begin{pmatrix} \Theta_1 \\ \vdots \\ \Theta_m \end{pmatrix}; \quad \boldsymbol{Q} = \begin{pmatrix} Q_1 \\ \vdots \\ Q_n \end{pmatrix} \tag{7.2}$$

Where $\boldsymbol{\Theta}(\mathrm{j}\omega)$ and $\boldsymbol{Q}(\mathrm{j}\omega)$ are column vectors of temperature and powers at each relevant point, respectively, $\boldsymbol{Z}_{\mathbf{th}}(\mathrm{j}\omega)$ is the matrix of all the cross-coupling impedances between the relevant points and n and m are the number of power dissipation points and relevant temperatures respectively.

The effect that the intensity of active cooling has on the temperature response is also considered. In our example, active cooling from a blower with a controllable angular velocity, δ, is used. To demonstrate the techniques presented in this Chapter, a test system of D2PAK devices mounted on an insulated metal substrate (IMS) board is constructed. The power dissipation in each device is controlled by an external circuit and a K-type thermocouple is attached to the tab of each device to monitor the temperature. A speed-controlled blower is aimed to blow over all devices. An external computer controls and monitors all elements. The arrangement is shown in Fig. 7.2.

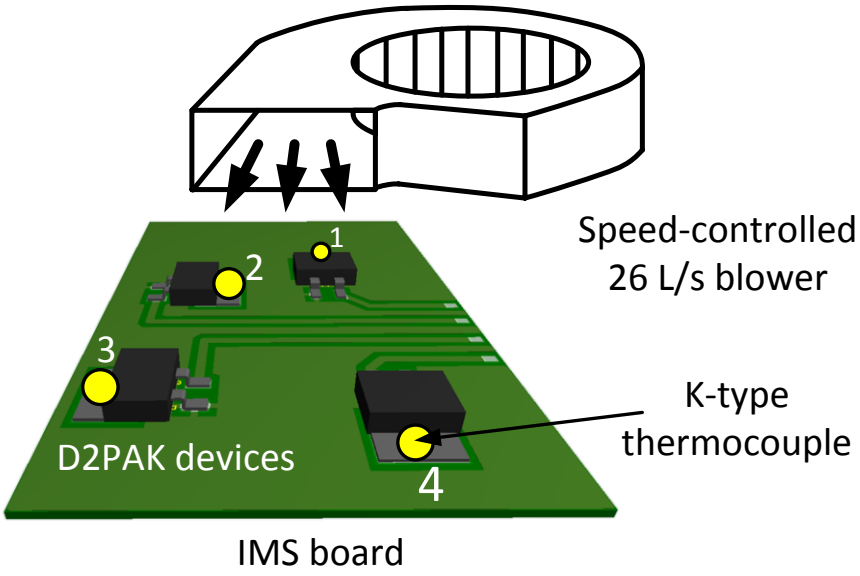

Fig. 7.2 Experimental arrangement

When active cooling is present, the cross-coupling matrix $\boldsymbol{Z}_{\mathbf{th}}$ is not only a function of frequency but is also a function of the level of active cooling—blower speed in this case. The matrix is therefore denoted $\boldsymbol{Z}_{\mathbf{th}}(\mathrm{j}\omega, \delta)$. Determining the elements of this matrix over all blower speeds and frequencies produces a complete model of the system.

7.2.2 Generation of cross-coupling characteristics using pseudorandom binary sequences

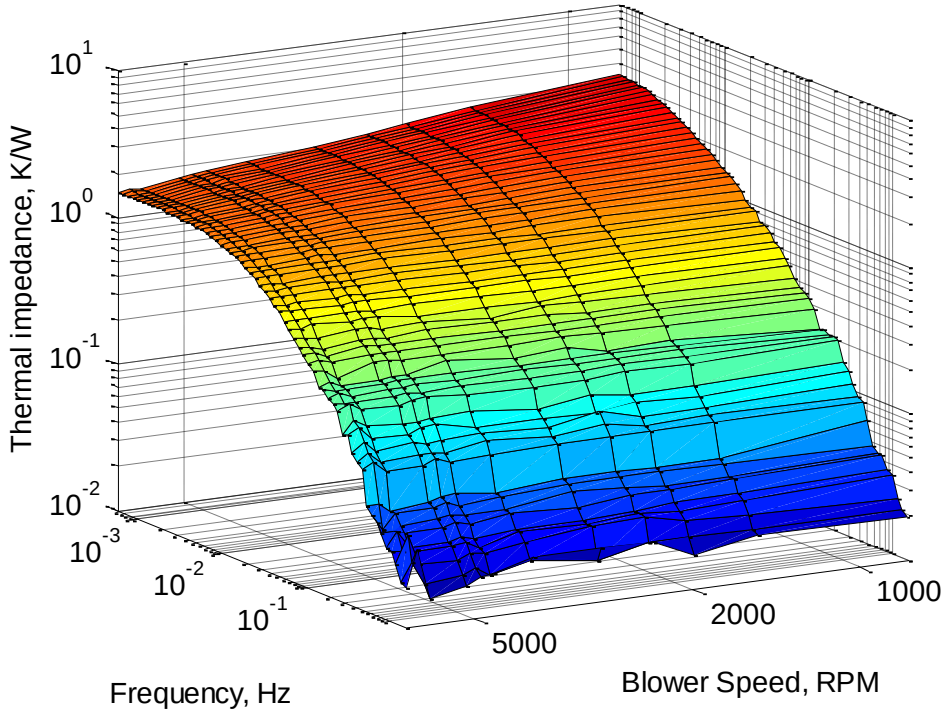

(a)

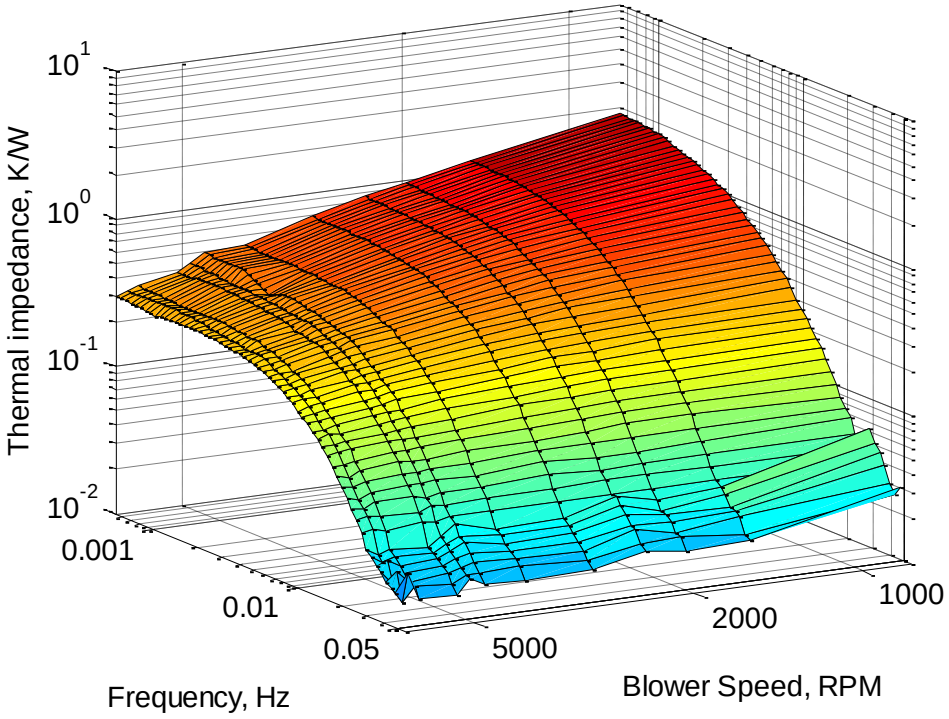

(b)

Fig. 7.3 Three dimensional Bode plots of the auto-coupling of device 1 over the available range of blower speeds and frequencies. Note that all axes are logarithmic. (a) $Z_{\mathrm{th}\cdot 1\to 1}$. (b) $Z_{\mathrm{th}\cdot 3\to 1}$.

The improved PRBS technique proposed in Chapter V is used to characterise the cross-coupling between devices. Two synchronised 7-bit PRBS signals, clocked at 64 mHz and 3.5 Hz, respectively, are mixed together using the *AND* operator. This results in a test signal which generates a valid frequency response between 0.50 mHz and 1.5 Hz. Characterisation was performed over a range of blower speeds between 0 and 6600 revolutions per minute (RPM). The results at each blower speed can be presented in a Bode plot. An extra dimension is added to the graph to show the effect of blower speed. Fig. 7.3 therefore shows a three-dimensional (3-D) Bode plot of $Z_{\mathrm{th}\cdot 1\to 1}(\mathrm{j}\omega,\delta)$ and $Z_{\mathrm{th}\cdot 3\to 1}(\mathrm{j}\omega,\delta)$.

Similar 3-D Bode plots can be produced for the cross-coupling between all combinations of devices. On the logarithmic scales shown, there is a linear relationship between blower speed and thermal impedance, particularly at higher and lower frequencies. In mid-range frequencies, there is some deviation from linearity due to the corner frequency of the impedance changing with blower speed. By using logarithmic interpolation, the full standard Bode plot at any blower speed can be estimated from a characterisation with a limited number of blower speeds. Frequency response curves for each blower speed are subsequently used to generate empirical models of the system, as described in section 7.3.2.

7.3 Estimation of the temperature response

7.3.1 Estimation using cross-coupling frequency response directly

The temperature response at a single device due to dissipation in another device can be directly calculated in the frequency domain using (7.3).

$$\Theta_y(\mathrm{j}\omega,\delta) = Z_{\mathrm{th}\cdot x\to y}(\mathrm{j}\omega,\delta)\cdot Q_x(\mathrm{j}\omega) \tag{7.3}$$

This direct technique requires the input power waveform to be converted into the frequency domain and the response to be converted back into the time domain. These conversions can be achieved using the discrete Fourier transform (DFT). The required conversions to and from the frequency domain are given in (7.4) and (7.5) respectively.

$$\theta_y(t) = \mathcal{F}^{-1}\big(Q_x(\mathrm{j}\omega)\big) \tag{7.4}$$

$$Q_x(\mathrm{j}\omega) = \mathcal{F}\big(q_x(t)\big) \tag{7.5}$$

Where $\mathcal{F}$ represents the discrete Fourier transform ($\mathcal{F}^{-1}$ is the inverse), and θ and q represent the time-domain temperature and power respectively.

7.3.2 Representation of cross-coupling as a digital filter

In Chapter VI it was shown that Fourier transforms are computationally intensive and require a significant number of mathematical operations when used for temperature estimation. Using the technique developed in Chapter VI to reduce the computational complexity, the model of the system is converted to a digital infinite impulse response (IIR) filter. Using a digital filter has the advantage of being computationally light since only a few real number calculations need be performed for each time sample. A digital filter is fitted to each element of the cross-coupling matrix using the MATLAB function `invfreqz` which finds filter vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ such that Z_{th} in (7.6) is the best approximation to the measured cross-coupling [7.9]. This process effectively converts frequency-domain $Z_{\mathrm{th}}(\mathrm{j}\omega)$ to a digital filter with the equivalent transfer function in the $\mathcal{Z}$-domain, $Z_{\mathrm{th}}(\mathrm{z})$. A different filter will result for each level of active cooling and each pair of power input and temperature measurement points.

$$Z_{\mathrm{th}}(z) = \frac{\sum_{\lambda=0}^{G} b_\lambda z^{-\lambda}}{1 + \sum_{\lambda=1}^{H} a_\lambda z^{-\lambda}}; \quad z = e^{-\mathrm{j}2\pi\left(\frac{\omega}{\omega_\mathrm{s}}\right)} \tag{7.6}$$

Where ω_s is the sampling frequency, and G and H are the polynomial orders of filter vectors $\boldsymbol{b}$ and $\boldsymbol{a}$, respectively. The output of the filter to an arbitrary input sequence can be calculated using (7.7). This equation can also be expressed as the time-domain difference equations shown in (7.8) and (7.9) [7.10]. In this form, the input power at each time step can be evaluated individually, with the system state stored as a state vector $\boldsymbol{s} = (s_0, \ \cdots, \ s_{L-1})$ where L is the greater of G and H ($a_\lambda = 0$ for $\lambda > G$ and $b_\lambda = 0$ for $\lambda > H$).

$$\theta[i] = \sum_{\lambda=1}^{G} a_\lambda \theta[i-\lambda] + \sum_{\lambda=0}^{H} b_\lambda q[i-\lambda] \tag{7.7}$$

$$\theta[i] = q[i] \cdot b_0 + s_0[i-1] \tag{7.8}$$

$$\boldsymbol{s}[i] = \begin{pmatrix} s_0 \\ \vdots \\ s_{L-2} \\ s_{L-1} \end{pmatrix}[i] = q[i] \cdot \begin{pmatrix} b_1 \\ \vdots \\ b_{L-1} \\ b_L \end{pmatrix} + \begin{pmatrix} s_1 \\ \vdots \\ s_{L-1} \\ 0 \end{pmatrix}[i-1] - \theta[i] \cdot \begin{pmatrix} a_1 \\ \vdots \\ a_{L-1} \\ a_L \end{pmatrix} \tag{7.9}$$

The new output temperature is calculated using (7.8) with this value then being fed back to allow the calculation of the new state using (7.9). These operations are repeated at each time step, i, with the state vector, $\boldsymbol{s}$, being stored. The response of a simple system consisting of coupling between temperature and power (and not under varying cooling) can be directly calculated by applying the digital filter as shown in Fig. 7.4.

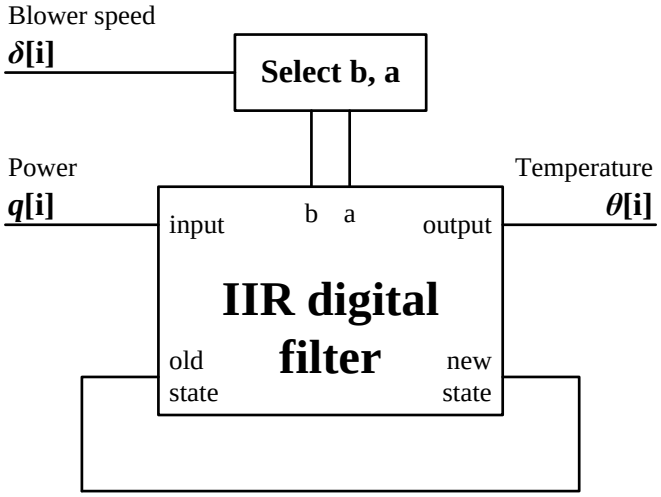

Fig. 7.4 Simple digital filter arrangement for temperature estimation with constant active cooling

7.4 Calculating the temperature response under dynamic active cooling

For statically cooled systems only a single impedance spectrum, $Z_{\mathrm{th}}(\mathrm{j}\omega)$, exists for each pair of devices. However, with active cooling, a separate $Z_{\mathrm{th}}(\mathrm{j}\omega, \delta)$ spectrum exists for each δ considered and therefore a separate filter, $Z_{\mathrm{th}}(z, \delta)$, with vectors $\boldsymbol{b}(\delta)$ and $\boldsymbol{a}(\delta)$

is required for each level of cooling. To estimate the temperature response under variable cooling, the appropriate filter must be used for each level of cooling.

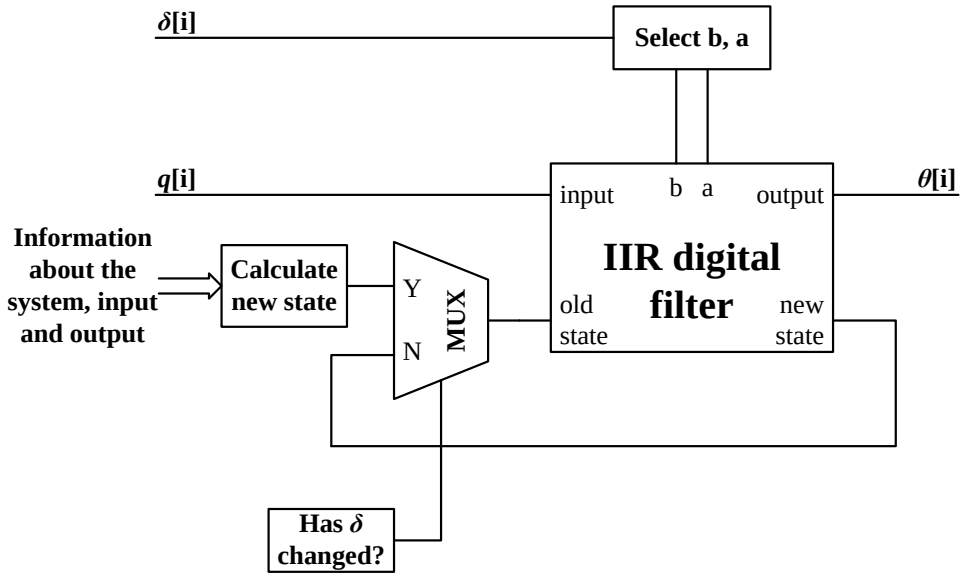

Fig. 7.5 Digital filter arrangement showing the need to recalculate the filter state vector after a change in active cooling

The difficulty posed by this approach is that there is no direct conversion between the internal state vectors of two different filters. A technique must therefore be developed to allow bumpless (i.e. smooth) transition between filters such that the modelled state of the system remains consistent as filter parameters $\boldsymbol{a}$ and $\boldsymbol{b}$ are substituted. Fig. 7.5 shows the procedure used. When there is no change in blower speed, the temperature estimation is simply the output of the digital filter to the power input (with the state vector zeroed initially); however, when there is a change in blower speed, a new digital filter is selected and a new state is calculated based on information about the system.

To limit the number of filters which must be fitted and stored, the range of blower speeds is quantised so that only a reasonable number of blower speeds are considered. When the blower speed changes, the filter in use is changed to the appropriate filter for the new blower speed. However, the internal filter state, $\boldsymbol{s}$, which is the filter's memory of its previous input, must be determined for the new filter such that the internal state reflects the current system state. It is not possible to directly convert the state of the previous filter to that for the new filter because the vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ may be significantly different, even for similar cross-coupling characteristics.

Techniques commonly used in control engineering for similar purposes are bumpless transfer from manual to automatic control [7.11] and gain scheduling for non-linear controllers whose operating domain has been divided into several approximately linear subdomains [7.12]. These methods are applicable where one or more closed-loop controllers are in operation. The proposed estimator is an open-loop system, however, and therefore these control techniques are not applicable. Bumpless transition at filter substitution is instead achieved using simple scaling. It is shown below that this approach produces results with minimal error compared to experimental verification.

Two easy-to-implement methods are available to determine the new state. In the first, the new filter can be set to its steady-state value. This state may be found mathematically and is scaled such that that the output temperature remains correct (i.e. without a bump) when the new filter is substituted. This approach has the disadvantage that the history of the input power to the system, stored in the filter state, is lost. Alternatively, to counteract this problem, the filter state can be determined by providing the filter with the same input waveform that the previous filter received. The input and output values to the filter can then be scaled to ensure the correct output is produced at the point of substitution to ensure bumplessness. Because the new filter has processed the entire input waveform, its state vector more accurately represents the true state of the system. The properties of these two techniques are described below.

7.4.1 Steady-state assumption method

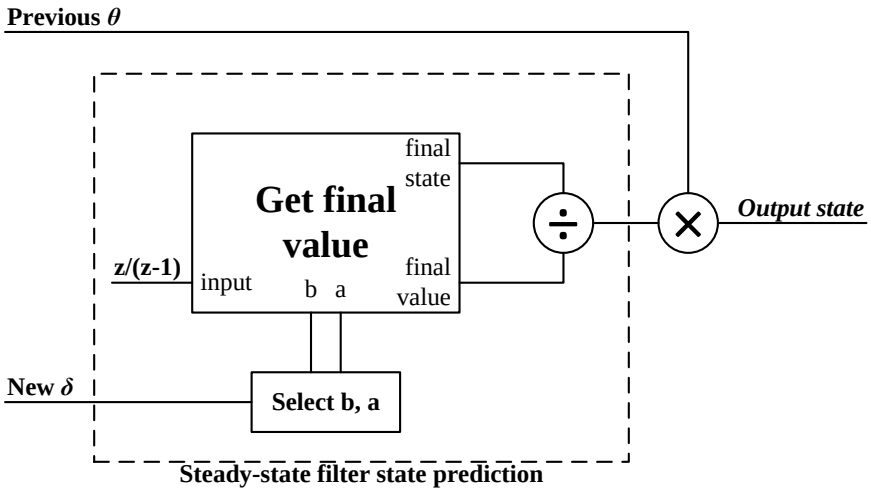

Fig. 7.6 Filter state recalculation technique for the steady-state assumption method. This replaces *calculate new state* in Fig. 7.5.

The simplest method to generate the internal state of a new filter is to initialise the state vector, $\boldsymbol{s}$, to the appropriate steady state level for the last output of the previous filter. This vector is found using the $\mathcal{Z}$-domain final value theorem [7.13] which states that the final value of the sequence can be calculated as

$$\theta[\infty] = \lim_{t\to\infty} \theta(t) = \lim_{z\to 1}(z-1)\Theta(z) = \lim_{z\to 1}(z-1)Z_{\text{th}}(z)Q(z) \tag{7.10}$$

To calculate the final value steady-state in steady state, $q(t)$ is 1 W for all t, therefore [7.10]

$$Q(z) = \frac{z}{z-1} \tag{7.11}$$

Having determined $\theta[\infty]$, the values of the of the state vector can be calculated by solving (7.9) for a constant state vector and the final temperature, as shown in (7.12).

$$\begin{aligned} s_{L-1} &= b_L + 0 - \theta[\infty]a_L \\ s_{L-2} &= b_{L-1} + s_{L-1} - \theta[\infty]a_{L-1} \\ &\vdots \\ s_0 &= b_1 + s_1 - \theta[\infty]a_1 \end{aligned} \tag{7.12}$$

The technique is shown in Fig. 7.6. The state vector is normalised by dividing the state by the final temperature, θ. Upon filter change, the normalised state vector of the new filter is scaled to the correct value by multiplying by the final output temperature of the previous filter. Scaling in this way is valid because digital filters are linear [7.10].

This steady-state method assumes that when the blower speed is changed, the system is in steady state. Because the memory of previous input is lost, this assumption is limited to cases where the filter has stabilised or where there is no significant delay between power input and response. The method is useful, however, because the steady-state vector, $\boldsymbol{s}$, can be calculated offline, eliminating the need for significant online computations.

To illustrate the operation of the technique in practice, device 1 in Fig. 7.2 was set to dissipate selected power waveforms and levels of active cooling while the temperatures of all devices were monitored. The temperature at each device was also estimated using the filter technique described above for a blower with speeds of between 0 and 6600 revolutions per minute (RPM). The model was quantised to one filter for each multiple of 50 RPM.

Two types of input were selected to illustrate the abilities and limitations of the technique, as described below.

1. A power waveform consisting of three repeats of the European driving cycle waveform [7.14] was chosen due its wide range of powers and changes of rates (for the sake of simplicity, power is taken as proportional to the velocity parameter in the cycle). A blower waveform consisting mainly of gradual changes was applied during the cycles, meaning a new state must regularly be calculated for the filter. These inputs represent a difficult real-world problem for the estimator.
2. A set of inputs consisting of irregular and rapid arbitrary changes in power and blower speeds was chosen. For these inputs, only occasional state recalculations are required and the consequent error in calculating the filter state is smaller.

For both waveforms, the power input was applied to device 1 with the temperature responses measured at devices 1 and 3 (i.e. the auto- and cross-coupling responses respectively) measured. The results are shown in Fig. 7.7.

For the gradual changes in Fig. 7.7(a), the agreement with practical results is excellent for device 1 (the auto-coupling case) with the root-mean-square error (RMSE) between the estimation and practical results being 2.3 K. In the case of the device 3 (cross-coupling), however, there is significant difference between estimated and practical results demonstrated by an RMSE of 7.8 K. This inaccuracy is caused by the loss of memory due to the steady-state assumption method removing the effect of past frequency input from the filter's internal state. In the case of auto-coupling, the time constants of the thermal response are short and therefore the effect of inputs some time ago is limited. However, when cross-coupling is considered, the time constants are

longer and the full effect of a change in input power is not given time to appear before there is another filter change.

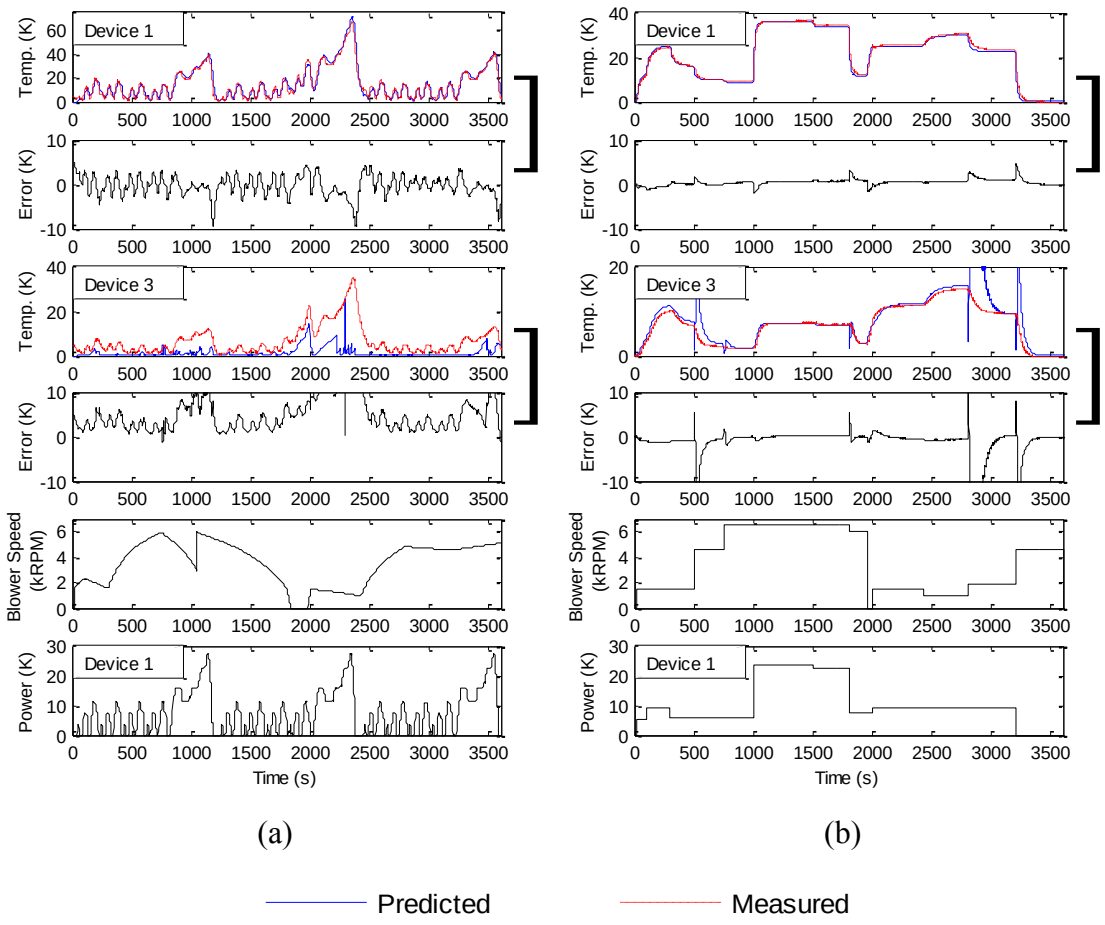

Fig. 7.7 Estimated and actual temperature responses at device 1 (auto-coupling) and device 3 (cross-coupling) due to power dissipation in device 1 under variable blower speed for the steady-state assumption method. Large square brackets between graphs link a device temperature result with its error. Results for system under (a) gradual frequent changes of blower speed, and (b) pronounced infrequent blower speed changes.

For frequently changing input, erasing the memory by repeatedly reinitialising the filter to steady state invalidates the estimation. The bumps in the cross-coupling case are a result of the apparent step change in input seen by the filter. The issues are highlighted where only irregular sudden changes in blower speed occur, such as in Fig. 7.7(b). In this case, the error is smaller in both the auto- and cross-coupling cases (with RMSEs of 0.8 and 6.8 K, respectively) because there is less memory loss. The error in the cross-

coupling is mainly from unwanted bumps. Owing to the inaccuracy caused by the loss of memory, a technique which does not make broad assumptions of previous inputs is required for accurate temperature estimations under constantly changing cooling in the case of cross-coupling.

7.4.2 Scaled input assumption method

An alternative approach is to assume that the new filter has experienced the same input as the old filter. The frequency content of the new filter's state vector will therefore be the same as the previous filter, although the filter state may have little numerical resemblance. To achieve this effect, state vectors for all filters are maintained irrespective of whether they are presently required. These vectors are updated at each time step, and a normalisation procedure is implemented where there is a change of blower speed. The full procedure for the scaled input assumption method is shown in Fig. 7.8.

The full range of blower speeds is quantised into discrete levels δ_1 to δ_n to limit the number of filters (and therefore memory) required. Each filter receives the same input but, because each has its own vectors $\boldsymbol{b}$ and $\boldsymbol{a}$, each also maintains a separate state vector, $\boldsymbol{s}$. The outputs of all the filters are multiplexed and the output from the filter whose δ is closest to the present level is selected as the active temperature estimation output. If δ does not change, the method is identical to the single filter method described in section 7.3.2, except that other filters are silently calculating their response and updating their state vectors.

When a change of blower speed occurs (and thus the '*multiply?*' flag in Fig. 7.8 is set), all the filters have their output normalised to the estimated temperature immediately preceding the change. To achieve this, a scaling factor is calculated to produce the correct output temperature immediately following the change of filter. The input power must also be scaled by the reciprocal of the output scaling factor to maintain correct frequency response of the filter unit. This method is more computationally intensive than the steady-state assumption method because it requires several filters to be simultaneously maintained. However, owing to the long time constants involved in thermal systems, a new temperature value is only occasionally input (every 4.6 seconds

in our example). This allows sufficient time to recalculate each filter using a typical microcontroller. The computational requirements are described in section 7.4.4.

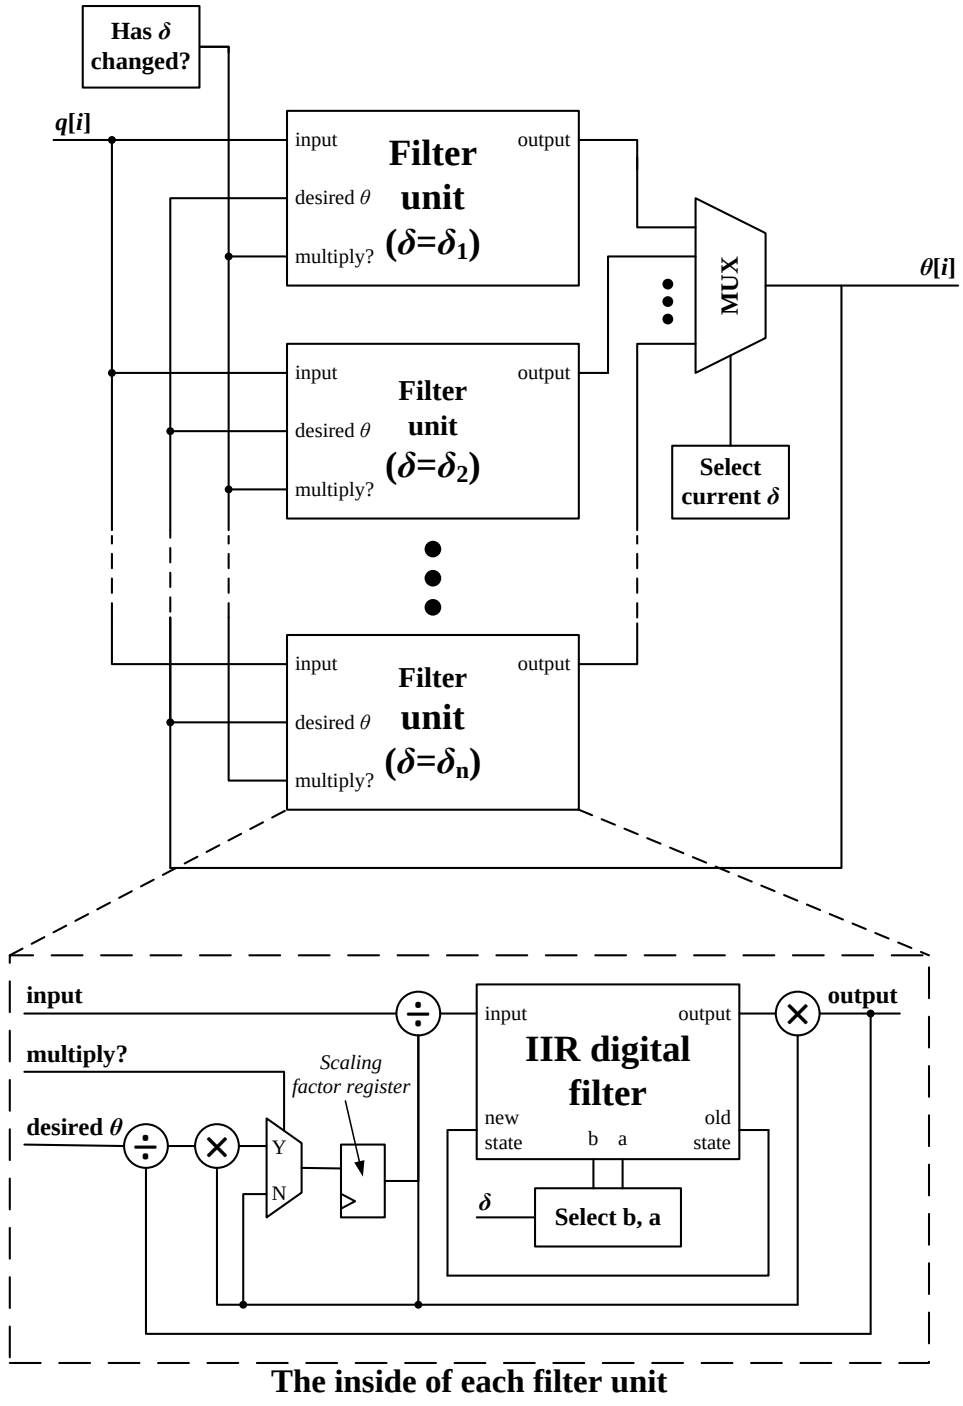

Fig. 7.8 Digital filtering in the scaled input assumption method.

For verification, the scaled input assumption method was applied to the same arbitrary power and blower waveforms as in the previous section. Results are shown in Fig. 7.9. For the auto-coupling response of device 1, the method shows little improvement over

the steady-state method, with similar RMSEs of 1.6 K and 1.6 K for Fig. 7.9(a) and (b) respectively. For cross-coupling, however, there is a significant improvement in the closeness of the estimation for the regularly changing blower speed seen in Fig. 7.9(a). The error is similar to that of the auto-coupling response at 0.7 K, a 90% reduction over the steady-state technique. In addition, the shapes of the curves are similar. The irregular power and blower speed test in Fig. 7.9(b) shows similar results to the steady-state assumption technique although the elimination of unwanted bumps reduces the RMSE to 0.5 K. These results confirm that with separated power and measurement points in addition to inputs involving regular changes in blower speed, there is an increased need for complex filter conversion techniques.

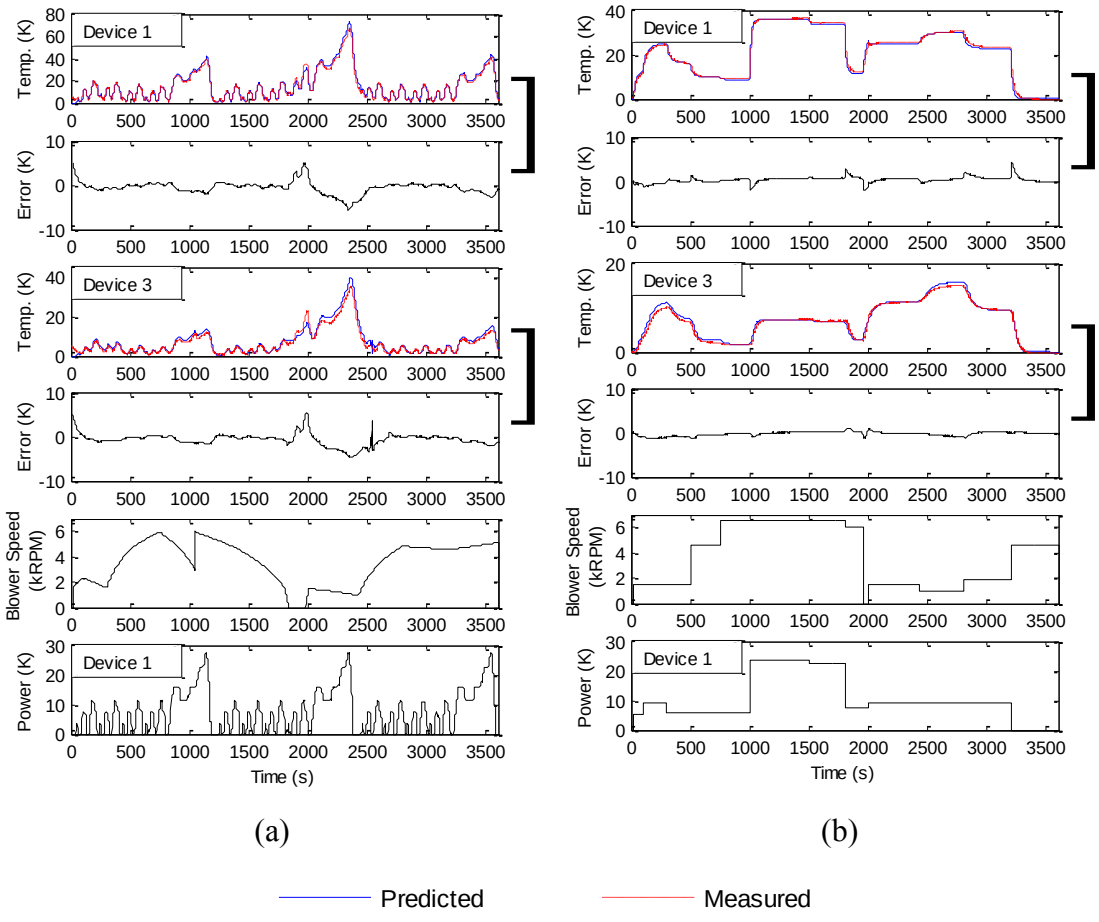

Fig. 7.9 Estimated and actual temperature responses at device 1 (auto-coupling) and device 3 (cross-coupling) due to power dissipation in device 1 under variable blower speed for the scaled input method. (a) Under gradual frequent changes of blower speed. (b) Under pronounced infrequent blower speed changes.

7.4.3 Temperature response with multiple devices dissipating

Results presented so far are for a single device dissipating power. However, the techniques presented in this Chapter are equally applicable in cases where multiple devices are dissipating. Fig. 7.10 shows estimated and actual temperature at device 1 where both devices 1 and 3 are dissipating the European driving cycle based power waveform under slow-changing cooling. The estimated response is calculated using (7.2) by summing the individual estimated temperature response at device 1 due to dissipation at device 1 alone and device 3 alone.

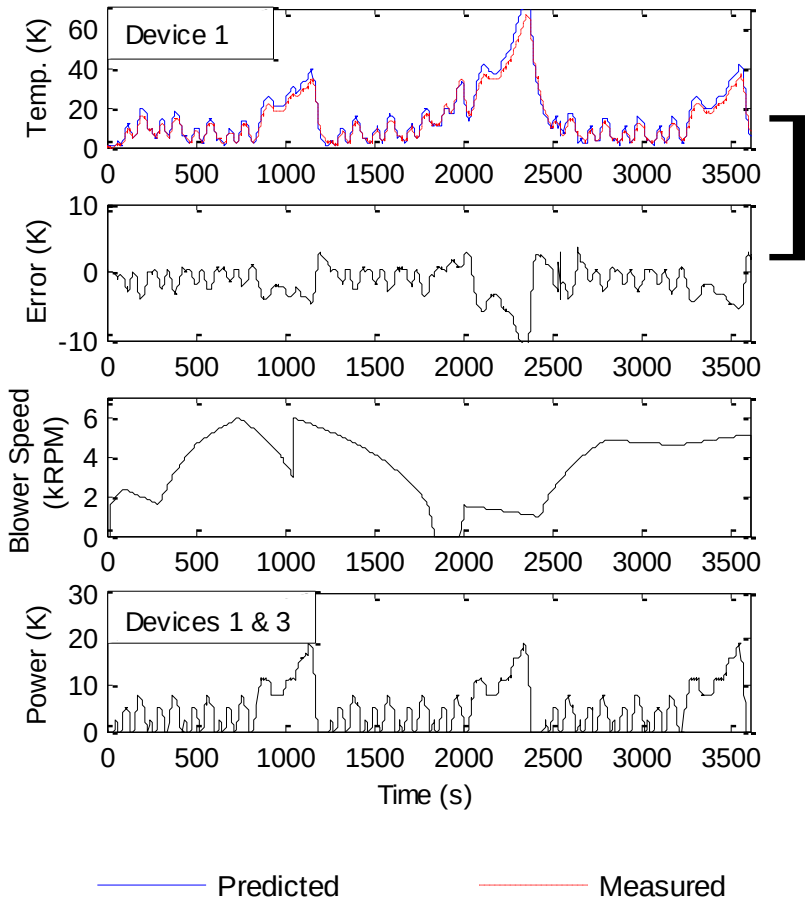

Fig. 7.10 Estimated and actual temperature responses at device 1 due to equal simultaneous power dissipation in devices 1 and 3 under variable blower speed

The agreement between estimated and practical results is excellent and is comparable to the results for a single device, demonstrating the technique's viability for multiple device temperature estimation. The RMSE is 2.9 K, which represents an insignificant error for most practical purposes.

7.4.4 Computational requirements

Method	Steady-state	Scaled input
Volatile storage (no. of values)	6	924
Read-only storage (no. of values)	1452	1452
Multiplications/divisions per second	3.3	499

Table 7.1 Computational requirements per device pair for the two proposed methods (4.5 s sampling period; 132 quantised cooling levels)

Table 7.1 shows the computational requirement for each $\boldsymbol{Z}_{\mathbf{th}\cdot x\to y}$ if the estimation technique were implemented on a microcontroller. Read-only storage, where the filter vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ are stored, is identical for both methods because the same filters are used. The scaled input assumption method requires significantly more volatile storage and multiplication/division operations to maintain the additional state vectors. However, because these requirements are typically fulfilled by modern microcontrollers and the scaled input assumption method produces better results, it is the preferred technique.

7.5 Auto-coupling active cooling control using feedback from the model

It has been shown that it is possible to estimate the temperature across devices due to real-time power dissipation data. In this section, the estimations are used to control an active cooling system in order to maintain system temperatures at a target level. Steady-state method estimations are used to control the temperature of a single device system (i.e. this section considers the auto-coupling case). The section thus provides a proof-of-concept for estimation-based feedback control.

Active temperature control has many applications in thermal management. One important use which is focussed on in this Chapter is the improved lifetime of power electronic components through reduction of thermal cycling. Depending on the type of device used, thermal fatigue due to thermal cycling can be the critical consideration in device (and therefore system) failure. When the proposed controller is used, cooling is

adjusted to maintain the device temperature at a high level. At periods of low power dissipation, cooling is minimised to reduce change in temperature; similarly, for high powers cooling is fully activated. The result is that device temperature varies little and thermal transients (and hence cycling) are reduced. This approach is not without limitations. The average temperature of devices is increased and this has a negative impact on lifetime [7.15], especially for non-semiconductor components such as capacitors [7.16], due to the increased rate of chemical reactions. A trade-off must therefore be made between the reduced thermal cycling and increased average temperature. Lifetime models for electronic devices can usually be generated which will indicate the optimum compromise for a particular arrangement [7.1][7.17][7.18].

For a thermally optimised system, it may be assumed that for a significant amount of time the system will run hot, since cool systems often indicate that excessive heatsinking has unnecessarily been incorporated into the design. Periods of low load will reduce temperature transiently; the temperature will increase again rapid when full load resumes if controlled active cooling is not used. Removing these thermal transients will eliminate the thermal cycling fatigue they cause.

7.5.1 PID tuning

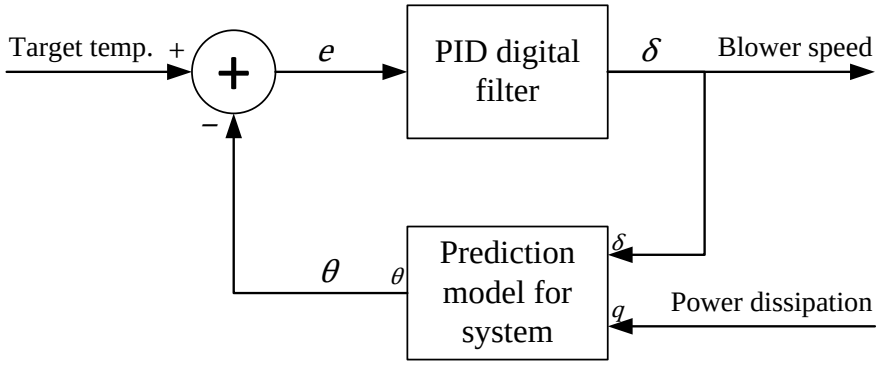

Fig. 7.11 Temperature control arrangement

Controlling temperature usually means monitoring the temperature at a critical point on the circuit and adjusting cooling, often using a relay controller with hysteresis. This requires temperature sensors to be embedded at critical points in the circuit, which represents increased manufacturing cost. It is therefore preferable to use the estimator proposed earlier in this Chapter as the source of temperature data. In this section, a PID-based active cooling controller which takes its feedback from the estimator is therefore

proposed. The PID controller is implemented as a digital filter meaning the whole system can be realised using a set of difference equations.

Fig. 7.11 shows the temperature control arrangement. The estimator is used to calculate device temperature in real-time from the power dissipation and blower speed. This temperature is subtracted from the target temperature to provide a temperature error which acts as an input to the PID controller, which in turn generates the blower speed demand. In order to implement this system, the PID controller must first be tuned. Because an empirical model of the system has already been developed, the controller can be tuned offline using only the temperature estimator developed earlier in the Chapter, removing the need for further practical experimentation. To demonstrate the effectiveness of the arrangement, the controller was tuned using the steady-state assumption temperature estimator developed in section 7.4.1 using the relay method [7.19] under a target temperature of 30 K above ambient. The steady-state method was used due to the simplicity of its implementation and its good results for the auto-coupling case. In simulation, the system is heated at a constant power of 15 W and cooling is provided at either of two discrete intensities: blower speeds of 0 and 4000 RPM in this case. If the system temperature falls below the target temperature, the less intense cooling level (0 RPM) is used while for over-temperature, the more intense cooling level (4000 RPM) is used. Fig. 7.12 shows the temperature and cooling response of the system undergoing relay tuning.

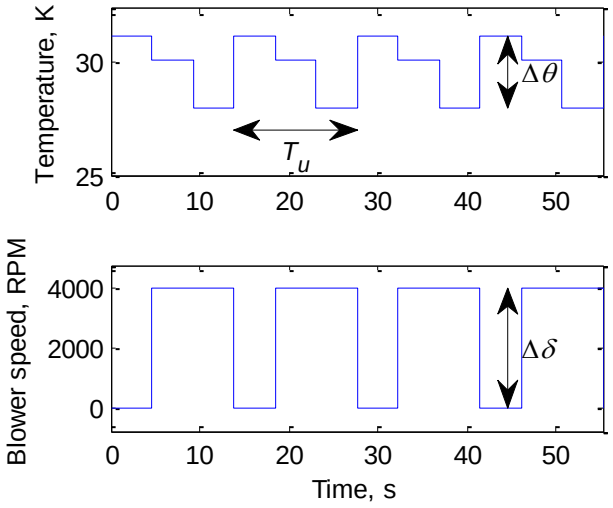

Fig. 7.12 Temperature and blower speeds of the estimator under relay tuning. The ultimate period, T_u, and the amplitudes of the waveforms are labelled.

The PID parameters are determined from these waveforms using the quantities labelled in Fig. 7.12. The ultimate period, T_{u}, is the period of oscillation observed on the temperature response and the ultimate gain, K_{u}, is proportional to the ratio of waveform amplitudes according to

$$K_{\mathrm{u}} = \frac{4}{\pi} \cdot \frac{\Delta\delta}{\Delta\theta} \tag{7.13}$$

Where $\Delta\theta$ and $\Delta\delta$ are the amplitudes of the variations in temperature and blower speed, respectively. Using the Ziegler-Nichols tuning rules [7.20], the parameters of the PID controller can be selected, as given in (7.14) to (7.16).

$$K_{\mathrm{p}} = 0.6 \times K_u \tag{7.14}$$

$$T_{\mathrm{i}} = \frac{T_{\mathrm{u}}}{2} \tag{7.15}$$

$$T_{\mathrm{d}} = \frac{P_{\mathrm{u}}}{8} \tag{7.16}$$

Where K_p, T_i and T_d are the gain, integration period and differential period of the PID controller respectively. The standard form of the PID controller with these parameters is shown in the Laplace domain in (7.17). To reduce the computational intensity required in the resulting system, the PID controller was implemented as a digital filter. The integral component was mapped to a discrete-time form using trapezoidal mapping while the differentiator used backwards Euler mapping as shown in (7.18) and (7.19), respectively. These types of mapping are used to ensure the controller implementation is stable and causal.

$$\mathcal{L}(\delta) = \mathcal{L}(e) \cdot K_\mathrm{p}\left(1 + \frac{1}{T_\mathrm{i}} \cdot \frac{1}{s} + \frac{T_\mathrm{d}}{1/s}\right) \tag{7.17}$$

$$\frac{1}{T_\mathrm{i}} \cdot \frac{1}{s} \quad \rightarrow \quad \frac{T_\mathrm{s}}{2T_\mathrm{i}} \cdot \frac{z+1}{z-1} \tag{7.18}$$

$$\frac{T_\mathrm{d}}{1/s} \rightarrow \frac{T_\mathrm{d}}{T_\mathrm{s}} \cdot \frac{1}{z/(z-1)} \tag{7.19}$$

Where T_s is the sampling rate. Using these mappings, the PID controller can be transformed into the $\mathcal{Z}$-domain. By rearranging the equation in terms of z^{-1} as shown in (7.20), the PID controller can be seen to be a rational function; that is to say it can be expressed as the quotient of polynomials in z^{-1}. The PID controller is therefore a digital IIR filter [7.10], and can be expressed in the time domain as shown in (7.21) to (7.24). This can be implemented using similar difference equations to those given in section **7.3.2**, noting that in the resulting digital filter, $\boldsymbol{b}$ and $\boldsymbol{a}$ are the coefficients of z^{-1} in the numerator and denominator respectively. In this case, $\boldsymbol{a} = (1, \quad -1)$ and $\boldsymbol{b}$ is as defined below.

$$\mathcal{Z}(\delta) = \mathcal{Z}(e) \cdot K_\mathrm{p} \frac{(1 - z^{-1}) + \frac{T_\mathrm{s}}{2T_\mathrm{i}}(1 + z^{-1}) + \frac{T_\mathrm{d}}{T_\mathrm{s}}(1 - z^{-1})^2}{1 - z^{-1}} \tag{7.20}$$

$$\delta[k] = b_0 e[k] + b_1 e[k-1] + b_2 e[k-2] + \delta[k-1] \tag{7.21}$$

$$b_0 = K_\mathrm{p}\left(1 + \frac{T_\mathrm{s}}{2T_\mathrm{i}} + \frac{T_\mathrm{d}}{T_\mathrm{s}}\right) \tag{7.22}$$

$$b_1 = K_\mathrm{p}\left(-1 + \frac{T_\mathrm{s}}{2T_\mathrm{i}} + 2\frac{T_\mathrm{d}}{T_\mathrm{s}}\right) \tag{7.23}$$

$$b_2 = K_\mathrm{p}\frac{T_\mathrm{d}}{T_\mathrm{s}} \tag{7.24}$$

In the practical system, the value of δ is limited to the blower speeds that the system is capable of delivering, $0 \leq \delta \leq 6600$ RPM for the blower used in this Chapter. Where the blower speed demand goes outside these limits, the blower will be limited to the nearest practical value. In this case, the system is temporarily uncontrollable using the available cooling and therefore the accuracy of temperature control is reduced. Mitigation for this problem is described below.

7.5.2 Practical implementation of PID control

To verify the accuracy of the proposed temperature controller, a PID controller was fitted to $Z_{\mathrm{th}\cdot 1\to 1}$ resulting in $T_\mathrm{u} = 13.8$ s and $K_\mathrm{u} = 1350$ (K·min)$^{-1}$. The estimator (using the steady-state assumption method) and PID controller were implemented as digital filters according to (7.14) to (7.16) and (7.21) to (7.24) and coupled together as shown in Fig. 7.11. Device 1 in the experimental arrangement was set to dissipate a typical waveform of power (in this case, the EU driving cycle with a static offset) while feedback from the temperature estimator was fed into the PID controller to produce a cooling demand. The demanded cooling was applied to the system and the actual temperature response measured.

This experiment was carried out with a temperature target of 40 K above ambient. Results, shown in Fig. 7.13, are plotted for estimator output and measured practical results. Grey areas on the temperature plots indicate that the blower speed is at its maximum and minimum and therefore, in these areas, the system is temporarily uncontrollable.

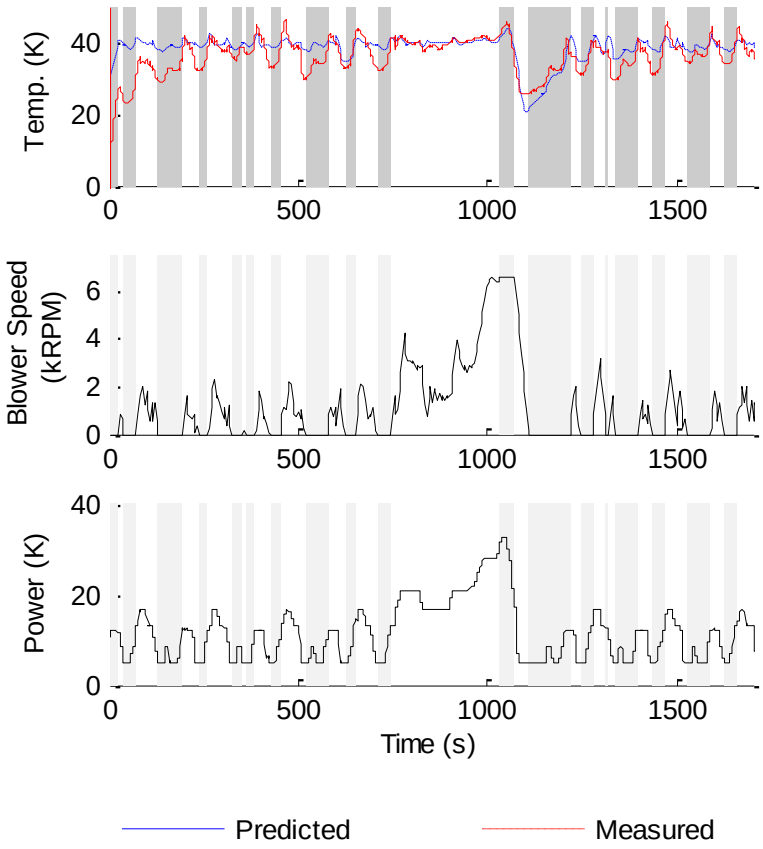

Fig. 7.13 Temperature response of the thermal system under arbitrary power input with PID-controlled active cooling and a target temperature of 40 K above ambient

The temperature control is good in the controllable regions, matching the target temperature with 4 K perturbations. The perturbations are larger for the practical result compared to the theoretical result. This demonstrates that for a PID controller with finite sampling speed the response is less accurate where rapid changes in cooling are required, due to the response time of the physical blower. In addition, at the times where the output of the PID controller demands a blower speed outside the practical range, the PID controller's state would be invalid if it continued to be updated as normal because a large integration error would accumulate, giving integral windup. Instead, when the maximum blower speed is exceeded, the state of the PID controller is not updated until the demand is within the practical range. Similarly, when a negative blower speed is requested, the PID controller's state is zeroed until a positive speed is requested. This is demonstrated particularly in Fig. 7.13 where, during the long controllable region

between 750 and 1050 s, both theoretical and practical temperature control is high quality.

7.5.3 Comparison to uncontrolled cooling

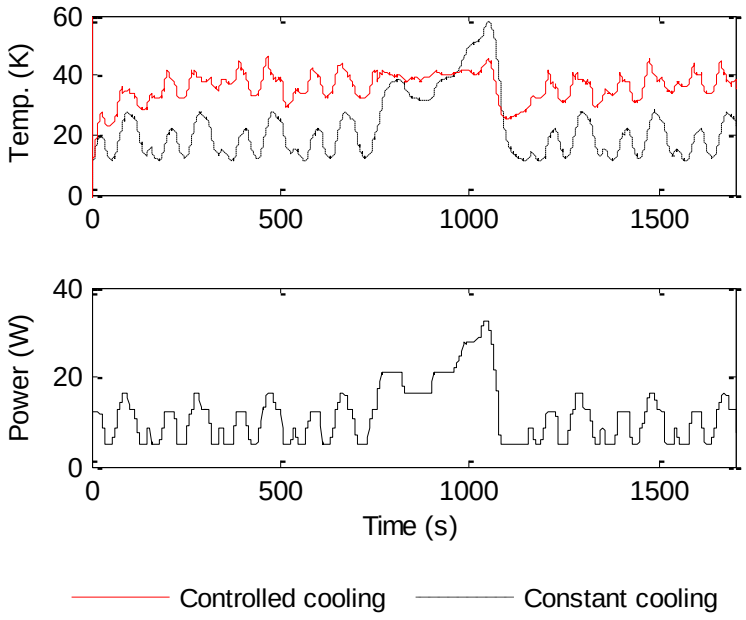

Fig. 7.14 Comparison of controlled and controlled cooling

The purpose of the controller reported is to reduce rapid temperature transients under varying load, which may lead to thermal fatigue or reduced accuracy in temperature sensitive systems. To verify the effectiveness of the proposed technique, the practical results of the demonstrator under the proposed control from Fig. 7.13 are compared to the temperature response that was measured when the system was excited by the same power waveform under a constant 2900 RPM blower speed cooling regime. The results are compared in Fig. 7.14 and demonstrate that the proposed technique reduces thermal cycling significantly, maintaining temperatures closer to the set point even with a power demand where the system is not always controllable.

7.6 Chapter conclusions

A computationally efficient temperature estimator for a power electronics system under varying levels of active cooling has been proposed. The estimator uses data from offline mixed PRBS-based characterisation to produce a set of digital IIR filters which model the dynamics of the system under quantised levels of active cooling. Two techniques are

developed to convert the internal state of one filter to another when the cooling level changes. The first assumes that the system is in steady-state at the instant where cooling changes, a typically valid assumption for systems with infrequent changes in power dissipation or cooling level. However, greater accuracy is achieved using the second technique, which duplicates the frequency content of the existing filter's state by exposing it to the same input. The estimation techniques, which demonstrate accurate tracking of temperature with minimal error when compared with the experimental data, can be used as part of a cooling control system. Such a system is demonstrated for the single device auto-coupling case, showing that thermal cycling can be reduced without the need for device-attached temperature sensors.

7.7 References

[7.1] H. Lu, C. Bailey and C. Yin, '*Design for reliability of power electronics modules*', Microelectronics Reliability, vol. 49, pp. 1250-1255, 2009.

[7.2] P. Ning, G. Lei, F. Wang and K. D. T. Ngo, '*Selection of heatsink and fan for high-temperature power modules under weight constraint*', *in Proc* Applied Power Electronics Conference and Exposition (APEC), 2008, pp. 192-198.

[7.3] L. Meysenc, M. Jylhakallio and P. Barbosa, '*Power electronics cooling effectiveness versus thermal inertia*', IEEE Transactions on Power Electronics, vol. 20, pp. 687-693, 2005.

[7.4] L. Fried, '*Prediction of Temperatures in Forced-Convection Cooled Electronic Equipment*', IRE Transactions on Component Parts, vol. 5, pp. 102-107, 1958.

[7.5] C. Lundquist and V. P. Carey, '*Microprocessor-based adaptive thermal control for an air-cooled computer CPU module*', *in Proc* Semiconductor Thermal Measurement and Management, 2001, pp. 168-173.

[7.6] G. Xiong, M. Lu, C. L. Chen, B. P. Wang and D. Kehl, '*Numerical optimization of a power electronics cooling assembly*', *in Proc* Applied Power Electronics Conference and Exposition (APEC), 2001, pp. 1068-1073 vol.2.

[7.7] P. J. Rodgers, V. C. Eveloy and M. R. D. Davies, '*An Experimental Assessment of Numerical Predictive Accuracy for Electronic Component Heat Transfer in Forced Convection—Part I: Experimental Methods and Numerical Modeling*', Journal of Electronic Packaging, vol. 125, pp. 67-75, 2003.

[7.8] P. J. Rodgers, V. C. Eveloy and M. R. Davies, '*An Experimental Assessment of Numerical Predictive Accuracy for Electronic Component Heat Transfer in Forced Convection—Part II: Results and Discussion*', Journal of Electronic Packaging, vol. 125, pp. 76-83, 2003.

[7.9] MathWorks. (2013, 8 July 2013). *MATLAB 2013a documentation: invfreqz*.

[7.10] A. V. S. Oppenheim, R.W., *Discrete-Time Signal Processing*. New Jersey: Prentice Hall, 1999.

[7.11] R. Hanus, M. Kinnaert and J. L. Henrotte, '*Conditioning technique, a general anti-windup and bumpless transfer method*', Automatica, vol. 23, pp. 729-739, 1987.

[7.12] J. D. Bendtsen, J. Stoustrup and K. Trangbaek, '*Bumpless transfer between observer-based gain scheduled controllers*', International Journal of Control, vol. 78, pp. 491-504, 2005.

[7.13] E. W. Kamen, *Introduction to signals and systems*. New York: Macmillan, 1990.

[7.14] *EEC Journal Officiel No. C81*, 1990.

[7.15] M. Held, P. Jacob, G. Nicoletti, P. Scacco and M. H. Poech, '*Fast power cycling test of IGBT modules in traction application*', *in Proc* Power Electronics and Drive Systems (PEDS), 1997, pp. 425-430 vol.1.

[7.16] P. Spanik, M. Frivaldsky and A. Kanovsky, '*Life time of the electrolytic capacitors in power applications*', *in Proc* ELEKTRO, 2014, pp. 233-238.

[7.17] Y. Yang, R. Master, G. Refai-Ahmed and M. Touzelbaev, '*Transient Frequency-Domain Thermal Measurements With Applications to Electronic Packaging*', IEEE Transactions on Components, Packaging and Manufacturing Technology, vol. 2, pp. 448-456, 2012.

[7.18] H. Lu, T. Tilford and D. R. Newcombe, '*Lifetime Prediction for Power Electronics Module Substrate Mount-down Solder Interconnect*', *in Proc* High Density packaging and Microsystem Integration (HDP), 2007.

[7.19] D. I. Wilson, '*Relay-based PID tuning*', Automation and Control, Auckland University of Technology, New Zealand, pp. 10-11, 2005.

[7.20] J. Ziegler and N. Nichols, "Optimum settings for automatic controllers," in *Proceedings of the Annual Meeting of the American Society of Mechanical Engineers* vol. 64, ed. New York, 1942, pp. 759-768.

Chapter VIII

Estimating die temperatures in a multiple device system following PRBS-based characterisation

This Chapter presents a technique to estimate the die temperature of a MOSFET based on an empirical model which is derived following an offline thermal characterisation of a system. The characterisation uses the method for near-simultaneous measurement of die of the system temperature during controlled power dissipation presented in Chapter III. Using the pseudorandom binary sequence-based system identification approach, the thermal transfer impedance, or cross-coupling, between the dice of two devices on shared cooling is determined. For each pair of devices, an infinite impulse response (IIR) digital filter is fitted to both the auto- and cross-coupling. By inputting power data to the digital filters, the temperature response of the semiconductor die in each of the devices can be estimated. Experimental verification shows excellent agreement between measured and estimated temperature responses. For comparison, the surface temperature is also estimated. The differences demonstrate the advantages of using die temperature estimations. Results show that die temperature estimations have accuracy and computational requirements which are comparable to those of the surface temperature estimations made in Chapter VI. The usefulness of the technique for estimating die temperatures in real time without the need for on-die sensors is thus confirmed.

8.1 Introduction

Chapter VI presented a method to estimate the case temperature in real-time following an offline thermal characterisation. Knowledge of real-time temperatures has many applications including predicting mechanical failure [8.1], manipulating the temperature profile of devices to prevent overheating and reduce thermal cycling [8.2], and cooling control [8.3]. However, the efficacy of these applications would be improved if, instead of relying on case temperature, the temperatures of the semiconductor dice were used. This is because in most cases it is damage to the active junction on the semiconductor

die which leads to system failure through mechanisms including bond wire lift-off [8.4][8.5], ruptured wire link to the die [8.5] and dopant diffusion [8.6]. Die temperatures are therefore the variables which should preferably be monitored and controlled. Case temperatures are only an approximation to junction and die temperature and the relationship between the two varies with packaging, cooling and input power.

A number of techniques can be used to determine the temperatures of the junctions on semiconductor dice by measuring a temperature sensitive parameter (several of which are described in Chapter II). However, these techniques typically require exclusive use of the semiconductor which makes it challenging to take die temperature readings of a working device in an operational system. One solution is to add specialised die temperature measurement circuitry to the system and sample temperature sensitive parameters from voltage and current transients during a switching event. Unfortunately, this solution introduces increased per-unit cost and complexity to a system—a lower cost and easier to implement approach is desirable.

In Chapter VI, a case temperature estimation methodology was proposed which eliminated the need for per-device temperature sensors in a power converter. In this Chapter, that solution is further developed to estimate die temperatures in an electronic power system. Although a number of techniques for the same purpose have been published in the literature (see Chapter II), often only single-device systems are considered. Practical systems typically consist of several devices mounted on shared cooling. This Chapter therefore presents a technique to estimate the die temperature of each device in a multiple device system in real-time. The characterisation process is conducted offline meaning the solution does not require additional circuitry in operation. A comparison with the case temperature results presented in Chapter VI is also provided to demonstrate the advantages of this die temperature-based estimation technique.

8.2 Characterisation of the thermal auto- and cross-coupling

Because a semiconductor die is significantly more sensitive to its own power dissipation than to the power dissipation of its neighbours, the thermal auto-coupling is more

significant when considering the die temperature than the case temperature. Accurate measurement is therefore vital; however, practical measurement of the thermal auto-coupling presents a number of issues:-

1. Taking a temperature reading of the die during power dissipation is difficult because many die temperature measurement techniques require the exclusive use of the die.
2. The parameters which are typically measured to calculate die temperature (e.g. forward diode voltage, MOSFET threshold voltage or on-state resistance) have relatively low temperature sensitivity, making inferring the die temperature susceptible to measurement noise.
3. If problem 1 is avoided by eliminating power dissipation during the measurement, the instantaneous power dissipation is affected which has an impact on results.

In order to address these issues, the near-simultaneous temperature measurement and power controller circuit reported in Chapter III is used. This technique is not truly simultaneous, because power is interrupted to take measurements; however, by keeping the measurement time small and using a precise noise-resilient analogue-to-digital converter, good die temperature measurements can be taken seemingly simultaneously without significantly affecting results.

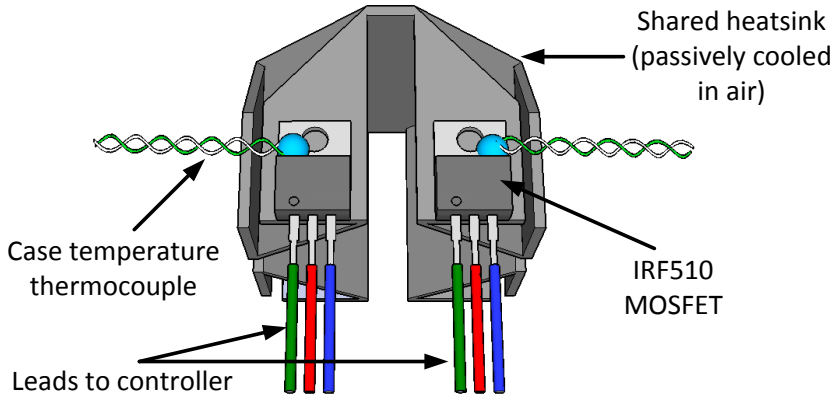

Fig. 8.1 Experimental arrangement for characterisation and die temperature estimation

An experimental arrangement consisting of two MOSFETs mounted on a single heatsink is used to demonstrate the techniques described. This arrangement is shown in Fig. 8.1. A system with multiple devices thermally coupled to each other is analysed

because it is representative of real-world situations; for instance, a power module typically consists of several power devices on a common backplate. Power dissipation in one device causes heating both in that device (due to the thermal auto-coupling) and in nearby devices (due to cross-coupling).

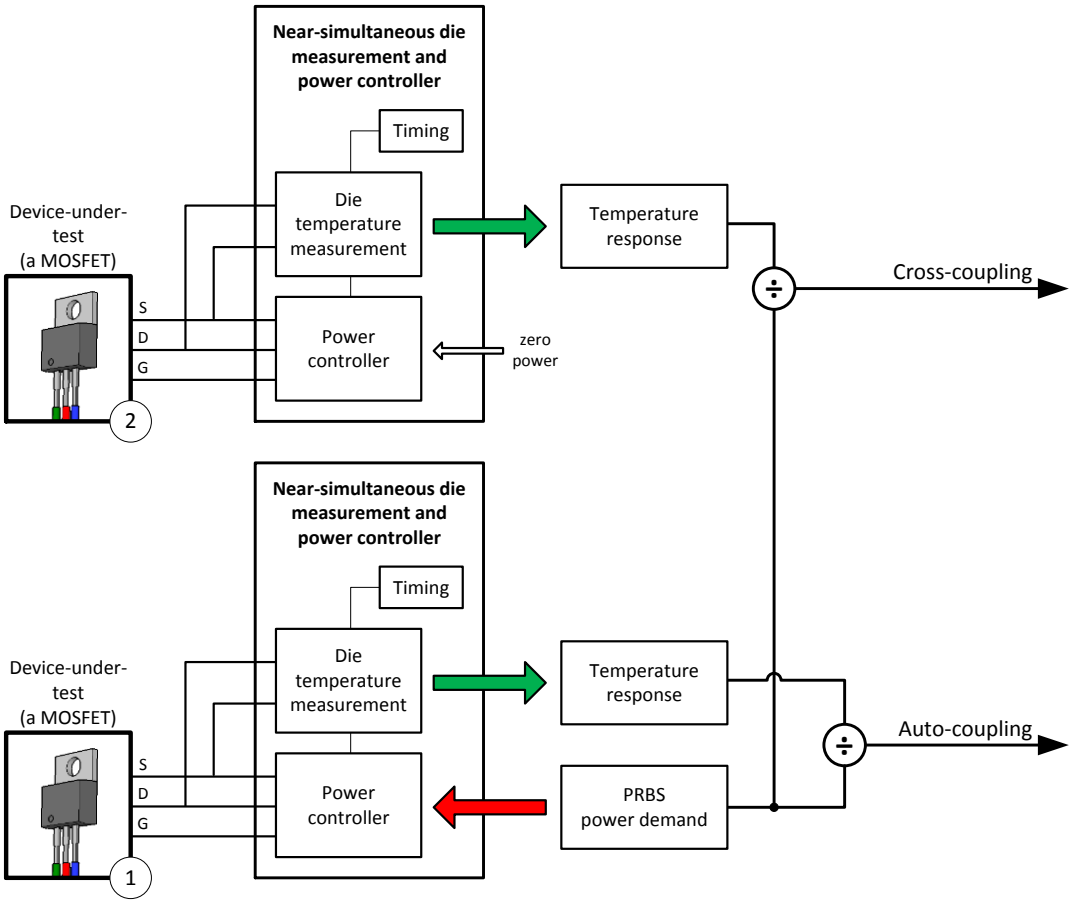

Fig. 8.2 Characterisation of the cross- and auto-coupling

The characterisation of the auto- and cross-coupling is performed using the technique described in Chapters IV and VI. The arrangement shown in Fig. 8.2 is used. A PRBS power demand is delivered to the power controller which in turn controls the dissipation in the device-under-test. For auto-coupling, the die temperature measurement module in the same controller is connected to the device and measures the temperature response. For cross-coupling characterisation, both devices are attached to separate near-simultaneous controllers. The power dissipation is applied to one device while the temperature of the other is recorded. This arrangement provides data for thermal characterisation.

In addition to auto- and cross-coupling, the die-to-case coupling is characterised in order to highlight the differences between the external case temperature-based results presented in Chapter VI and the die temperature results produced in this Chapter. Therefore, during each thermal characterisation three temperatures are taken:-

1. The die temperature of the dissipating device (to determine the auto-coupling)
2. The case temperature of the dissipating device (to determine the case-to-die coupling)
3. The die temperature of the other device (to determine the cross-coupling)

8.2.1 Characterisation results

Before measurements are taken, both measuring circuits and MOSFETs were calibrated using the technique described in Chapter III. The thermal characteristics were then determined using the simple PRBS technique. To achieve sufficient bandwidth, two 8-bit PRBS power inputs were used in turn, clocked at 0.05 Hz and 0.5 Hz respectively. To reduce the effect of noise, several repeats of each sequence were averaged together. This resulted in a valid frequency range of 0.20 mHz to 190 mHz. In the cross-coupling case, the upper frequency was reduced to 24 mHz since noise was a problem above this frequency. This has negligible effect on results because the cross-coupling is small above this frequency.

A microcontroller-based PRBS waveform generator was used as an input to the power controller to apply each of the signals. Once running, a settling period was allowed for any transient effects to pass before the measurements were taken. The results from each PRBS clock frequency were joined together to produce a single characteristic. Noise affected portions of the results were excised. The thermal coupling results are shown as Bode plots in Fig. 8.3.

Results show that the thermal coupling is highest at the die under test, with a substantial reduction between die and case. The auto-coupling shows increased impedance compared to the die-to-case coupling at low frequency and has a reduced roll off at higher frequencies. The die temperature is therefore significantly more sensitive to high power transients, which have high magnitude at higher frequencies, than the case temperature. The extent of the difference between the die and case characteristics

highlights the size of error that can be caused by using the case temperature as an indicator of die temperature instead of using the die temperature itself.

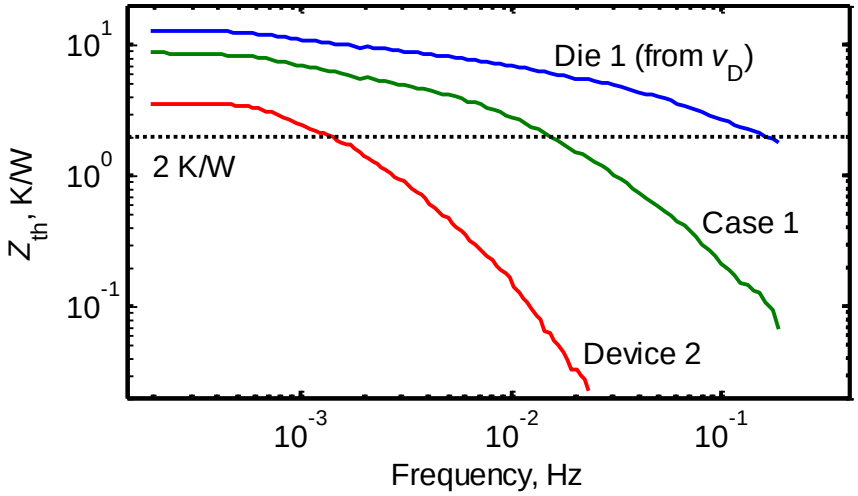

(a)

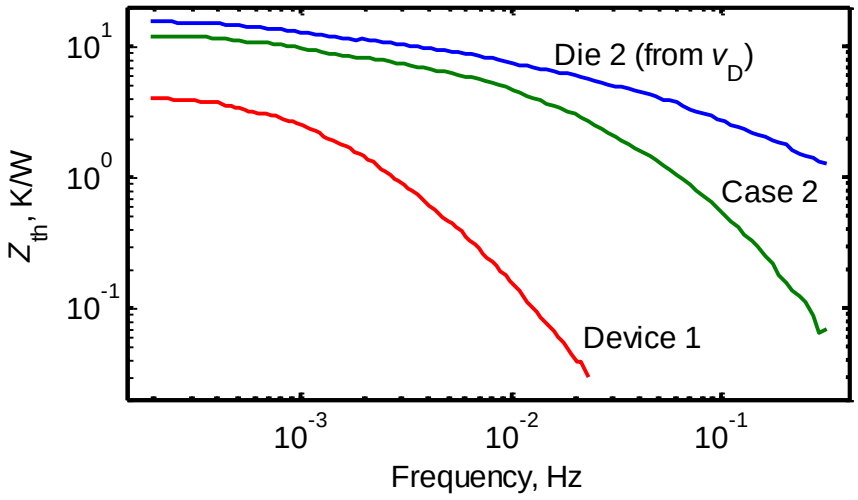

(b)

Fig. 8.3 Thermal coupling Bode plots from (a) device 1 and (b) device 2

The cross-coupling between the two devices is significantly lower than the die-to-case coupling. This is expected because the thermal impedance reduces with distance. The cross-coupling has a similar roll off to, but lower corner frequency than, the die-to-case coupling. It therefore has reduced sensitivity to high frequency inputs. For example, the thermal coupling between the die of device 1 and (i) itself at 160 mHz, (ii) the case of device 1 at 15 mHz, and (iii) the other device at 1.4 mHz are identical at 2.0 K/W. There is therefore an order of magnitude difference between the frequency required to cause a

given perturbation at the die and the frequency required at the case. A further order of magnitude is required for the same effect at the other device. For systems with high frequency components in their power waveforms, this means that larger temperature perturbations are seen at the die compared to the case. In addition, the higher DC component in the auto-coupling case causes a higher average temperature. Real-time estimation of the die temperature is therefore important in cases of rapidly varying power dissipation.

8.3 Temperature estimation

The system's thermal dynamics can be described using the frequency-dependent cross-coupling matrix, $\boldsymbol{Z}_{\mathbf{th}}(\mathrm{j}\omega)$, which is related to temperature and power by (8.1) and (8.2).

$$\boldsymbol{\Theta}(\mathrm{j}\omega) = \boldsymbol{Z}_{\mathbf{th}}(\mathrm{j}\omega) \cdot \boldsymbol{Q}(\mathrm{j}\omega) \tag{8.1}$$

$$\boldsymbol{Z}_{\mathbf{th}} = \begin{pmatrix} Z_{\mathrm{th}\cdot 1\to 1} & Z_{\mathrm{th}\cdot 2\to 1} \\ Z_{\mathrm{th}\cdot 1\to 2} & Z_{\mathrm{th}\cdot 2\to 2} \end{pmatrix}; \quad \boldsymbol{\Theta} = \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix}; \quad \boldsymbol{Q} = \begin{pmatrix} Q_1 \\ Q_2 \end{pmatrix} \tag{8.2}$$

Where $\Theta_x(\mathrm{j}\omega)$ and $Q_x(\mathrm{j}\omega)$ are the frequency-domain temperature representations of temperature and power respectively at die of device *x*, respectively, and $Z_{\mathrm{th}\cdot x\to y}(\mathrm{j}\omega)$ is the cross-coupling between devices *x* and *y*. To avoid the need to implement frequency domain mathematics, the computationally efficient infinite impulse response (IIR) digital filter method is used. The frequency domain characteristics are transformed into the $\mathcal{Z}$-domain using the MATLAB function `invfreqz` [8.7]. Equations (8.3) to (8.5) show the transformation.

$$\mathcal{Z}\{\boldsymbol{\Theta}(\mathrm{j}\omega)\} = \mathcal{Z}\{\boldsymbol{Z}_{\mathbf{th}}(\mathrm{j}\omega)\} \cdot \mathcal{Z}\{\boldsymbol{Q}(\mathrm{j}\omega)\} \tag{8.3}$$

$$\boldsymbol{\Theta}(z) = \boldsymbol{Z}_{\mathbf{th}}(z) \cdot \boldsymbol{Q}(z) \tag{8.4}$$

$$Z_{\mathrm{th}}(z) = \frac{\sum_{\lambda=0}^{4} b_\lambda z^{-\lambda}}{1 + \sum_{\lambda=1}^{7} a_\lambda z^{-\lambda}}; \quad z = e^{-\mathrm{j}2\pi\left(\frac{\omega}{\omega_\mathrm{s}}\right)} \tag{8.5}$$

Where z is the $\mathcal{Z}$-operator variable, ω_s is the angular sampling frequency and $\boldsymbol{a}$ and $\boldsymbol{b}$ are the IIR filter coefficients. The characterisation procedure, which determines $\boldsymbol{a}$ and $\boldsymbol{b}$, is computationally intensive; however, it is performed offline only once per system.

8.3.1 Implementation of temperature estimation

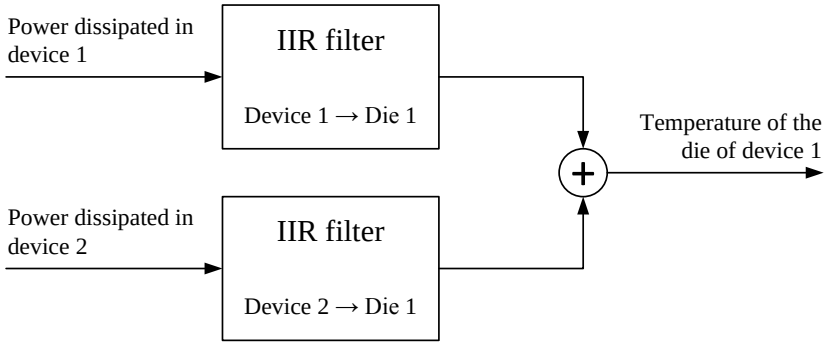

Fig. 8.4 Procedure for calculating die temperature estimation from power dissipation data

The temperature estimator from Chapter VI is implemented for die temperature estimation. Two IIR filters are required to estimate each temperature response. In order to estimate the temperature at device 1, the contribution from device 1 (the auto-coupling) and the contribution from device 2 (the cross-coupling) are calculated separately and added together to produce the overall response. This arrangement is shown in Fig. 8.4.

To verify the proposed techniques, both devices in the example system shown in Fig. 8.1 are excited with a power waveform based on the European community's standard driving cycle waveform [8.8]. This was selected as the test power waveform for both devices because it contains a wide range of levels of power and rates of change. The input power waveform to device 2 was delayed to allow the differences between the auto- and cross-coupling responses to be identified. A double speed version of the waveform was also used to highlight the effect of rapid power change. For simplicity, the input power is taken as proportional to the velocity parameter in the driving cycle.

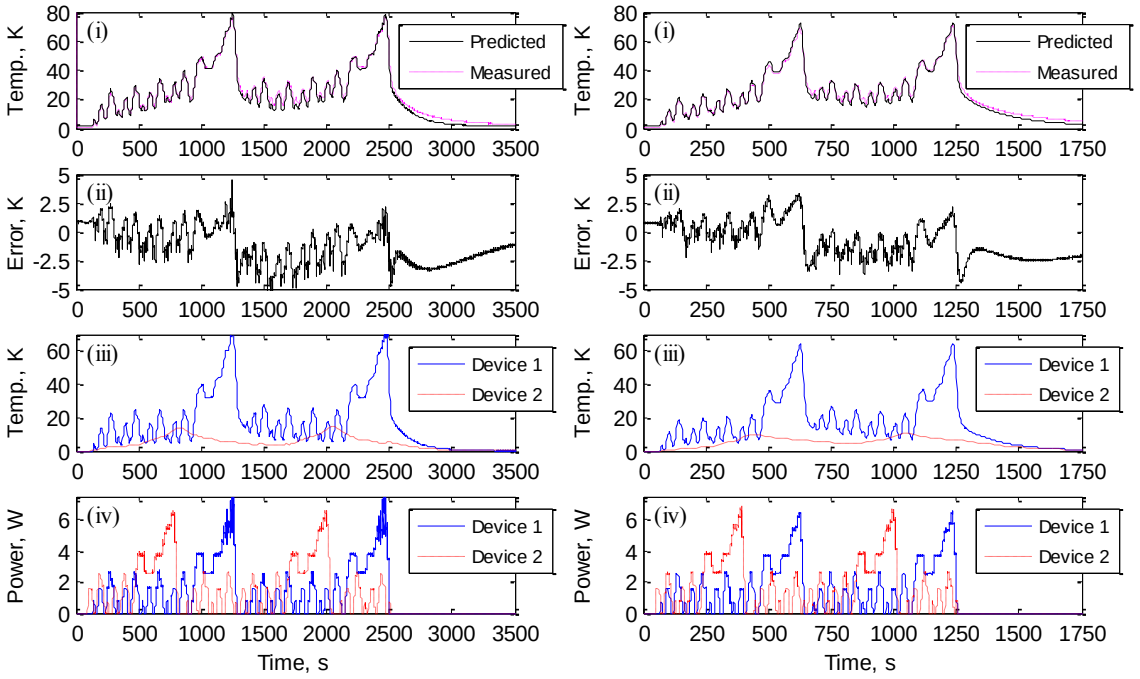

Fig. 8.5 Die temperature response of device 1 at (a) standard rate, and (b) twice standard rate. The graphs show: (i) estimated versus measured temperatures, (ii) error between estimated and measured temperatures, (iii) the temperature response due to dissipation in device 1 and device 2 separately, and (iv) the power input to the devices.

The power waveforms were applied to the devices and the die temperature of device 1 was measured. Graph (i) of Fig. 8.5 shows the estimated and measured temperature response of the die of device 1, with temperatures expressed relative to ambient. Estimation results show excellent agreement with measured temperature as shown by the error in graph (ii). There is no significant difference in accuracy between the results from standard and double rate driving cycle tests shown in Fig. 8.5(a) and (b) respectively. They have root-mean-square errors between measured and estimated results of 2.3 K and 1.8 K for the standard and double rate power inputs, respectively. When there is a rapid temperature change, errors are greatest due to a delay in the estimator's response. In practical circumstances, this is unimportant because the correct peak temperatures are calculated meaning calculations of thermal cycling fatigue and similar properties are not affected.

Graph (iii) shows the temperature response at the die of device 1 due to power dissipation in devices 1 and 2 separately, illustrating the difference between the auto-coupling and cross-coupling responses. The results are summed together to calculate the estimation in graph (i). The auto-coupling response (the response due to device 1) exhibits high amplitude and rapid reaction to changes in input power. As a result, it dominates the shape of the overall estimation; however, the cross-coupling response (due to device 2) has a slow-changing but important influence on the estimation. By taking into account both devices, an accurate overall temperature response can be produced.

8.3.2 Comparison to surface temperature estimation

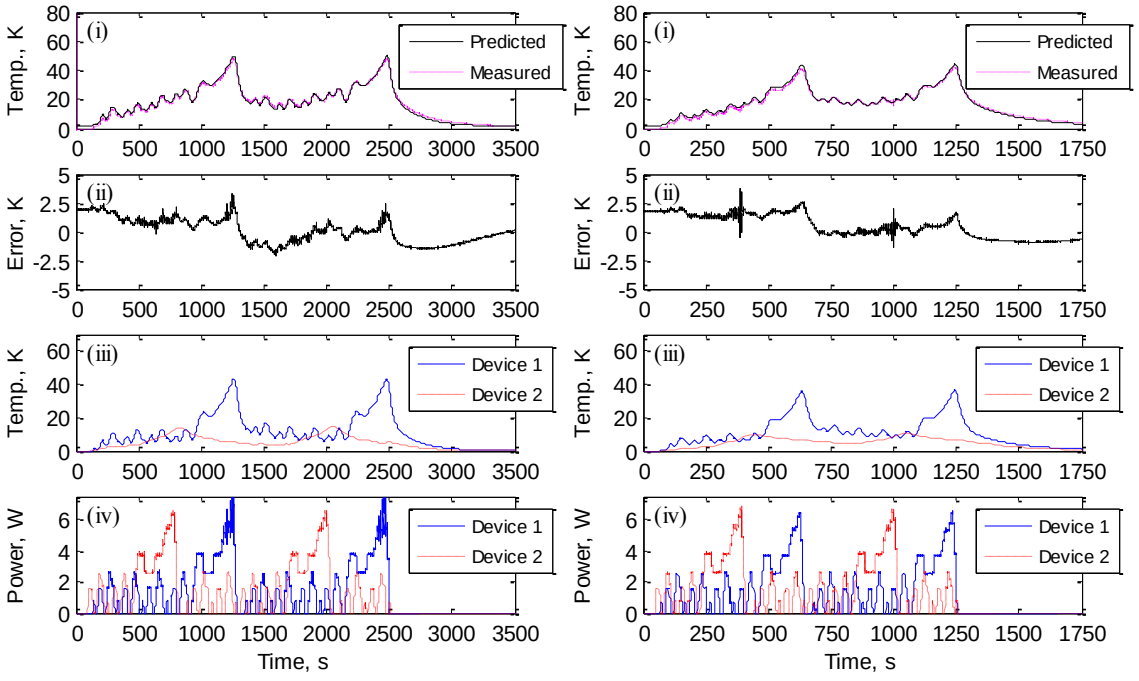

Fig. 8.6 Case temperature response of device 1 at (a) standard rate, and (b) twice standard rate. The graphs show: (i) estimated versus measured temperatures, (ii) error between estimated and measured temperatures, (iii) the temperature response due to dissipation in device 1 and device 2 separately, and (iv) the power input to the devices.

The case temperature response was estimated using the technique presented in Chapter VI to allow comparison between that technique and the die temperature estimation technique presented in this Chapter. The results are shown in Fig. 8.6. The estimations have a root-mean-square error of 1.51 K and 1.23 K compared to the measured response for the standard and double rate power inputs respectively. This error is comparable to the error from die temperature estimations despite being numerically smaller because the average temperature is also smaller. The computational requirements to fit the filter to the data are similar because both techniques use the same order of digital filter. The case temperature estimator has a reduced sampling rate due to the lower frequency cut-off, shown in Fig. 8.3, hence reducing computational requirements. The computational requirements of the die temperature estimator remain modest, however, because an IIR filter can be implemented as a set of difference equations. The die temperature estimation technique therefore has comparable accuracy to the case temperature estimator.

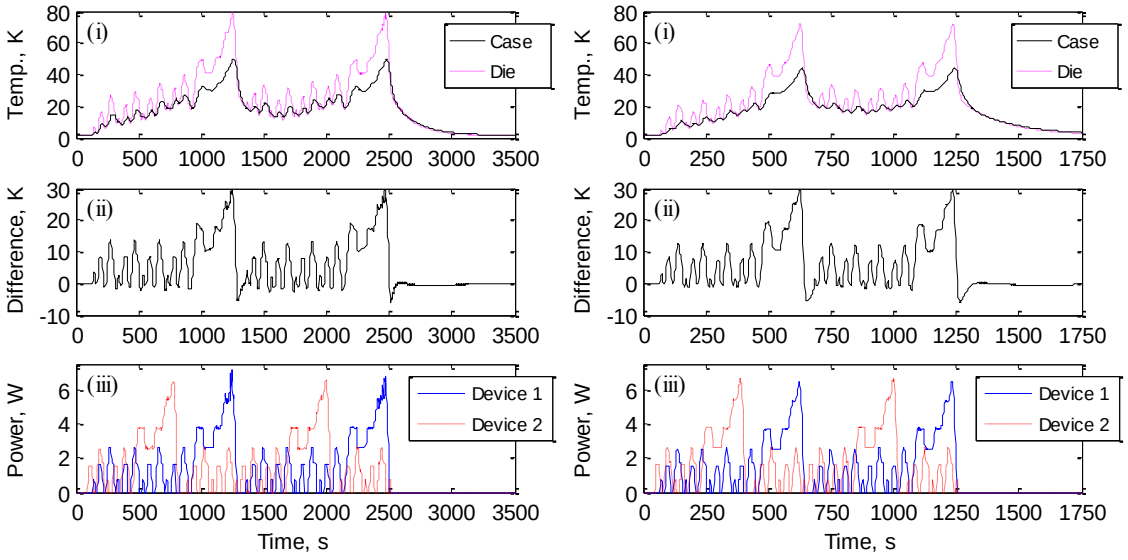

Fig. 8.7 Comparison of case and die temperature responses of device 1 at (a) standard rate, and (b) twice standard rate. The graphs show: (i) case versus die temperatures, (ii) difference between case and die temperatures, and (iv) the power input to the devices.

The differences between the case and die temperature estimations are shown in Fig. 8.7. As expected, there is a significant reduction in perturbations for the case temperature compared to the die temperature. Because of this, the die temperature also has a higher average. The difference between the case and die temperatures, shown in graph (ii), is significant with root-mean-square differences of 9.13 K and 9.12 K for the standard and double rate inputs, respectively. The difference is greatest during power dissipation where differences and transients of up to 30 K are observed. This demonstrates the importance of using die temperatures in thermal management systems because otherwise these effects would be missed.

Overall, the results show that the proposed estimator can produce accurate die temperature estimations following an offline empirical cross-coupling characterisation. By using a digital filter-based technique, the estimator can be implemented on an on-board microcontroller. Die temperature predictions are more useful for thermal management than case temperature estimations because they take account of effects from transient load changes. When thermal management is implemented, a microcontroller is required regardless of whether die temperatures are directly measured or inferred using an estimator. For this reason, and because microcontrollers are normally embedded on high-end power converters in any event [8.9], implementation of the technique requires negligible additional system cost.

8.4 Chapter conclusions

A technique which can estimate the temperature response of the semiconductor dice in a multiple device power system has been presented. By combining the near-simultaneous die temperature measurement and power controller from Chapter III with the case temperature estimator from Chapter VI, a new estimator for die temperatures was developed. The PRBS technique was used to develop a non-parametric empirical model of the thermal system from which a set of IIR digital filters were fitted. The filters are then used to estimate the temperature response of the dice of two devices on shared cooling with each dissipating a different waveform. The usefulness of the technique compared to the case temperature estimation is considered, and the importance of using the die temperature for thermal management is demonstrated. Experimental verification shows excellent agreement between measured and estimated temperature responses but

highlights the differences between approaches based on case and die temperatures, demonstrating the need for easy-to-calculate real-time die temperature estimations.

8.5 References

[8.1] D. Hirschmann, D. Tissen, S. Schroder and R. W. de Doncker, '*Reliability Prediction for Inverters in Hybrid Electrical Vehicles*', IEEE Transactions on Power Electronics, vol. 22, pp. 2511-2517, 2007.

[8.2] D. A. Murdock, J. E. R. Torres, J. J. Connors and R. D. Lorenz, '*Active thermal control of power electronic modules*', IEEE Transactions on Industry Applications, vol. 42, pp. 552-558, 2006.

[8.3] B. Wrzecionko, D. Bortis and J. W. Kolar, '*A 120°C Ambient Temperature Forced Air-Cooled Normally-off SiC JFET Automotive Inverter System*', IEEE Transactions on Power Electronics, vol. 29, pp. 2345-2358, 2014.

[8.4] K. B. Pedersen and K. Pedersen, '*Bond wire lift-off in IGBT modules due to thermomechanical induced stress*', *in Proc* Power Electronics for Distributed Generation Systems (PEDG), 2012, pp. 519-526.

[8.5] R. Wu, F. Blaabjerg, H. Wang and M. Liserre, '*Overview of catastrophic failures of freewheeling diodes in power electronic circuits*', Microelectronics Reliability, vol. 53, pp. 1788-1792, 2013.

[8.6] D. Cheney, E. Douglas, L. Liu, C.-F. Lo, B. Gila, F. Ren and S. Pearton, '*Degradation Mechanisms for GaN and GaAs High Speed Transistors*', Materials, vol. 5, pp. 2498-2520, 2012.

[8.7] MathWorks. (2013, 8 July 2013). *MATLAB 2013a documentation: invfreqz*.

[8.8] *EEC Journal Officiel No. C81*, 1990.

[8.9] S. Sirisukprasert, L. Jih-Sheng and L. Tian-Hua, '*Optimum harmonic reduction with a wide range of modulation indexes for multilevel converters*', IEEE Transactions on Industrial Electronics, vol. 49, pp. 875-881, 2002.

Chapter IX

Conclusion

In this Chapter, the work presented in the thesis is summarised and discussed as a whole. All the chapters are then brought together to provide an overall conclusion. Proposals for further work which expand this thesis are also suggested.

9.1 Discussion

This thesis has proposed a range of techniques to improve thermal model extraction for power electronics systems, estimate and predict temperatures in real time and use the estimations to control cooling.

As stated in Chapter I, the principal motivation of this work is the impetus for smaller, more power dense power electronics systems, which has come about due to the pervasiveness of power electronics in a range of applications, notably transport. Thermal management is identified as one vital area for improvement to achieve increased power density. The thesis therefore focusses on the thermal management aspects of power electronics, and offers a contribution to the state of the art, as described below.

To facilitate the research, the state of the art as reported in existing literature is presented in Chapter II. The various types of model commonly used to analyse thermal systems are described. The models are divided into two categories, namely parametric models, which are based on physical attributes of the system and include lumped parameter models and finite element analysis; and non-parametric models, which are based on empirical thermal characterisation and include compact thermal models and thermal impedance spectroscopy-based characterisation. Various methods of temperature measurement are discussed including external methods such as thermocouples and infrared sensors, and die temperature estimation methods such as estimation from the diode voltage drop and from turn-on time. Techniques for the measurement of and styles of presentation of thermal impedances are also reported.

Methods for determining model parameters from the thermal impedance are reported and the pseudorandom binary sequence (PRBS) technique is discussed in detail. Taken as a whole, this Chapter provides a foundation for the following chapters.

Chapter III describes the methodologies used for experimental control in the subsequent chapters. A novel technique to control the power dissipation in an active semiconductor device is presented first. This technique is used throughout the thesis as a means for thermal impedance measurement and the verification of temperature estimation accuracy. Methods for logging and processing data used in the thesis and typical experimental set-ups are then described. Finally, a new technique is presented to measure the temperature of the semiconductor die in a MOSFET undergoing near-simultaneous power dissipation. The technique modifies the earlier power controller to allow a small bias current to be periodically driven between the source and drain, allowing the diode voltage to be measured. The die temperature is then inferred from this voltage.

Techniques to determine the thermal impedance of a system are presented in Chapter IV. An analysis of the Cauer network, commonly used for thermal modelling, is presented. By comparing results from the model with experimentally-derived results, the minimum order of complexity required to model thermal transfer impedance is determined. This analysis uses results from pure sine wave spectroscopy, which itself is analysed to show the difficulties in measuring thermal impedances due to the inability to produce bipolar power dissipation in devices. It is shown that a substantial settling period is required for each measurement, demonstrating the advantages of wide bandwidth characterisation processes such as PRBS. A mathematical analysis of the PRBS is presented, describing its production and frequency-domain attributes. Its limitations in determining the gain of a noisy system are analysed and the minimum gain which can be identified is determined in terms of the noise present.

A novel technique to improve the noise-resilience and bandwidth of a PRBS whilst incurring only a modest increase in complexity is presented in Chapter V. The technique involves the mathematical mixing of two PRBS signals clocked at different frequencies. Signals can be mixed by a variety of mathematical operators and it is shown that the *AND* operator is most effective because it produces the signal with the most noise

resilience for a given average input level. A comprehensive mathematical analysis of the technique is provided. The technique is verified on a thermal system and results show significant improvement over traditional PRBS techniques.

Chapter VI uses the PRBS technique to characterise the thermal cross-coupling between devices in a multiple device system resembling an H-bridge. Using the technique, a separate Bode plot of the cross coupling between each pair of devices is drawn. A technique is developed which estimates the temperature response of the thermal system using frequency domain analysis and the empirical system model derived using PRBS. Although accurate, this technique is computationally intensive especially where real-time estimations are desired. An alternative technique which converts each Bode plot into an infinite impulse response (IIR) digital filter is therefore developed. The new technique produces estimations with equivalent accuracy while being computationally light and well-suited to microcontroller-based real-time implementations. In addition, accurate temperature estimations based on assumed future power dissipation can easily be produced.

The techniques developed in Chapters V and VI are combined in Chapter VII to produce an on-line temperature estimation methodology based on an off-line mixed-PRBS thermal characterisation. The scope of the estimation technique is expanded to estimate the temperature response of a thermal system under dynamically changing cooling. As in Chapter VI, IIR digital filters are produced for each pair of devices and used to estimate the temperature response; however, in this Chapter, separate filters are produced for each level of dynamic cooling over a quantised range. When the level of cooling changes, the filters in use are substituted for filters reflecting the new cooling level and their internal states are recalculated. Two options for recalculation are considered and a method which maintains the state of every possible filter is recognised as accurate in most cases. The developed technique is then applied to cooling control and PID control (tuned using the estimator) is implemented to reduce thermal cycling.

The temperature characterisation and estimation techniques developed in Chapters IV to VII use device surface temperatures, measured using thermocouples, as the temperatures of interest. In reality, it is the junction temperature which is most relevant in the typical applications for which temperature estimation is useful. Chapter VIII

therefore applies these characterisation and estimation techniques using the die temperature. The near-simultaneous dissipation and die temperature measurement technique developed in Chapter III is used during the characterisation and verification stages. The surface temperature-based method from Chapter VI is also performed, and the improvement offered by the die temperature method is highlighted. A die temperature estimator is implemented with similar computational complexity, but improved relevance, compared to the technique in Chapter VI.

The research encompassed in this thesis has been published in five journal and three conference papers to date, with a further two journal papers in preparation, demonstrating the relevance of the findings to the power electronics community.

9.2 Conclusions

This thesis set out to develop thermal measurement and management techniques to improve the power density of next generation power electronics. It subsequently developed a range of techniques which provide these improvements. Novel power control and temperature measurement techniques were developed for experimental control while a new analysis of PRBS-based spectroscopy was provided to identify the limitations of noisy measurements. These facilitated a novel mixed PRBS characterisation method and novel temperature estimator for a dynamically cooled multiple device system. The thermal measurement techniques presented can be used for improved thermal management by providing low-cost and accurate temperature data across all devices on a system. Improved management is demonstrated when a temperature controller is fitted to the estimator to reduce thermal cycling in a practical system.

Taken together, these techniques will allow design engineers to develop more power-dense systems through better lifetime estimation and reduced safety-margins in system design. These improvements are facilitated by the increased availability of real-time temperature data and easy-to-calculate estimations of future temperatures based on per-device power dissipation as a result of this work. Where temperatures are likely to exceed design parameters, system adjustments can be used to change system load or cooling in order to reduce temperature and thermal cycling to improve reliability.

Alternatively, a smaller, more power dense system can be constructed with the same reliability.

9.3 Further work

Although this thesis presents a coherent body of work which offers improvement over the state of the art, it is recognised that further improvements could be made to increase the impact of the work. The following subsections summarise a number of areas for further development which leads on from this thesis.

9.3.1 Reliability modelling and experimental verification

Reduced thermal cycling and higher average temperatures in power electronics are some of the outcomes which could result from the application of the techniques described in this thesis. The overall effect of the changes on wider power electronics applications is an area of further work which will allow the impact of the proposed techniques to be evaluated. To perform this evaluation, a reliability model can be developed for the modified system. It is an increase in reliability for an otherwise unchanged system that will lead to the increased power density proposed in order to maintain reliability at previous levels. The reliability model will take into account the effect of thermal cycling and the average temperature of each device on the overall reliability of the system. A technique for on-line lifetime prediction could be developed to allow the prediction of remaining system life given estimated and predicted temperature response and thermal cycling profiles of the devices.

Experimental verification of the consequential improvements to system reliability as a result of the techniques proposed in this thesis could additionally be undertaken. Verification would consider the reliability of a system implemented with the proposed temperature estimators, controlled cooling and the suggested load-limiting technique against a typical system in commercial use. This verification could also provide data to construct the lifetime models described above.

9.3.2 Fully cross-coupled cooling controller

Chapter VII reported a temperature control technique based on a PID controller for a single device. While this has clear applications for a single device cooling system, many

real-world systems consist of multiple devices with differing input power waveforms. The controller could therefore be extended to work with the fully cross-coupled temperature estimator developed in the first half of Chapter VII.

Such a system would need to use a lifetime model for each device on the system, taking the temperature estimations as an input to the model. Cooling would then be controlled to increase the longevity of devices through manipulation of temperature. A further extension to the technique could be implemented by considering the case of per-device modulated cooling where the intensity of a cooling system can be changed for one device independently of the others. Such a control system could therefore manipulate device temperatures independently.

The accuracy of estimations could be improved using measured temperature data from the system, such as from a central temperature sensor or indirect measurements of temperature-dependent component values. Kalman filtering could be used to combine the data from multiple sensors in order to improve estimator accuracy and take account of changes in ambient conditions.

9.3.3 Implementation of load limiting techniques

One of the advantages offered by the real-time temperature estimator is the ability to estimate the future temperature response of a system to simulated input. It was proposed in Chapter VI that such estimators could be used to adjust system load and hence avert damaging over-temperature conditions. The implementation of such a system is another area for further development.

One example might be domestic electric vehicles. Under normal load, demand is modest with high demand only occurring for short periods, for instance, during acceleration as part of an overtaking manoeuvre. However, under abnormal conditions such as high velocity on a steep incline, high power dissipation in components is sustained and may lead to over-temperature conditions. An embedded controller could estimate the likelihood of devices exceeding thermal ratings and hence mitigate this risk. If the controller predicts that thermal ratings will soon be exceeded if high power dissipation is sustained, the maximum wheel torque (and therefore speed) of the vehicle

can be reduced to a safe level. Because such conditions are rare in practice, the overall impact on driving quality will be low. Safety implications will have to be considered but a managed reduction in speed is unlikely to pose significant problems. The reduction in system size due to reduced cooling demands will allow the efficiency and capital cost of the vehicle to be reduced.

9.3.4 Application of the die temperature estimator to cooling and load control

The cooling control in Chapter VI was implemented for estimated surface temperatures of the power devices. In practice, it is control of the temperature of the active junction on the semiconductor die which will lead to the improvement in system reliability. The die temperature estimation techniques proposed in Chapter VII are therefore relevant to the subject of cooling control.

One limitation of this technique is the long time constants observed for a step change in cooling. These constants are usually longer than the constant observed for a change in die temperature due to a step change in input power. This limitation is a barrier to die temperature control because the effect of a change in cooling cannot match the speed of a power step. To mitigate this problem, it may be possible to pre-empt a step change in power by increasing cooling prior to the change (supposing there is sufficient warning). Research into the practicality and implementation of this suggestion is an area for further work.