

Real-Time Temperature Estimation in a Multiple Device Power Electronics System Subject to Dynamic Cooling

Jonathan N. Davidson, David A. Stone, Martin P. Foster, and Daniel T. Gladwin

Abstract—This paper presents a technique to estimate the temperature of each power electronic device in 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 quantized range of active cooling levels using a technique based on pseudorandom binary sequences. The technique is illustrated by an application to the case temperatures of power devices. For each cooling level and pair of devices, a sixth-order 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 infrequent changes in power dissipation and cooling, and a more accurate method, which requires increased memory and processing capacity. Results show that the temperature can be estimated with low error using a system which is suitable for integration on an embedded processor.

Index Terms—Estimation, infinite impulse response (IIR) digital filters, pseudonoise processes, temperature, thermal variables measurement.

I. INTRODUCTION

DEMANDS on the thermal management of power converters are increasing because of a strong industrial impetus for smaller, more efficient power electronics systems [1], [2]. Design engineers must, therefore, build systems using physically smaller components with less heatsinking. However, the extent to which system designs can be miniaturized is limited by the rated maximum component temperatures, which may not be exceeded during normal operation. In addition, the life expectancy of power converters is degraded by thermal cycling, where damage from multiple heating and cooling events accumulates until system failure results [3]. A simple method of mitigation, 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 [4]. 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 prolonged operating periods. This is an unusual situation, however, and, under practical conditions, the system operates significantly below the worst case. The design for worst case leads to increased cost due to excess heatsink requirements and increased energy usage in the active cooling system. Second, where active cooling is not controlled, temperatures rapidly increase and decrease with system load, which can lead to significant thermal cycling [5].

In automotive applications, designing for thermal conditions below worst case is a reasonable approach if measures are taken to mitigate the risk from overheating semiconductor devices. Such measures include reducing the load and increasing cooling if the device temperatures would otherwise be exceeded. In electric vehicles, load reduction could be achieved by a temporary lessening of wheel torque or speed, where a small reduction will have limited impact and the mechanism to achieve it can safely be incorporated into the vehicle management system for most categories of vehicle. To allow effective mitigation, a technique to monitor the temperature of critical devices is required. While temperature sensors could provide the required information, they represent an additional per-unit cost to the system. Instead, in this paper, we propose a mathematical real-time temperature estimator based on the power dissipation in devices within the system. Additionally, the effect of variable cooling such as from a fan or blower is considered.

Early work on estimating and predicting the temperature of components under variable cooling used an analytical approach, such as that proposed by Fried [6]. This type of modeling 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. More recently, Lundquist and Carey [7] use an analytical model to predict the future temperature of a computer processor under a prospective level of forced air cooling. Together with direct measurement of the actual processor temperature, the technique is used to control the temperature at an optimal level. In this study, we eliminate the need for direct measurement in similar systems by implementing accurate temperature prediction.

Musallam and Johnson [8] propose using compact models between each input–output pair in a thermal system with multiple points of interest. Using a combination of finite-element analysis (FEA) and direct experimental characterization, the thermal

Manuscript received September 3, 2014; revised April 30, 2015; accepted June 2, 2015. Date of publication June 9, 2015; date of current version November 30, 2015. This work was supported by the EPSRC under Grant EP/G037477/1. Recommended for publication by Associate Editor T. M. Lebey.

The authors are with the Department of Electronic and Electrical Engineering, The University of Sheffield, Sheffield S1 3JD, U.K. (e-mail: jonathan.davidson@sheffield.ac.uk; d.a.stone@sheffield.ac.uk; m.p.foster@sheffield.ac.uk; d.gladwin@sheffield.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TPEL.2015.2443034

transfer function between each input–output pair is calculated in the Laplace domain. From this data, a matrix is constructed that describes the relationship between every heat source input and every temperature-of-interest output. By transforming the model into a discrete-time expression, a computationally efficient real-time temperature estimator is developed. Experimental validation confirms the accuracy of the estimations; however, the work is limited to thermally linear systems, where the transfer functions are unchanged throughout operation. The technique cannot, therefore, be used directly in a system with variable cooling.

A solution that overcomes this problem is proposed by James *et al.* [9], who consider the effect that variable forced-air cooling has on thermal dynamics. The authors propose that for a multiple device system, a series of thermal impedance matrices can be found, where each matrix describes the system dynamics at a given cooling level. However, for the system examined, it is found that the model can be simplified by neglecting the dependency of thermal cross coupling on the level of cooling, thus simplifying the mathematics. The authors identify this as a limitation of their approach. In the present paper, an extended approach that overcomes this simplification and explicitly considers the effect of cooling level on the thermal cross coupling is therefore proposed.

This paper reports a measurement-based characterization of the system. Alternatives to this approach include the offline modeling of a system using FEA and computational fluid dynamics [10]. 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 modeling of turbulent flow is difficult [11]. It is shown by Rodgers *et al.* [12] that more practical thermal analysis is preferable in many cases. In this paper, we therefore present an experiment-based characterization of a power electronic system under controlled cooling.

Previously [13], we presented a method to estimate the temperature over several devices coupled on a heatsink by characterizing the cross coupling between them. In this paper, a similar technique is used wherein a pseudorandom binary sequence (PRBS) power dissipation waveform is applied to each power device in turn, and the resulting temperature response at all devices is measured. Due to the special nature of the PRBS (which we will explore later in this paper), it is possible to determine the cross coupling (i.e., the thermal transfer impedance) between each pair of devices over a band of frequencies. In this study, the cross coupling is used to estimate the temperature response to a given power input and cooling level using a novel technique that results in accurate estimation. The particular topology studied in this study concerns temperature estimations with an update rate of 1 Hz or less, and is, therefore, sensitive to temperature changes due to load, rather than individual switching events. To illustrate the technique, the study is applied to device case temperatures, although the presented estimation techniques could equally be applied to the junction temperature following a junction-temperature-based characterization.




Fig. 1. Norton equivalent electrical circuit between heat dissipation point x and temperature measurement point y .

II. THERMAL SYSTEM CHARACTERISATION

A. Cross Coupling

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, we are 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 that dissipate high power and are damaged by overheating. To model the thermal cross coupling, we consider the problem 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. 1.

We use Q_x to denote the power dissipated at point x , Θ_y for the temperature at point y and $Z_{th.x \rightarrow y}$ for the cross coupling between points x and y . To simplify the mathematics, we use these quantities in the frequency domain, which we signify 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 (1) and (2). It is possible to use simple point-by-point multiplication in these calculations because the mathematics is carried out in the frequency domain

$$\Theta(j\omega) = Z_{th}(j\omega) \cdot Q(j\omega) \quad (1)$$

$$Z_{th} = \begin{pmatrix} Z_{th.1 \rightarrow 1} & \cdots & Z_{th.n \rightarrow 1} \\ \vdots & \ddots & \vdots \\ Z_{th.1 \rightarrow m} & \cdots & Z_{th.n \rightarrow m} \end{pmatrix}; \Theta = \begin{pmatrix} \Theta_1 \\ \vdots \\ \Theta_m \end{pmatrix} \quad (2)$$

$$Q = \begin{pmatrix} Q_1 \\ \vdots \\ Q_n \end{pmatrix}$$

where $\Theta(j\omega)$ and $Q(j\omega)$ are column vectors of temperature and powers at each relevant point, respectively, $Z_{th}(j\omega)$ is the matrix of every cross coupling impedance between each set of 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, we use




Fig. 2. Experimental arrangement.




Fig. 3. Frequency domain representation of an 8-bit maximum length PRBS signal clocked at 1 Hz. The response was calculated using the DFT.

active cooling from a blower with a controllable angular frequency δ . To demonstrate the techniques presented in this paper, a test system of D2PAK devices mounted on an insulated metal substrate board is constructed. The power dissipation in each device is controlled by an external circuit. In order to perform the characterization measurements and to verify the temperature estimations generated by the proposed technique, a K-type thermocouple is attached to the tab of each device to monitor the temperature. A speed-controlled blower is positioned to blow over all devices and an external computer controls and monitors all elements. The physical layout is shown in Fig. 2.

When active cooling is present, the cross-coupling matrix Z_{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 as $Z_{th}(j\omega, \delta)$. Determining the elements of this matrix over all blower speeds and frequencies produces a complete model of the system.

B. Generation of Cross-Coupling Characteristics using PRBSs

In previous work [13], [14], we have demonstrated that PRBS can be used to characterize the thermal cross coupling between devices. Using the same principles, the cross coupling is calculated over a range of blower speeds in this study. The PRBS technique is essentially a practical form of band-limited white noise system identification [15]. It uses a PRBS test signal, which has almost uniform frequency content over its bandwidth as shown in Fig. 3. For thermal analysis, the relationship between power and temperature under given operating conditions (e.g., cooling level, chemical phase of materials) is approximately linear because heat transfer is a diffusion process described by a linear equation [8]. Therefore, when a PRBS power signal is applied to the thermal system, the temperature response spectrum

reflects the system's frequency response over the test signal's bandwidth.

A PRBS signal is applied to each device in turn and the temperature responses between every pair of devices are measured using a thermocouple and logged by computer. This experiment is repeated over a range of blower speeds to fully characterize the cross coupling. The cross-coupling thermal impedance $Z_{th \cdot x \rightarrow y}$ between devices x and y at any frequency and blower speed can be calculated according to

$$Z_{th \cdot x \rightarrow y}(j\omega, \delta) = \frac{\Theta_y(j\omega, \delta)}{Q_x(j\omega)} \quad (3)$$

where Q_x is an n -bit PRBS power signal applied to device x which, when clocked at frequency ω_c , produces a valid temperature response Θ_y at device y over the bandwidth [14] in

$$\frac{\omega_c}{2^n - 1} \leq \omega \leq \frac{\omega_c}{2.3}. \quad (4)$$

In this paper, a 9-bit ($n = 9$) PRBS is used clocked at 0.25 Hz, leading to a valid frequency response between 0.5 and 110 mHz. The experimental arrangement to perform the characterization is shown in Fig. 4. A microcontroller generates a PRBS voltage waveform (using the linear feedback shift register method with taps on the fifth and ninth registers to produce a maximum length sequence). This voltage controls a current source that determines the power dissipation in the D2PAK resistors through control of I^2R losses. The temperatures of the devices are measured using a thermocouple and both power input and temperature output are logged by computer. The characterization procedure is repeated at 22 blower speeds distributed with approximate logarithmic spacing in the range of 0 and 6600 revolutions per minute (rpm). Experiment control and data processing are handled by a PC running LABVIEW and MATLAB.

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. 5, therefore, shows a three-dimensional (3-D) Bode plot of $Z_{th \cdot 1 \rightarrow 1}(j\omega, \delta)$ and $Z_{th \cdot 3 \rightarrow 1}(j\omega, \delta)$. The $Z_{th \cdot 3 \rightarrow 1}(j\omega, \delta)$ plot shows increased noise at high frequency due to the low-thermal impedance from cross-coupling characteristics. The effect is negligible because the temperature response of a system is determined mainly by the lower frequency region.

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 characterization 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 III-B.

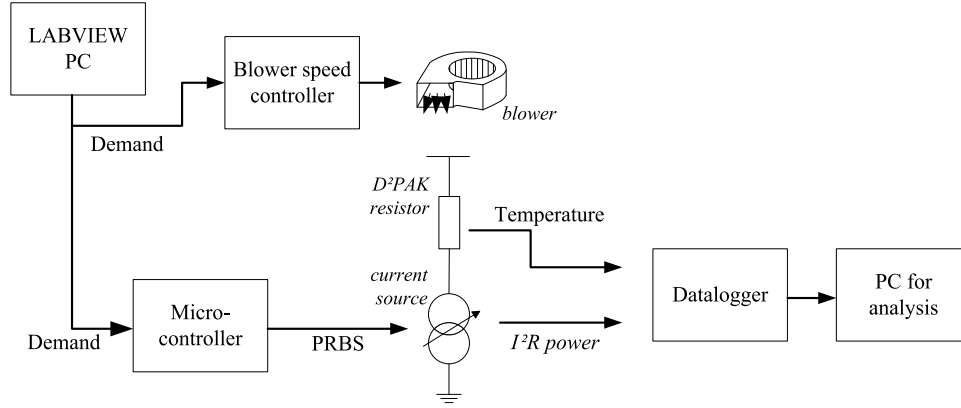


Fig. 4. Experimental setup for characterization of devices.

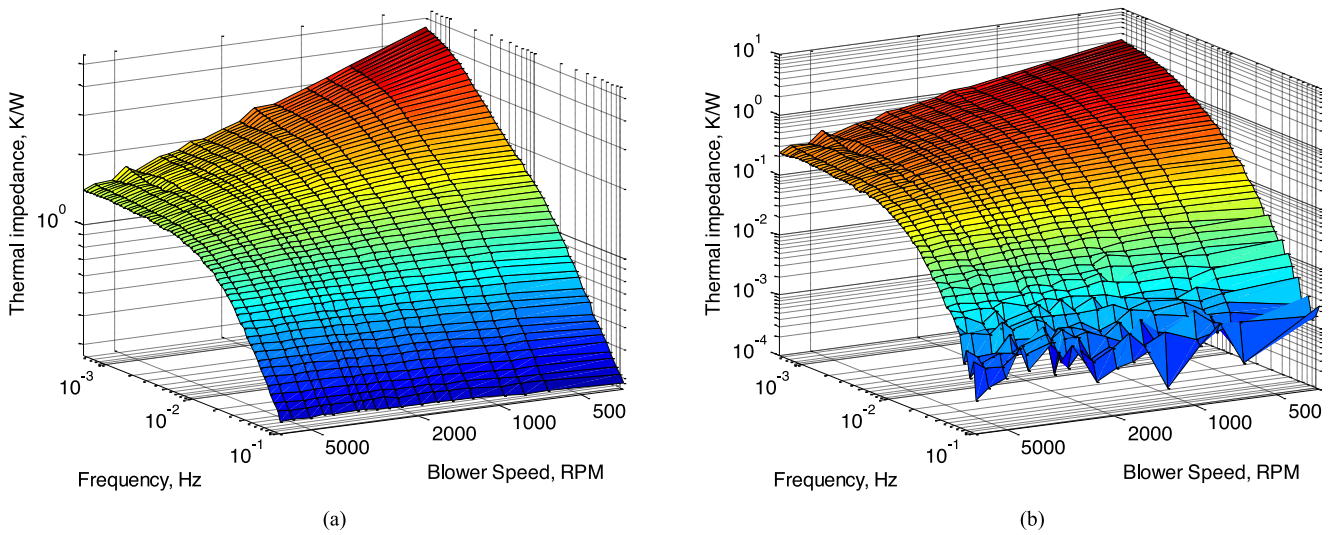


Fig. 5. Three-dimensional Bode plots of the autocoupling of device 1 over the available range of blower speeds and frequencies. Note that all axes are logarithmic. (a) $Z_{th.1 \rightarrow 1}$. (b) $Z_{th.3 \rightarrow 1}$.

III. ESTIMATION OF THE TEMPERATURE RESPONSE

A. Estimation Using the Cross-Coupling Frequency Response Directly

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

$$\Theta_y(j\omega, \delta) = Z_{th.x \rightarrow y}(j\omega, \delta) \cdot Q_x(j\omega). \quad (5)$$

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 (6) and (7), respectively,

$$\theta_y(t) = \mathcal{F}^{-1}(Q_x(j\omega)) \quad (6)$$

$$Q_x(j\omega) = \mathcal{F}(q_x(t)) \quad (7)$$

where \mathcal{F} represents the DFT (\mathcal{F}^{-1} is the inverse), and θ and q represent the time-domain temperature and power, respectively.

B. Representation of Cross Coupling as a Digital Filter

In [13], it has been shown that Fourier transforms are computationally intensive and require a significant number of mathematical operations. 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 a and b such that Z_{th} in (8) is the best approximation to the measured cross coupling [16]. This process effectively converts frequency-domain $Z_{th}(j\omega)$ to a digital filter with the equivalent transfer function in the z -domain, $Z_{th}(z)$. A different filter will result for each level of active cooling and each pair of power input and temperature measurement points

$$Z_{th}(z) = \frac{\sum_{\lambda=0}^H b_{\lambda} z^{-\lambda}}{1 + \sum_{\lambda=1}^G a_{\lambda} z^{-\lambda}}; \quad z = e^{-j2\pi(\frac{\omega}{\omega_s})} \quad (8)$$

where ω_s is the sampling frequency, and H and G are the polynomial orders of filter vectors b and a , respectively. The output

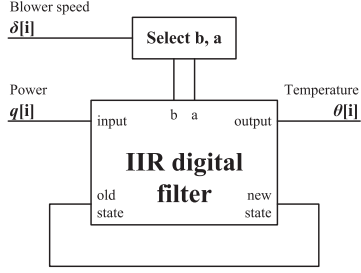


Fig. 6. Simple digital filter arrangement for temperature estimation with constant active cooling.

of the filter to an arbitrary input sequence can be calculated using (9). This equation can also be expressed as the time-domain difference equations shown in (10) and (11) [17]. In this form, the input power at each time step can be evaluated individually, with the system state stored as a state vector $s = (s_0, \dots, s_{L-1})$, where L is the greater of G and H ($a_\lambda = 0$ for $\lambda > G$ and $b_\lambda = 0$ for $\lambda > H$). In this study, $G = 3$ and $H = 6$ are selected as suitable filter lengths since these are the smallest values produce accurate results

$$\theta[i] = \sum_{\lambda=1}^G a_\lambda \theta[i-\lambda] + \sum_{\lambda=0}^H b_\lambda q[i-\lambda] \quad (9)$$

$$\theta[i] = q[i] \cdot b_0 + s_0[i-1] \quad (10)$$

$$s[i] = \begin{pmatrix} s_0 \\ \vdots \\ s_{L-2} \\ s_{L-1} \end{pmatrix} [i] = q[i] \cdot \begin{pmatrix} b_1 \\ \vdots \\ b_{L-1} \\ b_L \end{pmatrix} + \begin{pmatrix} s_1 \\ \vdots \\ s_{L-1} \\ 0 \end{pmatrix} [i-1] - \theta[i] \cdot \begin{pmatrix} a_1 \\ \vdots \\ a_{L-1} \\ a_L \end{pmatrix}. \quad (11)$$

The output temperature at the next time step is calculated using (10) with this value then being fed back to allow the calculation of the new state using (11). These operations are repeated at each time step i with the state vector 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. 6.

IV. CALCULATING THE TEMPERATURE RESPONSE UNDER DYNAMIC 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 modeled state of the system remains consistent as filter parameters b and a are substituted. Fig. 7 shows the procedure used.

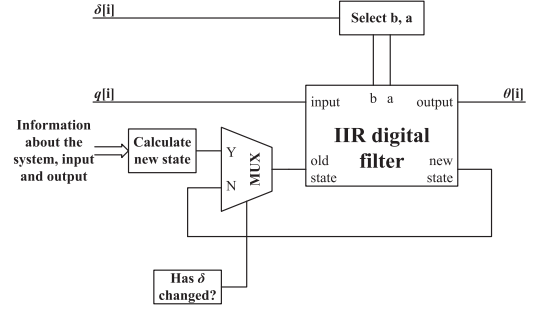


Fig. 7. Digital filter arrangement showing the need to recalculate the filter state vector after a change in active cooling.

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.

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 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 b and a 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 [18] and gain scheduling for nonlinear controllers whose operating domain has been divided into several approximately linear subdomains [19]. 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 advanced control techniques such as these are not required. 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.

A. Proposed Methods for Bumpless Digital Filter Transition

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 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 later.

1) *Steady-State Assumption Method*: The simplest method to generate the internal state of a new filter is to initialize the state

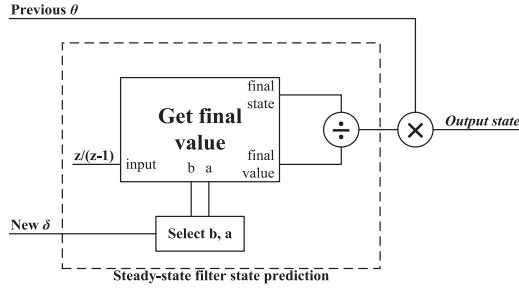


Fig. 8. Filter state recalculation technique for the steady-state assumption method. This replaces *calculate new state* in Fig. 7.

vector s to the appropriate steady-state level for the last output of the previous filter. This vector is found using the z -domain final value theorem [20], which states that the final value of the sequence can be calculated as

$$\begin{aligned}\theta[\infty] &= \lim_{t \rightarrow \infty} \theta(t) = \lim_{z \rightarrow 1} (z-1) \Theta(z) \\ &= \lim_{z \rightarrow 1} (z-1) Z_{th}(z)Q(z).\end{aligned}\quad (12)$$

To calculate the final value steady state in steady state, $q(t)$ is 1 W for all t , therefore [17]

$$Q(z) = \frac{z}{z-1}.\quad (13)$$

Having determined $\theta[\infty]$, the values of the of the state vector can be calculated by solving (11) for a constant state vector and the final temperature, as shown in

$$\begin{aligned}s_{L-1} &= b_L + 0 - \theta[\infty] a_L \\ s_{L-2} &= b_{L-1} + s_{L-1} - \theta[\infty] a_{L-1} \\ &\vdots \\ s_0 &= b_1 + s_1 - \theta[\infty] a_1.\end{aligned}\quad (14)$$

The technique is shown in Fig. 8. The state vector is normalized by dividing the state by the final temperature, θ . Upon filter change, the normalized 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 [17].

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 stabilized or where there is no significant delay between power input and response. The method is useful, however, because the steady-state vector s can be calculated efficiently.

To illustrate the operation of the technique in practice, device 1 shown in Fig. 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 rpm.

Two types of input were selected to illustrate the abilities and limitations of the technique, as described later.

- 1) A power waveform consisting of three repeats of the European driving cycle waveform [21] 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. 9.

For the gradual changes in Fig. 9(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 1.9 K. In the case of the device 3 (cross coupling), however, there is a significant difference between estimated and practical results demonstrated by an RMSE of 4.9 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. For frequently changing input, erasing the memory by repeatedly reinitializing the filter to steady state invalidates the estimations. The issues are highlighted where only irregular sudden changes in blower speed occur, such as in Fig. 9(b). In this case, the error is similar in both the auto and cross coupling cases (with RMSEs of 2.1 and 1.3 K, respectively) because there is no significant memory loss. For this reason, a technique that does not erase memory of previous inputs is required for accurate temperature estimations under constantly changing cooling in the case of cross coupling.

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 old and new filter states 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 normalization procedure is implemented, where there is a change of blower speed. The full procedure for the scaled input assumption method is shown in Fig. 10.

The full range of blower speeds is quantized 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 b and a , each also maintains a separate state vector s . The outputs of all the filters are multiplexed and the output from the filter whose δ is closest to the present level

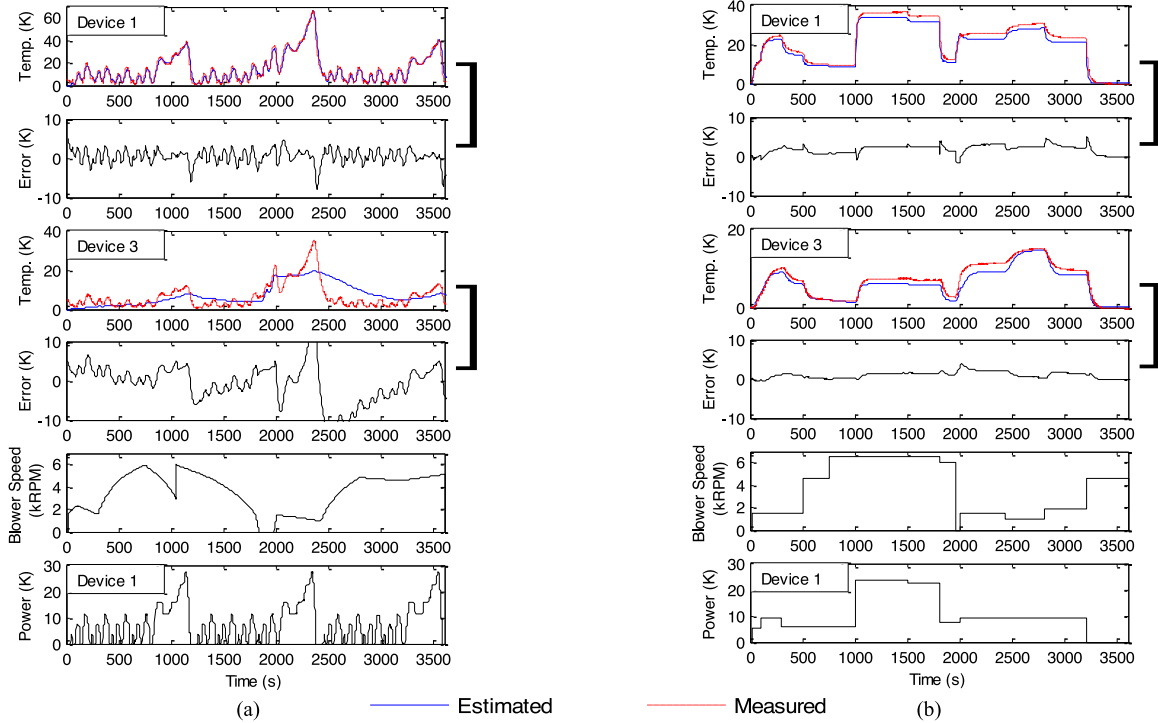


Fig. 9. Estimated and actual temperature responses at device 1 (auto coupling) and device 3 (cross coupling) due to power dissipation in device 1 under variable blower speed for the steady-state assumption method. Large square brackets between graphs link a device temperature result with its error. Results for system under (a) gradual frequent changes of blower speed, and (b) pronounced infrequent blower speed changes.

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 III-B, 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. 10 is set), all the filters have their output normalized 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. The scaled assumption 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 s in our example). This allows sufficient time to recalculate each filter using a typical microcontroller. The computational requirements are described in Section V.

For verification, the scaled input assumption method was applied to the same arbitrary power and blower waveforms as in the previous section. The model was quantized to one filter for each multiple of 50 rpm to reduce memory and processing requirements. Results are shown in Fig. 11. For the auto coupling response of device 1, the method shows little improvement over the steady-state method, with similar RMSEs of 1.2 and 1.4 K for Fig. 11(a) and (b), respectively. For cross coupling, however, there is a significant improvement in the closeness of the estimation for the regularly changing

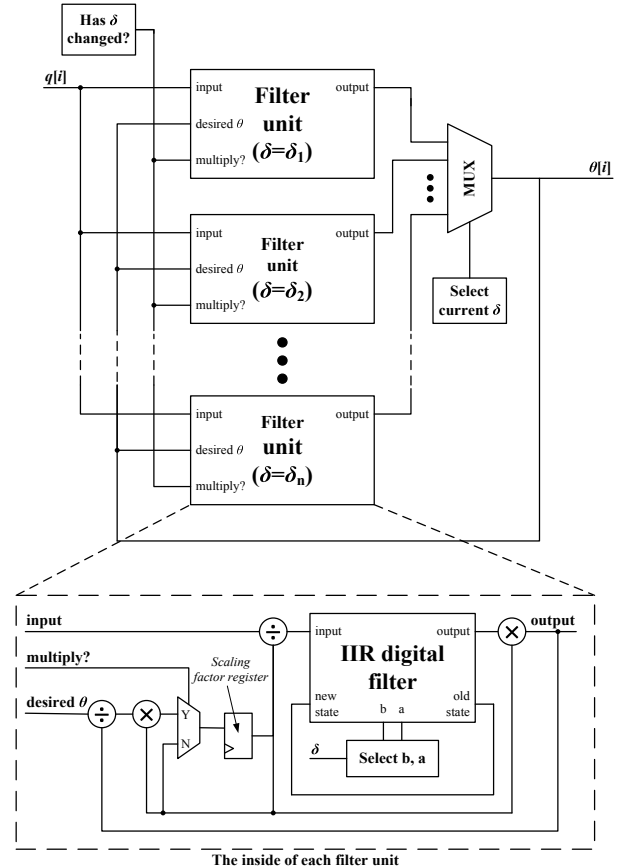


Fig. 10. Digital filtering in the scaled input assumption method.

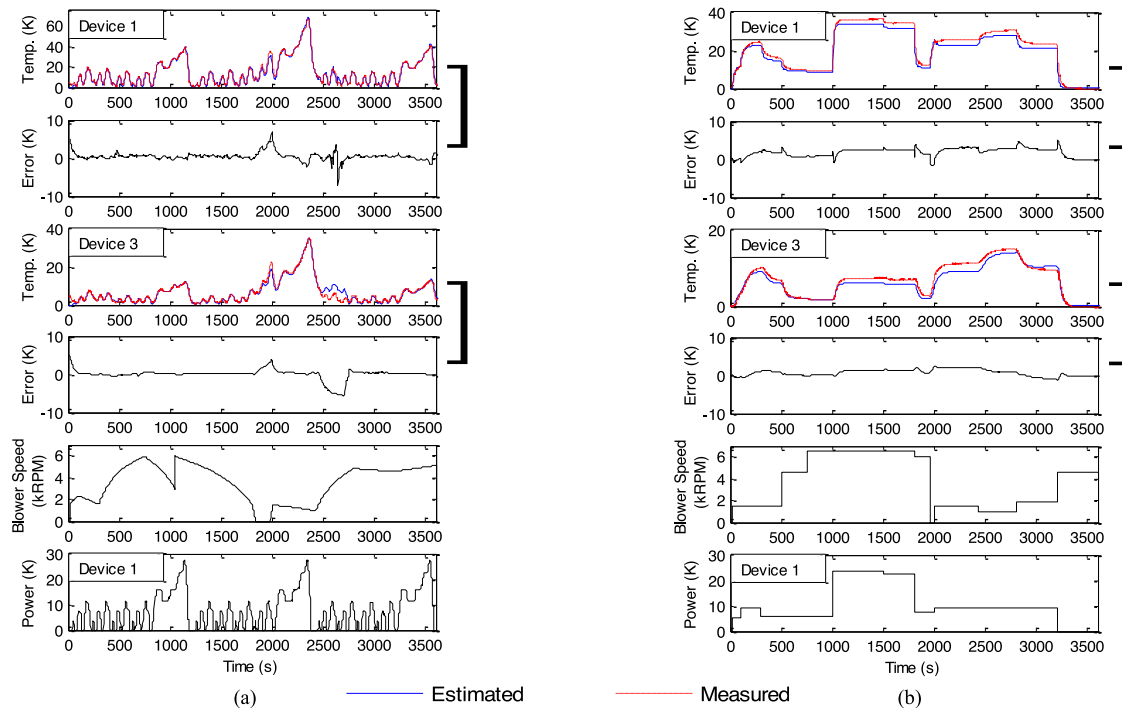


Fig. 11. Estimated and actual temperature responses at device 1 (auto coupling) and device 3 (cross coupling) due to power dissipation in device 1 under variable blower speed for the scaled input method. (a) Under gradual frequent changes of blower speed. (b) Under pronounced infrequent blower speed changes.

blower speed seen in Fig. 11(a). The error is similar to that of the auto coupling response at 1.4 K, a 70% reduction over the steady-state technique. In addition, the shapes of the curves are similar. The irregular power and blower speed test in Fig. 11(b) shows similar results to the steady-state assumption technique (RMSE of 1.1 K), confirming that regular changes in blower speed increase the need for complex filter conversion.

B. Temperature Response With Multiple Devices Dissipating

Results presented so far are for a single device dissipating power. However, the techniques presented in this paper are equally applicable in cases where multiple devices are dissipating. Fig. 12 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 (2) by adding the individual estimated temperature response at device 1 due to dissipation at device 1 alone to the response due to device 3 alone. For this estimation, the scaled input assumption method was used.

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. Errors compared to practical results are small with occasional increases during periods of rapid temperature change. The profile of the temperature estimated and practical responses are the same with correctly estimated peaks and gradients, making the results appropriate for reliability modeling and cooling control. This is confirmed by the RMSE of 2.0 K, which represents an insignificant error for most practical purposes.




Fig. 12. Estimated and actual temperature responses at device 1 due to equal simultaneous power dissipation in devices 1 and 3 under variable blower speed.

V. PRACTICAL IMPLEMENTATION ON A MICROCONTROLLER

The proposed estimation technique was implemented on a standard low-cost microcontroller to demonstrate that the technique can successfully be implemented for a typical system. The method was implemented in C++ on an ATmega1284P microcontroller and compiled using the AVR-GCC programme. The microcontroller is a single-cycle 8-bit Harvard architecture machine with a two-cycle hardware multiplier but no hardware support for division or floating point operations (which are

TABLE I
COMPUTATIONAL REQUIREMENTS PER DEVICE PAIR FOR THE
PROPOSED METHODS

Method	Steady-state	Scaled input	Unit
Programme memory required (for code)	5262	5724	byte
Flash data memory required (for filter coefficients)	5808	5808	byte
Read-write memory required (for static variables)	30	4230	byte
Floating point operations per sample			
(multiplications)	13	1855	
(divisions)	2	139	
(faster operations, e.g., addition)	16	1804	
Processing time per sample (at 8-MHz clock)			
(average)	1.0	87	ms
(maximum)	1.7	225	ms
Maximum sample rate (worst case)	590	4.4	S/s
RMSE compared to MATLAB code	3.4	64	mK

Conditions: 4.6-s sampling period; 132 quantized cooling levels for both methods.

instead implemented in software). This low cost and moderate performance microcontroller was selected to illustrate the low-computational requirements of the proposed technique. Indeed, although capable of 20 MHz, the microcontroller was clocked using the internal 8-MHz RC timer to further demonstrate that the proposed technique can be implemented on basic hardware. Where the requirements of a particular application require a higher iteration rate, higher specification microcontrollers are economically available.

The C++ code was a functionally identical rewrite of the MATLAB development code, with platform alterations to account for the limitations of a microcontroller. In particular, the polynomial coefficient vectors b and a are stored in the flash memory (which can be written to by the software) and all data is stored as single-precision (32-bit) floating point numbers. For test purposes, the response to the European driving cycle power with gradually but infrequently changing blower speed inputs, as used in Fig. 9(a) in Fig. 11(a), was calculated using the microcontroller. Near identical results were obtained (with a RMSE between MATLAB-based results and microcontroller generated results of less than 63 mK, as a result of floating point precision differences).

Table I shows the computational resources required to implement the program and calculate the temperature response on the microcontroller. Flash data memory, where the filter vectors b and a are stored, is identical for both methods because the same filters are used. The scaled input assumption method requires significantly more read–write memory and floating point operations per sample to maintain the additional state vectors needed by that method. The difference between the program memory requirements is insignificant because the code base is very similar. The scaled input assumption method is significantly slower (allowing only 4.4 samples/s for the worst case, compared to 590); however, this is more than sufficient for typical systems such as the demonstrator system. To achieve greater speeds, a dedicated digital signals processing controller could instead be used. Because the computational requirements of both methods are fulfilled by typical microcontrollers, the scaled input assumption method is preferred since it produces better results.

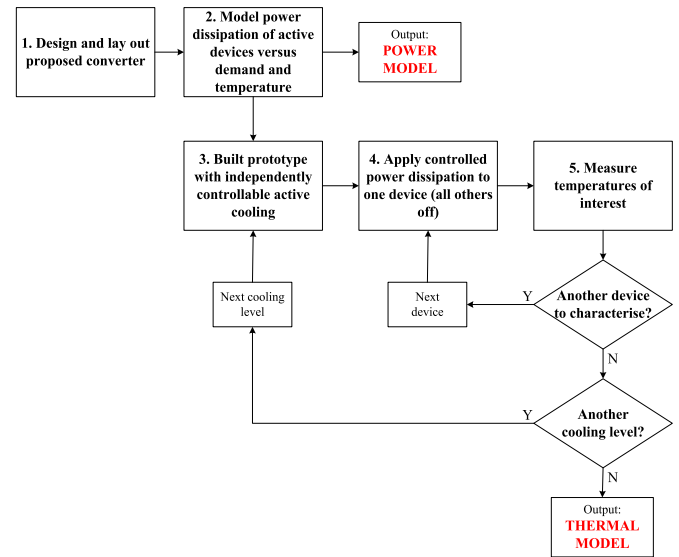


Fig. 13. Design stage procedure to generate models and characteristics.

VI. PROCEDURE FOR APPLICATION TO A POWER CONVERTER

The techniques discussed in the paper are applicable to a wide range of thermal systems under variable active cooling. The procedure for implementing the technique can be divided into two parts: (A) design stage modeling and characterization and (B) implementation in the operational system. The considerations for these stages are described later.

A. Design Stage Modeling and Characterization

Fig. 13 presents a suggested procedure for implementing the technique for a typical power converter. Initially, the power converter is designed and laid out according to normal engineering principles. Because the proposed temperature estimation technique requires real-time power dissipation data as an input, an electrical power model of the circuit is required. This model can be developed in a number of ways and the literature has extensive examples of viable approaches. For example, Capua and Femia [22] produce a computationally efficient model for a MOSFET switching circuit by first principles analysis of the structure of a MOSFET. Orabi and Shawky [23] take this approach further by calculating the power losses of a MOSFET with respect to load, a procedure that could be applied to the present study to calculate losses in the active devices with respect to demand on the circuit. Whatever approach is taken, a circuit-specific model is generated, which relates demand on the power converter prototype to power losses in each device.

In order to generate the thermal characteristics required for the proposed temperature estimator, a prototype of the system must be built (stage 3 in Fig. 13). This prototype must be modifiable in such a way that the power devices and cooling can be controlled independently of the converter to allow thermal characterization. This might be achieved through the use of jumpers on the circuit board, or by delaying installation of the gate drive circuitry in the prototype. Additionally, electrical access is required to the terminals of the MOSFETs for power control. One option to achieve this connectivity would be to make electrical contact

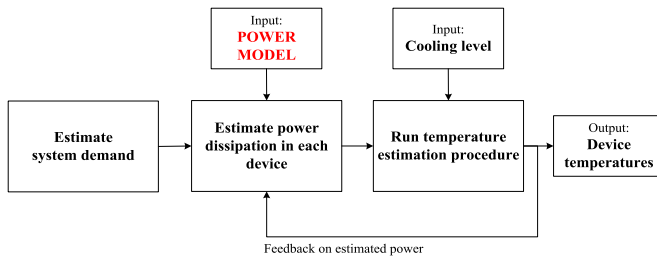


Fig. 14. Program flowchart for each iteration of temperature estimation on the operational system.

with test points on the prototype PCB using a modified bed-of-nails tester.

Following the procedure outlined in Section II-B, the thermal cross coupling of the system is characterized over an appropriate range of cooling levels. Because an active device is being used instead of the illustrative resistor used in Section II-B, an additional circuit, such as the one reported in [24], is required to control the power in each device. This circuit is required only for characterization and will, therefore, be external to the prototype. The temperature response of each device may be measured using a nonintrusive method such as with a single pixel thermal camera. Using the procedure described in Section II, the thermal characteristics are produced and this forms the thermal model of the system.

B. Implementation Within the Operational System

Having generated the models on a prototype system, the temperature estimation technique can be realized on the operational system. The flowchart of the required program is shown in Fig. 14. At each time step, the demand on the system is estimated. Depending on the application, this might be through measurement of the output current (e.g., on a simple power converter) or through measurement of the work done (e.g., water throughput when designing a drive for a pump). Using the power model developed as part of the design stage modeling, estimates of the power in each device are produced. Since losses in devices may be temperature dependent, the estimated temperature is used as an additional input to the model. The temperature estimation procedure described in Section IV and implemented in Section V is then employed to produce real-time temperature data, which is, hence, available for such purposes as cooling control, demand management or for any other purpose.

VII. DISCUSSION

A. Implementation Decisions

This paper has proposed a method of estimating the temperature of devices under variable cooling. The technique is applicable in several areas, notably for air- and liquid-cooled power electronics in drive systems. To implement the estimator in practice, two decisions must be made. First, which of the proposed methods will be used for the transition between filters must be chosen. It was shown in Section IV that for cases where the response is mainly due to auto coupling or where changes in cooling level are infrequent, the simpler steady-state assumption

method may suffice. However, where there is a delay between input and output (in the case of cross coupling, for example) and frequent changes in cooling level, a more involved method is required. The engineer must decide whether the additional computational requirements are justified by the increased accuracy. That is a complex application-specific question involving the budget for the particular system versus the desire for tight thermal control. We have seen that both methods can be implemented on a cheap 8-bit microprocessor running on a low-speed RC clock. There is, therefore, a strong preference for the scaled input method except in the case high sample rate systems consisting of many devices. Alternatively, thermal management of complex systems could be implementation on a digital signals processor to allow increased processing rates through accelerated mathematics.

The second decision is the level of quantization desired between cooling levels. An increased number of steps offers greater model acuity at the expense of processor requirements (memory and sample rate), which scale linearly with number of discrete cooling levels. Additionally, for very finely quantized models, transitions between filters become undesirably frequent. In this case, transitions may occur where there is no significant difference in actual cooling or merely as a result of noise in cooling level measurement. The degree of quantization is, therefore, also application specific.

B. Comparison to Related Work

Techniques for device temperature estimation and prediction on a system under variable cooling have received relatively little attention in the recent literature, with many authors preferring to use a thermal model that assumes consistent cooling [8], [9]. These methods lose accuracy under variable cooling. Alternatively, authors use methods that measure the temperatures under variable cooling directly, without an estimator (one such example is [25]). These methods, on the other hand, increase the per-unit system cost through the increased requirement for sensors and signal conditioning. Finally, alternative proposals that do produce temperature estimations under variable cooling often calculate the change in the parameters of thermal equivalent circuits and produce a SPICE model [26]. Although these estimation techniques generate results of similar accuracy to the technique proposed in this paper, the requirement to produce an accurate equivalent circuit model and identify which parameters are affected by cooling increases the design effort required for implementation.

VIII. CONCLUSION

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 PRBS-based characterization to produce a set of digital IIR filters that model the dynamics of the system under quantized 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 demonstrate accurate tracking of temperature with minimal error when compared with the experimental data.

REFERENCES

- [1] M. Gerber, J. A. Ferreira, N. Seliger, and I. W. Hofsjager, "Integral 3-D thermal, electrical and mechanical design of an automotive DC/DC converter," *IEEE Trans. Power Electron.*, vol. 20, no. 3, pp. 566–575, May 2005.
- [2] V. Smet, F. Forest, J.-J. Huselstein, F. Richardeau, Z. Khatir, S. Lefebvre, and M. Berkani, "Ageing and failure modes of IGBT modules in high-temperature power cycling," *IEEE Trans. Ind. Electron.*, vol. 58, no. 10, pp. 4931–4941, Oct. 2011.
- [3] H. Lu, C. Bailey, and C. Yin, "Design for reliability of power electronics modules," *Microelectron. Rel.*, vol. 49, no. 9–11, pp. 1250–1255, Sep.–Nov. 2009.
- [4] 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. 23rd Annu. IEEE Conf. Appl. Power Electron. Conf. Expo.*, Feb. 2008, pp. 192–198.
- [5] L. Meysenc, M. Jylhakallio, and P. Barbosa, "Power electronics cooling effectiveness versus thermal inertia," *IEEE Trans. Power Electron.*, vol. 20, no. 3, pp. 687,693, May 2005.
- [6] L. Fried, "Prediction of temperatures in forced-convection cooled electronic equipment," *IRE Trans. Compon. Parts*, vol. CR-5, no. 2, pp. 102–107, Jun. 1958.
- [7] C. Lundquist and V. P. Carey, "Microprocessor-based adaptive thermal control for an air-cooled computer CPU module," in *Proc. 17th Annu. IEEE Symp. Semicond. Therm. Meas. Manag.*, 2001, pp. 168–173.
- [8] M. Musallam and C. M. Johnson, "Real-time compact thermal models for health management of power electronics," *IEEE Trans. Power Electron.*, vol. 25, no. 6, pp. 1416–1425, Jun. 2010.
- [9] G. C. James, V. Pickert, and M. Cade, "A thermal model for a multichip device with changing cooling conditions," in *Proc. 4th IET Conf. Power Electron., Mach. Drives*, 2008, pp. 310–314.
- [10] G. Xiong, M. Lu, C.-L. Chen, B. P. Wang, and D. Kehl, "Numerical optimization of a power electronics cooling assembly," in *Proc. IEEE 16th Annu. Appl. Power Electron. Conf. Expo.*, 2001, vol. 2, pp. 1068–1073.
- [11] P. J. Rodgers, V. R. C. Evely, 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," *J. Electron. Packag.*, vol. 125, pp. 67–75, 2003.
- [12] P. J. Rodgers, V. R. C. Evely, and M. R. Davies, "An experimental assessment of numerical predictive accuracy for electronic component heat transfer in forced convection—part II: Results and discussion," *J. Electron. Packag.*, vol. 125, pp. 76–83, 2003.
- [13] J. N. Davidson, D. A. Stone, and M. P. Foster, "Real-time prediction of power electronic device temperatures using PRBS-generated frequency-domain thermal cross-coupling characteristics," *IEEE Trans. Power Electron.*, vol. 30, no. 6, pp. 2950–2961, Jun. 2015.
- [14] J. N. Davidson, D. A. Stone, M. P. Foster, and D. T. Gladwin, "Improved bandwidth and noise resilience in thermal impedance spectroscopy by mixing PRBS signals," *IEEE Trans. Power Electron.*, vol. 29, no. 9, pp. 4817–4828, Sep. 2014.
- [15] B. Miao, R. Zane, and D. Maksimovic, "System identification of power converters with digital control through cross-correlation methods," *IEEE Trans. Power Electron.*, vol. 20, no. 5, pp. 1093–1099, Sep. 2005.
- [16] MathWorks. (2013, Jul. 8). MATLAB 2013a documentation: Invfreq [Online]. Available: <http://www.mathworks.co.uk/help/signal/ref/invfreq.html>.
- [17] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1999.
- [18] R. Hanus, M. Kinnaert, and J.-L. Henrotte, "Conditioning technique, a general anti-windup and bumpless transfer method," *Automatica*, vol. 23, no. 6, pp. 729–739, Nov. 1987.
- [19] J. D. Bendtsen, J. Stoustrup, and K. Trangbaek, "Bumpless transfer between observer-based gain scheduled controller," *Int. J. Control*, vol. 78, no. 7, pp. 491–504, 2005.
- [20] E. W. Kamen, *Introduction to Signals and Systems*. New York, NY, USA: Macmillan, 1990, pp. 333–334.
- [21] European Community, "Directive 90/C81/01," EEC Journal Official No. C81, 30 Mar. 1990, p. 110.
- [22] G. Di Capua and N. Femia, "A versatile method for MOSFET commutation analysis in switching power converter design," *IEEE Trans. Power Electron.*, vol. 29, no. 3, pp. 920–935, Feb. 2014.
- [23] M. Orabi and A. Shawky, "Proposed switching losses model for integrated point-of-load synchronous buck converters," *IEEE Trans. Power Electron.*, vol. 30, no. 9, pp. 5136–5150, Sep. 2015.
- [24] J. N. Davidson, D. A. Stone, and M. P. Foster, "Arbitrary waveform power controller for thermal measurements in semiconductor devices," *Electron. Lett.*, vol. 48, no. 7, pp. 400–402, 2012.
- [25] C. E. Bash, C. D. Patel, and R. K. Sharma, "Dynamic thermal management of air cooled data centers," in *Proc. 10th Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst.*, 2006, pp. 8–452.
- [26] M. März and P. Nance, *Thermal Modeling of Power Electronic Systems*. Munich, Germany: Infineon Technologies AG, 2000.



Jonathan N. Davidson received the M.Eng. degree in electronic engineering and the Ph.D. degree in thermal modeling and management from the University of Sheffield, Sheffield, U.K., in 2010 and 2015, respectively.

He is currently a Research Associate at the University of Sheffield. His research interests include thermal modeling and management of power electronics, and design of piezoelectric power converters.

Dr. Davidson is a Member of the Institution of Engineering and Technology.



David A. Stone received the B.Eng. degree in electronic engineering from the University of Sheffield, Sheffield, U.K., in 1984 and the Ph.D. degree from Liverpool University, Liverpool, U.K., in 1989.

He then returned to the University of Sheffield as a Member of academic staff specializing in power electronics and machine drive systems. His current research interests include hybrid-electric vehicles, battery management, electromagnetic compatibility, and novel lamp ballasts for low-pressure fluorescent lamps.



Martin P. Foster received the B.Eng. degree in electronic and electrical engineering, the M.Sc. (Eng.) degree in control systems, and the Ph.D. degree for his thesis "Analysis and Design of High-order Resonant Power Converters" from the University of Sheffield, Sheffield, U.K., in 1998, 2000, and 2003, respectively.

In 2003, he became a Member of academic staff at Sheffield specializing in power electronic systems and was made a Senior Lecturer in 2010 and then a Reader in 2014. His current research interests include

the modeling and control of switching power converters, resonant power supplies, multilevel converters, battery management, piezoelectric transformers, power electronic packaging, and autonomous aerospace vehicles.



Daniel T. Gladwin received the M.Eng. (Hons.) degree in electronic engineering (computer architecture) and the Ph.D. degree in automated control structure design and optimization using evolutionary computing from the University of Sheffield, Sheffield, U.K., in 2004 and 2009 respectively.

In 2012, he was appointed as a Lecturer at Sheffield. His active research is into power system control, power electronics, energy storage and management, embedded systems, and intelligent systems.