ISSUE: May 2020

Inductorless Heater Switch-Mode Control Enhances Configurability Of Server Motherboards

by Viktor Vogman, Power Conversion Consulting, Olympia, Wash.

The increase in server complexity has resulted in a growing need for better and easier motherboard configurability. To meet this requirement, server motherboard (MB) developers recently unveiled a new technology [1, 2] that facilitates motherboard reconfiguration with a so-called rework grid array (RGA) interposer. Such an interposer is composed of a ball grid array (BGA) or an IC package support structure, which has electric heaters embedded in its foundation layer.

The heaters supply heat locally to reflow solder and thereby enable attachment or detachment of the IC or the processor package. Due to its simplicity, this RGA-based reconfiguration method provides significant cost advantages over the use of traditional electrical/mechanical sockets or expensive specialized equipment for de-soldering and re-soldering.[3]

The power required to activate the RGA heater can be supplied from the MB power supply, e.g. a 12-V bus. If control of the heater temperature, i.e. reflow temperature profile, can be provided with high precision, the local heat method would compete well with expensive and complex BGA rework systems. Minimizing the cost and size of the heater control circuitry would allow it to be placed on the interposer with the electric heaters, making the RGA interposer fully autonomous.

This article studies aspects of the electrical design of the RGA heater control in order to make this circuitry highly efficient, as well as cost- and size-optimized. This is done ultimately with the goal of enabling and simplifying implementation of the new MB reconfiguration technology.

The RGA Interposer Structure

A side view of an RGA power interposer residing between a motherboard and a BGA package is shown in Fig. 1. An embedded heater and isolated temperature sensor (optional) are shown with thick and thin yellow dashed lines, respectively.

Examples of the heater trace routing and Pb-free classification reflow profile according to IPC/JEDEC J-STD-020 are shown in Fig. 2 parts a and b, respectively. Different heater segments may have different temperatures, depending on cooling conditions. Equalizing those temperatures may require use of several independently controlled heater zones.

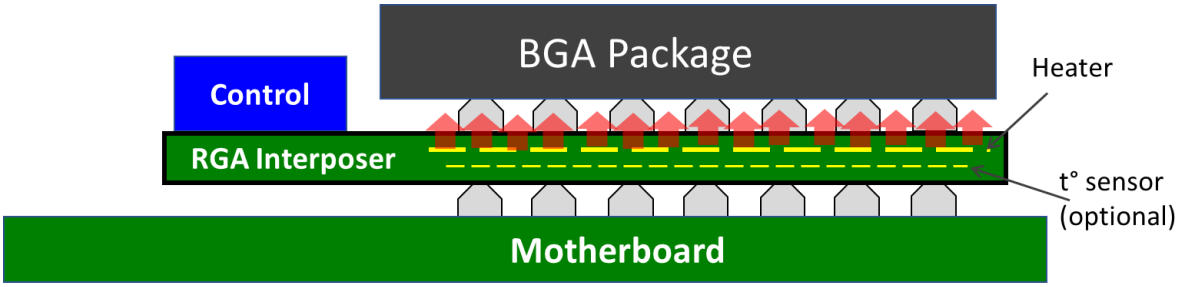

Fig. 1. A side view of an RGA power interposer residing between a motherboard and a BGA package.

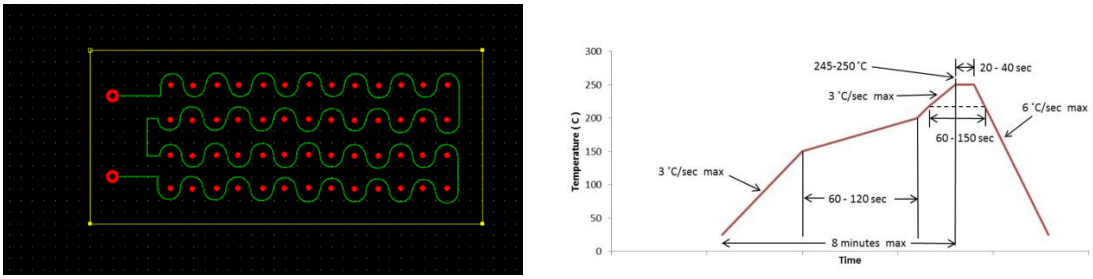

(a) *(b)*

Fig. 2. Examples of the heater trace routing (a) and Pb-free classification reflow profile according to IPC/JEDEC J-STD-020 (b).

Inductorless PWM Control

Control of the heater temperature is facilitated by control of the current through the heater trace. A straightforward solution that can accomplish this is the use of a switching voltage regulator (VR) capable of providing heater current control over a wide range.

The main drawback of this approach is that a VR requires the use of magnetics—power inductors, along with large numbers of semiconductor components and decoupling caps. These components impact the solution cost, which becomes especially noticeable in a multizone case since the power stage parts count is proportional to the number of controlled heater zones.

Heater temperature variation is a slow process due to thermal inertia, so the heater temperature may be considered constant within short, e.g. a few millisecond time interval. In other words, as far as temperature variation is concerned, the heater trace can act like a large-time-constant low-pass filter in a switching regulator, removing switching frequency harmonics from a pulse wave.

This feature presents an opportunity for the implementation of practically lossless and low-parts-count heater control, which can be provided by a series solid-state switch—a power MOSFET. In this power stage topology, the power delivered to the heater can be controlled by varying the switch-on duty cycle, without using additional active or passive components.

An illustrative block diagram of such a control arrangement is shown in Fig. 3. In this diagram a heater temperature control signal supplied by an MCU is compared with the heater temperature sensor signal.

Amplified by the error amplifier EA, the error signal fed to the PWM input controls switch Q1's on duty cycle, changing the power delivered to the heater and making the heater temperature follow the temperature set by the MCU signal "T set".

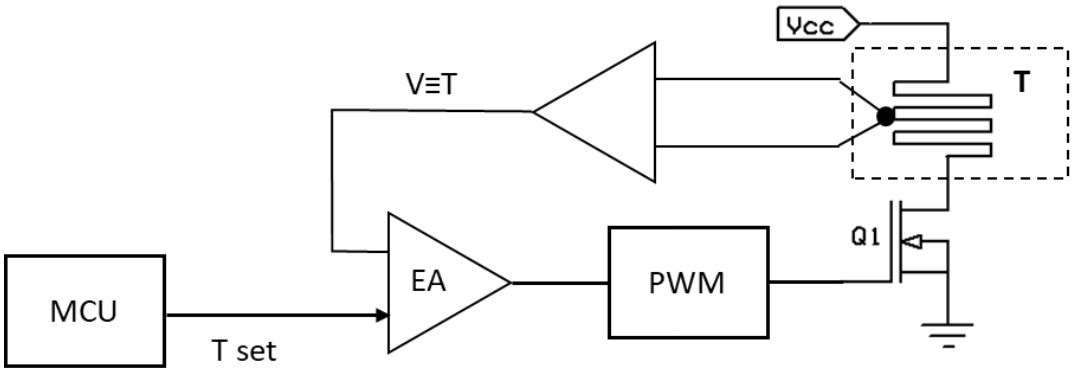

Fig. 3. Functional block-diagram of closed-loop control for a single-zone-heater. Generated by an MCU, the heater-temperature control signal is compared with the heater-temperature-sensor signal. The amplified error signal then controls the on duty cycle of switch Q1, causing the heater temperature to follow the temperature set by the MCU.

Heater Control Considerations

Let's determine the controller power level and operating conditions, i.e. the switch peak current and its active (switch-on) duty cycle, to obtain the required temperature.

In a pulse-wave case, heater resistance R_H can be determined as a ratio of supply voltage V_{cc} to the heater current magnitude expressed via RMS current level, *Irms*:

$$R_H = \frac{V_{CC}}{I_m} = V_{CC} \cdot \frac{\sqrt{D}}{Irms} \qquad (1)$$

where *D* is the duty cycle of the current pulses.

The expression for *Irms* is published in the IPC-2221 standard[4] section 6.2:

$$Irms = k \cdot \Delta T^{0.44} \cdot A^{0.725} \qquad (2)$$

where ΔT is the required temperature rise, *A* is the trace cross section in square mils and *k* is a constant such that $k = 0.048$ for outer layers and $k = 0.024$ for inner layers. Substituting it into equation (1) we find that:

$$R_H = V_{CC} \cdot \frac{\sqrt{D}}{Irms} = \frac{V_{CC} \cdot \sqrt{D}}{k \cdot \Delta T^{0.44} \cdot A^{0.725}} \qquad (3)$$

The trace geometry (thickness, width and length) can be engineered based on package mounting (solder) pad layout: conducting copper pad locations, the distance between consecutive balls on a BGA package (pitch size) and pad diameter. For standard BGA pads and pitch size[5] the heater trace width can vary between 3 and 7 mils. Considering standard PCB copper weight specs: 0.5 oz/ft^2 and 1.0 oz/ft^2, which correspond to copper thicknesses of 0.7 mils and 1.4 mils, respectively, we can define the heater trace cross-sectional area range as 2 square mils < *A* < 10 square mils.

Plots of the required heater RMS current as a function of trace cross-sectional area and the temperature rise obtained with expression (2) are given in Fig. 3a. These graphs show that for a given range of cross-sectional area, solder melting temperatures can be reached at relatively low RMS currents.

Heater resistance R_H at elevated temperatures can also be derived from Pouillet's law and expressed as a function of trace length *L*, its cross-sectional area *A*, and copper electrical resistivity ρ (0.017 ohm·mm^2/m):

$$R_H = \frac{\rho(1 + \alpha\Delta T) \cdot L}{0.000645 \cdot A} \qquad (4)$$

where α is the copper thermal coefficient of resistance (α = 0.004°C^{-1}), trace length L is given in meters and cross-sectional area A is given in square mils.

Per IPC-2221, setting k = 0.024 for inner layers and equating expressions (3) and (4) we find a relationship between the switch-on duty cycle and heater trace length for a given temperature rise ΔT, supply voltage, and trace cross-sectional area:

$$D(\Delta T, L, A) = \left[\frac{37.2L \cdot \Delta T^{0.44} \cdot \rho \cdot (1 + \alpha\Delta T)}{V_{CC} \cdot A^{0.275}}\right]^2 \qquad (5)$$

If $D \le 1$ the heater with selected trace length *L* can reach the required temperature at a given V_{CC} level. If $D > 1$ the heater trace needs to be split into several sections having smaller lengths. Under identical cooling conditions these sections can be electrically paralleled and controlled by a single switch or, if they are expected to have different cooling conditions, by individual switches.

The dependence of the duty cycle on the trace length and cross-sectional area is shown in Fig. 3b for V_{CC} = 12 V and ΔT = 230°C, which corresponds to the maximum solder melting point. The dependencies presented in these plots can be used as a reference for a heater PWM controller selection. Values of the D = 100% crossing point abscissas correspond to the absolute maximum trace length that can be used for a selected cross-sectional area. To have a control margin in real applications the trace length needs to be 10% to 20% below the absolute maximum. Such a control margin can compensate for cooling conditions or ambient temperature variations.

Equation 5 and Fig. 3b show that the switch-on duty cycle increases with trace length. The physical interpretation is that for a given supply voltage, as trace length and thereby resistance increases, maintaining the same RMS current and temperature rise can be facilitated simply by expanding the fraction of the period over which the switch is active. Thus, equation 5 defines feasibility boundaries of the inductorless PWM heater control method and allows a designer to form heater zones to get the required temperature rise if a single-zone heater appears to be unfeasible.

The defined RMS current, switch-on duty cycle and given supply voltage allow us to determine a switch peak current I_m and average power level P_H needed to provide required heater zone temperature rise ΔT_{max}. These characteristics for a pulse wave can be described by the following equations:

$$I_{m(\Delta T=\Delta T_{max})} = \frac{Irms}{\sqrt{D(L, A, \Delta T_{max})}} \qquad (6)$$

$$P_{H(\Delta T=\Delta T_{max})} = V_{CC} \cdot Irms\sqrt{D(L, A, \Delta T_{max})} \qquad (7)$$

where I_{RMS} and D are defined by equations (2) and (5), respectively.

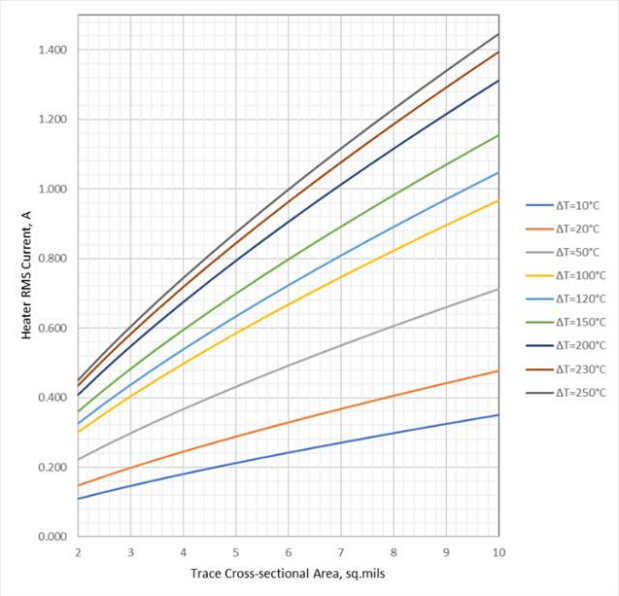

(a) Reguired heater RMS current vs. trace cross-sectional area at different temperature rises.

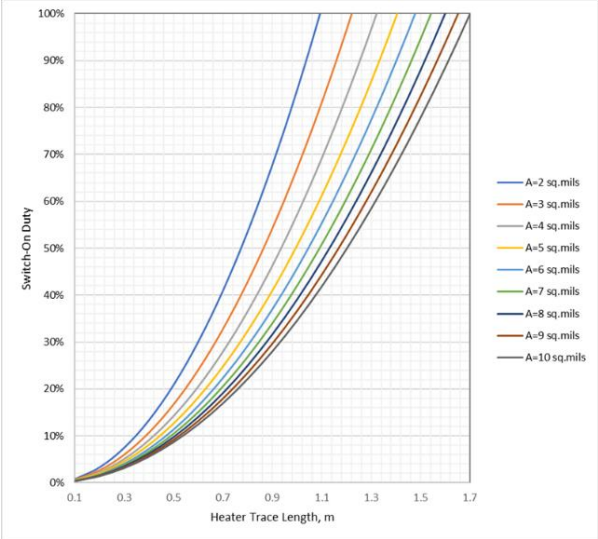

(b) Switch-on duty cycle vs. heater trace length at different trace cross-sectional areas (ΔT = 230°C).

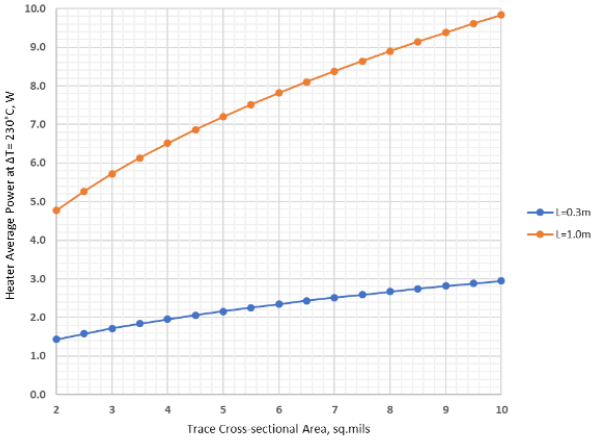

(c) Heater average power vs. trace cross-sectional area at solder melting point (ΔT = 230°C).

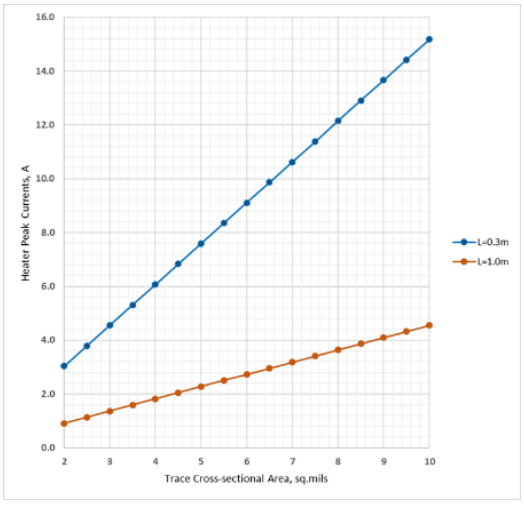

(d) Heater peak current vs. trace cross-sectional area at the beginning of the reflow profile execution (ΔT = 0°C).

Fig. 3. Heater and switch operating conditions (Vcc = 12 V) that can be used for determining control signal characteristics and MOSFET ratings selection.

At the beginning of the reflow profile execution the trace is typically at room temperature, which is why the expressions for maximum peak current and peak power must include a correction factor $(1 + \alpha\Delta T)$ that reflects

the trace resistance reduction. For selecting supply maximum current and peak power ratings, the following expressions should be used:

$$I_{m(\Delta T=0)} = \frac{Irms(1+\alpha\Delta T)}{\sqrt{D}}$$

$$P_{H,m(\Delta T=0)} = V_{CC} \cdot I_{m(\Delta T=0)}$$

Average power and peak current ratings for two typical zone lengths—0.3 m and 1 m—are plotted as a function of trace cross-sectional area in Fig. 3 parts c and d, respectively. These graphs present a useful model of heater operating mode and allow us to make some practical recommendations:

- The lower the cross-sectional area of the heater trace, the lower the average and peak power required to reach the melting point.
- The average power required to reach the melting point temperature in one heater zone does not exceed 10 W and can be provided with one 20-A-rated MOSFET.

Similar graphs for different supply voltages can be generated using equations (1) through (5). If the available supply voltages are not high enough to generate a sufficient RMS current in a single zone heater, then separate, shorter heater zone traces can be used.

Even with a single-zone heater the use of closed-loop operation is supposed to compensate for ambient temperature variations, and PCB thickness and package heatsinking effects. However, some heater areas' temperatures can still differ from projected values due to variation in cooling conditions for different heater area locations.

Consequently, the heater trace zone configuration and trace width selection may need to be fine-tuned at the pilot build stage of RGA interposer development. This fine-tuning will provide uniform temperatures across the mounting pad and the highest success rate for pin attachment and detachment. The pilot build will also help determine whether the duty cycle margins required for cooling condition variations are sufficient for every heater zone.

Heater Temperature Sensing Options

To execute a required reflow profile, each heater zone temperature needs to be accurately monitored. There are two relatively simple and accurate methods of sensing the heater temperature:

1. An embedded temperature sensor (a thin trace) located in the closest proximity to the corresponding heater zone. Its location is illustrated in Fig. 1 with a thin yellow dashed line.

2. A precision current-sense resistor connected in series with the MOSFET source terminal.

If the RGA control signal, forming the required reflow profile, is provided by an MCU, then both of these methods constitute low-cost solutions.

Embedded Temperature Sensor

An embedded temperature-sensor trace acts as a thermistor having the same thermal coefficient as a heater. Passing a small current I_S through such a trace with a precision current-source IC will generate a voltage drop V_T across it that is proportional to the current level and resistance of the trace at a given temperature $R_S(T)$:

$$V_T = I_s \cdot R_s(T) = I_s \cdot R_{s0} \cdot (1+\alpha\Delta T)$$

where R_{s0} is the sensor trace resistance at room temperature. From this expression we can determine the heater temperature as a function of V_T:

$$T(V_T) = \Delta T + 20^{\circ} = \frac{\left(\frac{V_T}{I_s \cdot R_{s0}} - 1\right)}{\propto} + 20^{\circ}$$

The advantage of the embedded sensor approach is that closed-loop control can be provided by a standard signal op amp comparing an analog-temperature-setting voltage signal with a sensor voltage V_T without any additional signal processing. When a detached heater temperature sensor is used, the temperature monitoring delay, which is related to heat propagation from heater to sensor, needs to be characterized and taken into account.

Series Resistive Sensor

An obvious advantage of the series-resistive-sensor method is that it has no inertia: the sensor monitors the heater itself, which is why it does not have a monitoring time delay. Also, the current sensing element already exists since it is an integral part of practically any short-circuit protected switching power MOSFET stage.

With a series resistive sensor, the actual heater resistance can be derived from the heater current magnitude I_m. This magnitude can be detected with a standard precision op amp rectifier or MCU. The heater temperature as a function of the detector output voltage V_m can be expressed as follows:

$$T(V_m) = \frac{\left(\frac{r_s \cdot V_{CC}}{R_{H0} \cdot V_m} - 1\right)}{\propto} + 20^{\circ}$$

where R_{H0} is the heater trace resistance at room temperature. Since the sensed signal (V_m) in the series-resistive-sensor case is inversely proportional to heater temperature, the error op amp's (see Fig. 2) inverting and non-inverting inputs are supposed to be swapped as compared to the embedded sensor case.

Minimizing Power Losses In the Control Switch

If the heater control circuitry is placed near the heater, as depicted in Fig. 1, it is critical to minimize the power dissipated in the power MOSFET. That's because the control circuitry needs to operate in elevated temperature conditions, which are caused by the heat radiated from the heater. To minimize the switch conduction loss the power MOSFET needs to have the lowest possible on-resistance.

Because the MOSFET load is inductive it can generate large voltage spikes when the MOSFET is switched off in an abrupt fashion. Such spikes would require using MOSFETs with a higher voltage rating, which in turn leads to them having a larger on-resistance. This means that the requirement for minimizing conduction losses can be met if the maximum MOSFET drain-source voltage rating can be kept as close to the V_{CC} level as possible.

There are three basic ways of suppressing the inductive voltage spikes in hard-switching topologies:

1. Clamping the spike with a diode and low-ESR decoupling caps positioned right next to the power MOSFET.

2. Using snubber circuits absorbing the energy stored in the inductance.

3. Slowing down the falling edge rate so that the EMF $L_H di/dt$ is limited to a projected low level (Fig. 4).

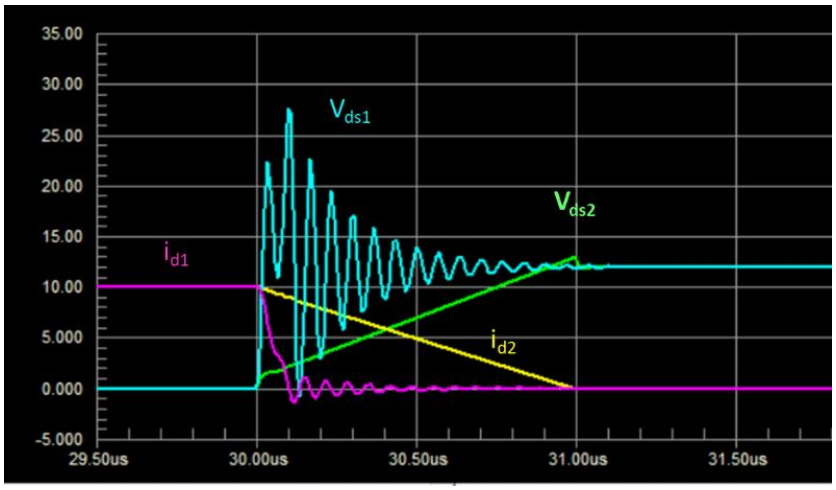

Fig. 4. Simulation of the power MOSFET turn-off with the original (V_{ds1}, i_{d1}) and slowed down gate-signal falling edge (V_{ds2}, i_{d2}). In the inductorless power stage topology, the voltage spike produced by the heater control switch can be drastically reduced with a shallower falling edge on the gate signal without any impact on the power stage dimensions.

The first two methods lead to the increased size and cost of the control circuitry. The third method may cause a noticeable increase in switching losses. However, in the inductorless switch-mode topology this obstacle can be overcome by switching frequency reduction without any increase in the power stage dimensions.

Limiting the EMF generated by the heater series inductance L_H at a level of ΔV requires di/dt = $\Delta V/L_H$. Assuming a linear current ramp and taking into account that $I_m = I_{RMS}/\sqrt{D}$, we can determine the required duration of the falling edge:

$$t_f = \frac{I_m}{di/dt} = \frac{Irms \cdot L_H}{\Delta V \cdot \sqrt{D}}$$

For example, for L_H =1 μH, D = 0.1, I = 0.5 A, ΔV = 1 V we find that t_f = 1.58 μs.

Let's evaluate the upper limit of switching frequency which would guarantee that the power loss impact will be negligible. For linear voltage and current ramps, the average power loss associated with the MOSFET switching can be determined by the equation:

$$P_{SW.avg} = f_{SW} \cdot \int_0^{t_f} I_m \left(1 - \frac{t}{t_f}\right) \cdot V_{CC} \cdot \frac{t}{t_f} \, \mathrm{dt}$$

where f_{sw} is the switching frequency. Assuming $P_{sw.avg}$ = 0.001 $V_{cc} \cdot I_m$, which corresponds to less than 0.1% efficiency impact, and solving this equation for f_{sw} we find:

$$f_{SW.max} = \frac{0.006}{t_f}$$

For the example above: $f_{sw.max}$ = 6.32 kHz.

Experimental Results

An experiment was carried out with the following objectives:

- Evaluate the validity of equations (5) through (7) and verify that for given heater parameters—length (L) and cross-sectional area (A)—actual heater control operating conditions—duty cycle (D) and average power level (P_H)—match the projected values and result in the required temperature rise.
- Determine areas of considerable (>10°C) temperature deviations from the projected level, which may require using variable-width heater traces or creating separate heater zones.
- A proof of concept that the interposer can be realized as a standalone device—i.e. a standard solder alloy used on the solder pads can melt while the control circuitry, located on the same support structure, can operate within its normal operating temperature range.

The experiment was conducted with a single-zone-heater RGA interposer. The heater control circuitry supply voltage was set at 12 Vdc. The transistor power switch was implemented with the BSC010NE2LS (25-V, 1-mΩ) power MOSFET providing low power dissipation.

Hysteretic control using a series-resistive current sensor was keeping the heater current magnitude between preset upper and lower boundaries, corresponding to the heater temperature levels set manually by an external potentiometer. Heater trace parameters were copper thickness = ½ oz/ft^2 (0.7 mils or 0.0178 mm), width = 5 mils and length L = 93.3 cm.

Fig. 5 shows the experimental interposer board after six hours of continuous operation and its IR camera images in active mode. A comparison between predicted and measured critical operating point data at ΔT = 0°C and ΔT = 230°C is given in the table below.

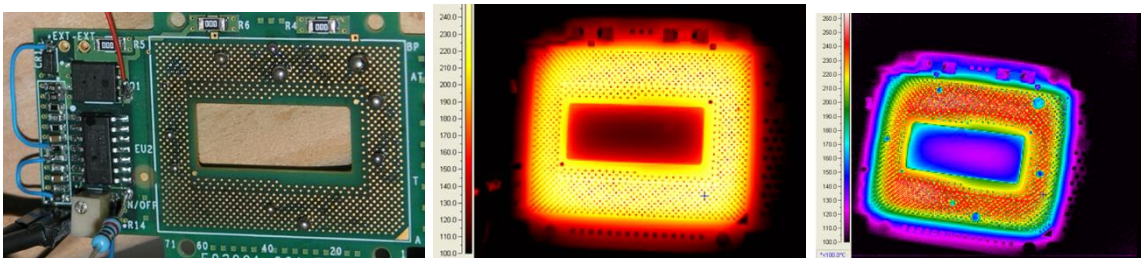

(a) *(b)* *(c)*

Fig. 5. Experimental interposer after 6-hour thermal run (a) and its IR images in active mode at ΔT = 220°C (b) and 230°C (c).

Table. Measured versus calculated data for the experimental interposer.

	Calculated	Measured
I_m (ΔT = 0°C), A	1.71	1.73
I_m (ΔT = 230°C), A	0.89	0.93
Irms (ΔT = 230°C), A	0.651	0.69
Duty (ΔT = 230°C)	0.535	0.55
P_H (ΔT= 230°C), W	5.71	6.14

The experiment shows that while average heater temperature measurements agree well with the projected

values there are areas of considerable (about 10°C) temperature deviations from the anticipated level. Equalizing temperatures on the edges and at the center of the interposer may require either creation of a separate heater zone or using a heater with variable trace width.

Let's evaluate the prospect of using a variable-width heater trace. Based on equation 2 for the same RMS current flowing through a two-width trace the following equation is valid:

$$Irms = k \cdot \Delta T_1^{0.44} A_1^{0.725} = k \cdot \Delta T_2^{0.44} A_2^{0.725}$$

where ΔT_1, A_1 and ΔT_2, A_2 are temperature rises and cross-sectional areas of the trace sections with widths W_1 and W_2, respectively. From this equation we find:

$$\frac{\Delta T_1}{\Delta T_2} = \left(\frac{W_2}{W_1}\right)^{1.65}$$

This expression demonstrates that for a given temperature deviation, which is less than 5% of the melting temperature, the trace width ratio needs to be less than 1.03, which is impractical for thin (several mils width) traces. This means that when temperature deviations are considerable, but much smaller than the melting temperature, splitting the heater into two zones is recommended for better consistency of the heater temperature across the IC package footprint.

Measurement results also reveal that the required surface temperature can be achieved with comparatively low power. Components in the heater control circuitry remain at safe temperatures that do not exceed 110°C.

The experiment used a SN/AG3/CU0.5 solder alloy with a melting point of 217°C to 220°C.

Conclusions

An IC package support structure with embedded heaters represents a novel assembly option that facilitates the removal and reattachment of the IC package without the need for expensive sockets and factory machinery.

To power the heater, an inductorless topology offers a means to reduce the size of the heater controller. It also enables the use of a low voltage-rated, low $R_{DS(ON)}$ MOSFET and thereby provides extra-low power dissipation in the power stage. PWM control of the inductorless power stage is capable of providing successful reflow: it can produce the desired temperature profile within an IC package mounting pad at comparatively low power levels that are consistent with projections.

When temperature deviations of some heater segments are considerable but much smaller than the solder melting temperature, better temperature consistency can be achieved across the heater segments by splitting the heater into several zones.

The recommended design guidelines provided in this article can be used for developing low-cost heater control circuitry with minimal power dissipation such that this circuitry is capable of operating in close proximity to the heater and can be an integral part of the interposer.

Future Work

One possible follow-up to this work would be a study of how to select an appropriate solder paste for the rework grid array interposer and how to optimize the reflow temperature profile for maximizing the yield of high-quality solder joints. Another possible study could explore the package impact on temperature variations across the mounting pad. Finally, an implementation of the required reflow temperature profile using automated MCU closed-loop control would be valuable for easing implementation of the RGA interposer method.

References

1. United States Patent Application Publication US 20170178994A1, Jun. 22, 2017
2. "Rework grid array interposer with direct power" by Russell S. Aoki et al., US Patent 10,211,120, February 19, 2019.
3. TF2800 BGA/SMD Rework System for Extra Large Boards.
4. IPC-2221A Generic Standard on Printed Board Design. May 2003.
5. Ball Grid Array (BGA) Packaging.

About The Author

Viktor Vogman currently works at Power Conversion Consulting as an analog design engineer, specializing in the design of various power test tools for ac and dc power delivery applications. Prior to this, he spent over 20 years at Intel, focused on hardware engineering and power delivery architectures. Viktor obtained an MS degree in Radio Communication, Television and Multimedia Technology and a PhD in Power Electronics from the Saint Petersburg University of Telecommunications, Russia. Vogman holds over 50 U.S. and foreign patents and has authored over 20 articles on various aspects of power delivery and analog design.

For articles relating to powering servers, see How2Power's Design Guide, locate the "Application" category and select "Data Centers".