Package Thermal Challenges Due to Changing Mobile System Form Factors

Cameron Nelson, Jesse Galloway
Amkor Technology
2045 East Innovation Circle
Tempe, Arizona
Cameron.Nelson@Amkor.com

Abstract

Mobile platforms have driven semiconductor package form factors from several times the die area to sizes approaching die size. To make matters even more challenging, the package thickness, traditionally on the order of 1-3 mm, has been reduced to thicknesses less than 0.40 mm. In mobile applications, it is size that drives package design. There has been a price paid for the package shrinkage in terms of its thermal performance. Copper, previously used for power and ground planes, and multiple layers of traces, has been replaced with extremely fine traces built in fewer layers. The body size previously used to promote the spreading of heat from the die is now left with two options for heat flow: either out the top (immediately above the die) or through the bottom (below the die) of the package. Improvements to the chip architecture in terms of power efficiency is one of the few remaining options for thermal enhancement at the package-level. Further thermal enhancement should focus on the system level, as this is where the greatest opportunities exist. While many papers have focused on the thermal challenges associated with the system-level, few have translated these constraints to challenges at the package level. This study investigates the historical evolution of mobile platforms and their impact on packaging thermal challenges. Metrics for evaluating the optimization of packages for the mobile space will also be discussed.

Nomenclature

CABGA	ChipArray® Ball Grid Array
FCBGA	Flip Chip Ball Grid Array
fcCSP	Flip Chip-Chip Scale Package
FCLBGA	Flip Chip Lidded Ball Grid Array
LQFP	Low Profile Quad Flat Package
MLF®	Micro Lead Frame
SiP	System in Package
SWIFT®	Silicon Wafer Integrated Fan-out Technology
TEPBGA	Thermally Enhanced Plastic Ball Grid Array
WLCSP	Wafer Level Chip Scale Package
WLFO	Wafer Level Fan Out
Θ_{JA}	Junction-to-ambient Thermal Resistance
Θ_{JB}	Junction-to-board Thermal Resistance
Θ_{JC}	Junction-to-case Thermal Resistance

1 Introduction

Since the introduction of mobile platforms, one of their primary drivers has been form factor. Clearly, the portability of a device is directly related to its size and weight. Therefore, the design of all components of the mobile devices has been influenced by the need to reduce these two parameters. Early devices were mostly limited to very few functions. Cell phones could only make and receive calls, and pagers were an early form of text messaging. With limited functionality, and therefore computational demand, thermal concern was minimal, if present at all. As technology progressed, with improved battery life and energy density, more functionality was inevitably added to devices. The increased demand for computational power ultimately led to thermal challenges emerging in the mobile platforms. Packaging technologies that were once used to support this magnitude of computational power could simply not comply with the confined form factor desired in the emerging portable devices. In this paper we will discuss packaging thermal challenges as a result of the constraints presented by modern mobile platforms.

While the goal of this paper is not to focus on thermal management challenges at the system level, some issues will be discussed to expose the relationship and constraints to package level thermal challenges. Arguably one of the most critical thermal challenges at the system level of mobile platforms is regarding the management of external surface temperature. Since this category of devices is intended to be held by hand, it is important to keep the surface temperature below a certain threshold to maintain user comfort. Chiriac [3] and Wagner [6] present similar methods for characterizing the thermal efficiency of mobile platforms in terms of surface temperature distribution. Chiriac [3] defines a dimensionless parameter, the Coefficient of Thermal Spreading (CTS) that can be used as a figure of merit for comparing the effectiveness of a thermal design. The parameter is a function of the average surface temperature relative to the maximum surface temperature. A thermally ideal design would have no external hotspots, and the surface temperature would be isothermal. Figure 1 shows thermal imaging of a modern smartphone undergoing a benchmark test to load its CPU. On the backside of the device a significant hotspot can be seen above the location of the applications processor.

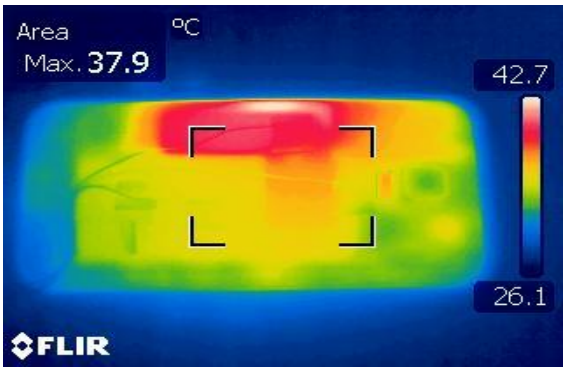

Figure 1: Surface temperature of a high-end smartphone. The high-power applications processor produces a significant hotspot on the external surface.

Assuming an isothermal surface temperature under natural convection conditions, the maximum sustained power a mobile device can achieve is directly related to its surface area and surface emissivity. Figure 2 shows theoretical maximum power dissipation for various modern mobile devices based on their external dimensions. These values will significantly overestimate the maximum sustained power of the actual device since the surface temperature will certainly not be isothermal.

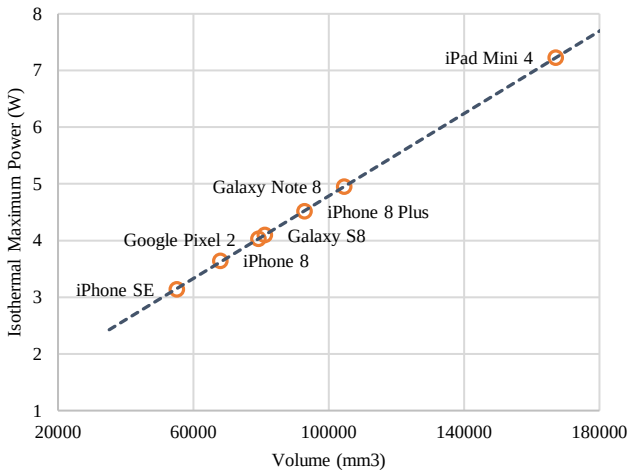

Figure 2: Theoretical steady-state thermal capacity versus total body volume for ideal isothermal surface temperature.

At the package level, heat spreading through conduction is the principal mechanism for effective thermal management. With no constraints on size, an optimal package design for a fixed heat source (silicon die) size would have a significantly larger footprint than its source and be constructed to maximize high conductivity material content. Conversely, the least optimal package design would be one where the heat source is equal to the footprint. This situation would clearly offer no heat spreading advantage. Of similar importance is the consideration of the package's through-plane resistance: the direction normal to the spreading direction, through the top or bottom of the package. This direction is where the heat would flow into a printed circuit board (PCB) or heat sink. In a laminate ball grid array (BGA) style package, the through-plane resistance to the bottom side of the package is a function of the buildup layer thicknesses, copper density, and via density/geometry, as well as the spreading resistance itself.

2 Mobile Device Design and Form Factor

It wasn't until the late 1990's that mobile phones became affordable and popular among the general consumer. This was the point in time when the cell phone became truly pocket-sized. This pocket-sized handheld form factor has trended all the way to today's modern smartphones. A major challenge in the design of a mobile phone is the optimization of this mechanical stack up [1]. Early pocket-sized mobile phones ranged in thickness from 20 to 30 mm while the most recent smartphones are less than 10 mm (See Figure 3). The primary drivers in this regard are the display size and battery size. As display technology progressed and cost was reduced, large color displays became a standard offering. With larger displays covering more of the body size, the phone's internal circuits had to be relocated, moving components behind the display itself. This put further strain on its total thickness.

Early mobile phones used external, removable batteries. This configuration did not restrict the internal design of the electronic components. In the latest smartphones, the batteries are embedded within these devices and occupy a large percentage of the internal volume. This leaves very restricted space for all other electronic components. The logic board of a modern smartphone, will typically account for only about 25% of the internal area of the device. The battery will comprise about 40% of internal area, and various components and sensors occupy the remaining 35% (See Figure 4). Although battery technology has continuously improved since the first mobile phones, the increasing demand for power means the battery dimensions are still a primary factor driving the device's form factor.

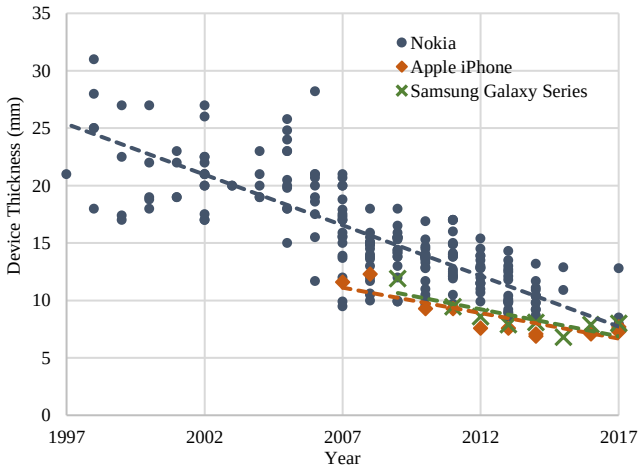

Figure 3: Historical progression of mobile phone thickness. [2]

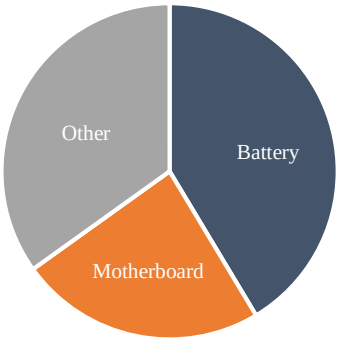

Figure 4: Area composition of motherboard (logic board), battery, and other components in a modern smartphone.

As the number of features and functionality of mobile phones expanded, modern high-end smartphones became packed with a high number of packages and integrated circuits

(ICs). (See Figure 5.) The high number of packages claiming area on the PCB adds further constraint to the individual package footprint. Keeping the total package cost to a minimum is critical for the device's profitability for the manufacturer and affordability to the consumer. The growth in low-cost packaging solutions has allowed manufactures to implement technologies that would have been impossible years ago. Wafer level package technologies, such as the Wafer Level Chip Scale Package (WLCSP) are highly favored for their low-cost, small footprint, and high I/O count. While the total number of packages in smartphones has steadily increased, so too has the percentage of packages comprised of WLCSPs due to the aforementioned advantages, see Figure 5.

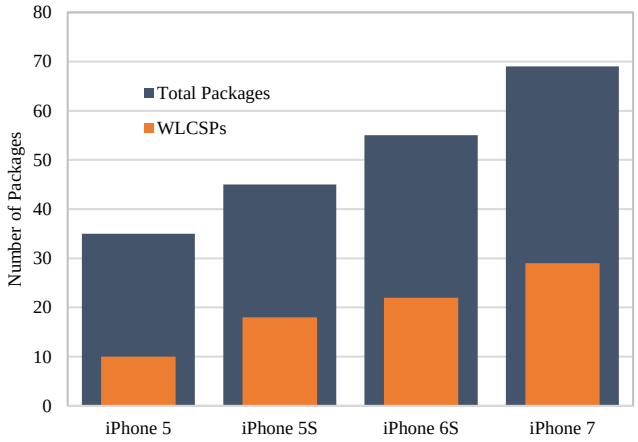

Figure 5: Total number of packages and Wafer Level Chip Scale Packages (WLCSPs) (shown inside the total bar) in recent Apple iPhones.

Figure 6 shows a cross section of a modern high-end smartphone through the application processor. The device has an overall thickness of less than 10 mm. The mechanical stack-up at the point of cross section consists of the aluminum body with a thickness of ~1 mm, ~1 - 2 mm thickness for the logic board, ~3 - 4 mm thickness for the display and touch screen, and only about 1 mm for the applications processor. The rest of the total thickness is attributed to mechanical tolerance and air gaps. It is clearly seen that there is little air voids within the stack up, eliminating any possibility for active cooling solutions (air movers). It is also evident that there is a severe thickness constraint on the package.

It is common for the logic board to have a long narrow form with a high aspect ratio. At its widest, the board is typically almost equal to the width of the applications processor, which is often the largest package on the board. It is also characteristic for these boards to have high layer counts with high copper densities, see Figure 7 and 8. The width severely bottlenecks the heat dissipation from the applications processor because the heat can only flow in two lateral directions. However, high density vias below the package (see Figure 7) in the board can aid in conducting heat to the many copper plane layers.

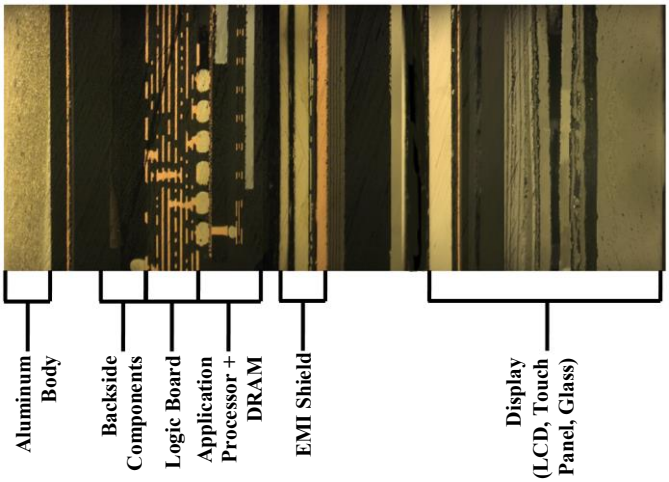

Figure 6: Cross section of a recent smartphone.

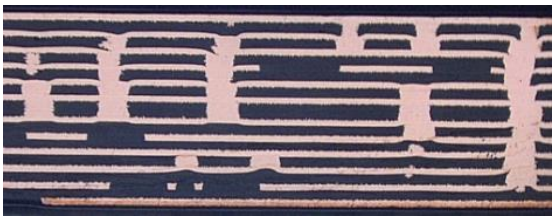

Figure 7: Logic board cross section of a recent smartphone.

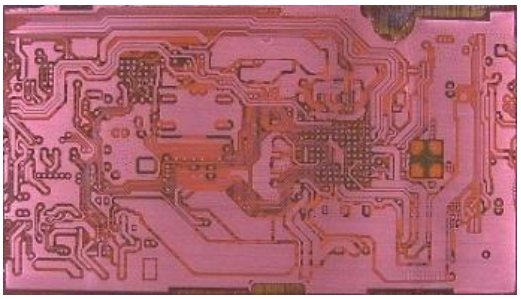

Figure 8: Planar section of logic board of a recent smartphone showing high copper content.

3.0 Packaging Thermal Challenges – Physical Geometry and Heat Flow Paths

In nearly all applications, primary drivers of package design are cost and size. These requirements are of utmost importance in mobile platforms. As discussed in the previous sections, with a dwindling allowance of internal space dedicated to packages, size constraints present significant pressure on the package's physical design. With mobile applications as a driving force, packaging technology has been steadily progressing to chip scale, where the body size or footprint is nearly equal to the area of the contained silicon chip. See Figure 9.

While the footprint is driven smaller, at the same time, input/output (I/O) requirements are increasing. With the core purpose of a package being to redistribute power and signals from a silicon chip, the I/O count is another key element of a package's design. Higher I/O densities inevitably require finer traces and thinner layers. Higher operating frequencies and

signal integrity requirements have also shaped the advancement of package substrate technologies.

Traditionally, laminate buildup technologies were the primary substrate technology for high I/O count wirebond and flip chip packages. However, this technology has reached limits in terms of minimum feature size due to manufacturing limitations. In response to this bottleneck, wafer level build up technologies have been developed to achieve significantly smaller feature sizes allowing much higher density design possibilities. As discussed in the previous section, the WLCSP, an example of a wafer build up technology, has tremendously grown in popularity in mobile platforms.

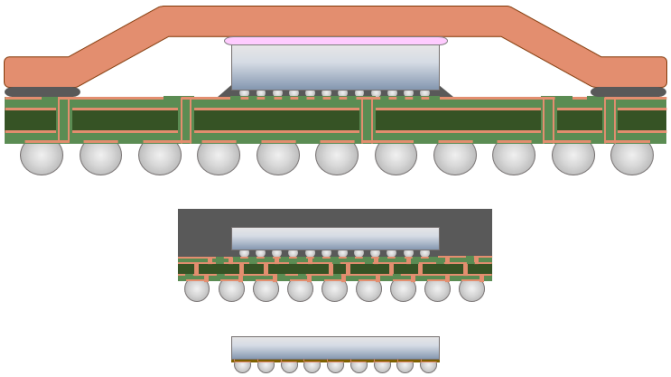

Figure 9: Illustrations of flip chip package technology evolution to chip scale. Top: flip chip lidded ball grid array (FCLBGA), middle: Flip Chip chip scale package (fcCSP), bottom: WLCSP.

In chip scale packages, with the die equal to (or nearly equal to) the body size, there exists limited internal heat spreading. Only two primary options for conduction heat flow direction are available; through the package top, or through the package bottom. Conversely, a package with high heat spreading capability would have a die size that is significantly smaller than the body size. This would provide the largest possible heat spreading advantage. Packages in this category can have thick copper planes within the substrate, lids, and copper die pads. With chip scale packaging, by nature, it's not possible to achieve these spreading advantages.

Many traditional laminate BGA packages can suffer from high through-plane thermal resistance. Typically in the laminate substrate, with the dielectric buildup material having a low thermal conductivity, a major factor in the through-plane resistance is the presence and quantity of vias. In package technologies approaching chip scale size, while the footprint decreases so does the overall thickness. This, in turn, results in thinner layers of dielectric buildup in the laminate and ultimately a lower through-plane resistance. At the extreme of the chip scale packaging is the Wafer Level Chip Scale Package (WLCSP). Since the heat generating component (the silicon die) makes up what is the package's body, there is no possibility of heat spreading at the package level. Also, since these packages have typically only a single, thin redistribution layer, they have a very low through-plane resistance. Other wafer build up technologies, such as Amkor's Silicon Wafer Integrated Fan-out Technology (SWIFT®), contain thin, high-density copper buildup layers and micro-vias which provide a very low through-plane resistance.

The heat flow path exiting the package depends on the package's actual application configuration. When a package is mounted to a PCB with no heat sinking or interactions on the package topside, the majority of heat will typically dissipate into the PCB. The heat conduction represents a much more effective heat flow path compared to the convection and radiation present on the package's topside. As discussed in the previous section, a typical logic board, at its widest, is typically nearly equal to the width of the applications processor package itself, restricting the heat flow to only two lateral directions.

Conversely, when a package is interfaced with a heat sinking mechanism to its side opposite the PCB, heat will also dissipate in this direction. Depending on the efficiency of the heat sinking mechanism, a majority of heat will also be drawn from the package in this direction. A heat sinking mechanism can range from finned heat sinks to the device chassis or any high thermal conductivity component that can interface the package. The system-level thermal management solution should drive package design in terms of the most effective direction for heat dissipation from the package.

3.1 Packaging Thermal Challenges – Thermal Capacitance

The majority of processors in modern smartphones are unable to achieve maximum power for a sustained period of time without exceeding a surface temperature limit. Fortunately, most are typically not required to function at full power continuously anyway. Instead, they use computational bursts over multiple cores for short periods when demand is needed. How efficiently the package can manage the surge of heat during these cycles is critical to the maximum duty cycle and temperatures that can be achieved. Figure 10 demonstrates the transient temperature response for various package types considering they each have the same die size and power dissipation. Other design parameters, such as their body size and construction, are typical for the specific package.

A package with a low thermal mass, such as the WLCSP, can be seen to have more rapid temperature rise compared to larger packages. The FCLBGA with its large copper lid, has the highest thermal capacitance and therefore the die temperature responds the slowest. While the steady-state temperatures of two different packages can be relatively close, their transient response curves can vary considerably resulting in significantly different duty cycle capabilities. This is a function of the geometry and material properties through which the heat propagates over time. The rate of change of temperature is dependent on the specific heat and density of the material that the heat is propagating through at that moment in time. Therefore it is expected to have the temperature rate of change vary over time.

The thermal time constant is a metric used to evaluate the transient performance of a package. See Figure 11. The thermal time constant is equal to the time required for the package to reach 1-1/e, about 63.2%, of the final asymptotic temperature. The thermal time constant is strongly dependent on the boundary conditions, not just the package itself. These are including, but not limited to, the motherboard design, enclosure, and ambient conditions. Therefore, when comparing

values, it is key to only consider packages under the exact same conditions. Figure 11 shows that chip-scale packages clearly have a disadvantage in their thermal capacitance due to their low thermal mass. A high thermal time constant allows greater duty cycle "on" times. For a given power requirement, if the time constant of the package is high enough, the device may be able to theoretically achieve a 100% duty cycle without exceeding temperature limits.

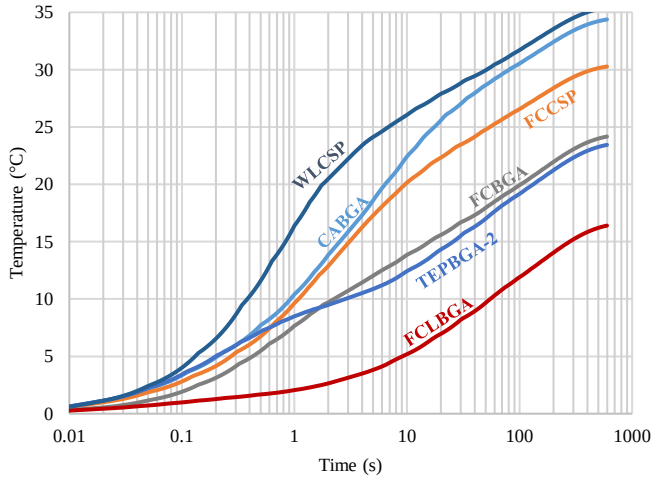

Figure 10: Transient temperature response for various package types. Data uses fixed silicon size and power dissipation.

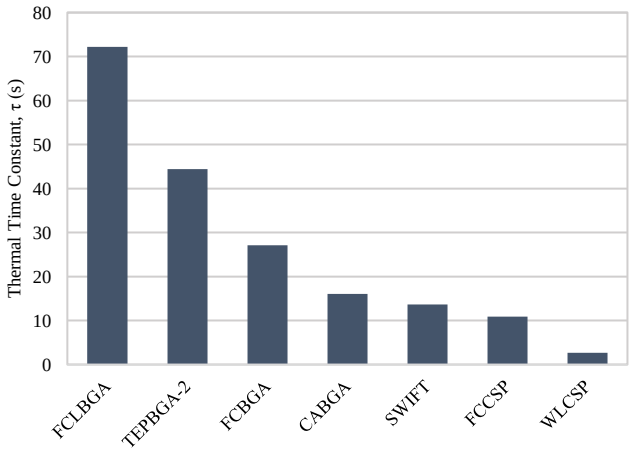

Figure 11: Thermal time constant for various package types.

3.2 Packaging Thermal Challenges – Silicon Thickness

Trends in package thinning inevitably lead to the thinning of the silicon chip itself. This presents a range of thermal challenges at the package level. Silicon, having a relatively high thermal conductivity, is quite effective at spreading heat due to hotspots in the design of the chip's power map. Multi-core as well as system on chip (SoC) architecture can present significant variations in the power density over the area of the silicon. The issue of hotspots is more prevalent on thin silicon because it is not able to effectively spread heat away. Advanced package technologies use silicon at thicknesses from 100 µm down to 50 µm. Without the spreading advantage of the thick silicon, mitigating hotspots is much more challenging. Figure 12 demonstrates the effect of silicon thinning on maximum power that can be achieved to maintain a given temperature limit. A FCBGA package was simulated using a finite element analysis (FEA) software with a single hotspot located near the corner of the die to represent a single core running in a processor. Considering a 775-µm thick silicon die, when thinned to 50 µm, the maximum power that can be dissipated from the hotspot (core) was reduced by almost 65% percent.

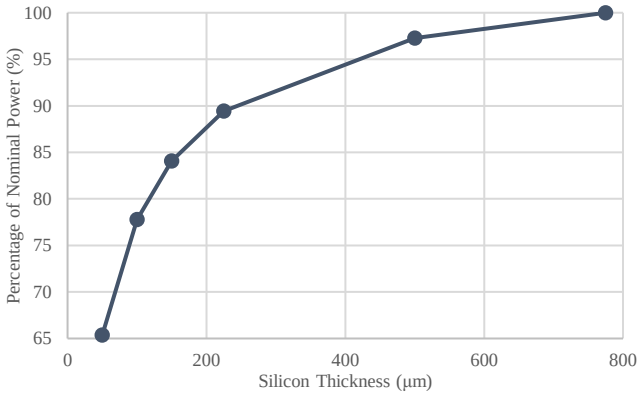

Figure 12: Impact of silicon thickness on single hotspot (core) power for a FCBGA package.

The silicon thickness also affects the transient performance of the package. The thicker the silicon, the higher its thermal capacitance to absorb small timescale transient pulses of power. Figure 13 shows the impact die thickness has on the package transient temperature response of a sample package. In the early time region, temperature rises at a much faster rate due to the lower thermal capacitance of the thinner silicon. At later times, once the heat has propagated beyond the silicon, the temperature differences between the curves can be seen to stabilize. At the extremes of this data, 775 µm versus 50 µm, the thin silicon reaches 20°C about 50% sooner than the thick silicon.

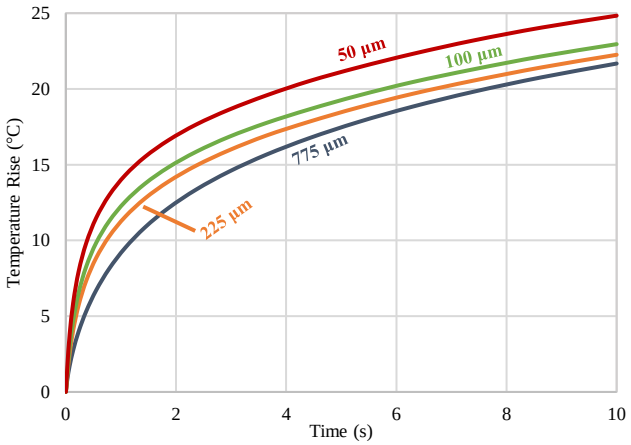

Figure 13: Silicon thickness impact on transient temperature response.

4 Package Evaluation

In addition to the transient properties discussed in the previous sections, when selecting a package for a mobile application, there are other characteristics to consider. Some of the most important factors of a package in a mobile device are cost, size, I/O count, and thermal performance. There certainly are tradeoffs between these factors, and a balance must be achieved to meet the goals of the package's intended application(s). When designing a package, the initial bottlenecks confronted are cost and I/O count. The device simply cannot function if there is insufficient I/O for the intended design or if it cannot be manufactured cost effectively.

Predicting package thermal performance in an end-use environment is not straightforward. A detailed understanding of all components, materials, power dissipation, and thermal behavior would be required to accurately predict a package's thermal performance in its end-use application [5]. Since this information is difficult to obtain and varies widely, thermal characterization metrics must be used to compare package performance and optimization. It should be expected that a package will often perform differently compared to its data sheet thermal resistances.

Package thermal performance is characterized using standard metrics of thermal resistance representing the various heat flow paths from the junction, or die. To standardize these metrics, they are measured using guidelines from JEDEC standards JESD51 series of documents [10]. Junction-to-ambient (Θ_{JA}) (JESD51-2A), junction-to-board (Θ_{JB}) (JESD51-8), and junction-to-case (Θ_{JC}), thermal resistances represent heat flowing along specific paths out of the package. Θ_{JB} is commonly preferred for characterizing packages in the mobile phone space because packages normally do not have external heatsinking present on the top surfaces and therefore will conduct a majority of their heat into the motherboard. When external heatsinking is present, such as in larger mobile platforms, Θ_{JC} is typically a more useful metric. The junction-to-board or junction-to-case resistance, by itself, does not provide insight into the maximum power dissipation a package can handle. Again, this would be highly dependent on the package's environment, the type of PCB it is mounted to, the enclosure, and any implemented heat sinking components.

An important design consideration for packages in the mobile space is the I/O count, the number of connections the package can support. A recent trend in packaging technology favors increasing I/O count while decreasing the package footprint. High I/O densities require advanced manufacturing techniques at the board level to support the fine pitch interconnects from the package. For complex ICs in a mobile application, an ideal package would have a high I/O count, low junction-to-board thermal resistance, and a small footprint. By fixing die size and power dissipation, Figure 14 shows the relationship between Θ_{JB} and footprint area for various packages. As a function of I/O count and package footprint, the I/O density, when high, is a desirable characteristic for packages in the mobile space. An ideal package would have a low junction-to-board resistance with a high I/O density, see Figure 15. These parameters, along with an understanding desired total power dissipation, should be used to evaluate packages for mobile applications.

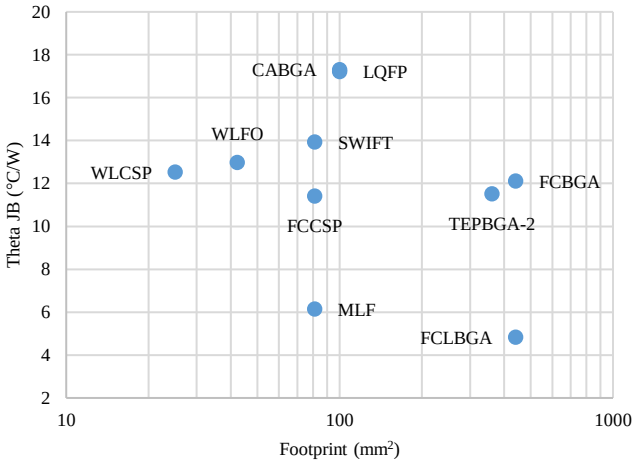

Figure 14: Junction-to-board thermal resistance versus package footprint for various packages with fixed die size and power dissipation.

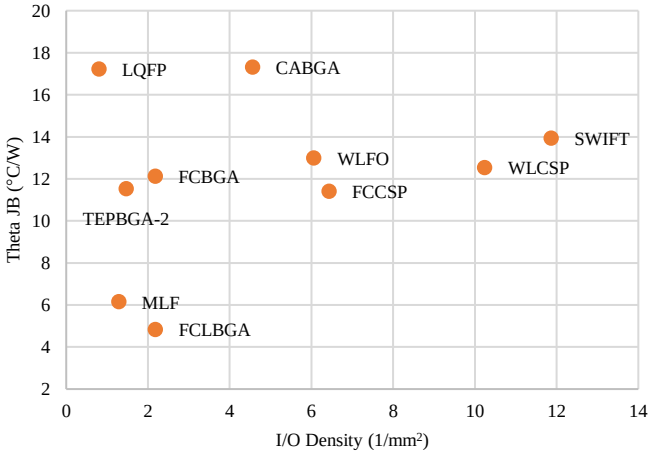

Figure 15: Junction-to-board thermal resistance versus package I/O density with fixed die size and power dissipation.

5.0 Thermal Enhancement – Materials & Package Structure

As package size trends towards chip scale, the opportunity for package level thermal enhancement in terms of geometry and material properties diminishes. As discussed in a previous section, when the body to silicon size ratio is high, meaning the body is much larger than the silicon die, there is the highest opportunity for improving the heat spreading away from the silicon. This can be accomplished in a variety of ways. Typical methods include but are not limited to, heat spreaders (lids), high thermal conductivity mold compounds (for overmolded packages), and increased thickness inner metal planes. As the body to die ratio decreases, the effectiveness of these enchantments decreases to zero. When the die size is nearly equal to that of the body size, a heat spreader offers no advantage because there is simply nowhere to spread the heat.

Similar to body to die ratio, the overall package thickness, when large, has the greatest margin for thermal enhancement. It should be noted that the margin of thermal enhancement

opportunity does not necessarily translate to the most thermally optimal design. Thermal enhancement margin simply implies it has a higher capacity for improvement relative to its original state. For example, in a chip scale laminate BGA package, where the die is nearly equal to size of the body, a thick 4-layer laminate buildup will have more opportunity for thermal enhancement compared to a thin 2-layer buildup. However, since the thin 2-layer buildup package will have less thermal resistance from die to BGA because of the reduced number of layers and material thicknesses, it will ultimately have better thermal performance compared to the 4-layer configuration. So, although there may be little room for thermal enhancement, the design or configuration of the package can itself represent a thermally optimal design. Figure 16 shows sample thermal enhancement margin opportunities for various package types trending to chip scale.

Figure 16 values are relative to specific package sizes in a typical state, and would change depending on body size, die size, and other geometry specifics. A value of 1 on the chart indicates that the package design has no opportunity of thermal enhancement. This metric is obtained by comparing a standard package configuration to a fully thermally enhanced configuration with all available enhancement options. The wafer level package in the Figure 16 has no margin for thermal enhancement because: (1) the silicon is equal to the body size so there is no opportunity for spreading heat, and (2) with interconnects that can attach directly to the silicon, the thermal resistance along this path is already very low.

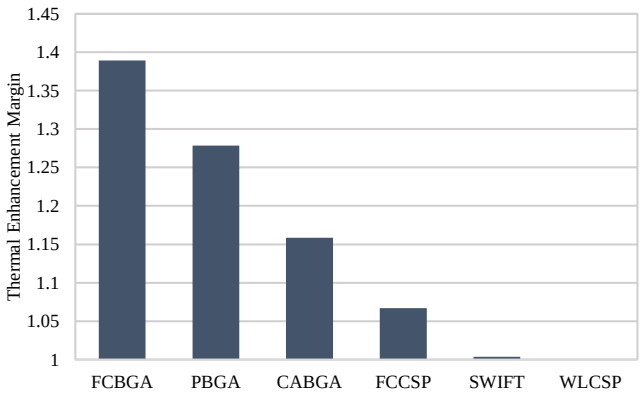

Figure 16: Realized margin of thermal enhancement relative to a standard configuration of that package type.

5.1 Thermal Enhancement – Processor Efficiency

An applications processor and communications processor in synchronous design dissipates power primarily due to switching power losses and leakage power losses, especially at lower supply voltages. Switching power loss is a function of capacitance, frequency, and the square of the supply voltage. Therefore, power loss is more sensitive to changes in voltage rather than capacitance and frequency. Power loss due to leakage is caused by leakage current within the silicon itself and consists of dynamic and static components. As the transistor size continues to decrease, power loss due to leakage current increases [4].

Even outside of the mobile space, nearly all high power processors are unable to maintain maximum power dissipation for extended periods of time. However, this scenario, for most applications is not required. Instead, the workload and thus power dissipation is highly dynamic, allowing thermal solutions to focus on the average rather than maximum power dissipation [7]. High power applications processors in modern smartphones implement multi-core architectures which provide multi-threading and efficiency advantages. Donald [8] demonstrated various methods for exploiting the distributed design of multi-core processors to improve thermal management. Their research showed that thermally aware core migration policies can dramatically increase performance through hotspot balancing [8].

One of the most basic dynamic thermal management techniques for processors is known as clock gating [7]. When the processor reaches a critical temperature, dynamic operations are suspended until temperatures fall back below a certain threshold. However, this sacrifices the user experience by disabling computing power. Adaptive voltage and frequency scaling (AVFS) is another common method for reducing power dissipation. AVFS greatly improves power efficiency of the device by dynamically adjusting the voltage and frequency depending on the current workload demand. Other power management techniques include, dynamic frequency scaling (DFS), dynamic voltage scaling (DVS), multiple supply voltage (MSV), and power supply shut-off (PSO) (to reduce leakage power losses) [4].

Figure 17 demonstrates a power throttling technique from measurements of a modern smartphone undergoing a benchmark test of its application processor. The phone was disassembled, fitted with temperature sensors at specific locations including the application processor, and reassembled back to original functionality. This allowed normal operation of the phone while being able to make temperature measurements of its internal components. The power throttling can be observed once a specific temperature threshold is reached. At this point the power and performance is reduced and temperature can now be successfully maintained below the thermal limit. [9].

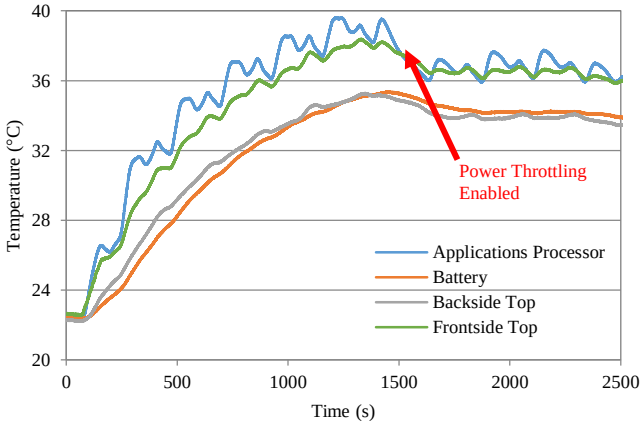

Figure 17: Applications processor temperature response measurements during benchmark workload. Power is throttled once thermal limit is reached. [9]

5.2 Thermal Enhancement – System Level

While the goal of this section is not to evaluate possibilities for thermal enhancement at the system level, some of the most common as well as promising options will be discussed to elevate the aforementioned discussions of package-level thermal challenges. As battery size in devices continues to grow, the motherboard size correspondingly continues to shrink. Although the motherboards in modern smartphones typically contain high densities of copper, new heat flow paths need to be considered for effective thermal management.

Since mobile platforms favor slim designs, active cooling solutions are nearly impossible given the available internal volume. In addition, with such limited internal air volume, any natural convections effects are almost negligible. Therefore the majority of mobile platforms are cooled passively, and improving passive cooling is mainly accomplished by improving heat spreading [12]. New techniques for thermal management include the implementation of heat spreading materials such as pyrolytic graphite sheets. Manufactured in thin sheets, the anisotropic crystal structure of the material provides an extremely high thermal conductivity in plane. Xiong et al. [11] investigated thin graphite heat spreaders in mobile platforms with thermal conductivities from 425 W/m*K to 1000 W/m*K. Both showed significant improvement in mitigating hotspots on the external surface when compared to traditional copper heat spreaders, reducing maximum temperatures by over 35%.

Heat pipes are commonly found in larger mobile platforms such as laptops and tablets, however only recently has their miniaturization led to implementation in smartphones, see Figure 18. Other advanced techniques for thermal management include the use of phase change materials in mobile platforms to mitigate transient temperature spikes. While phase change materials don't dramatically increase maximum power dissipation overall, they do offer improvement in the package transient cycles. Scott [15] demonstrated that phase change materials can delay time to peak temperature limits by nearly 2x compared to normal conditions.

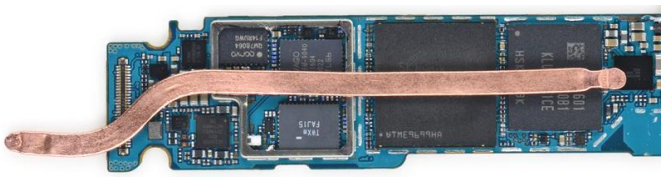

Figure 18: Micro heat pipe interfacing an applications processor in a modern smartphone [13].

6 Conclusion

The pressure to reduce the mobile phone's form factor while simultaneously increasing functionality, has driven package technology to extremely thin thicknesses, small footprint areas, high interconnect densities, and low cost. With the current trend driving towards almost entirely chip scale packages, there remains little room for thermal enhancement at the package level. Relative to their size, chip scale packages represent an optimal thermal design, with a balance between footprint and I/O count. At the scale of chip scale packages, the motherboard and the system itself, has effectively become the package. The largest opportunity for thermal enhancement should be focused on the system level and power efficiency through electrical optimization.

When designing packages for the mobile space, engineers should consider the thermal capacitance, thermal resistance, and I/O density to optimize the package for its intended application. In addition, the system-level thermal management solution should drive package design in terms of the most effective direction for heat dissipation from the package.

References

1. Y. Hang and H. Kabban, "Thermal management in mobile devices: challenges and solutions," 2015 31st Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), San Jose, CA, 2015, pp. 46-49.
2. https://www.gsmarena.com
3. V. Chiriac et al., "A figure of merit for mobile device thermal management," 2016 15th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2016.
4. Y. Huh, "Future direction of power management in mobile devices," IEEE Asian Solid-State Circuits Conference 2011, Jeju, 2011, pp. 1-4.
5. A. Carroll and G. Heiser, "An Analysis of Power Consumption in a Smartphone," 2010 USENIX Annual Technical Conference, 2010, pp. 21-35.
6. G. Wagner and W. Maltz, "Comparing Tablet Natural Convection Cooling Efficiency," Engineering Edge, Volume 3, Issue 1.
7. D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," Proceedings HPCA Seventh International Symposium on High-Performance Computer Architecture, Monterrey, 2001, pp. 171-182.
8. J. Donald and M. Martonosi, "Techniques for Multicore Thermal Management: Classification and New Exploration," 33rd International Symposium on Computer Architecture (ISCA'06), Boston, MA, 2006, pp. 78-88.
9. P. Fosnot, "Thermal Design in Mobile Application Space," 2016 32nd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), San Jose, CA, 2016.
10. JEDEC Standard JESD51 Series Documents. Electronic Industries Association.
11. Y. Xiong et al., "Thermal tests and analysis of thin graphite heat spreader for hot spot reduction in handheld devices," 2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Orlando, FL, 2008, pp. 583-590.
12. M. Carbone, "Influence of temperature control limits on passively cooled computing system performance," 2015 31st Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), San Jose, CA, 2015, pp. 41-45.
13. https://www.ifixit.com. Samsung Galaxy S7 Teardown. https://creativecommons.org/licenses/by-nc-sa/3.0/legalcode
14. Q. Li et al., "Technical challenges and novel passive cooling technologies for ultra-thin notebooks," 2017 16th

IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, 2017, pp. 1069-1074.
15. R.Scott et al., "Thermal Management of Portable Electronics using Phase Change Materials: Initial Experiments," First Pacific Rim Thermal Engineering Conference, PRTEC 2016.

© 2018, Amkor Technology, Inc. All rights reserved. SWIFT, ChipArray, and MicroLeadFrame are registered trademarks of Amkor Technology, Inc.