

THERMAL MANAGEMENT OF POWER ELECTRONIC BUILDING BLOCKS

By
William A. Stinnett

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University in
partial fulfillment of the requirements for the degree of

Master of Science
in
Mechanical Engineering

D.J. Nelson, Chairman
C. L. Dancey
A. A. Elshabini

10 February 1999
Blacksburg, VA

Keywords: Thermal Management, Power Electronics, PEBB, MCM

Copyright 1999, William A. Stinnett

THERMAL MANAGEMENT OF POWER ELECTRONIC BUILDING BLOCKS

William A. Stinnett

Abstract

Development of Power Electronic Building Block (PEBB) modules, initiated through the Office of Naval Research (ONR), is a promising enabling technology which will promote future electrical power systems. Key in this development is the thermal design of a PEBB packaging scheme that will manage the module's high heat dissipation levels. As temperatures in electronics are closely associated with operating efficiency and failure rates, management of thermal loads is necessary to ensure proper and reliable device performance.

The current work investigates the thermal design requirements for a preliminary PEBB module developed by the NSF Center for Power Electronics Systems (CPES) at Virginia Tech. This module locates four primary heat-generating devices onto a copper bonded substrate in a multi-chip module format. The thermal impact of several design variables (including heat sink quality, substrate material, device spacing, and substrate and metallization thickness) are modeled within the multi-layer thermal analysis software TAMS™. Model results are in the form of metal layer surface temperatures that closely represent the device junction temperatures. Other design constraints such as electrical and material characteristics are also considered in the thermal design.

Design results indicate for the device heat dissipation levels that a low resistance heat sink coupled with a high conductivity substrate, such as aluminum nitride, are required for acceptable device junction temperatures. Substrate performance, in the form of a spreading resistance component, will be negatively affected by a lower quality heat sink. Both forced air and cold plate cooling methods were found acceptable; factors such as environment, cost and integration will determine which solution is most feasible. Maximum surface temperatures can be lowered somewhat through adjustment of device spacing. However, this reduction was small compared to the impact on parasitic capacitance. Additionally, there is some thermal benefit to thicker high-conductivity substrates, whereas lower conductivity substrates will increase the maximum surface temperature. Thicker copper layers will prove beneficial though this benefit is not as great for higher conductivity substrates.

Also discussed are the on-going and future development efforts that are expected to require thermal consideration. These consist of a top-level thermal bus for additional heat removal, the use of metal matrix composites and concepts for multi-module integration.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor, Doug Nelson, for his virtuous patience and continued support throughout this ‘long and winding road’. In addition, I would like to thank my family, especially mom and dad, for their unwavering emotional support during the times when all seemed lost; they were there when I needed them the most. Lastly I would like to thank my friends and colleagues, including those at Goodyear, for their support, flexibility, and understanding. The albatross has now been lifted.

CONTENTS

ACKNOWLEDGEMENTS.....	iii
TABLE OF CONTENTS	iv
FIGURES.....	vi
TABLES	viii
Chapter 1: Background.....	1
1.1 Introduction.....	1
1.2 Thermal Management.....	1
1.3 Reliability.....	2
1.4 Power Electronics.....	4
1.4.1 Multi-chip Modules (MCM).....	4
1.4.2 Power Electronics Operation.....	6
1.4 PEBB Initiative.....	8
1.5 Thesis Objectives.....	9
Chapter 2: Design Strategy.....	10
2.1 Previous Research	10
2.2 PEBB Specifications	12
2.3 Design Software	14
2.4 Design Criteria	17
Chapter 3: Single PEBB Thermal Management.....	20
3.1 Design Approach.....	20
3.2 Heat Sink Selection	21
3.2.1 Preliminary Analysis	21
3.2.2 General Requirements	24
3.2.3 Results.....	29
3.3 Substrate Material.....	32
3.3.1 Preliminary Analysis	33
3.3.2 Candidate Materials.....	37
3.4 Substrate Size and Device Spacing.....	38
3.4.1 Substrate Size / Device Placement	39
3.4.2 Device Pair Lateral Spacing.....	45
3.5 Substrate and Metal Layer Thickness.....	47
3.5.1 Substrate Thickness	47
3.5.2 Metal Layer Thickness.....	48
3.6 Single PEBB Thermal Management Summary.....	51

CONTENTS

Chapter 4: On-going and Future Development.....	54
4.1 Top Path Thermal Bus.....	54
4.2 Metal Matrix Composites (MMC)	55
4.3 Multi-module Integration.....	57
Chapter 5: Conclusions and Recommendations	59
APPENDIX A: Heat Spreading	62
APPENDIX B: Geometry Iteration Summary.....	64
APPENDIX C: Preliminary Iteration Details	65
APPENDIX D: BASIC Code	70
APPENDIX E: Material Properties.....	74
BIBLIOGRAPHY	75
VITA	78

FIGURES

Figure 1: Variation of failure rate with component age	3
Figure 2: Multichip Module (MCM).....	5
Figure 3: Fundamental power electronics schematic – single phase	7
Figure 4: Fundamental power electronics schematic – three phase.....	7
Figure 5: Basic PEBB electrical schematic, power stage.....	12
Figure 6: Original PEBB layout.....	14
Figure 7: TAMS model schematic	15
Figure 8: TAMS source data detail	16
Figure 9: PEBB thermal model design variables.....	19
Figure 10: PEBB thermal model – lower convection coefficient (h2) detail	22
Figure 11: Original geometry thermal map, alumina substrate, natural convection.....	30
Figure 12: Original geometry thermal map, alumina substrate, forced air heat sink.....	31
Figure 13: Original geometry thermal map, alumina substrate, cold plate cooling.....	32
Figure 14: Original geometry thermal map, $k_{\text{substrate}} = 200 \text{ W/mK}$, natural convection ...	33
Figure 15: Original geometry thermal map, $k_{\text{substrate}} = 200 \text{ W/mK}$, forced air heat sink ..	34
Figure 16: Simplified spreading resistance network.....	36
Figure 17: Iteration 4 overall geometry.....	41
Figure 18: Iteration 4 thermal map, AlN substrate, forced air heat sink.....	41
Figure 19: Iteration 5 overall geometry.....	42
Figure 20: Iteration 5 thermal map, AlN substrate, forced air heat sink.....	42
Figure 21: Final iteration geometry.....	43
Figure 22: Final iteration thermal map, AlN substrate, forced air heat sink	43
Figure 23: Final iteration thermal map, alumina substrate, forced air heat sink	44
Figure 24: Final iteration thermal map, AlN substrate, cold plate cooling	44
Figure 25: Device lateral spacing range, final iteration substrate dimensions	46
Figure 26: Max. surface temp. vs. device lateral spacing	46
Figure 27: Max. surface temp. vs. substrate thickness.....	47
Figure 28: Max. surface temp. vs. copper thickness	48
Figure 29: Max. surface temp. vs. device spacing, single and double metallization.....	49
Figure 30: Max. surface temp. vs. substrate and copper thickness, AlN substrate.....	50
Figure 31: Max. surface temp. vs. substrate and copper thickness, alumina substrate.....	50
Figure 32: Bonded fin, forced air heat sink concept	52
Figure 33: Cold plate cooling concept	53
Figure 34: MPIPPS thermal design details.....	55
Figure 35: Metal matrix composite thermal model.....	56
Figure 36: Metal matrix composite thermal map, forced air heat sink	56
Figure 37: Multi-module solution using forced air, convoluted fin heat sink.....	58
Figure 38: General spreading resistance model.....	63
Figure 39: Simplified heat spreading (Biot number)	63
Figure 40: Original design, device spacing detail.....	65
Figure 41: Original design, overall geometry.....	65
Figure 42: Iteration 1 overall geometry.....	66
Figure 43: Iteration 1 thermal map, AlN substrate, forced air heat sink.....	66

FIGURES

Figure 44: Iteration 2 overall geometry.....	67
Figure 45: Iteration 2 thermal map, AlN substrate, forced air heat sink.....	67
Figure 46: Iteration 3 overall geometry.....	68
Figure 47: Iteration 3 thermal map, AlN substrate, forced air heat sink.....	68
Figure 48: Orig. geometry thermal map, alumina substrate, natural convection, $h_1=0$	69

TABLES

Table 1: Failure rates of common hybrid circuit elements.....	4
Table 2: MCM cooling evolution.....	6
Table 3: PEBB device heat transfer rates.....	13
Table 4: TAMS required source data.....	16
Table 5: Transition material characteristics.....	22
Table 6: PEBB substrate and metal layer characteristics.....	24
Table 7: Heat sink general requirements comparison.....	28
Table 8: Typical heat sink thermal resistances.....	29
Table 9: Spreading resistance for substrate (including metal layers), original geometry.....	35
Table 10: PEBB substrate materials property comparison.....	38

CHAPTER 1: BACKGROUND

1.1 Introduction

In years past, electrical devices that handled significant levels of power were typically packaged with large enclosures and fairly large heat sinks or base plates. At the time, the average efficiencies of these devices may not have been particularly good and no real method, other than experience, was widely used for predicting the size of heat sink required for maintaining a particular device temperature and operating life. Consequently, a thermal solution became one of using the largest heat sink and enclosure that would fit in the space provided. While the capability certainly existed for basic analyses, the current state-of-the-art in manufacturing technology didn't typically, at the time, have the capability of producing a cooling problem that couldn't be solved with a simple aluminum mass.

One of the first papers written on electrical device cooling was by Cockcroft (1925), who examined the temperature distribution inside the laminations of a transformer. This effort caused an influx of other work over the next twenty years, and by the time World War II had ended and the electronics age began, nearly all electronics development utilized some form of thermal analysis. An area of particular importance was the transistor.

As technology and competition expanded globally, the transistor evolved into one of the most important developments of the twentieth century. As part of the semiconductor family, its important role in computer central processing units (CPUs), hybrid microelectronics, power electronics, and countless other electronic devices, is without question. Along the way, as manufacturing capabilities progressed at an alarming rate, able to etch transducers that are just $0.35\ \mu\text{m}$ square, as in the Intel Pentium™ II (Intel, 1998), it is clear that thermal analysis needed to play an important role.

The intention of this thesis is to document the preliminary thermal design of an advanced power electronics device. This effort is coupled with the on-going efforts of the NSF Center for Power Electronics Systems (CPES), located at Virginia Tech, in the development of a Power Electronic Building Block (PEBB) module.

1.2 Thermal Management

In any electrical system, as in mechanical systems, the amount of heat generated can be found from the operating efficiency and power input. Lacking proper design, if the electrical system is not 100% efficient, operation can lead to thermally induced failure. For instance, a 100 kW electric motor that is 90% efficient will generate 10 kW in heat when operating continuously at full load. Left uncooled, this motor will overheat and eventually fail due to loss of magnetization of the rotor. In this system, a proper thermal management technique might consist of a closed-loop oil cooler that exchanges heat across a remote compact heat exchanger. If optimization has been pursued properly, such factors as pumping capacity, area of heat exchanger, and coolant type would all be design factors.

Most electrical and electronics development now employs thermal management, or thermal engineering, as a primary tool in aiding the engineer to a solution in packaging design. Elevated temperatures can adversely affect electronic device operation, reliability, power-handling capability, and achievable packaging density. The useful life of individual devices can be directly correlated to the number and magnitude of temperature fluctuations, in addition to temperature-induced creep, corrosion, and electromigration (Seraphim, et al., 1988). Therefore, the amount and means by which heat is transferred to the environment must be carefully considered. A balanced engineering approach is necessary when considering the many factors involved, and design studies are usually conducted with cost and space restriction limitations as standard procedure. Final designs are typically in the form of an optimized configuration, which minimizes device temperatures with an intelligent heat transfer solution.

With the aid of modern computers and advanced analysis software, engineers can accurately predict the temperature rise of a particular device from known power output, material properties, geometry, and operating environment. In the design of an effective packaging solution, however, it is often necessary to conduct iterations on this analysis, where systematic perturbations are applied to such variables as metallization thickness, device placement, heat sink quality, etc., until a balanced compromise has been achieved. Other factors such as packaging costs and numerous electrical issues must also be considered in this analysis. From this point, the expected reliability of the system can be predicted with statistical data such as component failure rates.

The evolution in electronics, and the associated need for thermal management, is not expected to slow any time soon. This is evident from an observation made by Intel's founder Gordon Moore, who, in 1965, noticed that semiconductor device density to date had doubled about every chip evolution, or 18 months. This simple observation is now known in the semiconductor industry as Moore's Law, and has held true ever since (Intel, 1998). Other observations have been equally insightful, recognizing that an exponential trend in the increase of semiconductor heat flux has existed over a number of years (Seraphim, et al., 1988). A rigorous approach to managing the heat generated in electronics is obviously required if reliability is to keep pace with manufacturing.

1.3 Reliability

The ultimate goal of thermal management, especially for electronics packaging, is the determination of reliability. In its formal definition, reliability is the "probability of a device or system performing without failure for the period intended under the operating conditions encountered" (Kraus and Bar-Cohen, 1983). The effective lives of electronics and other devices are typically characterized by the graph in Figure 1. The time t_b , called the burn-in time, is defined as the time at which the failures due to component malfunction, manufacturing defects, etc. are no longer observed. Beyond this time, failures occur at a much lower rate and in an unpredictable fashion. It is imperative that components be tested up to the burn-in time outside their intended system to avoid catastrophic failure. At time t_w , or wear-out time, failures occur at a much greater rate due to age, fatigue, or other malfunction.

The temperature-dependent failure rates for common hybrid circuit devices are shown in Table 1. Note the comparatively large failure rates and strong temperature sensitivity for densely packed devices (LSI) and high power devices (power transistors); both of these devices typically operate at higher temperatures.

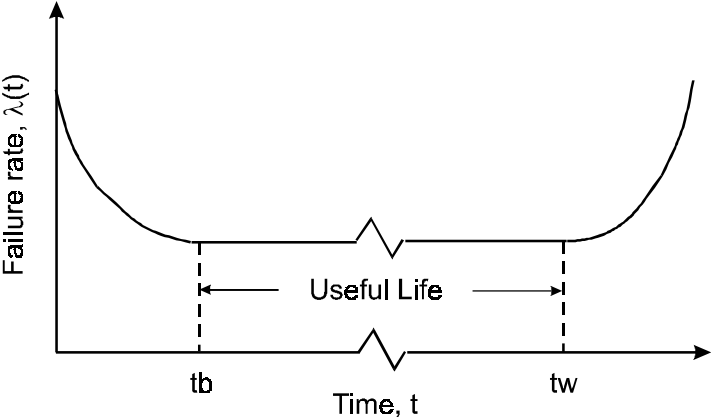


Figure 1 : Variation of failure rate with component age (Kraus and Bar-Cohen, 1983)

Table 1: Failure rates of common hybrid circuit elements, percent * 10⁶ per thousand (Kraus and Bar-Cohen, 1983)

Element	Temperature				
	25 °C	50 °C	75 °C	100 °C	125 °C
Thick film resistor	500	1,000	1,500	2,000	2,500
Chip capacitor	1,000	1,500	2,500	6,000	25,000
Wire bonds					
Au-Al ball	5	20	100	100	6,000
Al-Al	10	10	10	10	10
Au-Au	4	4	4	4	4
Transistor chips					
Low Power	100	300	900	2,700	7,000
Power	5,000	10,000	30,000	90,000	270,000
Diodes	100	300	900	2,700	7,000
Microcircuits					
Dual flip-flop	4,000	7,200	36,000	164,000	480,000
SSI (25 gates)	12,500	22,500	112,500	512,000	
LSI (100 gates)	50,000	90,000	450,000		

1.4 Power Electronics

One significant area where thermal management has played a vital role is in the area of power electronics. This area applies the fundamentals of the semiconductor to the intelligent and efficient conversion of both direct and alternating current to useful electrical power for use in applications such as three-phase electric vehicle inverters and uninterruptible power supplies (UPS). Devices of this type are also used in large-scale applications such as gas turbine powered electrical generating plants and drive systems on nuclear powered ships. The range of power levels these systems typically process is from the very low kilowatt to the high megawatt, either single or three phase. These applications represent the latest step in semiconductor application, where heat flux levels are even greater than those found in computer CPUs.

1.4.1 Multichip Modules (MCM)

Power electronic devices can be packaged into a form known as a Multichip Module (MCM). This device is so dubbed because it locates many individual transistor chips (dice) into a single module. MCMs are typically used for complex chip arrangements with many input and output (I/O) connections. The architecture of power electronics is simple in comparison, but the MCM format provides some unique advantages. By locating many transistors into the same package,

the length of interconnections are reduced, in addition to the chip to chip “time of flight”, thereby resulting in a more efficient design solution. There are other inherent advantages of this format, including device packaging density, manufacturing cost, switching speed capability, and reduced parasitic capacitance and resistance (Ginsberg, 1994). For power applications, it is possible to package the control circuitry within the module, greatly simplifying the integration effort. However, the much larger electrical processing capabilities creates problems for this type of MCM packaging, including the generation of excessive heat.

A very simple MCM assembly drawing is shown in Figure 2. Typically, the individual chips are in close thermal communication with a bottom-path thermal bus. This is one of the fundamental advantages of the MCM over the single chip module (SCM), which conducts heat through an array of solder connections (Simons, 1995). For the MCM, the chips are shown bonded to a copper metal layer that functions both thermally and electrically. This copper layer is, in turn, bonded to a ceramic substrate that electrically isolates the module; the bottom-most copper layer serves to spread the heat and match the coefficient of thermal expansion of a structural heat spreader. Multichip modules, in general, can have many different layers, chips (dice), electrical and thermal connections between layers, and input and output arrays.




Figure 2: Multichip module (MCM)

The evolution of the transistor from 1960 to the 1980s has realized a reduction in the required switching energy of more than four orders of magnitude, from 10^{-9} J to less than 10^{-13} J. Despite this reduction, the heat removal requirement for today’s chips has grown profoundly due to the increase in switching frequency and number of components that are integrated onto a single chip (Sherwani, et al., 1995). Specifically, the power dissipation for a single junction is found from:

$$P = \frac{1}{2} f_c V_{DD}^2 C \quad (1)$$

where f_c is the clock frequency of the circuit, V_{DD} is the switching voltage, and C is the capacitive load on the node. As an example, consider a system with 1000 nodes, a clock frequency of 100 MHz, an operating voltage of 3.3 V, and a capacitive load of 70 pF. The total power dissipation will be 38 W.

In analyzing or measuring the temperatures of devices on MCMs, or other semiconductor circuits, references are often made to the device operating or junction temperature. Formally, this is the temperature at the interface between the n-type and p-type semiconductors (Sherwani, et al., 1995). As this is difficult to actually measure, junction temperature more generally refers to average device temperature or the temperature at the bottom of the device. Junction temperature design limits vary from 80 ° to 180 °C, with the average limit about 120 °C.

The inherent difficulty with thermal management of MCMs is derived from the number of chips (or dice) located in close proximity to one another in a compact package, each operating at very high switching frequency (Sherwani, et al., 1995). In power electronics applications, there are generally many fewer devices, but each device has a much larger heat generation capability. Further complicating this matter are the applications for which MCMs are frequently developed, driven perhaps by their many advantages, where reliability and operating temperature requirements are very strict. As Table 2 indicates, the evolution of heat flux generated from state-of-the-art MCMs poses a significant challenge to the thermal packaging engineer. For comparison purposes, the preliminary specifications for the PEBB module are also shown.

Table 2: MCM cooling evolution (Simons, 1995)

IBM System	3081	ES/3090	ES/9000	PEBB*
Year	1980	1985	1990	1998
Max. Chip Power (W)	4	7	27	54♦
Max. Module Power (W)	300	600	2000	112
θ (K/W)	11.6	7.7	1.8	<1
Chip Flux (W/cm ²)	19	33	64	83
Module Flux (W/cm ²)	3.7	5.3	11.8	11.1

* preliminary ♦ IGBT devices

1.4.2 Power Electronics Operation

The operating principles behind many power electronic circuits are fundamentally the same (see Figure 3). Power is typically supplied to a load through a power processor, where load inputs are fed back to a controller. This controller, in turn, modulates the processor via an algorithm so that the load input more closely matches a reference. Because at least either the power input or output involves an alternating current where modulation of the current magnitude, frequency and phase are required, the power processor must be switched at variable frequencies through a transistor, which is simply a solid-state switching device. An example from this simplified schematic would be a computer uninterruptible power supply (UPS), where stored battery energy represents the direct current input and single-cycle, 60 Hz alternating current represents the output. The fundamentals of many other power devices can be represented by a similar schematic. For multiple current phases, such as those found in industrial three-phase motor

controllers, several of these type devices need to work concurrently, each controlling a single current phase (Figure 4).

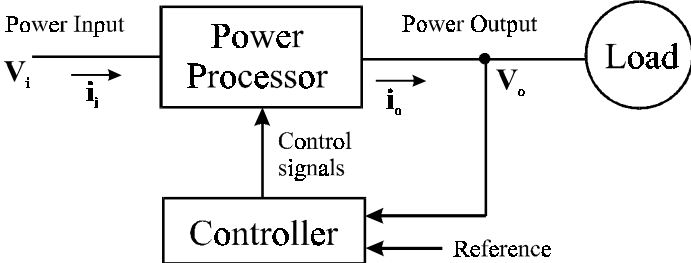


Figure 3: Fundamental power electronics schematic – single phase

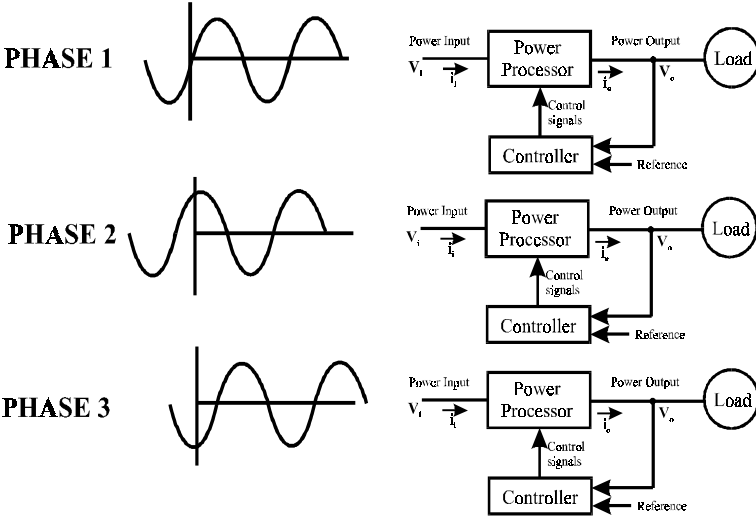


Figure 4: Fundamental power electronics schematic - three phase

1.5 PEBB Initiative

Until recently, there has been no serious initiative to take advantage of the inherent similarities in power electronics topologies. Historically, electrical engineers have designed unique solutions to individual power applications, often at the expense of high cost and long development time. These drawbacks have proven to be very restrictive for the development of systems that could use superior electrically controlled operation in favor of older, less robust mechanical methods. (As an example, the operation of many auxiliary rotating shafts on board naval vessels is typically controlled with hydraulic motors and valves. The use of electric power would allow much better control over speed and position, at higher levels of efficiency.) By realizing the similarities in power electronic topology, a modular system for such applications can be designed that reduces the overall cost and time to integration. This effort has been titled *Power Electronic Building Blocks* or PEBB in short. The application of the multi-chip module (MCM) format, with the advantages of integrated function and control in power modules, is the perhaps the best approach for the PEBB program.

The U.S. Navy initiatives, as detailed in the Office of Naval Research (ONR) SBIR 96.1, have been one of the first to call for such a modular power electronics solution. As outlined, the SBIR objective is to “develop enabling technologies for low cost, reliable, and easily manufactured modular electrical power systems for current and future Navy applications” (U.S. Navy, 1995). This effort requires an integrated systems-level approach for the modularization of power electronics components and packaging techniques. A large percentage of the Navy’s current machine control systems utilize either fully hydraulic or hybrid (electric and hydraulic) drives that typically require human involvement and are limited by reliability, size, and capability. Through the development of electric solutions, these systems can be made smaller, more robust, and capable of intelligent integrated and remote operation.

Potential application of PEBB technology will have vast influence in the industrial sector as well, as identified by the U.S. Office of Science and Technology Policy (Lee, et al., undated). These sectors include the energy storage, distribution, conditioning and transmission industries; pollution control; transportation; and communications. The transportation industry likely stands the most to gain from a modular power electronic device. As concerns about emissions and dependency on fossil fuels continue, the push toward high-efficiency electric vehicle development will be bolstered in a large part by the integration cost of electric motor controllers. In particular, the PNGV (Partnership for the Next Generation of Vehicles) goals of the USCAR (United States Council for Automotive Research), which is made up primarily of the big three auto manufacturers – Ford, GM, and Daimler-Chrysler, have initiated considerable development effort into electric and hybrid-electric vehicle drivetrain technologies. The PEBB has been targeted as a critical link to the commercial potential of these programs, and “could greatly reduce the cost of the drivetrain and increase the reliability of the vehicle’s electronic systems” (USCAR, 1999).

Virginia Tech and the newly-formed NSF Center for Power Electronics Systems (CPES), which is an evolution of the Virginia Power Electronics Center (VPEC), located within the Bradley Department of Electrical Engineering, have undertaken a lead role in developing a unique application of the PEBB concept. This application is called an Integrated Power Electronic

Module (IPEM) which is separate from the work on-going through ONR and numerous supporting industries, but parallels many of the long-term goals of PEBB. As part of this development, the Mechanical Engineering Department, under the direction of Professor D. J. Nelson, was requested to investigate the thermal design and analysis aspects of the packaging effort.

1.6 Thesis Objectives

The objectives of this effort are to investigate and optimize the thermal design requirements for locating four heat-generating elements (two IGBTs and two diodes) onto a copper-clad ceramic substrate such that ideal operation is possible and device junction temperatures are within acceptable limits. Constraints of the system are identified from several sources including cost, electrical, environmental, and manufacturing. This thesis will present a thermal design solution for packaging a single, first-stage PEBB multichip module with consideration of heat sink quality, substrate material and size, copper metallization thickness, and device placement. This effort was necessary for developing a fundamental understanding of the requirements for managing the high levels of heat output of the PEBB module.

Next, on-going and future developments for the PEBB thermal design will be summarized.

Short descriptions for each primary section follow:

- Background (Chapter 1) – discuss motivation for project, necessary background, and thesis objectives
- Design Strategy (Chapter 2) – cover previous research, project specifications, analysis tools, and design criteria
- Single PEBB Thermal Management (Chapter 3) – detail individual aspects of design criteria for single PEBB, present design summary
- On-going and Future Development (Chapter 4) – summarize current and future thermal design areas
- Conclusions and Recommendations (Chapter 5)

CHAPTER 2: DESIGN STRATEGY

2.1 Previous Research

Since multi chip modules (MCM) have been in wide spread use, numerous studies have been conducted to characterize or optimize the thermal impact of very specific aspects of their design, including the type of bonding, locations of thermal and electrical vias, and such details as lamination material properties. By contrast, there have been few simple yet fundamental thermal analyses, experiments, or design studies that have addressed the simultaneous issues of basic material geometry and properties, cooling schemes and component placement. There has been a fair amount of research published on individual aspects of this type of design, essentially pieces of a first principle approach. For the early stages of the PEBB project, a fundamental analysis was necessary to quickly gain a more complete insight into the parameters that play such an important role in this advanced development.

Kuhlman and Sehitoglu (1992) conducted analytical modeling of the temperature distribution of multichip modules using finite elements and compared the results to those obtained from the emissivity properties measured through an infrared camera. Design variations were conducted on bonding methods (i.e., wire bonding, TAB, flip-chip assemblies), substrate material, and device dielectric insulators. Results indicated, for bonding technique, that TAB and wire bond assemblies were more efficient than flip-chip types for both small to medium size I/O. For substrate materials selection, five different materials were confirmed to behave according to expectations. The authors recognized that advanced substrate materials such as silicon carbide and aluminum nitride were intelligent selections when considering thermal performance. Their conclusions also revealed that the polyimide dielectric, used to insulate the underside of the chip, imposes a significant penalty in terms of thermal resistance.

Lall et al. (1995) defined a simple superposition analysis technique for quantifying the temperature of the single quadrant junction of a generic MCM based on known power dissipation and junction-to-case thermal resistances. The authors' goal is to propose a technique to reduce the number of tests required to evaluate the thermal performance of an MCM. Experimental results are presented from different chip operating combinations under natural and forced convection. These results are compared to those obtained via superposition, where good correlation was observed. Future work plans include detailed computer modeling to further characterize the module temperature distribution.

Hahn et al. (1998) presented a design study for high power multichip modules that utilizes novel substrate embedding and minimization of junction-to-ambient thermal resistances. Through the use of planar embedding, a consistently low thermal resistance can be achieved between the heat sink and die, while also offering greater packaging and interconnection density, more reliably than standard attachment techniques. This embedding concept requires that fabrication be addressed such that heat removal can be achieved from back side of the die. Junction-to-ambient thermal resistances are further minimized through the use of adhesive interfaces, in lieu of thermal grease. Extremely high equivalent convection coefficients are attained through a liquid cooled, high-performance, micro-channel heat sink fabricated from aluminum nitride ceramic.

Hussein et al. (1991) compared experimental temperature measurements of four different hybrid circuits, comprised of bonded thick film resistors, with predicted results from a computer model. The model was defined in terms of several non-dimensional parameters, which were generated from model variables; these included device spacing, substrate and metallization thickness, material conductivities, and cooling medium quality. Analytical results were compared to those obtained from an infrared thermal imaging camera, which showed good correlation within measurement uncertainty. Also reported was the lack of significant thermal interaction (temperature rise due to spacing between two heat dissipating devices), which could be derived from sufficient device spacing.

In further investigation of thermal interaction, Hussein et al. (1992) developed a multi-layer thermal model to determine the impact of semiconductor spacing and other physical quantities on the junction temperatures at the surface of a copper clad ceramic substrate. Results were presented using previous non-dimensional parameters, and indicate that devices need to be placed further apart as copper thickness increases in order to avoid thermal interaction. It was revealed that an increase in copper thickness would lower surface temperatures for low conductivity substrates, but have the opposite effect for higher conductivity substrates, such as aluminum nitride or beryllia. In addition, lower heat sink thermal resistance results in lower temperatures for the same device spacing (less thermal interaction).

Yoo et al. (1996) used a naphthalene sublimation technique to experimentally investigate local heat transfer coefficients (h) around simulated single chips and two- and three-dimensional chip arrays. Several factors are identified as contributing to the surface-specific heat transfer coefficient – including air velocity, chip location, spacing distance of adjacent chips, and gap between chip and base plate. The single largest influence of heat transfer coefficient for all cases investigated was the chip to base plate gap distance. Additionally, for the same operating conditions, the two-dimensional case showed a slightly lower h than for the three-dimensional case, perhaps the result of a larger mixing effect.

Campbell et al. (1995) demonstrated the use of a hierarchical, three-dimensional thermal optimization routine, dubbed simulated “annealing”, for placement of heat generating chips in a given space. Simplistic examples were detailed. This automated technique iteratively locates a number of devices using a two step simulation that is intended to minimize computation time without considerably sacrificing accuracy. The two-level thermal analysis is performed first with a very basic non-matrix method that approximates the average temperature of the entire system, then individually alters the locations of the devices in coarse increments to find a lower average temperature. The second step involves the use of a finite difference technique that requires comparatively more processing time to fine-tune the solution using much smaller perturbations. Here, the temperatures of the devices are evaluated individually for an ideal minimum.

Lee et al. (1995) developed a geometric strategy that considers device placement for both reliability and routability using a “divide-and-conquer” approach. This complex technique involves forming clusters from sets of components with a so-called clustering algorithm, where identified hot elements are grouped with cooler elements and I/O components. After each cluster

is iteratively solved, it is assigned an area where the hot element is located at the geometric center of each area. Cooler elements are then located within the cluster by simultaneously evaluating sets of force-balance equations that are intended to represent the analog of an ideally distributed thermal load. At this point, an actual layout is generated from which a routing length is approximated for wiring efficiency evaluation. This routing length is claimed to be only slightly off ideal, from a min-cut comparison – due to the nature of the clustering algorithm. A thermal analysis is then conducted using TAMS™ (Ellison, 1983). Results have indicated that this technique is capable of designing for increased reliability without compromising connection efficiency.

2.2 PEBB Specifications

Early stages of the PEBB electrical development resulted in a preliminary design (VPEC, 1995) that located the primary power devices on a metal-bonded substrate. (The general MCM layout for a PEBB-type device is shown in Figure 2.) These device locations were identified while considering the parasitic inductance and capacitance concerns, the surface area requirements and chip wiring of the components given by the electrical schematic shown in Figure 5. At this stage, no significant thermal consideration was given to the details of the design. In this schematic there are four major devices identified, two IGBTs (Insulated Gate Bipolar Transistor) and two diodes, that produce the bulk of the heat transfer in the PEBB module.

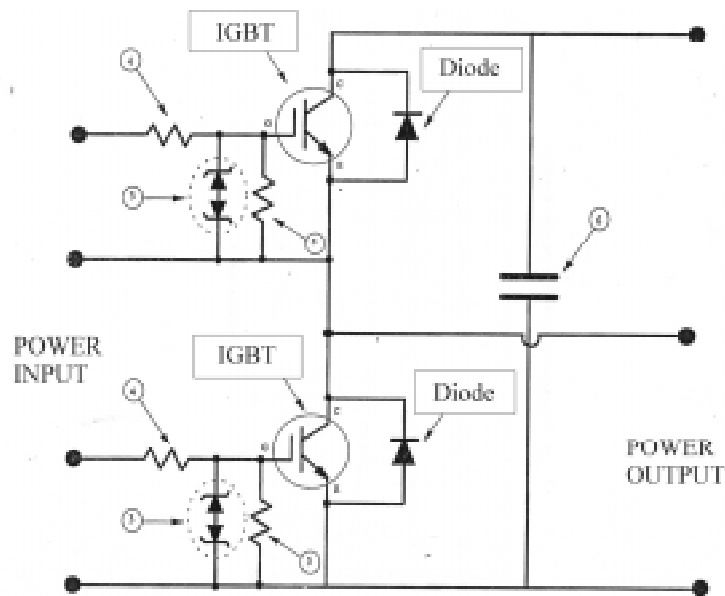


Figure 5: Basic PEBB electrical schematic, power stage

On a per-module basis, the power operating level is 3.33 kW and the design electrical efficiency is 96.8%. The corresponding heat transfer rates are given in Table 3 (VPEC, 1995):

Table 3: PEBB device heat transfer rates

Device	Qty.	Heat Transfer, ea (W)
IGBT	2	54
Diode	2	2
other components	-	negligible
Total		112 W

Details of the original VPEC layout (VPEC, 1995) including device locations are illustrated in Figure 6. The relative locations of the devices with respect to each other are given in Figure 40 from Appendix C; all dimensions are given in mm. Other elements of the schematic shown in Figure 5 are not illustrated. Metal layers were assumed continuous on both upper and lower surfaces. The metal-bonded substrate material, generally called direct bond copper (DBC), is a commonly used material in electronics applications and is also commercially available. This layout was used as a starting point for the design process.




Figure 6: Original PEBB layout

2.3 Design Software

The commercial software package TAMS™, Thermal Analyzer for Multilayer Structures, (Ellison, v. 3.1, 1991) was utilized in simulating and optimizing the thermal performance of the PEBB module. TAMS™ uses a Fourier series expansion of the steady-state, three-dimensional heat conduction equations to numerically solve for the surface temperature distribution at the geometric centers of each source (synonymous with junction or device temperatures). After separation of variables, the appropriate boundary conditions, and much effort, a numerical solution can be obtained for each node of the analysis that takes on the familiar form of the discrete Fourier series (Ellison, 1983). TAMS™ solves each node for the thermal effects of other sources using a fixed number of discrete Fourier series terms, which are set by the user. Convergence of the series can be checked graphically, and the number of terms adjusted if necessary.

Figure 7 represents a cross section of typical configuration for a TAMS™ simulation, in which a device generating heat Q_1 is located on the upper most surface and exposed to some ambient temperature, T_{1A} .




Figure 7: TAMS model schematic

and a convective heat transfer coefficient, h_1 . Up to four layers (t_1 , t_2 , t_3 , and t_4) can be simulated using TAMS™ – these layers represent the thicknesses of the different insulative, metallic or interlayers that make up the base surface of the model which, for this case, is a multichip module. The values k_1 through k_4 are the respective thermal conductivities (either isotropic or anisotropic) of the individual layers. Heat transfer at the lower surface, where a heat sink or thermal bus would be found, is characterized by an equivalent convective heat transfer coefficient, h_2 . This value can be solved by knowing the total thermal resistance, θ , from the heat sink and other elements, and the effective cross-sectional area (A_S) through which heat is transferred. The value A_S is same as the area of the four layers. The following equation illustrates the relationship:

$$h_2 = \frac{1}{\theta * A_S} \quad (2)$$

This assumption allows conductive boundary conditions to be analyzed as an equivalent convective surface with available thermal resistance information, which is commonly reported in most heat sink references.

In addition to the material and environment properties listed above, TAMS™ requires as input the individual source locations, heat generation rates and device dimensions, where rectangular source shapes are required (see Figure 8). Table 4 gives an example of the necessary source information, where distances are typically referenced in m, and Q is defined to be the heat generation for each source n in W:

Table 4: TAMS required source data

Source n	X location	X width	Y location	Y height	Q (W)
1	X1	DX1	Y1	DY1	Q1
2	X2	DX2	Y2	DY2	Q2
3	X3	DX3	Y3	DY3	Q3
etc.					




Figure 8: TAMS source data detail

To calculate the surface temperature distribution of the outer model layers, a grid of zero Q point sources may be defined that are superposed onto the existing devices (this was utilized in the PEBB design for temperature visualization).

The effect of interconnects, wire bonds, or stand-offs that lie outside the layers can be simulated through use of the thermal resistances, R1 through R4, as additional boundary conditions for the multi-layer model (see Figure 7). In this exercise, however, these were not significant.

The following limitations and assumptions are used in the TAMS™ model:

- a. The model exists at steady state.
- b. No thermal contact resistances exist between the substrate and metal layers
- c. The width and height dimensions of the model layers are limited to the same dimensions and must be continuous. TAMS™ will not allow the effects of pockets or thermal vias to be predicted.
- d. The sides of the different layers are limited to being adiabatic (insulated) only. This is valid when the material thicknesses are small in comparison to the overall surface area.
- e. Ambient temperature is assumed to be the same on both convective surfaces for ease of analysis ($T_{1A} = T_{2A} = T_A$). (TAMS™ will allow a correction for separate upper and lower surface temperatures, if needed.)
- f. Thermal conductivity for all materials is isotropic. TAMS™ does permit anisotropic conductivities, but this feature was not used in this work.
- g. Material properties are limited to being constant with temperature.
- h. Heat flux is assumed to be uniform across the device surface area. TAMS™ will allow simulation of a non-uniform heat flux by breaking the single source into a series of smaller sources.
- i. There are no heat losses from the surfaces due to radiation.

Three fundamental areas limit the overall accuracy of the modeling effort. First is the detail about which TAMS™ permits the physical problem to be modeled. Because of TAMS™ limited complexity, when compared to other finite element or finite difference computer codes, and the need for simplifying assumptions, not every problem characteristic can be modeled. For simple thermal analyses, however, where preliminary design studies are needed this is most often sufficient. Second is the extent to which the basic physical parameters such as material thermal conductivity and thermal resistances are known. This is typically no better than +/- 20%. Other parameters, such as convective coefficients, h , often have greater uncertainties. Third is the error that comes from truncation of the Fourier series used to converge the thermal model. As the number of terms is set by the user from within TAMS™, and can be qualitatively checked for convergence error, this contribution can be minimized for most physical problems. For very small heat sources, relative to substrate size, a larger number of terms are required for convergence. (This work uses relatively large heat sources.) Quantification of the total error from the TAMS™ model is difficult to predict with reasonable accuracy, in particular, because the individual component errors listed above are difficult to quantify.

TAMS™ is a well-proven tool for thermal analysis, as can be shown by the previous work in which comparisons with experimental results were correlated (Hussein, et al., 1991). These correlations were all shown to be in the range of their measurement uncertainty, approximately +/- 2 °C.

2.4 Design Criteria

Thermal management for MCMs such as the PEBB requires a multi-disciplinary design effort between the mechanical, electrical, and materials fields. To properly engineer a thermal solution all aspects of the design must be thoroughly considered.

The *general requirements* of the PEBB thermal design direction were decided at the outset of the project:

- a. Packaging should be kept as small as possible
- b. Design should provide for maximum life
- c. Costs should be minimized
- d. Flexible integration into a number of different environments
- e. Scalable with PEBB evolution
- f. Intrinsically robust

The *general requirements* of the design were considered from a system-wide perspective. Though not discussed in formal detail in Chapter 3 (an exception is the heat sink design section, which requires a separate sub-system analysis), these requirements are considered as a guideline for the design direction and are summarized in Chapter 5.

Implementation of the PEBB into its expected operating environment is key in understanding which thermal management design variables to consider first and to what level of depth the optimization should be executed. For the PEBB module, the expected operating conditions are almost entirely severe. For military applications, where an extreme operating environment may be found (such as the lack of ambient air in the confined space of a tank or the hull of a ship), the design direction might involve an approach into a remote cooling scheme, thicker substrates and copper films, or other implementations. As operating temperatures can be directly linked to the useful life of the device, cost vs. life trade-offs become an important consideration. For industrial applications, the environment may be somewhat less severe, where useful lives are readily achieved with multiple solutions and target costs dictate the design direction.

The *design variables* included in the optimization of the PEBB thermal model include (see Figure 9):

1. heat sink quality (h_2), cost effectiveness, and implementation issues
2. substrate material (k_s)
3. substrate size: height (A) and width (B)
4. device placement (X , Y) and spacing (C)
5. substrate thickness (t_s)
6. metal layer thickness (t_m)

These variables are the focus of this work and are discussed in detail in Chapter 3. Based on early collaboration with the electrical design team (VPEC), the PEBB thermal design was narrowed with the following assumptions: silicon devices with a maximum permissible junction temperature of 125 °C; very little natural convection on the upper surface of the PEBB ($h_1 = 10 \text{ W/m}^2\text{K}$); and a fixed ambient operating temperature of 30 °C. The metal layers on either side of the substrate were assumed to be copper, though future development may require other metals. In addition, processing limitations required a substrate size whose smallest dimension (A and B from Figure 9) was no greater than 50.8 mm.

The other specific design criteria that rise from the cross discipline nature of this work (from the electrical, mechanical, and materials fields) are called *system considerations*. While these are not addressed as hard design elements, they are nonetheless fundamental to the design process and are considered throughout the work. These considerations include:

- coefficients of thermal expansion (CTE) of adjacent bonded materials should be matched
- electrical parasitics (capacitive and inductive) should be minimized
- mechanical noise should be minimized
- design should provide for ease of manufacturing and integration
- safety should be maximized




Figure 9: PEBB thermal model design variables

CHAPTER 3: SINGLE PEBB THERMAL MANAGEMENT

3.1 Design Approach

For experimental analyses that involve the manipulation of numerous variables in determination of a particular trend, a systematic approach is desirable. This requires that the experimenter determine what variable, or combination of variables, has the greatest influence on a particular result. Complicating the procedure, however, the problem will often involve very large data sets, where the relationships are non-linear or otherwise hard to determine. Application of this analysis is most often needed for experimental trends that exist in complex manufacturing, chemical, or biological processes, and where conducting the number of otherwise necessary experiments would be prohibitively expensive (Montgomery, 1991). However, to avoid careless errors, even the most basic of experimental analyses should be methodical. To aid in this procedure a technique called *design of experiments (DOE)* can be used, of which books have been written, that streamlines the experimental process. The benefits of DOE are inherent in its ability to avoid redundant tests (thus shortening test time) and to help determine which elements contribute most to an outcome. It is not the intention of this thesis to discuss this topic in any detail, only to reference it as a method of experimental procedure.

Where computer models can be generated of a system, the benefits of DOE are not as apparent. Any more, simulations run on computer are very inexpensive compared to the costs required to conduct a laboratory experiment (in terms of labor, facilities, and equipment). There may be some instances where DOE is required even for computer experimentation, such as the study of geometry variation in CFD models. As such models generally require large amounts of computer resources, design of experiments would save considerable time to determine the impact of variable perturbations. For smaller models with less manipulated variables, a DOE approach would be unnecessary if the time needed to investigate and construct the test matrix was significant. In addition, for monotonic variable-result relationships and little or no cross correlation (as was expected in the PEBB design), a DOE may produce less than interesting results. For these reasons, design of experiments was not pursued for this thesis. Nonetheless, a systematic approach was still needed to accurately determine influential variables and to properly design a solution. For the PEBB thermal management project, the *design variables* listed in section 2.4 serves as a guideline for the concurrent aspects that were considered.

As can be expected, it was quite difficult to optimize each of the *design variables* individually without looking at the entire system closely. A concurrent engineering approach was used for the PEBB thermal design. Based on this, early design decisions were narrowed considerably through intuitive and careful inspection of the expected environment, in addition to the *general requirements* and *system considerations*. Once these decisions were made, it was more efficient to converge upon a final design. The next five sections will outline the individual aspects of the *design variables*, including the narrowing assumptions used, in attempt to thoroughly document the process.

3.2 Heat Sink Selection

The heat sink is perhaps the ultimate weapon in thermal management. An otherwise poorly designed electronic module could be brought to within acceptable temperature limits with a properly selected heat sink. This does not mean that the other *design variables* can be neglected if the cooling medium works sufficiently well. In the process of optimizing the heat sink for the application, it is necessary to maintain a balanced approach (i.e., concurrent engineering) such that the sink represents an equal part of a system-wide design.

In previous work, some electronic thermal analyses have utilized an isothermal boundary condition at the metal layer underside (David, 1977; Ellison, 1978; Negus and Yovanovich, 1987). While this greatly simplifies the analysis, it fails to take into account such elements as heat sink effectiveness, the bond between the substrate and sink, and ambient air temperature (Hussein, et al., 1992). This effort treats the heat sink, spreader, and bonding as separate external resistances with a particular ambient temperature. A much better representation is achieved with this approach.

3.2.1 Preliminary Analysis

As part of the collaborative effort with the electrical design team, the heat removal elements for the first generation PEBB module were to have taken the form illustrated in Figure 10. This architecture represents a common layout for a development MCM, where heat transfer efficiency is important – such as in a power electronics module. Here, the bottom-most metal layer of the MCM is brazed directly to a copper heat spreader. The heat spreader has two primary functions: first, to aid in spreading heat from the MCM to the heat sink, thus increasing the overall heat removal efficiency; and second, to act as structural reinforcement. The heat spreader is bolted to a heat sink through a thermal grease layer. The thermal grease serves to fill the air voids that occur at the heat spreader / heat sink interface, thus lowering the contact resistance – which was assumed to be negligible regardless. (It should be noted that the thermal grease layer is used for development efforts only, facilitating removal of the module. In service, the module would probably use a much higher conductivity permanent bond to the heat sink.) The characteristics of the transition materials (i.e., solder braze, heat spreader, and thermal grease) were set early in the design stages and are given in Table 5. The values used here are constant and represent typical conductivities and thicknesses for such a layout (Nelson, 1995).




Figure 10: PEBB thermal model – lower convection coefficient (h2) detail

Table 5: Transition material characteristics

Material	thickness (t) mm	Thermal Conductivity (k) W/mK
High Temp Solder Braze (B)	0.127	80
Heat Spreader (S)	3.175	400
Thermal Grease (G)	0.3	1

In considering the heat sink requirement for the PEBB module, the maximum permissible thermal resistance, from device junction to ambient environment was estimated from the basic one-dimensional heat conduction equation:

$$\theta_{\max} = \frac{T_j - T_a}{Q_{\text{total}}} \quad (3)$$

where T_j is the maximum junction temperature of 125 °C (given from section 2.4), the ambient temperature, T_a , is 30 °C, and Q_{total} is the sum of the components from Table 3, which is 112 W. This results in a maximum thermal resistance, θ_{\max} , of 0.85 K/W. It should be noted that this resistance includes the lumped resistive component that comes from heat spreading. (This is

essentially two- and three-dimensional conduction effects included with one-dimensional thermal resistance; a more complete explanation is given in Appendix A.) Because the heat transfer for all devices is summed, and the total area of the substrate is assumed to be fully utilized (uniform heat flux), the thermal resistance, θ_{\max} , is only an approximation. With the assumptions used, it is quite possible that a smaller thermal resistance would be required if the module has non-uniform heat flux or if there is significant thermal interaction, as was expected to be the case.

The actual total thermal resistance considered in the modeling effort for the lower boundary condition, which includes the components from Figure 10 that lie below the bottom-most metal layer, is designated θ_{h2} . This value is comprised of the transition materials given in Table 5 and the heat sink:

$$\theta_{h2} = \theta_{\text{transition}} + \theta_{\text{sink}} \quad (4)$$

The thermal resistance of the individual transition materials was approximated from the one-dimensional equation:

$$\theta = \frac{t}{k * A_s} \quad (5)$$

where the values for thickness, t , and thermal conductivity, k , are known, and A_s is the area through which heat is conducted, which is equal to the area of the metal bonded substrate. The total thermal resistance for the transition materials is then:

$$\theta_{\text{transition}} = \theta(B) + \theta(S) + \theta(G) \quad (6)$$

The designations B, S, and G represent the contributions of the high-temperature braze, heat spreader and thermal grease, respectively. As iterations are conducted on substrate size, the contribution of $\theta_{\text{transition}}$ will vary. These are listed in Appendix B.

In order to calculate the approximate heat sink thermal resistance, θ_{sink} , necessary to maintain the maximum permissible junction temperature of 125 °C, a comparison must be drawn to θ_{\max} . Here, the total thermal resistance, from junction to ambient is considered. Expanding Equation 5 (see Figure 10) to include the substrate and metal layer properties gives:

$$\theta_{\text{junction-ambient}} = (\theta_{\text{metal}} + \theta_{\text{substrate}} + \theta_{\text{metal}}) + \theta_{h2} \quad (7)$$

where θ_{metal} and $\theta_{\text{substrate}}$ are calculated from Equation 6 using the following first-iteration material characteristics which were provided by the PEBB electrical design team (Table 6):

Table 6: PEBB substrate and metal layer characteristics

Material	thickness (t) mm	Thermal Conductivity (k) W/mK
Substrate (Alumina, 96%)	0.635	20
Metal layer (Copper)	0.305	400

By substituting Equation 4 into Equation 7 and equating θ_{\max} with $\theta_{\text{junction-ambient}}$, the maximum sink thermal resistance, $\theta_{\text{sink-max}}$, can be found:

$$\theta_{\text{sink-max}} = \theta_{\max} - (\theta_{\text{metal}} + \theta_{\text{substrate}} + \theta_{\text{metal}}) - \theta_{\text{transition}} \quad (8)$$

which represents the largest heat sink thermal resistance that will support the module heat dissipation of 112 W. Solving this equation with original design properties given previously, including a substrate size, a surface areas, A_S , of 1008.1mm^2 (31.75mm square), $\theta_{\text{sink-max}}$ is found to be 0.51 K/W. This value represents the thermal resistance that might be found in a natural convection, high-fin-density heat sink of the type described in Table 8. This heat sink has a top surface approximately 120mm square (14400mm^2), which is significantly larger than the PEBB substrate area. It was assumed that the spreading resistance component for the sink is included in the total sink resistance.

The value for $\theta_{\text{sink-max}}$ assumes that the junction temperature will be maintained at 125 °C for the assumptions used. However, lower chip operating temperatures will permit devices to operate longer and more reliably, as indicated in section 1.3. For this reason, this resistance value was considered an upper limit and actual sink thermal resistances should be as low as possible, balanced with other *design variables* detailed in the next section.

3.2.2 General Requirements

The objectives of the heat sink selection were drawn from the *general requirements* of the PEBB thermal design, presented in section 2.4. A detailed discussion is presented here because of the relative complexity of this element of the design. These requirements can be expanded upon as follows:

- a. Packaging should be kept as small as possible

The primary contributor to the overall volume and mass of the packaged PEBB module is from the heat sink selection; as such, more stands to be gained from optimization of heat sink size. Every effort was taken to pursue as compact a heat sink as was practical. Special consideration was paid to both widely used and advanced technologies.

b. Design should provide for maximum life

Because device operating temperatures are directly linked to their failure rates, the heat sink choice should provide the lowest thermal resistance possible. This thermal resistance is generally proportional to the amount of heat transfer surface area that ultimately exchanges heat across ambient air. Heat sink solutions are available that will extract significantly larger heat flux levels than that from PEBB. However, this would not provide a balanced engineering design considering all requirements and would be excessively expensive to integrate.

c. Costs should be minimized

These costs will include those incurred from development in addition to hardware. Generally, a higher cost is associated with better heat removal capability. Some solutions, such as top-path thermal conduction modules, may require significant electrical design considerations for operation. The intent for this requirement was to integrate a heat sink solution with minimal intrusion into the module.

d. Flexible integration into a number of different environments

Expected operating environments for the PEBB module include military and industrial installations, where exposure to elevated temperatures and difficult integration issues are primary considerations. This may be the case in the confined spaces of a ship's hull or the interior of a tank. Depending on the application, remote operation with a separate heat exchanger, such as cold plate cooling, may be needed to take advantage of comparatively favorable ambient conditions. Stand-alone operation would always need an adequate supply of ambient air. Ideally, the heat sink should permit ideal operation in any number of volatile environments.

e. Scalable with PEBB evolution

Ideally, the heat sink solution should be compatible with future generations of PEBB modules that will process larger amounts of electrical power or handle multiple-phase loads (see Figure 4). An intelligent solution would be designed around the concept of a thermal bus, which would permit scaling based upon the amount of heat to be removed. Ideally, this would have two distinct characteristics: the ability to lower the thermal resistance within the same geometry, and the ability to link adjacent bus sections together to accommodate larger surface areas. For traditional convective applications, this could be accomplished with heat sinks that allow parallel flow paths and variable flow velocities. However, it is difficult to expect more heat removal from the same geometry in these applications because air is limited in heat capacity when compared to liquids, and flow velocities are generally restricted by fan capacity and heat sink geometry. These ideal characteristics would be achieved more readily with a liquid cooling solution, where cold plates can be linked together through simple plumbing alterations and elevated heat flux levels can be managed with larger pumps or fluids with greater heat capacity.

f. Intrinsically robust

The heat transfer solution should be designed such that the system can withstand shock, stress, and intermittent periods where design limits will be exceeded (i.e., ambient temperature, module heat flux); and should have a fail-to-safe contingency if temperatures become too high. This is especially important for forced-convection heat sinks where fans or coolant pumps fail and lead to elevated temperatures and, ultimately, device failure. A heat sink must be selected such that it can provide a reasonable factor of safety. In addition, for these applications, it might be necessary to install a thermocouple, air velocity sensor or fan operation monitor to reduce or shut down the system in the event of a loss in heat removal capacity.

In considering the primary route for heat removal in the PEBB, it was clearly intended to take place through the bottom of the module, as evidenced by the near-adiabatic value of h_1 given in section 2.4. The current design of the module locates the I/O and control and driver board (see Figure 2) at the upper most surface of the module, restricting any access to an additional thermal path. With future development, a top-level thermal bus could be investigated.

There are a fair number of design choices that might have been pursued in selecting an appropriate heat sink for this application. Such techniques as direct liquid immersion, hybrid air / liquid cooling, pool boiling, and liquid jet impingement are all solutions that have been proven to work exceptionally well at removing heat (Kraus and Bar-Cohen, 1983). However, detailing these alternatives is somewhat beyond the intent of this thesis and each would require a lengthy development period to work effectively with the PEBB module. The plan here is to provide an overview of the readily adapted heat transfer options, and to quantify their expected performance from reference sources and engineering data. It is not the intent to focus on heat sink design details or elaborate cost functions.

Based on the aforementioned *general requirements*, there are four heat sink classifications, narrowed from the numerous alternatives, which are suggested as viable options:

Air, Natural convection: worthwhile when buoyancy-induced flows can maintain the components at acceptable temperatures. This method is very simple and cost effective yet is generally only suitable for lower power density applications. There are no electrical power requirements. Heat sink types are numerous and utilize different design approaches to increase the heat transfer surface area, such as serrated extrusions, bonded fins, and convoluted fins (Figure 41 in Appendix B). Sufficient ambient air must be provided to maintain performance.

Air, Forced convection: cost effective and capable of removing a greater amount of heat than by natural convection but requires a large volumetric air supply. This method can be readily integrated into most environments. Sinks used for natural convection can be adapted for forced convection with the addition of a fan. However, failure of the fan could lead to very high temperatures and module failure if sink capacity is fully utilized; fan operation, airflow or device temperatures should be monitored. Noise can also be a concern.

Liquid, single-phase, Forced convection: on the basis of the much greater thermal capacity of liquids over air, offers the greatest heat removal capability at the expense of higher cost and more difficult integration. Power requirements could be substantially higher than forced convection of air, and continuous operation is critical for module survival. Plumbing is necessary from the heat source to a separate heat exchanger, where the ultimate sink of ambient air is still required for operation. Remote operation in hot environments is possible. Coolant loops can be combined for thermal control of multiple modules. Scaling is also readily achieved with larger pumps and interchangeable cold plates.

Liquid, two-phase, Natural convection: generally packaged as heat pipes, they are fairly new to the electronics packaging industry. They function in a similar manner to refrigeration units, utilizing the phase change energy of common liquids in a closed-loop for heat transfer. There are no moving parts and no power requirements. However, the condenser section must be located somewhat near the evaporator, due to the limited length of the pipes. Heat removal capacity is quite high, but limited by the size of the pipes and ability to return cooled liquid to the evaporator. The condenser section can be cooled with either natural or forced convection. For compact packaging they provide very distinct advantages. Orientation of the pipes could degrade performance.

For the first generation PEBB, these heat sink types were subjectively compared for heat removal capacity, cost, integration flexibility, and scalability. These characteristics are listed in Table 7.

Table 7: Heat sink general requirements comparison (Wakefield, 1996)

Heat Sink type	Thermal resistance	Cost	Integration Flexibility	Scalability
<i>Air, Natural convection</i>				
Extruded	Medium-High	Low	Fair	Good
Bonded Fin	Medium	Medium	Fair	Good
Convolute Fin	Medium	Medium	Fair	Good
<i>Air, Forced convection</i>				
Extruded	Medium	Low	Good	Fair-Good
Bonded Fin	Low	Medium	Good	Fair-Good
Convolute Fin	Low	Medium	Good	Fair-Good
<i>Liquid, single phase</i>				
Cold Plate	Very Low	High	Excellent	Excellent
<i>Liquid, two phase</i>				
Heat Pipes	Very Low	Medium-High	Fair-Good	Fair

Of the three different heat sinks listed for convection of air, both the convoluted and bonded fin types offer a very large surface area to weight and volume ratio when compared to the extruded fin heat sink. This is because extruded shapes are generally limited in aspect ratio (fin height to gap between adjacent fins) to about 8. Bonded fin sinks can have aspect ratios of 25 or greater (Wakefield, 1996). Convoluted fins, which are essentially compact heat exchangers similar to those found in automobile radiators, utilize an entirely different manufacturing process, and can have exceptionally high surface area to weight ratios. In a forced convection configuration, however, these sinks have a much higher pressure drop than bonded fin sinks and require larger or push-pull fan arrangements (Cowan, 1995).

Heat pipes offer a unique option for heat removal, but have fundamental limitations that make cold plate cooling more attractive in comparison. First, the limited length of the pipes and orientation requirements requires that the condenser and evaporator sections be located in close proximity to one another. This would permit limited remote operation; however, when compared to cold plate cooling, heat pipes have significantly less capability at nearly the same cost. Second, because of their closed loop, fixed-geometry nature, scaling the heat pipes for larger geometries or heat flux is not as flexible as cold plate systems, where parts can be readily interchanged. An increase in capacity would require an entirely new assembly at a significant cost.

A range of heat sink options was reduced from the many design choices; these were selected as solutions that offered the best tradeoff of design requirements discussed previously. The natural convection sink is given as a minimally acceptable solution; its thermal resistance of 0.41 K/W is just less than the $\theta_{\text{sink-max}}$ value of 0.51 K/W given earlier. Thermal resistances of the representative sinks are shown in Table 8, along with the total resistance used for the lower boundary condition. The values for θ_{h2} were used in the calculation of the *design variable* $h2$ that was modeled within TAMS™. The value for θ_{h2} varies based on each geometry iteration, due to the area component in the thermal resistance calculation from Equation 5. The difference between θ_{h2} and θ_{sink} comes from the contribution of the transition materials from Equation 4.

Table 8: Typical heat sink thermal resistances (Wakefield, 1996)

Heat sink type	θ_{sink} , K/W	θ_{h2} , K/W
Natural convection, air, bonded fin	0.41	0.72
Forced convection, air, bonded fin	0.10	0.41
Liquid cooling, cold plate	0.05	0.36

3.2.3 Results

The original PEBB specification, given in section 2.2, was used to benchmark the performance of the module. The three representative heat sink thermal resistances were converted to an equivalent convection coefficient, $h2$, along with the transition materials, as outlined in section 3.2.2. The results are illustrated in the form of a two-dimensional thermal map of the upper-most copper surface. These maps were created within Excel, where output data from TAMS™ was translated into Excel format with a BASIC program (see Appendix D). Temperatures reported were from the matrix of zero Q sources that were superposed onto the module surface. The x and y coordinates are broken up into an even number of segments, representing each zero Q source. Constant temperature bands (isotherms) are shown at 5 °C increments, with larger gradient areas designated by closer band lines. The legend gives the corresponding temperature range.




Figure 11: Original geometry thermal map, alumina substrate, natural convection

Figure 11 shows the original geometry PEBB module, modeled with a natural convection lower boundary condition h_2 (from Equation 2) of $1386 \text{ W/m}^2\text{K}$. This value was calculated using the area of the substrate / copper clad and the thermal resistance of the natural convection heat sink given in Table 8. The rather small substrate area results in the apparently large convection coefficient. As shown, a maximum copper surface temperature of $164.5 \text{ }^\circ\text{C}$ is achieved. (When node points on the plot coincide near the devices, the surface temperatures can be assumed to closely represent the junction or operating temperatures.) This is considerably higher than the estimated temperature for this heat sink resistance of $125 \text{ }^\circ\text{C}$. The difference is, in part, due to the fact that the assumptions used previously of constant heat flux and little thermal interaction are not valid.

When forced convection is used, as in Figure 12, the maximum temperature drops to $126.4 \text{ }^\circ\text{C}$. This is nearly a $40 \text{ }^\circ\text{C}$ reduction from the natural convection case, but still in excess of the maximum permissible temperature.

Using liquid cooling (Figure 13) reduces the temperature further, to $119.9 \text{ }^\circ\text{C}$, but the improvement is only marginal. Seemingly, the added cost of this system over forced convection outweighs its benefit; consideration must also be given to the other *general requirements* found in Table 7 before any judgment can be made.




Figure 12: Original geometry thermal map, alumina substrate, forced air heat sink

The results of the original geometry PEBB using the ranges of values for h_2 show temperatures that exceed or are slightly less than the maximum permissible for silicon devices. Other *design variables* must be exercised if the operating temperatures are to be lowered significantly within acceptable limits. As this preliminary design was established with consideration of electrical issues only, there are substantial improvements expected with the proper manipulation of *design variables* such as substrate material properties, substrate size and device placement, and metal layer thickness.




Figure 13: Original geometry thermal map, alumina substrate, cold plate cooling

3.3 Substrate Material

Substrates are pivotal in electronic packaging from an electrical, structural, and thermal design perspective. For this effort, it was specifically important to consider the thermal effect of candidate substrate materials to aid in the design process. A balanced approach with the *general requirements* for the PEBB was necessary as well.

The selection of alumina for the substrate material, from the original PEBB specification, represents the most common choice for MCM applications. This is based on its optimal combination of properties such as strength, electrical resistivity, and dielectric loss (parasitic capacitance). It also has a much higher operating temperature and thermal conductivity than most plastics and composites used in traditional electronics applications. In applications with very high heat flux levels, however, the thermal conductivity of alumina may not be sufficient to maintain proper device operating temperatures (Seraphim, et al., 1989).

Other substrate materials are available that offer many of the same advantages of alumina, but with higher thermal conductivity. Such materials as beryllia (BeO), aluminum nitride (AlN), and even diamond are being utilized with ever increasing frequency to manage the higher packaging and power densities of modern electronics.

3.3.1 Preliminary Analysis

The thermal effect of a higher k substrate material was determined simply by altering the variable within TAMS™, and plotting the results for the original geometry as a comparison. Here a thermal conductivity of 200 W/mK, such as that from beryllia or aluminum nitride, was used in the model. Figures 14 and 15 show the thermal maps using natural and forced convection of air, respectively. The figures show a reduction in the maximum matrix temperature of approximately 20 °C for both cases. With a forced air heat sink, the operating temperatures of the devices are within acceptable bounds which illustrates that substrate selection is an important consideration for the PEBB thermal design.

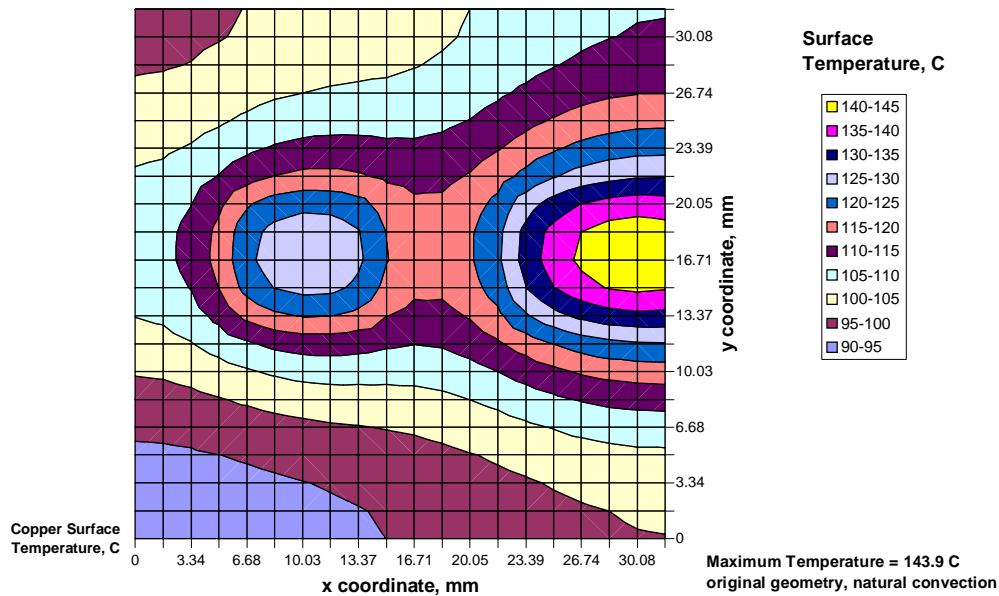


Figure 14: Original geometry thermal map, $k_{\text{substrate}} = 200 \text{ W/mK}$, natural convection




Figure 15: Original geometry thermal map, $k_{\text{substrate}} = 200 \text{ W/mK}$, forced air heat sink

An estimate of the required substrate thermal characteristics cannot be quantified readily because of the two- and three-dimensional conduction (or spreading) effects. The use of the one-dimensional thermal resistance relationship (Equation 5) alone would overestimate the acceptable resistance and substrate thickness or underestimate the required conductivity. It is possible, however, to estimate a representative spreading resistance component for the substrate layer from the three-dimensional TAMS™ results and consideration of a simplified thermal resistance network. From these results, the effect of device and substrate *design variables* (such as thermal conductivity, device size and placement, etc.) in addition to heat sink interaction effects on spreading resistance can be characterized. Related to the current work, this was used as a qualitative reference (Nelson and Sayers, 1992; Kraus and Bar-Cohen, 1983) to better understand some of the interactive influences at the device and substrate level.

In the simplified network diagram (Figure 16), the value θ_{sp} represents the spreading or constriction resistance as ‘seen’ by all devices. (It is shown linked to one device only for ease of explanation.) As θ_{sp} increases, the amount of two- and three-dimensional heat conduction will also increase. The variable θ_{msm} is the one-dimensional resistance for substrate layer, which is modeled in series with the spreading resistance (Kraus and Bar-Cohen, 1983). T_{avg} is defined to be the mean temperature of the substrate, which isn’t necessarily linked to a discrete position in the volume. θ_{h2} is the resistance (from Table 8) for the lower boundary condition used in the TAMS™ model. Spreading for the lower boundary is assumed to be included in the value θ_{h2} . The heat transfer, Q_{T} , is the dissipation sum from the individual devices (Table 3). Solving the simplified network:

$$Q_T = \frac{\Delta T}{\theta_{tot}} = \frac{T_{max} - T_{amb}}{\theta_{sp} + \theta_{msm} + \theta_{h2}} \quad (9)$$

then,

$$\theta_{sp} = \frac{T_{max} - T_{amb}}{Q_T} - \theta_{msm} - \theta_{h2} \quad (10)$$

where T_{amb} is the ambient temperature (30 °C) and T_{max} is assumed to be the maximum surface temperature on the upper surface (usually located at a device junction). This assumption of a localized high temperature could misrepresent the bulk spreading resistance if the devices are not very close together. The spreading resistance values are only used as a reference from which trends can be established, however, and absolute accuracy is not imperative.

The validity of these equations assumes that the heat loss from natural convection, $h1$, out the top of module is small compared to the total heat transfer. As a check, the maximum surface temperature from a model with a very small value $h1$ (~0), Figure 48 in Appendix C, was compared to that from a model with otherwise similar properties from Figure 11. A difference of less than 0.5% reinforced this assumption.

Spreading resistance components for the original PEBB geometry are found in Table 9. Values given for θ_{msm} , are calculated for the whole substrate / metal layer composite area.

Table 9: Spreading resistance for substrates (including metal layers), original geometry

Substrate	Convection	T_{max} K	θ_{h2} K/W	θ_{msm} K/W	θ_{sp} K/W
Alumina	Natural	164.5	0.7158	0.033	0.452
Alumina	Forced	126.4	0.4058	0.033	0.422
Alumina	Cold Plate	119.9	0.3558	0.033	0.414
BeO / AlN	Natural	143.9	0.7158	0.0047	0.296
BeO / AlN	Forced	107.1	0.4058	0.0047	0.278
BeO / AlN	Cold Pate	100.9	0.3558	0.0047	0.273

Results presented above show the effects of substrate material and heat sink quality on apparent spreading resistance for the original PEBB geometry. These resistances were used as a reference from which details regarding the effects on spreading could be inferred.

The first three entries shown above display the same alumina substrate with three different heat sinks. As θ_{h2} decreases with improved heat sink quality, the values for θ_{sp} also decrease,

indicating that heat conduction becomes more one-dimensional through the substrate. This implies that a change in external resistance will affect the level of heat spreading and, hence, the spreading resistance (this is explained in more detail in Appendix A). Based on this effect special consideration should be paid to the interaction between the sink and rest of the module. This is particularly important when considering the robustness of the design, where a reduction in cooling capacity would also include an increase in the substrate spreading resistance as well as device temperature.

When heat sink quality is assumed constant and substrate thermal conductivity is varied, a larger effect on spreading resistance can be seen. This effect is not as great as the difference in substrate resistances, indicated by θ_{msm} , would suggest. For the forced convection case, model results indicate that a substrate with 86% less thermal resistance will only offer 38% less spreading resistance when adapted into the original PEBB design.

As efforts are made to reduce the thermal resistance of the individual components of the PEBB, the resistance of the entire system must be carefully considered. As an example, the operation of a high conductivity substrate will be hindered with spreading effects when it is coupled with a high resistance heat sink (see Appendix A). The most effective design will match the general performance characteristics for all heat removal components. For the PEBB work, this indicates that both heat sink and substrate material should be of low resistance for maximum heat removal.

Other contributors such as substrate size and device placement can be identified as potential contributors to the spreading resistance θ_{sp} by conducting simple A-B comparisons similar to the manner presented. This would give a more complete picture of the interactions involved so that the performance of the substrate can be fully optimized within the module. For the PEBB effort, however, this exhaustive effort was not pursued.




Figure 16: Simplified spreading resistance network

3.3.2 Candidate Materials

In general, an ideal substrate material will have certain characteristics:

- Low dielectric constant
- High thermal conductivity
- Good electrical insulating properties
- Matched coefficient of thermal expansion (CTE) to bonded metal and devices
- High strength and toughness
- Low toxicity
- High operating temperature and thermal shock resistance
- Low cost
- Good machinability
- Smooth surface finish and low camber

There are numerous electronic substrate materials but perhaps only a few that are suitable for a high power, high temperature application such as PEBB. Fundamental in this determination was the thermal conductivity of the material. Based on early design decisions, three popular substrate materials were identified as potential candidates for the PEBB thermal design; these were alumina (Al_2O_3), beryllia (BeO), and aluminum nitride (AlN). Diamond was also recognized as being suitable for high power applications. Though its mechanical, thermal and electrical insulating properties are far superior to any of the other substrate materials, the cost of diamond currently makes its use somewhat restrictive (Gordon, et al., 1995).

A summary of the mechanical and electrical material properties of common substrates can be found in Appendix E. Table 10 illustrates the comparative qualities of the substrate materials, where the range of comparative qualities is base-lined at a single '+' score and scaled relatively.

Table 10: PEBB substrate materials property comparison (Seraphim, et al., 1988)

Property	alumina (Al₂O₃)	beryllia (BeO)	aluminum nitride (AlN)	diamond
Dielectric constant*	+	+	+	+
Thermal conductivity	-	++	++	+++
Electrical Resistance	+	+	+	+++
CTE matched to silicon	-	-	+	+
Strength, toughness	+	+	+	+++
Toxicity*	+	--	+	+
Cost*	+	-	+	---
Machinability	+	-	+	--

* lower is scored as better

From the substrate characteristics listed above, both beryllia (BeO) and diamond were discounted for possible application in the PEBB thermal design. Despite having excellent thermal conductivity, BeO was eliminated due to its potential toxicity and very high cost, about ten times that of alumina (Seraphim, et al., 1989). Aluminum nitride (AlN) offers nearly the same conductivity as BeO at a lower cost and without any recognized toxicity. The only recognized drawback of AlN is that the coefficient of thermal expansion (CTE) is not as well matched to that of copper, as is BeO. However, the CTE of AlN matches that of silicon better than either BeO or alumina; this is important when considering the direct bond of silicon die on substrates in MCMs. Diamond, on the other hand, is a nearly ideal substrate material but is still not readily available at a reasonable cost. It is expected that integration issues with diamond are not entirely resolved due to its limited applications. Based on these observations, both alumina and aluminum nitride are only considered as possible design solutions.

3.4 Substrate Size and Device Spacing

The original PEBB geometry, as mentioned earlier, was designed with limited consideration of the thermal requirements necessary for reliable operation. Locating the devices onto the substrate did, however, take into account manufacturing, system integration, and electrical issues such as parasitic capacitance, the details of which are beyond the scope of this thesis. In

proceeding with the thermal design, these issues were also closely considered in the form of *system considerations*. Based on meetings with the electrical design team, fundamental limitations were imposed on the layout of the module:

- The relative placement in a single diode / IGBT pair cannot be changed (i.e., diode cannot be moved relative to its neighboring IGBT).
- The lateral separation of each left- and right-side diode / IGBT pair should be kept to a minimum; this is required to minimize parasitic capacitance.
- Devices must be located such that I/O and control signals can be routed without interference with other devices or metallization
- The smallest dimension of the substrate must be no larger than 50.8mm (2 in) due to metallization processing limitations

From these restrictions, iterative design changes were made within TAMS™ on the *design variables* of substrate size (height and width) and device placement. Relative device spacing was also modified to assess the thermal effects.

3.4.1 Substrate Size / Device Placement

The ideal scenario for this type of design is to place the heat generating devices about the substrate so that the heat flux is as uniform as possible. This will give the lowest surface temperature for given Q, substrate size, and environment condition by minimizing the thermal resistance from source to sink and the level of thermal interaction. However, this is seldom achievable because of other *system considerations* involved.

Starting with the original PEBB geometry, a total of six iterations were conducted on the substrate size and device placement. For each iteration the devices were not moved relative to each other, but instead, as a unit about the substrate. These iterations were conducted in simple steps: first to understand the effects of geometry changes and ensure that relevant *system considerations* could be met, such as manufacturing, integration and parasitics. Then, the design was further optimized to find the smallest geometry that resulted in acceptably low operating temperatures. The models were analyzed using TAMS™, with an aluminum nitride (AlN) substrate and a forced air heat sink environment, as discussed previously. Thermal resistances from representative heat sinks (Table 8) were calculated as equivalent convective coefficients, h_2 , based on the substrate area using Equation 2. A summary of these coefficients, including other iteration details, can be found in Appendix B.

Figures and results of the first three iterations are given in Figures 42 through 47 in Appendix C. These initial steps were conducted as preliminary analyses with the electrical design team to understand how the geometry changes would affect the temperature and electrical layout of the module. Details are provided strictly for completeness. Close collaboration was needed to ensure that the devices could be coupled to I/O and control connections, and to verify that processing limitations were not being exceeded. Progressing from the original geometry, iteration 1 appended 5.15 mm to the ‘A’ (see Figure 9) dimension such that the devices were now centered on the substrate in the x direction. This resulted in a reduction in the maximum

matrix surface temperature of 14.3 °C, from 107.1 to 92.8 °C. This large reduction came from moving the devices away from the right side of the module, which is considered an adiabatic boundary. Iteration 2 increased the 'A' dimension an additional 10 mm, while keeping the devices centered laterally, and resulted in an additional 3.7 °C improvement. Iteration 3 increased the 'B' dimension by 10 mm, and moved the devices up by 5 mm such that a similar amount of additional substrate lied on both sides of the devices (y direction). This resulted in a reduction of 1.6 °C. Note that the thermal maps for iterations 1, 2, and 3 all appear fairly uniform in both the x and y direction, unlike the original geometry given in Figure 13. The slight lack of uniformity in the y direction is the result of unequal distance from the IGBTs, which are the major heat producing devices, to the edge of the substrate.

The next three design steps are illustrated in the following figures. Iteration 4 (Figures 17 and 18) locates the approximate thermal center of the devices, at the vertical center of the substrate, and adds 2.7 mm to the 'B' dimension. The net improvement is only 0.8 °C over the previous iteration, which indicates that the conductive heat removal to the lower convective boundary is nearly no longer being improved by the increase in module area, and device thermal interactions are dominating. Iteration 5 (Figures 19 and 20) increases the substrate dimensions to 50.8 mm square (which represents a design limitation given in section 2.4) and centers the device composite both vertically and horizontally onto the substrate. Modeled surface temperatures show that the maximum increased by 0.3 °C for this iteration, which indicates that the added substrate surface area is at its useful limit for the design conditions. Iteration 6 (Figures 21 and 22), the final iteration, was modeled to determine the influence of centering the dominant heat producers (IGBTs) on the same substrate size; the result was a reduction of 1.1 °C over iteration 5, and a small improvement (0.8 °C) over iteration 4. Based on these results, this step was assumed to be nearing optimum, and just meets the design limitation of 50.8 mm maximum for the smallest dimension. The slight lack of vertical symmetry is due to the small heat contribution of the diodes.

In observation of the results from the final iteration, it appears as if some benefit could be gained by modeling other substrates that have an even larger 'A' dimension. However, due to the apparent diminishing returns of surface temperature from substrate size, and the fact that 50.8 mm (2 in) square substrates are readily available (VPEC, 19995), no additional modeling iterations were conducted for larger substrate sizes.

Figure 23 gives the results for the final iteration using an alumina substrate. The maximum temperature was found to be 103.1 °C, which is 17.2 °C higher than that from AlN. Acceptability of this temperature would be determined through failure rate information, such as that discussed earlier. Using Table 1 as a guideline, for the temperatures modeled, the power devices mounted to alumina would have a failure rate that is two to three times higher than that of AlN.

With cold plate cooling (Figure 24), the final iteration maximum matrix temperature is nearly 7 °C less than with forced air. When considering failure rate trends this solution may show that cold plate cooling is more attractive. The results for natural convection, from Appendix B, are on average, 35 – 40 °C higher than those for forced convection at all geometry iterations.




Figure 17: Iteration 4 overall geometry




Figure 18: Iteration 4 thermal map, AlN substrate, forced air heat sink




Figure 19: Iteration 5 overall geometry




Figure 20: Iteration 5 thermal map, AlN substrate, forced air heat sink

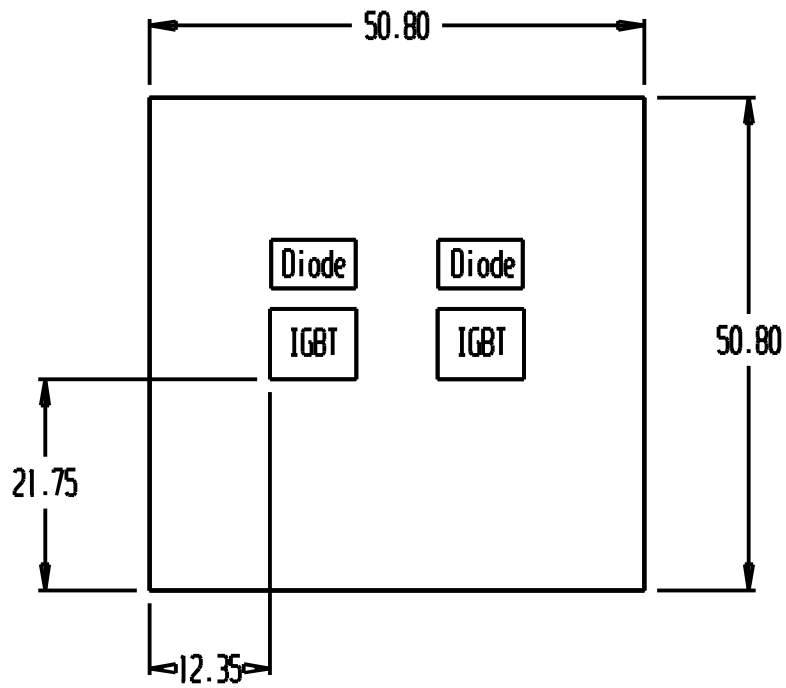


Figure 21: Final iteration geometry




Figure 22: Final iteration thermal map, AlN substrate, forced air heat sink




Figure 23: Final iteration thermal map, alumina substrate, forced air heat sink




Figure 24: Final iteration thermal map, AlN substrate, cold plate cooling

3.4.2 Device Pair Lateral Spacing

The effect of device spacing was investigated by adjusting the lateral separation of the device pairs from 1mm to 30mm (see Figure 25) for the final substrate size iteration. The intent here was to determine the optimum spacing, and to recommend a geometry change if a net benefit from all the *system considerations* could be concluded. The model was conducted with AlN substrate and a forced convection lower boundary condition. Maximum surface temperatures at the device centers, from the TAMS™ model, are found in Figure 26. These temperatures are defined differently than those from previous results, where the matrix of point sources at the upper surface may not have coincided with the actual source centers (and may be up to 2 °C less than the source maximum). For both alumina and AlN, minima occur at a spacing of about 16.5mm, which is approximately 2.5 °C less than the temperature from the original design of 8.3mm. This result is expected, as this spacing is equal to each device pair being centered on its own ‘half’ of the substrate area. At this distance it is intuitive that thermal interactions are at a minimum. The significance of this difference was determined by comparing the thermal advantage (failure rate, operating characteristics, etc.) with the implications on electrical, manufacturing, and system integration issues. In conclusion, the temperature improvement was deemed small compared to the negative impact on other design criteria.

Similarities in the general shape of the device spacing curve compare well with the thermal interaction results for low Biot numbers from the previous work (Hussein, et al., 1992). For the PEBB design, as device spacing increases, maximum temperature decreases up to a point where interactions with the adiabatic boundaries become significant. This differs somewhat from the previous results, which reported monotonic temperature trends by eliminating substrate dimensions from the analysis. (Within TAMS™ this can be accomplished by allowing the substrate to be very large in relation to the devices. Possible convergence problems could exist, however, and can be checked by setting a larger number of convergence terms.) To eliminate any effect of thermal interaction, based on results from the previous work, the non-dimensional spacing term D (defined to be the ratio of device width, assuming square devices, to device spacing) should have a value of about 10 or more. For the PEBB design, this would mean a spacing of at least 80mm. Up to this distance some thermal interaction would always exist.

Significant interaction can be seen for very small source spacing distances. For alumina, at lateral separations up to 2mm, the maximum temperature actually lies between adjacent IGBTs, instead of at the source centers. For AlN, distances up to 3mm will produce higher temperatures between the devices.

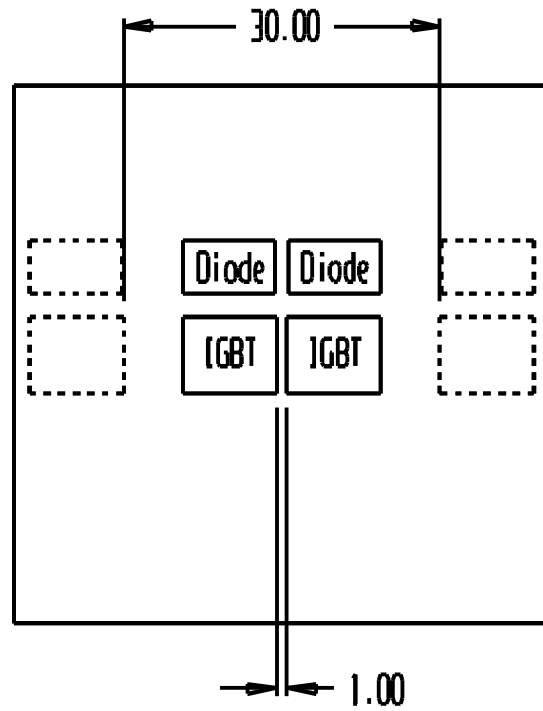


Figure 25: Device lateral spacing range, final iteration substrate dimensions

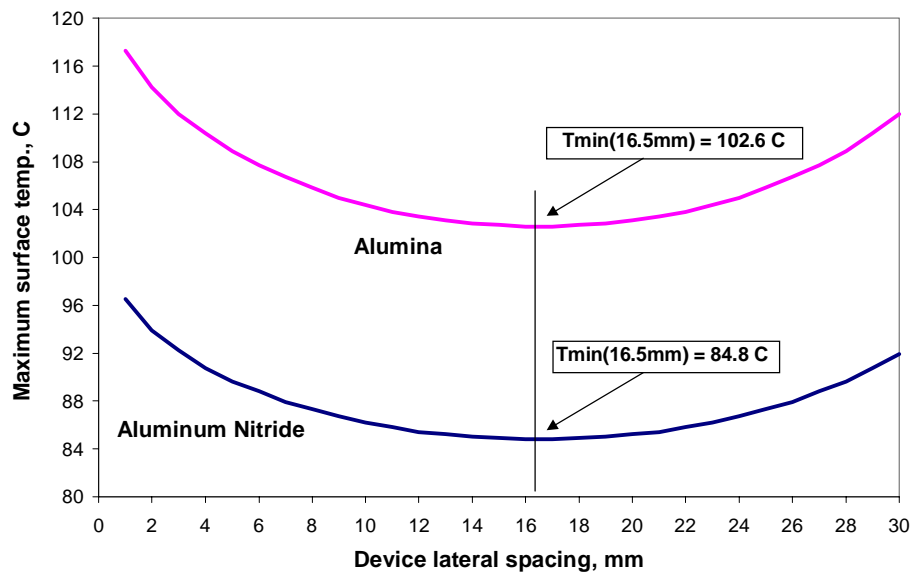


Figure 26: Maximum surface temperature (device center) vs. device lateral spacing

3.5 Substrate and Metal Layer Thickness

Proper selection of substrate and metal layer thicknesses is necessary when considering the design balance of acceptable temperatures, cost, and robustness. This final step of the model effort characterizes the maximum surface temperatures as a function of these material thicknesses.

3.5.1 Substrate Thickness

Thermal effects of substrate thickness on maximum temperature were also modeled within TAMS™. The models were run using both alumina and AlN substrate with 0.305mm (12 mil) copper clad, final iteration geometry, original design spacing, and a forced convection lower boundary condition. Maximum surface temperatures at the device centers are found in Figure 27. Results for the original substrate thickness are indicated.

These results shown are somewhat intuitive: for alumina, a greater thickness results in higher temperatures; the opposite is true for a high conductivity substrate such as AlN. From the slopes of the lines, it also appears as if there is a more to be gained by increasing the thickness of AlN, than by decreasing the thickness of alumina.

DBC materials are commercially available in a variety of substrates and thicknesses. Selection of a thick aluminum nitride substrate (at least as thick as the original design) seems to be a prudent choice for the PEBB module, considering the thermal load, general design requirements, and expected severe operating conditions where robustness is a concern.

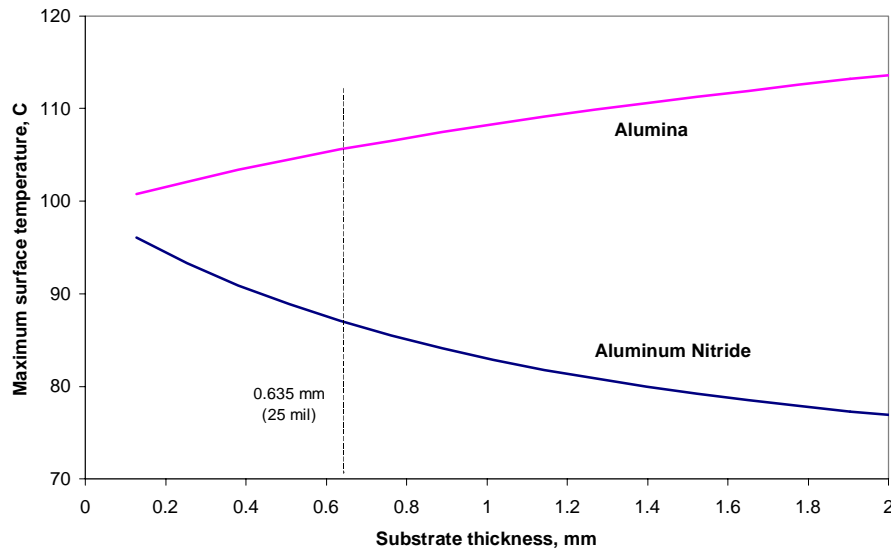


Figure 27: Maximum surface temperature (device center) vs. substrate thickness

3.5.2 Metal Layer Thickness

Metal layers are essential for heat spreading, in addition to being important for current and signal transmission. TAMS™ models were constructed to evaluate the effect of copper layer thickness on surface temperature. These models were run using alumina and AlN of 0.635mm (25 mil) thickness, the final iteration substrate geometry, and a forced convection lower boundary condition. Maximum surface temperatures at the device centers are found in Figure 28. Results for the original copper thickness are indicated.

For the substrate thickness used, both alumina and AlN benefit from increased copper thicknesses, and for very thick copper clad, greater than 1mm, the temperature difference is within 5 °C.

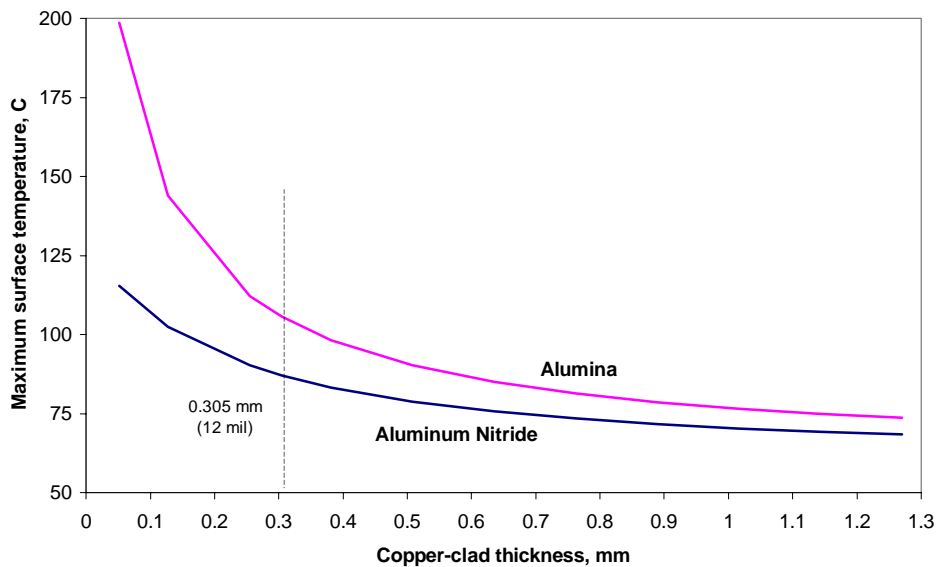


Figure 28: Maximum surface temperature (device center) vs. copper thickness

Actual MCMs such as the PEBB, or any printed circuit board, have sections of the metal layers etched away to form electrical connection paths. For the initial stage, it was difficult to determine the nature of these paths, other than their existence on the top layer only. The thermal impact of this partial layer was evaluated by modeling the substrate with a bottom metal layer only, and plotting these results with those from the double-sided model. The bounds set by either model represents a range of expected temperatures depending on the top-layer etched paths. These results are given for different device spacing, and are presented in Figure 29.

For alumina, the range is substantially wider than that of AlN, and surface temperatures could vary as much as 45 °C depending on the amount of metallization left on the upper surface. In fact, the upper bound for alumina pushes the surface temperatures well beyond the maximum acceptable for silicon devices. This result, in particular, reveals the inadequacy of alumina for the PEBB module. On the other hand, AlN is quite acceptable, with a variation of less than 15 °C. Properly spaced, devices mounted to AlN can be expected to operate at less than 100 °C.

More complete pictures are found in Figures 30 and 31 regarding the effect of both metal layer and substrate thickness on maximum surface temperature. These results are given for similar operating conditions and are intended to illustrate reasonable thickness ranges. For AlN (Figure 30), very thin substrate layers benefit significantly from increased copper layer thickness. As substrate thicknesses increase, the benefit of additional copper is quite small. The temperatures for alumina (Figure 31), on the other hand, can always be reduced by increasing the copper layer and decreasing the substrate layer thickness. For both substrate materials the thermal benefit of additional copper seems to outweigh the benefit of increased substrate thickness.

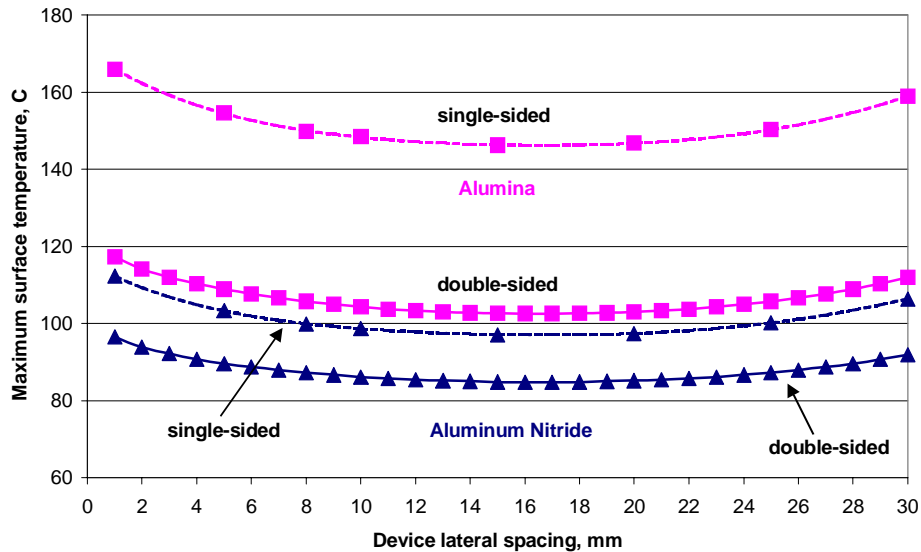


Figure 29: Maximum surface temperature (device center) vs. device spacing, effect of single and double metallization layer

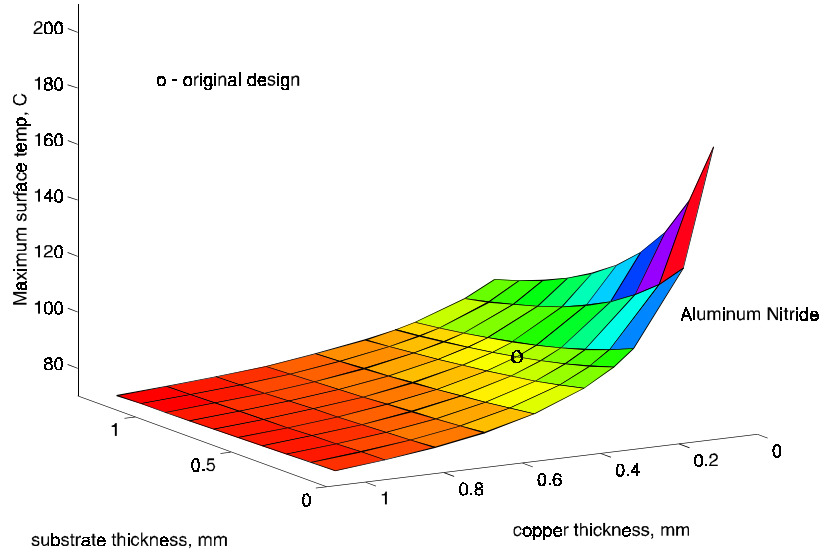


Figure 30: Maximum surface temperature (device center) vs. substrate and copper thickness, AlN substrate

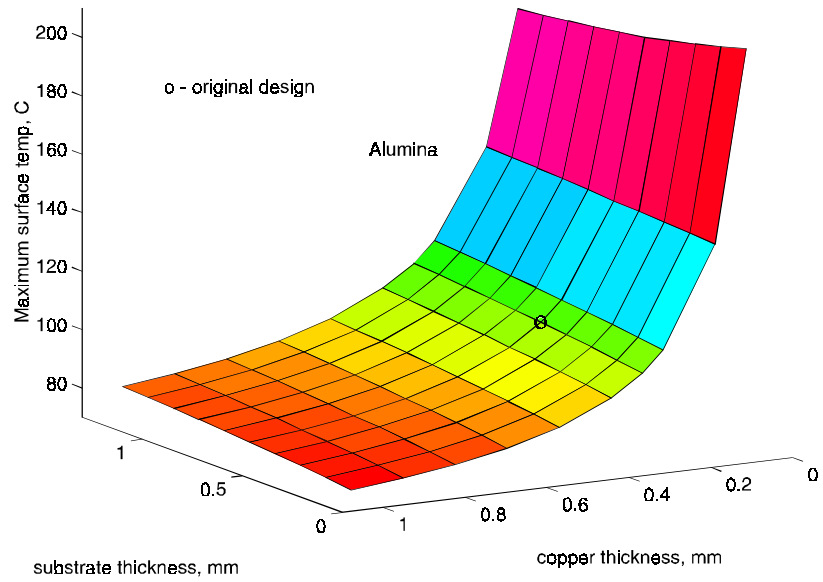


Figure 31: Maximum surface temperature (device center) vs. substrate and copper thickness, alumina substrate

3.6 Single PEBB Thermal Management Summary

Heat Sink (section 3.2)

Based upon the TAMS™ results, the use of a high-quality natural convection heat sink (air) produces unacceptably high surface temperatures for the PEBB power levels, even with a high conductivity substrate. A forced air heat sink does provide much lower thermal resistances, which result in acceptable surface temperatures. These sinks require a sufficient source of ambient air for proper operation, but are a cost-effective solution. Figure 32 illustrates a representative forced convection solution for the PEBB, where the heat sink dimensions are approximately 120mm (W) x 114mm (D) x 133mm (H). More severe environments could require a liquid cooling solution using a cold plate and a separate heat exchanger that takes advantage of remote ambient air. This solution is significantly more expensive than forced air, but can be readily scaled and requires comparatively less volume for the same heat removal capacity. The benefit of reduced device failure rates from lower operating temperatures is also a consideration. Figure 33 shows a representative cold plate PEBB thermal design, with more compact dimensions of 140mm (W) x 152mm (D) x 18mm (H). Both forced air and cold plate heat sinks are possible solutions for the PEBB, and will depend on the particular application.

Substrate Material (section 3.3)

The use of aluminum nitride (AlN) for PEBB seems quite clear due to its superior thermal conductivity, matched coefficient of thermal expansion to silicon, high strength, and low toxicity. Alumina resulted in unacceptably high surface temperatures. Beryllia and diamond were considered but discounted primarily due to toxicity and cost, respectively.

Substrate Size / Device Location (section 3.4)

Thermal maps from manipulation of substrate geometry and locations for the closely spaced devices indicate that the most uniform heat flux distribution arises from locating the vertical IGBT centers near the vertical center of the substrate. Alignment laterally is quite obvious – the line of symmetry for the device pairs should be aligned with the horizontal center of the substrate. Modeling efforts also indicate that a substrate size of 50.8mm square gives the lowest maximum surface temperature.

Device Lateral Spacing (section 3.4)

Maximum surface temperatures can be lowered somewhat (2.5 °C) by setting the IGBT lateral spacing distance at 16.5mm, which is greater than the 8.3mm spacing set from electrical design considerations. This improvement was concluded to be small when compared to the negative impact on other design areas, such as parasitic capacitance.

Substrate / Metal Layer Thickness (section 3.5)

There are thermal advantages (in addition to mechanical and electrical advantages) for both increased AlN substrate and metal layer thicknesses, with more benefit coming from thicker copper layers. The minimum material thicknesses are set from the original PEBB design. Increasing the thickness of a low conductivity substrate, such as alumina, results in higher surface temperatures. Availability of DBC substrates will limit the design options.




Figure 32: Bonded fin, forced air heat sink concept




Figure 33: Cold plate cooling concept