Thermal and Electrical Considerations for the Design of Highly-Integrated Point-of-Load Converters

Arthur Hugues Ball

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

Doctor of Philosophy

In

Electrical and Computer Engineering

Dr. Fred C. Lee – Chairman

Dr. Yilu Liu – Member

Dr. Guo Quan Lu – Member

Dr. Ming Xu – Member

November 18, 2008

Blacksburg, Virginia

Keywords: Point of load, converter, DC/DC, Integration, LTCC

Thermal and Electrical Considerations for the Design of Highly-Integrated Point-of-Load Converters

Arthur Hugues Ball

ABSTRACT

DC/DC Power converter design has been following a trend of reducing size while also increasing performance for the last several years. This push for higher power output and smaller footprint and profile requires integration and higher switching frequencies in order to continue. Higher frequencies require physical integration to eliminate problems induced by parasitics, which increase losses.

GE's Power Overlay and Philip's PCB integration schemes have been clear steps in the quest to reduce size with new system design techniques. However, both have downsides. GE Power Overlay embeds the devices inside a milled AlN ceramic cavity and then layers interconnections on top using polyimide dielectric interlayers. The milling of AlN ceramic is a very costly and time consuming task due to the brittleness of the material, and the interlayers add additional complexity to the fabrication process.

Philip's PCB integration was primarily aimed at integrating passives along with the PCB process for reduction of size. Inductor windings and capacitive layers were built up along with FR4 epoxy layers using typical PCB fabrication methods. However,

unlike GE's Power Overlay, the substrate material was several times lower in thermal conductivity which invariably has corresponding thermal penalties.

The work presented here reconciles the good of both integration techniques. Initially called Embedded Power, alumina ceramic was used as the substrate and rather than milling holes for the devices, holes were laser cut all the way through and interconnections were made by using solder masks and sputtered copper deposition, similar to GE's method. Integration of passives was done using LTCC ferrite to make an inductor of thin profile, rather than embedding cores and windings inside PCB. However, fabrication remained time consuming due to numerous solder masking and sputtering steps and thermal performance was not optimized due to the use of alumina ceramic coated in solder mask.

A revised design method called Stacked Power is presented in this dissertation that follows on the work of Embedded Power, but improves on it by simplifying fabrication through the elimination of thermally-restrictive solder mask layers, as well as time consuming sputtering and electroplating of copper interconnections. Instead, AlN Direct Bonded Copper is used as a multifunctional material thanks to its many-times-greater thermal conductivity than PCB or alumina, solderable device dies are implemented in a vertical fashion, and interconnections are simply made using copper straps soldered into place. For applications where moisture contamination and breakdown isolation are potential problems, dip conformal coating can easily be applied, replacing laborious solder masking.

The work in this dissertation describes the fabrication methodology for Stacked Power and demonstrates its layering concept of integration, along with its thermal

advantages, in the form of point-of-load buck converters that achieve super-high levels of power density in the smallest of volumes and require no more thermal management than modest, if any, airflow. The added cost incurred with aluminum nitride is traded for distinct advantages in terms of low-profile, low airflow requirements for the available power output, capability of natural convection for use in locations where fans are prohibitive and compact size for ease of implementation.

Acknowledgements

I would like to sincerely thank my advisor, Dr. Fred C. Lee for giving me the opportunity to do this work, for his guidance and vision, and for the invaluable life lessons he taught me through it all. Also, thanks to Dr. Daan van Wyk for his insightful questions and comments that get to the fundamental issues in a most direct and memorable way. I also thank the National Science Foundation who generously inducted me as an NSF fellow to be able to pursue this research, as well as the Bradley Department of Electrical and Computer Engineering for awarding me the Pratt Fellowship in honor of my academic credentials.

I could not have completed this work in such a short amount of time were it not for Ying-Feng Pang and Dion Minter who were instrumental in helping me with complex I-DEAS thermal simulations used for PCB characterization, so a very special thanks to both of them.

A lot of this work was a team effort and included Michele Lim and David Gilham who helped me with the LTCC inductor design and fabrication respectively. Also, Jinghai Zhou, Jess Calata, Zhenxian Liang, Daniel Huff, Yu Meng, Douglas Sterk, Jerry Francis, Qiang Li and Julu Sun helped with all sorts of issues and were great people to bounce ideas off of. And finally Yan Dong was an integral part of the making of the coupled-inductor and its testing for the 2-phase POL converter module. A warm thanks to all of them for making this enabling technology live up to its name.

My family is an integral part of this work despite their indirect association. My mother Anne and father Frank have always been there in the sense that their confidence in me was a guiding force for my perseverance and determination to get

through all four of my college degrees. And having my brother Philippe going through medical school at the same time gave me someone to relate to on a day-by-day basis.

I would like to dedicate this dissertation to my nephew and godson Xavier Alexandre Ball. My brother Philippe and his beautiful wife Brandy have this most wonderful son to whom I dedicate this work and effort in honor of his love of life. He will go far and I want this to be a part of his first steps.

But this dissertation might not have come about were it not for my partner Brad Whitney. He was there for me on the days when I doubted myself, when I felt like it wasn't worth the worry and anguish, but also when times were great. Without him to share it with on a daily basis, this entire experience would have lost an important part of its meaning for me.

And last but not least, I want to acknowledge my entire PhD committee who gave me valuable insight into how to make my dissertation better than it was. It would not have been as meaningful without their time and dedication.

This work was supported primarily by ERC Program of the National Science Foundation under Award Number ECC-9731677

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

1.1 Converter Design Overview

Thermal management is becoming one of the more critical tasks in the design of power electronic systems, including Point of Load (POL) converters and their specialized subset of Voltage Regulator Modules (VRM). In many cases today, thermal considerations are simply verified only after a complete electrical design is implemented. This sequential way of treating thermal “design” is fine as long as power density levels are low, resulting in comfortable thermal margins. But with increasing power densities, many designs are now thermally limited and rely exclusively on a very large heat sink to keep them from overheating. Iterations need to be done in many cases between the electrical and thermal designs to end up with a solution that is thermally capable without consuming excessive real estate. Since thermal limitations have become a real barrier to further increasing power density in power converters, their characterization is critical for understanding how to design a power electronic system that can work efficiently within these limitations.

Non-isolated POL converters are becoming more and more common as power electronics start to infiltrate all aspects of modern life. The range of possible applications is shown in Figure 1.1 and will soon reach $4 billion of market value. This makes the POL in a great position to reduce global energy needs through higher efficiency than the large power supplies of yesterday.

Figure 1.1. The wide range of POL applications.

One of the main methods for size reduction is the increase in switching frequency to reduce the size of necessary passives. But an increase in frequency requires careful layout to keep the parasitic inductances and resistances from undermining the advantages. As a result of this, packaging has also become a distinct issue since the reduction of parasitics requires closer proximity of components to shorten traces. This alone leads to thermal problems but also a question of how to interconnect and integrate all these components together in the most effective manner. New concepts like Intel's Bumpless Build-Up Layer (BBUL) aim to not only integrate components together in a fully compatible process but also keep size and profile low. The downside is that the build-up on silicon is not thermally conducive and so high performance with minimal thermal management is still in the distant future.

In Figure 1.2, many DC/DC converters are shown, representing a few types of topologies, to show the range of performance in the DC/DC converter field. The data point numbers are the same as their respective number in the list of references at the

end of this dissertation. Below 6 MHz, all the converters shown on the map are buck converters. Above 6 MHz, there is a mixture of Quasi-resonant topology (data point 37), boost (25 and 26), Class E (28-30), hetero-junction bipolar transistor RF circuits (27, 34, 36), and ultra-high frequency monolithic bucks (6, 39).

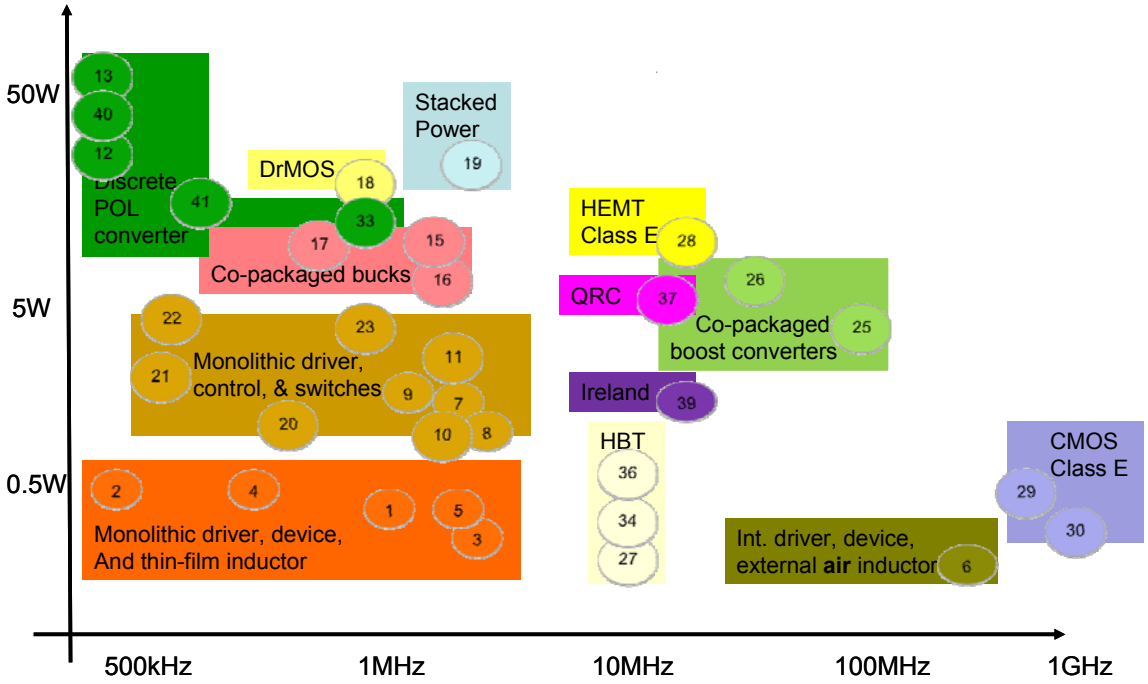

Figure 1.2. Overall packaging roadmap data, power versus frequency

Taking a look at a few examples of current industry trends, we can identify the current marketing positions and technical limitations. Today, the majority of converters are still discrete solutions (dark green data points) without a lot of effort going into pure integration. It appears that this is cost driven, at least for now. Here we will discuss briefly a couple examples of discrete converters since they represent the vast majority of high-power POLs today. In section 1.2, we will focus on aspects of integration for making higher-performance POLs.

The Artesyn PTH05020 is an example of the typical high-current POL converter currently seen on the market and is shown in Figure 1.3. It operates at a very-low 340 kHz switching frequency (putting it off the map), which is a range also seen in many of its competitors and represents the “classic” design method for a buck today. This low frequency has the advantage of less-critical circuit layout due to reduced parasitic effects, lower device switching losses and ease of control implementation compared to higher-frequency alternatives. The downsides are that the inductor is very large and physically dominates the converter. Also, a larger number of output capacitors are needed and the lack of integration results in an overall significant size increase over what is possible today. As a result, its power density is a low 65 W/in^3 at 5 V - 1.2 V conversion. Efficiency is a respectable 84% in this case and the large size reduces thermal issues, allowing a high 22 A output without the use of a heat sink.

Figure 1.3. Typical example of 20-25A discrete POL.
Used with permission from Emerson Technology

Next, we look at a high-current example with higher power density is the Power One ZY7120, data point #33. This module operates at a much higher frequency of 1 MHz, which is still uncommon in the POL marketplace today, especially in the 20 A output range, but typifies the trend. This converter’s power density is up to 100 W/in^3, and thanks to this higher frequency, we see a smaller low-profile inductor is used and all

necessary capacitors are onboard. The PCB is used as the primary thermal cooling mechanism through the use of many thermal vias, as seen in Figure 1.4, with the heat flow path designed to be out of the bottom of the POL's vertical orientation into the motherboard.

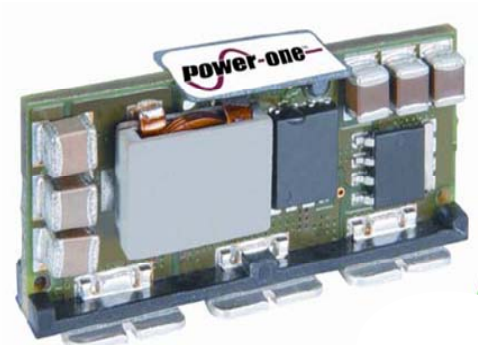

Figure 1.4. Example of high-frequency, high-current POL.
Used with permission from Power One

Looking at Figure 1.4, we see that large numbers of thermal vias have been added to the PCB to extract heat. However, the heat has a difficult time convecting off the surface of the PCB since the majority of the surface area is covered by the large inductor, many bulk capacitors and the drivers. Nevertheless, the large physical size of the module improves its heat capacity enough to make thermal issues less of a concern. Convection off the surface of the board is not very good since the majority of the surface area is covered by components, and the PCB itself, despite the large number of thermal vias, will still have a low composite thermal conductivity on the order of 35 W/m^2K (discussed in Chapter 2). The thermal bar connections at the base are designed to be soldered to the motherboard for conduction cooling to alleviate these issues but then the customer is expected to be able to accommodate this extra heat on

the main platform – a luxury that is getting increasingly difficult to handle in today's system designs.

1.2 POL Integration

Now that we have seen a couple typical discrete examples, we are going to get into the main interest of this dissertation: integration. Before going further, a discussion about the term "integration" is necessary. This word is now hackneyed and ubiquitous, representing many different design situations with no real clarity. Ideally it would mean that *all* the components are processed at the silicon level. This means that there would be a process that could make devices, interconnections, inductors and capacitors all at one time. However, reality is still very far from this concept.

POL design varies substantially depending on the output power needed from the circuit. For low-power applications (< 5W), on-chip converters are tiny and their active components are integrated based on CMOS technology for a monolithic solution. Although this integration scheme is not really applicable at higher power levels, the barriers and challenges do overlap to some extent so taking a look at these low-power converters is worthwhile.

For higher-power applications, many products described as being "integrated" are, more accurately, "co-packaged" components where each component (inductor, FETs, controller, drivers, etc.) is done using a different process and then they are combined together in an open-frame module, or encased in plastic, with a small overall size/footprint. We shall look at the four "technology groups" one by one:

1. monolithically-integrated active components,

2. monolithically-integrated active with integrated passives,
3. co-packaged active components
4. co-packaged active with integrated passives.

Figure 1.5 shows data points in graphical form representing integration efforts for 12V, and lower, buck converters. The black line denotes the division of power level between monolithic and co-packaged buck converters. The data points above the line are co-packaged solutions of either active, or active + passive integration schemes. The data points below the line demonstrate either monolithic active, or monolithic active + passive integration. Here we also see that power levels vary widely but that monolithic solutions are only available below a few Watts of power output.

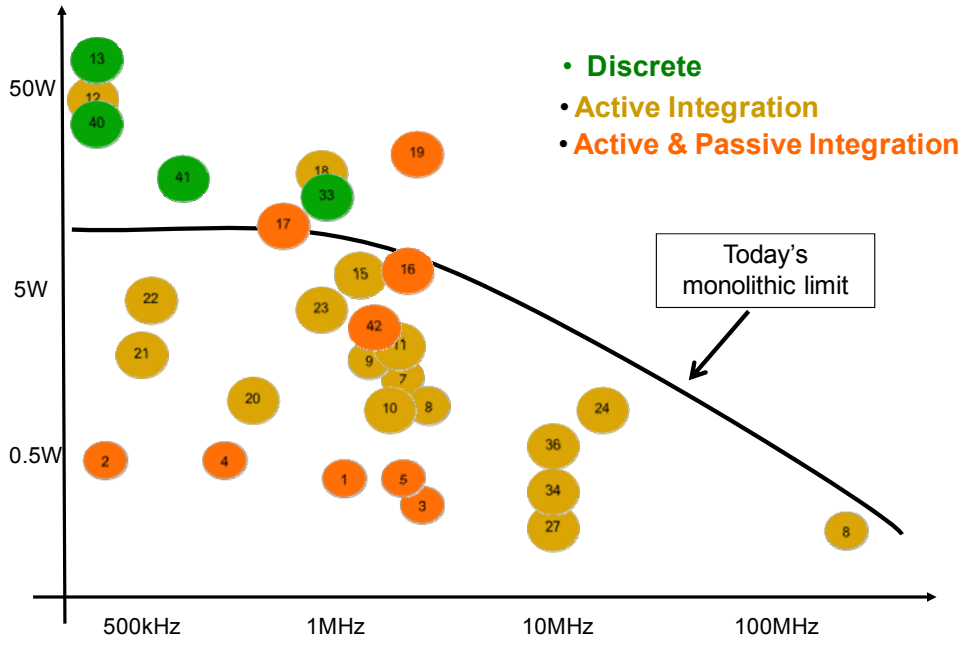

Figure 1.5. Output power versus frequency for low-voltage DC/DC buck converters

Above the 6 W power level and below 3 MHz, data backed up by market demands suggests that it is more practical and cost effective to use discrete components co-packaged into a module. This leads to a physically larger module, like the one shown in Figure 1.4, which relaxes thermal requirements compared with small integrated modules of a similar power level, as discussed later in this section. It is important to note that convection cooling is directly tied to the available surface area of the component, and the heat flux concentration of each component. This is partly so that thermal requirements can be more relaxed in this case since the package has greater surface area to dissipate heat at higher power levels. However, above 3MHz, interconnection parasitics become too much to handle and smaller monolithic components are the only clear solution beyond that, as of today. Therefore, the playing field is split into two distinct parts – the half where monolithic packaging is practical and the half where it is not.

The data shown in the graph is a mix of industry examples and research work currently conducted. This information will be discussed in this section and the next, alternating between research and industry as appropriate, based on the technological groups. Figure 1.6 shows the demarcation between industry and research converters based on output current and switching frequency.

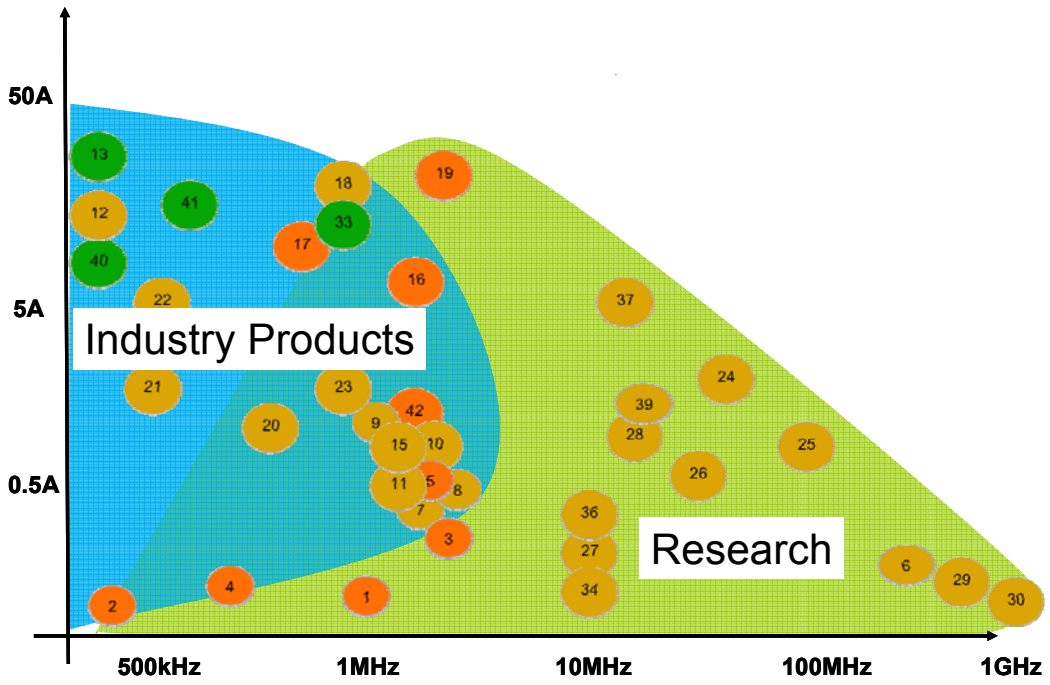

Figure 1.6. Output current versus frequency for industry products and academic research on DC/DC converters.

Monolithic integration is much closer to the ideal but then, integrating an inductor with the monolithic active stage presents some challenges. Today, there are many products that feature integrated CMOS active components with a separate conventional inductor. Output power levels are low due to the CMOS fabrication processes used. Figure 1.7 shows a few examples of monolithically integrated active stages that use an external commercial inductor.

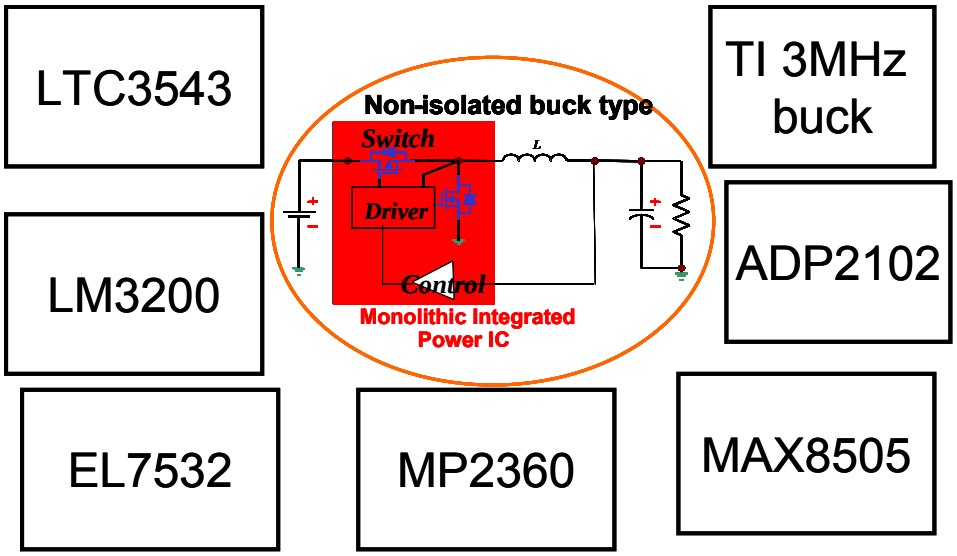

Figure 1.7. Converter modules featuring monolithic active integration.

The next step in the integration process is monolithically-integrated active with integrated passives. This work is mostly current research work, taking active integration that has been proven in industry, and integrating an inductor with it. The method of inductor integration is thin-film application of the ferrite on the silicon. The down-side is that thin-film has high losses and thus power output is lower than what is achievable with the previous case of only integrating the active components and using a normal commercial inductor.

A few examples of this have been published in research literature and are shown in Table 1.1. The reference numbers correspond to the data points shown in Figures 1.5 and 1.6. It presents a summary of key design parameters for some of these regulators published in literature since 1996. They have an input voltage of no more than 4V and output voltage between 1 and 3.3V. Note the low output current levels and wide range of switching frequencies.

Table 1.1

Research on CMOS bucks with thin-film inductors

	[1]	[2]	[3]	[4]	[5]	[6]
phases	1	1	1	1	1	4
Duty cycle	0.83	0.66	0.75	0.56	0.75	0.75
F (MHz)	1.6	0.5	3	0.75	1.8	233
L (uH)	3	10	1	15.2	1	0.002
C (uF)	?	47	1	21.6	?	0.003
Imax (A)	0.3	0.25	0.33	0.25	0.3	0.3

Industry products are just now showing up and they use Low-Temperature Co-fired Ceramic (LTCC) for their inductor design. There are really only two commercial examples to date: one from Fuji and one from National Semiconductor. The Fuji FB6861J has an embedded LTCC micro inductor with metallizations deposited on top for integrated circuit connections. The module is tiny at 2.4 x 2.4 x 1 mm and only 23 mg weight. Efficiency is quite high at a 90% with a 3.6V to 1.8V conversion, but only 200mA maximum output. It uses fixed on-time control and is packaged in the SON 8-pin form factor. A sibling is shown in Figure. Then there is the National Semiconductor module LM3218, operating at 2 MHz and has a 650mA output current capability (data point and reference #42). It too uses an LTCC inductor. We shall return to these examples later on.

Moving up in power level, co-packaged active components with external inductor have become very popular recently. Compared to monolithic devices, the larger die areas of co-packaged devices allow for higher output currents. Generally, the FETs and driver are combined together for optimization of the driver/device system since low parasitics are easily achievable this way. These modules are generally known as DrMOS. In some cases, the controller is also integrated for a more complete solution. Only large-value capacitors and the inductor are necessary externally to make the solution work. Examples are available from IR's IPOWIR, ON Semi., Power One's Maxyz X300 series and Renesas DrMOS.

The DrMOS from Renesas is quite a feat of design. The DrMOS is represented by data point 18, which has the highest current *and* highest frequency among the industry products shown. They are able to achieve 30 A continuous by packaging a highly-optimized driver with the two switches in a small package. Internal interconnections are still using wirebonds but future generations will go with planar interconnections as demonstrated by CPES work in recent years [19, 43, 121-124]. Additionally, they are actively pursuing co-packaging a decoupling capacitor with the module to add one more step of integration for enhanced performance, once again as demonstrated by CPES work [19].

Some modules on the market state they have an 'integrated' inductor whereby the inductor is inside a plastic encapsulant along with the rest of the circuit. This inductor may be of purely conventional design, i.e. a wire wound inductor in a ferrite core - and thus does not represent true physical integration. A closer approximation would be to use an alternate means of inductor fabrication where it is functioning as a

chassis for the remainder of the circuit. As a result, the word “integration” as used in this dissertation is essentially equivalent to “co-packaging” when speaking of modules capable of output more than 1A of current.

One of the best examples found on the commercial market today is the Linear Technology LTM4600 series, which houses every part of the converter, from power stage with inductor, to control, to drivers. It is shown in Figure 1.8. A very-wide input range of 4.5V – 28V and up to 14A peak output current is possible from a module that is only 15 x 15 x 3 mm. Essentially the only external parts needed are input and output capacitors, along with an output voltage trim resistor. Numerous pads on the underside allow for heat to escape through the bottom side of the package into the motherboard.

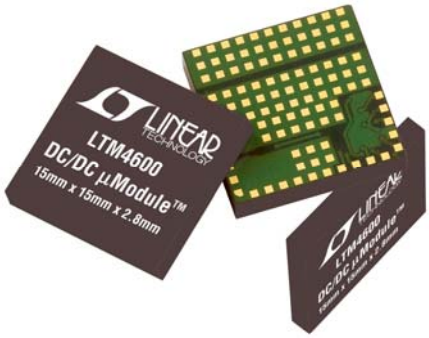

Figure 1.8. Linear Tech LTM4600 module with all components co-packaged together. Backside pad interconnections are also shown.
Used with permission from Linear Technology

Co-packaged designs with something other than a conventional wound inductor are not available so far in industry but research is being conducted in this area to boost output power and performance to greater levels. The concept for getting through the challenges this next level imposes is the basis of this dissertation, and will be introduced in the next section with details to follow in the remainder of the chapters.

1.3 Overcoming Integration Challenges

The goals of this work are to investigate, assess, design and implement a methodology for integration of functional parts of a buck converter into a small effective package that addresses electrical and thermal considerations concurrently. The push for integration is brought about by the desire for high power density and small footprint and is mainly achieved by raising the switching frequency to the extent that we can use the inductor as a substrate for the active stage and reduce the amount of necessary output capacitance. However, the higher frequency can increase switching losses dramatically and thus requires careful circuit layout to reduce circuit parasitics. For this example, I will introduce a buck POL with a voltage conversion of 12V and 5V to 1.2V and 1.3 MHz switching frequency that will demonstrate these size-reduction techniques as a system design method called “Stacked Power.”

Raising the switching frequency to reduce size of the module can impose its own problems as circuit and component parasitics become more and more critical as frequencies increase [102]. This will be discussed in detail in Section 4.2. So far, we have made use of integration to reduce certain layout parasitics including: device power loop inductance, snubber capacitor placement and input decoupling loop inductance. Electrical (Saber), thermal (I-DEAS) and electromagnetic (Maxwell 3D) finite element (FEA) models were made for rapid analysis and to find improved layouts quickly and efficiently.

Another barrier to integration is thermal management. Many of today’s high-power point of load converters are thermally limited. They require being derated at

higher ambient temperatures to prevent overheating. As the packages get smaller, less and less surface area is present for cooling which causes temperatures to rise rapidly. Add to that the increased switching losses from higher switching frequencies and an increasing ambient temperature in server farms, and you have a big thermal problem to deal with. Figure 1.9 shows this trend and how it affects output power versus switching frequency. Chapter 5 describes these effects in detail.

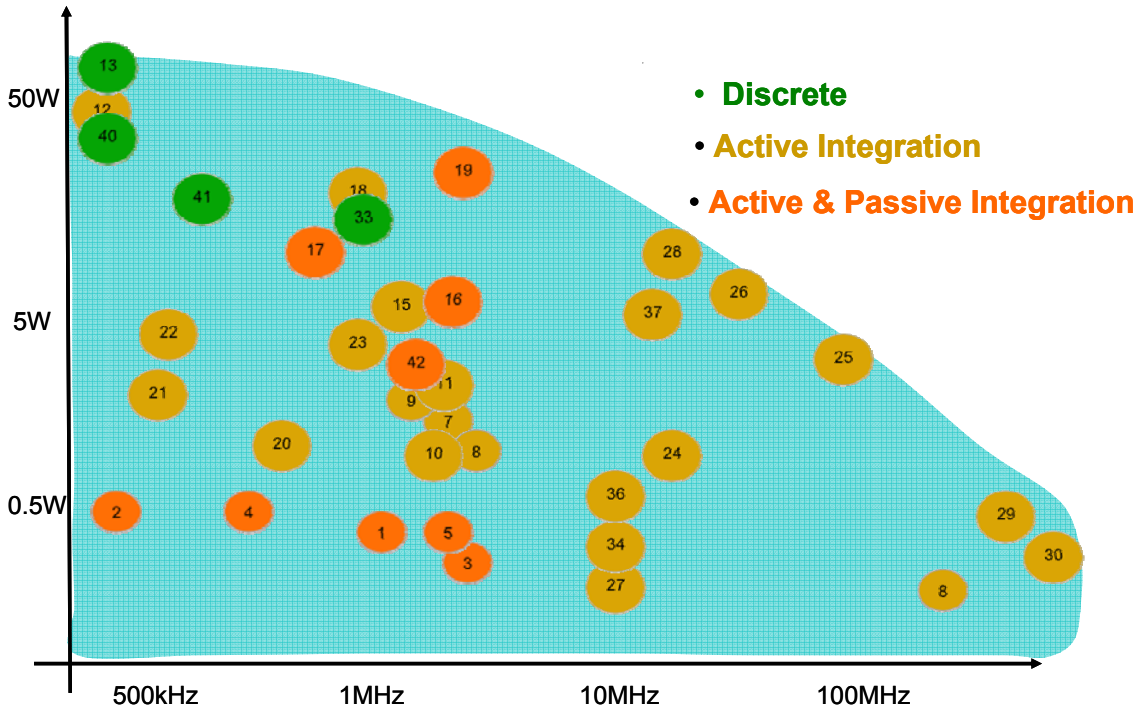

Figure 1.9. Data map showing thermal limit of low-voltage DC/DC converters

A look at the data trend makes it clear that thermal issues are roadblocks to high-power integration. The downward slope of the curve indicates that new technologies have higher losses, which limits their output power capability. In order to overcome the obstacle, the materials used to make the POL need to become multi-functional in order

to get maximum packaging thermal efficiency, which in turn allows best use of the external thermal management available.

The substrate material plays a key role in determining how much integration can be performed because it represents the largest surface for convection cooling. If the package uses a poor thermal conductor for its substrate, its already-limited surface area will cripple the performance of the circuit. There is a strong need to utilize all the available package surfaces for maximum cooling in order to handle the increased switching loss that many highly-compact converters face. This is applicable even when a heat sink is to be used. If the heat loss cannot be removed effectively, the package cannot be made any smaller.

Many of the packages shown in the map require the use of additional cooling mechanisms beyond the package's own convection capability, particularly as the size decreases. Any package is subject to thermal limitations which are imposed by the power loss of the components. The extent is always dependant on the airflow rate present and the surface area of the component. Component datasheets show derating curves which represent the maximum amount of output current (or power) that the package can handle in a given environment (consisting of both ambient temperature and airflow rate).

The ones that do not represent integration are the largest in size, like the discrete examples given in section 1.1. Others vary depending on power level. Some of them need a finned heat sink mounted on top, others have copper heat bars that solder to the board, and some have thermal pads – the latter two rely on their motherboard to

remove a large part of the heat through a very-low thermal resistance path so the package can be made quite small.

Some of them are not an open-frame design but are instead encapsulated in black plastic for aesthetic reasons, as well as for confidentiality. Unfortunately, this plastic comes with the price of inherently poor thermal characteristics which will degrade the effectiveness of a heat sink and so the heat sink size will have to be oversized accordingly. Also, this decrease in effectiveness will cause the design to be predominantly cooled via the thermal pad connection on the bottom side.

However, cooling through the motherboard can yield mixed results. There are many applications where the motherboard is already at thermal capacity and should have temperature limits imposed on it of around 90°C [113]. Intel has issued bulletins for desktop fan boards to ensure that motherboard temperatures do not exceed 46°C without a large fan running at full speed to cool it [114]. Section 5.3 shows an analytical case study on the impact of surface mount converters on the motherboard temperature and ambient environment.

Therefore, to use the motherboard as a heat sink can present serious thermal risks left for the system design engineer (i.e., end user) to solve. It is better if the converter can handle as much of its own heat as possible by convecting it off its own surfaces rather than counting on conduction cooling into the motherboard to be the dominant mechanism. As we will see in Chapter 5, convection and conduction cooling mechanisms work in tandem, and convection should always be maximized before conduction because ultimately, the heat expelled must reach ambient temperature no matter what.

Excellent thermal design and low-parasitic packaging with high efficiency can extend the feasibility of a monolithic solution, at least in spirit, into a co-packaged solution capable of much higher output power. Cost-effective monolithic buck solutions remain low in power output, either because of inductor design limitations (data points #1-5) or because of power loss and subsequent thermal issues in a small package (data points #7-10). The former group uses thin-film inductors that are made from a variety of materials and are integrated in a number of ways as described in the inductor process and materials sections of this report. The latter are faced with switching losses from the tiny active components, which inherently have limited current handling capability and low breakdown voltage.

With this in mind, we now have the context to introduce our target zone data point #19 – above and beyond the performance of the nearest examples just described. This is graphically shown in Figure 1.10. We have developed a new system design method for addressing all these issues simultaneously. After extensive research into thermal limitations of VRMs and other converters, it was determined that the main barrier to integration and size reduction is actually the PCB itself. The FR4 epoxy used to make the majority of PCBs today is simply a poor thermal conductor (on the order of 4 W/m°K). It worked just fine for decades but we are now reaching extremely high power density levels that are hindered by the PCBs' low heat conduction. As a result, hotspots form where the devices are connected to the board which greatly reduces the potential output current level in order to keep the converter thermally stable.

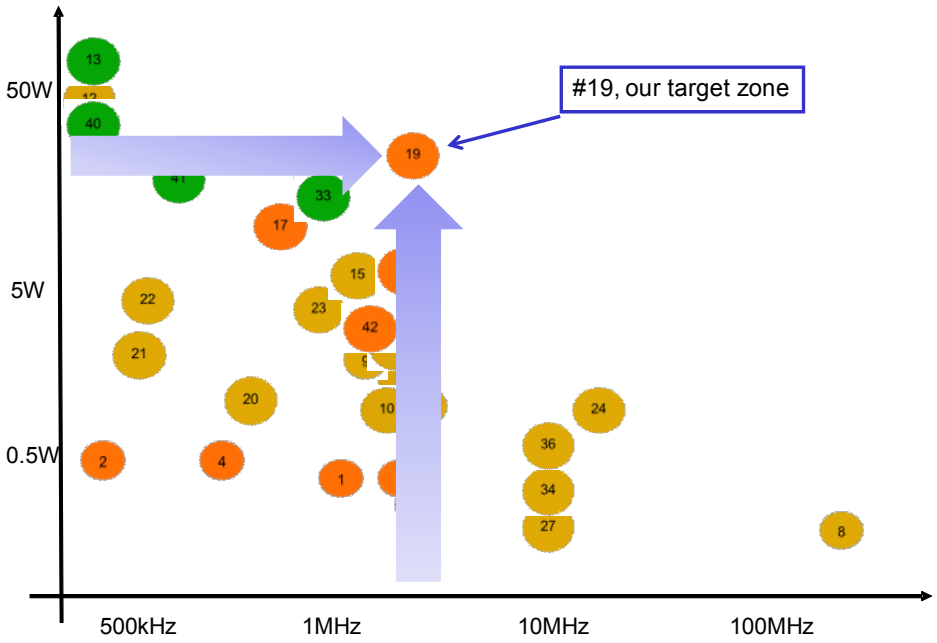

Figure 1.10. Buck converter map showing the high-level goals of Stacked Power

What has typically been done in these sorts of situations is to either add a heat sink to the hotspot, add thicker copper traces or vias, and/or add a high conductivity substrate such as Direct Bonded Copper (DBC) or Inter-Metallic Substrate (IMS) to the PCB. What we are proposing is to simply *replace* the PCB with the conductive material. In this manner, nothing will have to be added to the system – the thermal management is integrated into the converter. This will allow for a significant improvement in functional efficiency since not only will the substrate support traces but it will also allow for significantly greater heat distribution. This work began at CPES by Yu Meng and Dr. Zhenxian Liang with Embedded Power [86]. A modified approach adapted to solderable device dies is described in this dissertation, called “Stacked Power,” and is shown in Figure 1.11 along with its circuit schematic in Figure 1.12.

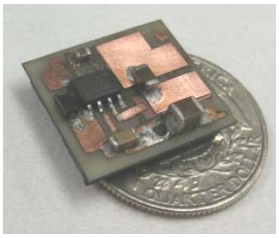

Figure 1.11. Photo of actual CPES Stacked Power converter on a 25¢ coin.

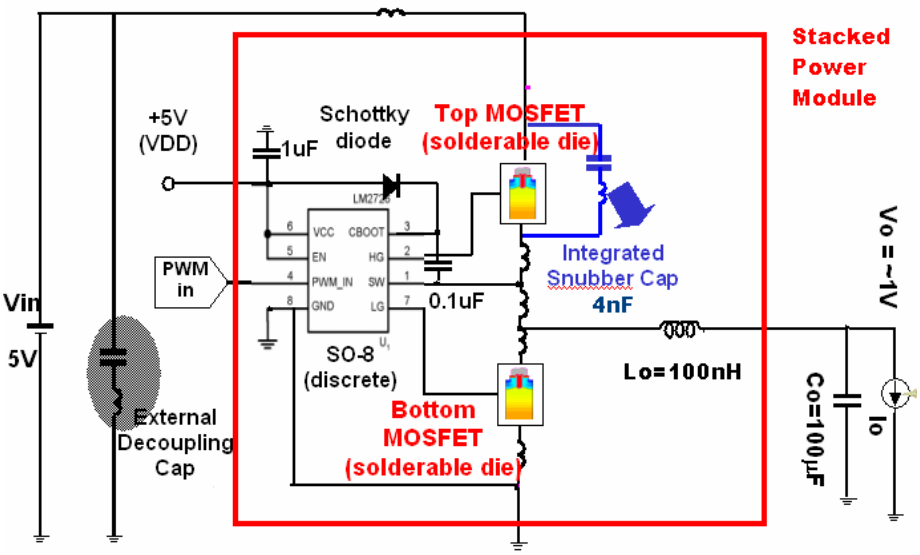

Figure 1.12. Circuit schematic of Stacked Power POL

The uniform heat spreading will increase convection efficiency of the substrate to the ambient since over 50% of its surface is not covered by components thanks to the dies being embedded inside the substrate. This will reduce peak temperatures by up to 40% (discussed in Chapters 3 and 4). The alternative to reducing peak temperatures would be running the converter at peak temperature but with correspondingly higher output currents than a PCB version. The extra loss can also be partitioned such that high switching frequencies can be achieved, in addition to high current. Also, the rapid heat removal from hot spots allows for smaller packaging since extra surface area is no

longer necessary, which shortens signal paths, which reduces parasitics, and finally yields less loss at high frequencies. See Sections 4.2, 4.4 and 5.2 for detailed information on all these aspects.

Not only that, but Stacked Power is so called because of its structure. In order to have multiple active layers stacked on top of each other for both low parasitics and small package size, the switches are ***embedded*** inside a layer of ceramic. The typical method for using DBC is to have it as a heat spreading substrate and the circuits sit on top of it. In this case, we are laser cutting holes inside the ceramic carrier and use epoxy to keep them in place. This allows us to layer more components on top and bottom, which is something that can't be done with conventional PCB substrates. Figure 1.13 shows a schematic of the structure and Figure 1.14 shows the typical use of PCB with its obvious drawback.

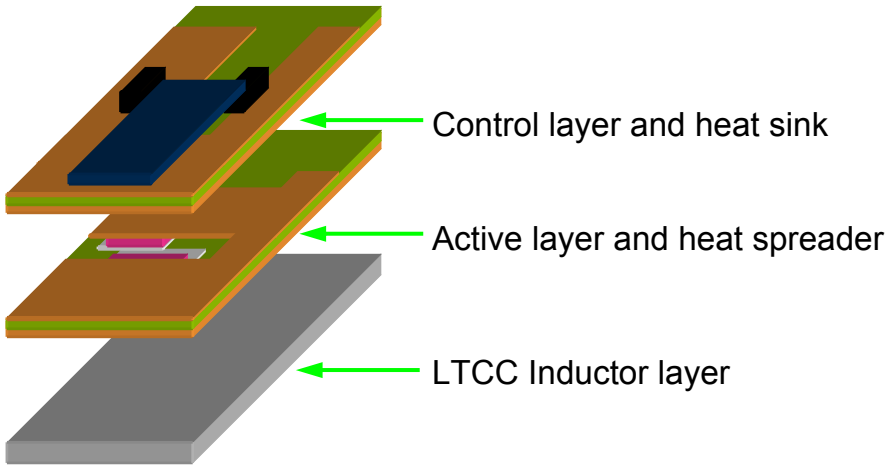

Figure 1.13. Schematic of Stacked Power layers.

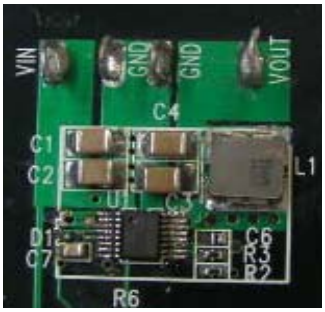

Figure 1.14. Typical use of PCB substrate with all the components mounted on top, preventing layering of substrates.

This design method allows us to take active and passive components and co-package them in a new method conducive to higher power. Similar work has been done by using planar interconnections to achieve modules capable of handling 600V [121]. In addition, further work demonstrated capability at the 1kW power levels [122] using similar planar co-packaging. Also flip-chip technology on a flexible substrate has been demonstrated for size reduction of module and yet be able to handle high power levels as well [125]. But these high power applications weren't seeking to ultimate thermal performance as much as high voltage capability in a small package size.

For POL converters below 50W discussed in this work, the layering of active components on top of the inductor is in the same vein as the Fuji and National Semiconductor POLs shown above – but the great heat spreading capability of the substrate and larger dies allow for much higher output power for broader applications (~20A output compared with ~650mA output).

In addition, Stacked Power differs from other high-power DC/DC modules in the way interconnections are made. There are two main methods for making interconnections: the classic method is using wire bonds and then the planar

metallizations. CPES has been a strong proponent for planar metallizations for some time, arguing that the reduced parasitic effects are well worth the manufacturing change from wirebonds [19, 43, 86]. Figure 1.15 shows an example of how the planar interconnections not only reduce inherent trace parasitic inductance and resistance but also allow for smaller packaging with the potential of layering, bringing with it additional parasitic reduction [121]. Planar interconnects also offer improved thermal performance compared with traditional wirebonds [122-124].

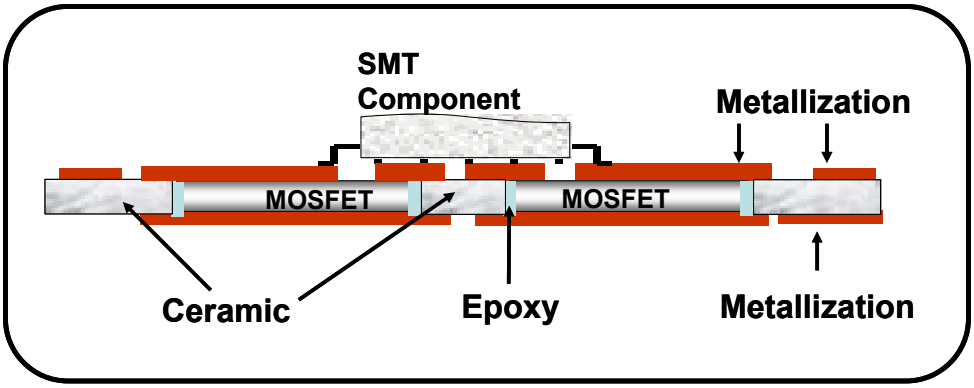

Figure 1.15. Metallized interconnections to replace wire-bonding techniques [121]

GE's Power Overlay concept (new version of "Chip On Flex") is very similar to CPES' Embedded Power where interlayers made of polyimide or polymeric adhesive are used to separate metal and substrate because the traces are deposited [131]. Power Overlay makes the power modules using a highly complex procedure that involves a number of interlayers and laser cutting steps and *then*, places the module onto a substrate for interconnection.

Stacked Power, on the other hand, is different in that it does not require all these extra interlayers because the traces are already solidly attached to the ceramic substrate and interconnections are made with copper straps. There is also no need for

milling the ceramic, which is a very tedious and time consuming process. Taking advantage of vertical MOSFETs of today, all that is need with Stacked Power is a hole cut all the way through, rather than a cavity, and the dies are epoxied in place (discussed in Chapter 3). Also, finding more reliable methods of making AlN consistent in thermal conductivity have been developed in recent years [116]. This has made this material available whereas previously to obtain this level of thermal conductivity, environmentally-unfriendly beryllium oxide had to be used. So advancements in manufacturing and control of aluminum nitride oxides is what has made this material available for high density integration such as Stacked Power.

Philips has demonstrated the integration of passive components into the printed circuit board (PCB) as embedded passive integrated circuits ("emPIC" as they call it) [134]. The goals are to obtain higher power density and to develop it for highly-automated manufacturing typically found with PCBs, by using layered construction. The magnetic components are thus designed to be very thin and are realized using Maglam magnetic material for the core. This material is a ferrite polymer compound that can be adapted for compatibility with the PCB laminating process. Figures 1.16 shows what this structure looks like and Figure 1.17 shows the inductor layer's structure.

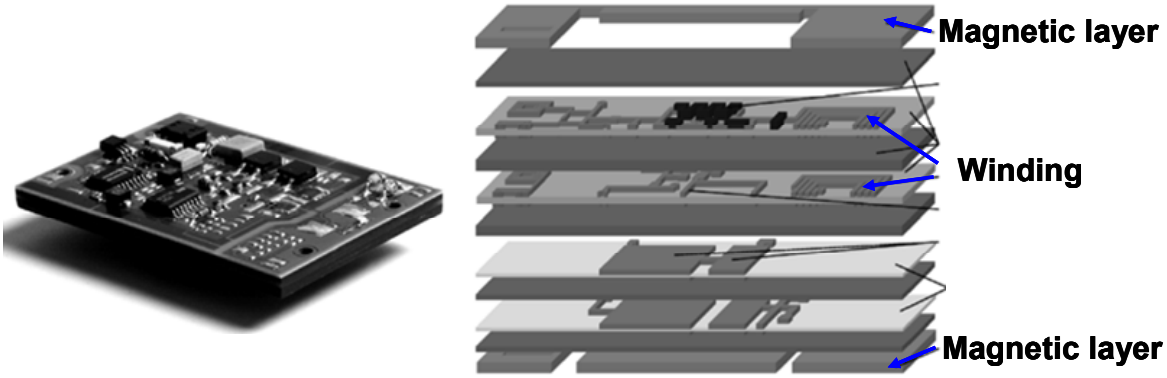

Figure 1.16. Philips integrated PCB structure [134].

A 60W converter was designed and fabricated using this technology [134]. Its transformer is entirely integrated in the PCB as well as 11 capacitors. Standard PCB lamination processes are used for the layering of integrated components. The circuit needs an area of 55 x 85 mm with a PCB thickness of 4 mm. Up to 82% efficiency was demonstrated, which indicates a significant amount of heat (13W) that must be dissipated, and as will be shown in this work, the low thermal conductivity of FR4 is a limitation if the losses are in inner layers. Dispensing with FR4 is therefore an obvious benefit. No thermal performance data was shown for this module.

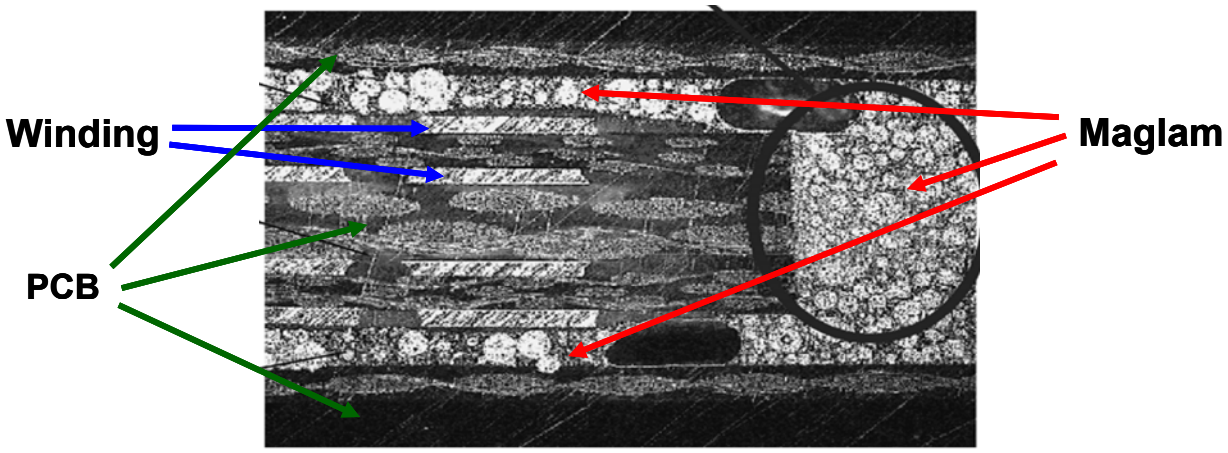

Figure 1.17. Inductor structure used in Philips design, surrounded by PCB material [134].

Stacked Power builds on this layering concept but uses different materials to achieve the goal. Since the dies are co-packaged with a ceramic substrate using no polyimide coatings or FR4 interlayers, it has the benefit of reducing the manufacturing steps and allowing for hybrid hermetic packaging, such as that used in aerospace modules. This is an advantage because there are no off-gassing materials used in the process. This is essential to fundamentally meet stringent requirements such as MIL-PRF-38534 Class H and K standards, with proper design. In this case, the use of conformal coating and gap filler materials is rendered moot.

Many examples in industry have yet to follow these design concepts. Most of them still rely on wire-bond technology such as shown in the following Figure 1.18. This concept is quite old and has thus been fully tested so its reliability metrics are fully known – which is a real advantage over an emerging packaging method such as Stacked Power. However it took some time to obtain a reliable method of attachment for wire bonds because the thin wires are highly susceptible to thermal expansion which

causes high stresses at the constraint points. The pads see this cyclical stress and eventually there is pad "lift off," causing a failure. This issue has been overcome and work at CPES has demonstrated in several different designs that this is not as much as issue with planar interconnects because heat extraction and transfer is so much faster [123-124]. As a result, some effort is currently developing in industry to move towards planar interconnections to take advantage of these inherent benefits [117].

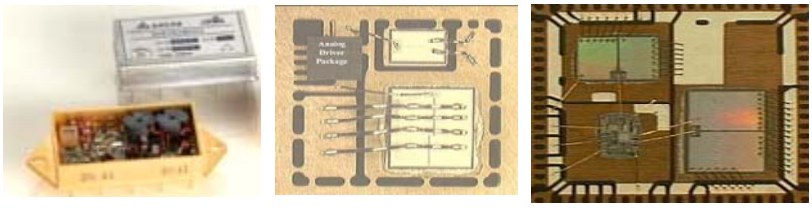

Figure 1.18. Three typical examples of recent commercial products that use wirebond interconnections.

The closest industry example so far to what Stacked Power is a module made by Fuji Electric. Their latest is called the μDD FB6861J and is a no-lead package of 2.4 x 2.4 x 1 mm buck converter with integrated LTCC inductor. Only capacitors need to be externally added for operation. Traces are deposited on the surface of the inductor so that the active IC layer can interface on top. The inductor is made with windings wrapped around an LTCC core.

However, as remarkable as this module is, it is only able to output 300mA due to low efficiency, particularly when the duty cycle drops below 28%. The small size of the module and the fact that the materials all have poor thermal conductivity will reduce the maximum possible capability of this circuit to applications of very low power – mainly RF applications rather than Power Electronics.

All these issues to improve system performance will be addressed in this dissertation. We will start off with system-level analysis, followed by component-level analysis and finally finish off with environment-level analysis since no converter operates in a vacuum. In order to consider all these issues methodically, one must take a step back and look at the problems in full context. Only then will solutions emerge that will take advantage of all the possibilities available to the converter is in its intended operating environment. Now, we will consider one of the highest-performance POLs on the market: the VRM. The VRM is important because its stringent system specifications offer critical insight into the challenges involved in very-high power applications.

1.4 VRM Integration

Voltage Regulator Modules (VRM), are high current converters and their design to date is discrete due to this high power. Their distinction in the POL category is due to transient performance being of paramount priority because of the demands that modern microprocessors require for ultra-high processing speeds. As a result, control of the output impedance of the VRM is critical and the output currents need to be very high (today's levels are around 130A output!). Many other POLs do not necessarily have these issues and instead target small size. This desire to reduce size is what leads us to increasing levels of component and system integration.

The VRM can be physically in the motherboard (known as VRDown) or as a plug-in (known as VRM). In desktops today, VRDs are used fairly exclusively due to their low cost since they can be made by motherboard manufacturers. Figure 1.19 shows the large amount of board area occupied by today's VRDs and there is strong desire to reduce this space requirement to make computers that are more compact. In workstations and servers, VRMs are more common due to their better cooling capability for these higher power systems and the fact that the budget is a little less low-cost oriented. Understanding their thermal barriers will lead to ways of dealing with heat removal barriers that have recently become a great concern.

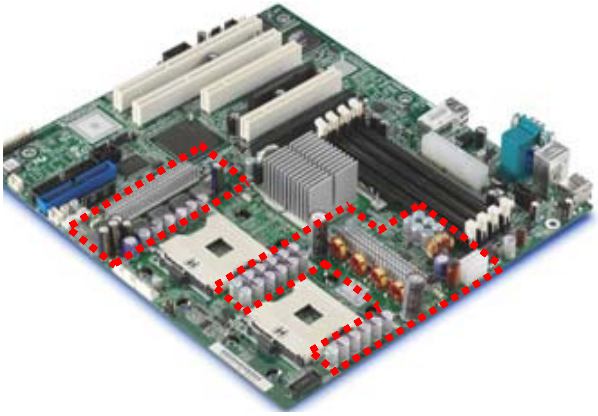

Figure 1.19. Dashed lines show the large real estate usage by the VRD.

With high frequency operation of the latest VRM technology at or near the MHz range, the reduction of parasitics becomes a key electrical challenge. The main parasitics I speak of are the result of trace inductances and resistances so a trace-minimization scheme is always in order for a good design. However, this generally conflicts with thermal reasoning that dictates that heat sources should be spaced out to reduce crowding of hotspots. This tug-of-war between electrical and thermal

requirements needs to be reconciled in some way to obtain a reliable, effective and hopefully optimal design. How to do this effectively is the question at hand.

As an example of the possibilities afforded by integration, we take a look at the difficulties of supplying power to a computer's microprocessor. The trace parasitic inductances between the VRM and the CPU are a significant source of performance loss. The inductance of the traces decreases transient response so it is desirable to have the VRM as close to the CPU as possible. A possible sequence for integrating the CPU with its power supply is shown in Figure 1.20 as three successive stages. The top graphic is what we currently have today, followed by two more steps, each showing an increased level of integration, ending with the CPU and VRM integrated together for optimal performance since all parasitics are removed from the power delivery path.

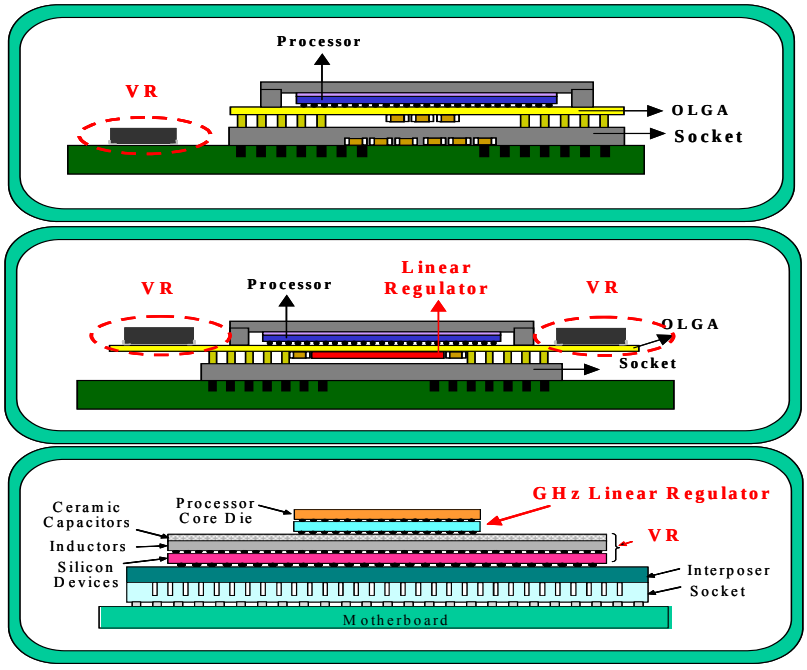

Figure 1.20. Step-by-step approach for system integration.

The baseline case on which the 3 stages are based is similar to what was outlined previously where a VRM is plugged into the motherboard and the power delivery is routed in the motherboard to the underside of the CPU as shown in Figure 1.21. The first step is shown as the top portion of Figure 1.20 and in Figure 1.21 and is what is used today. In this case, the VR is a module that remains plugged into the motherboard. The VR is characterized, along with the system parasitics, and determination of how far the VRM can be pushed has been assessed.

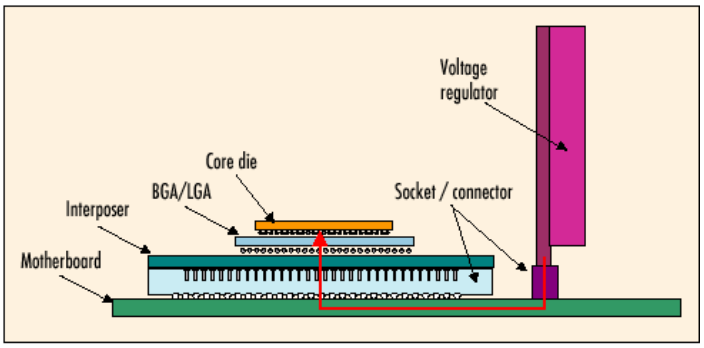

Figure 1.21. Today's power-management structure with long power delivery path

For step 2, the VR would move to the OLGA layer so that the path parasitics can be minimized to aid in achieving the next generation Pentium chip's high di/dt slew rates of up to 200 A/ns and yet maintain tight regulation on the 0.8 V output voltages. This portion can be achieved by means of a two-stage converter design using a switched-capacitor first stage to step the voltage down from 12V to a bus of 6V and then a second stage would bring the voltage down to the necessary 1V range. This will be looked at in detail in subsequent chapters.

Step 3 would be the ultimate goal in which the CPU would be packaged with the VR all in one so that the path lengths will be as short as possible. This would prove to be a huge thermal problem and so novel cooling techniques would have to be looked

into. Some of these techniques have been shown to be feasibly manufacturable on their own but their implementation in integrated electronics still has many unresolved issues that will require future work.

1.5 Materials for Converter Design

With miniaturization and improvements in device performance, the development of power electronics systems has progressed to a state where active devices' impact on the system's size and cost has been taken over by passive components. Magnetic components have been one of the bulkiest components which take up significant circuit surface real estate. With the increase in popularity in low profile electronics, development in planar magnetics has taken up an increasingly important role in today's power electronics research and development. In view of the numerous techniques available for inductor integration, the integration techniques have been classified as below:

1. On-silicon (chip level integration)
2. Organic substrate (package level integration)
3. Thick film core and winding (package level integration)
4. Integratable winding with commercial core (semi-discrete level)

For chip level integration, the inductor / transformer is built on the silicon die. The processes are compatible with silicon processing. For package level integration, the magnetic component is built on a substrate where the active device can be mounted on. Package level integration can be classified into two categories. One is integrating the inductor / transformer in organic substrate where the processing temperature is low (<300 °C), another is integrating the component using thick film technology where the processing temperature for the component is high (> 600 °C). The mechanical properties of the organic substrate are usually not compatible with the silicon die, i.e. FR4 has a coefficient of thermal expansion (CTE) of 17 ppm / °C. For thick film technology, the materials involved can support higher operating temperature and has a CTE between 4 to 7 ppm / °C.

Another material for inductor manufacturing is Low Temperature Co-fired Ceramic (LTCC) technology. The material comes in the form of ferrite sheets or, more often, as a paste that gets screen-printed. The ferrite sheet type is more conducive to high current capability and so it he one discussed here. The sheets can be cut easily with scissors to the desired shape. The layers are between 60 and 100 microns thick and are available in several relative permeabilities ranging from 50 to 500. Multiple layers are put together and then pressed to achieve the desired thickness. Once the pressed components is thick enough, it is sintered at 900°C according to a complex firing profile that changes for each type of ferrite tape. When it comes out of the oven, the tapes are fused together and are now hard ferrite. Shrinkage rates are around 30% as the binder burns out of the tape, rendering the structure hard and brittle.

A technique more suitable for making inductors which can support higher current is developed using ferrite sheets [44,45], following the initial work reported in [46]. This technique also proves it can produce components with lower losses compared to using the screen printing technique [47]. Due to the non-linear behavior of the inductor, improvement in light load efficiency can be achieved [48], and the extent of improvement can be controlled by varying the geometry of the inductor. So not only is high light-load efficiency achievable with this material, but also low profile inductors for applications with little real estate.

Integrating inductors on silicon has been discussed through the use of thin film materials [49]. In this case, inductances of 1uH at 5MHz has been achieved, with DC resistances as low as 150mΩ and Q factors up to 18. Coupled-inductor technology has been able to push load currents to 10A. However the main limitation to the success of this approach is the high DC resistance of the materials. This represents a very large I^2R loss that eventually makes the feasibility of thin-film inductors questionable in high current applications. Figure 1.22 shows the inductance gap between conventional wire-wound inductors and those found in literature. Achieving the levels of inductance demonstrated by commercial wire-wound inductors using thin-film technology would lead to much higher resistances. As you can see in the graph, the inductance is low and the resistance is high for thin-film inductors compared to typical wire-wound technology.

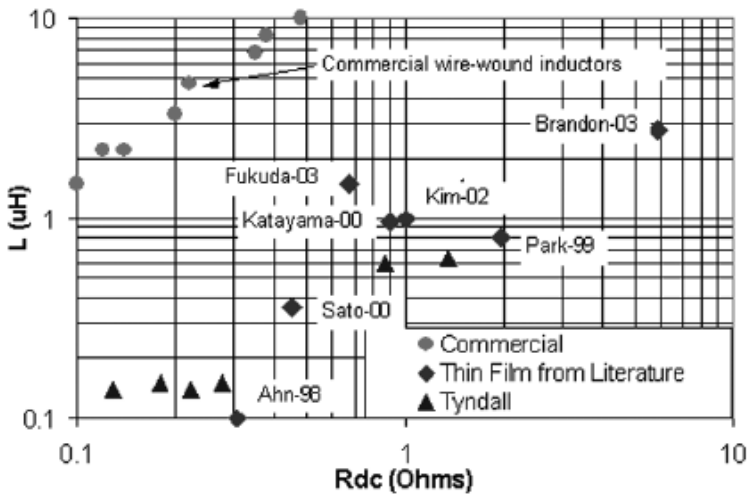

Figure 1.22. Thin film inductor resistances compared to commercial products [49].
© [2005] IEEE

As far as hardware is concerned, there is a lot of recent interest since thermal problems are becoming all the more prevalent. Thermoelectric materials are an interesting possibility although they have not really been associated with a microelectronic cooling application thus far. Recent advances in new material development to increase the efficiency of thermoelectrics are covered in [50]. Derivations from the base material of bismuth telluride are used to make them work. With a small voltage applied to them, they are able to transport heat away from a hot side to a cold side by means of high thermal conductivity and high electrical resistivity. However, their effectiveness has never been shown to be excellent and their cost is prohibitive for low-power DC/DC applications despite the continued interest.

Microchannels are also candidates, where there is the room to make convection cooling feasible, and have even been demonstrated in water-cooled versions for electronic applications with very high heat removal rates [51]. However their tiny dimensions have high pumping losses so their efficiency is not that great. Miniature

heat pipes, in comparison, with a capillary wick have shown great efficiency [52]. Micro heat pipes are even smaller and thus have to be wickless to simplify manufacturing. They operate just like an air conditioning system except that they are entirely passive. Their fabrication has been demonstrated by means of vapor deposition processes [53] and micro machining [54] and is shown in Figure 1.23.

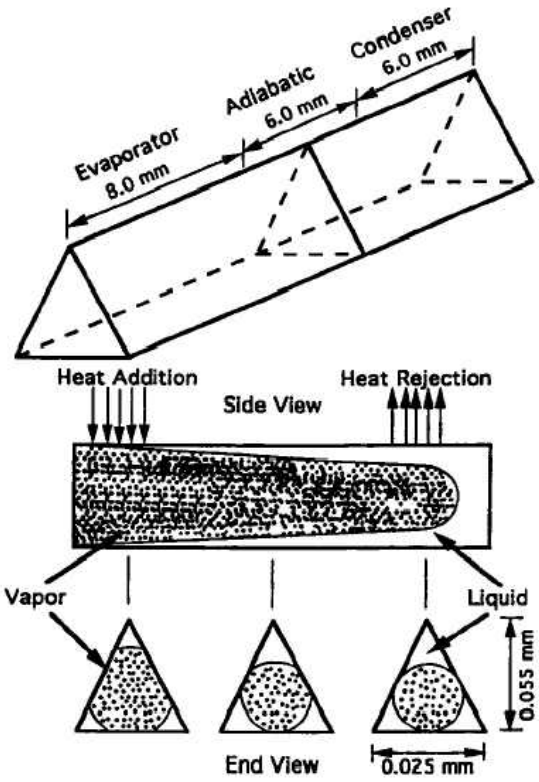

Figure 1.23. Example of micro heat pipe design for converter cooling [53].

Heat pipes in conjunction with thermal vias/sinks could be an excellent possibility for DC/DC converters of low power [55] and provide functional, as well as structural, integration [56]. Integrating micro-heat pipes into the Step 3 CPU/VRM module could alleviate the thermal mass issues [57,58]. However, heat pipe cost and long term

reliability remain big hurdles to their use in low-cost applications though they have proved effective at cooling systems, such as laptop computers. These are some promising cooling technologies that have yet to demonstrate their application in a fully integrated electronics system such as Stage 3 and so the need for fundamental research in their implementation is critical [59].

But if the available surface area for cooling is sufficiently well exploited, there may not be any need for these more exotic, expensive and complex structures. For example, previous work found in literature has shown that DBC is highly effective at heat removal over standard PCB [60-63] and can effectively be integrated with additional cooling mechanisms such as heat pipes. Aluminum nitride DBC has been combined with encapsulation techniques for water cooling of high-power devices [64] due to its exceptional thermal conductivity, compared to the more conventional alumina DBC (150 W/m°K and 36 W/m°K respectively).

In order to satisfy requirements of simplicity (for cost reduction), ease of integration, mechanical compatibility with other materials, and flexibility, the work discussed in this dissertation selected aluminum nitride DBC as the cooling method of choice. In terms of low-power systems, chips are typically mounted on a patterned copper layer of DBC such as for Multi-Chip Modules (MCM) [65], [66]. However, in all cases, the devices are mounted ***on top*** of DBC and the reverse side is left for attaching a heat sink, the ceramic providing electrical isolation. In this work, we show a method where the devices are ***inside*** the DBC in order to eliminate the need for a heat sink, integrate thermal management, reduce profile, simplify system design and increase power density of a low-power compact POL, all at one time. This makes the DBC a true

multifunctional material to achieve our goals for high-levels of integration at high power output.

1.6 Thermal Modeling

Thermal research for electrical circuits has a well-established foundation in literature due to today's high-power-density initiatives. Of the many papers covering this huge topic, several are notable here due to direct reference to components that could make up future integrated power supply (IPS) modules and/or describe processes necessary for good thermal design leading to an IPS.

Thermal design in conjunction with electrical design is highly necessary at this point in time, as discussed, and so a good thermal model to complement the electrical one is in order. Fraunhofer Institute's Thermal System Modeling paper [67] is a reference paper for general overview and covers modeling methods for representing thermal systems with Cauer RC networks. The basic lumped model approach is covered in [b]. Dynamic behavior [68] also needs to be covered, especially in component-level analysis – a necessary first step to a good system model. The preferred method for this analysis is Finite Element Analysis (FEA) due to the complex nature and large quantity of variables as described in [69]. Specific to DC-DC converters, [70, 71] cover the details involved in the FEA process.

Packaging is a big issue with parasitics with thermal issues largely being the cause [72, 73]. Thermal modeling of magnetic planar components will be required soon in VRs and is well covered in [74-76]. Determining heat transfer coefficients and spreading resistances is needed to determine what is a good design and is described in

[77, 78]. Thermal constraints have to be balanced with EMI requirements and their coupled transient solution can be very involved but is a necessary part of the design process [79, 80]. Combining all of these different analyses is done by means of FEA modeling [81] so that an excellent overall design can be achieved and goals met [59]. These are just a few of the many papers relevant to this subject but they collectively represent the current state-of-the-art methods needed for today's electro-thermal modeling.

1.7 Dissertation Synopsis

In order to assess potential levels of integration, it is necessary to be clear as to the power level and the role the particular power supply will have. Even within POL converters, the environment and power levels can change substantially. Some POLs are in the 25W-100W range (like VRMs) while others are in the 3W range. The degree of integration for each of these power levels can be considerably different, although a lot of the barriers, benefits and challenges are common to both and will be discussed in this dissertation.

In recent literature, a lot of attention has been placed on the low-power level due to its high integration potential. This low-power range can have switching frequencies in the tens or hundreds of MHz which allows the use of CMOS circuits for on-silicon-chip active integration. Input voltage levels are typically in the range of 1.2 to 5V and output voltages from 0.9 to 3.3V, but in all cases, a high duty cycle is very desirable to keep losses, and thus temperatures, under control. Switching frequencies vary substantially,

from 500kHz to 233MHz so far, using CMOS technology from 0.25um to 90nm. However, output currents are consistently less than 0.5A and it will not be long before much higher currents will be desirable.

The next higher-power level converters are above the capability of CMOS circuits, and thus switching frequencies are reduced to compensate for less integration. The power levels here are greater than 15W and often greater than 25W. Integration for this range is not covered nearly as extensively in current literature but its necessity cannot be overlooked. As a result, the work described in this report is based on this higher-power range.

The dissertation overview is as follows:

- The PCB is the standard substrate material for converters today
- A generalized system-level converter Finite Element Analysis modeling methodology is developed based on a VRM that uses PCB (Chapter 2)
- The model is used to expose the thermal limitations of PCB (Chapter 2)
- A system design methods named “Stacked Power” demonstrates how using DBC instead of PCB can relax these thermal limitations (Chapter 3)
- Mechanical and electrical techniques for taking advantage of the improved thermal performance demonstrate:

1. The extent to which more-even distribution of heat across the substrate surface allows components to be physically closer together with lower operating temperatures (Chapter 3)

2. How embedding devices inside the DBC substrate takes better advantage of today's vertical MOSFET structures (Chapter 3)
3. How bullets 1 and 2, when combined together, allow for a reduction in converter footprint compared to similar conventional POLs (Chapter 4)
4. How the improved thermal performance of Stacked Power allows the end user more flexibility in integrating the POL into his main system architecture compared with other design methods (Chapter 5)
5. How Stacked Power can be extended to a 2-stage coupled-inductor buck design for state-of-the-art integration at the 45W power level (Chapter 6)

Chapter 2: PCB Modeling and Characterization

2.1 Electrical Aspects

To start off the process of investigation for improving electrical and thermal package performance, a CPES-designed state-of-the-art VRM was chosen as the initial circuit design. This circuit was designed by fellow CPES student Jinghai Zhou, et al. The hardware is shown in Figure 2.1. The topology is a 1 MHz non-isolated self-driven, 12V to 1.3V, DC-DC converter, that has a full bridge front-end with a buck following a 3:1 transformer [82]. The transformer allows the circuit to have larger duty ratios than would be typical for this large conversion. It is packaged in a 1U form factor, has 2 pairs of interleaved phases for ripple cancellation, and is capable of 100A output. A schematic of the converter is shown in Figure 2.2.

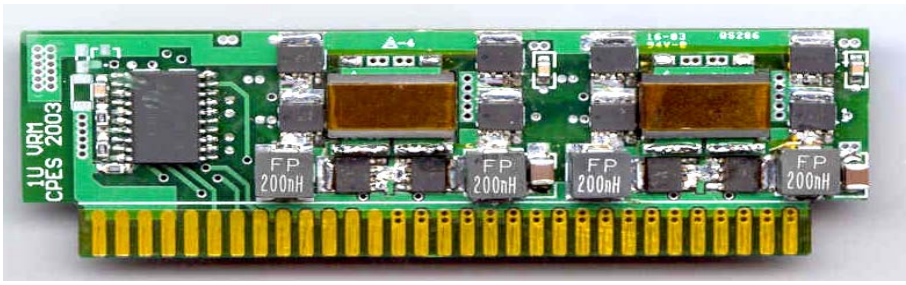

Figure 2.1 CPES-designed 1 MHz self-driven VRM in 1U form factor [82]

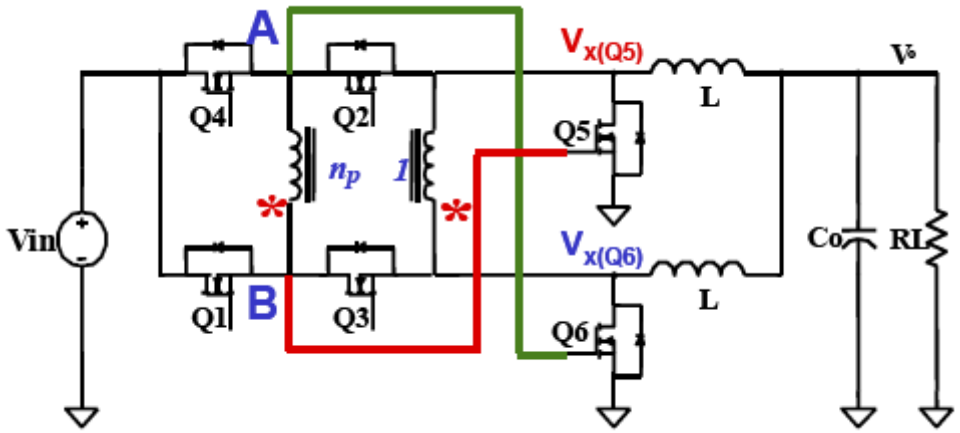

Figure 2.2. Schematic of the 2 phases of the self-driven VRM.

Using a phase-shifted buck front-end allows the use of the transformer concept to this non-isolated application. As such, the extreme duty cycle (of around 0.1 today) is extended and some of the benefits gained are that both the switching loss and body diode reverse recovery loss are dramatically reduced. The additional devices necessary to make this work would normally complicate the driving scheme significantly. However, the secondary voltage waveforms are just right for self-driving the synchronous rectifiers which eliminates this potential issue, while at the same time reducing loss. This leads to high efficiency compared with a typical 4-phase-interleaved buck, as shown in Figure 2.3. Reducing heat loss by means of high electrical efficiency is the first step towards making a thermally viable solution.

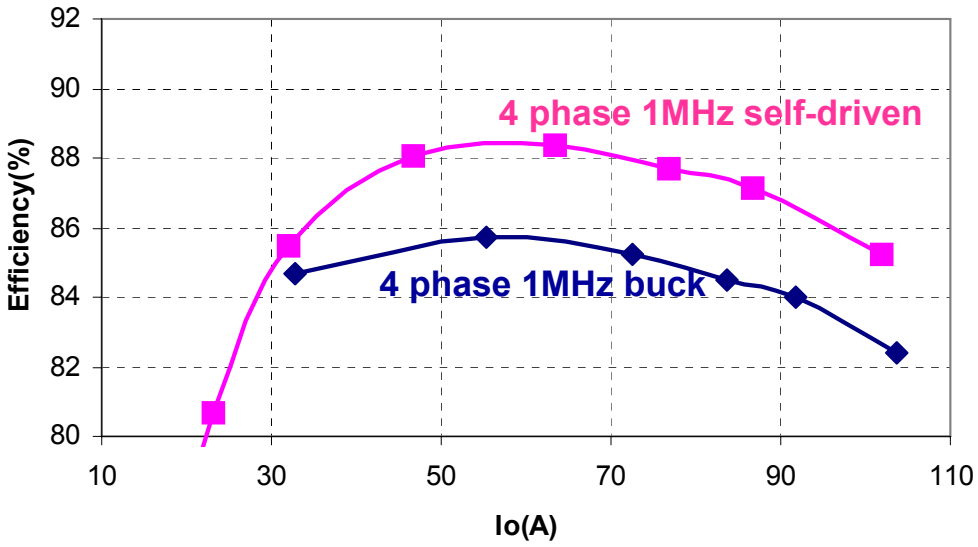

Figure 2.3. Efficiency advantage of the CPES self-driven VRM

The electrical model loss breakdown of the VRM is shown in Table 2.1. As can be seen from the efficiency curve at full load, the major loss contributors shown in Table 2.1 collectively represent 96% (21W) of the total losses of the VRM. The remaining losses (0.8W) stem from gate drivers, controller, trace DCR, and inductor and core losses but are considered to have negligible impact on the PCB since they are distributed over large surfaces areas and stem from components that have a high component-to-board thermal resistance, which means that their heat will mostly be convected directly off their own surfaces rather than absorbed by the PCB.

For the complete 4-phase circuit, there are 8 synchronous rectifiers and 8 full-bridge switches consisting of 4 top and 4 bottom switches. Their power loss was obtained through electrical simulation and verified with hardware experiments. Simulations were done using SABER from Synopsys, with trace AC inductance and DC resistance parasitics obtained from Maxwell's Q3D parameter extraction tool.

The transformer has a 12-layer winding structure and is shown in Figure 2.4 [120]. Since the PCB was designed as a 6-layer board, the extra 6-layers were made from extra pieces of PCB that were soldered on the base PCB and will be shown in later figures. Figure 2.4 shows how the windings are interleaved between primary and secondary sides with two ounce copper in each of the inner layers, containing two sets of primary windings in parallel shown in blue (dark color, also shown in figure's text), and six sets of secondary windings in parallel shown in red (lighter color, also shown in figure's text). Without the transformers, the extended duty cycle and self-driven benefits would be lost and the circuit would become a buck converter with redundant top switches.

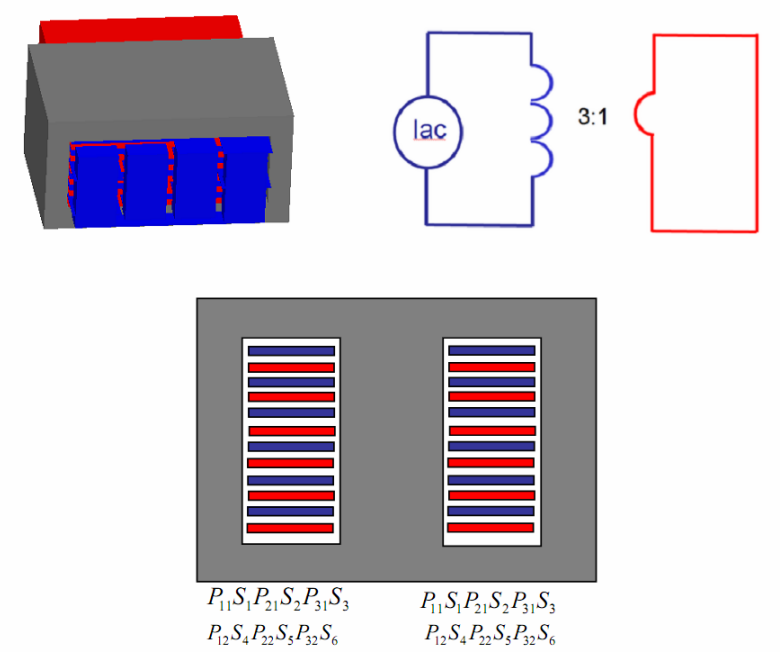

Figure 2.4. 12-layer transformer structure showing how the windings are interleaved [120].

The transformer winding resistance and its leakage inductance between the primary side and the secondary side are measured with an impedance analyzer, and

the values reflected to the primary side are 20mΩ and 60nH, respectively. The devices inductances are calculated in device-model subroutines obtained from Hitachi (models HAT2168 and HAT2165). The loss values obtained are used as inputs in the I-DEAS thermal simulation to see the heat distribution and characterize the location and degree of hotspots. The SABER device models are set to run at 100°C for electro-thermal compensation. The thermal simulations will be discussed in detail in the next section.

Table 2.1
Loss breakdown for each part of the VRM

Part	# of parts	Loss each (W)
Top Q	4	0.67
Bottom Q	4	0.54
SR	8	1.38
Windings	2 sets	1.26

The high switching frequency also allows us to achieve higher bandwidth, with careful design, in order to reduce the number of output capacitors required for a give transient specification. Through research and simulation done by Jinghai Zhou and Julu Sun [82], it has been demonstrated that at a bandwidth of 130 kHz is where OSCON and ceramic capacitor numbers become equal, and at a bandwidth of 390 kHz, no bulk capacitors are required to meet the VRM 11 transient specification! This is shown in Figure 2.5.

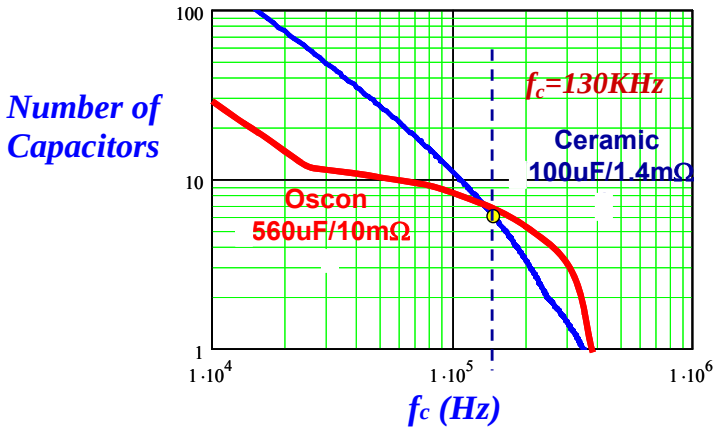

Figure 2.5. Wider bandwidth reduces number of bulk capacitors necessary, whether OSCON or ceramic. f_c is the bandwidth where they are equal.

Once this level of bandwidth is achieved, the power path between VR and microprocessor becomes highly simplified, to the point of having what is shown in Figure 2.6. However, achieving this level of bandwidth is difficult and as we know, half the switching frequency is the theoretical limit so 1/3 the switching frequency is about as much as can be done successfully today. This would mean a switching frequency of around 1.3 MHz is necessary to obtain the level of bandwidth for eliminating the bulk capacitors. For this reason, 1.3 MHz will be the switching frequency used for each module made in the course of this dissertation with the assumption that in closed-loop form, bandwidth would be 390 kHz to eliminate bulk capacitors.

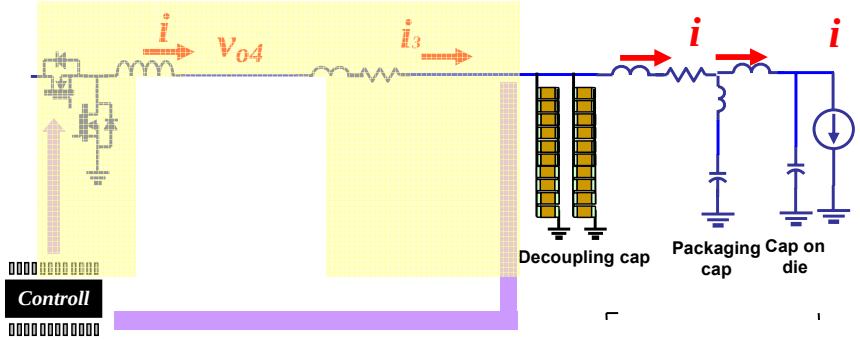

Figure 2.6. The power path once a bandwidth of 390 kHz has been achieved.

The entire circuit layout was purely based on electrical considerations. The goal was to get the highest electrical efficiency so the components were laid out accordingly, which certainly is one part of thermal design, but thermal layout considerations did not get factored in. For example, the synchronous rectifier pairs seen right below the transformers are at the output of the board – but they have high losses and are side by side, each with a second pair for the remaining two phases directly beneath them, on the opposite side of the board. This, as will be shown in the thermal modeling part of this work, is a big thermal limitation. It keeps the high current path short but also causes a heat loss concentration.

Also, the inductors flank either side of synchronous to, again, keep the high current path short, but they block critical airflow over the largest heat dissipation part of the VRM. And finally, thermal vias were not used due to the large areas the transformer windings took up inside the PCB, which obviously cannot be shorted by thermal vias.

The VRM schematic for half the circuit (2 phases) is shown in Figure 2.7. On this schematic is a bold/thick (purple) line shown over a trace loop. This loop, which goes through the synchronous rectifier devices and the transformer secondaries, must be kept as short as possible to maximize the decreasing di/dt when Q5 cuts off so that its body diode conduction loss is minimized. The efficiency of this converter is highly dependant on keeping this loop, in particular, as small as possible. Therefore any changes made in layout have this constraint, which poses the thermal problem of having to locate the synchronous rectifiers close to each other. More will be discussed about this in the next few sections.

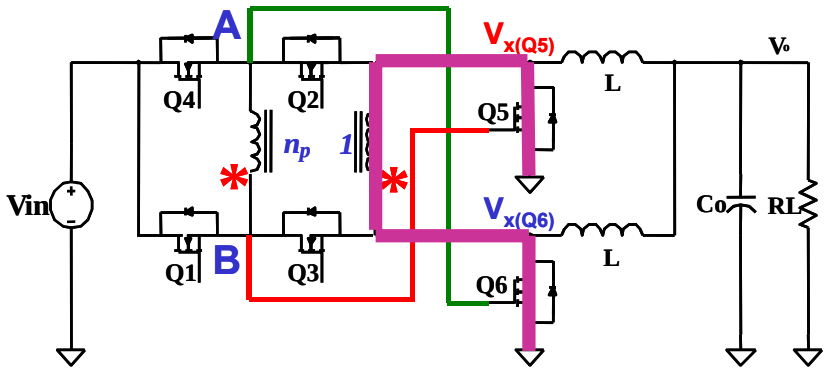

Figure 2.7. 2-phase schematic of Self-Driven VRM circuit

Part of this research involves determining what the best way is to model such a system. The term “best” is used in a sense that the model must be simple and yet yield accurate results which is more complex that it may first seem. One must realize that in order to model a VRM from a thermal perspective, the Printed Circuit Board (PCB) is a very critical part of the cooling system, acting as a heat spreader first, followed by acting as a heat sink. As a result, a very good understanding of how to go about modeling the PCB is critical to getting a reliable model with which to simulate future design changes.

2.2 PCB Modeling Methodology

We know that the PCB is made of complex copper trace paths on each layer. Modeling these in detail would be extremely time-consuming due to the large file size that would be generated and consequent lengthy simulation times. In order to accomplish this task more efficiently, a simplified way of representing the traces is in order. In a VRM, like any other modern switching converter, there are control trace

layers, power flow layers, and ground planes. These are the essential parts of a multi-layer PCB for converters and must be taken into account as such.

The control layers can have lower amounts of copper due to small trace widths and numerous gaps when compared to the power traces. Using the Protel PCB files as a guide, each layer is assessed for copper density. What was needed was the *ratio* of copper in the control layers to that of the power layers. Looking at the CPES VRM layers, I assessed the following:

- There are 6 layers, each 2 oz. (70 μm thick) copper
- The control layers are very packed with traces
- Each layer is used to its maximum for good power density
- Power and ground layers alternate
- The control layers are estimated to have about 80% the copper of the power traces

The goal here is to model the VRM introduced in section 2.1 and see how close it is to running at full 100A output without resorting to a bulky and expensive heat sink [92,93]. A close look must be taken at the PCB in order to model it. The quantity of copper is the main criterion, which is based what function the layers play. The partitioning of the power and control layers is not all that critical since each layer is used to its fullest. However, there is some difference and so it will be taken into account. There are two main parts to this PCB: the signal layers which basically represent the 4 top layers are not full copper planes (assumed 85% copper), followed by the ground plane, and then a high-current bottom layer. These last two layers are very nearly full planes (assumed 100% copper), and thus contain more copper than the others, so

there will be a thermal conductivity difference between the front 4 signal layers and the back two high-copper content layers. Figure 2.8 shows this split in the PCB.

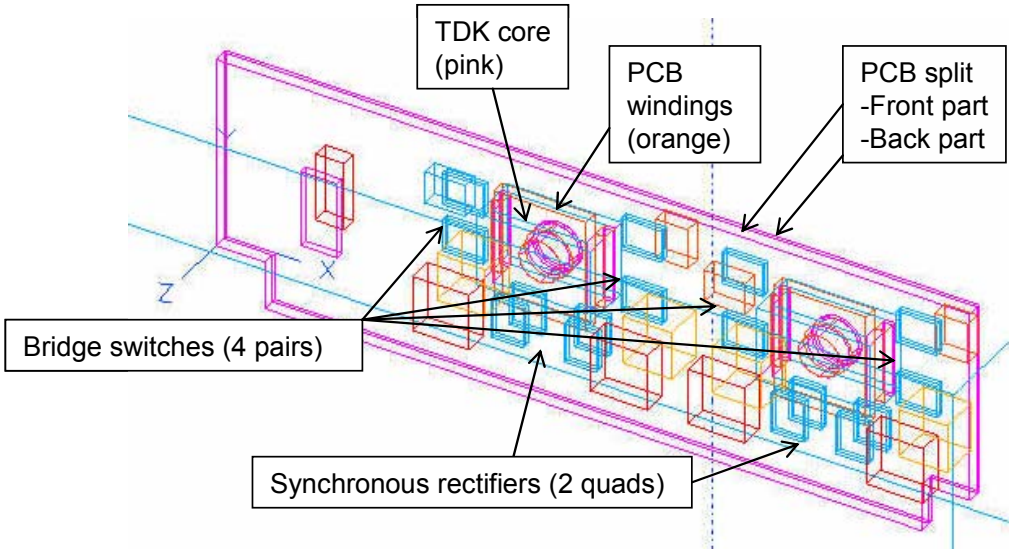

Figure 2.8. Wireframe drawing showing PCB split and heat sources

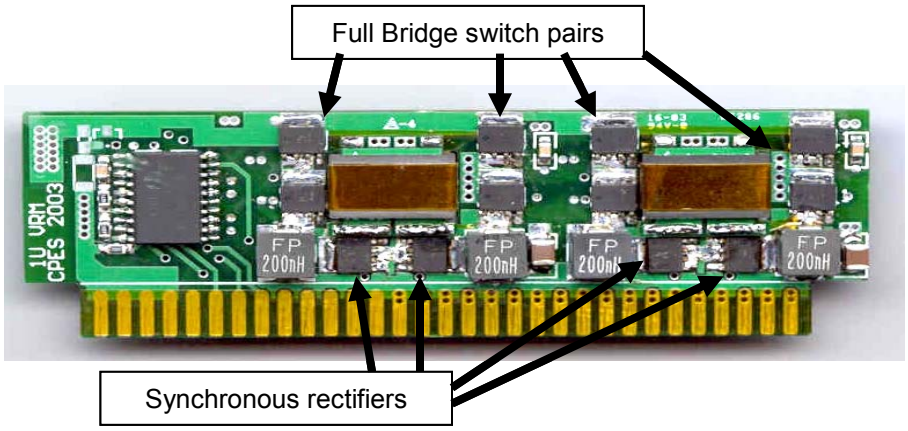

Figure 2.9. Actual hardware that corresponds to the model in Figure 2.8. In addition to the transformers and windings, the FETs are also heat sources and shown here. Losses were described in Table 2.1.

The next step is to obtain an isotropic (homogenous) thermal conductivity value to combine the two materials that make up the PCB, which will be the value assigned in the I-DEAS simulation software to the volume of the PCB. These *k* values are based on the material involved. In this case, there are two main values: The in-plane (parallel) k_p

and the through-plane normal (series) k_n which have to be combined in order to enter the value into a simulation tool such as I-DEAS. The k_n equation represents a weighted harmonic mean of the thicknesses associated with the conductivities, whereas the k_p equation is simply the weighted mean. There are three ways that these two values have been combined in the literature [83] with their respective tendencies in composite k value estimation:

$$\frac{k_n + k_p}{2} \Leftarrow \boxed{\text{Very high}} \qquad (1)$$

$$\sqrt{k_n \cdot k_p} \Leftarrow \boxed{\text{Medium}} \qquad (2)$$

$$\frac{2 \cdot (k_n \cdot k_p)}{k_n + k_p} \Leftarrow \boxed{\text{Very low}} \qquad (3)$$

The formulas used for the calculation of each k value are based on weighted mean and are as follows [83]:

$$k_n = \sum_{i=1}^{N} t_i \Big/ \sum_{i=1}^{N} t_i / k_i \qquad (4)$$

$$k_p = \sum_{i=1}^{N} k_i t_i \Big/ \sum_{i=1}^{N} t_i \qquad (5)$$

where t_i is the thickness of the layer and k_i its thermal conductivity for N number of layers. These calculations were carried out for the CPES VRM just described. The copper layers were assumed to be of uniform copper quantity since the thermal conductivity of copper is 760 times great than that of FR4. The values used for each layer are shown in Table 2.2.

Table 2.2
Parameters for thermal conductivity calculations

Layer	Thickness, ti, (m)	Conduc., ki, (W/mK)	ki*ti	ti/ki
Copper 1	7.E-05	380	2.66E-02	1.84E-07
FR4	3.E-04	0.5	1.50E-04	6.00E-04
Copper 2	7.E-05	380	2.66E-02	1.84E-07
FR4	3.E-04	0.5	1.50E-04	6.00E-04
Copper 3	7.E-05	380	2.66E-02	1.84E-07
FR4	3.E-04	0.5	1.50E-04	6.00E-04
Copper 4	7.E-05	380	2.66E-02	1.84E-07
FR4	3.E-04	0.5	1.50E-04	6.00E-04
Copper 5	7.E-05	380	2.66E-02	1.84E-07
FR4	3.E-04	0.5	1.50E-04	6.00E-04
Copper 6	7.E-05	380	2.66E-02	1.84E-07

The results give (assuming each copper layer has equal copper density):

- Arithmetic mean: 42.1 W/m°K
- Geometric mean: 7.3 W/m°K
- Harmonic mean: 1.3 W/m°K

With such a large variation in values, which is correct? According to [83,84], the geometric mean has been most often used but is not very accurate and [91] says that the typical value is on the order of 40 W/m°K. Therefore, it has been found that the geometric mean underestimates the actual conductivity by at least two to four times [83]. As a matter of fact, the actual through-plane thermal conductivity values for various sample board configurations were significantly closer to the arithmetic mean overall. This was predicted by analytical methods and verified by means of flash measurements [83]. The end result is that using the typical geometric mean as the guideline will lead to thermal system over-design with modern PCBs which are made using software that automatically maximizes the amount of copper on each layer.

So we have a general idea that the our VRM board's conductivity is closer to the arithmetic mean than the geometric mean, which would put us in the thirty- something range. How critical is this value anyway? How much does it change the system temperatures? In order to find answers to these questions, we need to model the PCB in 3D Computational Fluid Dynamic (CFD) software like I-DEAS and do a sensitivity study.

There are generally three approaches used to model a PCB [84]. These are outlined in Figure 2.10. The first strategy is actually not a model in the strict sense of the word because all the traces are drawn just as they are in the real circuit. This requires large amounts of time due to the complexity of the trace patterns and the subsequent large file size makes for very lengthy simulation times and increased chances of convergence failures. The second strategy is commonly used in system-level board modeling by representing the entire board as a composite isotropic *k* value as outlined earlier. This however can lead to error due to over-simplification of the system. The third strategy uses some thermal theory to determine that the surface layers are important to the cooling process and so only they are drawn and the remaining (internal) trace layers use a composite value as in strategy #2.

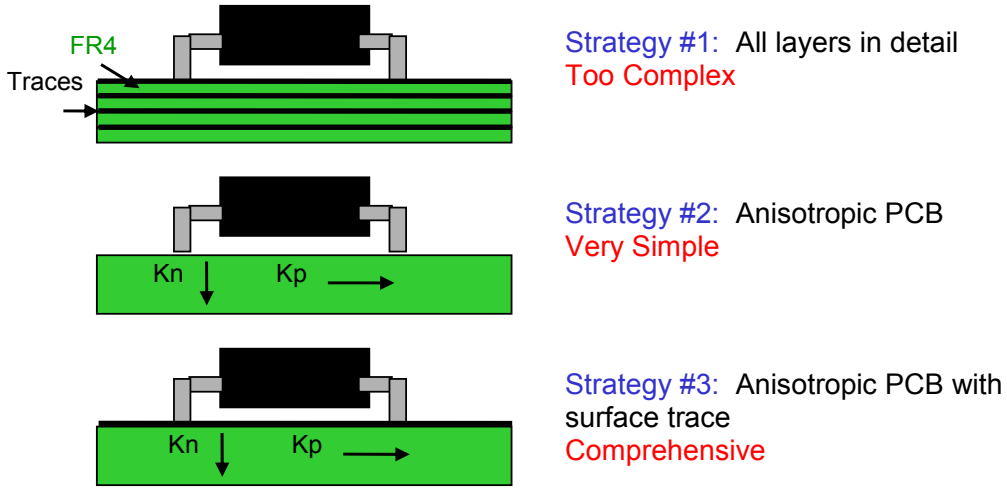

Figure 2.10. PCB modeling strategies.

Two models were made using I-DEAS FEA simulation tool. The first uses strategy #2 from Figure 2.10 and the second uses strategy #3. In strategy #3, the topside MOSFETS are the blue ones (see also Figure 2.9) and are all connected on a common trace as shown in Figure 2.11. This green trace doesn't exactly represent the electrical one of course but from a thermal point-of-view, it is a valid assumption due to heat conduction into the copper that is directly connected to the package pins and pad. Same goes for the synchronous rectifiers shown in red at the bottom of the board. These red MOSFETs are in parallel with four more on the backside of the board, directly in the same alignment so that they "stack." As we will see in the next section, Strategy #3 much more closely approximated the actual hardware results.

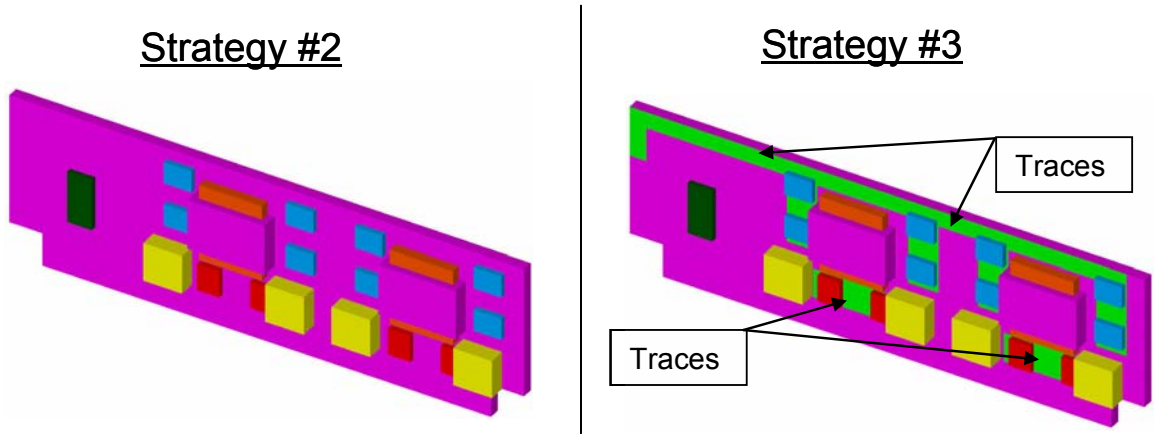

Figure 2.11 The two model strategies. #3 has surface traces added.

For FEA simulations, there has to be a well-defined boundary limit for the airflow so that CFD calculations can be done. This is accomplished by means of a control volume "airbox." The airbox used in this case is a rectangle 155 x 70 x 70 mm with the VRM located in the middle bottom half. The air mesh size (using ESC_Air material property in the I-DEAS library) was kept at a maximum of 0.75 mm; small enough to maintain accuracy but large enough that the model size did not hamper simulation convergence or run out of virtual memory.

The thermal model (includes the VRM and all the components on it) has 37,000 nodes and 190,000 thermal elements and the flow model (the air mesh) has 82,000 nodes and 420,000 fluid elements. The fluid element count was kept constant in different simulations so that a viable temperature comparison can be made between them. Figure 2.12 shows the model in the airbox and the direction of air flow. 105 CFM (5e7 mm3/s or 984 LFM *at zero pressure*) was nominal because this was the air flow rate of the actual cooling fan available on the test bench and later subsequently used for experimental verification.

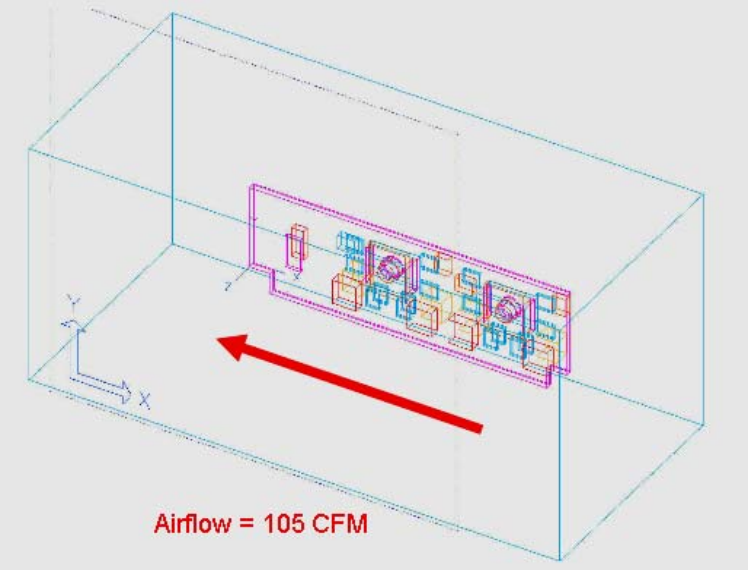

Figure 2.12. Complete model enclosed in airbox control volume.

I-DEAS software was used to try model strategies #2 and #3 to quantify the differences between these two approaches [94,95]. Experimental verification in the next section shows that model strategy #3 is very accurate considering all the simplifying assumptions made, which demonstrates that the assumptions are valid ones. I-DEAS Model of strategy #3 is shown in Figure 2.13. Strategy #2 over-estimated the peak temperatures by 10.5%, which proves that surfaces traces play a large role in PCB cooling.

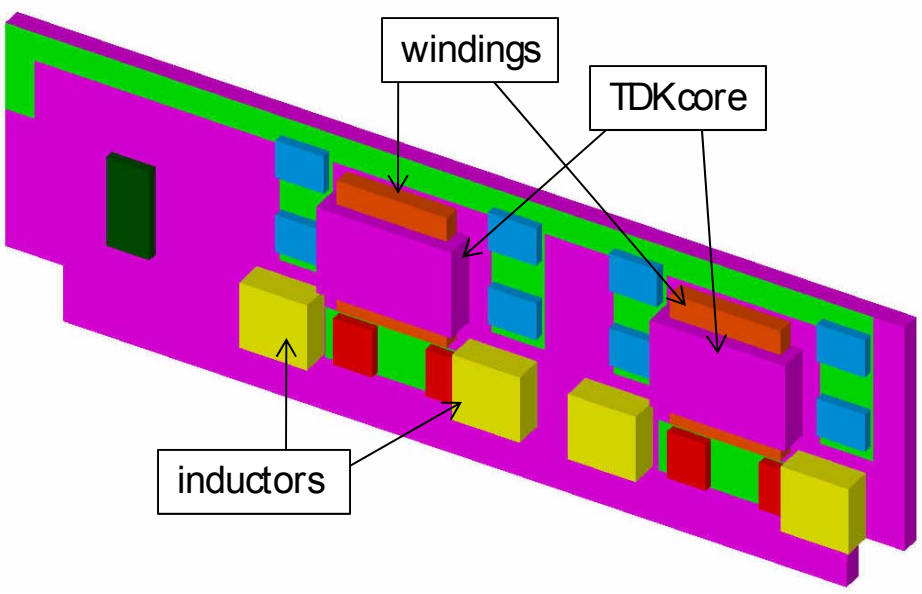

Figure 2.13. Various parts of the VRM as modeled in I-DEAS

The isotropic thermal conductivity values for the internal layers of the PCB were analytically calculated using the values in Table 2.2 to be 32 W/m°K for the front part/section and 36 W/m°K for the back part/section. The sections are shown in Figure 2.8. These represent the nominal values for this PCB with 85% copper coverage for the front 3 inner layers and 100% for the back ground plane. The surface traces are modeled and so are not a part of this calculation. The need for two values is due to the fact the back plane had a full copper ground plane which has more copper content than the signal traces of the front 4 layers.

However, to determine the impact of different conductivities, a series of simulations were run with different values of copper coverage per layer. Table 2.3 shows these results and also shows that the conductivity values only need to be somewhat accurate since the temperature spread is not very significant – particularly when the thermal conductivity values are over 30 W/m°K. The 50% coverage point is

equivalent to 1 oz. (35μm) copper thickness so we see that a “doubling” of copper thickness yields only a 5% drop in peak temperatures. This is due to the fact that the FR4 ultimately limits the maximum values that can be attained.

Table 2.3
Parametric analysis of PCB isotropic conductivity values on temperature

Cu coverage Front part	Cu coverage Back part	K front In W/m°K	K back In W/m°K	T maximum In °C	T average In °C
30%	40%	12	16	87.6	72.3
55%	65%	22	26	79.9	69.1
85%	100%	32	36	76.7	67.6
100%	115%	36	41	75.6	67.0

The LFPAK MOSFETs were modeled in a similar fashion as the PCB wherein they are divided into three parts, each with their own thermal conductivity [84]. In this case however, no isotropic composite value needs to be calculated since each part is made up of a single homogenous material. Figure 2.14 shows the three materials in the actual simulation wireframe part:

1. The top half is made of black plastic encapsulant
2. The bottom half is dominated by the copper pad of the LFPAK
3. The middle layer is the lead frame

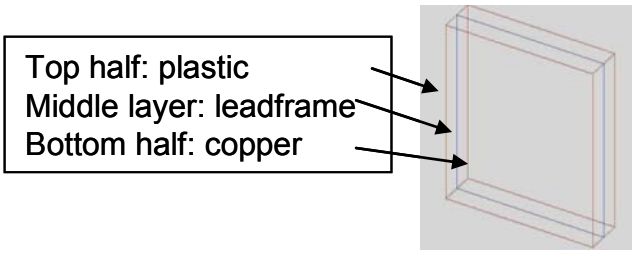

Figure 2.14. Device model in 3 simple parts. Heat load assigned to bottom half of devices as a volumetric loss.

The devices used are Hitachi HAT2165 and HAT2168 and so all the dimensions for the three parts came from those datasheets. This simplified approach (Figure 2.13) to modeling transistors is valid for system-level analysis but would need to be more detailed for accurate component-level results. Following are all the boundary condition details necessary to run an FEA simulation:

- MOSFET is 4 mm by 5 mm with 0.2 mm thick lead frame layer in the middle
- The thermal coupling resistance of LFPAK to board is 0.35 °C/W
- The heat load was 0.67W for top HAT2168, 0.54W for bottom HAT2168, and 1.38W for HAT2165 (each in all cases)
- and 263 W/m°C for the lead frame layer
- Inductor thermal coupling resistance (θ) is 4.9 °C/W (heat is absorbed from PCB) and analytically calculated using Equation 9 and values of Table 2.4.
- Transformer winding block coupling resistance is 1 °C/W

If the thermal coupling resistance for a component is not given in the datasheet, it can be analytically calculated based on the conducted temperature drop through a material interface. Heat conduction (q_k) is proportional to the temperature gradient (dT/dx) times the area of contact (A), through which it is transferred,

$$q_k \propto -A\frac{dT}{dx} \tag{6}$$

where T is the local temperature and x is the distance in the direction of heat flow. The minus sign indicates that heat gradient always goes from a high temperature to a lower temperature. The rate at which the heat flows is dependant on the conductivity of the material and is generally denoted as k,

$$q_k = -kA\frac{dT}{dx} \tag{7}$$

The SI units for q are in Watts and for thermal conductivity, Watts per meter per degree Kelvin. For a one-dimensional heat flow through a plane, the rate of heat conduction is simplified by solving (2) and making the conductivity independent of temperature,

$$q_k = \frac{\Delta T}{L/Ak} \tag{8}$$

Defining the denominator, we have thermal coupling resistance,

$$R_k = \frac{L}{Ak} \tag{9}$$

Table 2.4
Calculation values for the output inductors

L (m)	A (m2)	k (W/mK)	θ Resistance (°C/W)
0.001	4.5E-06	51	4.9

Calculation of the thermal resistance of LFPAK device to board of 0.35 °C/W was first done directly from datasheet information assuming that all the heat would flow through the drain since it represents 97.7% of the contact area. Any heat that flows through the pins can be considered added cushioning. Once again, there is no need for a composite isotropic value since the interface is all made up one a single uniform material: solder. Then, literature was consulted [86] and it was found that interface resistance using Sn-37Pb solder is actually up to 250% higher than calculated. The calculated value for the drain pad area, 0.1 mm thick solder layer, and solder thermal conductivity of 51 W/m°K yielded a low resistance of 0.14 °C/W. This multiplied by the 2.5x increase factor yields 0.35 °C/W, correlating with Figure 2.15 [99].

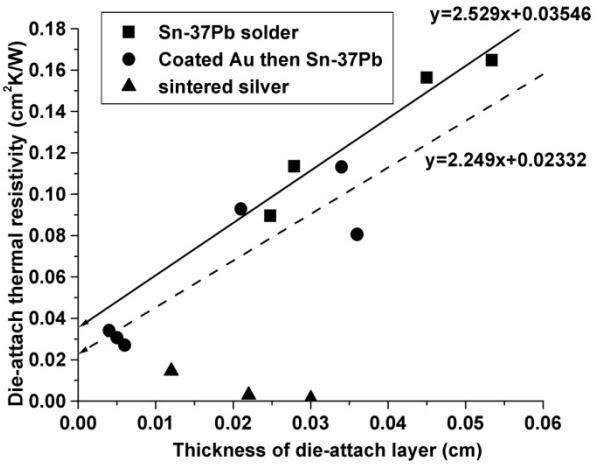

Figure 2.15. Properties of die attach materials [99]

Besides the transistors, the transformer winding losses were also included and all losses for each were shown in Table 2.1. The remaining parts were simply modeled as blocks with no heat dissipation, though the inductors had a coupling resistance of 5 °C/W because they absorb more heat from the PCB than they put into it. They have to be geometrically included in the model because they modify the airflow around the board by causing localized perturbations that result in turbulent airflow regimes at the board surface, and thus, lower airflow rates in certain spots. The inductors were subsequently removed by magnetic integration with the transformers, and the improvement is shown in the next section. Table 2.5 shows the material conductivity values that defined each material used in the I-DEAS FEA simulations [90,116]. This information is used by the software to calculate heat flow. These conductivity values are material properties and are not calculated using kn and kp as the PCB was because the material is homogenous (i.e., not made up of more than one material like PCB is).

Table 2.5
Material conductivity values used for I-DEAS FEA simulations [90,116].

Material	**Conductivity**
Leadframes	263 W/m°C
Silicon dies	125 W/m°C
Ferrite cores	4 W/m°C
Encapsulants	1 W/m°C

The following section will experimentally demonstrate the accuracy of this thermal modeling strategy. So the generalized system modeling methodology is as follows:

- LFPAK Devices can be modeled in 3 parts: plastic top, leadframe layer, copper base
- PCB must have surface traces modeled
- PCB inner layers have an isotropic thermal conductivity values where in-plane and through-plane conductivities are combined using arithmetic mean
- Solder thermal resistance couplings should be multiplied by 2.5 for improved accuracy

2.3 Model Results and Verification

A point to mention is the fact that many modern PCBs have a non-conductive dielectric conformal coating on the outer surfaces. This coating is to prevent moisture contamination, oxidation of the copper, and to provide some electrical isolation. This coating is generally either acrylic or polyurethane and according to IPC-CC-830 standards, the thickness is to fall in the range of 30 – 130 um (with 50um ±20um being typical from PCB makers like 4pcb.com) and the thermal conductivity is on the order of 2 W/m°K. The simulation results shown in Figure 2.6 assume a 50um solder mask of 2 W/m°K, modeled as a sheet mesh on the board surfaces since its thickness is too thin

to mesh properly with a 3D mesh structure. The impact of removing the coating reduces the peak temperature to 74.2°C from 76.7°C. The impact is not very large on a VRM (3.4%) because the board-to-ambient thermal resistance is on the order of 30°C/W and the coating only represents 0.1°C/W at most (Equation 9), because of how thin it is. Were the board to have a lower board-to-ambient thermal resistance, then the impact would be more substantial.

Using the nominal case PCB conductivity values of 32 and 36 W/m°C for the top and bottom parts respectively (as described in the previous section) and an ambient of 25°C, the simulation result is shown in Figure 2.16 using the I-DEAS VRM model shown thus far. The solver algorithm used to solve the model temperatures was the Fixed Viscosity Flow Model solver, meshed as described in the previous section, with the simulation parameters as summarized in Table 2.6. This solver method is very robust and does not have the potential convergence problems of the k-ε solver so for a model of this complexity at small scale, it is the wisest choice. The maximum temperature is 76.7°C, average is 67.6°C and two of the inductors are not shown so as to make the heat distribution more clearly.

The 105 CFM airflow fan used in the simulation was modeled using the I-DEAS library file which actually contained the exact model fan we used on the test bench. The maximum airflow rate is 1.41e4 mm/s (10.41 m/s) and the average inside the control volume is 9.332e3 mm/s (9.332 m/s). This shows that the pressure build-up in the control volume was very weak since the inlet and outlets are close to each other (155 mm apart) and that the actual fan blade diameter is close to the box opening (88 mm

versus 70 mm respectively). The actual operating point is shown in the next section in more detail.

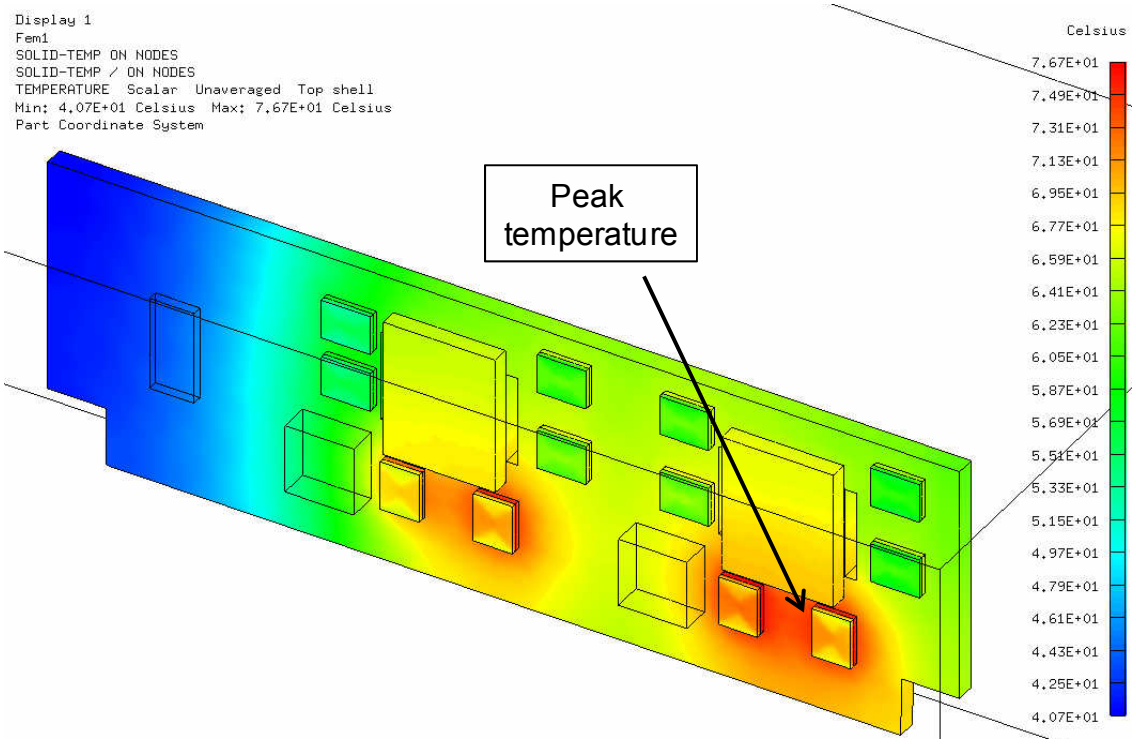

Figure 2.16. Thermal map results of Strategy #3.

Table 2.6
Simulation values for thermal map of Strategy #3

Component	Quantity	Loss (W each)	θ Resistance to PCB (°C/W)
Top FETs	4	0.67	0.35
Bottom FETs	4	0.54	0.35
Rectifiers	8	1.38	0.35
Winding blocks	2	1.26	1
inductors	4	n/a	4.9

So far, the simulation results' merits have been relative to each other. In order to establish an absolute point of reference, an experiment is conducted in which the VRM is tested in actual circumstances using a control airflow volume that is the same as used

in the I-DEAS simulation so that the component temperatures can be correlated. The test board used for this experiment is shown in Figure 5.1 with the VRM plugged in and the thermocouples attached. The control volume box used to channel the airflow is placed on this test board with a slot cut out of the bottom so that the socket can be accessible, shown in Figures 2.17 and 2.18.

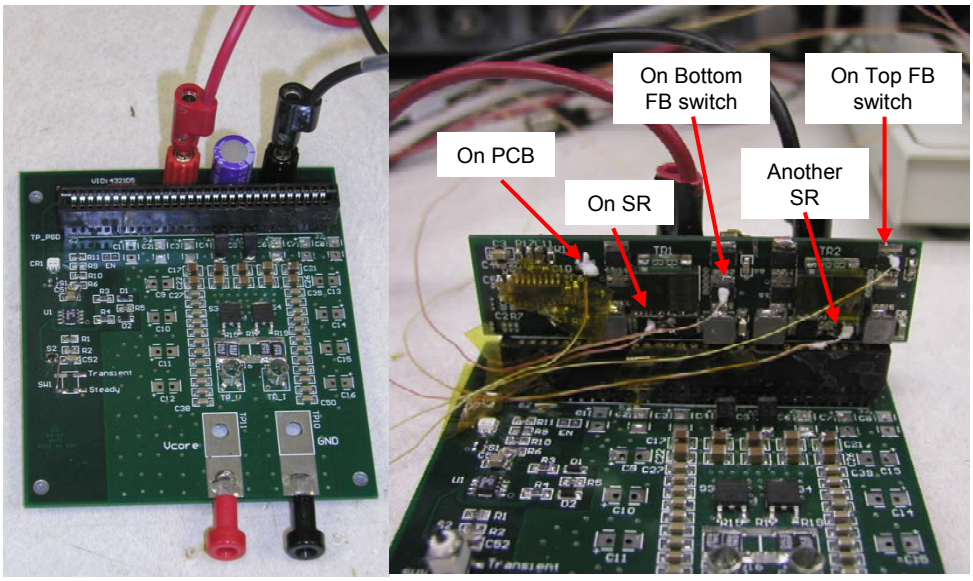

Figure 2.17. VRM with test board and thermocouples

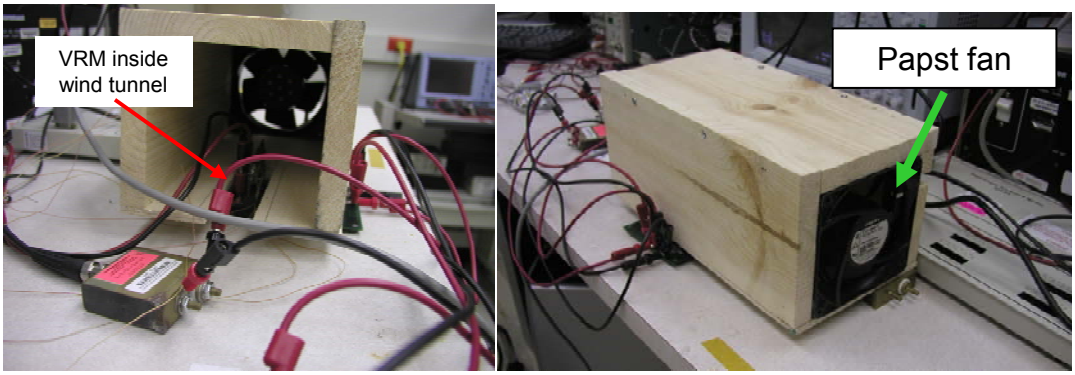

Figure 2.18. Wind tunnel with VRM inside it (left) and with test board protruding out from underside to make power connections (right)

The fan used is the one visible in Figure 2.18 and is a Papst TYP 4600X capable of 105 CFM airflow at zero pressure, according to its datasheet. This exact fan's flow curve from the I-DEAS library has been used in all these simulations for accuracy (discussed more in next section). The thermocouple monitoring system is the Cambridge Accusense TCM-24 Thermocouple Temperature Monitor where 7 thermocouples are connected as shown in Figure 2.19 and described in Table 2.5.

The results of this experiment correlate rather closely with simulation. The results of the temperatures are graphed in Figure 2.15. Referring to Table 2.7, simulation temperatures are very close for the output rectifier devices, less close for the top-side MOSFETS, and lower for the PCB spot. Figure 2.20 shows these locations on the actual board for reference. It seems that in actuality, the PCB gets hotter on the right side than the simulation shows (a thermal mass issue related to, once again, the fact FR4 has such low thermal conductivity). This accounts for the right-most rectifier and right-most MOSFET being hotter than shown in simulation. The interface thermal resistance of 0.35 °C/W is apparently close however since devices *not* on the right half have closer temperatures with simulation.

Overall, the simulation results are close to experiment when thermocouple errors and simple simulation assumptions are factored in – particularly since this system model has many small parts on a single PCB substrate. Much higher accuracy can be obtained when the individual parts on a PCB are on the same order of magnitude in size. The modeling techniques demonstrated here are proven to be valid concepts that can be applied in future PCB modeling of various electronic boards.

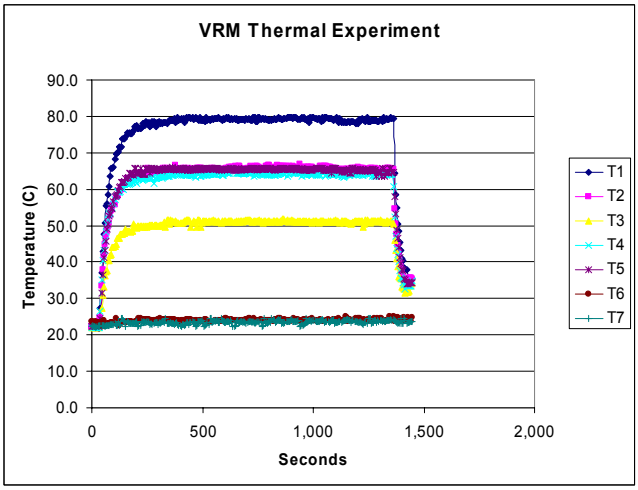

Figure 2.19. Board thermocouple temperatures show thermal equilibrium.

Table 2.7
Comparison of simulated and experimental temperature results

Thermocouple	Location	Simulation (°C)	Experiment (°C)
1	Right-most HAT2165 rectifier	77	79
2	Left-most HAT2165 rectifier	67	67
3	Central HAT2168 MOSFET	57	52
4	Right-most HAT2168 MOSFET	59	63
5	PCB left of transformer	56	65
6	Ambient on bench	25	25
7	In box's air stream	NA	24

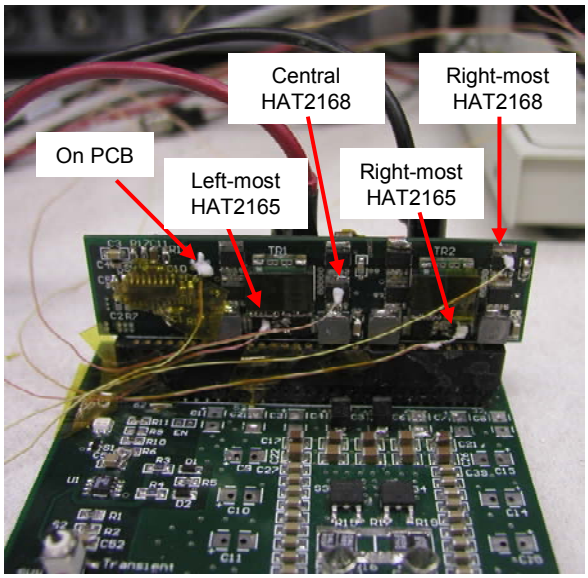

Figure 2.20. Showing the thermocouple locations on the actual board

An important point here is that the ambient conditions were of only 25°C which is much lower than typical VRM environment. The usual assumed ambient temperature is 55°C so to make an equivalent judgment; 32°C needs to be added to all the temperatures shown in Table 2.5. This would mean the peak temperature in a realistic environment is 111°C, which is close to the typical maximum of 125°C. And, the converter was not running at full current in this case – only 70A instead of 100A. This demonstrates that this state-of-the-art VRM requires a high-speed cooling fan of at least 400 LFM for safe operation despite the industry-leading efficiency, as will be discussed in the next section. Thermal issues are omnipresent in VRM design and a potential solution to this problem will be shown in subsequent chapters of this dissertation.

2.4 Thermal Analysis

Thermal design for a converter must first start off with high electrical efficiency. Tackling this challenge has been shown in the previous section where we have chosen a circuit that has high-efficiency capability – better than what has been demonstrated thus far in VRM commercial applications. Once this has been set, the next step is to determine how this converter fares in its operating environment.

Once the power loss of the circuit is known, determining the operating temperature of any converter depends on its environment, and more specifically, on three factors:

- Airflow rate and temperature
- Orientation of the converter relative to airflow
- Interconnection of the converter with the rest of the system

All three of these factors are inextricably linked and cannot be separated. If the ambient temperature increases, the airflow must also increase to maintain the same motherboard temperatures. If the orientation of the board relative to airflow is perpendicular, then hotspots will form on the backside. If the long side of the board is axial to airflow, then maximum cooling will be possible.

When subjected to these environmental conditions, the VRM design using Strategy #3 (the one with surface traces modeled and isotropic thermal conductivity of the internal PCB layers) and the losses shown in Table 2.1, showed in the previous sections, follows the simulated curve shown in Figure 2.21, relative to airflow rate for a given ambient temperature of 55°C, which is considered typical in computer and

telecommunication systems. As we can see in the typical volumetric airflow range of 5e6 mm3/s (equivalent to 1 m/s or 200 LFM) with the control volume inlet area of 70 x 70 mm), the VRM design is too hot to function reliably as currently designed. The maximum temperatures are those of the leadframes of the rectifiers, since we have seen that they are the hottest components on the PCB, and the average temperatures are for the entire VRM.

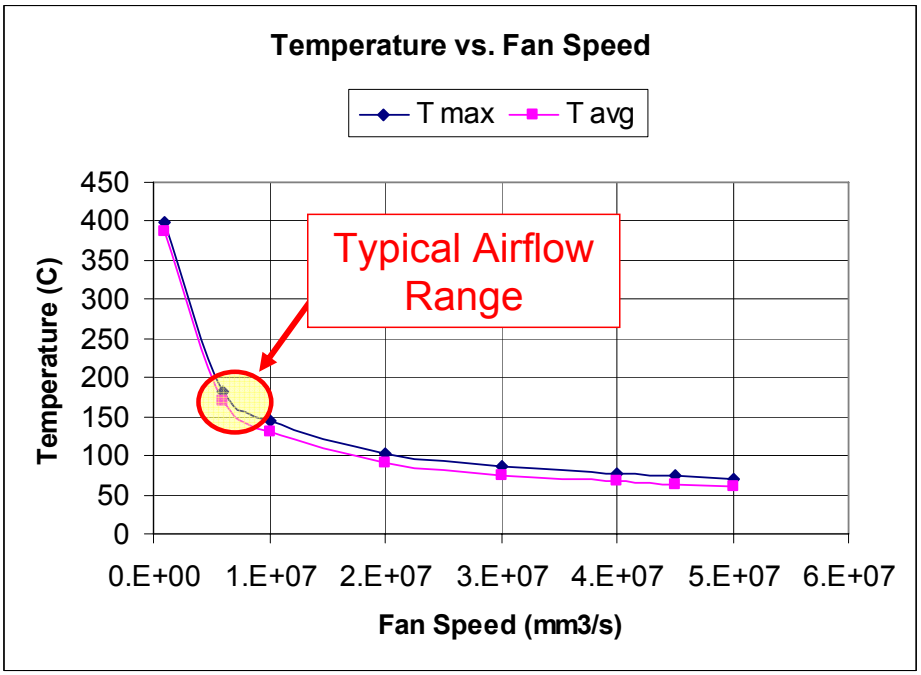

Figure 2.21. VRM temperature relative to airflow

The reason for the shape of this curve is what is called the heat transfer coefficient (HTC), which is based on a number of air properties, as well as geometrical calculations of the converter surface (neglecting surface components) and has units of $W/m^{2^{\circ}}K$. The HTC is the principle metric for convection cooling effectiveness and so is frequently assessed in mechanical engineering analysis. The equation is based on empirical data for highest accuracy because modeling of air flow is subjected to much

more random fluctuations than is typically encountered with electrical analytically-described systems. Therefore, thermal models are not as accurate as electrical models, in general.

The simulated plot of the HTC versus the temperature for strategy #3 is shown in Figure 2.22 [95] using I-DEAS NX software. As we can see, the average HTC climbs fairly linearly with increases in fan speed where as the localized maximum HTC makes an abrupt increase beyond 1e7 mm^3/s (2 m/s or 400 LFM in this case) fan speed and continues to climb more-or-less linearly from there. Comparing this to the average HTC, we see that the difference increases with fan speed, indicating that there is more of a localized cooling effect due to a turbulent flow regime at higher fan speeds that reduces the effects of the boundary layer.

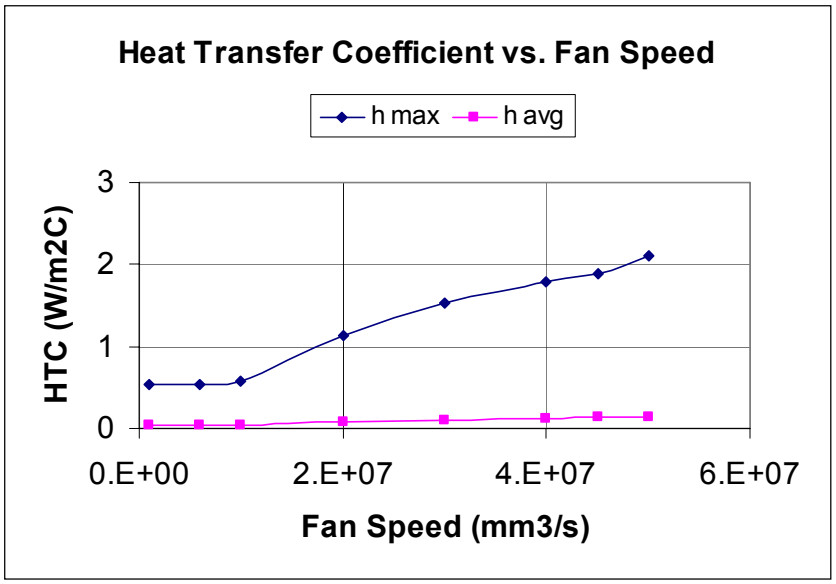

Figure 2.22. Simulation result of HTC versus fan speed.

The classical empirical formula for HTC in $W/m^{2\circ}C$ over a flat surface is shown below [89] and will be discussed in great detail in Chapter 5. This empirical formula is

valid for Reynolds numbers below 200,000 and in this case, the Reynolds number ranges from 4000 – 9000 for airflow rates between 200 and 400 LFM (1 and 2 m/s). The airflow is nearly always in the turbulent regime for electronic enclosures. The heat transfer coefficient varies with air speed and converter surface area, whereas heat loss is also dependent on the change in temperature between the surface and the ambient air temperature.

$$h = \frac{0.664\left(\frac{U\rho \cdot x_a}{\mu}\right)^{0.5} \mathrm{Pr}^{0.3333}\, k}{\frac{x_a x_p}{2x_a + 2x_p}} \tag{10}$$

$$q = (x_a x_p) h \Delta T \tag{11}$$

Where,

- U is airflow speed (m/s),
- ρ is air density (kg/m^3),
- x_a is axial length (relative to airflow direction) of the converter (m)
- μ is absolute viscosity of air (Ns/m^2)
- x_p is the perpendicular length of the converter (m)
- Pr is the Prandtl number which is generally 0.71 (no units)
- k is the thermal conductivity of air, in this case 0.032 W/m°C
- $\Delta \mathrm{T}$ is difference between ambient and average surface temperatures. For this example, it is 125°C - 85°C = 30°C.

I-DEAS software uses the MAYA thermal engine for the Electrical System Cooling algorithm FEA calculation method where the heat transfer coefficient, h, is defined based on finite element incremental calculations, but follows a similar concept as the classical empirical formulations [131]:

$$h = \frac{\rho C_p u^*}{T_{LS}^+} \tag{12}$$

Where,

- C_p is specific heat

$$u^* = \sqrt{\frac{\tau_w}{\rho}} \tag{13}$$

$$\tau_w = \mu \frac{\partial V}{\partial n} \tag{14}$$

- $\frac{\partial V}{\partial n}$ is velocity gradient normal to the wall
- T_{LS}^+ is the standard thermal wall function for low speed flows (< 4m/s, 800 LFM) defined as,

$$T_{LS}^+ = (\Pr y_f^+)e^{-\Gamma} + \left(\frac{1}{\kappa}\ln(\Pr y_f^+) + C\right)e^{-1/\Gamma} \tag{15}$$

Where,

- κ is a dimensionless constant,

$$\Gamma = \frac{0.01(\Pr y_f^+)^4}{1 + 5\Pr^3 y^+} \tag{16}$$

Where,

- y_f is the normal distance from the wall to the near wall mesh node defined as,
- y^+ is a dimensionless velocity defined as, alsng with its normal equivalent y_f^+
- y is some distance from the wall depending on mesh size

$$y^+ = \frac{\rho y\, u^*}{\mu} \tag{17}$$

$$y_f^+ = \frac{\rho y_f u^*}{\mu} \tag{18}$$

- C is a dimentionless constant based on the Prandtl number

$$C = (3.85\,\mathrm{Pr}^{1/3} - 1.3)^2 \tag{19}$$

Where,

- Pr is the standard definition of the Prandtl number
- Cp is the specific heat of the fluid (air)

$$\mathrm{Pr} = \frac{\mu C_p}{k} \tag{20}$$

The various parameters used for these solutions are obtained from a user input, such as shown in Table 2.6, or obtained from a thermal textbook for air properties [89] or calculated internally by the FEA program to develop a converged solution.

Taking a closer look at the airflow on the board, there are some complicated mechanisms at work. A fundamental principle of airflow is that friction is maximum right at the surface. As you move away from the surface, any surface, the friction decreases until you reach what is called “free stream” air flow. At this point, the flow is at maximum velocity but at the surface, it is zero at the microscopic level. This is why driving a dusty car on the highway doesn’t make it clean. The dust doesn’t blow off the surface because airflow is essentially zero at that particle size. Figure 2.23 shows this in pictorial fashion.

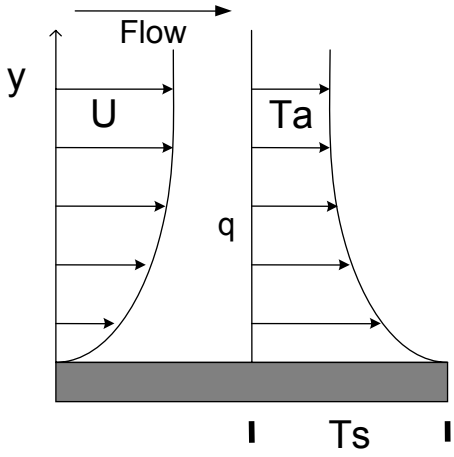

Figure 2.23. How airflow speed and temperature are affected by a surface.

This transition zone between the surface and free stream air speed is called the boundary layer. Its thickness is dependent on airflow speed and surface roughness/obstructions and has a certain thermal resistance associated with it since heat removal rate is directly, but nonlinearly, proportional to airflow speed. This resistance plays a key role in determining surface temperatures of a converter. Thermal resistance (*Θa*) is defined as the slope of *delta T* versus heat loss (*q*) curve. Figure 2.24 shows a schematic representation of these effects.

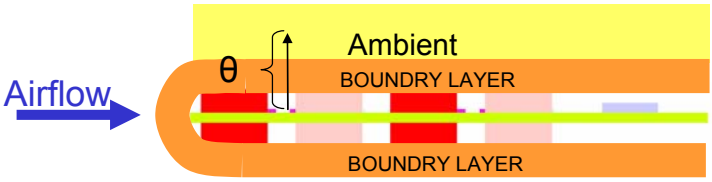

Figure 2.24. Top view of boundary layer around VRM and the relationship between temperature difference and heat loss

At the surface of the board, there are components that obstruct airflow. This causes the airflow to be in shambles at the surface rather than linear across it. This

“turbulent” airflow is essentially a given in just about any force-air cooling situation of a converter. What it means is that maximum theoretical convective cooling efficiency is decreased because the air velocity is different at every point/location on the board. The ranges of values are often more than one order of magnitude, depending on the size of surface obstructions and the nominal airflow rate.

So for example, using the same I-DEAS simulation software, Figure 2.25 shows simulated air velocity results for a 400 LFM (2 m/s) nominal airflow point-of-load converter simulation, which results in a large range of 40 LFM to 400 LFM at the boundary layer of the converter! The large box around the converter board is the control volume used by the FEA calculations, which need a closed-form environment to generate finite solutions. The airflow enters the right side of the box and exits on the left. You can see that corners and the front and back of components on the board see 10x less airflow than what is specified nominally. This has a large impact on the temperature of those parts. The smaller the parts, the less resistance the airflow path sees, which lowers the thermal resistance of the boundary layer and allows the converter to run cooler (or at higher output power for a given temperature).

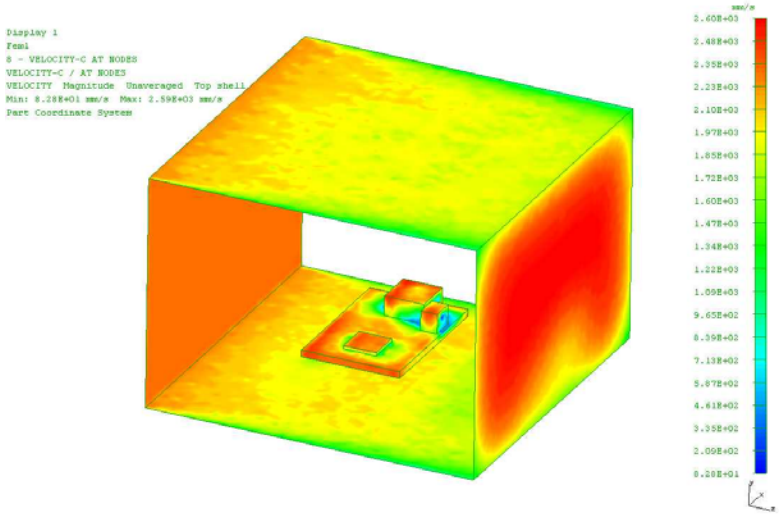

Figure 2.25. Air velocity map for a converter at 400LFM nominal airflow rate.

As a result of all this, the thermal design of a converter needs to take into account the physical location and size of surface components relative to the direction of airflow for a low-resistance path, the orientation of the converter relative to airflow, the interconnections between the converter and motherboard (more on this later), exposed surface area available for convection, power loss, and the airflow rate itself in the operating environment.

Another issue to consider with PCB is the internal layers. Since they are relatively insulated from the outside ambient air by the FR4 interlayers, they are susceptible to high heat loads. The presence of copper inner layers reduces temperatures up to a point – 4 internal layers are about the most before an increase in layer temperature is apparent. This is graphically shown in Figure 2.26 where 5 internal layers actually have more heat build-up than 3 internal layers! This is of course dependent on surface copper and number of vias that interconnect the layers. For this particular model, which is the same verified one (using strategy #3) shown in the

previous section, there are no thermal vias – only interconnect vias. But this shows that if you are to have more than 4 internal layers, you have better make room for plenty of thermal vias and surface copper traces to mitigate this potential problem.

The copper thickness also plays a role with reducing temperatures and this too is shown in Figure 2.26 by the downward slope of all three curves. As can be seen from the graph, doubling the copper thickness results in an approximate 5% drop in temperatures [94,95]. The logarithmic trend of the temperature curves follows Equation 21 as a general guideline and can be reliably used for boards with greater than 1 oz. copper weight (35um thick).

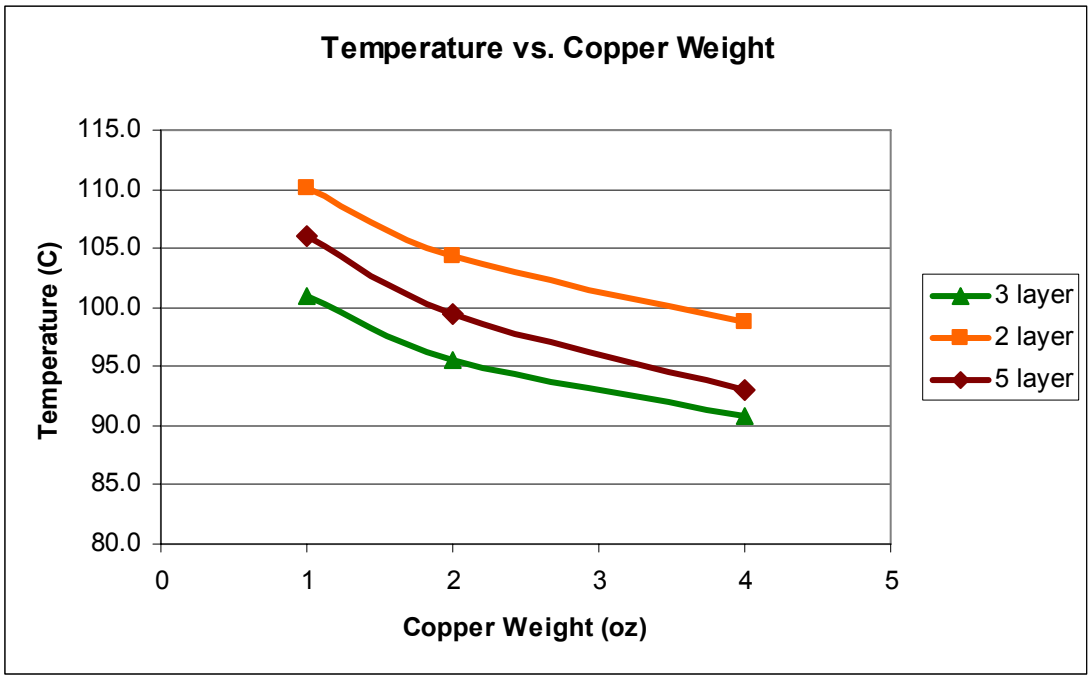

Figure 2.26. Impact of number of layers and copper weight on temperature.

$$T_{new} = T_{old} - 8 \cdot LN(C_{weight}) \quad (21)$$

Where, T_{new} = lower temperature due to thicker copper

T_{old} = measured peak temperature of thinner copper board

C_{weight} = copper thickness in ounces

The increase in temperatures due to thermal mass build-up is shown differently in Figure 2.27. Here was see the U-shaped curve of optimization. Layouts that have few thermal vias will be subjected to this curve. This is a side effect of using a substrate (FR4 interlayers) that does not dissipate heat. If more than 3 internal layers are desired, large numbers of thermal vias will have to be planned for to avoid overheating issues. More is discussed about this issue in Chapter 5.

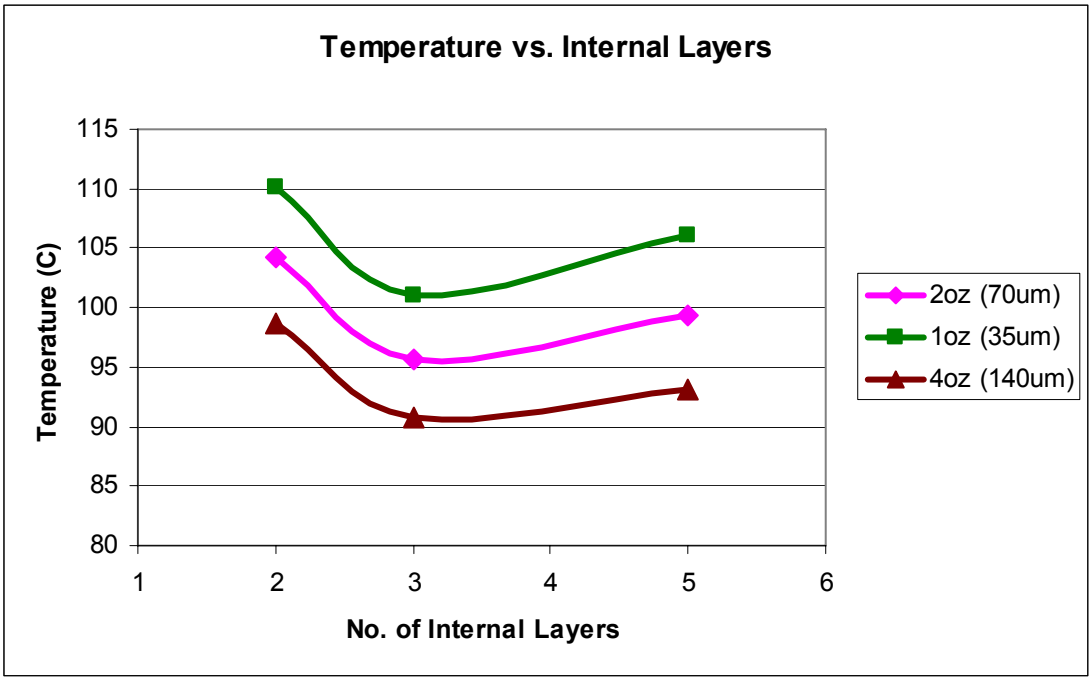

Figure 2.27. Impact of number of layers and weight on temperature.

2.5 Using the Thermal Model

In order to prioritize the thermal layout design rules, it is necessary to determine the impact of the conductivity of the PCB substrate relative to the impact of convection cooling off its surfaces. A sensitivity study of the board as a system is done using the Strategy #3 I-DEAS model to compare the ratio of these two general thermal mechanisms [95]. The outcome will be valid for PCB modules of a similar size as the VRM (93 x 24 mm). The methodology used consists of the typical statistical sensitivity study formulations, shown below, with a data point thermal conductivity increment of 5 W/m°C [90]. The nominal value of sensitivity in this case is 32 W/m°C since it is the PCB isotropic conductivity value for the majority of the PCB (see page 61 for details).

$$X_i^+ = \frac{\partial T^+}{\partial \beta^+} \cong \frac{\Delta T^+}{\Delta \beta_{Si}^+} \tag{22}$$

where

$$\Delta T^+ = \left| \frac{T_N(\beta_{Ni}) - T_S(\beta_{Si})}{T_N(\beta_{Ni}) - T_\infty} \right| \tag{23}$$

$$\beta_{Ni} = \text{nominal value of sensitivity}$$

$$\beta_{Si} = \beta_{Ni} + 5 \text{ W/mC} \tag{24}$$

$$\Delta\beta_{Si} = 5 \text{ W/mC}$$

$$\Delta\beta_{Si}^+ = \left| \frac{\Delta\beta_{Si}}{\beta_{Ni}} \right| \tag{25}$$

Performing these operations makes the assumption that the conduction and convection tendencies are linear. In order to ensure the results are valid, Figure 2.28 shows the variation in the thermal conductivity values and Figure 2.29 shows the variation in the fan speed according to the fan curve for the Papst 4600 fan curve used

in these simulations (it was already programmed in the I-DEAS fan library files and so it simply had to be selected to exactly match the same fan make and model available on the bench in the lab).

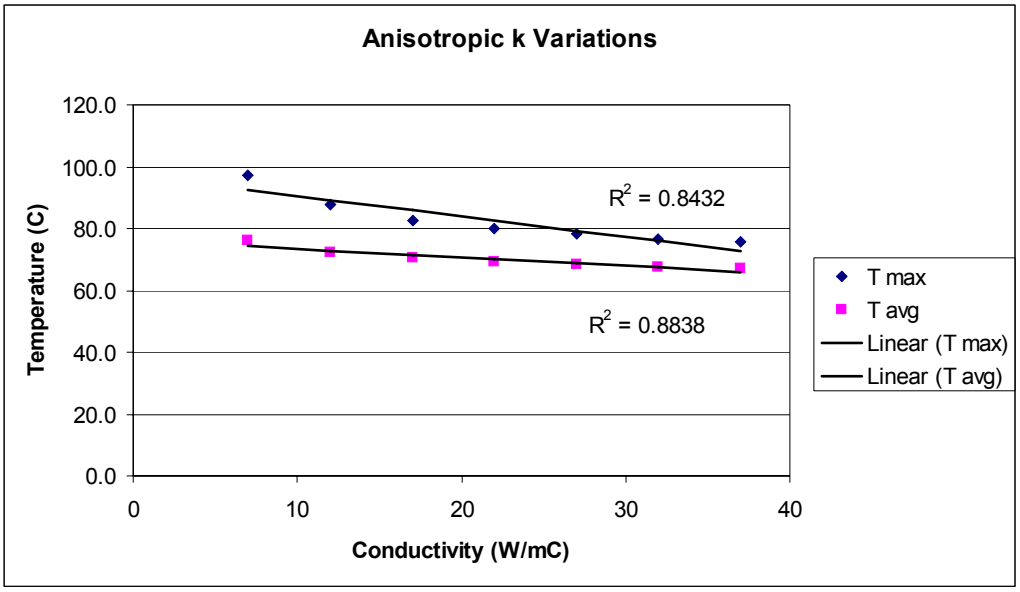

Figure 2.28. Variations of temperature over a wide range of thermal conductivity remains linear.

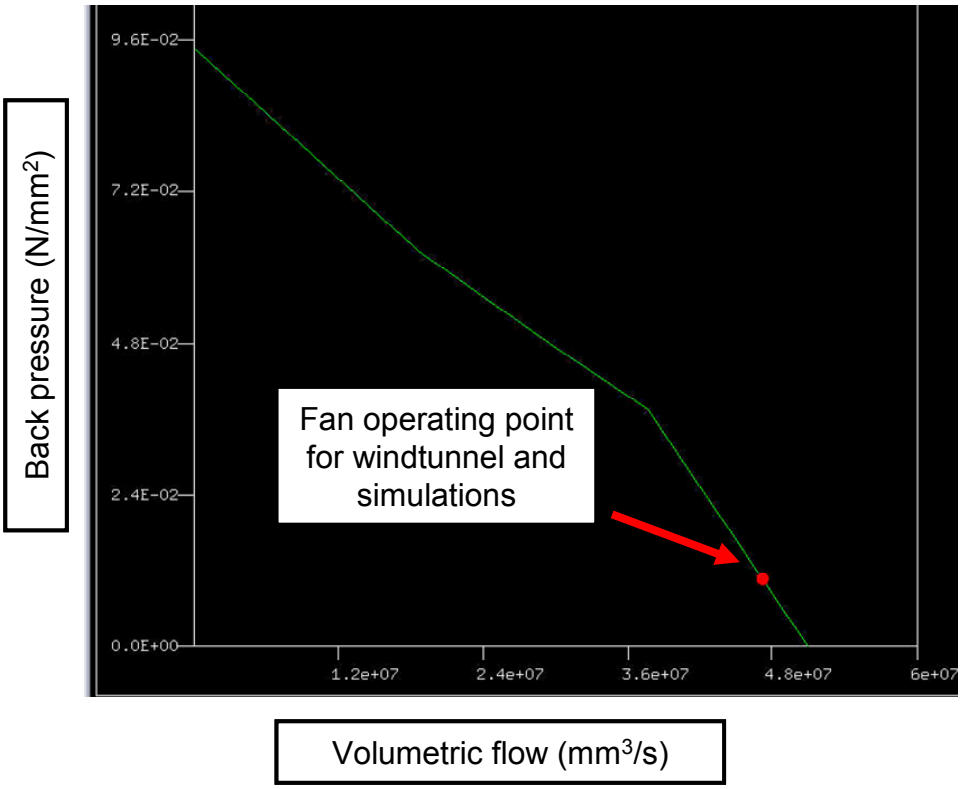

Figure 2.29. Fan curve of the Papst 4600 is linear over the region we are interested in (4e7 to 5e7 mm^3/s). Dot on the curve is the actual operating point in the control volume.

As can be seen from these graphs, the regions of interest are quite linear and so the results of a sensitivity study based on these would be accurate enough to give us an idea as to the ratio of impacts between conduction and convection make the biggest thermal cooling impact. In addition, the data point spacing has to be equal for both to make the comparison fair so the fan speed data points are calculated from the conductivity spacing and are shown in Table 2.8.

Table 2.8
Data point spacing for a fair comparison

Fan Speed	Max Temperature
5e7 mm^3/s	76.7 °C
4.32e7 mm^3/s	81.7 °C

The result of the analytical sensitivity study is shown in Figure 2.30 which shows an over 5 times higher sensitivity to convection (airflow) than conduction (PCB conductivity). This is a very interesting result that can be used to better our insight in a couple different ways. First, the use of the PCB modeling methodology developed in this project can be applied to boards with varying copper trace numbers and copper content with low deviation from valid results. This allows a claim to be made that increasing the number of layers in the board will *not* have a very significant impact on component temperatures (not to mention the fact that added traces do essentially nothing for increasing surface area). This effect was already alluded to and recognized in the previous section. This makes sense due to the poor thermal conductivity of FR4 preventing copper traces from releasing heat effectively, as shown in Table 2.6. However, this applies only to the internal traces because the two surface traces, as demonstrated in the previous section, contribute significantly to heat removal. So effectively, the internal traces are blocked from really transferring heat by the neighboring FR4 layers. Since airflow is independent of the internal PCB trace modeling methodology, the technique presented here is valid for a wide array of PCB boards, so long as they are of similar overall surface area.

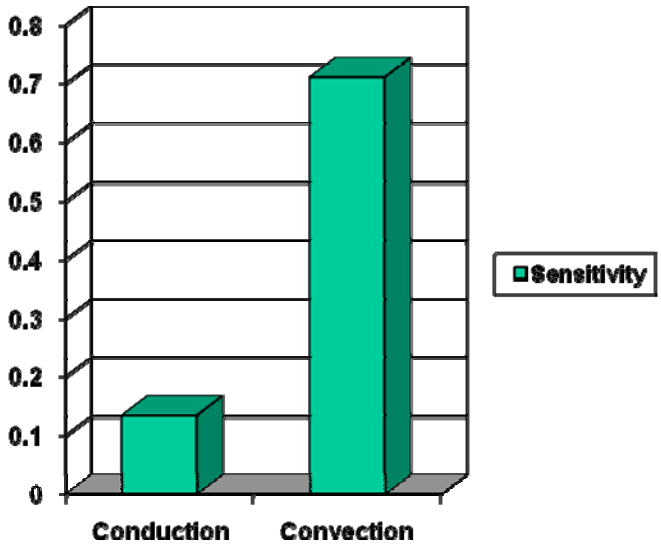

Figure 2.30. Result of sensitivity of VRM temperature with respect to airflow speed (convection) and PCB thermal conductivity (conduction).

In addition, we can use the sensitivity study results for determining the best approach to organizing our priorities for cooling our board. As such, convection plays such an important role in the cooling of the PCB that we need to use it to its fullest capability. Going back to the previous discussion, we have two ways to cool the hottest components: by convection and conduction. If we could combine these two, we would have a better approach for cooling a system such as this VRM. Let us look at the situation: We have a high-performance LFPAK device, chosen for its excellent electrical characteristics, with a thermal resistance of around 1 ºC/W from junction to board. The junction to case (top surface) resistance is on the order of 20 ºC/W (about the same as for SO-8) and so 95% of the heat will be conducted into the PCB instead of transmitted to the surface of the device. This heat flow priority path is described in [91]. The DirectFET is more even in its distribution since part of its design goal was to shift the imbalance but still 66% of the heat is conducted to the board. With all this heat

going to the PCB, why cool the device from the topside? Why not improve the cooling from the bottom side instead since convection plays such a significant role?

Using the thermal model for typical full load conditions, we made 4 different versions to see what improvements can be had [96]. Figure 2.31 shows the baseline case and Figure 2.32 shows the final design along with actual hardware. Both simulations are run under the same conditions: at full 100A output (21W loss, assuming same efficiency which is basically true), 55°C, and 400 LFM airflow. In order to reduce the temperatures, the design includes many thermal vias along the top edge of the board. We would have liked to have added even more but could not due to the PCB windings of the transformers.

The vias contributed to increasing the isotropic thermal conductivity of the board. Based on the number and density used in the hardware, it has been shown that they contribute to a 16% increase. The I-DEAS model PCB was repartitioned for those via "zones" and those paritions were assigned a thermal conductivity values of 42 W/m°C instead of the 36 W/m°C nominal [97,98]. Losses were essentially unchaged due to similar hardware efficiency so refer to Table 2.1 for those values. In this way, the model results can be compared to each other from a strict thermal and airflow perspective.

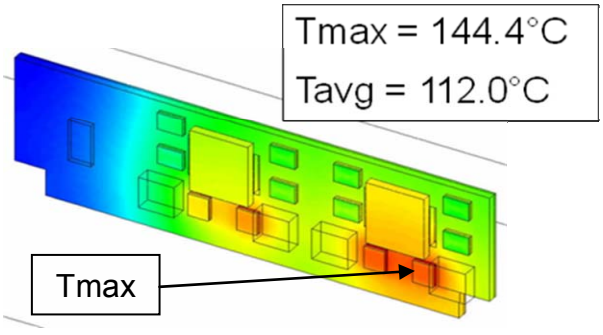

Figure 2.31. Original hardware case at full load and this time under typical environmental conditions.

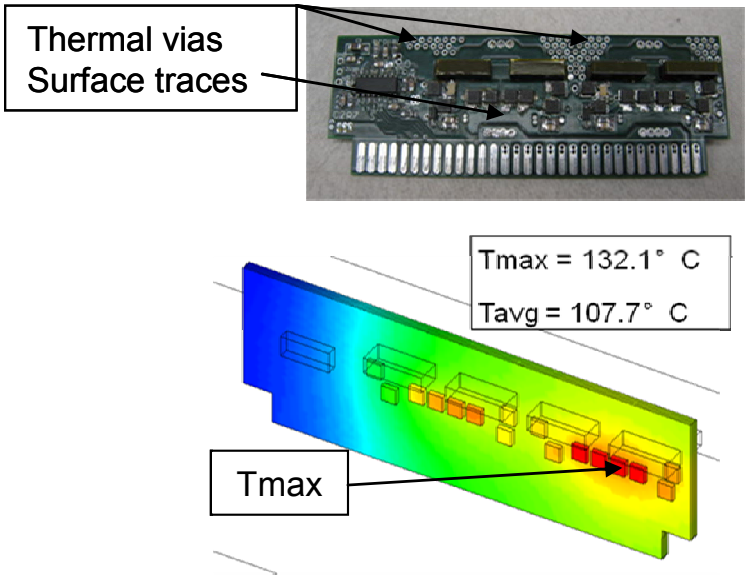

Figure 2.32. Model results from baseline case to final design with hardware [96].

The temperatures also came down thanks to more large surface copper traces, which as shown in the previous section, a doubling of the surface area increases the HTC by 50% locally. And the final improvement came from, significantly reduced airflow path impedance thanks to removing the inductors that were next to the SRs, and using smaller/lower-profile devices. This increased the local airflow rate by 23% on the surface of the SRs, which are, and remain, the hotspots in this VRM design.

However, this is still not enough of an improvement to meet our target of 125°C maximum temperature under these typical conditions. Either a lower ambient temperature, or a higher airflow rate, or a specialized shroud (as is often done in workstations for this very reason) is necessary to get the temperatures down another 10 degrees. But these are environmental factors that the system design engineer must take into account.

In terms of generalized modeling methodology steps, several key points have been demonstrated by using a VRM as an example. Our process has shown that, for a system-level analysis the following is true:

- The modeling method of the PCB must include the surface traces for a 10% gain in accuracy, as proven by the hardware verification (Table 2.5).
- The poorly-thermally-conductive FR4 layers block the internal traces' effects and so the internal layers can be modeled as a lumped thermal conductivity based on a single calculated isotropic value (Equations 4 and 5, Section 2.2).
- This internal isotropic value does not have to be very accurate since it represents only 20% of a comparable change in airflow rate for a surface of approximately 20 cm^2 in axial flow (Figure 2.30).
- Actual MOSFETs themselves can be modeled as blocks provided the lead frame is included due to its 260 times higher thermal conductivity relative to the surrounding encapsulant (Table 2.5).
- The die does not have to be modeled but the lead frame should be since LFPAKs and other modern packages use different materials top and bottom of the package (Figure 2.14).
- Components that have large profiles compared to MOSFETs, like inductors, need to be modeled even if they give off negligible heat since they modify the airflow over the board (Figures 2.13 and 2.25).

The electrical considerations are closely tied to the thermal ones. The layout of components can only contribute so much to heat interaction via conduction (Figure

2.30). Components that block the airflow can be moved an electrically-negligible amount that thermally makes a sufficient difference as shown in Figures 2.31 and 2.32 where reducing the airflow path impendace by 16% by choosing low-profile components and integrating the inductors resulted in an 8% decrease in peak temperatures and 4% in average temperatures. Surface traces are the main heat paths for convection – if little or no surface copper is present, then FR4 will not distribute the heat due to its 760 times lower thermal conductivity than copper (Table 2.2).

The maximum board current carrying capability is mainly dependent on surface area as well as airflow rate (Figure 2.21) and so must be considered as a coupled problem that is best solved by FEA methods since analytical ones can have excessive error when generalized to differing geometries [81,84].

2.6 Comments on PCB Limitations

So from this study, it is clear that the PCB leads to hotspots which reduces convection efficiency and prevents this VRM module from running at full power without a high speed fan to cool it. The ideal situation would be to have a substrate (since it is the circuit component with the most surface area) that has infinite thermal conductivity. This would lead to near-uniform surface temperatures, thereby maximizing thermodynamic efficiency and reducing the peak temperature to that of the average temperature. Figure 2.33 shows two simulations done under the exact same loss and airflow conditions except that one has a 4-layer 2,1,1,2 oz. (70, 35, 35, 70 μm thick) FR4 PCB substrate with an internal isotropic thermal-conductivity of 30 W/m°K,

determined by Equations 4 & 5 and described in Section 2.2. assuming full surface copper, and the other has a 150 W/m°K aluminum nitride Direct Bonded Copper substrate also with full surface copper. The surfaces are all copper because it is assumed that the dies are embedded inside the substrate for both cases in order to make a fair substrate comparison.

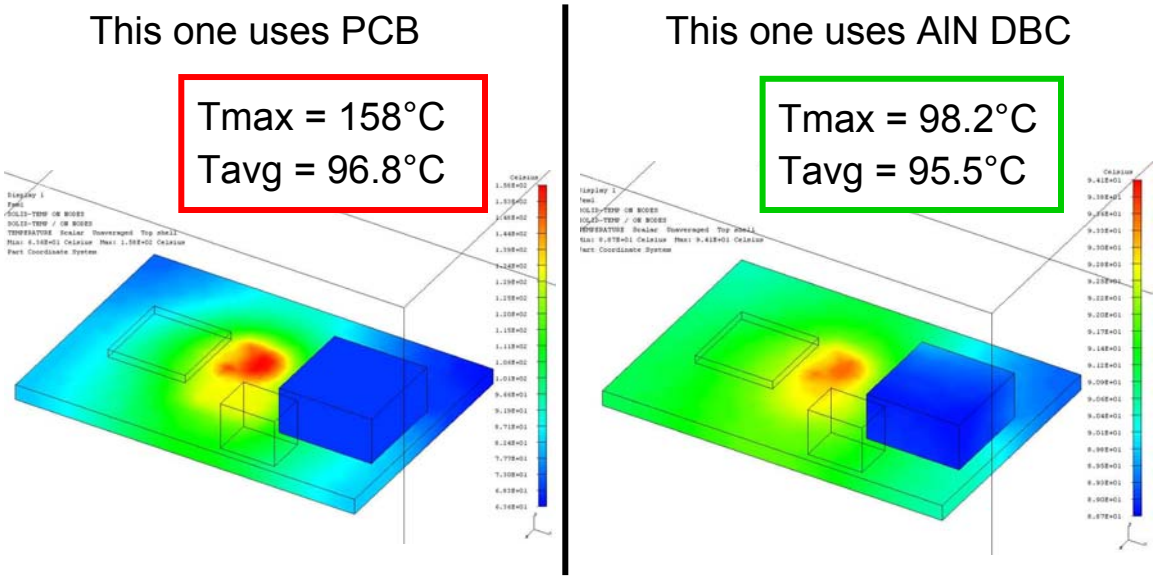

Figure 2.33. Impact of changing substrate material on peak temperatures.

Note that the module with AlN DBC has 40% lower temperatures than the PCB version! This is because AlN DBC has ~6x the thermal conductivity of 4-layer PCB (150 W/m°K compared to ~ 30 W/m°K) and so the uniform heat distribution essentially makes the peak temperature very close to the average temperature – an ideal scenario. As the available surface area of modules decreases, more has to be done to get the heat out to the ambient air with the shortest path. Convection off the module surface is the shortest path, rather than having the heat conduct first into the motherboard and then convected off the motherboard to the ambient.

A heat sink could be added to the PCB board to reduce the temperatures but this will invariably add significant size to the package, thereby reducing power density. A heat sink will also increase the profile height of the module, making it more difficult to place in confined locations inside the end user's final system, whether it be a computer, TV, or telecom rack. Smaller size with higher performance is common requirement in today's systems.

In addition, the PCB is not conducive to embedding components inside it, short of windings. This leads to the situation where components can only be placed on the two surfaces of the PCB, which precludes the effective stacking of components as layers. The ability of being able to have multiple layers of substrates stack directly on top of each other, each with much-higher heat spreading capabilities than FR4, widens the performance capabilities to unparalleled heights. In addition, we saw in the last section that convection plays a huge role in the cooling of the devices through the substrate surface area. This means that the lower the thermal resistance of the device interconnections to the substrate, the more we can take advantage of the 5x greater impact convection affords us.

So the PCB prevents one from maximizing the available surface area for cooling, which increases temperatures, which requires more distance between hot spots, which decreases the level of integration, which lowers the switching frequency, which lowers the power density, which increases the real estate required, which makes it harder for the customer to use, and increases the cooling needs as a whole. Changing the PCB out with a higher-conductivity material makes a much bigger impact that it first seems.

This leads us to wonder what the best substrate material to use is. There aren't many other options really – at least economically feasible ones, which rules out graphite and diamond. Insulated Metal Substrates (IMS) are Direct Bonded Copper's (DBC) only real competitors today [85]. There are many types of IMS but the majority fall into three main categories:

1. Metal polymer composite
2. Selectively-anodized aluminum
3. Copper-clad Invar, ceramic and molybdenum

Figure 2.27 shows a graph of delta T versus output power for a given substrate size [91]. The numbers shown in the boxes represent the lines' equations, which are in fact the equivalent thermal resistances of the substrates. We can see that in this graph, the 60 mil aluminum-clad IMS is the best with a board-to-ambient resistance of 1.5 °C/W.

How does DBC fare compared to the IMS? Alumina DBC results in a board-to-ambient resistance of 8.4 °C/W (a peak of 112.7 °C in the same simulations as shown in Figure 2.34) for a single sheet. Aluminum nitride (AlN) DBC, at a bit more cost, has a super-low 1.27 °C/W, the lowest of them all. As part of future work, a thermal cycling study should be performed to assess long-term reliability and Meant Time Between Failure (MTBF) ratings for the different substrate materials.

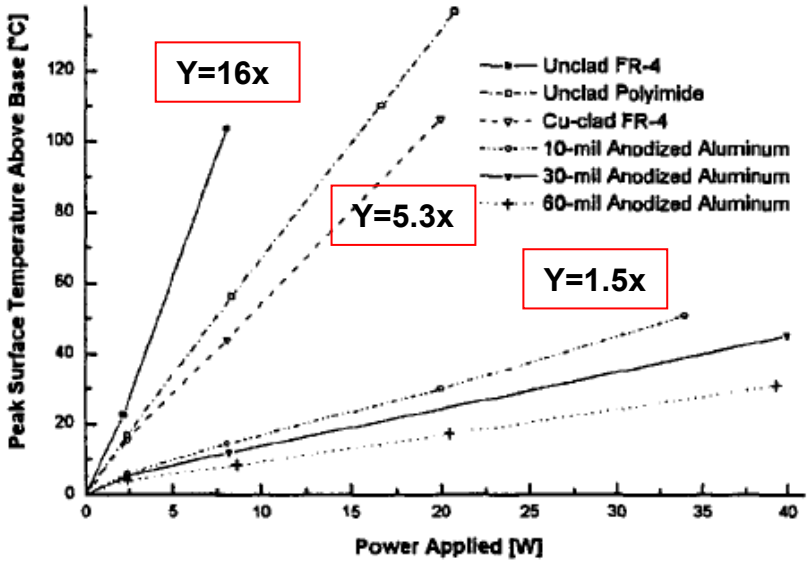

Figure 2.34. Comparison between various IMS substrates [91].

2.7 Summary of Modeling Methodology

In summary, an FEA model was developed based on our CPES-designed VRM introduced at the beginning of this chapter. Figure 2.1 shows the actual hardware and Figure 2.16 shows the corresponding thermal model. This model was then experimentally verified (Table 2.5), and great care was taken to ensure that the airflow regime witnessed by the real hardware was the same as in simulations (Figures 2.12 and 2.18). This was achieved by using a fan whose appropriate fan curve characteristic was found in the FEA software's library files (Figure 2.30).

With good accuracy established with the model, it was then used to assess the impact of changes on layout – without having to make all new hardware each time. This

is the main benefit of having a validated FEA thermal model. Four new layouts were simulated and the most advantageous one is shown in Figure 2.33. The use of the model allowed us to avoid making three less-satisfactory hardware designs and leap-frog right to the best we could do [96]. It was also used to demonstrate that multi-layer PCBs with more than 4 internal layers can have internal temperatures be higher for the middle layers as for the outer ones, again due to the 760x lower thermal conductivity of FR4 compared to the copper traces which leads to heat being trapped in the middle layers (Figures 2.26 and 2.27). It is best to layout the board such that high current traces are on the outer layers to avoid this potential problem.

The model's accuracy limit is not determined by current level or temperature because these values are simple numerical calculations. The model limitations are rather based on "feature size," which is the smallest component that needs good meshing accuracy, like FETs for example. The meshing for the FETs was based on a 1mm mesh size. If a smaller feature than an SO8 packaged FET needs to be simulated, then the mesh size will also have to be smaller, not only internal to the FET but also for the air mesh surrounding its exterior for adequate heat transfer accuracy.

It was determined through I-DEAS simulation that the air mesh minimum size should be no more than twice the smallest part mesh. However, the limitation comes into play when air mesh size goes below 1mm because then the number of nodes required for good accuracy becomes too large (e.g., greater than 1 million for a control volume like the one shown here) to properly converge in a reasonable amount of time. As mesh size decreases, simulation time increases rapidly due to the cubed nature of volume. For the 1mm air mesh size of the 155 x 70 x 70 mm control volume shown in

these simulations requires about 9 hours of simulation time using an Intel Pentium 2 desktop computer.

Another model limitation is the airflow rate. The model shown here uses the Fixed Viscosity Flow Solver, as discussed earlier. This model is only accurate down to approximately 50 LFM (0.25 m/s). Below this, the k-e flow solver is necessary.

Modeling of the PCB using isotropic lumped thermal conductivity values should be based on the arithmetic mean of the series and parallel sums of individual material thermal conductivities. This is because the thermal conductivities are additive quantities, rather than multiplicative, from the point of view of heat propagation. What changes is simply the rate of propogation inside each layer. Were the effects multiplicative, then the arithmetic mean would be more appropriate.

The lumped thermal conductivity only needs to be calculated for circumstances where the layers of differing materials extend out to the outer limits of the components – such as PCB or DBC. Components such as FETs are not that way because the heat loss is directly coupled to the lead frame, which exits the package and leaves the plastic casing behind (as for an SO-8 package). In this case, heat does not need to travel through the plastic so including it into the package thermal conductivity is not necessary.

The necessary steps to develop a simple, yet accurate, generalized system-level thermal model are summarized by the following bullets:

- The modeling method of the PCB must include the surface traces for a 10% gain in accuracy, as proven by the hardware verification (Table 2.5).

- The key equations are Equations 4 and 5 which allow for a simple analytical calculation to be made to lump all the internal layers of a multi-layer PCB into a single isotropic thermal conductivity value that can be entered into any FEA thermal software of your choice.
- This internal isotropic value does not have to be very accurate since it represents only 20% of a comparable change in airflow rate for any 1U form-factor converter in axial flow (Figure 2.31).
- Actual MOSFETs themselves can be modeled as blocks provided the lead frame is included due to its 260x increase in thermal conductivity relative to the surrounding encapsulant (Table 2.4).
- The die does not have to be modeled but the lead frame should be since LFPAKs and other modern packages use different materials top and bottom of the package (Figure 2.14).
- Components that have large profiles compared to MOSFETs, like inductors, need to be modeled even if they give off negligible heat since they modify the airflow over the board (Figures 2.13 and 2.25).

Chapter 3: Integrated POL Design

3.1 Overcoming Barriers for Integration

As we have seen in Chapter 2, the main barrier to integration is the substrate material used today: FR4 PCB. This material has a low thermal conductivity (~4 W/m°K) and so its only real use in the circuit design is as a support structure for components and interconnections. This was valid enough for the last 30 years but for today's high power densities, it is simply no longer enough. As we move towards more and more integration in a package, the substrate can play a key role and leap to the forefront of improvement with some thought.

So why do we want to make an integrated design? As discussed in Chapter 1, the obvious benefits are: small footprint and low profile thanks to the reduction of circuit parasitics, which allow for higher switching frequencies to shrink passive components' size. However, there are many levels and types of integration today. Some attempts have been to integrate everything on silicon, while some are hybrid packages for higher power, as discussed in Chapter 1. Which is used depends on the output current level and duty cycle. The final conversions are typically done in two steps: 5V-1.2V and 1.2V to 0.9V. The low voltage case has been targeted for chip-level integration and will be discussed here because its barriers and limitations give good insight on the problems with integrating for low duty cycles (i.e., high switching loss and large voltage conversions).

The ITRS2003 road map [101] for small hand-held electronics is shown in Figure 3.1. Power levels will be steadily increasing, and output current even more so. The main thrust for integrating on silicon comes from this road map and the low power levels needed today, and in the near future, allow that to happen. However, as we have seen with all other power supply road maps, power levels only increase and eventually, high power densities at high current output levels will be required. For now, we will take a look at the integration work that has been done for this road map.

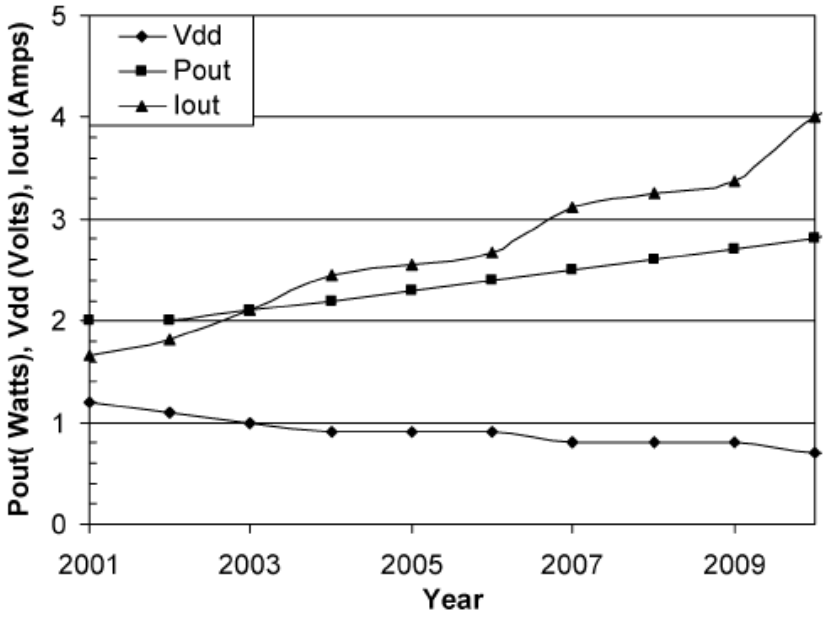

Figure 3.1 2003 ITRS roadmap for small hand-held electronics [49]

A step-down regulator can be implemented three different ways: as a switched capacitor charge pump, linear regulator or switching DC/DC converter. The limitation with the charge pump concept is low efficiency and large size for high output currents, even with bulk CMOS implementation and low 100mA output current. A linear regulator has a low efficiency of < Vout/Vin so it isn't very practical either despite its benefit of

very fast regulation and low silicon usage. The switching converter tends to be bulky with off-chip inductors and bulk capacitors but can have high efficiency.

With switching frequencies in the 10MHz range, the estimated inductor size in about 130nH and capacitance needed is 0.6uF for a 3W converter which should be compatible with integration on silicon and the Power Supply On Chip concept (PSOC) [49, 101]. One of the main loss mechanisms with inductors is their impedance so it has to be as low as possible. Inductance values of 20 – 40nH/mm2 have been achieved with a 1Ω winding resistance. This is acceptable for a 3 or 4A converter but not for anything higher than that.

Thin film magnetic material offers a 5x lower power loss density at 10MHz than ferrite and is compatible with semiconductor processing. This technology can reduce resistances to 120mΩ for a 120nH inductor; have a Q factor of 15 for an inductance of 350nH; and reach 1A of current with a 1μH inductor. All these values will still be limitations at power levels higher than 3W.

At even higher switching frequencies, it has been reported in literature that the 100 to 300MHz range can yield higher levels of integration but so far, only for very low output currents – less than 400mA. See Table 1.1 for comparisons and references. However, this frequency level allows for a reduction of three orders of magnitude for both capacitors and inductor compared to the typical 300kHz design. This allows for output decoupling capacitors to be integrated on chip and eliminates the need for inductors with magnetic cores. However, these inductors remain discrete, soldered-on SMT air cores, creating a sort of hybrid chip integration due to today's inductor technology limitations.

But at these high frequencies, output current is limited and the duty cycle must be high. The 1.2V to 0.9V conversion gives a large duty cycle of 0.75 so that high efficiency can be maintained without having to resort to transformer-based (isolated) buck topologies. However, this still poses problems for on-chip air core inductors since their Q factor is typically in the 5-10 range which is too low for efficiently transferring power. Not to mention that even low-current inductors with the latest technology take up a substantial amount of silicon area.

Capacitors are also a problem for on-chip integration. To be able to do it successfully, the total capacitance would have to be in the several nanofarad range to keep their size manageable and leakage loss low. One way this capacitance can be reduced without having to pay a large ripple-voltage penalty is to use a multi-phase interleaved buck topology. This also has the benefit of reducing inductor sizes since ripple current is also lower. However, enough capacitance is still necessary to maintain adequate regulation. If the capacitance level is in the nanofarad range, the required response time is in the nanosecond range which imposes a very high switching frequency and challenging control techniques.

Meanwhile, these converters with 1.2V inputs have to get their power from another converter. What about the integration of this converter, which is typically 5V to 1.2V? It turns out that this type of converter integration is hardly covered or looked at all in the literature. The reason for this is unclear but it may have to do with the fact that there is more available space allocated to high-power converters than there is for on-die versions. However the general trend is to reduce the size of everything so it is only a matter of time before these converters will also have to be smaller and more powerful.

In a laptop for instance, the amount of available real estate is shrinking rapidly so there is clearly a need for a highly-integrated high-current DC/DC converter that can output significantly more than 500mA.

Thermal considerations are a very important part of this work since it has been found in Chapters 1 and 2 that thermal issues present the main barrier towards further integration. Rather than handling these considerations after-the-fact with a bulky and expensive heat sink or requiring the end user to make a large space available on a multi-layer PCB with thick copper layers for necessary cooling (e.g., JEDEC EIA/JESD 51), we chose to come up with a system design that will find an elegant and compact generalizable solution that requires no heat sink or fan.

Our thermal analysis using I-DEAS 3D FEA software showed that the main limitation with the conventional approach is that the PCB's thermal conductivity is just too low to count on it for real "double-sided" cooling of the devices. This really leaves just a single heat-removal mechanism for cooling them: The use of a heat sink on top of the devices, particularly with the common SO-8 packages who have a junction-to-board thermal resistance on the order of 20°C/W, as depicted in datasheets. Even a low device junction-to-board thermal resistance such as LFPAK or DirectFETs (~ 2°C/W) is little help and can actually limit device performance by clamping the maximum temperature of the devices to the PCB's lower maximum temperature. This is a detail that is often overlooked to the detriment of absolute performance.

So in summary of the last two chapters, the general practice in today's typical POL module designs is as follows:

- ❖ Low-current (<12A) packages are migrating towards a plastic-encapsulated module with thermal pad interconnection to motherboard
    - o Most use monolithic active layers to reduce size and increase frequency
    - o Thin-film inductor is used in the research settings
    - o The motherboard is expected to handle the modules waste heat
- ❖ Higher current (between 12 and 20A) packages are either pin-interface or thermal pad-interface with the motherboard.
    - o These tend to mostly be co-packaged designs for higher current
    - o Inductor is co-packaged with the other components
- ❖ Current levels beyond 20A are nearly all through-hold mount
    - o Most are open-frame modules with low switching frequency
    - o Use PCB substrates so thermal performance is based on large board area/footprint for convection cooling directly off the module
    - o Inductor is a discrete conventional ferrite of large size.

3.2 "Stacked Power" Process

This new methodology is divided into two main parts. One is integration of the active components and the second is integration of the passive components. In the first, the devices used are bare dies (IR DirectFET IRF6633 and IRF6691 for the top and bottom switches respectively) that are embedded in a single layer of ceramic. Along with the MOSFETs, a snubber capacitor for the top switch is also embedded in

the ceramic carrier. The rest of this buck converter uses discrete parts for the inductor, driver, and output and input capacitors.

The vision of this project is to be able to integrate all the functional building blocks of a compact Point of Load converter into a small efficient module. A conceptual diagram of this vision is shown in Figure 3.2 where each layer will have its own job contributing to the overall cause. The key conceptual element here is the fact that integration is being done in the z-axis. Many of the converters shown in Chapter 1 demonstrate lateral integration. This is not the best way to reduce footprint, and thus more real estate on the motherboard is allocated than necessary. Vertical integration allows for a 3D volumetric concept to space utilization which greatly increases the performance envelope of a small converter.

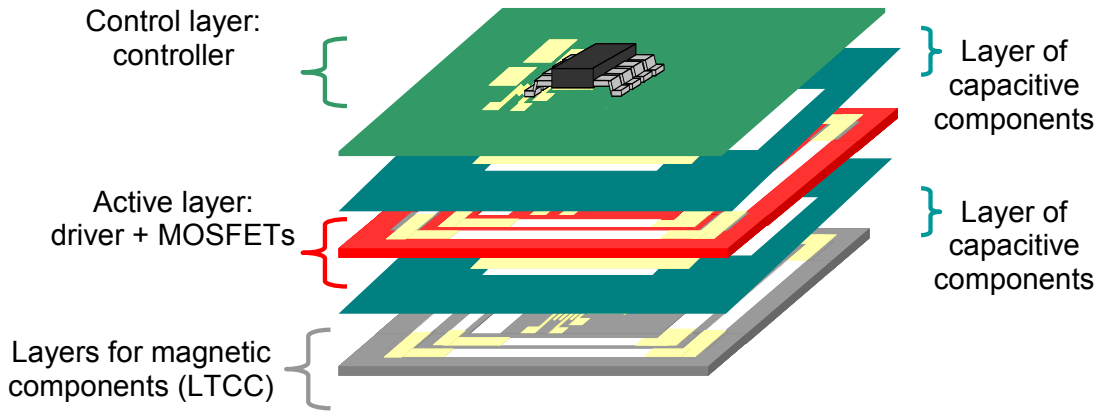

Figure 3.2 Overall vision for 3D integrated POL.

The achievement in the previous years has focused on an active layer using "Embedded Power." The results were not very good in terms of efficiency, with a peak of 76% and a modest output current capability of 9 A. Figure 3.3 shows the

complicated fabrication methodology used for Embedded Power based on alumina ceramic carrier. The interlayer was a dielectric coating used to isolate the sputtered and plated metal from neighboring traces and substrate.

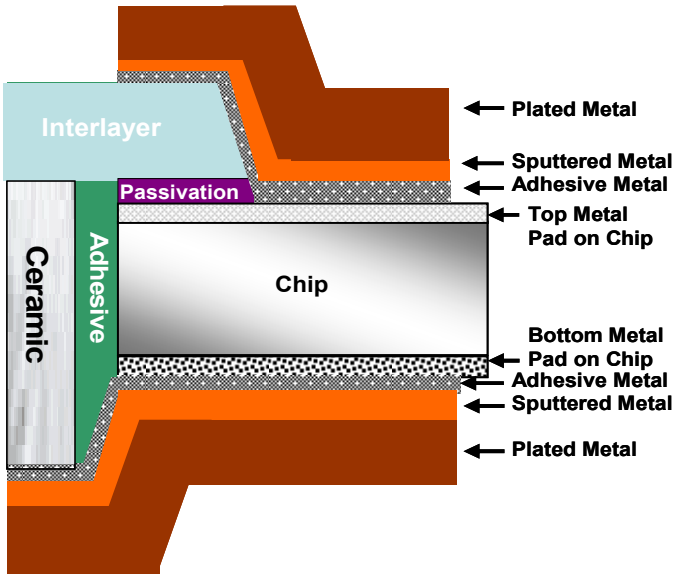

Figure 3.3. The many process layers necessary for Embedded Power.

Figure 3.4 shows the delamination issue where the copper traces lift off from the surface of the dies. Consulting a packaging handbook [116] and discussing the issue with packaging engineers at VPT Inc. and CPES, it is rather evident that the problem was due to attempting to sputter material on a die finished in aluminum, strictly intended for wirebonding. The wirebonding process uses ultrasonic application of the wire on the die surface in order to penetrate the inevitable oxidation that instantly forms on aluminum when it isn't in a reducing environment or vacuum. Even with careful plasma cleaning, oxidation would occur in handling between the plasma cleaner and sputter chamber, which would prevent the titanium from properly adhering. It would be possible to have a module be successful but the yield would be extremely low – and this was

proven when only 1 out of 9 samples actually worked, and then that one did for only 2 power cycles before failing.

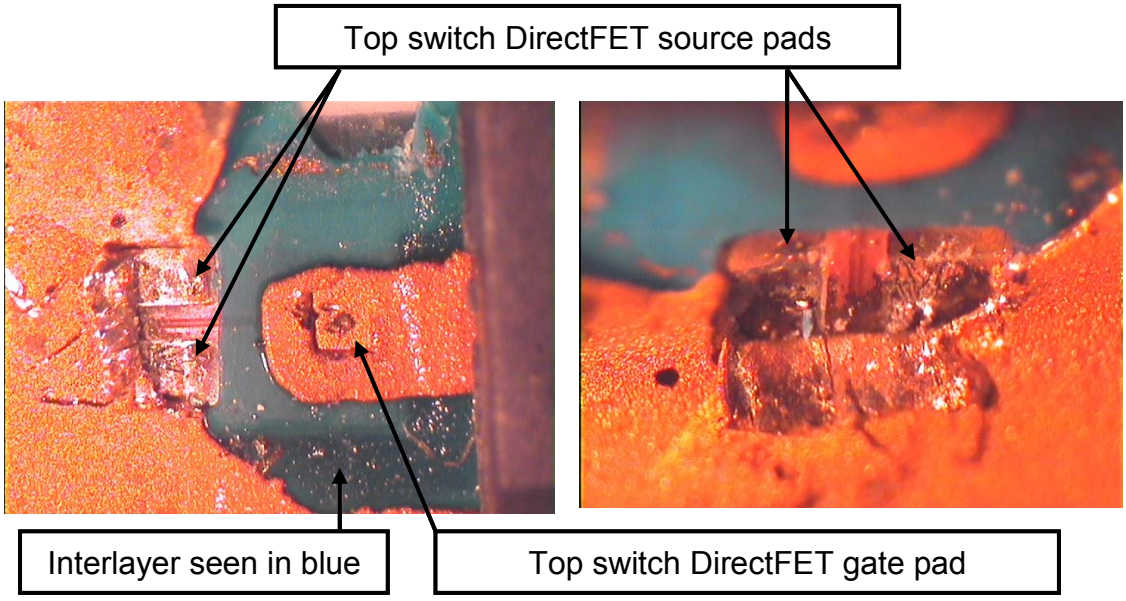

Figure 3.4. An example of the die attach delamination of Embedded Power

Stacked Power does away with this problem by using dies that are fabricated with the intention of soldering rather than wirebonding. Therefore, sputtering is no longer necessary because the newly-developed DirectFETs are solderable dies, and more options are coming out, such as the PolarPAK from Vishay. The devices used in this case are International Rectifier IRF6633 for the top switch and IRF6691 for the bottom switch. Now with these new devices, the extra process steps for sputtering are no longer required. Thus, the fabrication process has been entirely revised and now does away with sputtering of the metal and, consequently, its failure mode. The name has been changed from Embedded Power to Stacked Power to make this distinction clear.

Now instead of depositing the copper traces onto the ceramic, Direct Bonded Copper substrate is used so the negative is removed from the copper surfaces to leave very-well-bonded copper traces on the ceramic. Holes are still cut into the ceramic for embedding the bare dies and snubber capacitor but now instead of using sputtered metal to make the connection, another layer of DBC with the appropriate copper interconnection paths etched out of the bottom side is placed on top and soldered together using the very simple and well known reflow process with solder paste. This is schematically shown in Figure 3.5. For the outer sides of the module where there is no mating DBC layer, copper straps are used to make the interconnections. These are better than the traditional wirebond techniques because the wider and thicker copper reduces parasitic inductance and resistance, and as a result, are being seen used more and more in commercial applications, such as the next generation DrMOS from Renesas [117]. More will be discussed about this in the following sections. Figure 3.6 shows a simplified schematic representation of this new approach.

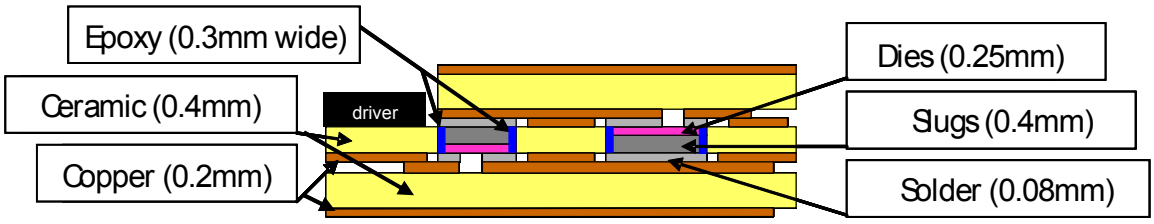

Figure 3.5. Schematic of embedded dies (in pink), heat slugs (dark gray) and 3 DBC layers

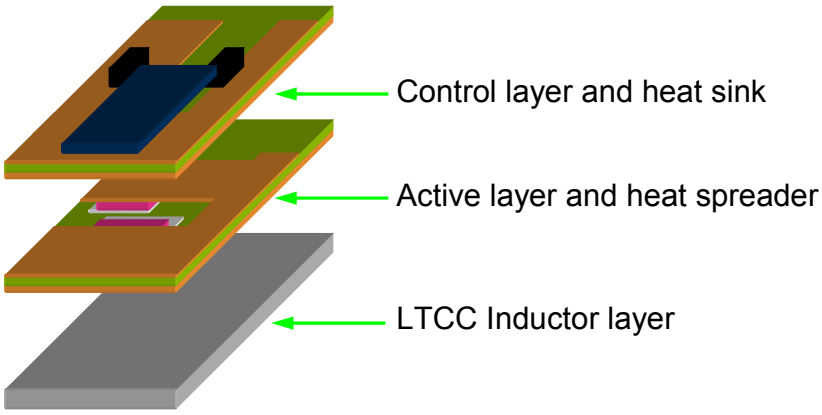

Figure 3.6. Schematic of embedded dies (in pink), heat spreading layers in green, and inductor substrate

There are other benefits to using DBC to make the interconnections, besides ease of implementation and reliability. One of the main ones is that we get integrated thermal management built-in. The ceramic is specified as AlN, and no longer as Al_2O_3, so thermal conductivity has increased by almost 600%. With high circuit efficiency and a cooling fan of moderate size, we can eliminate the use of a heat sink which will increase power density, lower cost and allow for a very-low profile.

Another benefit of using DBC is its heat spreading effect. AlN ceramic conducts heat 38 times better than FR4 PCB material so conductive cooling can happen in *all* directions at the same time. Heat is quickly spread through a large volume rather restricted to internal layers. At the surface, airflow path impedance is greatly reduced compared to conventional PCB designs thanks to the flat surfaces which further improve cooling efficiency. Section 4.4 will discuss these advantages in great detail.

Stacked Power uses another type of ceramic as the chassis for the active stage: Low-Temperature Co-fired Ceramic. This is actually ferrite particles mixed with a ceramic tape material. The thin, flexible tape layers are stacked together in various

shapes, pressed, and then fired in an oven to create a hard ferrite structure. This structure has the right properties to have the semiconductor layer right on top of it. This stacking of the layers allows for a relatively high-power (between 5W and 40W or so) converter to be integrated together to achieve excellent mechanical rigidity and high power density.

3.3 New active layer philosophy

By changing the substrate material and using it in a different manner, small size with excellent thermal handling is possible – to the extent that the need for a heat sink and a fan can be eliminated. This high level of performance results from a new process called “Stacked Power” which was devised to solve electrical and thermal issues from a system-level design point-of-view.

This design method uses aluminum nitride (AlN) Direct Bonded Copper (DBC) rather than PCB as the substrate. One of the main reasons is that we can now integrate thermal management as well as components! AlN ceramic’s thermal conductivity is about 150 W/m°K (whereas FR4 PCB is on the order of 4W/m°K) allowing for excellent heat distribution on the surface which increases the effective convection area greatly. Also, embedding the devices *inside* the ceramic allows for flat surfaces onto which other converter functions can be layered, yielding z-axis integration in a low-profile “stacked” arrangement. With high circuit efficiency, we can eliminate the

use of a heat sink which increases power density significantly and reduces converter profile.

All these advantages from using a new type of substrate, when compared to conventional PCB, culminate into the thermal packaging “figure of merit” of hotspot-to-package ratio now being many times larger. When using a substrate with a low thermal conductivity, such as FR4, heat doesn’t spread very quickly, creating hotspots in and around the main power dissipation devices – typically the MOSFETs. These hotspots not only limit the potential current output but they also reduce the package effectiveness. By going to a ceramic substrate, the hotspot is effectively spread over the entire package. The hotspot-to-package-area ratio has thus gone from about 20% to 85%. This improvement directly translates into either higher current output or smaller package size, whichever is deemed appropriate for the application at hand.

This new methodology allows for heat removal paths in all directions, as well as taking advantage of very low thermal resistances since ceramic does not limit the maximum temperature of the devices like PCB does. All connections between devices and heat sink (i.e., substrate) are less than 0.5°C/W thanks to the use of planar interconnections, so heat flow is highly effective and the ceramic spreads it efficiently over the entire surface for maximum convection transfer to the ambient. The two methods are schematically compared in Figure 3.7. Details on the fabrication techniques for the new thermally-improved Stacked Power method will be shown in the following sections.

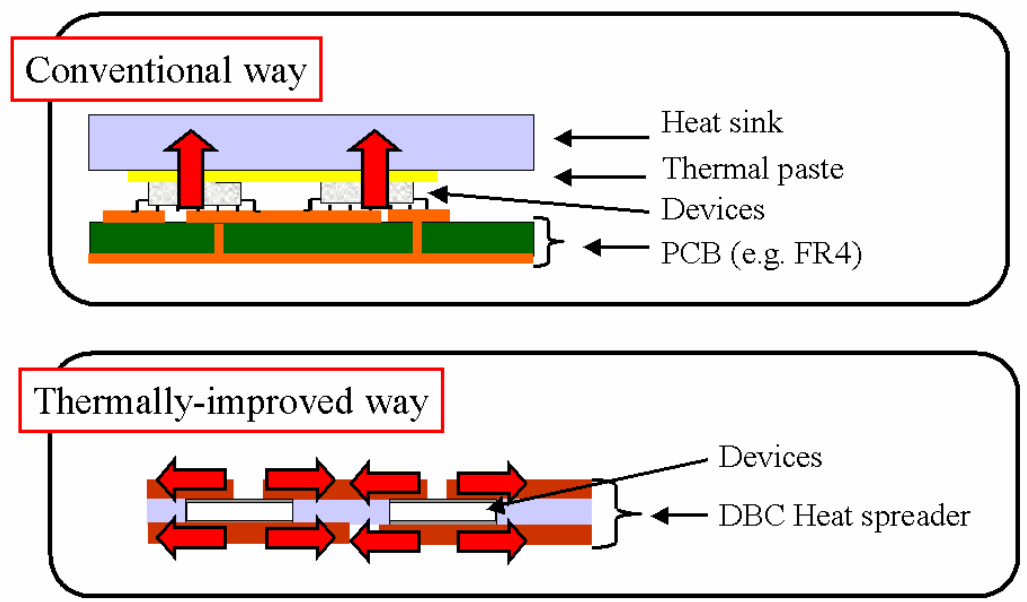

Figure 3.7. The new method allows for greatly improved thermal paths and multifunctional built-in thermal management

When a heat sink is used, thermal paste interfaces the device surface with the heat sink. Unfortunately, thermal paste is a poor thermal conductor so the primary heat flow path is hindered from the outset in a conventional design. In addition, the heat sink requires a significant airflow rate in order to have adequate pressure to force air into the narrow gaps between fins and counteract the associated pumping losses. This increases cost, size and complexity of the system.

In situations where only the PCB is used as a heat sink, the situation is even worse. The main heat removal path is still only in one direction: into the PCB because FETs in an SO8 form factor have a ~5 times lower thermal resistance to the PCB than to the the top of the plastic case of the devices itself. This method can only adequately reduce thermal resistance to ambient by requiring a large surface area around the module to be exposed to airflow cooling. This increases system size and complicates the layout of auxiliary components and circuits.

On the other hand, the thermally-improved Stacked Power method solves all these problems by requiring little air flow, reduces thermal resistances by more than 10x, and improves thermal conductivity of the substrate by 40x for true double-sided cooling. This module can be located in a cramped location without having to be derated and requires little additional board area for cooling thanks to its built-in, and largely self-sufficient, thermal management.

3.4 High-Frequency Layout Considerations

It is clear that one of the primary goals of integration should be high efficiency – from thermal as well as electrical aspects. Less loss directly translates to less heat and thus higher output capability – or smaller package size, both of which are beneficial. The use of electrical circuit simulation and experience in laying out high-frequency switching circuits in PCB have taught us there are several electrical design principles that need to be used to make an excellent high-efficiency converter.

One of them concerns the external decoupling capacitor at the input. For high circuit efficiency, it has to be physically located as close to the devices as possible to reduce the input loop parasitic inductance – and not necessarily close to the input source as commonly believed. In Stacked Power, the devices are oriented and placed so as to have extremely low parasitic loop inductance via the input decoupling cap because it is physically located right at the devices. In order to achieve this, it is necessary to flip the MOSFETs relative to each other so that the drain of the top switch

and the source of the bottom switch are on the same side of the board. Figure 3.8 shows this loop and 3.9 shows the impact that two slightly different layouts with different parasitic loop inductances has on the switch node waveforms

.

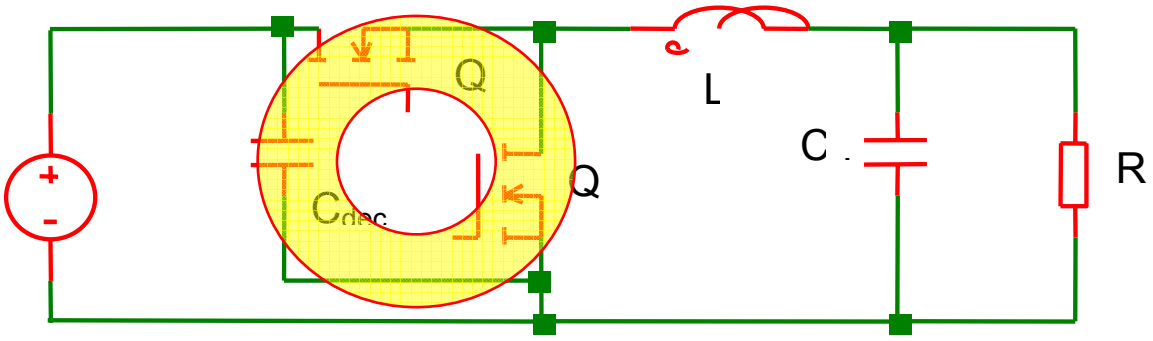

Figure 3.8. Input decoupling loop should be as small as possible.

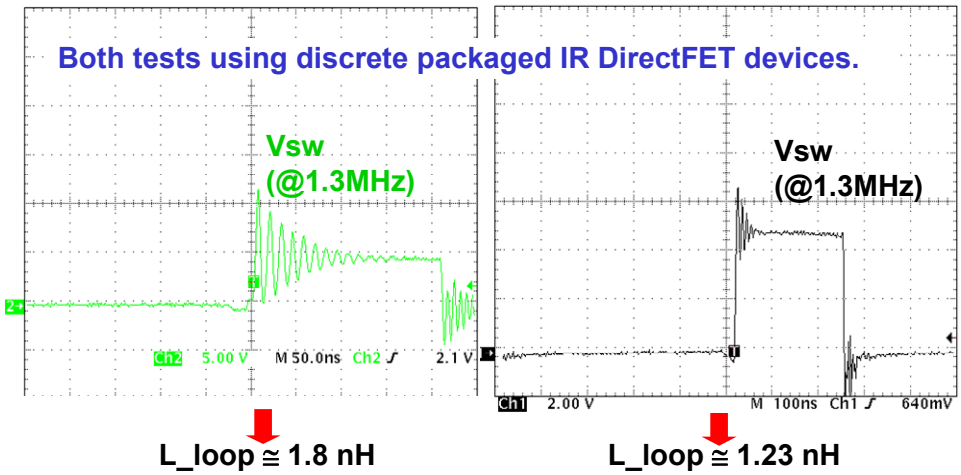

Figure 3.9. Effect of input decoupling capacitor on parasitic inductance. Comparisons made with two different PCB layouts using conventional parts.

A value over 1nH causes significant ringing at the switch node (where the two MOSFETs and inductor are connected to each other) as shown in Figure 3.9 for two different PCB layouts. Maxwell 3D simulations show that with Stacked Power and its

embedded devices, we can cut this value down by almost half, to 0.82 nH, as shown conceptually in Figure 3.10. So the physical location of this decoupling capacitor is one layout constraint that must be taken into account for high efficiency at high switching speeds. This information has recently been adopted for use in a commercial application for POL converter [117]. Figure 3.11 shows the measured waveform on the actual hardware with very little ringing present.

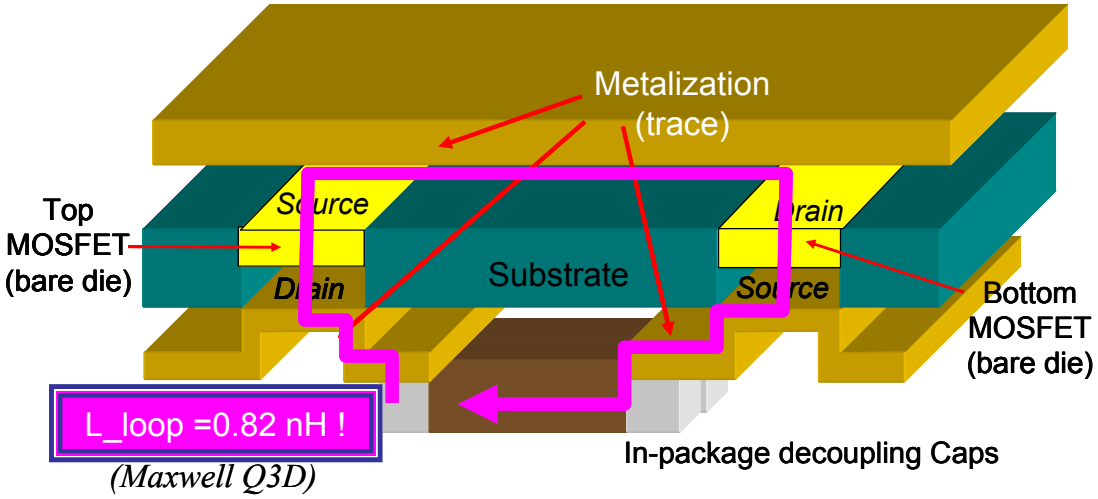

Figure 3.10. The input decoupling loop that must be minimized.

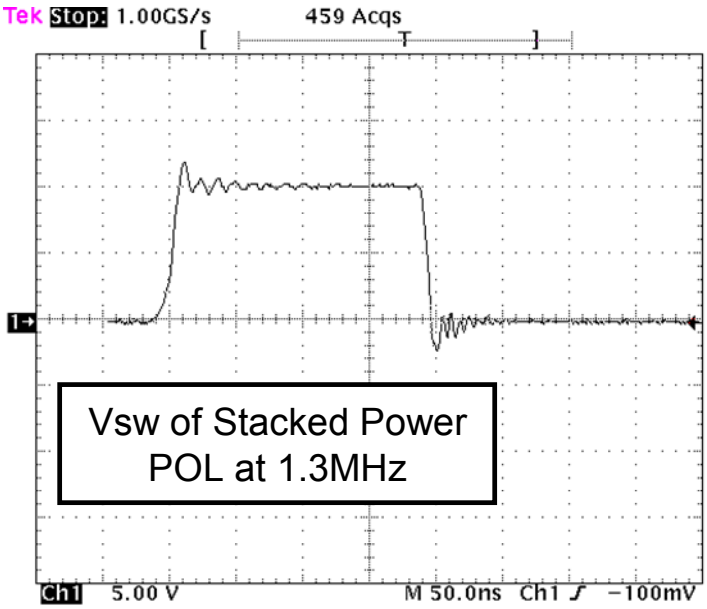

Figure 3.11. Vsw node voltage showing little ringing when the decoupling capacitor is on the switches.

Also, electrical simulations based on the schematic of Figure 3.12 tell us a snubber cap across the top switch will improve efficiency by about 1% at heavy load but that the capacitor tolerance must be very close to 4nF for optimal performance [43]. For the schematic of Figure 3.12, Figure 3.13 shows how the snubber slows down the Vds rise at turn off so the current has more time to decrease, thereby reducing loss. We integrated a 4nF snubber wafer capacitor for the top switch in our first generation layout. Figure 3.14 shows three Medici simulations run using physical models of the buck circuit shown in Figure 3.12 for three cases: without snubber, with a 4nF capacitor across the top switch, and with an 8nF capacitor across the top switch [43]. The 4nF case shows the greatest peak reduction and the 8nF shows that a tight tolerance on 4nF is necessary since 8nF has a 47% increase in peak current over the 4nF case.

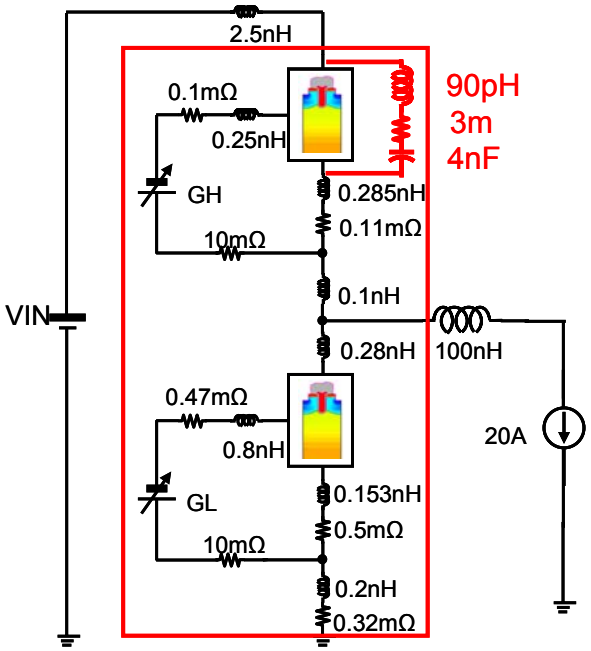

Figure 3.12. Simulation schematic with parasitics and snubber capacitor shown in red.

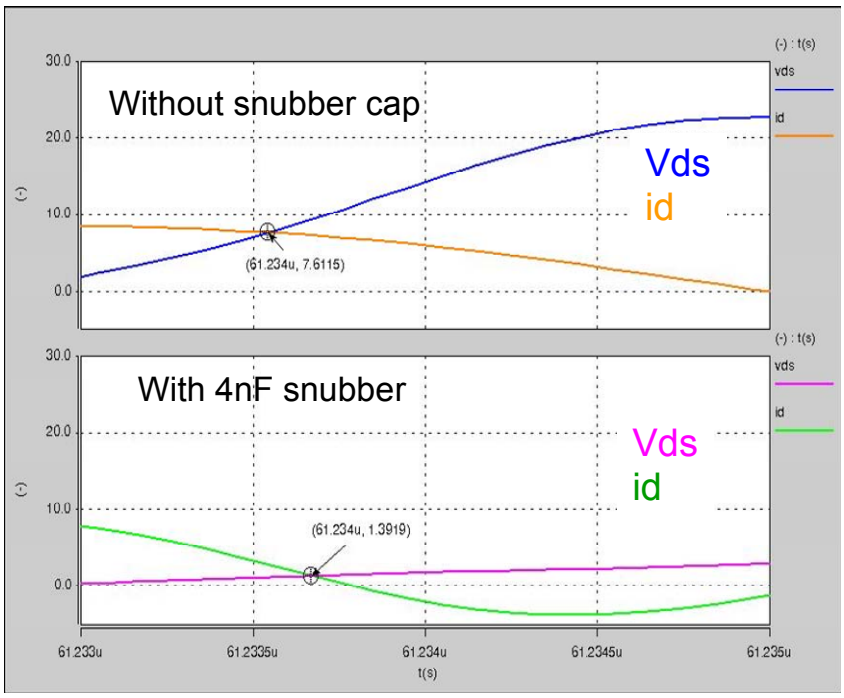

Figure 3.13. Simulation showing Vds rising slower at turn-off with snubber in place.

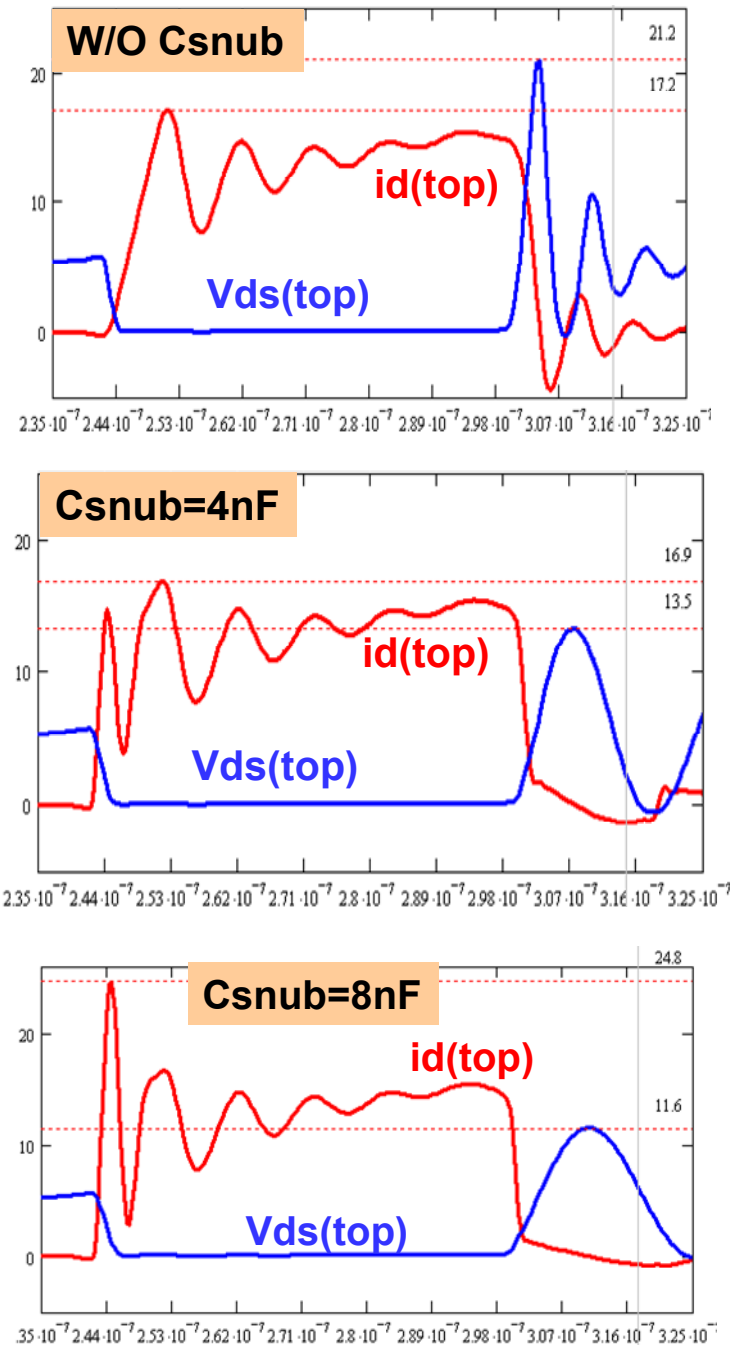

Figure 3.14. Three Medici simulation graphs showing impact of snubber capacitor in a buck circuit. Note that 4nF has greatest peak reduction for the circuit shown in Figure 3.11 [43].

However, in actual hardware testing, this improvement was not readily visible. Figure 3.15 shows actual efficiency measurements using 0603 form factor capacitors so that they may be placed immediately next to the top switch. This test was conducted on

a buck PCB board using Si7106 top switch and IRF6691 synchronous rectifier, operating at 1MHz and converting 6V-1.2V. Without a snubber cap in place, the converter achieved 20A at 87.0% efficiency. With a 4nF snubber (two stacked 2nF 0603 capacitors) capacitor, effiecient went to 87.2%, but this difference in on par with measurement error even though it was the highest efficiency test at full load.

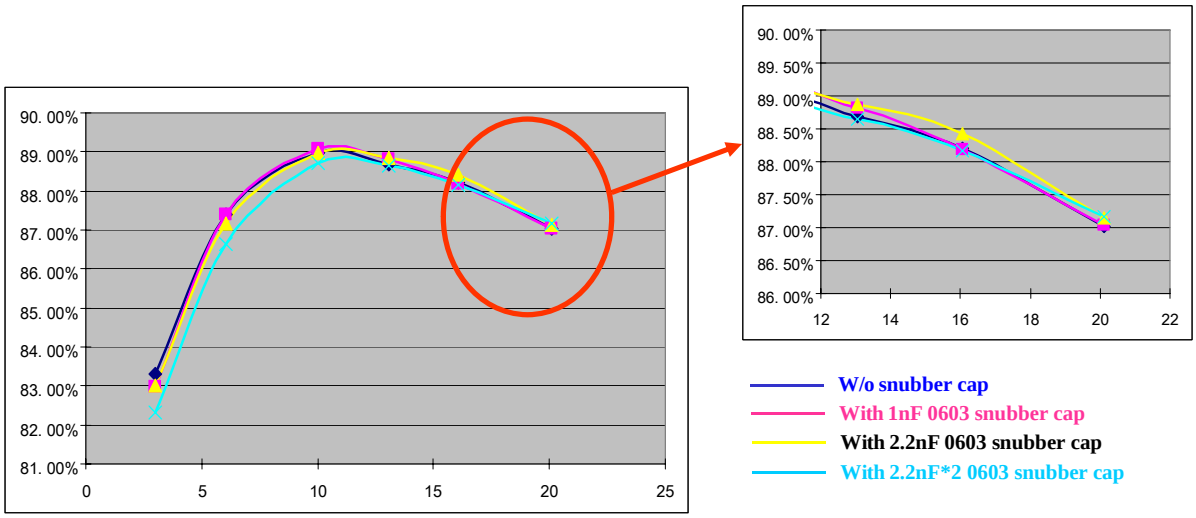

Figure 3.15. Hardware efficiency test with top switch snubber capacitors.

Furthermore, the MOSFET drivers have several requirements in order to deliver the highest performance. The latest drivers have the capability of driving extremely high di/dt and so are just as sensitive to circuit parasitics as the devices themselves. One of these drivers is the LM27222 which has superlative performance specifications and is used as the primary example here although any other high-speed driver will benefit from these good layout practices.

Figure 3.16 shows a loop in red, overlaid on the driver datasheet diagram, which must be kept small in order to drive the top switch as fast as possible. The large voltage conversion ratios required in today's applications mean the top switch has to

contend with very short on-times. In order to get the most from its short burst, the parasitic inductance of this loop much be kept to a minimum.

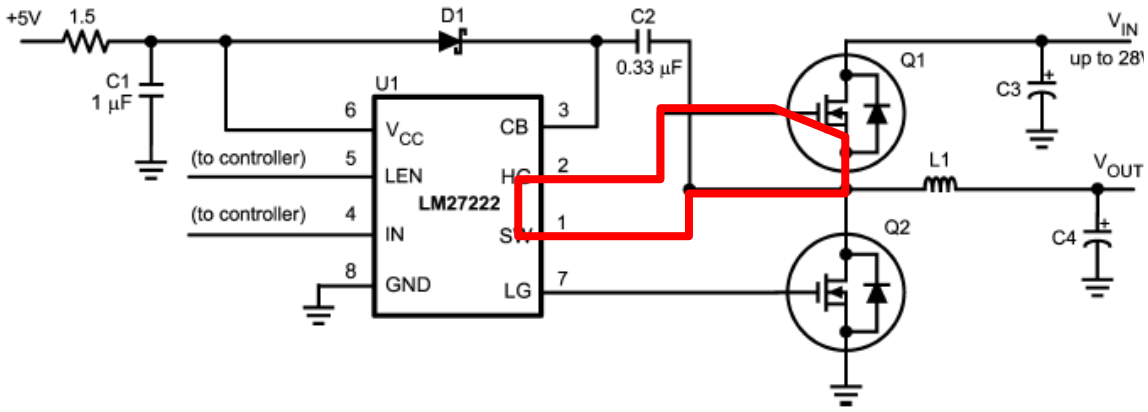

Figure 3.16. Top switch driver loop.

In particular, Figure 3.17 shows a portion of this gate drive loop that must be connected very close to the source of the top switch so that the parasitic inductance between the source and the switch node is out of the gate drive loop as much as possible. Otherwise, the high di/dt in the power path will cause an opposing voltage proportional to the inductance of the trace, reducing the available driver voltage if too much parasitic ends up in the gate drive path.

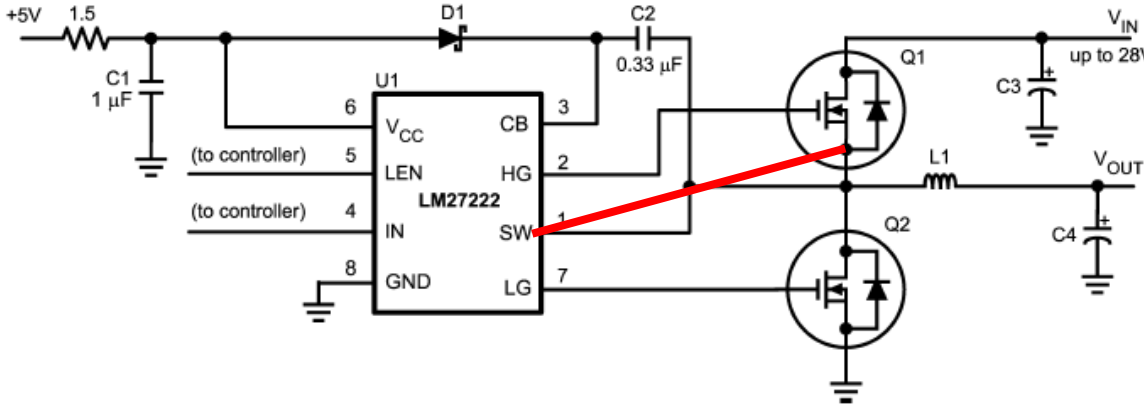

Figure 3.17. The most important path of the top switch driver loop.

The two previous layout requirements are necessary due to the short (~10%) on time of the top switch and the corresponding need for high di/dt capability to push switching frequencies. This leaves the bottom switch of the buck converter to do the lion's share current carrying capability since its duty ratio can often be a complimentary 90%. As a result, the high current path shown by the curved arrow in Figure 3.18 must have low DC resistance to minimize I^2R losses.

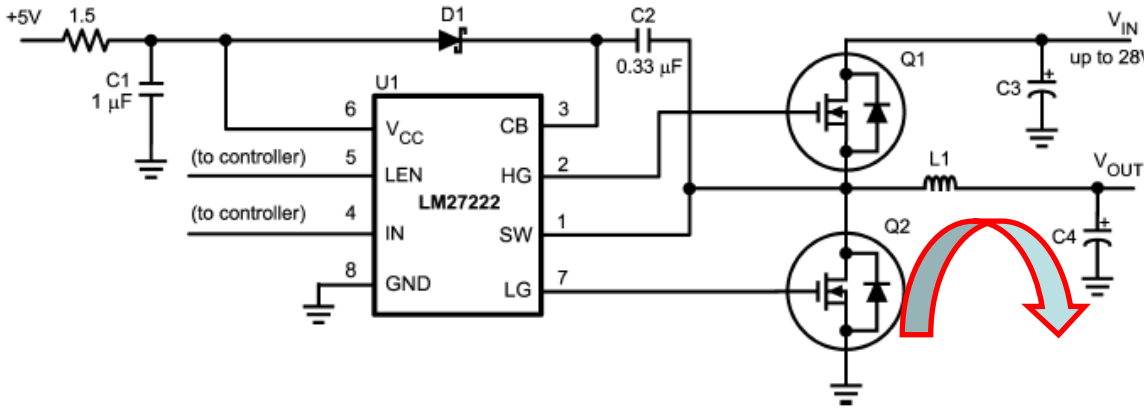

Figure 3.18. The high-current path that needs low resistance.

Other important layout factors have to deal with the current return paths. When laying out a high-frequency board, it is important to realize that any signal beyond 10kHz or so will return following the path of least *impedance* – and not the path of least resistance. Generally speaking, for a high di/dt or high frequency signal the path of least impedance will be directly below the signal trace. As a result, it is best to have a ground plane below the signal paths that remains continuous throughout. This is not always possible however and if there has to be a cut in the ground plane, the use of

"image rails" or wire jumpers will allow the return current to remain close to the signal trace, thereby reducing parasitic inductance of the path.

3.5 Active Layer Fabrication

This integration methodology places the devices side-by-side in one layer of ceramic. The devices are placed so as to have extremely low parasitic loop inductance via the input decoupling cap which is placed as closely to the devices as is possible, as described previously. This is determined by means of simulations using Ansoft's Maxwell Q3D FEA tool. Also, I-DEAS 3D FEA software is used to assess the thermal design before fabrication. Input from these simulated results is used to make a good electrical model using Synopsys' Saber Electrical Simulation software, complete with parasitic values for the traces and interconnections, to assess where the power losses are, and thus the thermal hotspots. Using this information, we can make a proper layout using Autodesk's AutoCAD software based on this information. These software simulations steps are graphically shown in Figure 3.19.

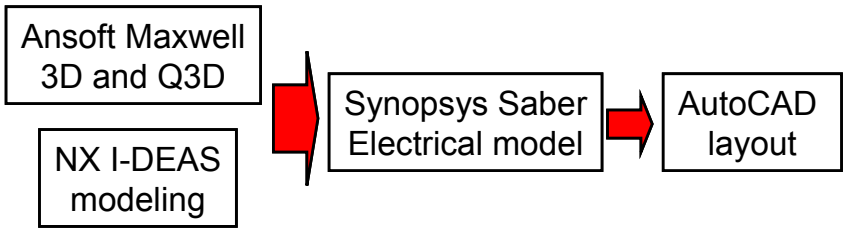

Figure 3.19. Simulation inputs used for determining best layout

A big thank you goes to David Gilham who was an instrumental part in refining the fabrication steps for Stacked Power. In order to help clarify the complex number of

steps involved, a flow chart is shown in Figure 3.20 to give a concise and simplified overview of what it takes to make a Stacked Power module. Detailed descriptions of each step follow in the subsequent pages.

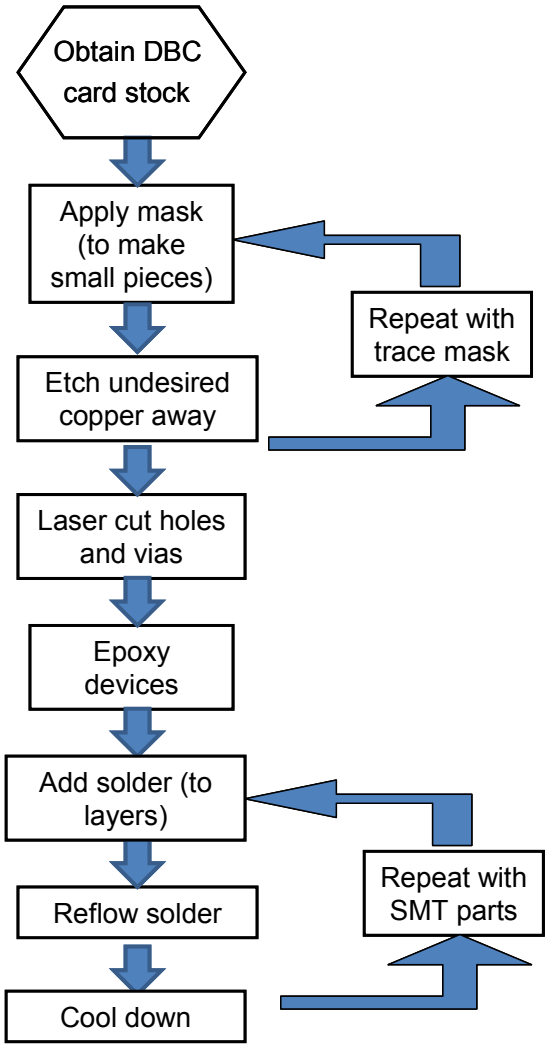

Figure 3.20. Generalized flow chart for Stacked Power fabrication.

To start off the process, the DBC must be cut into manageable pieces for manufacturing. The 8/15/8 mil AlN DBC used here comes typically in a 5”x7” sheet called a “card.” Out of this card, small pieces (in this case, the middle and bottom layers in this example are 18 x 18 mm) are cut using a laser cutter and wet etcher to

begin the fabrication process. The example used here is one of the more complex modules made for this work, wherein 3 layers of DBC are used. This complicated structure was chosen for this part so that all potential fabrication steps are shown. The middle layer is the one that houses the devices and the top and bottom pieces are used for extracting the heat from the dies, as well as shielding from the LTCC inductor that will be the bottom substrate for the entire module (the process for it is shown later). Figure 3.21 shows the process steps for cutting one half of the card into the small pieces necessary to put the module together.

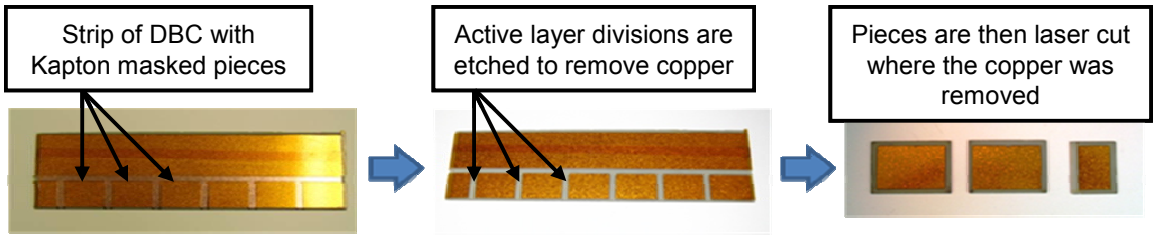

Figure 3.21. Cutting DBC stock for starting Stacked Power process.

The laser cutter used is the Resonetics Micromachining System G150 using a CO_2 Laser Router System and A3200 software, with a maximum output power of 150 W. It is shown in Figure 3.22. Cutting AlN is actually quite tricky because of its high thermal conductivity and brittleness requires less than maximum power and "low" cutting speed with "high" number of pluses to prevent cracking. The active layer fabrication yield has been on the order of 90% provided the layers do not get dropped on the floor, in which case cracking of the DBC where no copper is present (e.g., the corners) is often the result. The 10% failures were due to using incorrect laser cutting parameters where localized heating next to a copper trace caused cracking of 2

substrates out of approximately 20 boards made. The laser cutter values that ensure excellent yield are shown below:

For the 15 mil ceramic used here, it generally takes 4 to 5 laser passes to cut all the way through the material. After lots of experimenting, we found the best parameters to use are:

- Power Level = 80 W,
- Pulse Number = 600,
- Pulse Spacing = 0.0004 inch,
- Lasing Velocity = 0.1 inch/sec.

Figure 3.22. Resonetics Laser Cutter

The next step involves using the aforementioned AutoCAD mask file, which are loaded directly into the laser cutter for patterning of the Kapton tape mask on the DBC (Steps 1, 2, 3 in Figure 3.24). Once cut, the undesired-half of the Kapton mask is

removed so the remaining Kapton keeps the desired traces from being washed off the DBC during ferric chloride wet etching (Steps 3, 4, 5 in Figure 3.24). This etching is done using a Kepro Bench-Top Developer and Etcher, model BTD 201B shown in Figure 3.23. Etching takes about 10 minutes per side of the DBC (one complete turn of the timer) when laying on a suspended platter inside the etcher such that it is only exposed to the spray and not fully immersed in the ferric chloride. The etching resolution and limitation we can achieve with this process, for 8 mil copper cladding, is 0.3 mm between features for adequate mask removal and 0.3 mm minimum feature size (e.g., trace width). Any narrower traces will be unetched excessively for adequate trace continuity.

Figure 3.23. Kepro wet etcher used to remove copper from the DBC.

Once out of the etcher, only the masked copper is left on the AlN ceramic. Then we return to the laser cutter so that we can cut the holes in the DBC to embed the dies for the middle layer (Step 7 Figure 3.24), and cut the outer edges off the top and bottom

layers (Steps 9 and 10 Figure 3.24). All the steps up to this point in the fabrication are shown in Figure 3.24 for each of the three layers, in vertical arrangement with the process flow going downwards. The middle layer requires the most process steps because of the embedding of the devices, which the other two layers do not have. Also, note that not only do the extra layers extract heat from the dies, they also provide complementary trace patterning on the mating sides of the DBC layers to handle the necessary circuit interconnections to bridge the epoxy that surrounds the devices. This makes multi-functional use of this expensive material. A single-layer DBC module consists of only fabricating what is called here the “middle” layer and then the interconnections are made with copper straps.

When cutting the pieces with the laser cutter, it is important to pay attention to registration so that alignment of the traces will be correct. If the pieces are not properly aligned with respect to each other, then traces could short out internally and lead to failure during electrical testing. The procedure for this is two fold: First, in the AutoCAD file, corner features for the layers must be kept the same on both sides of each piece. This aligns the top and bottom copper layers of each piece of DBC to each other and is easily accomplished using stacked wireframe views in AutoCAD. Then, these same features are used to align the layers of DBC to each other, because the via holes reference these features. The vias holes are 1 mm in diameter for a pin diameter of 0.8 mm. This 0.2 mm margin is sufficient tolerance for alignment across the 25 mm diagonal distance between opposite via holes. The addition of the pins ensures the pieces go together properly. For the case of the little top piece of DBC, wireframe

feature alignment using AutoCAD is sufficient since the size of the piece is small, and therefore error is small as well.

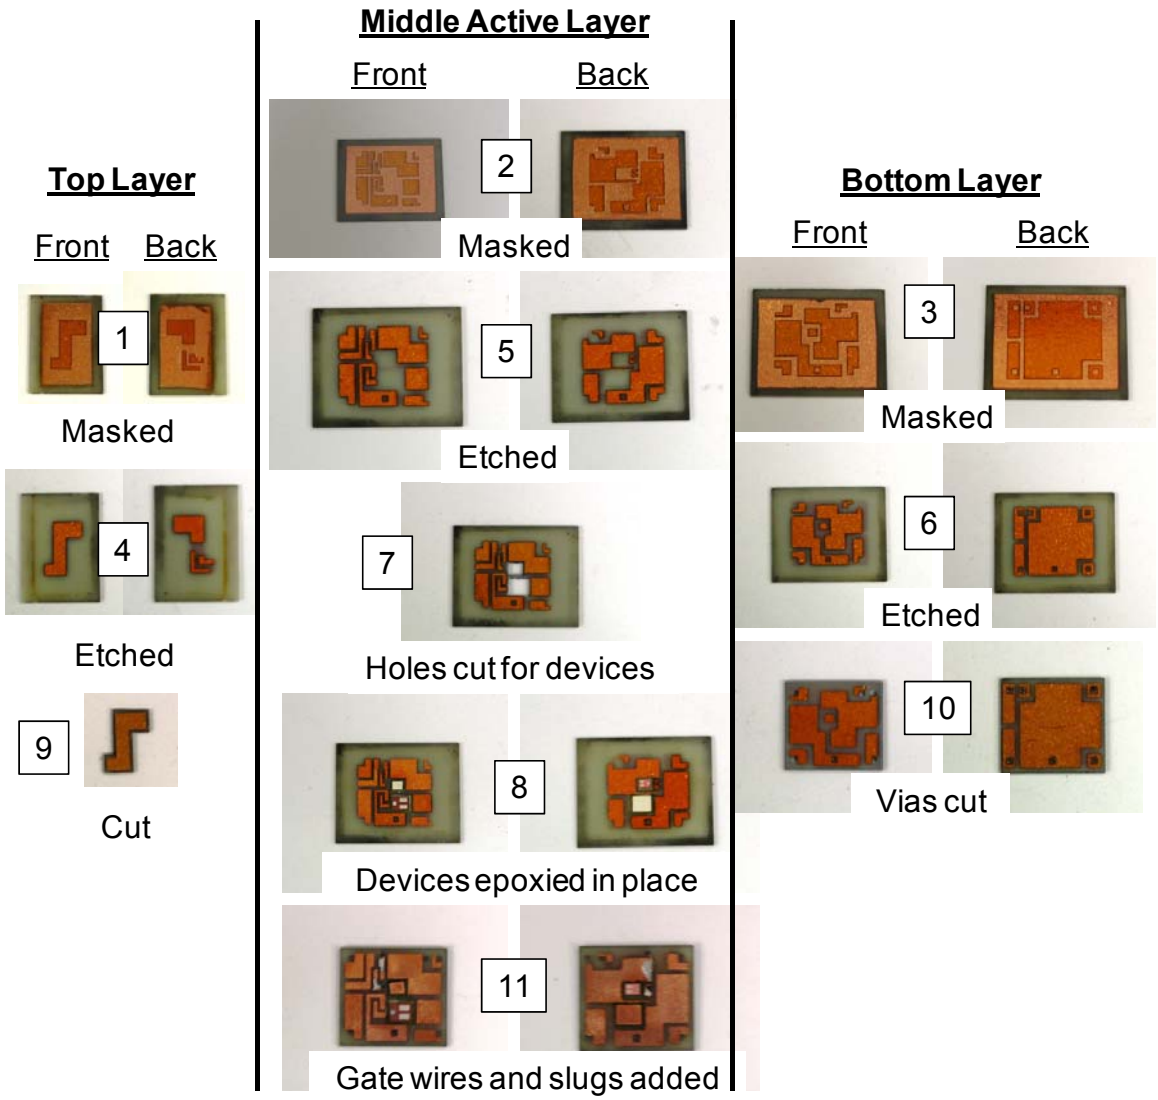

Figure 3.24. The three DBC layer process pictures, each one showing both sides of each layer side-by-side, and process steps follow systematically.

The bare dies are then expoxied in place using EP3HT from Masterbond (Step 8 in Figure 3.24). The dies used are unpackaged IRF6633 for the top switch and IRF6691 for the bottom switch, but the Vishay PolarPaks could also be used. The dies are 2 x 2.6 mm (78 x 102 mils) and 3.4 x 4 mm (133 x 159 mils) respectively and the

gap around the devices is 0.4 mm (15 mils). In order to do mount the dies, we apply a layer of Kapton tape on the entire piece of DBC to cover the die holes, turn it over, and stick the die with the gate/source side into the hole, against the tape. The stickiness of the Kapton tape holds the dies in place and maintains their coplanarity so that there is no tilt of the die relative to the DBC surface. This ensures the die is flat and level with the surface copper on the gate/source side of the FET. The epoxy is applied around the FET and is kept in place by the Kapton tape.

The epoxy application process to fill the gap around the FET is done using a pressurized syringe to dispense the epoxy in a uniform way with constant flow. The epoxy comes available directly from Masterbond already in a syringe just for this type of application. The Kapton tape is left in place during curing to keep the epoxy in the well around the FET. The curing cycle is 150°C for 7 minutes with a warm-up rate of 1°C/min. Cool down is passively done in the oven with the door closed for at least 1 hour. It all takes place in a computer-controlled oven, which is shown in Figure 3.25. When the board has completed its cool down cycle, the Kapton tape is easily removed, leaving a die that is coplanar with the surface copper, epoxied in place so that it cannot fall out or shift during solder reflow. This guarantees a proper mating surface for the additional stacked layers/interconnects.

Figure 3.25. Programmable oven for curing the epoxy. Shown with the front door open (it opens upwards).

The DirectFET dies are only 10 mil thick but the ceramic plus the thickness of the traces totals 31 mils (8+15+8) so there is a 21 mil hole that needs to be filled to make both sides of the die flush with the DBC surface traces, as shown in Figure 3.26. Then cut out two copper slugs from a 8 mil sheet of copper, the same size as the device drain, and layer them on top of the drain-side with low-temperature 43Pb/43Sn/14Bi solder paste, which is also dispensed using a pressurized syringe, in between each layer. Each of the two layers of solder represents 2.5 mil thickness after reflow. So the 10 mil die along with two 8 mil copper slugs and two layers of solder builds up the die to be flush with the surrounding DBC copper traces (Step 11 in Figure 3.24). Then 32 gauge via wires are fed through the vias for gate interconnections through 0.3 mm holes and soldered in place. A schematic cross-section showing the dies, slugs, copper and ceramic is shown in Figure 3.26.

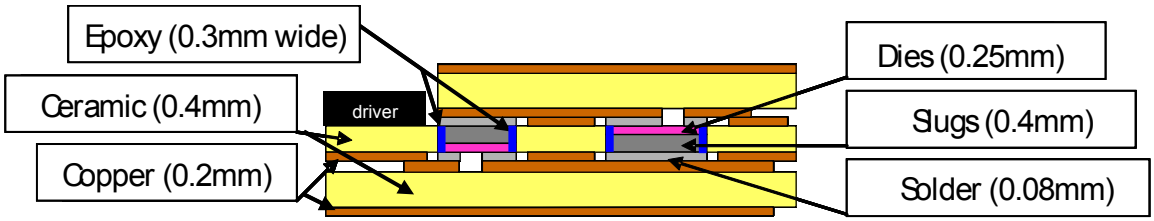

Figure 3.26. Schematic of embedded dies (in pink), heat slugs (dark gray) and 3 DBC layers.

Now that the dies are in place and each layer is cut to size, the layers can be assembled together. As stated before, the top and bottom layers of DBC are made with the interconnection patterns to connect the devices to their traces, and are "stacked" on top and bottom of the active layer. These extra layers not only act as interconnections but also as heat extractors, heat spreaders, efficient convection surfaces and electrical shields. As such, they represent the integrated thermal management that will eliminate the need for an external, bulky, expensive heat sink.

Once again, the same low-temperature 43Pb/43Sn/14Bi solder that was used to attach the slugs to the dies is applied to the mating traces between each layer. The bottom layer is laid down first, the solder is applied to its top, followed by the middle layer put down on it, apply solder to its top traces, and finally the top piece is put on the middle layer. Only a thin coating is necessary as excessive application of the solder can result in traces getting shorted between layers during the reflow process, which leads to obviously complicated problems. Also, solder should be applied down the middle of each trace so as to reduce the risk of voids forming. The solder will naturally remain humped up in the middle and when a DBC layer is put on it, it will spread out evenly. These final steps are shown in Figure 3.27.

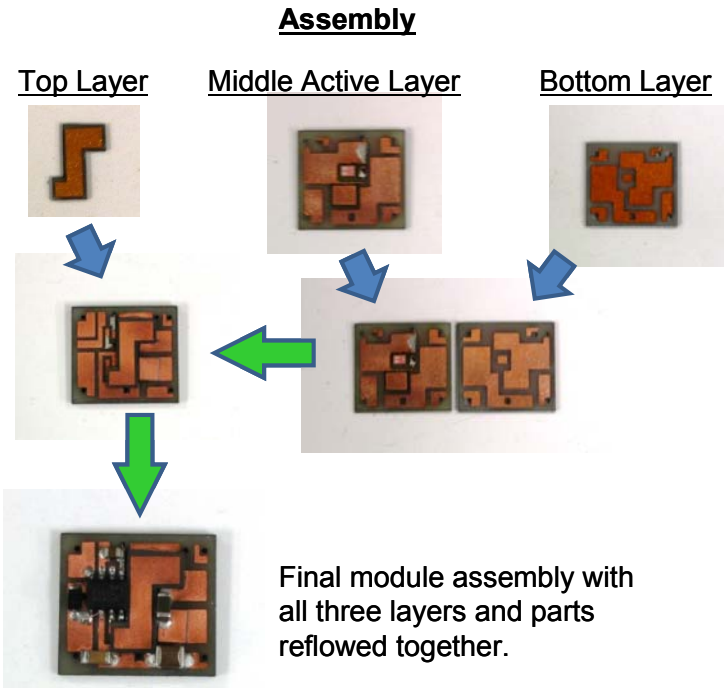

Figure 3.27. Final assembly procedure for the layers and SMT parts.

We use a precision reflow machine to solder the layers together, as well as discrete surface-mount parts such as the driver (1 x LM27222), bootstrap capacitor (1 x 0.33uF MLCC) and diode (1 x 16V B3K), and input (1 x 22uF MLCC) and output capacitors (1 x 100uF MLCC). Again, we use low temperature 43Pb/43Sn/14Bi solder because it is easy to reflow and eliminates the risk of overheating the bare dies during soldering since its liquidus temperature is 163°C. This is an important issue to keep in mind because the AlN DBC substrate has a much higher thermal conductivity than conventional PCB, and the interconnections to the devices are of very low thermal resistance, so heat will transmit to the devices much more than with a PCB application. This also means that homologous temperatures (operating temperature over melting temperature) are well beyond 0.5, but this is typical for many applications today in which

case >0.8 has been observed with no significant reliability issues and 0.87 considered the limit for SnPb eutectic [116], which is right where Stacked Power operates with the low-temperature solder. Under these conditions, both creep and fatigue process exist and operate interactively. Solders containing 14% bismuth have a higher creep rate due to the low-moderate creep resistance but their high tensile strength and low temperature performance offset creep with good fatigue properties [116]. However, solders with higher quantities of bismuth are not recommended due to their high brittleness [116]. Future work needs to be done to determine actual reliability using low-temperature Bi solder.

Figure 3.28 shows the Sikama Reflow machine that is used and the five platters that can be individually set to the proper temperature in accordance with the EFD Inc. 43NCLR silver paste profile [109] recommended for this paste, as shown in Figure 3.29. The 5 platter temperatures are: 70°C, 90°C, 120°C, 160°C, 195°C, followed by a cool down on a 25°C ambient resting platter. The rate of speed is set to 20 in/min but a pause on the 195°C is necessary when more than 2 layers are reflowed at one time. The pause varies depending on the number of components but 20 seconds is typical in order to keep voids at a minimum [108]. A visual inspection of the board will make it clear when the solder has fully melted because the components will suddenly settle down and then they will stop moving equally suddenly.

Figure 3.28. Sikama reflow machine.

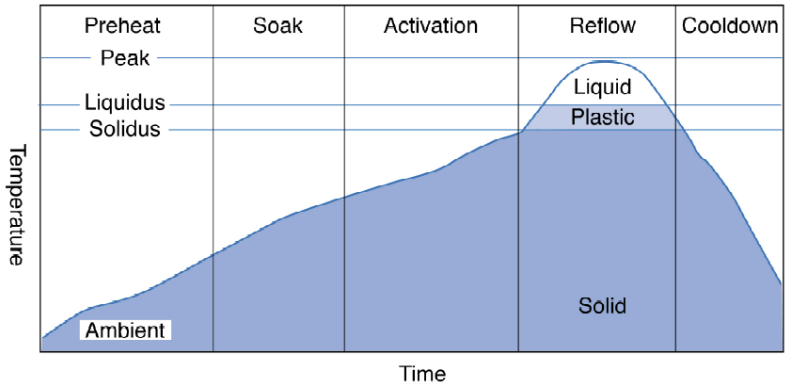

Figure 3.29. Reflow profile for the 43Pb/43Sn/14Bi solder: Solidus, liquidus, and peak temperatures are 144°C, 163°C, and 193°C, respectively [109].

According to International Rectifier, 40% voiding in the DirectFET soldering process is acceptable, and 30% passes all qualification tests, but that in nearly all cases, 20% is typical so long as the reflow profile for the selected solder was adhered to. Figure 3.30 shows a top-view X-ray photo of a DirectFET with 18% voiding on th left and an equally acceptable 3° side-view tilt on the right [107]. In fabrication of Stacked Power modules, tilt was not witnessed because of the cured epoxy holding them in place and voiding is assumed to be 20% or less, since IR's and the solder paste manufactuer's recommendations were closely followed. The JEDEC voiding specification is 25% [107]. However, as it pertains to Stacked Power, this is an area

where more work needs to be done in the future. The dies are put down using a die bonder as shown in Figure 3.31.

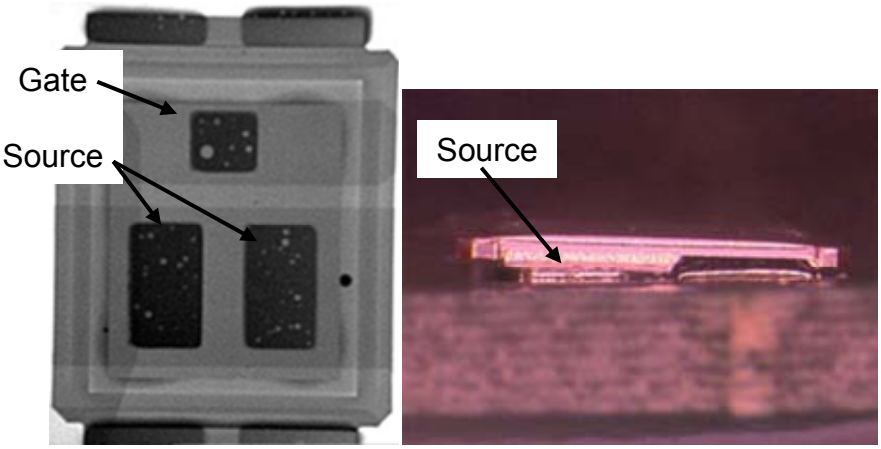

Figure 3.30. IR's acceptable DirectFET voiding and tilt photos.

Load stresses on the layers are rather uniform since DBC is constained on both sides by the same material: copper. In the three layer stack shown in detail in Figure 3.24, the middle layer is also constrained on both sides, though the top side only has a small piece of DBC whereas the bottom side has a full layer. This may lead to thermal cycling issues that have yet to be determined. This will be part of future work as well.

Figure 3.31. Semiconductor Equipment Corp. Model 860 Eagle Omni Bonder.

Figure 3.32 shows what the layout looks like with the snubber capacitor for the top switch in place. It is physically located 0.3mm from the top switch to ensure lowest possible interconnection inductance. With the snubber capacitor, we were able to obtain about a 0.3% improvement at heavy load but at light load we paid a penalty of over 1% so subsequent generations were dispensed of the snubber capacitor so that light load efficiency could be maximized and the circuit made more compact [103].

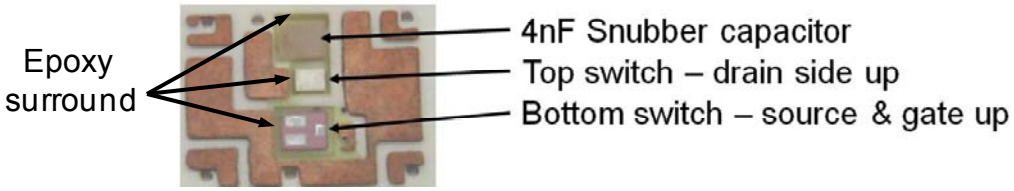

Figure 3.32. Integrated snubber capacitor wafer along with the two switches.

One of the main benefits of using unpackaged dies and copper strap interconnections is the thermal resistances from devices to board (i.e., interconnection) are extremely low. The top switch is smaller than the bottom switch, as seen above, and thus it has higher resistances because of its decreased area, but nonetheless, they are all still as low as possible considering the pad sizes are fixed by the device manufacturer. Table 3.2 shows the calculated values which will be used as input for thermal simulation analysis shown later in the chapter. The conductivity values are of solder for source connections and a weighted average of copper slug and solder for the drains. The gate pad connections are neglected since the source will handle more heat and has much larger surface area in comparison.

The equations used to calculate the resistances of the IRF6619 (Top) and IRF6691 (bottom) were shown earlier as Equation (9) for the individual thermal conductivity of each material (solder and slug layers), combined once again using

Equations (4) and (5). The two layers of 8 mil copper slugs along with two layers of 2.5 mil solder (0.08 mm each, after reflow), as described for Figure 3.33, are combined using Equation (1) and yield the drain conductivity of 200 W/m°C shown in Table 2.1. The solder thickness was obtained by soldering 4 DirectFETs onto a layer of DBC and measuring the thickness before and after adding and reflowing solder. The average thickness change for the four samples was 0.08 mm. The source conductivities are simply a solder interface so this calculation is obtained directly from Equation (9). The corresponding thicknesses in millimeters for Table 3.1 are shown in Figure 3.33.

Table 3.1
Thermal resistance values for DirectFET interconnections

Device side	**Length**	**Area**	**Conductivity**	**Resistance**
Top – source/gate	0.08 mm	1.41 mm^2	51 W/m°C	1.11 °C/W
Top – drain	0.8 mm	4.5 mm^2	200 W/m°C	0.89 °C/W
Bot – source/gate	0.08 mm	2.81 mm^2	51 W/m°C	0.56 °C/W
Bot – drain	0.8 mm	10.8 mm^2	200 W/m°C	0.37 °C/W

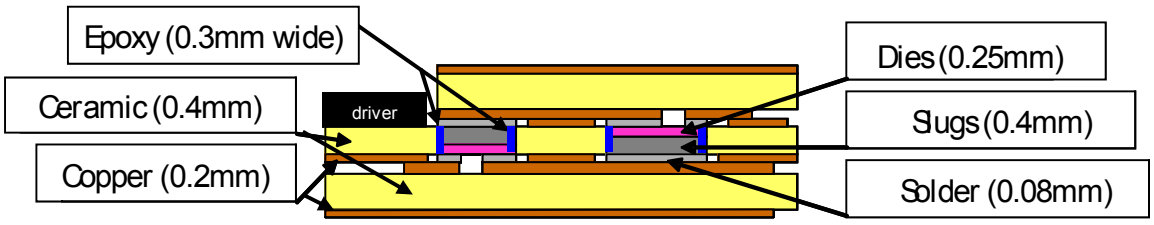

Figure 3.33. Schematic of embedded dies (in pink), heat slugs (dark gray) and 3 DBC layers with their associated thicknesses in mm

Figure 3.34 shows one of the FETs soldered to the DBC board. The alumina used for this test is 1.1 mm thick (43.25 mils) and the IRF6619 DirectFET is 0.34 mm

thick (13.5 mils), for a total of 1.44 mm (56.8 mils). After soldering, the total thickness was 1.52 mm (60 mils) which yields a thickness of 0.08 mm (3.25 mils) for the solder. All four samples were within 4% of this value, demonstrating very good repeatability.

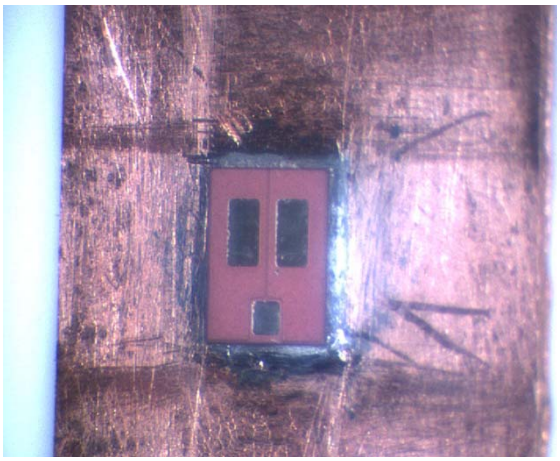

Figure 3.34. A DirectFET soldered to alumina DBC for solder thickness test.

Figure 3.35 shows comparisons between three Renesas packaging techniques [110], going from SO8 packaging with wirebonds, to LFPAK with copper straps, and finally combining devices and driver in a wirebond package, each with less packaging related loss and lower parasitics as the packaging method advances. The next logical step would be to have the devices integrated in a package like the Dr.MOS but using copper straps instead of wirebonds – and this is precisely where Stacked Power's advantages become clear.

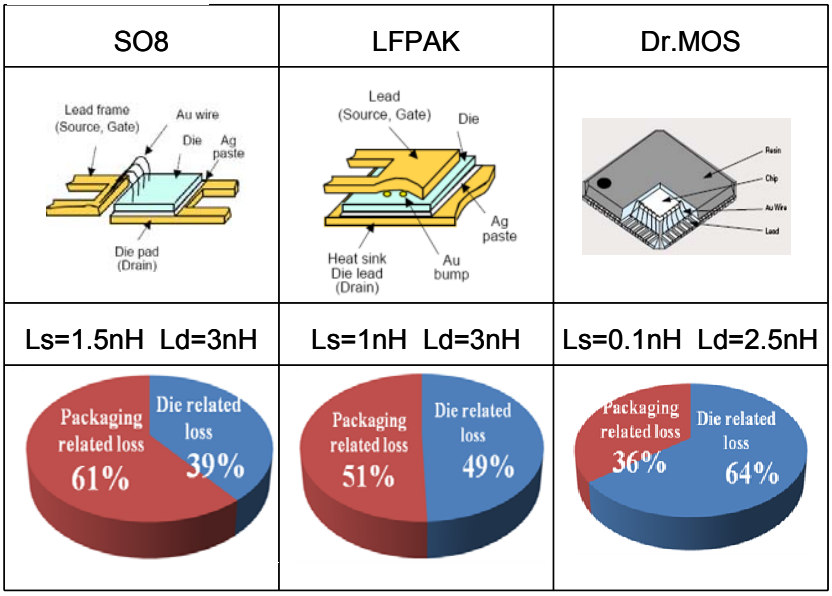

Figure 3.35. Different packaging methods for devices [110]

Comparing the device thermal resistance values using Stacked Power with conventional device package/implementation methods shows the improvement gained by having integrated dies with no packaging and the use of copper strap interconnections to interconnect dies to substrate. Figure 3.36 shows the typical device packages found today for 20-30V switching MOSFETs, all of which are "thermally-enhanced" in comparison to the classic SO-8 package.

Figure 3.36. Range of typical packaged FETs for buck applications. SO-8, copperstrap, PowerPak, and DirectFET shown respectively.

Table 3.2 compares the device package characteristics for each, including thermal resistances, of the typical MOSFET packages [111]. It is clear that the

DirectFET has the best characteristics and it has achieved this status due to the elimination wirebonds that were the main source of package resistance, and the elimination of plastic encapsulant that was the main source of thermal limitations. Stacked Power took these concepts and extended them to a converter-wide solution, which yields the same benefits when compared with typical PCB-based converters and/or wirebond technology used in hybrid modules today.

Table 3.2.
Package characteristics for each packaged FET type [111]

Package type	R_{DS}on of package (mΩ)	Source Inductance (nH)	Rth junction_pcb mounted (°C/W)	Rth junction_case top (°C/W)
SO-8	1.6	1.5	11	18
CopperStrap	1	0.8	10	15
PowerPak	0.8	0.8	3	10
DirectFET	0.15	<0.1	1	1.4

Figure 3.37 shows the complete first generation active layer using a single layer of DBC. Yield for this process is very good after a couple initial attempts. Excluding the first two prototypes which were used to determine all the process parameters just discussed, the following 6 made functioned perfectly without any need for reworking, which suggests excellent repeatibility and process control. This is one of the big advantages of Stacked Power over Embedded Power. The converter runs open loop for now and uses the LM27222 driver from National Semiconductor. All capacitors are MLCC type. The final size is 18 x 18 x 5 mm. The copper strap interconnections for a single layer version module are visible here [103].

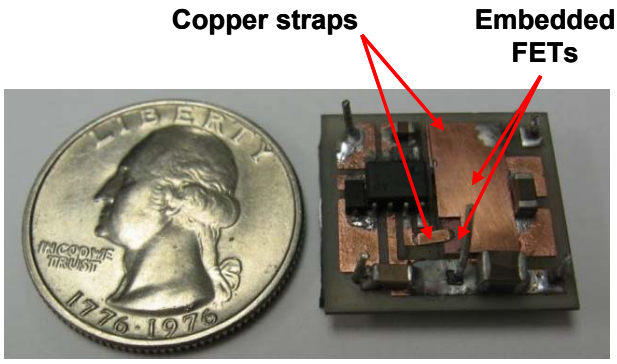

Figure 3.37. Generation 1 active layer.

3.6 Passive Layer Design and Fabrication

With miniaturization and improvements in device performance, the development of power electronics systems has progressed to a state where active devices' impact on the system's size and cost has been taken over by passive components. Magnetic components have been one of the bulkiest components which take up significant circuit surface real estate. With the increase in popularity in low profile electronics, development in planar magnetics has taken up an increasingly important role in today's power electronics research and development. In view of the numerous techniques available for inductor integration, the integration techniques have been classified as below:

1. On-silicon (chip level integration)
2. Organic substrate (package level integration)

3. Thick film core and winding (package level integration)
4. Integratable winding with commercial core (semi-discrete level)

For chip level integration, the inductor / transformer is built on the silicon die. The processes are compatible with silicon processing. For package level integration, the magnetic component is built on a substrate where the active device can be mounted on. Package level integration can be classified into two categories. One is integrating the inductor / transformer in organic substrate where the processing temperature is low (<300 °C), another is integrating the component using thick film technology where the processing temperature for the component is high (> 600 °C). The mechanical properties of the organic substrate are usually not compatible with the silicon die, i.e. FR4 has a coefficient of thermal expansion (CTE) of 17 ppm / °C. For thick film technology, the materials involved can support higher operating temperature and have a CTE between 4 to 7 ppm / °C, which is a good range because these values are close to that of silicon.

For the passive inductor layer, LTCC technology is used. This work was done by fellow collegue Michele Lim. The inductor and the shield are integrated to form a low profile substrate to be integrated with the active layer described above. As opposed to using the usual screen printing technique predominantly used in LTCC technology, a technique which is more suitable for making inductors which can support higher current is developed [44,45], following the initial work reported in [46]. This technique also proves how to produce components with lower losses compared to using the screen printing technique [47]. Due to the non-linear behavior of the inductor, improvement in

light load efficiency can be achieved [48-49], and the extent of improvement can be controlled by varying the geometry of the inductor.

A model of its non-linear behavior has been developed empirically [44-46, 128] and Equation (26) is shown for an infinitely wide core. The first half of the equation is the definition of the relative permeability of the material because it is non-linear and changes depending on the level of bias current. Hence the formula is based on bias current and was empirically obtained. Initial attempts were made using LTCC ferrite tape with a permeability of 50 at zero bias so this formula was derived specifically for this material.

The right half of the equation represents geometrical relationships based on the structure shown in Figure 3.38, whose values are used for equation inputs. The equation uses mks units. The result is 1.77 uH/m at 16A load for a 2.5 mm trace width (the value of “w” in Figure 3.38). Corners in the trace are determined to have half the inductance of the unit length which total to an additional 8.8nH for the 4 corners [46] as shown in Figure 3.39. The terminations are included in the trace length as shown in Figure 3.40. The vertical portion of the termination for each connection to the DBC active layer is neglected since it is only 1 mm for each, and outside of the ferrite core. Based on these calculations, Figures 3.39 and 3.40 shows the dimensions of the inductor design.

Permeability is defined as a function of bias current

Geometrical term for designing size

$$\frac{L}{l} = 10^{4w(I_{DC}+8)+0.0366(46.38-I_{DC})} * \frac{\mu_0}{2\pi} \ln\left(\frac{\frac{w+e}{2}+2g+\sqrt{\frac{w^2+e^2}{2}+4g^2+2g(w+e)}}{\frac{w+e}{2}+\sqrt{\frac{w^2+e^2}{2}}}\right) \quad (26)$$

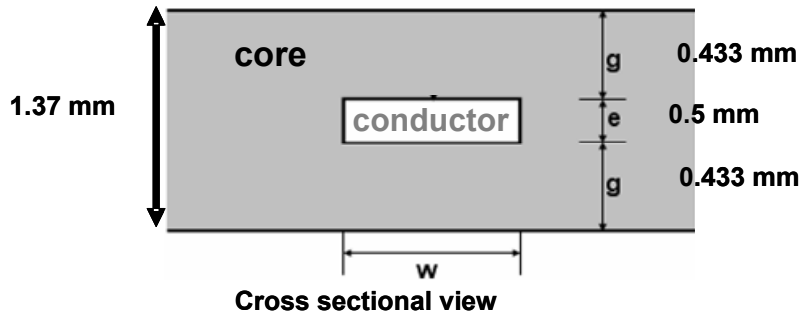

Figure 3.38. The structure for using the inductance formula, assuming infinitely wide core [46].

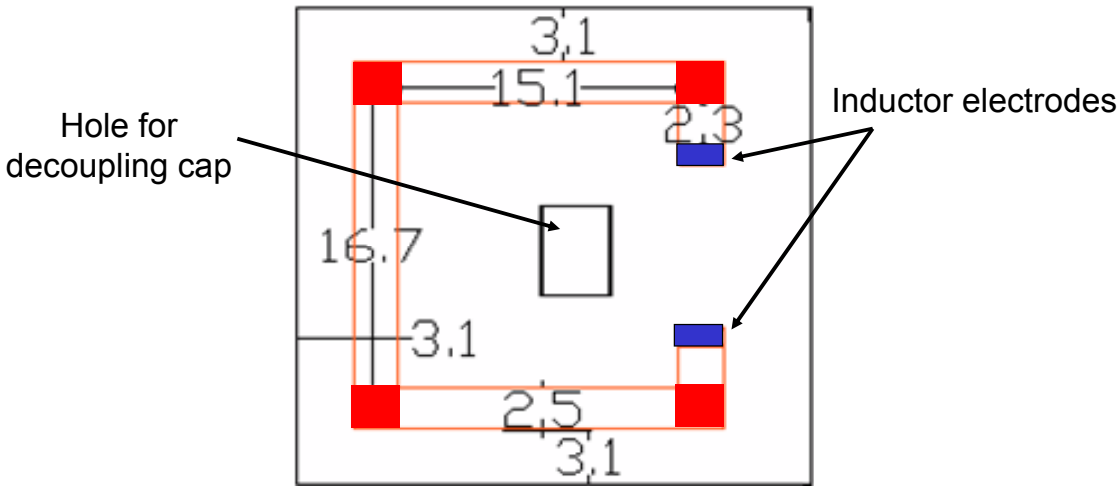

Figure 3.39. The lengths of each trace segment and the 4 corners identified [46].

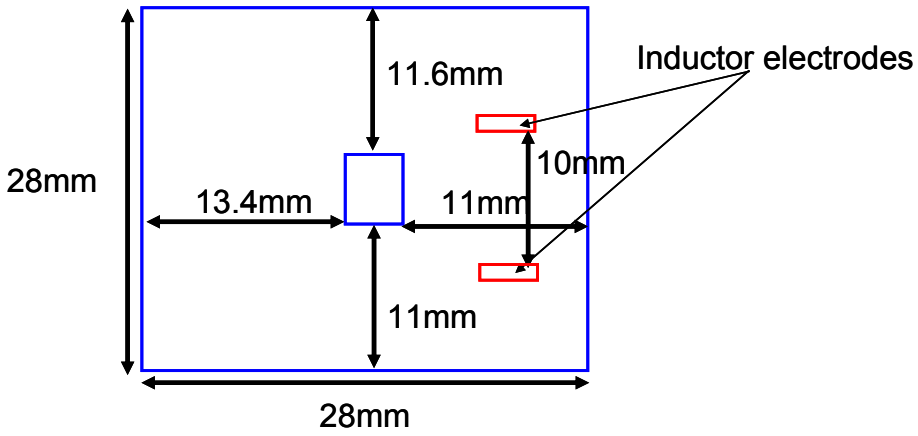

Figure 3.40. Inductor dimensions based on the inductance formula [46].

In order to combat the negative effect of using a magnetic material as a substrate [47, 127-129], the inclusion of a metal shield below the conductor traces is introduced. This shield is made here using the same silver paste as the conductor and is then covered by a layer of blue dielectric, as shown in Figure 3.41. The shield prevents the ferrite from adversely affecting the signals in the active layer located directly against its top side.

A 100 nH planar inductor with an integrated metal shield is designed and fabricated [47, 127, 129]. Figure 3.41 shows 4 photographs of the inductor at different stages of fabrication. The first picture is the embedded inductor winding prior to covering it with the top ferrite tape layer, second is with the top layer added, third is the planar inductor with integrated shield and finally the end product with blue insulator layer. Figure 3.42 shows the active layer combined with this prototype inductor and Figure 3.43 shows the voltage and current waveforms at 16A load.

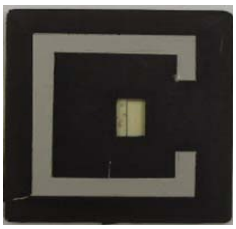

Inductor winding

Inductor top layer

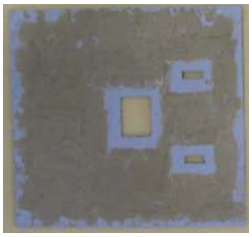

Insulator with shield

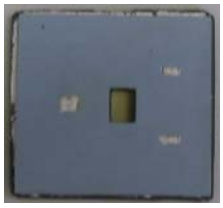

Inductor with shield

Figure 3.41. Our first LTCC inductor showing the internal trace, with top layer added, with shield paste, and final inductor [46].

The final results for the inductor values are as follows. The assumptions made for the corners and terminations, as well as the calculation of trace length, were experimentally verified to be within 6% accuracy in sum total at 16A load [46]. Inductor core loss for LTCC material has been described by [115] although it has not been applied to a practical inductor. As will be discussed in the next section, based on the system efficiency and taking device losses and inductor conduction losses into account, the loss discrepancy is approximately 1W (83% at 20A with 1.46W loss for the top switch, 1.41W for the bottom switch, and 0.85W for the inductor DCR). This 1W accounts for miscellaneous trace parasitic losses as well as core loss.

- Full load current = 16A
- Full load actual L = 105.5 nH
- Full load designed L = 99.8 nH
- No load actual L = 231.4 nH

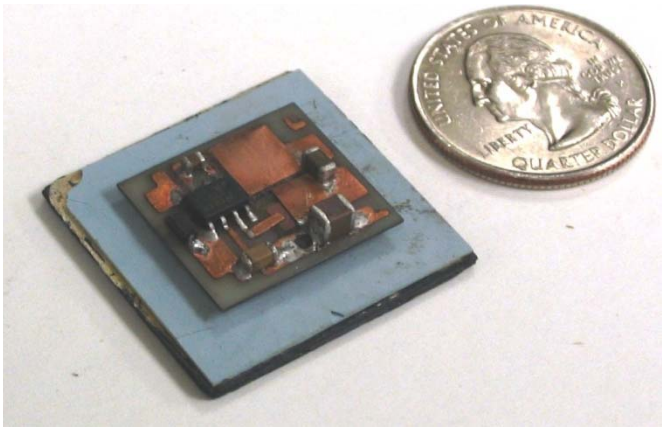

Figure 3.42. Generation 1 inductor with generation 1 active layer. The inductor acts as the substrate but is much larger than active layer [103].

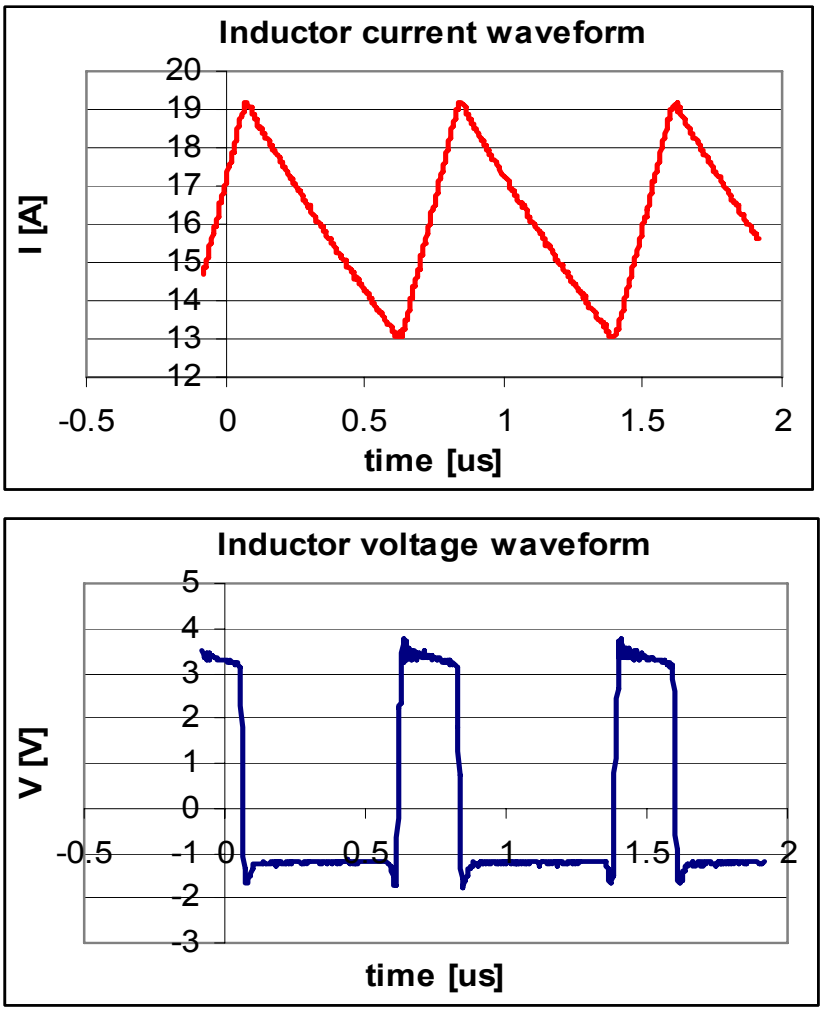

Figure 3.43. Inductor current and voltage waveforms at 16A load [46].

The inductance is non-linear based on current level, as discussed before. This is graphically shown in Figure 3.44 for various frequencies. It is clear that the inductor is able to handle high frequencies without any issues since the inductance curves essentially overlap. Note also that at light load, we obtain approximately 5 times more inductance than at heavy load. This is due to the nature of the LTCC material and the fact the inductor has no air gap.

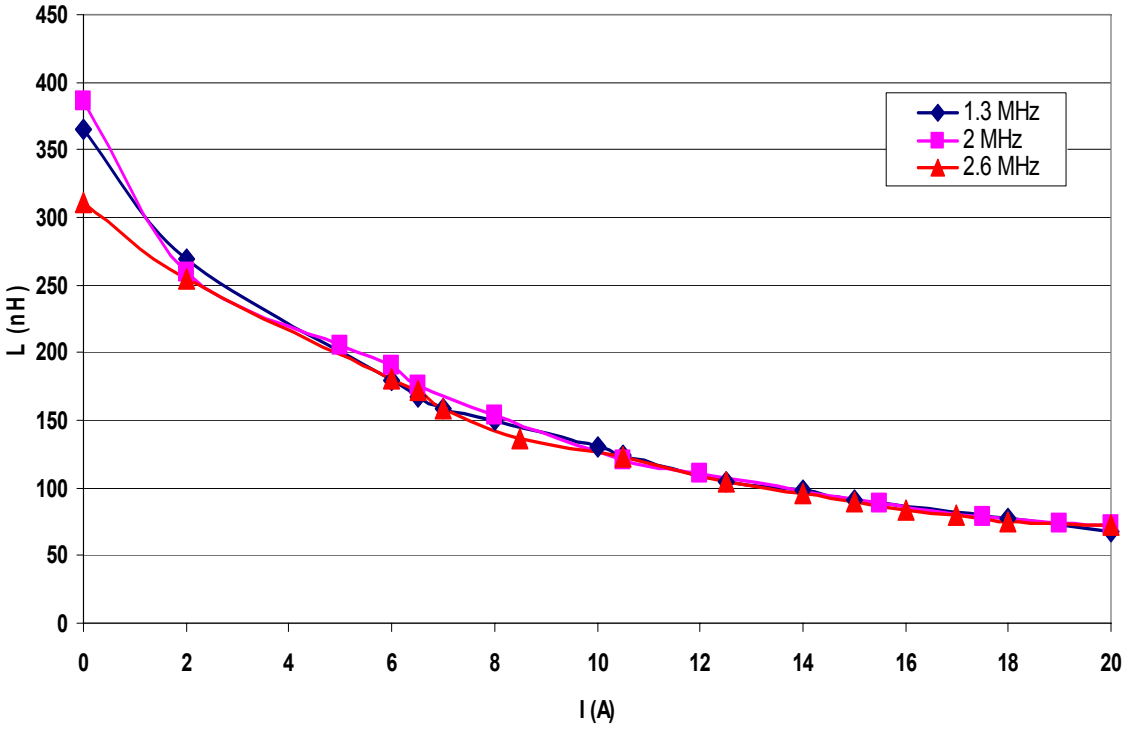

Figure 3.44. Graph of inductance relative to load current for various frequencies.

It must be noted however that this design method does not take material saturation into account. In the following section, a method using FEA simulations is shown that allows determination of the design's saturation point. So this method is best used as a starting point for the LTCC inductor design process, which is then followed by the more comprehensive FEA methodology described in the next section.

3.7 Stacked Power Test Results

Figure 3.45 compares the converter efficiency using a first generation 100 nH LTCC inductor on a first generation Stacked Power active layer, as well as a 100 nH discrete off-the-shelf inductor on a 6 layer PCB board using the same parts in conventional form.

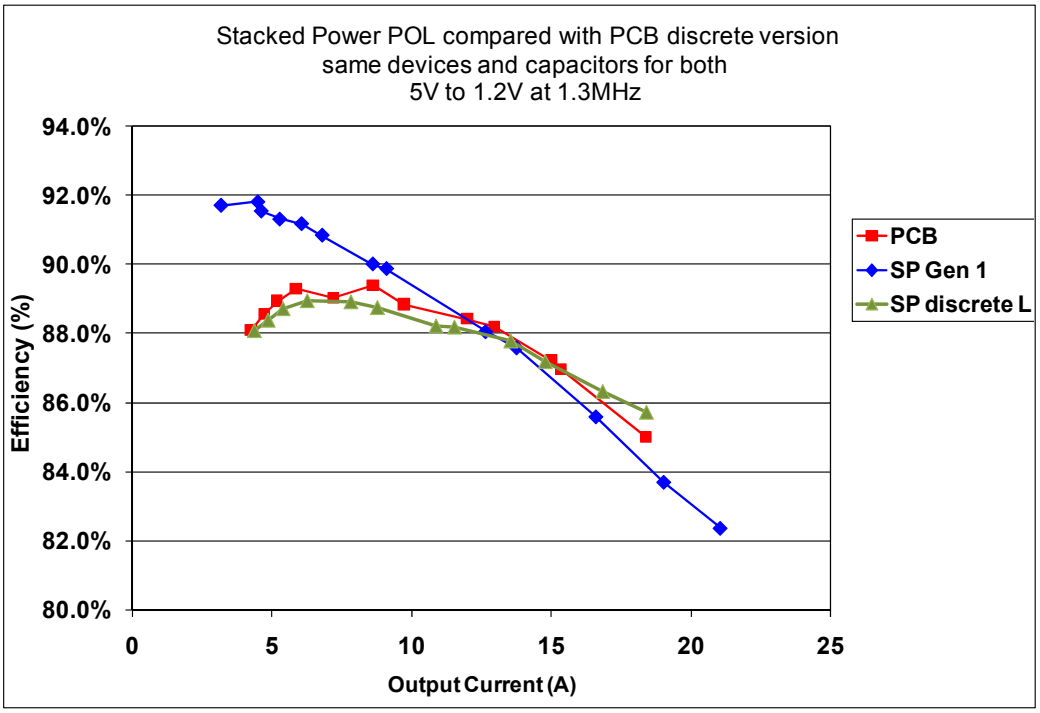

Figure 3.45. Generation 1 Stacked Power module compared with a discrete PCB converter using the same parts + commercial inductor

We see that light load efficiency rose significantly but at heavy load, we paid a penalty. After some analysis, it was determined that the culprit is a relatively high DC

resistance of about 2 mΩ. This compared to the 0.2 mΩ DCR of the commercial ferrite inductor yields a difference at heavy load.

In addition, the regular PCB discrete version has nearly the same efficiency up to 15A as the active layer using a discrete conventional ferrite inductor. This shows there is a lack of optimization in the layout. Using unpackaged dies embedded in the substrate should lead to efficiency improvements and yet these were not realized with this first attempt. The Stacked Power substrate already has a clear advantage over PCB in terms of thermal performance and layering capability (discussed in detail in Chapter 4) but that does not mean that an efficiency improvement below 15A should go untapped.

Therefore we are left with two main topics for improvement:

1. What must we do to make inductor improvements at heavy load?
2. What must be done to realize the potential improvements afforded by the capability of having a compact active stage with short signal paths?

These are our topics for the next chapter.

Chapter 4. Stacked Power Generations

4.1 Loss analysis of Generation 1

The performance of the first generation Stacked Power module shown in Figure 3.33 is excellent at light load due to the high inductance at that current level – approximately double what is available with a discrete commercial inductor with the same heavy-load inductance, which is to say 200nH. However the heavy-load efficiency is low due to DCR loss, as evidenced by the slope of the curve above 12A output current.

An efficiency test with the same layout but missing only the snubber capacitor was also performed to assess the impact of the capacitor. It was determined via simulation that it could give us a 1% efficiency improvement at 15A but this was not realized in practice with no real difference in the efficiency curve at that load current. However, there was a light load penalty of approximately 0.5% below 7A so the snubber cap was removed from subsequent generations. It appears the problem is that tolerance for a 4nF capacitor had to be too strict for the improvement to be realized in reality. If the actual value deviated by more then ± 50%, the advantages would be lost.

Taking a look at the input voltage shown in Figure 4.1, it is clear that there is more ringing and noise than we would like. The input decoupling capacitor is right on top of the devices so the problem is noise due to the top FET switching action. This is showing up because the Vin point is quite far from the capacitor due to the input being on top of the board but the capacitor is on the back side. Going through a via for the

input is not as short a path as it could be so in the future generations, this noise will be reduced and shown later in this chapter.

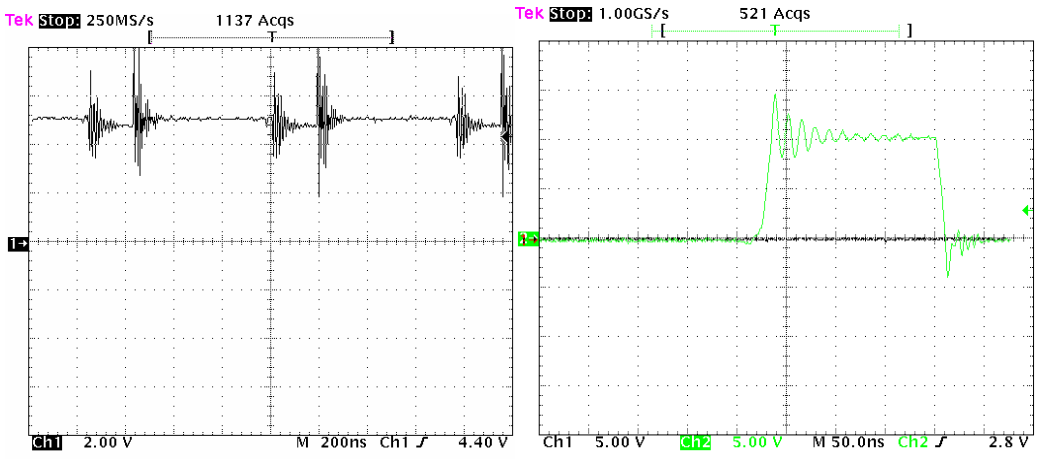

Figure 4.1. Graph of input voltage for the converter and Vds of the top switch with onboard capacitors

The gate drive signals are another place to look for parasitic issues. Figure 4.2 shows the waveforms for both switches, with the top switch waveform being taken with a differential probe since it does not reference ground. It is clear that the top switch turns on below the threshold voltage of the falling bottom switch, which is at the right moment. The ringing in the top switch implies some low-frequency parasitic which is also due to the large input loop because the power supply input is too far from the top switch. This issue will also be addressed in the redesign.

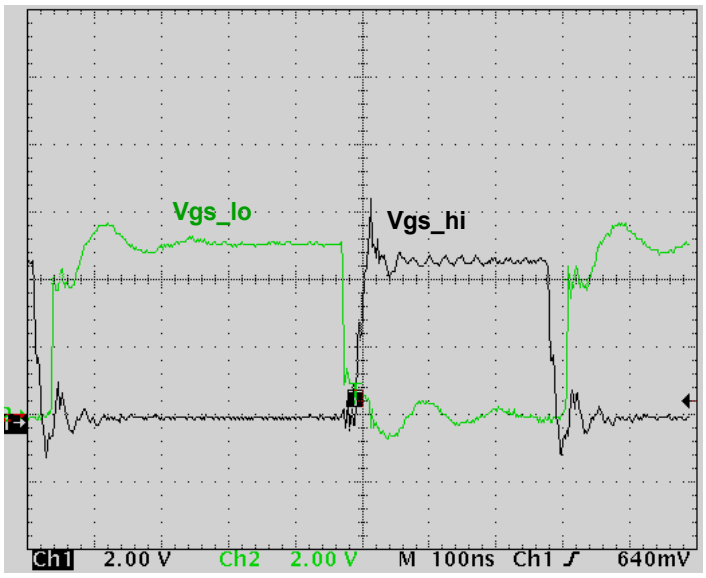

Figure 4.2. Gate drive waveforms for both switches.

The amount of heat loss at 20A and 83% efficiency is 4.9W, which is substantial for a module to handle, especially one of such small size as this one. Thermal performance was not very good as a result, with temperatures climbing towards the maximum allowable 125°C of device junctions in 200LFM airflow. Running without a heat sink in natural convection is out of the question above 14A for this module. In order to reach our thermal performance goals, higher heavy load efficiency is required.

4.2 Redesign Methodology

At this point it was clear the active layer can stand a few improvements but in particular, more work is needed on the inductor. The revised version needs to be smaller and more efficient. To help achieve this size reduction without a large penalty in

inductance value, we went from using a ferrite tape with a μr of 50 to one of 200. We also tried 500 but ran into saturation problems at high current so kept with 200.

Another necessary improvement would be to reduce the DC resistance of the silver paste conductor. In order to achieve this, we decided to go with a paste from Dupont that has 2.5x lower resistance. The goal was to reduce the size of the inductor, thereby reducing trace length, which would give us a correspondingly lower DC resistance than Generation 1. Commercial gapped inductors have an extremely low resistance, on the order of 0.2mΩ. But the smaller size would mean we need to have wide traces to keep the inductance from dropping too much. In order to assess how much reduction was possible, we resorted to Maxwell 3D FEA simulations.

The model used in Maxwell is shown in Figure 4.3. Along with the shorter trace length compared to Generation 1, the trace width was increased from 2.5 mm to 6 mm to reduce DCR (see Tables 4.3 and 4.4). A narrower trace curve at the bottom of the "U" shape of 4 mm gave us a longer path length in order to still reach 75 nH at full load of 16A. The thickness of the inductor was increased by 20% (to 1.75 mm) in order to keep losses manageable and the inductance sufficient. The Maxwell simulations allowed us to predict and verify whether or not this design structure could be made smaller and whether it would saturate or not.

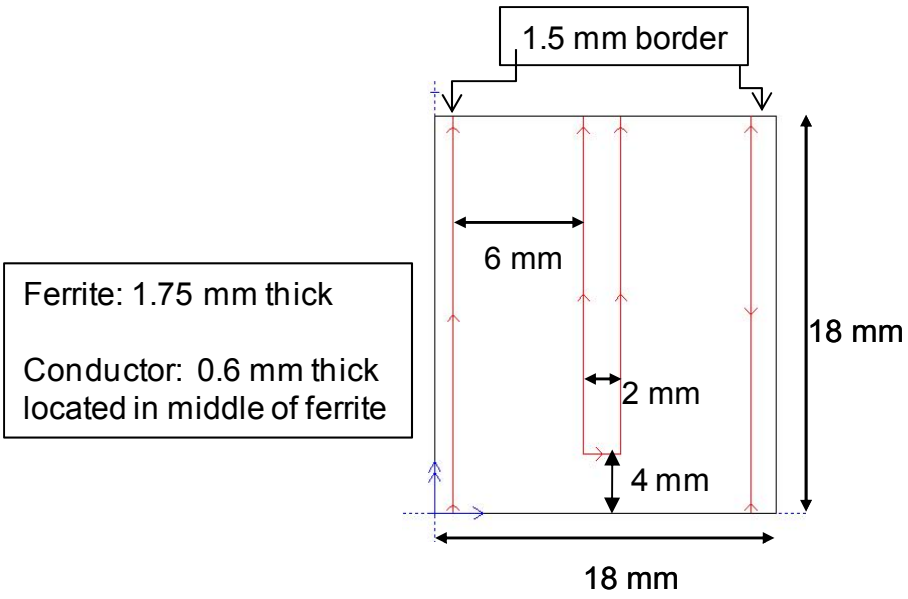

Figure 4.3. Trace geometry for Generation 2 inductor, as used in Maxwell simulation.

Maxwell 3D also allowed us to forego the tedious empirical data mining necessary to develop an inductance-per-unit-length formulation that was done for the u50 material shown previously. Instead, we could input the B-H curve of the u200 material and the simulator would solve all the design issues for us. Nevertheless, the u50 material equations are a good starting point for the inductor design that can then be followed up with FEA simulations for refinement/improvement.

The non-linear B-H curve was empirically obtained by fellow CPES student Nan Ying by using a very thin LTCC toroid connected to a sophisticated impedance analyzer shown in Figure 4.4 [112, 119]. Many discrete measurements were taken at 2 MHz to obtain the inductance versus current curve shown in Figure 4.5.

Figure 4.4. LTCC 1-turn inductor in buck circuit for characterization.

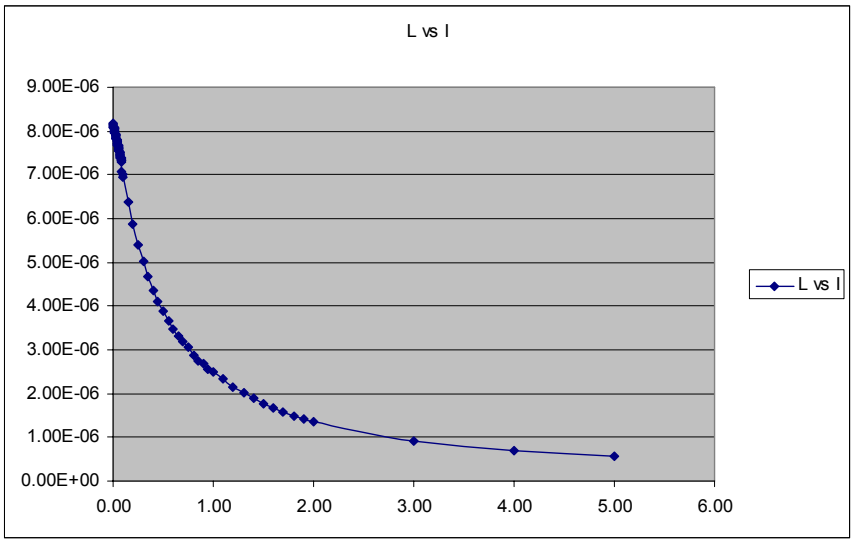

Figure 4.5. Inductance (H) versus current (A) for 1-turn LTCC inductor.

What is desired is the B-H loop so these two values must be found for each data point. Theoretical formulas are used with the understanding that nonlinearities and nonuniformities in the real components will introduce some error, but this is assumed sufficiently small to still yield a valid estimate.

First, we need to find Φ, which can be solved for using the inductances found by measurement in Figure 4.5 and Equation (20).

$$L = N\frac{d\phi}{di} \qquad (20)$$

Then Φ is used to solve for B based on the respective radii of the core.

$$\phi = \oint_A B \cdot dA = B(r2 - r1)t \qquad (21)$$

Now all that is left to be found is H which readily obtained from the current data points.

$$H = \frac{N \cdot i(t)}{2\pi r} \qquad (22)$$

The results are shown in the left-hand curve of Figure 4.6 [119]. This was a big advantage since we could now resort to magnetostatic simulations based on this non-linear curve to obtain accurate design data.

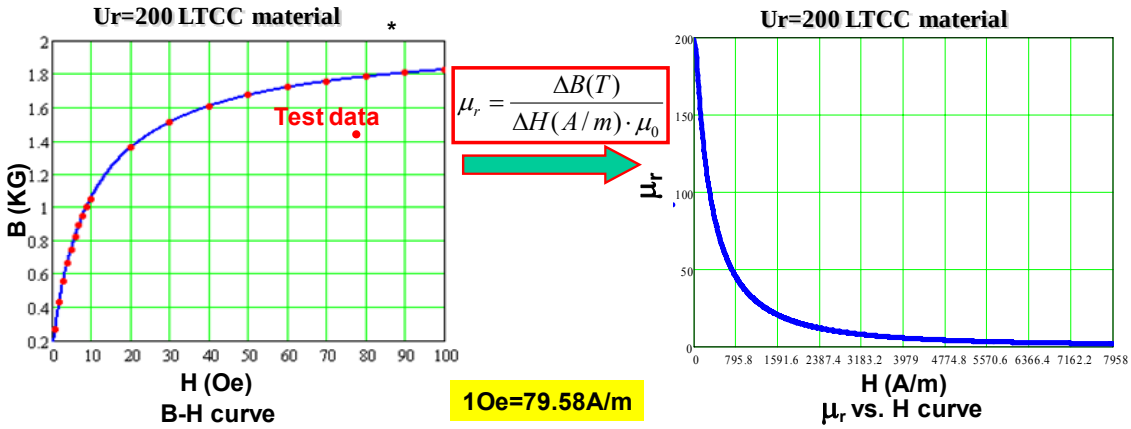

Figure 4.6. Left side is empirically-derived B-H curve and right side is Ur vs. H curve obtained by calculation with Equations 20-22 shown [119].

Maxwell 3D simulation conditions are as follows: The bias condition was 20ADC in order to ensure that output currents higher than 16A would still be within saturation

limits. The simulation sheet currents were defined as entering one side of the inductor and exiting the other, as shown in Figure 4.7, with electrodes that are the 6 mm width of the trace, by 1.5 mm long, and 0.6 mm tall [119]. The remainder of the dimensions needed are shown in Figure 4.7. The inductor was meshed using the automesh feature with the additional precaution of setting the seed size to 0.1 mm for the core. The target error was set at 0.1%. The analysis tool used is the Magnetostatic with results set to display an inductance matrix and flux density mapping. Refer to Section 5.2 of the Maxwell 3D v11 User Guide where the example shown follows the same steps as used in this simulation [132].

The u200 material has a saturation flux density of 0.3T and as we can see from Figure 4.7, 0.2T is the closest we get to it so there will be no saturation problems at full load using LTCC material with μr of 200 at zero bias current. Figure 4.8 shows the cross-sectional view of the actual inductor hardware once fabricated using the same method already described in Section 3.6.

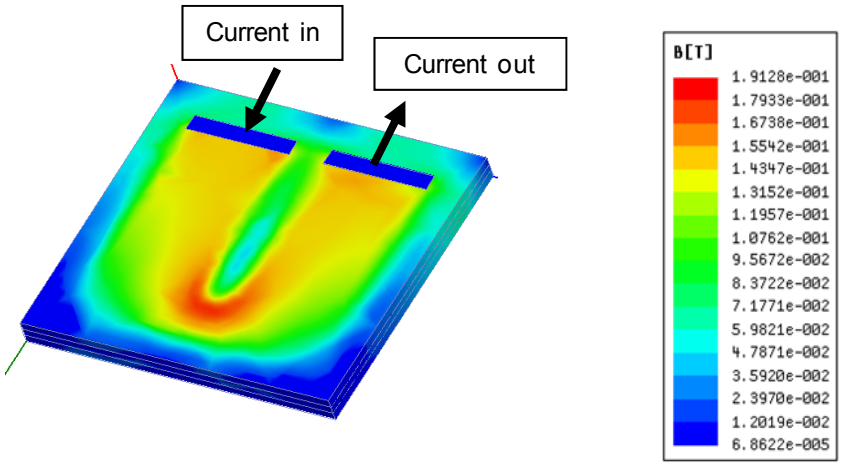

Figure 4.7. Maxwell 3D plot of flux density for the generation 2 inductor.

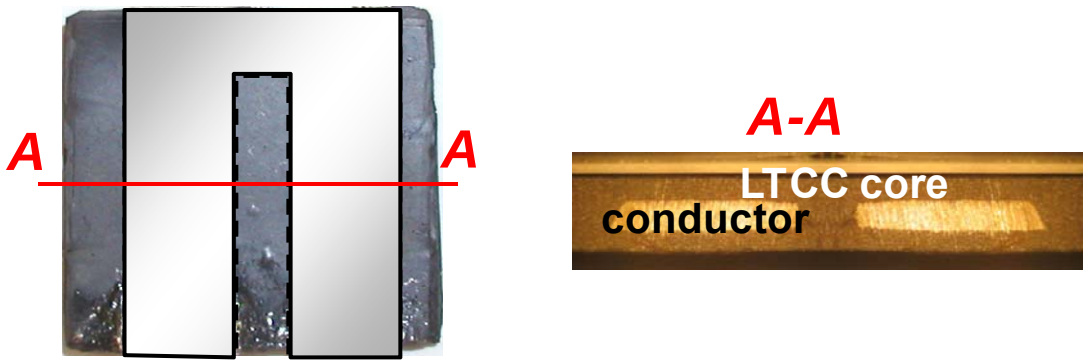

Figure 4.8. Cross-sectional view of actual inductor hardware showing conductor inside ferrite core [119].

With this new design and form-factor, we achieved a 2.8x reduction in volume, a 2.4x reduction in footprint, a 5.2x reduction in DC resistance, and a 1.6x increase in no-load inductance. The trade-off was a 12% reduction of inductance at heavy load. Figure 4.9 compares the two generations [104] and Table 4.1 compares the inductor values.

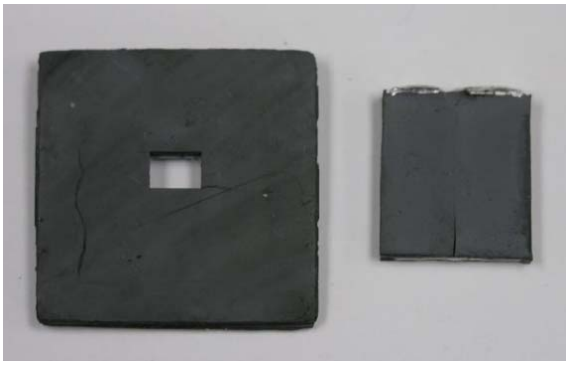

Figure 4.9. Generation 1 on the left and Generation 2 on the right.

Another benefit of going with the 200 material was that it accentuated the nonlinear inductance even more so that we could achieve with still higher efficiency at light load – a 4.5x increase over the full load inductance compared to a 2x increase for a

μr of 50. Figure 4.10 shows this nonlinear inductance effect with respect to output current for our two generations compared with a commercial 100nH ferrite inductor from Coilcraft and Table 4.1 shows the specification comparisons between all three.

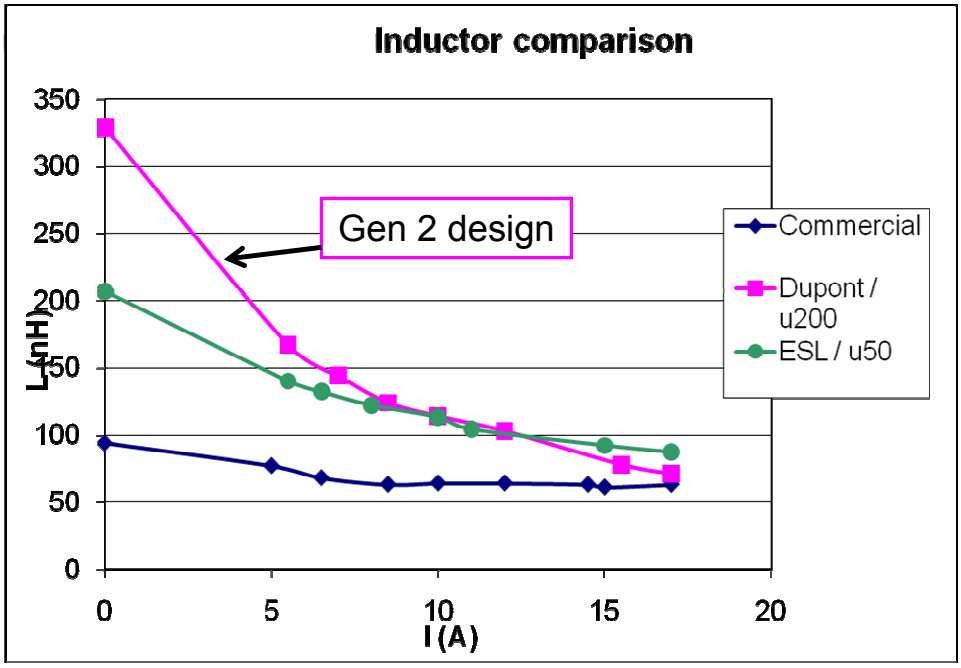

Figure 4.10. Inductance vs. output current for the two LTCC inductor versions.

Table 4.1
Comparison inductor values from actual hardware measurements.

Inductor	Size (mm)	DCR (mΩ)	Inductance (nH) at 17A
Generation 1 LTCC	28 x 28 x 1.4	2.5	86
Generation 2 LTCC	18 x 18 x 1.75	0.45	75
Commercial ferrite L	5 x 5 x 3	0.2	80

Active layer modifications include change the orientation of the MOSFETs so the decoupling cap is now on top of the board, leaving the back side flat for the inductor. This way there is no need for a hole in the inductor to allow the decoupling capacitor to fit through and the inductor can be a whole substrate layer for the active stage. This gave us the benefit of shrinking the input loop with Vin by 50% since it is now physically located closer to the top switch. This reduces the ringing and noise in the input and gate drive waveforms.

The methodology for improving the active layer is based on parasitic reduction theory. Whether for strip inductance or for loop inductance, in order to reduce the parasitic inductance effect, the ratio of trace length to trace width should be minimized [102, 116]. Therefore the traces should be made as short as possible and yet as wide as we have room for. These and other improvements were quantified using Maxwell Q3D and shown in the following paragraphs.

Maxwell Q3D allowed us to obtain the trace parasitics for the converter in order to make a model that was complete and representative of the actual circuit. The parasitics were determined by simulating each loop of the circuit using the same dimensions as the layout. The partitioning of the parasitics for a loop was not easy to determine exactly, of course, so a weighted average was used based on the length of the corresponding trace segments. Figure 4.11 shows the Generation 1 electrical model with the corresponding physical locations on the module.

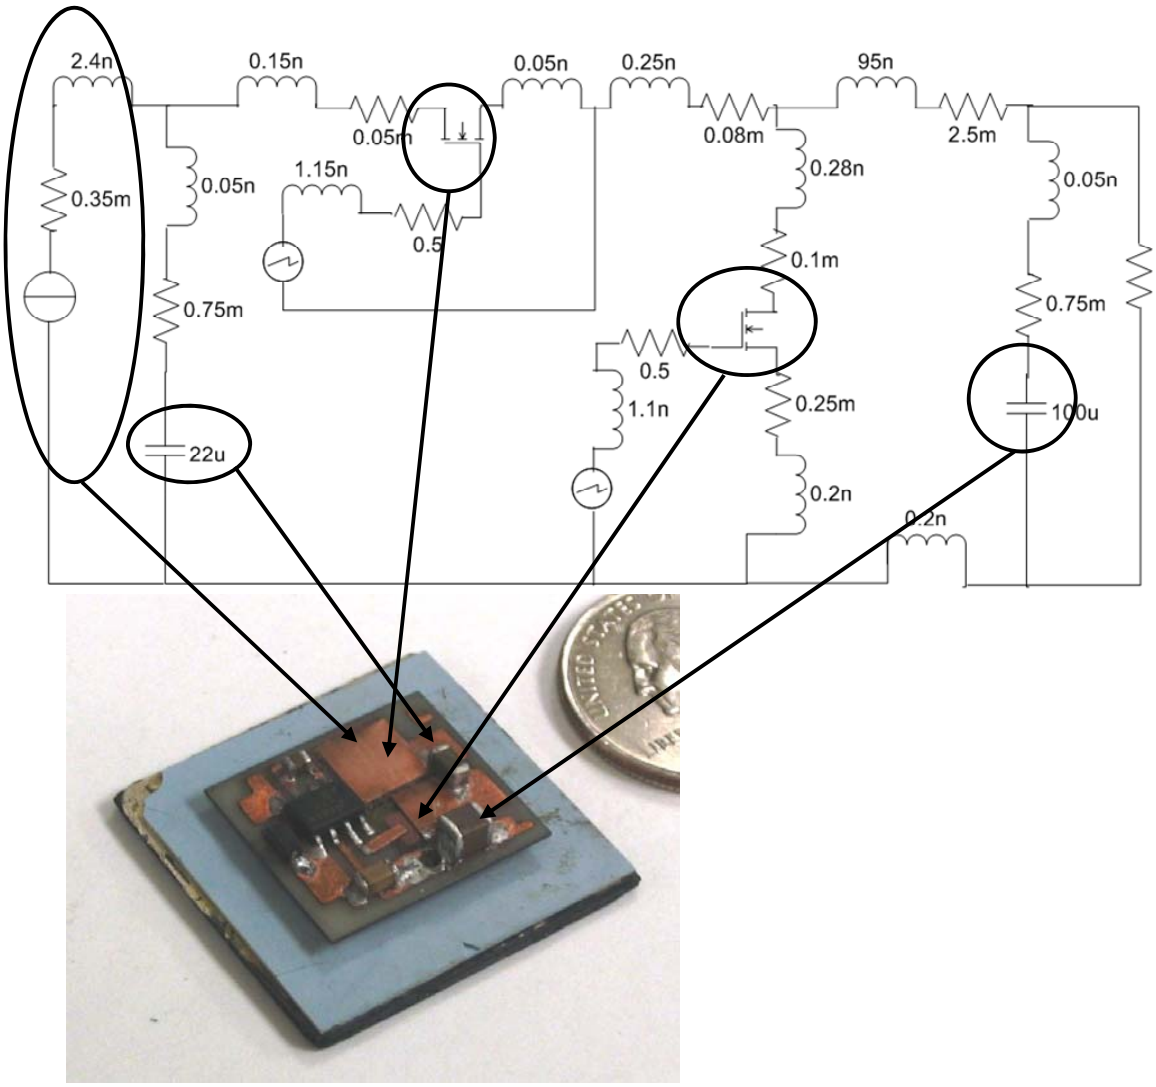

Figure 4.11. Generation 1 electrical model and corresponding physical locations

An improvement was the realization that the snubber capacitor for the top switch doesn't yield the expected outcome and that removing it allows the input voltage loop parasitic AC inductance and DC resistance to be reduced. Maxwell Q3D proved that small differences in layout can yield big changes in these parasitic values. Modeling the input board traces for the Generation 1 layout, we have the model shown in Figure 4.12,

with the input and output faces shown by arrows, and the simulation results for both parasitic values with same values as shown in the previous figure.

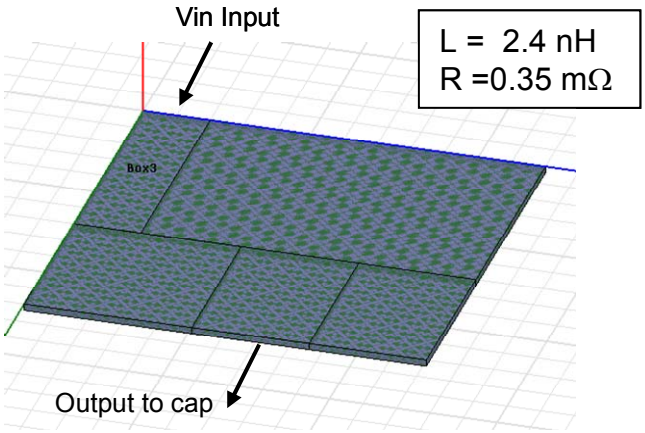

Figure 4.12. Generation 1 input voltage path

Rearranging the layout for Generation 2, although similar, yields a very different result. This is shown in Figure 4.13. Comparing the two, we have essentially halved both parasitic values, which as we shall see shortly, improve efficiency by about 0.2%. Every little bit is important.

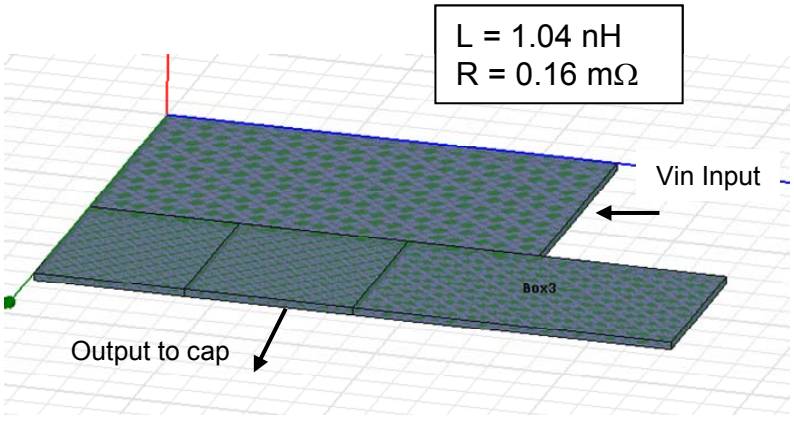

Figure 4.13. Generation 2 input voltage path

Another layout improvement was with the output loop. This is a high-current path so minimizing DCR is critical, as well as reducing the bottom switch's source inductance. The dominant DC resistance in this path stems from the inductor. The simulation "source" is defined at one end of the trace and the "sink" is at the other to represent the current path.

The structure's geometry is the same as what is found on the Generation 1 active layer, going from device drain (Va point), up through a via to the output capacitor. The switch is simply modeled as a copper trace, assuming the switch is on and its L_d and L_s inductances and $R_{ds(on)}$ resistance are subtracted out of the result. Maxwell calculates the inductance matrix for this path and the data is then entered in the Saber schematic accordingly. The other side of the loop consists of the inductor, whose simulation has already been shown. The Saber schematic for this path is shown in Figure 4.14.

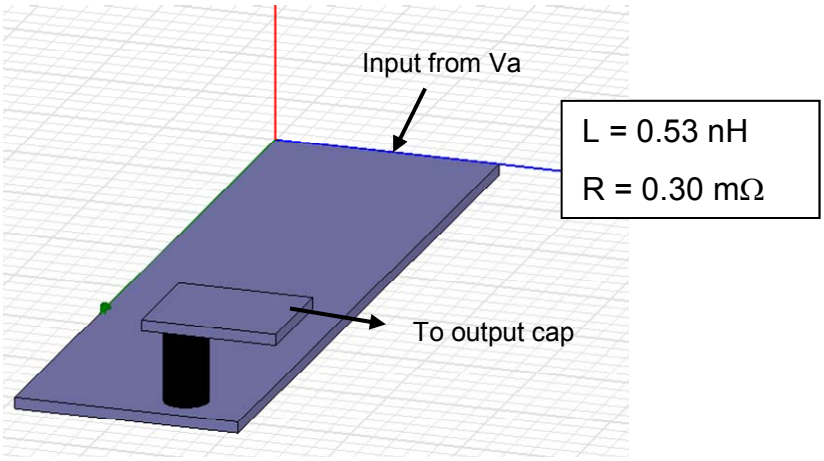

Figure 4.14. Generation 1 Maxwell Q3D high-current path.

For Generation 2, we want to minimize the DC resistance of this path, as well as the parasitic inductance around the bottom switch. The inductor redesign took care of the majority of the DCR reduction. Taking a look at the Maxwell result for current

density in Figure 4.15, we see the crowding occurs mainly in the via. Eliminating it would improve the circuit parasitics and shorten its length.

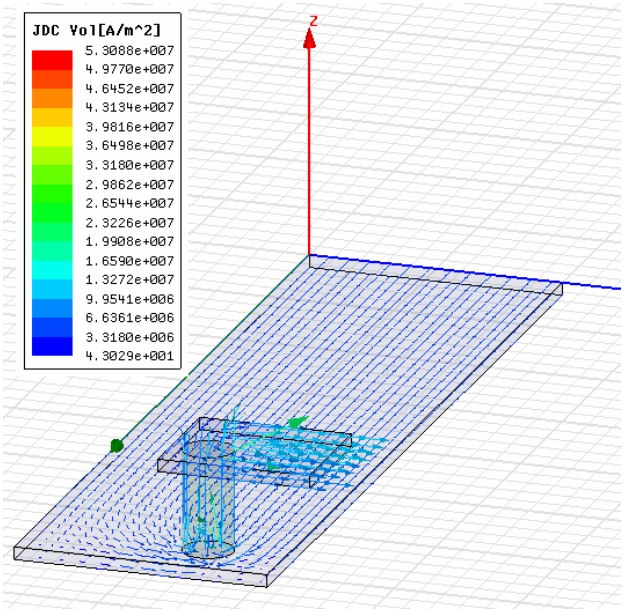

Figure 4.15. Flux crowding in the high-current via.

We could accomplish this if we flipped the MOSFET pair together. This would have an added benefit of moving the decoupling capacitor to the top rather than on the bottom of the board. This allows us to have a flat backside for the inductor layer so there's no need for a decoupling capacitor hole in its middle. Figure 4.16 shows the switch in schematic form.

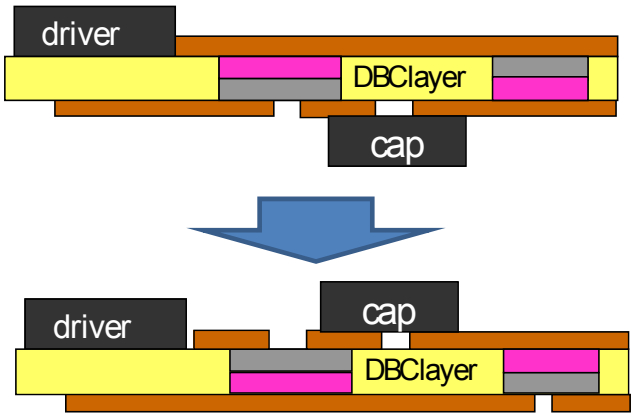

Figure 4.16. A simplified schematic showing the move of the decoupling capacitor to the top side of the board and the flipping of the devices.

Once the capacitor is on top, the high-current path is simplified by the removal of a via and a small reduction in distance between the source and ground. The result is shown in Figure 4.17: A resistance drop of 2/3 and the inductance drops by 1/3, roughly, because the small via has been removed and path shortened. This will be shown later as accounting for a 0.5% improvement in overall efficiency. It is the sum of many little changes that will yield a real performance increase. Such is life at 1.3MHz and beyond.

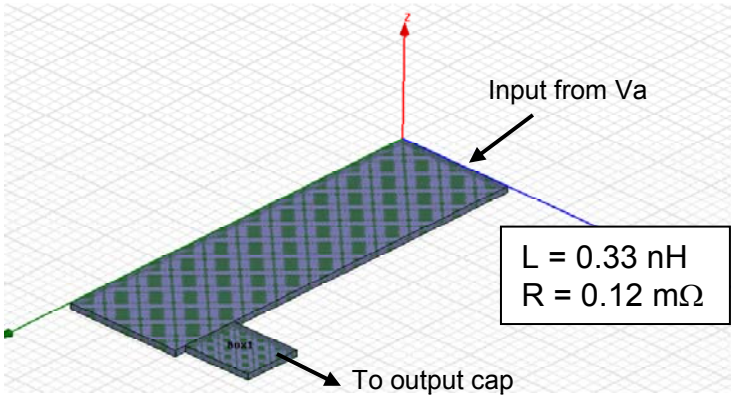

Figure 4.17. Generation 2 layout with reduced parasitic and no via

Similarly, by having the decoupling capacitor on the top side of the board, the path between Vsw (the point where both devices are connected together) and the inductor has been shortened. The parasitic inductance isn't the issue here since we are talking about a trace connected to an inductor, but rather the parasitic resistance since this trace is also part of the high current path. Eliminating the via in Generation 2 reduced the resistance of the trace by 50%. Figure 4.18 shows the final Generation 2 electrical model.

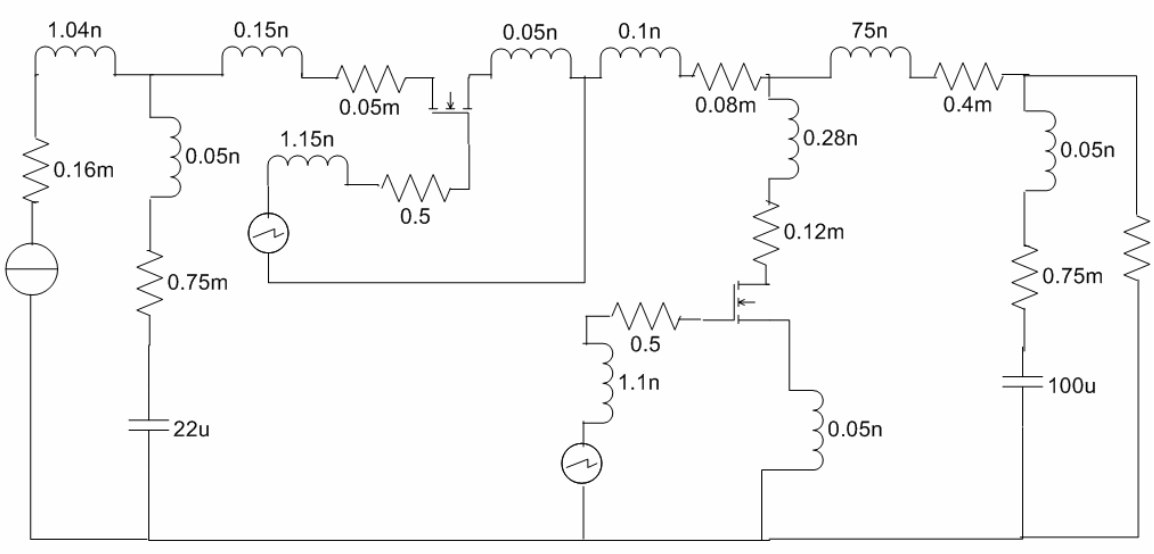

Figure 4.18. Generation 2 electrical model. The corresponding circuit parts are the same as in Figure 4.7.

All these parasitic values are entered in a Saber electrical simulation program that was compared with actual hardware test data once the Generation 2 converter was complete. This comparison is shown in Table 4.2. Metrology for the measured hardware was done using 4 Fluke digital multimeters set to "hold" positions to capture the input and output voltages and input and output currents when the 5v input was reached. The load was a calibrated 10mΩ shunt with a coaxial cable connection so that

output current could be accurately calculated from the mV readings. Voltages were taken directly off the board, connected at pins that were soldered to the board – and not using the actual power connectors on the board because high currents will lead to a voltage drop that will give inaccurate results if readings aren't taken through separate connections. The hardware environment was using 400LFM airflow, just like the simulation, and ambient was set at 25°C for both.

It is clear that the model accurately reflects the real circuit at all three current levels shown in the table. An error of +1.5%,-1% is about as good as can hope for considering the device models are not extremely complex. The models were obtained for the International Rectifier website and plugged in without any modifications. Nevertheless, the model can be used to accurately predict the individual improvements of each change made between Generation 1 and 2.

Table 4.2
Comparison of simulation and hardware efficiency results

Load current	Result type	Gen 1 efficiency	Gen 1 error	Gen 2 efficiency	Gen 2 error
10A	simulation	89.5%	+0.3%	91.1%	-1.0%
	hardware	89.2%		92.1%	
18A	simulation	85.0%	+0.7%	88.6%	+0.1%
	hardware	84.3%		88.5%	
21A	simulation	84.5%	+1.5%	87.6%	+0.1%
	hardware	83.0%		87.5%	

4.3 Generation 2 Improvement Analysis

Using the Saber simulation tool, the ripple waveforms for both Generation 1 and 2 are shown in Figure 4.19. The added inductance (about 20nH) of the Generation 1 inductor gives us less ripple than the Generation 2 inductor. The former is approximately 43% ripple and the latter is 53% ripple. Both of these figures are with a 100uF ceramic 1205 capacitor on-board and are low enough that for many applications, no additional external capacitors are needed – unlike many of today's small commercial POLs. Figure 4.20 shows the new Generation 2 board next to a US quarter coin for scale.

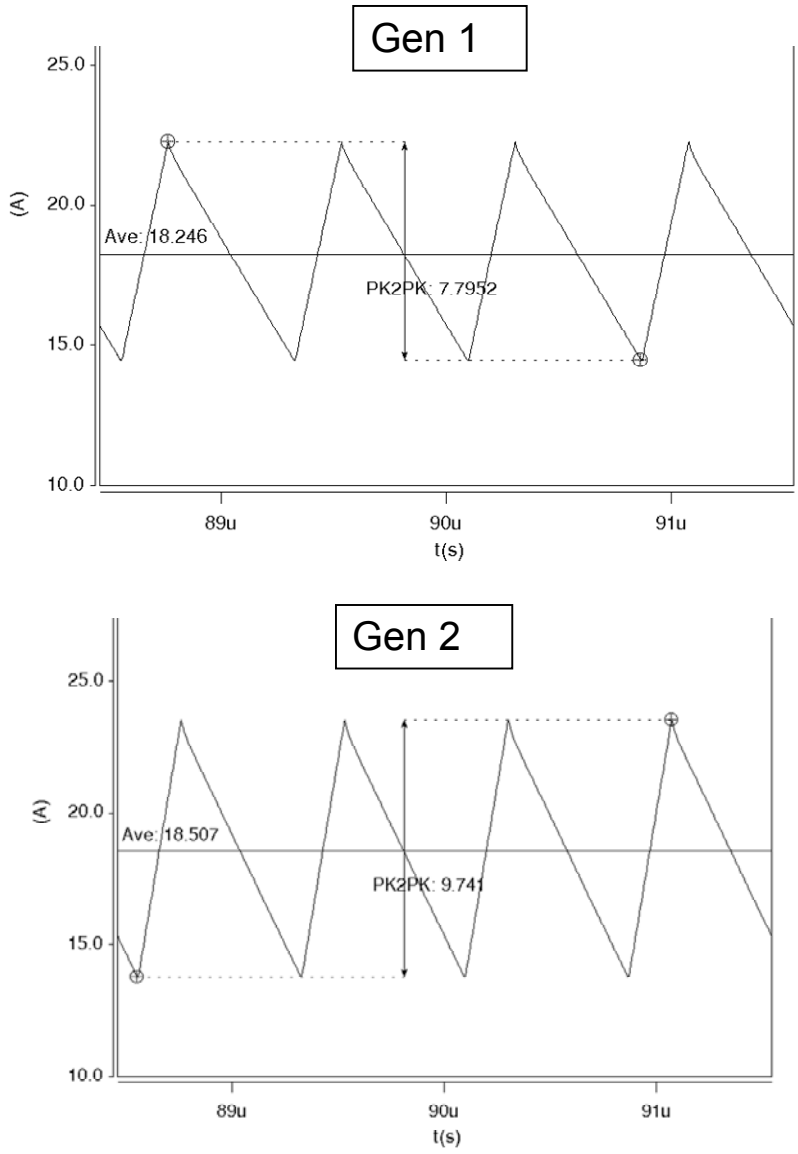

Figure 4.19. Simulated inductor current ripples for both generations.

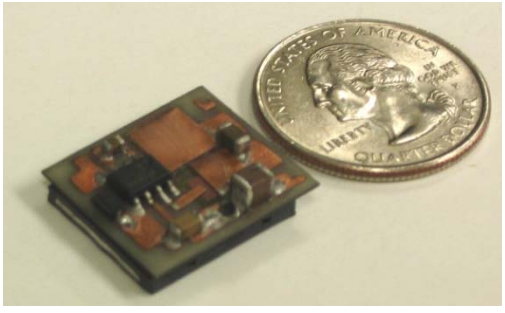

Figure 4.20. Generation 2 module. New active layer and a redesigned LTCC inductor.

The Saber simulation file allows us to determine the improvement breakdown of each modification we've made. Starting off with the baseline case of Generation 1, each modification was made in the simulation file by changing the parasitic values according to the changes Maxwell Q3D gave us, as described above. A simulation is run for each case at full load and the efficiency is calculated, along with the power loss. The circuit layout used for the Saber simulations is the same one shown in Figures 4.11 and 4.18. These include the gate drive losses, parasitic AC inductance, and parasitic DC inductance, in addition to using DirectFET IRF6619 and IRF6691 Saber device models obtained from the manufacturer's website. Inductor core loss was neglected. The results are for 18A output current and are shown in Table 4.3. They match very closely with actual hardware test results at that same output current.

Table 4.3
Comparison of simulation and hardware efficiency results of changes between Generation 1 and Generation 2 active layer designs

change	**η**	**Δ η**	**Δ q**
Gen 1	85.0%	----	----
Change inductor design	87.5%	2.5%	0.6 W
Shorten input voltage path/loop	87.7%	0.2%	0.05 W
Eliminate via and shorten ground path	88.2%	0.5%	0.13 W
Shorten Vsw-to-inductor trace	88.5%	0.3%	0.07 W
Total improvement	----	3.5%	0.85 W

Device loss change between Generation 1 and Generation 2 is shown by Saber simulation in Table 4.4. Figure 4.21 shows an example of the Synopsys Saber

simulation interval of how the switch losses were obtained. The interval is captured once steady-state has been reached (after 90us) and is the integral of drain-to-source voltage (Vds) and drain current (id), with the corresponding units of V*A*s. Multiplying the change by the switching frequency (1.3 MHz) yields the loss in Watts. This particular figure is for the top switch and shows gate-to-source voltage (Vgs) overlaid to indicate what parts are turn-on, conduction, and turn-off losses.

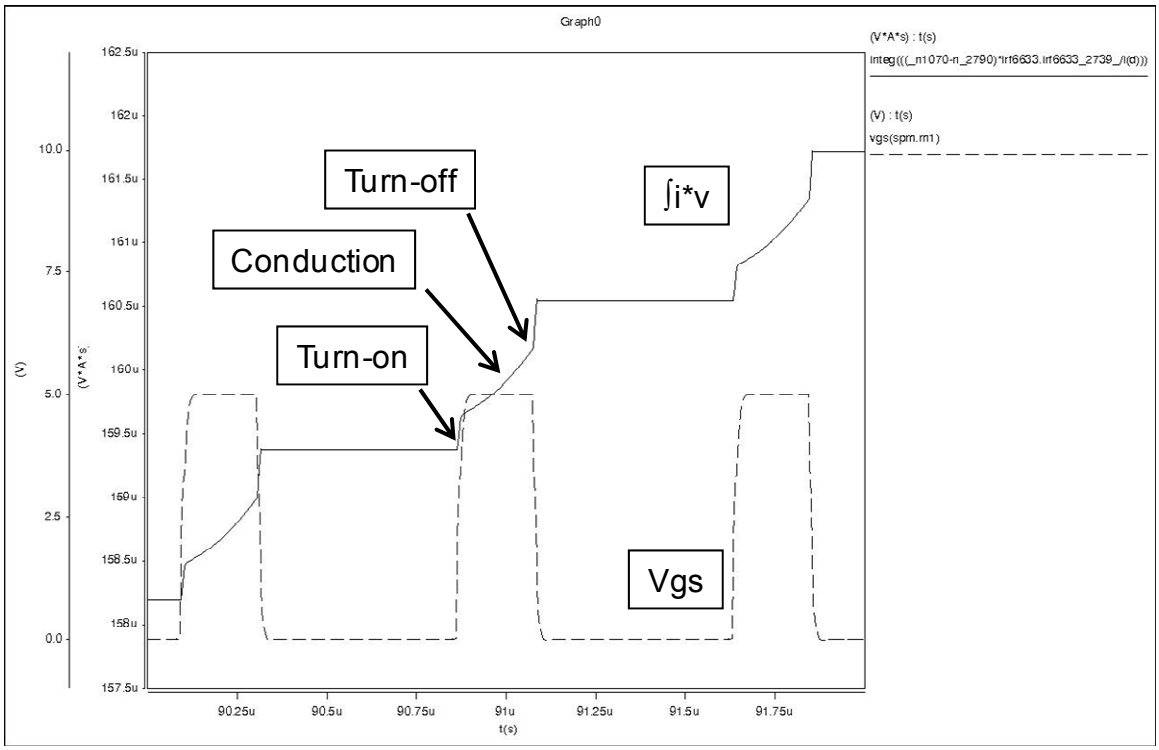

Figure 4.21. Simulation of the integral of voltage and current in the top switch.

Comparing these losses for the two devices in both generations, we see that the layout changes mainly impacted turn-off loss of the devices. Bottom switch losses primarily represent body diode losses and the reverse-recovery charge is attributed in part to both top and bottom switches. Conduction loss remained constant (same output current is used for both cases) and turn-on loss varied somewhat. The sum total losses

for Generation 1 are 3.07W and for Generation 2 are 2.66W. Total converter loss according to hardware is 3.88W for Generation 1, making miscellaneous losses (mostly I^2R loss) amount to 0.81W, compared with 2.95W and 0.29W for Generation 2 respectively. Inductor core loss is neglected but described in a generic form in [115]. Inductor DCR loss accounts for the bulk of the difference in these miscellaneous losses and it too is shown in Table 4.4.

Table 4.4
Simulated loss breakdown comparison at 18A output. All values in Watts.

	Turn-on	Conduction	Turn-off	Total
Generation 1 top switch	0.32	0.54	0.60	1.46
Generation 2 top switch	0.38	0.55	0.46	1.39
Generation 1 bottom switch	0.38	0.33	0.70	1.41
Generation 2 bottom switch	0.36	0.33	0.58	1.27
Generation 1 inductor	---	0.85	---	0.85
Generation 2 inductor	---	0.15	---	0.15

These efficiency improvements are shown in Figure 4.22, this time using actual hardware testing. We get a large improvement in efficiency of about 2% at light load and 9% at heavy load. We are now 4-5% higher than the discrete PCB buck using the same parts and circuit parameters. Also the input voltage waveform had less ringing in this version than in Generation 1 since the input loop parasitic inductance and resistance were reduced by 50%. This is shown in Figure 4.23 next to the original waveform for comparison.

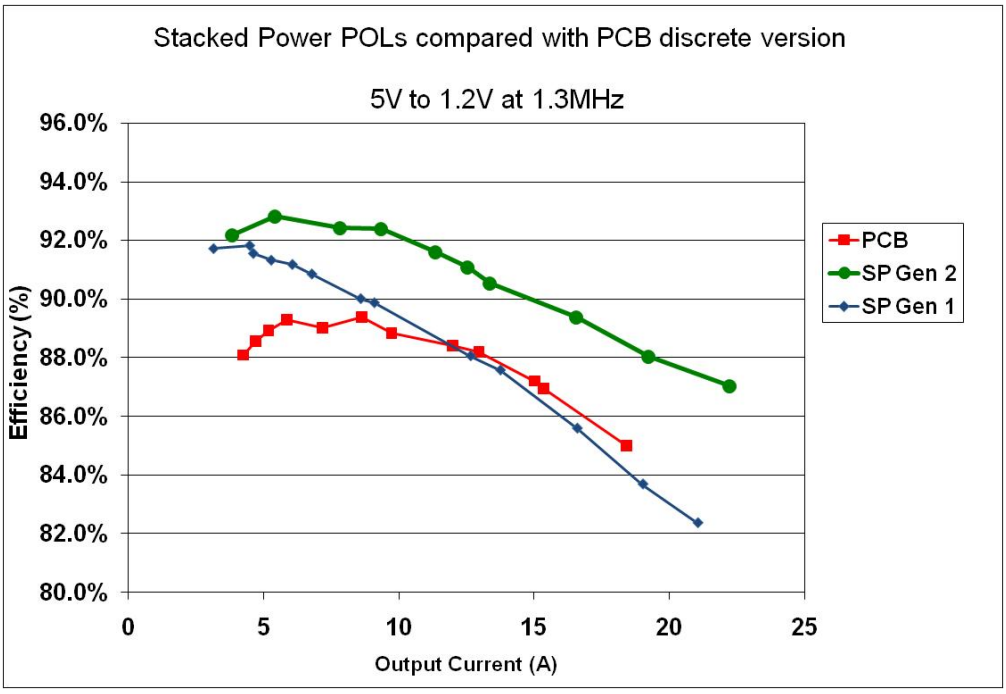

Figure 4.22. Generation 2 efficiency compared to generation 1 and the conventional PCB version

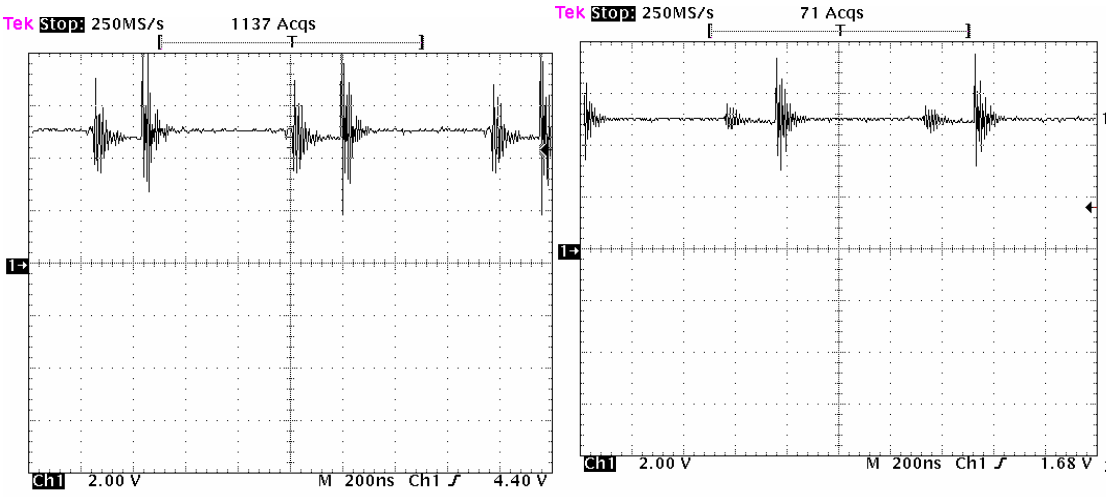

Figure 4.23. Input voltage: Generation 1 on the left and Generation 2 on the right

So what are the implications for this 1-phase DC/DC converter as a phase for VRM applications? We have seen the performance attainable with Stacked Power as a single phase module but what if it was to be extended to 4 phases, like most VRMs? This is a situation where each phase needs to achieve 25A for a total of 100A. Stacked Power has demonstrated that 25A is not a problem to attain so there is no reason it could not be extended to 4 phases. The efficiency of this module was tested at the 12V-1.2V conversion used in VRM applications and is shown in Figure 4.24, along with the 5V input test for comparison. These curves are not only the power stage, as shown above, but also include driver loss for a more complete picture. This efficiency level is very respectable compared to the CPES VRM discussed in Chapter 2, whose efficiency curve is repeated in Figure for comparison.

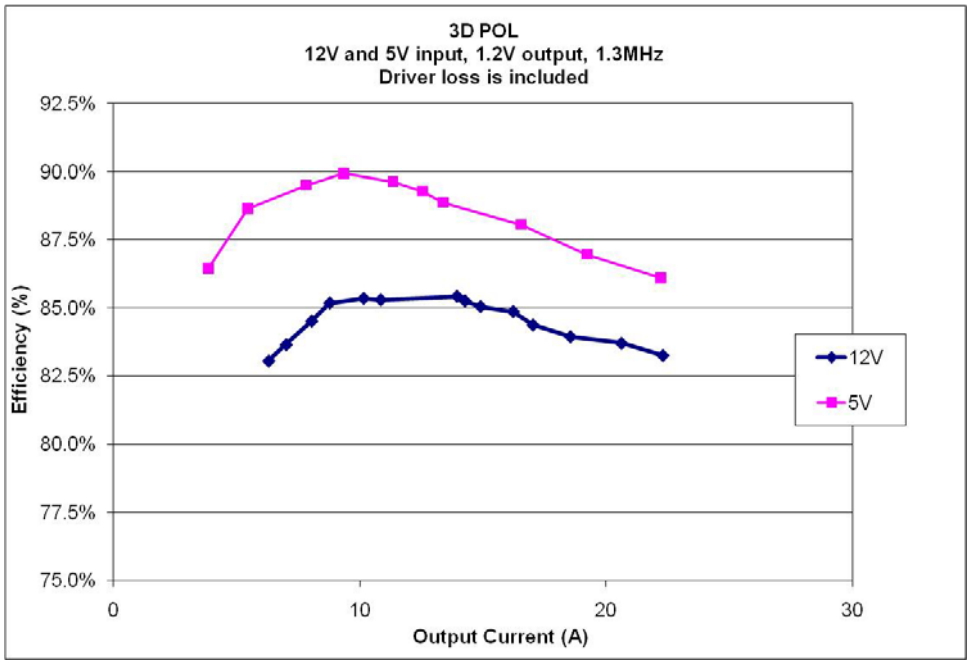

Figure 4.24. Efficiency curves for 12V and 5V input, including driver loss.

At 25A per phase, both Stacked Power and the specialized self-driven circuit have a very similar projected efficiency, the latter shown and discussed in Chapter 1. However, the Stacked Power module uses only 2 devices per phase, compared with an equivalent of 4 for the Self-Driven VRM so right here is a large advantage of Stacked Power in reduction of complexity and cost. Another reduction of complexity comes from the magnetics. The LTCC inductor is a simple layer that replaces the complex transformer arrangement (that requires through-hole mounting), as well as the discreet inductors used at the output. So 6 magnetic components can be replaced by 4 thin layers – and no holes are required in the substrate other than signal vias – using Stacked Power.

4.4 Thermal Analysis of Stacked Power

The other large improvement available with Stacked Power design is that of excellent thermal performance due to the replacement of FR4 with AlN ceramic. Commercially available modules of this size and power range are generally encapsulated in plastic, as will be discussed in the following chapter, and thus aren't capable of this level of thermal performance, demonstrating a 60% reduction in thermal capability compared to Stacked Power. Discussing this issue with various industry engineers, it has become clear that the only significant reason for using the plastic is to conceal company secrets. Telecom bricks still use an open-frame design to take

advantage of the thermal improvement over and encapsulated circuit so this issue seems limited to POLs under 20W output (see Chapter 1 for examples).

Plastic encapsulant has a very low thermal conductivity so that even if a heat sink was to be used on top for additional cooling of the component, its effectiveness would be greatly diminished due to the temperature drop through the plastic, and so the size of the heat sink would not be optimized, resulting in a larger size than necessary which reduces power density drastically.

In open-frame POL designs where the parts are exposed, the typical design uses FETs where in this case, they are the ones encased in plastic. As a result, the same issues ensue. Often the argument is that thermal vias in the PCB substrate can allow for the device to be cooled from both sides – off the PCB and off the device surface. Here we take an analytical look at how such an approach is limited due to the high thermal resistance of the device plastic case, the thermal pad for the heat sink, and the FR4 portion of the PCB (as discussed at length in Chapter 2).

Figure 4.25 shows the ultimate scenario in double-sided cooling: A heat sink on top of the device, as well as one on the back side of the PCB's thermal vias. This structure is analytically modeled with thermal resistances as shown. There are two heat flow paths – one where the FETs' die heat flows up through the case to the heat sink and then to the ambient air, and the other where the heat flows from the FETs' bottom side case through the PCB thermal vias, to the heat sink on the opposite side. We can setup a ratio between the two heat paths as shown in Equation (27). The heat-sink-to-ambient thermal resistance is not included in the ratio since it is assumed to be equal for both sides and thus cancels out.

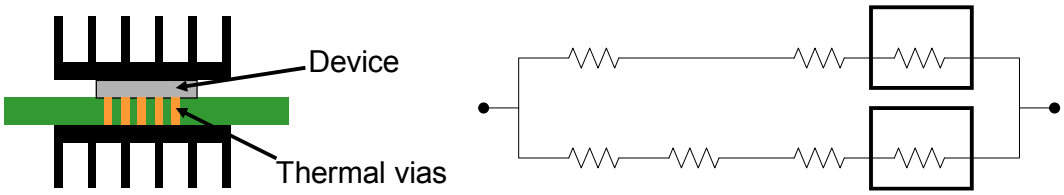

Figure 4.25. Double-sided cooling paths for a device on PCB

θjc

$$\theta_{ratio} = \frac{\theta_{jc} + \theta_{d-hs}}{\theta_{jb} + \theta_{vias} + \theta_{b-hs}} \quad (27)$$

Tdie

θjb

Where,

$$\theta = \frac{L}{Ak} \quad (28)$$

- Θjc = FET junction-to-case thermal resistance
- Θjb = FET junction-to-board thermal resistance
- Θvias = Thermal via in PCB thermal resistance
- Θd-hs = device case to heat sink thermal resistance
- Θb-hs = board to heat sink thermal resistance
- Θhs-a = heat sink to ambient air thermal resistance
- *L* = characteristic length of the material
- *A* = area of surface contact
- *k* = the thermal conductivity of the material

The thermal resistance values calculated for each type of material interface is shown in Figure 4.26. Thermal conductivity for each type of homogenous material has no need for an isotropic lump kn/kp value (refer to Chapter 2 if details on this process

are desired). However, the PCB thermal-via composite value is the only one determined in such a way, and here it is one with the highest achievable via density in 4-layer with the conventional 2,1,1,2 oz. (70, 35, 35, 70 μm thick) motherboard PCB substrate, at a value of 117 W/m°C [98], again for the assumption of a best-case scenario. The device-to-heat sink and board-to-heat sink thermal resistances are determined to be 8 W/m°C based on thermal pad materials commonly used today.

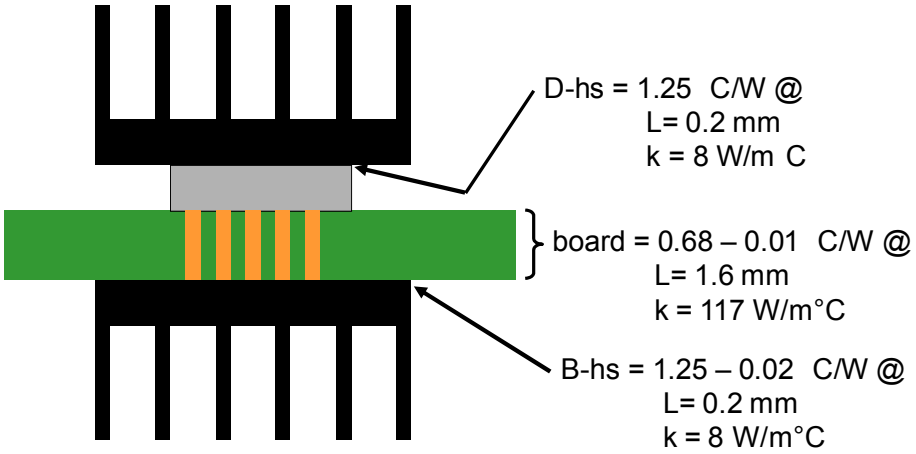

Figure 4.26. Interface resistances and associated variables

This leaves the device's internal thermal resistances as the final pieces necessary to assess the thermal resistance analysis. This of course will change depending on what devices are chosen for the application and the values are given in the device datasheets. Plastic-encapsulated SO8 devices will be adding on the order of 20 °C/W for each cooling path! Again, going with our best-case-scenario, we can choose the much-improved DirectFET or the Renesas LFPAK, both with junction-case thermal resistances of 3°C/W and junction-to-board resistances of 1°C/W – many times better than the classic SO8 package, as shown in Figure 4.27.

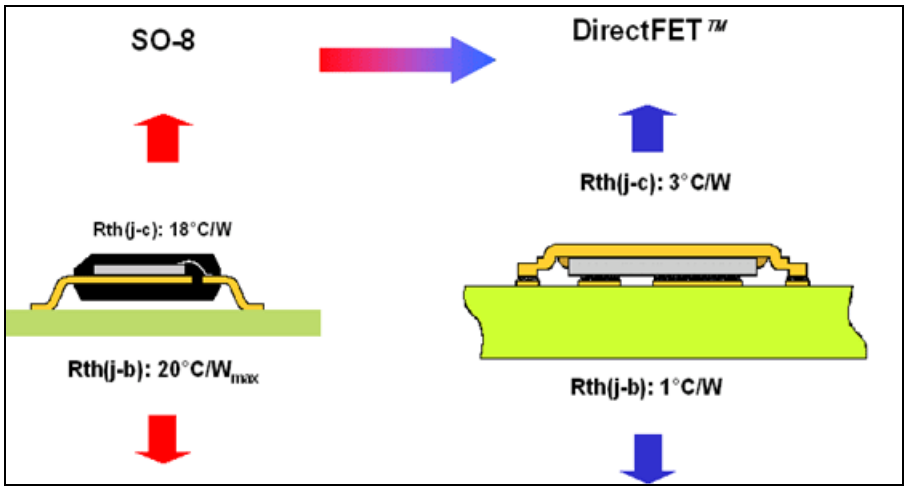

Figure 4.27. Comparison of in-package thermal resistances of SO8 and DirectFET

In so choosing these thermally improved devices, we can map out a range of resistance values depending on the contact area available. Of course, the contact area for the device to the heat sink is a fixed value, determined by the package size. As a result, there is only one answer for this resistance and is shown in Figure 4.28 as the diamond at approximately 4.2 °C/W for either DirectFET or LFPAK. This will be added to their respective junction-case thermal resistances of 3°C/W, and to the thermal pad interface at 1.25°C/W, for a total junction-to-heat sink value of ~8.5°C/W.

However, for the PCB side, the contact area between PCB and heat sink can vary widely depending on the PCB component layout. A heat sink cannot be installed overtop of drivers, controllers, capacitors, etc. and so a flat area must be made available to extract heat from the board. This can be very difficult to do and requires lots of constraints in the layout if a compact footprint is desired. As a result, Figure 4.28 shows a range of possible values depending on how much flat area is available for the heat sink to make contact. However this range quickly becomes asymptotic to the resistance value of the solder interface once the area becomes greater than 2e-4 m^2

(200 mm^2). Below this value, it goes up to a peak of 3°C/W at the area of the device footprint. For our best-case-scenario, a value of 2°C/W is appropriate since board area dedicated to thermal vias is always limited due to signal trace requirements.

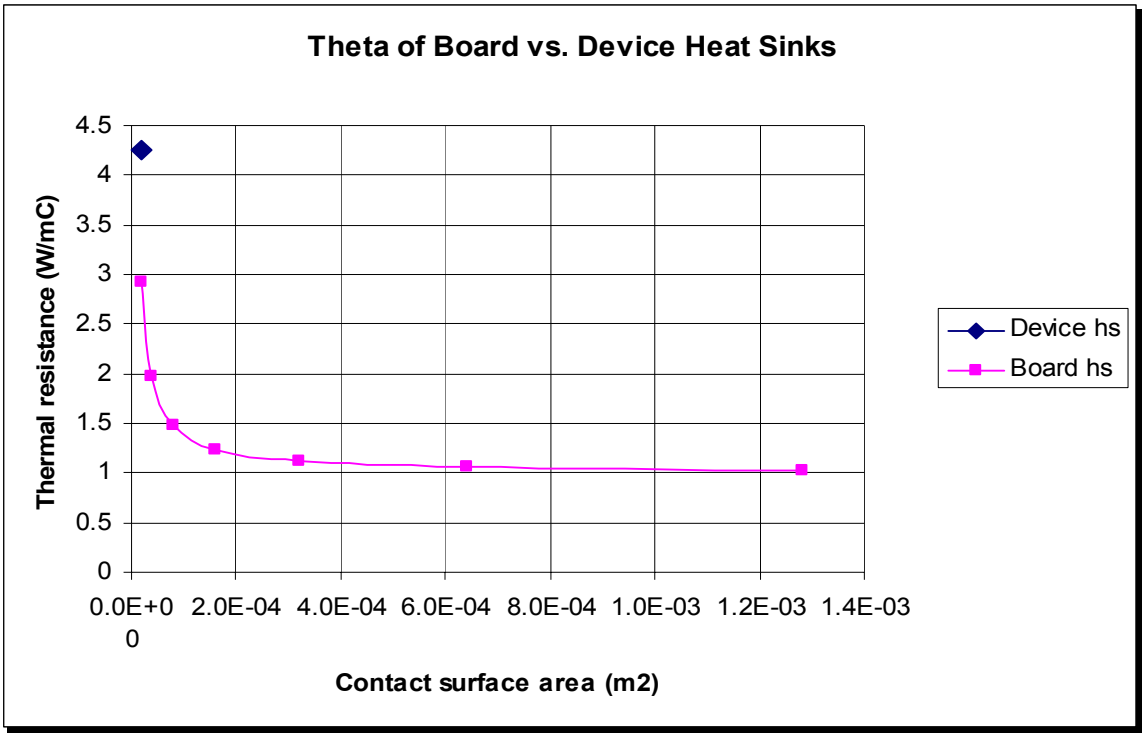

Figure 4.28. Heat sink interface thermal resistance

So on one side of the device, we have a total junction-to-heat sink value of 8.5°C/W as just discussed. On the other side, we have a total value of 1°C/W for junction-board, + 2°C/W for the PCB with maximum thermal vias, + 1.25°C/W for the thermal pad yields a total of 4.25°C/W – or about half of the top-side device cooling largely due to the dominant device package thermal resistance from junction-to-case. As a side note, spreading of the heat inside the PCB rather than strictly through the thermal vias is assumed negligible due to FR4's 800x greater thermal resistance (compared to copper vias).

In contrast, with Stacked Power, the heat sinking is actually the substrate itself so no additional heat sink is needed as long as there is natural convection available. This will be discussed more in Chapter 5 but for now, we want to compare the thermal analysis of the PCB just given to that of a comparable Stacked Power module. Figure 4.29 shows the cross-sectional view of the corresponding Stacked Power configuration.

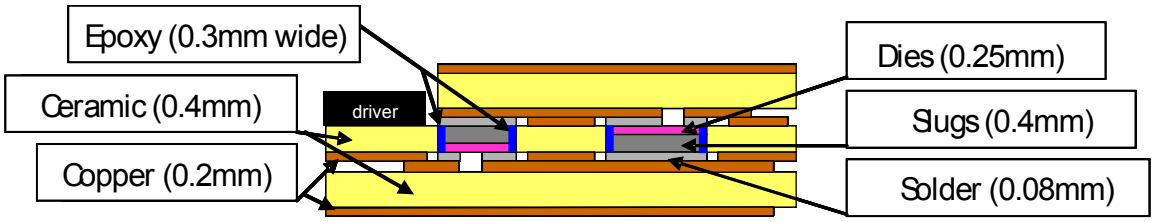

Figure 4.29. Cross-sectional view of Stacked Power module assembly

As shown in Table 3.2, the unpackaged die thermal resistances from device to substrate (heat sink) are on the order of 0.9°C/W for the small top switch to 0.4°C/W for the larger bottom switch. These are extremely low values since the interface is solder rather than a poorly-performing thermal pad. The device internal junction-to-board resistance for the DirectFET remains 1°C/W as before but the top side is reduced since the conductive epoxy used to connect the "can" to the drain is no longer needed because the can is not used in Stacked Power. Therefore, this thermal resistance is reduced from 3°C/W to the order of 0.5°C/W (based on double the surface area of the source/gate side of the die). The sum total yields a junction-heat sink thermal resistance for Stacked Power of 0.5°C/W (junction to drain) + 0.4°C/W (drain to copper interconnect) + 0.1°C/W (copper to AlN ceramic heat sink) = 1°C/W for the upper cooling path! This is 8.5x lower than the 8.5°C/W for the conventional heat sink method also using DirectFETs.

As for the other cooling path, we have the same scenario since the dies are embedded and have similar thermal resistances on top as on bottom (the top side has larger area but also has a heat slug to go through). So here we have 1°C/W for Stacked Power and 4.25°C/W for the convectional method using thermal vias. Between the two cooling paths, Stacked Power has over 12x better heat extraction from the dies (a simulation result was shown in Figure 2.33).

The stacking of multiple DBC layers can only improve the situation. The simulations shown thus far are for only one layer of DBC with copper strap interconnections but adding additional layers will increase the heat extraction due to the higher volume content of high-thermal-conductivity material. The heat capacity of the material is the first step to spreading it out for exposure to the cooling air flow. Heat capacity (also known as "thermal mass" in J/°K) is defined as,

$$C = V \cdot \rho \cdot c_p \tag{29}$$

Where,

V = volume in m^3

ρ = density of material in kg/m^3

c_p = specific heat in J/kg°K

Taking the fabrication example of Chapter 3 and looking at the heat capacity formula, it because clear that the prediction of an additional full layer of DBC plus a small "Z" piece on top represents an increase of 130% volume which directly translates into 130% increase in volume that can be heated since density remains constant. This is the same concept that governs "thermal lag" conditions where temperatures cannot change instantaneously due to heat capacity of the material. As will be shown in the

next few pages, adding the extra two layers of DBC allows for an 18% increase in output current capability for the same operating environment when compared to a lower-volume single-DBC-layer module since the heat loss remains the same for both. Of course this performance increase must be weighed against the added cost of extra material.

The heat transfer to the ambient air is dependent on many external factors: The surface area of the Stacked Power module, the orientation of the air flow relative to the surface, etc. These aspects are independent of the analysis shown here but will be discussed in great detail in Chapter 5. The module and the respective simulations assume no conformal coating or thermal gap filler. For applications where moisture contamination and/or breakdown isolation is needed, the module can be dip coated just as would be done for a PCB-based module. This will not affect thermal performance significantly, as demonstrated in Section 2.3.

In order to predict the thermal performance of the Stacked Power module in a more realistic way, FEA thermal simulations are used for great accuracy and modeling of the heat spreading effect (which is much more pronounced in AlN ceramic than FR4, as described in Chapters 2 and 3). The material conductivity values used for these simulations are shown in Table 4.5. The thermal resistances between dies and board were given in Table 3.2 and discussed just now. The losses used were given in Table 4.4.

Table 4.5
Material conductivity values used for I-DEAS FEA simulations [90, 116]

Material	Conductivity
AlN DBC	150 W/m°C
Silicon dies	125 W/m°C
Copper	380 W/m°C
Inductor	4 W/m°C

Using the same procedure as shown in Section 2.2 for the PCB isotropic conductivity values, the parameters for equations 4 and 5 for an AlN DBC layer are shown below in Table 4.6. Regardless of whether there is 1 layer of DBC or 3, the lump conductivity values are the same, being kn = 265 W/m°K and kp = 215.1 W/m°K. These values are combined using the arithmetic mean, as was done for PCB, although this time the harmonic mean is very close in value since the AlN and copper conductivities are both much closer in value than for FR4 and copper. This arithmetic mean is used in the I-DEAS thermal simulations as a lump conductivity value for the substrate.

Table 4.6
Isotropic thermal conductivity parameters for AlN DBC

Layer	thickness, ti	conductivity, ki	ki*ti	ti/ki
copper	2.00E-04	380	7.60E-02	5.26E-07
AlN	4.00E-04	150	6.00E-02	2.67E-06
copper	2.00E-04	380	7.60E-02	5.26E-07

- Arithmetic mean: 240.1 W/m°K
- Geometric mean: 21.9 W/m°K
- Harmonic mean: 237.5 W/m°K

Thermal performance of the Generation 2 module is considerably better than the Generation 1, thanks mostly to the much higher efficiency since both use the same DBC material and size of 18 x 18 x 1 mm. Inductor dimensions are from Figure 4.3. Figure 4.30 shows the I-DEAS thermal simulation model for this module. The box around the module is 50 wide x 35 tall x 50 mm deep and represents the simulation control volume for the calculation of closed-form boundary conditions. The volume is air meshed using 1mm non-deformed soild triangles of ESC-air, which is the I-DEAS material for air simulations. The output cap is a 1206 and the driver is a regular SO-8 package, both modeled as simple blocks. The dies are 2 x 2.6 mm (78 x 102 mils) and 3.4 x 4 mm (133 x 159 mils) respectively and the epoxy gap around the devices is 0.3 mm (15 mils) wide.

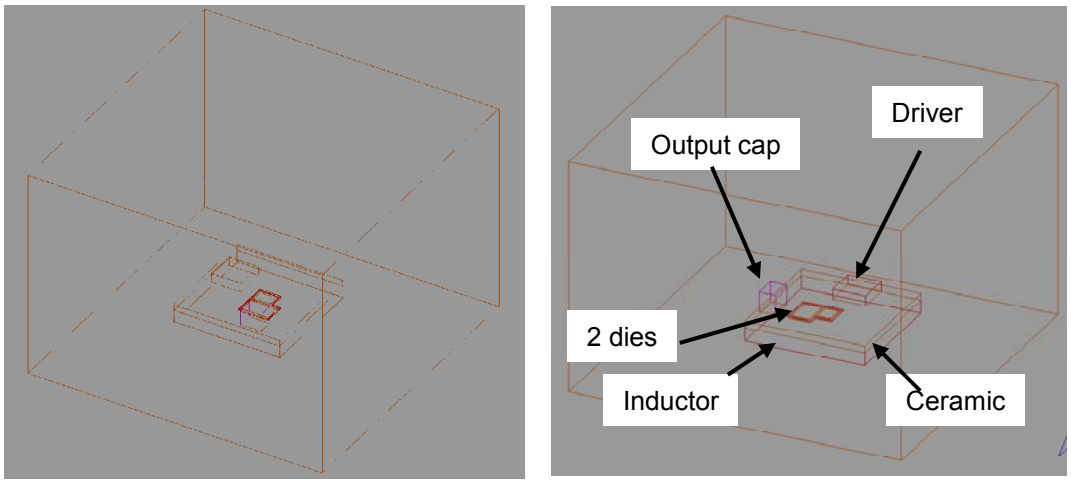

Figure 4.30. Two views of the wireframe model, rotated relative to each other. The dies are embedded inside the ceramic layer and surrounded by epoxy.

Using I-DEAS FEA thermal modeling, the predicted results for the Generation 2 module are shown in Figure 4.31. Simulation condition was at 18A and 23°C ambient to represent lab conditions and shows a peak of 87.8°C. The losses for this case were shown in Table 4.2 and are all modeled as volumetric losses. The hotspot is in the

middle of the DBC and represents the device losses. The inductor layer is directly beneath it and is approximately 10 degrees cooler than the DBC. Table 4.7 shows the simulation parameters used in the simulation. Equation 9 and Table 2.4 describe the procedure for obtaining the thermal coupling boundary resistances. The losses shown account for 97% of the losses indicated by the system efficiency.

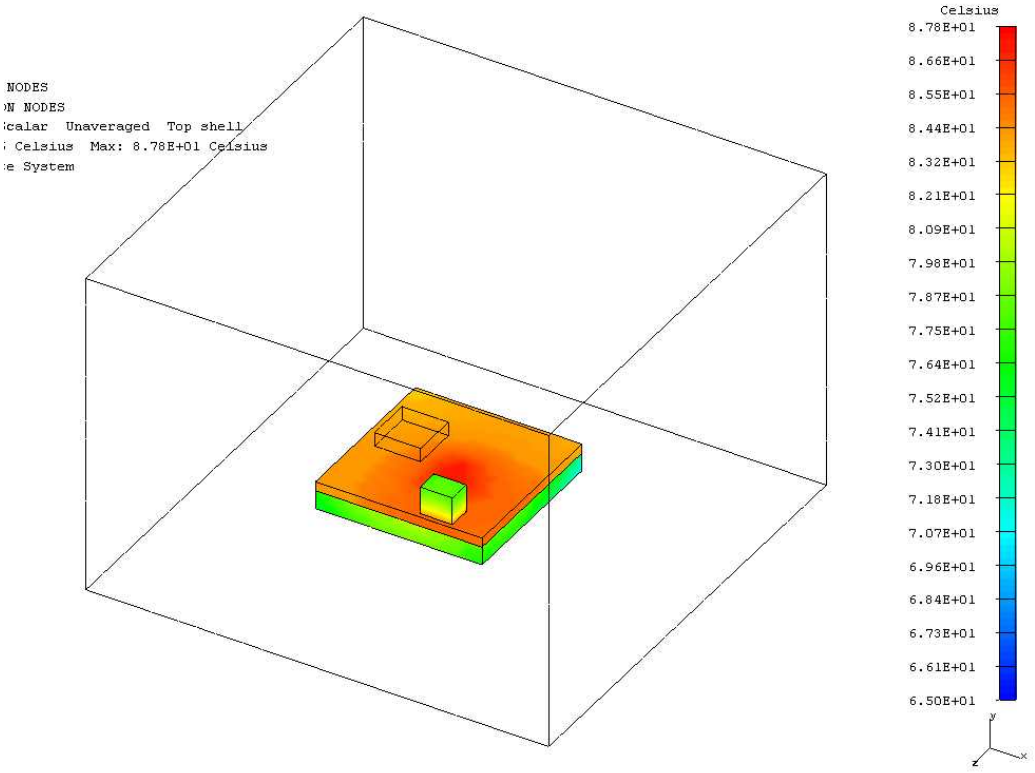

Figure 4.31. I-DEAS simulation of POL at 18A and 25°C ambient.

Table 4.7.
Simulation parameters [116].

Component	**Loss**	**Couplings**	**Conductivity**
Ceramic	0	isotropic	240 W/m°C
Top switch	1.39W	1.11, 0.89 °C/W	125 W/m°C
Bottom switch	1.27W	0.56, 0.37 °C/W	125 W/m°C
1206 capacitor	0	10 °C/W	4 W/m°C
Epoxy	0	0.5 °C/W	0.12 W/m°C
Inductor	0.15W	5 °C/W	4 W/m°C

Figure 4.32 shows a cross-sectional view of the simulation results. At this temperature scale, the epoxy surrounding the dies shows that it absorbs no heat due to its 2000 times lower thermal conductivity compared to the surrounding DBC material. How well heat spreads in the AlN DBC material is easily visible, as well as the low temperature of Generation 2 ferrite inductor.

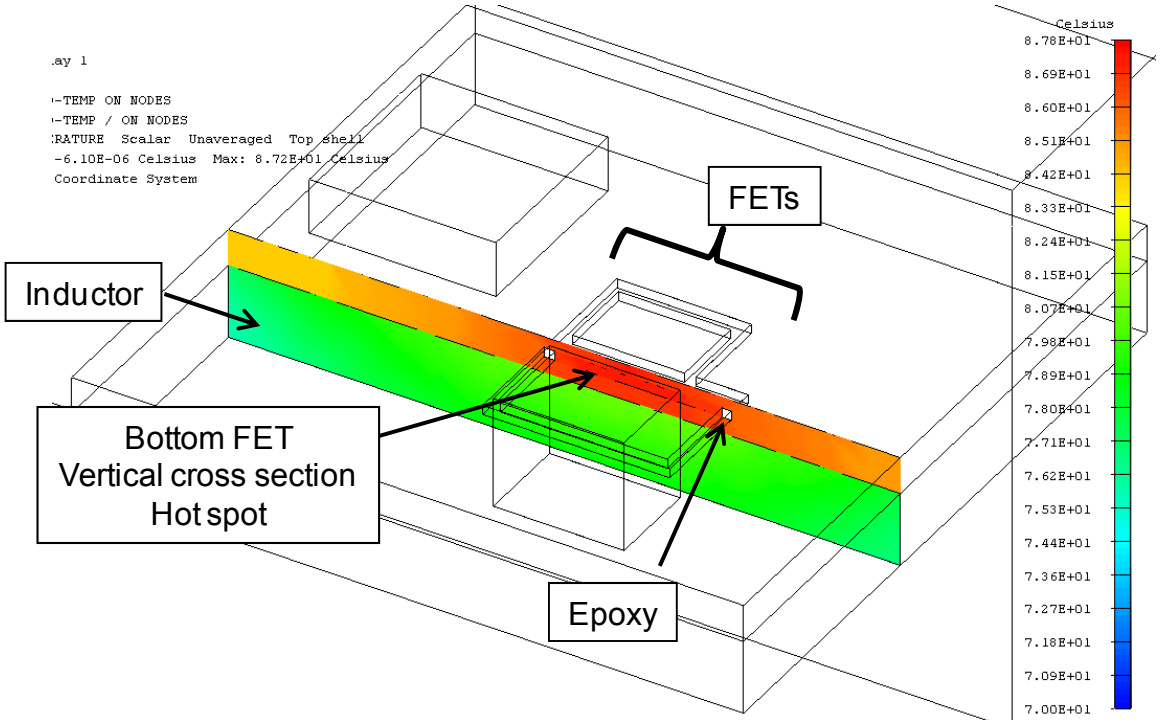

Figure 4.32. Cross-sectional view of thermal result of POL

The actual hardware Generation 2 module, while in operation on the bench in the lab, was photographed with an infrared thermal camera model FLIR S65 HS ThermaCAM, shown in Figure 4.33, to accurately map its thermal performance. The thermal camera's global emissivity value was set to 0.67 for copper, set according to the manual, and the reflected temperature was set equal to ambient: 23°C. Ideally the entire module would be painted with flat black paint in order to make the emissivity value be entire even over the entire surface but since this was not the case (because the paint would artificially raise the temperatures), the camera was calibrated using a type K thermocouple with excellent correlation at various points on the board surface. The thermocouple tip was dipped in thermal paste so as to significantly reduce contact resistance and thus improve accuracy. It was used to measure temperature on the substrate, on the FETs, and on the inductor. When a long-term measurement was

required, it was taped down with Kapton tape. Otherwise, it was held in location manually so as to reach deep within layers for maximum temperature readings.

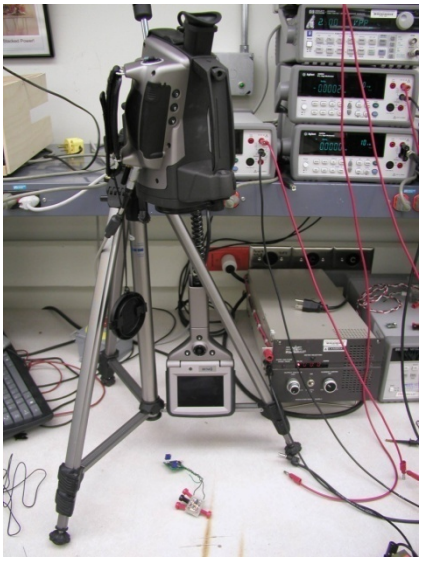

Figure 4.33. FLIR S65 HS thermal camera with power supplies and meters.

Figure 4.34 shows the resulting thermal map of the module. The typical ambient temperature inside servers is assumed to be 55°C. It would be wonderful for this temperature to be lower but cooling costs and the over-taxed air conditioning systems of server farms just cannot handle it. The lab is of course cooler than this, at 23°C. Therefore, the hardware thermal test was stopped prematurely to compensate for this difference. Here the module was running at 17A output with a peak temperature of 89°C. Adding the 32 degree ambient temperature difference essentially yields the maximum safe design limit of 120°C.

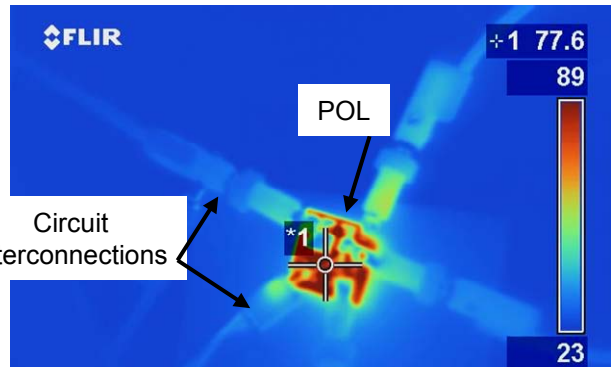

Figure 4.34. Thermal map of Generation 2 module

The inductor temperature was measured with a thermocouple since it was located on the bottom side of the module and not visible by the camera. However, inductor temperature rise was not very significant relative to the DBC temperature. A reading of 41°C for the conditions of Figure 4.34 was typical on the surface of the inductor with little variation when using a K-type thermocouple dabbed in a small amount of heat-transfer paste. An I^2R conduction loss calculation at 17A and 0.45 mΩ yields only a 0.13W loss which correlates with this low surface temperature.

As was mentioned in Chapter 3, the figure-of-merit of hotspot-to-package ratio is a convenient way of demonstrating the performance capability on hand. An encapsulated module or one with a PCB substrate will have a hotspot basically the size of the dies. The dies typically represent 20% of the surface area of the module so the ratio is only 0.2. But as can be seen in Figure, the hotspot is the size of the substrate with Stacked Power! A ratio of 1 is like having a chip-scale package, but in thermal terms. The advantage is multifunctional use of materials which reduces cost by

reducing the need for heat sinking and airflow, not to mention making the system design the POL is used in all the more convenient and simplified.

But the question now is can we do even better? Can we obtain even more output current in the same footprint and for the same temperatures? It is best to test this with hardware due to the complexity of natural convection mechanisms and thus the difficulty of getting accurate simulation results. However, simulations can be used as a good guide and it was determined via I-DEAS simulations that additional DBC material is necessary to improve the thermal capability in natural convection beyond 17A.

The epoxy walls surrounding the devices are thermally “shorted” by the nearly 500 times greater thermal conductivity of the solder/copper strap interconnect, and thus has no thermal impact. The heat is conducted away from the devices strictly through the interconnects, is then absorbed by the AlN substrate, spreads throughout the volume of the material, and finally the heat present on the outer surfaces is convected to the surrounding air.

Figure 4.35 shows an I-DEAS simulation of the Generation 2 module but with two layers of DBC. Previously (shown in Figure 4.31) we had peak temperatures of 87.8°C but by doubling the DBC thickness, and assuming a layer of 0.25 mm dielectric polyimide in between them (like the epoxy at 0.12 W/m°K, simulating dies that are *not* embedded), temperatures fell 6% to 82.8°C thanks to the increase in DBC material. This figure has its temperatures scaled to those of Figure 4.31 for relative comparison. This is also assuming unpackaged DirectFETs are used, which is a best case scenario. Later in this section, we will re-examine this simulation by using the same amount of

DBC material as shown in Figure 4.31 and see that temperatures will rise correspondingly.

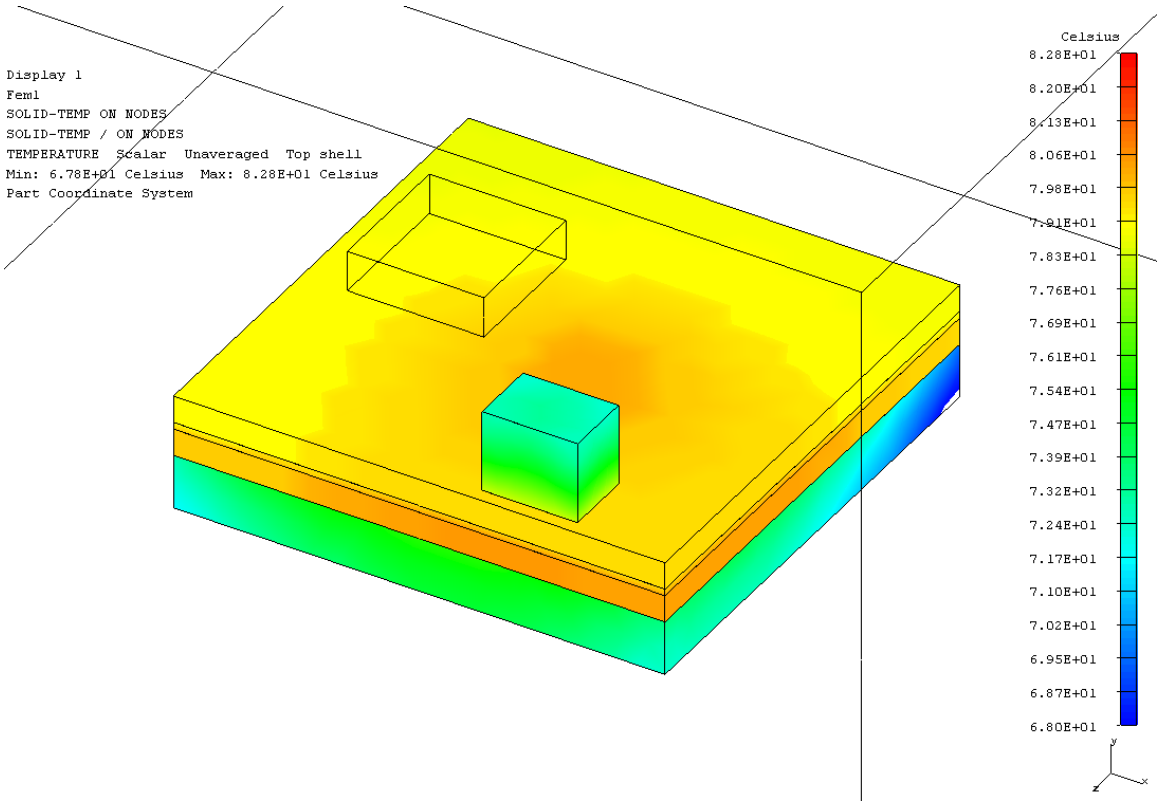

Figure 4.35. Impact on temperature of adding a second DBC layer.

With the Generation 2 hardware and again using the thermal camera, the circuit was run to a maximum of 121°C (assuming 55°C ambient) but that was only 17A load, as shown in the previous section. In a quest to develop a module that would not need a fan, or an added heat sink, and be able to reach our target 20A load, we took advantage of 3D FEA thermal simulations to lead us. It was determined that adding a second layer of DBC on the bottom of the active layer, equal in size, and a tiny piece on top of the devices to create a better thermal path to the remainder of the converter. In addition,

extra DBC has the potential of having its bottom layer of copper act as a shield for the inductor to good benefit [127, 129].

This module is shown in Figure 4.36 and has a similar layout as was shown for Generation 2 with the exception of the added DBC interconnect/heat spreader/shield layers described in great detail in Chapter 3's fabrication section. Its footprint is no larger than Generation 2 but the thickness is increased by about 1.2mm from the extra DBC added to handle the extra heat load from higher current operation.

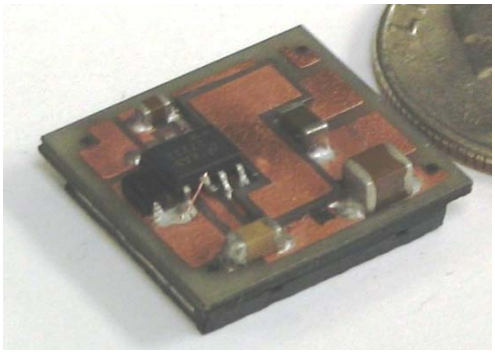

Figure 4.36. Generation 3 module, able to handle 20A output in natural convection conditions.

Verifying the simulations with a thermal camera, this module at the same 89°C as Figure 4.34 but delivering a full 20A output. The added DBC represents an additional cost but one that can be offset by eliminating the need for a fan for 20A operation. And if a fan is present, reliability and/or output current can be further increased. The bottom layer of DBC gives us an additional built-in feature of electrically isolated shield to keep the inductor layer from inducing parasitic loss in the neighboring circuit traces. In previous cases, this shield was made with silver paste on the inductor but getting this feature for "free" is even better. The thermal picture looked just like the one in Figure 4.34.

In order to put this into context, the same PCB board that was used in the efficiency comparisons (with the same devices – except packaged, and the same driver, and a commercial ferrite inductor) is shown here as a thermal map. The only real difference is that the PCB substrate is 3x larger (40x30mm) than the POL, and it uses 6 layers of 2 oz. (70 μm thick) copper. This is the typical structure of high-power PCB-based DC/DC modules so as to be a useful comparison.

The ambient conditions here were similar as for the POL case – no airflow except for natural condition and the only conduction cooling came from the interconnection wires. Figure 4.37 shows the module at close to the same peak temperature as in Figure 4.34 except that now the output current is a much-lower value of 14A. This is 30% less current despite the 3x larger board area. The extra area of the board just did not get fully utilized to allow for more output current. Under forced air conditions, more current would have been possible but not so in natural convection. All temperatures were calibrated with a thermocouple to ensure the emissivity of the camera was properly set.

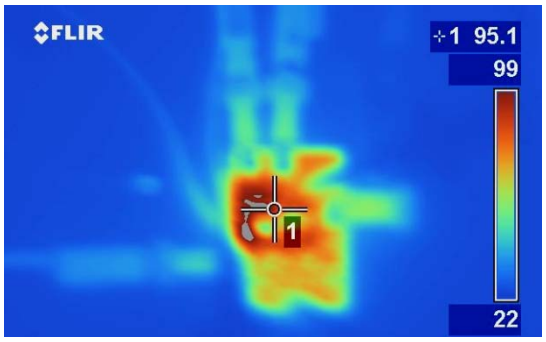

Figure 4.37. Thermal map of equivalent PCB version

Three design steps led us to a system solution that handles the electrical and thermal problems concurrently. Table 4.8 compares the three generations'

performance. The concept of layering control, active, and passive structures in a z-axis stack allows for a small footprint and high power density. In order to accomplish this, a method for embedding the devices inside DBC ceramic was shown. This frees up the top and bottom sides for stacking. An inductor layer was demonstrated, and design iterations shown, using LTCC material to complete the fully-ceramic POL solution.

Table 4.8.
Comparison of the three different POL generations

Generation:	1 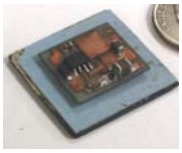	2 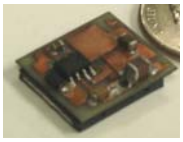	3 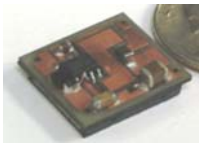
Thermal: natural convection and 50 C ambient	15A max	17A max	**20A**
Electrical:	91.5% at 5A 79.8% at 20A	90% at 10A 86% at 22A	90% at 10A 86% at 22A
Power density: In natural conv: At 20A & 200LFM:	28 x 28 x 7 mm 54 W/in^3 73 W/in^3	20 x 18 x 7 mm 225 W/in^3 260 W/in^3	20 x 18 x 7.5 mm 250 W/in^3 250 W/in^3

It has been shown that simulating a doubling of the DBC layer reduced temperatures by 6% and that also translated into a real improvement in the Generation 3 module. However, the Generation 2 case only had 1mm thickness of DBC with the dies embedded inside it. What if we were to compare a case where the dies are not embedded but rather sandwiched between equal total amounts of DBC (2 0.5mm

layers)? In Embedded Power, a dielectric solder mask was used between layers – Cookson Electronics Enthone Enplate DSR 3241BR. Assuming a polyimide of 0.25mm thickness between the layers, and the same overall quantity of DBC, how much improvement can be expected by having the dies embedded in a single layer of DBC? Figure 4.38 shows just that and reveals that a dielectric in between the DBC will reduce its cooling effectiveness also by about 6%. This thermal map is not to scale with the others so that the polyimide's low temperatures are visible.

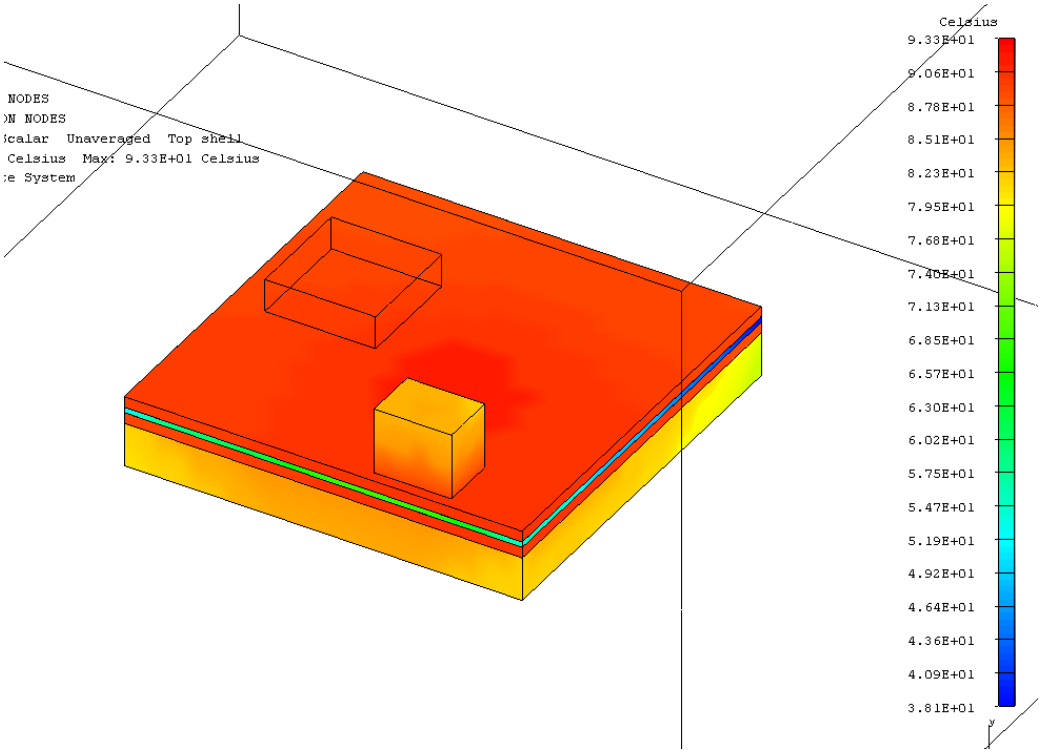

Figure 4.38. Equal thickness of DBC as Generation 2 simulation, but with polyimide layer in between two half-thickness layers of DBC.

Nevertheless, temperatures are still low considering the size and power output. This is because the FETs used in these simulations are DirectFETs which have thermally enhanced packaging. If the FETs used were the more conventional SO-8 or

plastic-encapsulated types, the peak temperatures would increase dramatically, by about 26°C due to their typical thermal resistance of 20°C/W, as discussed in Chapter 3 and Section 4.4, and loses being approximately 1.3W. Therefore, by using embedded unpackaged FETs compared with un-embedded conventional SO8 FETs sandwiched by DBC, a 25°C reduction is a drastic improvement that truly needs to be taken advantage of if utmost power density is desired.

Table 4.9.
Parametric simulation results

DBC structure	DBC thickness	FET type	Peak Temp
1 layer	1 mm	Unpackaged DirectFET	87.8°C
2 layers	2 x 0.5 mm	Packaged DirectFET	93.3°C
2 layers	2 x 1 mm	Packaged DirectFET	82.8°C
2 layers	2 x 0.5 mm	Packaged SO-8	119.3°C

So it is clear from Table 4.9 that having unpackaged embedded dies in the DBC allows for a decrease in temperatures compared with packaged DirectFETs sandwiched between an equal amount of DBC. Therefore, not only is there a packaging advantage by embedding the devices since profile is reduced and better mechanical strength, but there is also a 6% reduction in peak temperature. Also, doubling the amount of DBC also reduces temperatures by 6%. And finally, when using a FET package that is not thermally-enhanced like the DirectFET but rather of conventional SO-8 type, temperatures increase by 44%! The lesson here is that thermal resistance from the die to the substrate needs to be as low as possible and that using the maximum DBC allowed by the budget is advantageous.

What about the improvement anticipated of usinhg Stacked Power in VRM applications? As discussed in great detail in Chapter 2, the main issue of the VRM is the use of PCB material which causes hotspots that prematurely reduce its thermal capability. Using Stacked Power in a hypothetical case, simulations have been made to see what the contribution of AlN DBC could deliver. Figure 4.39 shows a comparison of the two cases, the typical 6-layer PCB one and the AlN DBC substrate one, both with same losses and components. The change in temperature is a large one – a 20% decrease in peak temperatures! Functionality between the two substrates is comparable although PCB windings would probably have to be made in PCB and then placed on top of the DBC substrate to facilitate fabrication – but this was actually done for the PCB case, as was shown in Chapter 2's case study already (since 12 transformer winding layers were needed on a 6 layer board), so this makes the two substrate materials even closer in terms of functionality.

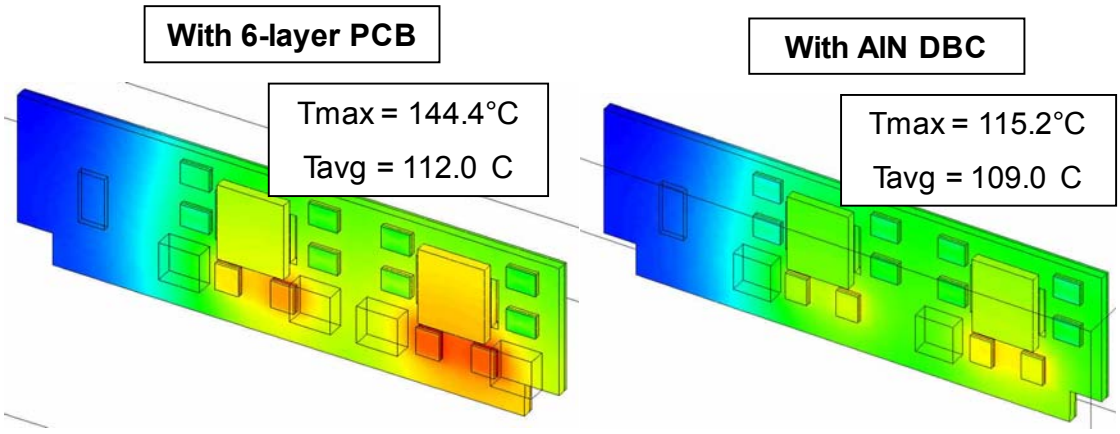

Figure 4.39. Comparison of PCB VRM with one made on DBC substrate.

Nevertheless, making a VRM using a circuit of this complexity would not be easy due to the need for multiple layers of DBC in order to have sufficient trace layers. It is estimated that 4 layers of DBC would be needed in order to have 5 trace layers and a number of vias would be needed, requiring complex registration methods. But it is not out of the question and embedding the MOSFETs in AlN DBC would make for a highly uniform temperature distribution, significantly improving the thermal performance of what would otherwise be SO-8 devices. SO-8 have 25 times higher junction-to-case thermal resistance than an unpackaged embedded FET such as the DirectFETs shown in this work.

So even when compared to ultra-high performance power supplies such as VRMs, the Stacked Power approach can be easily justified. The reduction of parasitics by means of embedding the devices to boost efficiency, and the high thermal conductivity of the substrate to reduce hotspots come together to deliver a simple solution that doesn't require large numbers of devices and can operate in more stringent ambient environments with less difficulty.

Comparing with commercially available products today, it would seem their power density is considerably higher. However, a closer look will reveal that these published figures do not include necessary external thermal management. Also, quoted efficiency figures are often for a much larger duty cycle ratio than what we have used here (24%) so this too must be kept in mind for fair comparisons.

The advantages and technical highlights of this methodology are:

1. Integrated thermal management *replaces*:
    1. PCB with high-performance ceramic substrate.
    2. Bulky and expensive external heat sink.
    3. Cooling fan up to 55°C ambient and 20A!
2. Active devices and inductor layer are integrated together
    - ❖ Integration allows for small size, low profile and short signal paths
3. Parasitics are reduced for less loss at high frequencies
    - ❖ High frequency reduces size of inductor and output capacitance
4. At 1.2 Vout, power density of 260 W/in^3 achieved so far
    - ❖ This includes *all* necessary thermal management!
5. High efficiency with ultra-low profile of 5 mm.
    - ❖ 12V-1.2V efficiency of 86% at 10A and 83% at 22A
    - ❖ 5V-1.2V efficiency of 90% at 10A and 86% at 22A

Chapter 5. Assessment of Environmental Mechanisms

5.1 Modeling of package environment

A POL converter's environment has a critical impact on circuit performance. The amount of airflow present, the ambient temperature, the proximity of other circuits that either block, modify or enhance airflow, the proximity of heat-producing circuits, the temperature of the platform (or motherboard) the POL is connected to, and the orientation of the POL relative to air stream and gravity can all play a significant role in whether or not the circuit can achieve maximum output reliably, and for lengthy life spans. Running it at anything less than full power is called "derating" and is generally based on airflow speed and ambient temperature only, but bear in mind that the other factors mentioned can be just as important in system-level implementation.

This information is relevant from the system-level engineer on down. In the past, thermal analysis was simply a hardware test once the circuit was designed but this is not the way to be competitive today. In order to make a high performance circuit in today's climate, it must be thermally efficient in order to push integration. For this to happen, the thermal design must be kept in mind from the very beginning of the design process. Some of the questions answered in this chapter are:

- How much heat can be dissipated by a particular type of package?
- What are the mechanisms involved in this limitation?
- How far can we push power output for a given efficiency and airflow?

- What type of mounting is best?

Modeling a package's environment requires a number of assumptions. Airflow is a highly complex mechanism with chaotic characteristics and the environment that the POL will be located in is as varied as imaginable. The air flow regime is generally "turbulent" and thus air flow rates are different at all points on the board. Also components that block air flow on the surface of the POL contribute to this uncertainty, as discussed in Section 2.4. I-DEAS FEA simulations were used to determine the reduction in airflow around blocks and are the basis for the assumptions given for Figure 5.1 below. But the model discussed in this chapter was developed using MathCAD software so analytical calculation methods could be used to derive results for various package sizes efficiently.

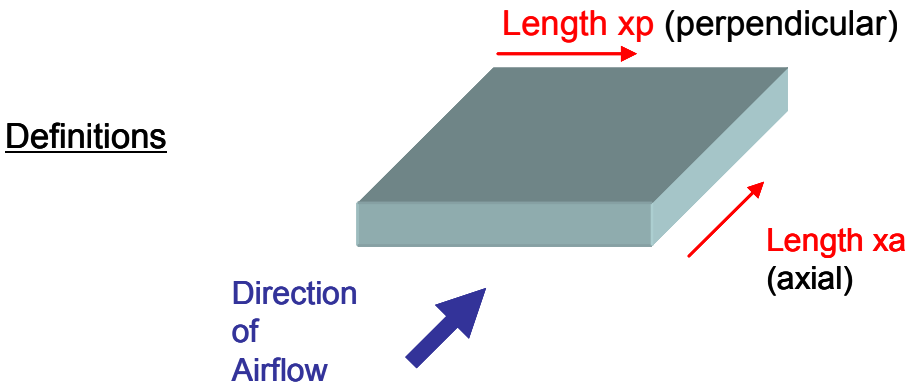

Figure 5.1. MathCAD model definitions

- Assumptions:
  - Uniform volumetric heat distribution inside the package
  - 200 LFM (1 m/s) uniform airflow
  - For bricks, sides get 50% of airflow

- For POL, sides get 75% of airflow
- For both, back gets 30% of airflow
- Ambient is 55°C

In addition to commercially-available POL packages, we will also look at telecom industry “bricks” since the smallest one (1/16 brick) is essentially the size of a typical POL buck and the quarter brick is typical of telecom applications. The temperature of the board in the model requires complications. It would be easiest to assume the board to be of uniform temperature but this is simply not the case in practice and assuming so is to introduce excessive error. As a result, the temperature of the board is assumed to be divided up into two parts, based on thermal maps found in industry literature. One part is called the “hotspot” and is representative of the device loss is determined to be a peak temperature of 130°C. The second part is the rest of the board (PCB or plastic, both of which have a low thermal conductivity which guarantees decoupling with the hotspot area) and is assumed to be 85°C.

The remainder of the modeled system is the motherboard that the converters are connected to. In this case, the size of the PCB they are mounted to would vary the potential cooling rates since it too sees airflow cooling, so it must be of fixed size to compare all the packages directly. US standards have already been implemented for industry to be able to do the same thing, called the JEDEC 51-7 standard [133]. This dictates the following PCB sizes depending on the converter being mounted to them and is the same standard I used in my modeling. There are two sizes to choose from, depending on the size of the module to be mounted. This is shown in Figure 5.2.

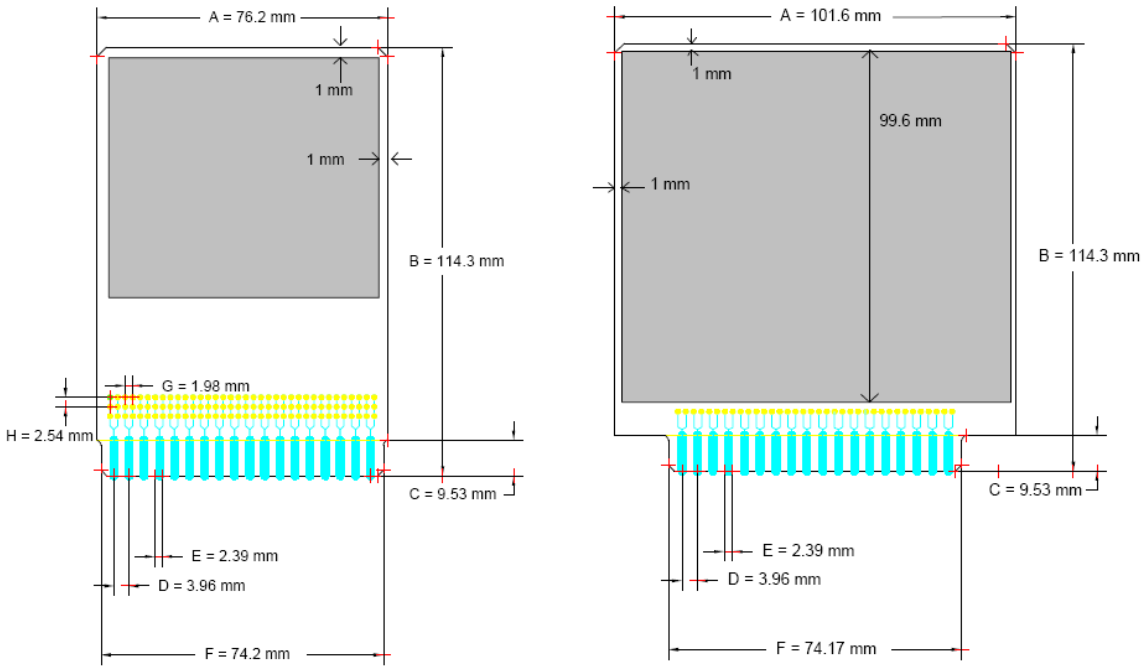

Figure 5.2. JEDEC 51-7 standard PCB sizes for POL converters used in thermal testing. One POL package per square PCB is used, the size of which is chosen based on the POL module length dimension [133].

For the analytical MathCAD model, equations must be used to calculate the theoretical thermal performance of the modules. The modules, as mentioned, are modeled as blocks where each side is assumed to be flat. In forced convection conditions for isothermal flat plates, analytical expressions are derived from empirical data. This formulation is based on a number of empirically-obtained "numbers" such as the Reynolds (R_{ea}), Prandtl (Pr) and Nusselt (Nu_a) numbers. These parameters are based on characterized properties of air and are used to obtain the average heat transfer coefficient (HTC or h_{ca}) in order to determine how much heat can be removed from the flat surfaces (q_{ca}). For the board surface temperature, the peak value should be used for Stacked Power since the substrate temperature is within a few degrees of peak. For PCB substrates, an average temperature should be used for better accuracy

and representation. In typical POL environmental conditions of 25-55°C ambient air and 100-400 LFM airflow, the following analytical forced-convection expressions can be used [89]:

$$R_{ea} := \frac{U \cdot \rho \cdot x_a}{\mu} \tag{30}$$

$$Pr := 0.71 \tag{31}$$

$$Nu_a := 0.664 R_{ea}^{0.5} \cdot Pr^{0.3333} \tag{32}$$

$$L_L := \frac{x_a \cdot x_p}{2 \cdot x_a + 2 \cdot x_p} \tag{33}$$

$$h_{ca} := \frac{Nu_a \cdot k}{L_L} \tag{34}$$

$$q_{ca} := (x_a \cdot x_p) \cdot h_{ca} \cdot \Delta T \tag{35}$$

Where,

- U is airflow speed (m/s),
- ρ is air density (0.968 kg/m^3),
- x_a is axial length (relative to airflow direction) of the converter (m)
- μ is absolute viscosity of air (20.8e-6 Ns/m^2)
- x_p is the perpendicular length of the converter (m)
- Pr is the Prandtl number which is generally 0.71 (no units)
- L_L is the characteristic length, calculated for each surface (m)
- k is the thermal conductivity of air, in this case 0.0293 W/m°C
- ΔT is difference between peak and ambient surface temperatures. For this example, it is 130°C - 25°C = 105°C.

Figure 5.3 shows the variation of average HTC (in units of W/m^2°C) and possible heat removal q (in Watts) for an 18 x 18 x 4 mm POL module, the typical size of one phase of a VRM or of a non-isolated POL, using Equations 26-31, under the conditions listed above, and for the MathCAD models shown. For these conditions, the Reynolds number ranges from 300-900 depending on which surface you are considering. This is a low value in absolute terms but this is because the surface lengths involved here are only on the order of 18 mm.

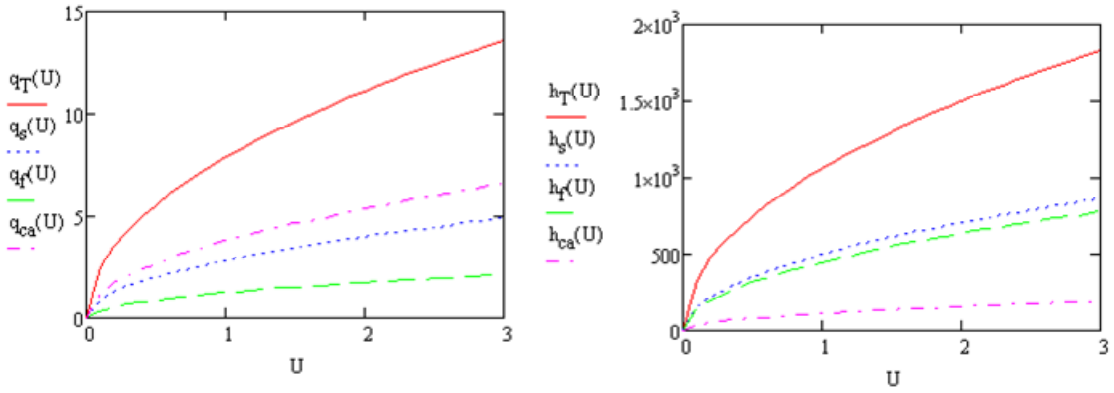

Figure 5.3. Variation of hca and qca with airflow speed (m/s), for the POL.

Where, q_{ca} and h_{ca} are off the POL's top surface,

q_s and h_s are off the POL's sides

q_f and h_f are off the front of the POL facing the airflow

q_T and h_T are the sum totals of all the surfaces

per the definitions of Figure 5.1

Maximizing the convective cooling of the board involves maximizing the average HTC so that the heat removal capability, q_{ca}, can also be maximized. In natural convection conditions, the HTC is generally around 20 W/m^{2}K but with forced

convection, it can be anywhere from 100 to several hundred, depending on the surface area of the board, and its location/orientation relative to airflow rate [116]. This means that HTC varies locally on the board, with the amount of surface obstructions present and airflow angle. Also note that heat removal rate is only a few Watts for these small modules.

5.2 Package Design Characteristics

There are two main methods of interfacing the converters with their motherboard. The first is the most common today and consists of the use of pins for mounting to the motherboard. These pins vary in size and length depending on the module in question and the orientation of the converter changes depending on the location of the pins. In some cases the pins are located at the edge and in the same plane as the substrate. This is called a “blade” arrangement because the converter module is perpendicular to the platform. The other way is to have the pins perpendicular to the converter so that it mounts horizontally like a surface mount component.

The second method is a thermal pad interface. This means there is a large area soldering pad on the bottom side of the converter. In this case, the converter mounts flush with the platform and must be reflowed due to the central location of this thermal pad. It is often tied to ground to double as a ground plane for the converter circuit and so care must be taken to ensure the thermal vias present on the platform are only connected to ground planes. Figure 5.4 shows three examples, each one of the three different mounting methods outlined above.

Figure 5.4. Showing surface mount with pins and surface mount with thermal pad, respectively.
Used with permission from Emerson Technologies and Linear Technology

These are the same two cooling mechanisms for all three types of packages. All of them have both convection and conduction cooling. Radiation cooling is a third possibility but in the typical environments, it is negligible compared to the other two. However, each of the three package types have varying ratios of convection and cooling rates, and each have advantages and disadvantages associated with them. Figure 5.5 shows a schematic representation of the two cooling mechanisms.

As you can see, both packages have red heat arrows coming off the PCB as well as their own surface but in the case of surface mount with pins, there is more convection off the package and less heat going into the PCB because of the increased thermal resistance of the pin connections. In the case of surface mount with pad, the much-larger area of the pad relative to the pins yields higher heat flow into the motherboard than is convected off the surface of the package.

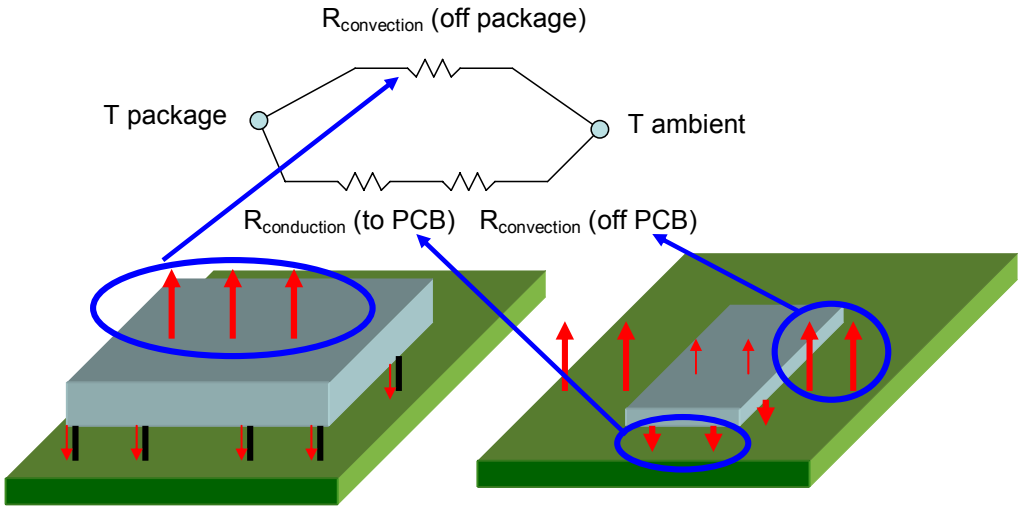

Figure 5.5. The two cooling mechanisms for each type of package.

The thermal maps use three examples based on real product hardware. The QFN58 package is represented by the Enpirion EN5365 POL module, my POL is my Generation 2 Stacked Power module and the 1/16 brick is a NetPower product. These three are shown together because they are of similar size, as shown in Figure 5.6.

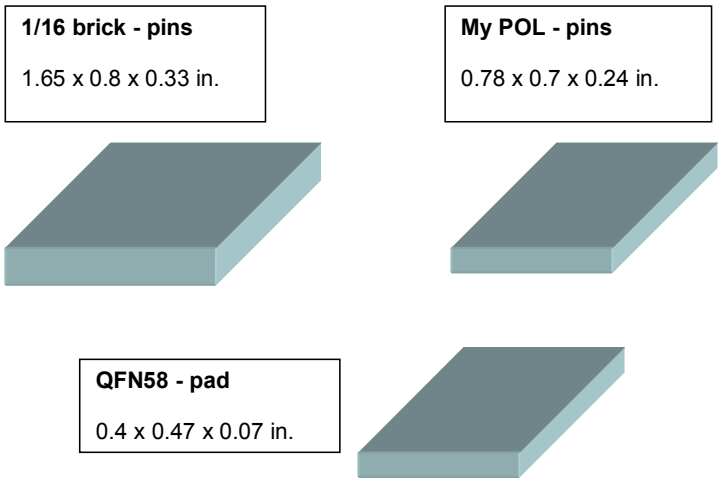

Figure 5.6. Package sizes used in the thermal maps.

A result of the thermal MathCAD model for the typical 200 LFM environmental airflow rate is shown for all three package types in Figure 5.7. The ambient condition in this case is 55°C and the motherboard the converters are at 85°C. However, the QFN and brick packages are assumed to be of low-thermal conductivity material: approximately 1 W/m°K for plastic and ~36 W/m°K for 6-layer PCB (see Chapter 2 for details), compared to 150 W/m°K for AlN DBC, and thus are assumed to have 20% hotspots as is typically seen in product thermal application notes. The 20% peak region is at the typical maximum temperature of 130°C and the rest of the surface is at 85°C average. The Stacked Power POL's DBC substrate has such high conductivity that from the thermal pictures shown in Chapter 4, it can be assumed to be uniform.

The contribution of conduction cooling in the various packages is a combination of thermal resistance through the motherboard interface (either pins in parallel or a thermal pad) and the convection resistance off the motherboard. For eight pins that extend 3 mm below the board and are 0.75 mm in diameter, made of tinned copper, the conduction resistance is on approximately 2.4 °C/W based on Equation 24. Then they connect to the PCB/motherboard which it is also convectively cooled – just like the POL surfaces. Therefore, Equations 25-31 are once again used for the PCB surface.

As shown in Figure 5.5, these two thermal resistances: 1. conduction to PCB and 2. convection off PCB, are added together and compared with direct convection off the POL surface. These represent the two parallel cooling paths shown. The thermal resistances are treated in the same manner as electrical resistances, so series ones add and parallel ones are the reciprocal of the reciprocal sum.

Equation 32 shows how to calculate the convection resistance of the PCB to the ambient, $R_{pcb\text{-}a}$. The heat transfer h_c is calculated using equation 26-30 just shown, and the given area of the JEDEC-51 standard size PCB described in Figure 5.2 [133].

$$R_{pcb-a} = \frac{1}{h_c A} \tag{36}$$

Then the conduction through to POL pins to the PCB is described by Equation 33, in much the same way as Equation 24.

$$R_{pin} = \frac{L}{Ak} \tag{37}$$

Then moving on to the POL itself, its convection to ambient resistance is defined as Equation 34 where we have the temperature difference between the surface of the converter and the ambient air divided by the sum of the heat that each side can dissipate.

$$R_{POL-a} = \frac{\Delta T}{q_{ca} + q_s + q_f} \tag{38}$$

So the total resistance is the parallel combination of the contributions from the PCB and the POL.

$$R_T = \frac{(R_{pcb-a} + R_{pin})R_{POL}}{(R_{pcb-a} + R_{pin}) + R_{POL}} \tag{39}$$

Where, q_{ca} and h_{ca} are off the POL's top surface,

q_s and h_s are off the POL's sides

q_f and h_f are off the front of the POL facing the airflow

per the definitions of Figure 5.1

Under these conditions, Figure 5.7 shows the large thermal pad of the QFN58 package allows it to dissipate more heat than even the much-larger 1/16th brick because the pad's area is significantly larger than the sum of the cross-sectional areas of the pins. Taking a look at the components of the total heat, you see that the 1/16th brick dissipates more through convection than conduction. On the other hand, the QFN package dissipates almost all its heat through conduction and very little through convection because the plastic casing causes localized hot spots and the module is so small that it has little surface area.

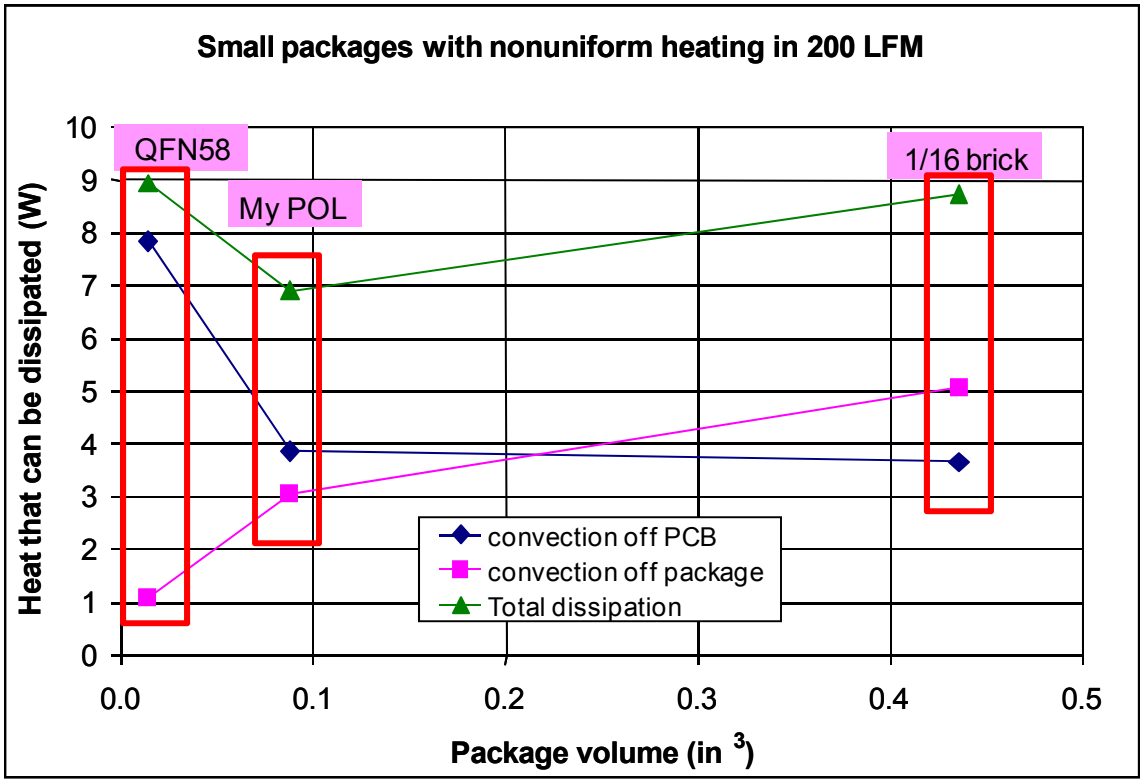

Figure 5.7. Theoretical heat that can be removed from package in 200LFM and 55°C ambient conditions obtained from MathCAD analytical calculations.

So the QFN package transmits 88% of its heat into the motherboard, which in this case follows JEDEC 51 Standards: 76.2 mm x 114.3 mm. However, the motherboard is also limited by how much heat it can convect, the extent of which depends on the motherboard's temperature difference with the ambient air (delta T). For this standardized size, as an example, Figure 5.8 shows the limitation. Comparing this with Figure 5.7 makes it clear that to get the most out of a thermal pad connection, the motherboard has to be 21°C hotter than the ambient for 8 W heat transfer off PCB, whereas with my POL, it only needs to be 10°C since the POL only requires 4 W off the PCB.

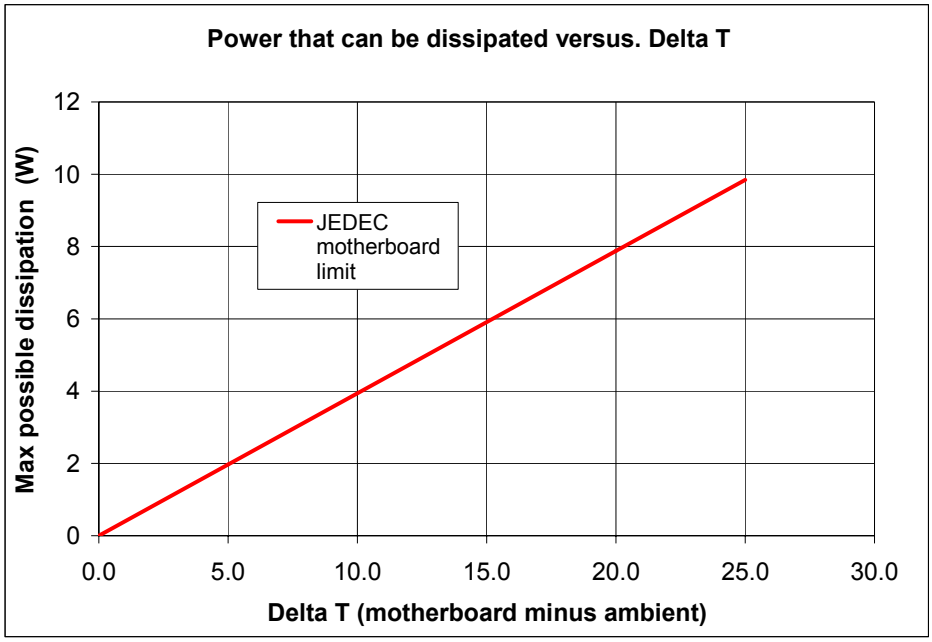

Figure 5.8. Dissipation limit of JEDEC-51-7 motherboard as obtained from MathCAD analytical calculations.

What this means is that if the motherboard has low heat transfer capability due to a low delta T or a small amount of surface area, the QFN package will have to be derated much sooner than the Stacked Power POL will. The longer derating can be

pushed back, the more flexibility the system designer has for thermal management, location of the module inside the bigger system, and more utilization of available capability.

Comparing the QFN package to my POL, the 72% increase in footprint is perfectly acceptable for the 300% higher module's convection capability you get in return. Also note that despite the 500% greater volume of the 1/16th brick, the Stacked Power POL is able to get within 81% of its power dissipation capability! And this is all because the PCB has been replaced by a substrate that doubles as a heat spreader.

The other packaging aspect to consider is the casing. Having a plastic encapsulating material on the module significantly reduces its cooling ability, as discussed. Same goes for using PCB substrates. Open-frame modules are very common in the telecom industry so there is no problem with this concept but, in the POL applications, a small plastic module has better marketing effectiveness, despite the potential thermal penalty just described. The end result is that package utilization is lower than it could be. If the plastic is necessary for reduction of moisture contamination and voltage breakdown, similar capability can be afforded by dip conformal coating the Stacked Power open-frame module and is commonly done for telecom applications of much higher voltage than POLs [116].

Comparing actual modules with their theoretical dissipation capability shows this to great effect in Figure 5.9. The QFN package can only be used to 1/3 of its capability and the 1/16th brick can achieve 2/3s but these are much lower than the 84% demonstrated by the Stacked Power module. The only reason Stacked Power

isn't 100% is due to the plastic and ceramic surface mount components, which reduce the available surface area for direct convection cooling – but this is basically inevitable.

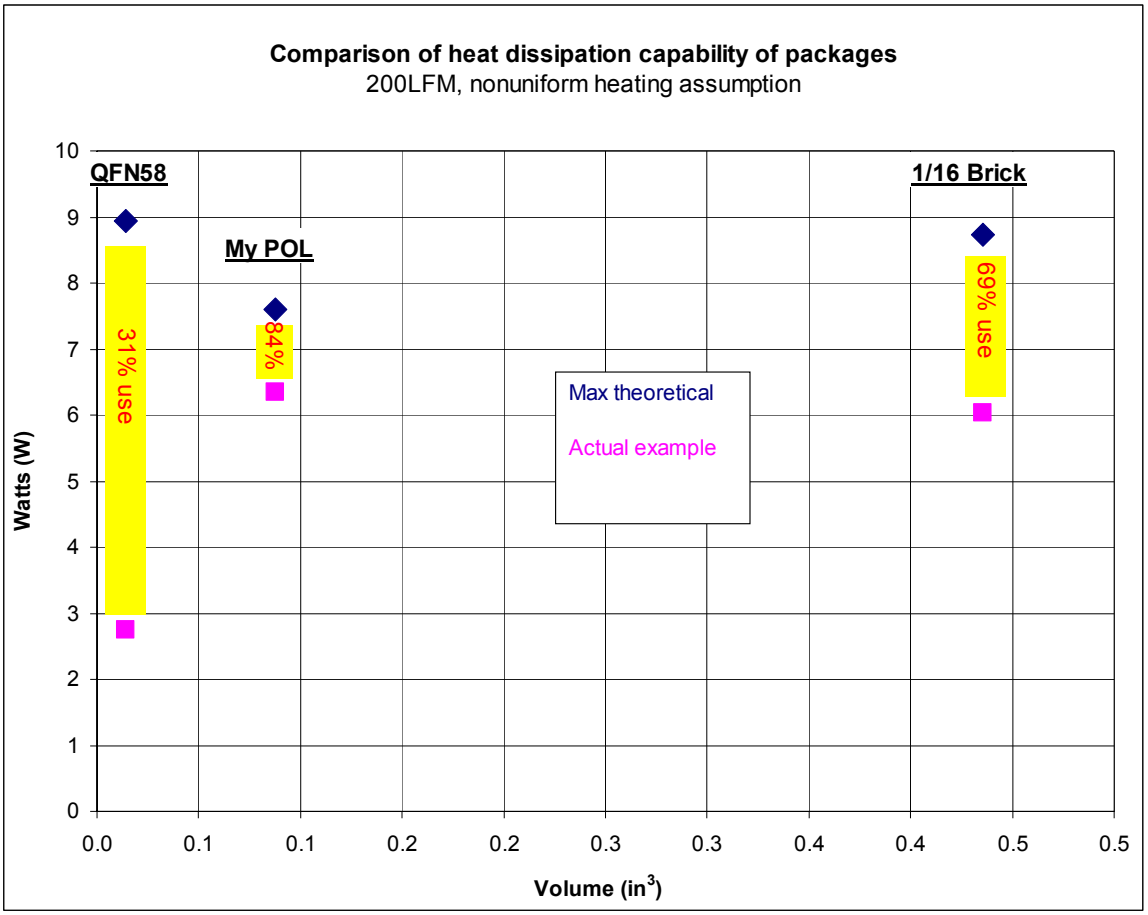

Figure 5.9. Disparity between package utilization of typical examples

Another interesting answer we can get from the MathCAD model is what is the best shape for a POL to have the highest convection efficiency? Assuming equal area, should it be square, rectangular, long axially, long in the perpendicular direction? Well it turns out that for 200 LFM and a delta T of 105°C, the amount of heat q_{ca} that can be dissipated, sweeping all possible length and width combinations

from 5 mm to 20 mm (so 400 mm^2 maximum surface area) with 1 mm increments, is highest, for the smallest footprint, when the module is narrow and long with airflow direction down the axial length.

This is shown in Figure 5.10 where, for two modules of equal area, a 5 mm long by 20 mm wide module can only dissipate 1.2 W, whereas a 20 mm long by 5 mm wide can do 4 W – for the same footprint! The equations used are numbers 26-31 defined in Section 5.1, with the exception of using range variables for both width and length and making the equations functions of both width and length. So to achieve a minimum footprint from a thermal perspective, it is best to have the module be long and narrow with airflow going in the long direction. This information will be used for the 2-phase POL design discussed in the following chapter where it will be designed to be long and thin so that axial airflow can have maximum benefit and thus extract maximum cooling capability, in order to maximally reduce package size.

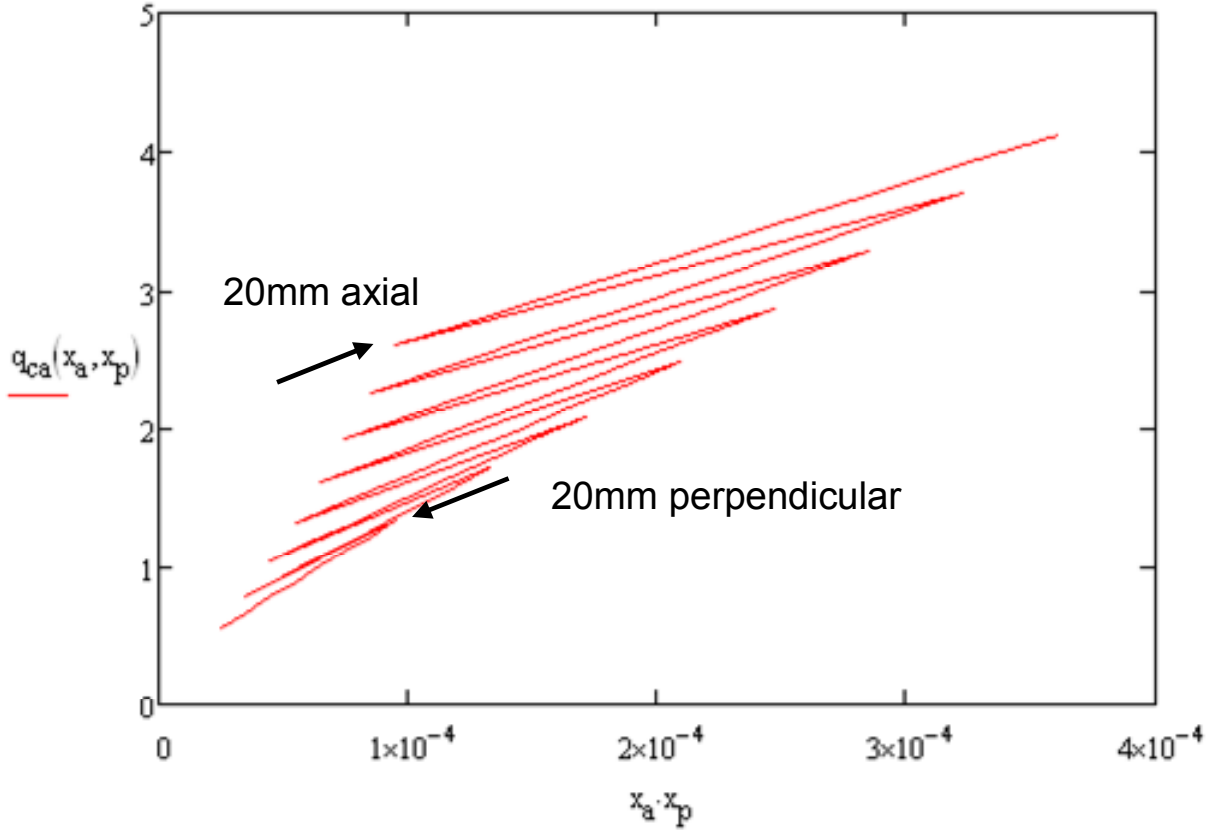

Figure 5.10. Power dissipation possible sweeping both length and width from 5mm to 20mm in increments of 2mm.

However, it must be pointed out that Stacked Power does have some potential drawbacks. For one, the thermal pad arrangement, provided the motherboard circumstances are favorable, can handle more heat loss than a pin arrangement. Also, the pins need to be chosen for high-current capability i.e., of sufficient thickness as to not cause too much I^2R loss. Although to be fair, many of the packages that use a thermal-pad still have very small interconnection points. In addition, telecom bricks generally use pins for even 100A output so the problem is apparently not a difficult one to overcome.

Manufacturability of Stacked Power would be very similar to what has already been done for many years in the aerospace power electronics field. Many aerospace power supplies are built on a ceramic substrate with thick-film interconnetions. Stacked Power is very similar so it is therefore not difficult to apply this technology today. The difference with Stacked Power is that the design concept is based on layering. By embedding active chips into the substrate, we can have additional layers on the top and bottom. This increases integration levels, decreases path lengths, reduces parasitic inductance and resistance, reduces switching loss, increases efficiency, improves thermal performance, reduces temperatures – and all this allows us to make it small in the first place.

5.3 Stacked Power Scalability

The Stacked Power process is a flexible one that can handle high-temperature and high-voltage applications. As such, it is amenable to a wide range of materials and devices. Here, as an example of this flexibility, a 10 A, 600 V SiC Schottky diode from CREE was packaged in CPES by Jon Claassens et al. [114] using the Stacked Power principles.

The silicon carbide die is 2.1 x 2.1 mm and 15 mil thick and poses a packaging problem because of its low Coefficient of Thermal Expansion (CTE); on the order of 4 ppm/°C. This low coefficient poses difficulties when finding conductors to mate it to the substrate. If the CTE of adjoining materials is too high compared with that of SiC, there

is a high probability that the die will crack in thermal cycling conditions because rigid interconnections will not allow for differential expansion rates – and silicon carbide is inherently brittle due to its high hardness capability.

Figure 5.11 shows the relationships between SiC and potential conductors in terms of electrical conductivity and CTE. The ideal would be as high a conductivity as possible and yet with a CTE equal to that of SiC. In this graph, it is clear that are only three main contenders even come close to realizing this. Molybdenum, tungsten and chrome are closest in CTE and although their conductivity is not very high, this can be tolerated more easily than a CTE mismatch.

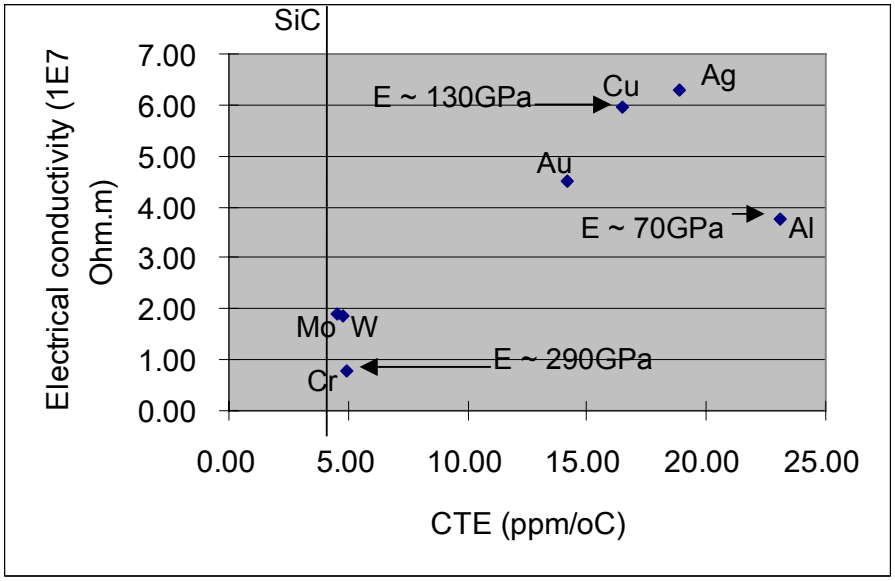

Figure 5.11. Comparison of CTE of conductor materials with SiC [114].

In addition to these considerations, the resistivity of the material in also important because an excessive voltage drop obviously increases the losses that must be dealt with, which directly impacts thermal performance. Molybdenum is chosen as the

interface material because it is easily obtained, has reasonable resistivity, low CTE, and can be easily made into the desired shape and thickness for the application.

Thermal conductivity is another important parameter but the Wiedemann-Franz law states that the ratio of electrical conductivity to thermal conductivity is more or less constant. Therefore, the material was chosen based on the electrical characteristics and CTE. Figure 5.12 shows the relationships between thermal and electrical conductivity for a range of metals.

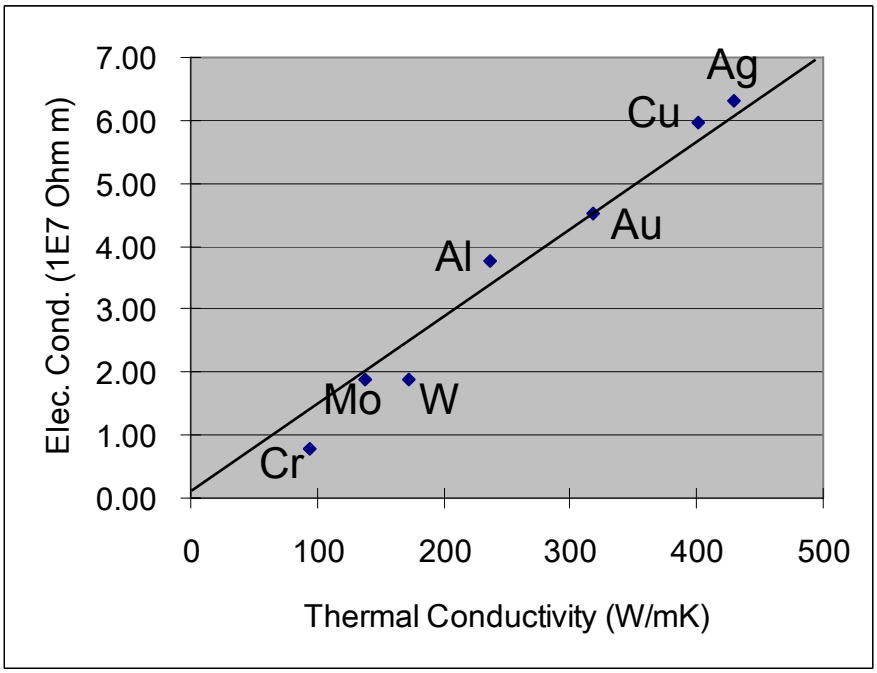

Figure 5.12. Wiedemann-Franz relationship of conductor materials [114].

The choice for substrate materials is essentially the same as what was discussed in Chapter 2. There are three main options:

- DBC (Direct Bonded Copper) on AlN or Alumina.
  - The AlN has great thermal conductivity

- DAB (Direct Bonded Aluminum)
- Hybrid technologies (printed Ag/Cu on low CTE ceramic)

AlN has the closest CTE match with SiC and so is the best choice for this high-power application (4.6 ppm/°C compared with 4.5 ppm/°C respectively). This will most likely be a huge advantage of Stacked Power as SiC use becomes more widespread. Molybdenum was chosen as a shimming material because of its CTE of 4 ppm/°C, which is again a close match. The package structure is shown in Figure 5.13.

The final material to be chosen for this application is the solder. Since the operating temperature is on the order of 250°C, the reliability of the solder is critical. Another important parameter is the softness so that even small CTE stresses can be absorbed by the solder interface. CPES has developed a nano-silver paste for this type of application that meets both of these criteria [87].

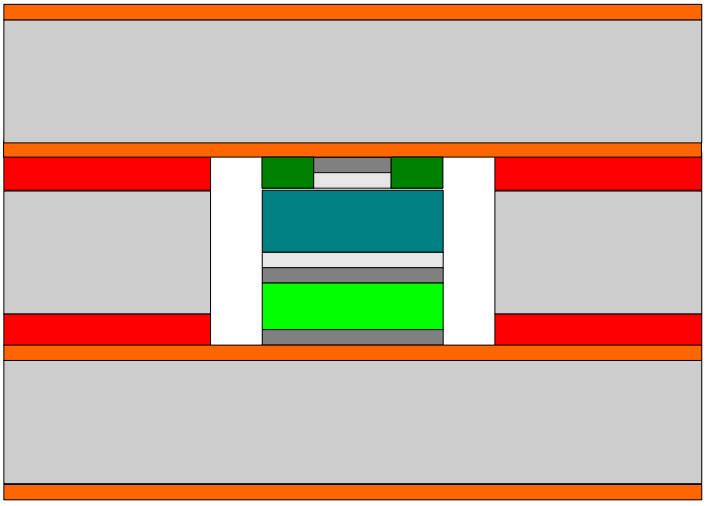

Figure 5.13. Schematic representation of the diode package [114].

A quick overview of the manufacturing process to make this Stacked Power module is similar to that of the low-voltage POL since both use embedded active components inside AlN DBC ceramic. The solderable dies are placed inside cut ceramic, attached to the shim, and interconnections are made with the nano-silver paste. The difference with the POL is the type of solder used – nano-silver paste was not used with the POL because the low-votlage FETs could not handle the 250°C curing temperature without failure, so a low-temperature PbSnBi solder was used instead. The SiC structure is schematically shown in Figure 5.14 along with an actual photo of the inner layer and final hardware.

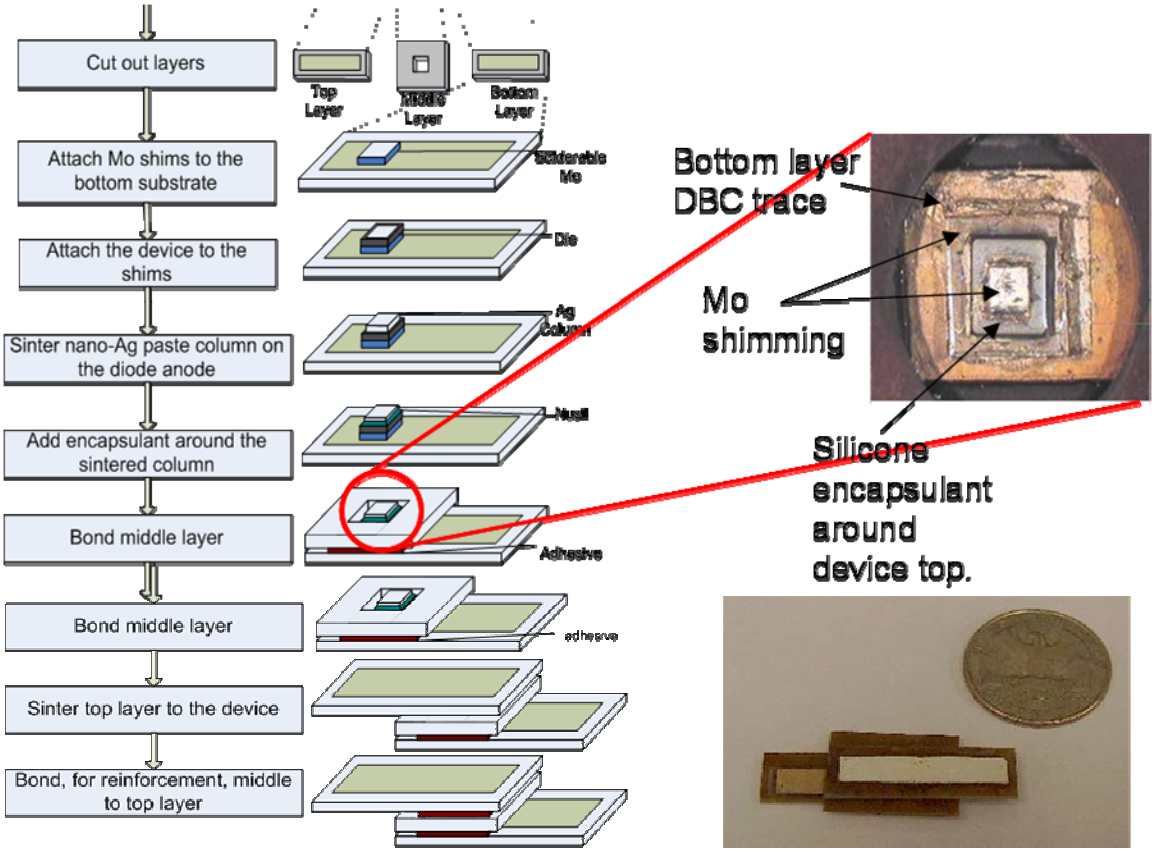

Figure 5.14. Process steps with actual pictures of hardware [114].

The final outcome is very positive, although more high-voltage testing would be needed in the future to determine the true voltage breakdown performance of this structure. Thermal cycling tests from -55°C to 250°C in the CPES thermal chamber revealed that 250°C operation with 50% duty cycles showed no change in the forward I-V characteristic of the diode. Stacked Power proved to be a worthy packaging concept for this example application.

5.4 Stacked Power and VRMs

The efficiency graphs for the Stacked Power 3D POL converter, including driver loss, are shown in Figure 5.15. As we can see, we get 12 V - 1.2 V efficiency of 87% at 14 A and 85% at 23 A and for the low voltage input case, 5 V - 1.2 V efficiency of 91% at 10 A and 88% at 24 A. These efficiencies are higher than what is currently found in industry, thanks to a design with ultra-low parasitics and thoughtful layout.

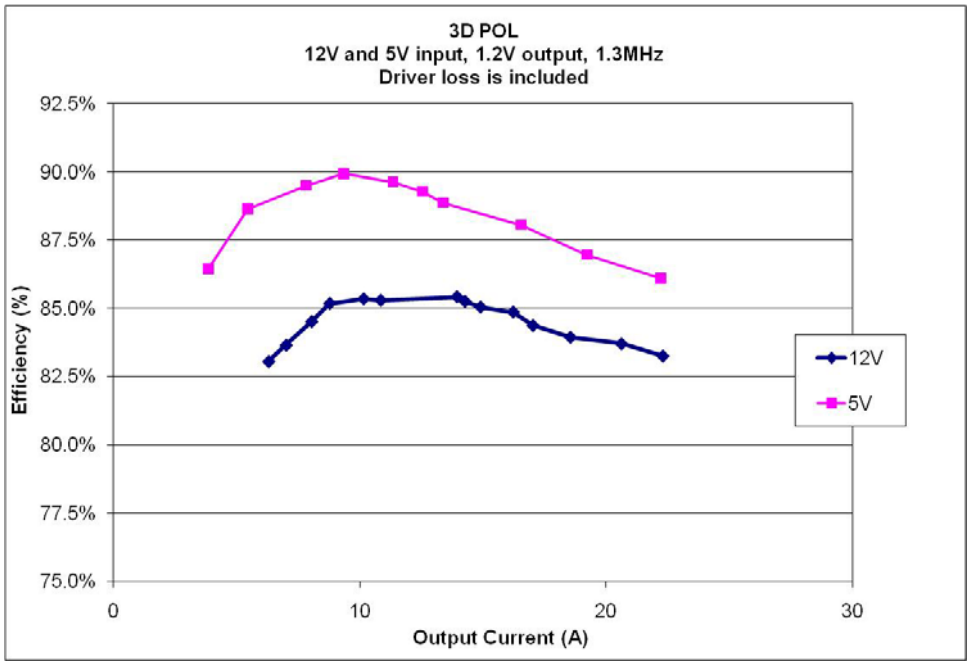

Figure 5.15. Efficiency curves for 12V and 5V input, including driver loss.

At a heavy load of 23A, efficiency is just above 83% which beats commercial state-of-the-art 12V input POLs, such as the Power One ZY7120 which achieves 81% at its 20A maximum output and the Artesyn PTH12020 which achieves 82% at its maximum of only 18A output. Even when looking at specialized VRM POLs, the Artesyn VRM10-105-12-EJ 4-phase module is good for 105A, or 26A per phase, and achieves 84% at 1.325V. So at 1.2V the efficiency would be on par with our integrated module and low-profile inductor.

What about the improvement anticipated of using Stacked Power in VRM applications? As discussed in great detail in Chapter 2, the main issue of the VRM is the use of PCB material which causes hotspots that prematurely reduce its thermal capability. Using Stacked Power in a hypothetical case, simulations have been made to see what the contribution of AIN DBC could deliver. Figure 5.16 shows a comparison of

the two cases, the typical 6-layer PCB one and the AlN DBC substrate one, both with same losses and components. The change in temperature is a large one – a 20% decrease in peak temperatures!

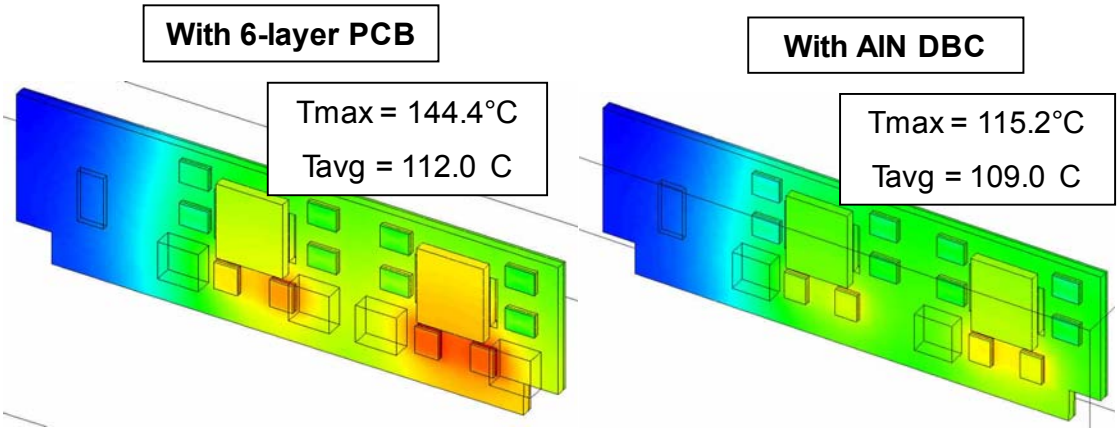

Figure 5.16. Comparison of PCB VRM with one made on DBC substrate.

So even for stringent applications like VRM, the converter does extremely well due to low top switch parasitics keeping switching losses in check even at these low duty cycles. Stacked Power module offers a convenient and highly effective solution to high-current applications since it not only offers full 12V-1.2V conversion capability but is also perfectly capable as the second stage of a two- stage design concept.

The two-stage concept is one where the 12V-1.2V conversion is accomplished in two parts. The first part is a 12V-6V conversion using a switched capacitor network operating at very high frequencies since it can do it efficiently thanks to the fact it has no magnetics. It features extremely high power density of 2500 W/in^3, 97% efficiency, 45% footprint reduction compared with typical designs. This CPES design by Dr. Ming Xu is shown in Figure 5.17 [113].

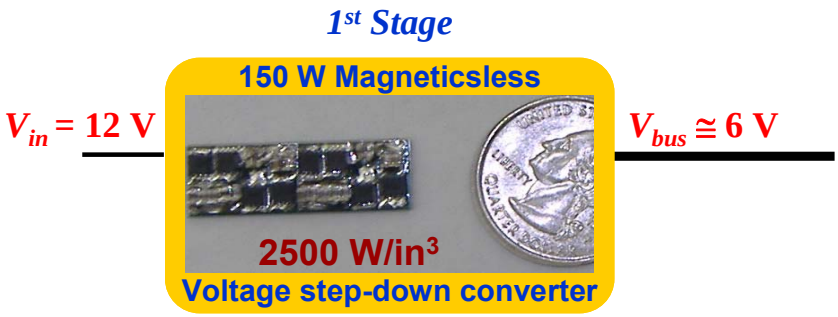

Figure 5.17. First stage in a 2-stage conversion concept for VRM and other high power applications [113].

The second part is a high-current stage that operates at lower frequency for highest efficiency and can be multiple modules, each optimized for their own load. One example of this is the Stacked Power module. Figure 5.18 shows the general structure for a second-stage implementation, with various possible POLs intended for a computer system.

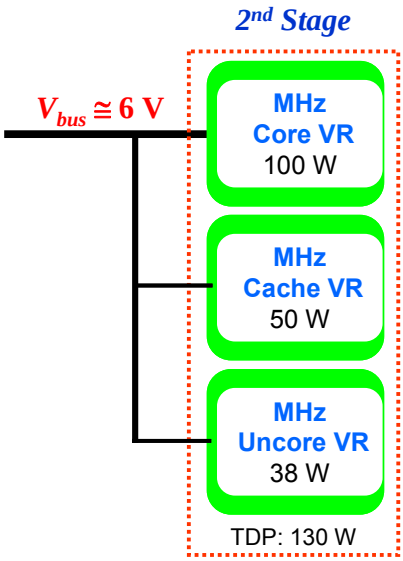

Figure 5.18. The second stage implementation with POLs for each type of load [113].

Stacked Power fabrication techniques are very well suited to this second-stage concept. As demonstrated in Chapters 3 and 4, a high-performance module of small

size can be made for point-of-load applications that need to be close to their load for optimal performance. This load can be various possibilities as shown in Figure 5.18, and in many cases, the power loss of this load can mean that the ambient environment in its vicinity has high temperatures. Therefore is best to have a POL that can effectively handle heat, as Stacked Power has demonstrated in Chapter 4.

Chapter 6. 2-Phase Design with Coupled Inductor

6.1 Coupled Inductor Design Philosophy

As discussed in Chapter 1 and 4, the integration issue with inductors today is their size. Integrating the active components is popular today in the low-power sector, as evidenced in Chapter 1. The external inductor however, tends to be much larger in size compared to the integrated modules and thus much harder to integrate, as shown in Figure 6.1. Finding ways to make the inductor smaller is a real challenge in power electronics because the power level is high. Miniature inductors have been around in radio frequency applications such as cell phones and wireless communication components but they operate only at very low power. For power electronics applications, this is still a very new field.

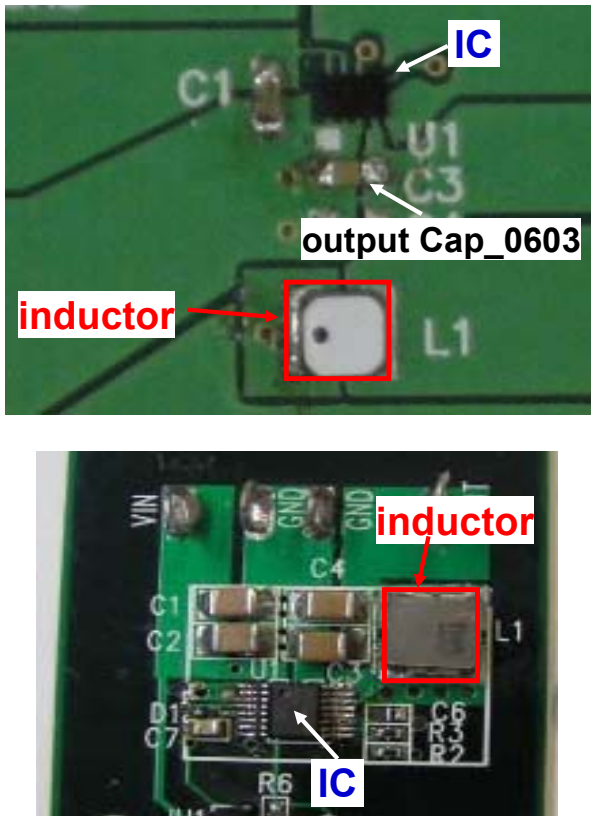

Figure 6.1. Demoboards for two integrated converters with their respective output inductors shown for size comparison.

The ideal integration scheme would be to have that inductor size similar to that of the silicon, known as chip-scale integration. This has only recently been demonstrated in industry and only for very low output current levels (<0.65A) so far, as discussed in Chapter 1. What we will demonstrate will be a similar inductor-to-active-component ratio but for 40A (!).

So what about at high power? The Stacked Power approach has demonstrated a highly effective method of combining all the components together for high power in a small package. In addition, the layering approach gives room for easy integration of multiple phases together using a unit cell structure active layer. The multi-phase

approach has been successfully used for several years, thanks to its introduction to industry by CPES. Figure 6.2 shows how the multi-phase schematic is implemented for high-current buck applications.

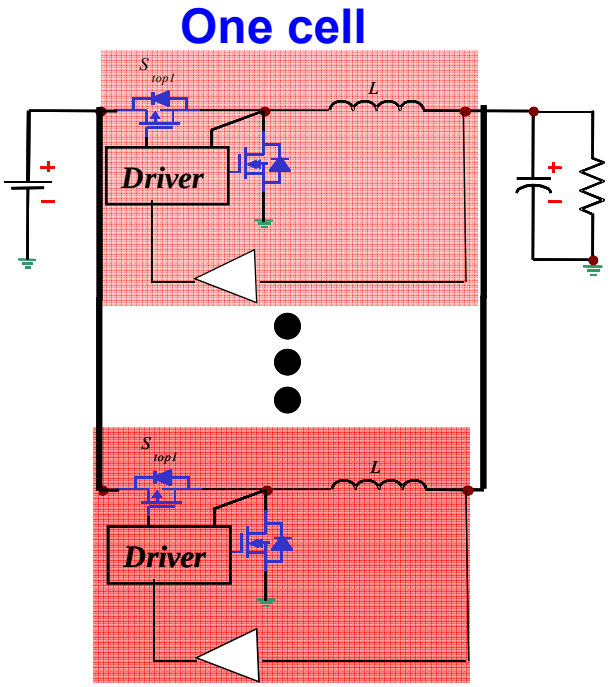

Figure 6.2. The multi-phase approach using modules with integrated inductor.

In order for this concept to work effectively, the inductor must be small. Can the LTCC process be used to make multiple inductors in one single unit cell? It turns out that it can be made very effective with careful design and fabrication.

6.2 Coupled Inductor Design Methodology

The general design concept is this. Let's also assume that 2 phases are necessary to reach the desired output current. The amount of coupling between the 2

phases is called the coupling coefficient (generally shown as “α”). We want to have inverse coupling (so the sign is negative) so that we can reduce the size of the necessary magnetics since the fluxes cancel in the outer legs. This concept and its associated schematic are shown in Figure 6.3 [119].

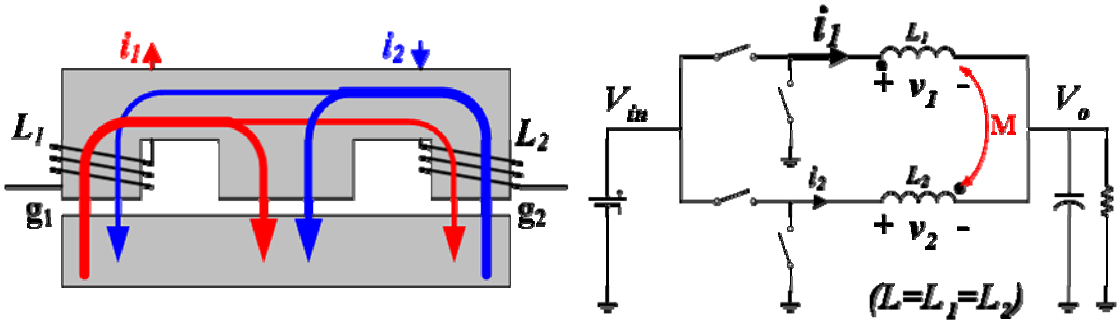

Figure 6.3. 2-phase coupled-buck concept

Perfect coupling is equal to a value of 1, which offers no inductance, essentially making it a transformer, negates its use as a buck inductor even though it would allow for fastest transient response. Therefore, the coefficient should be less than 1 but how to determine how much less?

A given application has a transient specification that must be met. This specification is used as a guideline to determine the lowest transient inductance (L_{tr}) necessary to meet the specification. The less this inductance is, the faster the signal can change, assuming adequate control bandwidth. This has another benefit of reducing the quantity of output bulk capacitance, which reduces cost and board space.

So you need lower inductance to meet transient requirements but high enough steady-state inductance (L_{ss}) to keep the ripple in check and have enough output

current capability to meet the requirements. Figure shows the waveforms and equations that determine L_{ss} and L_{tr} using the schematic of Figure 6.4 [119].

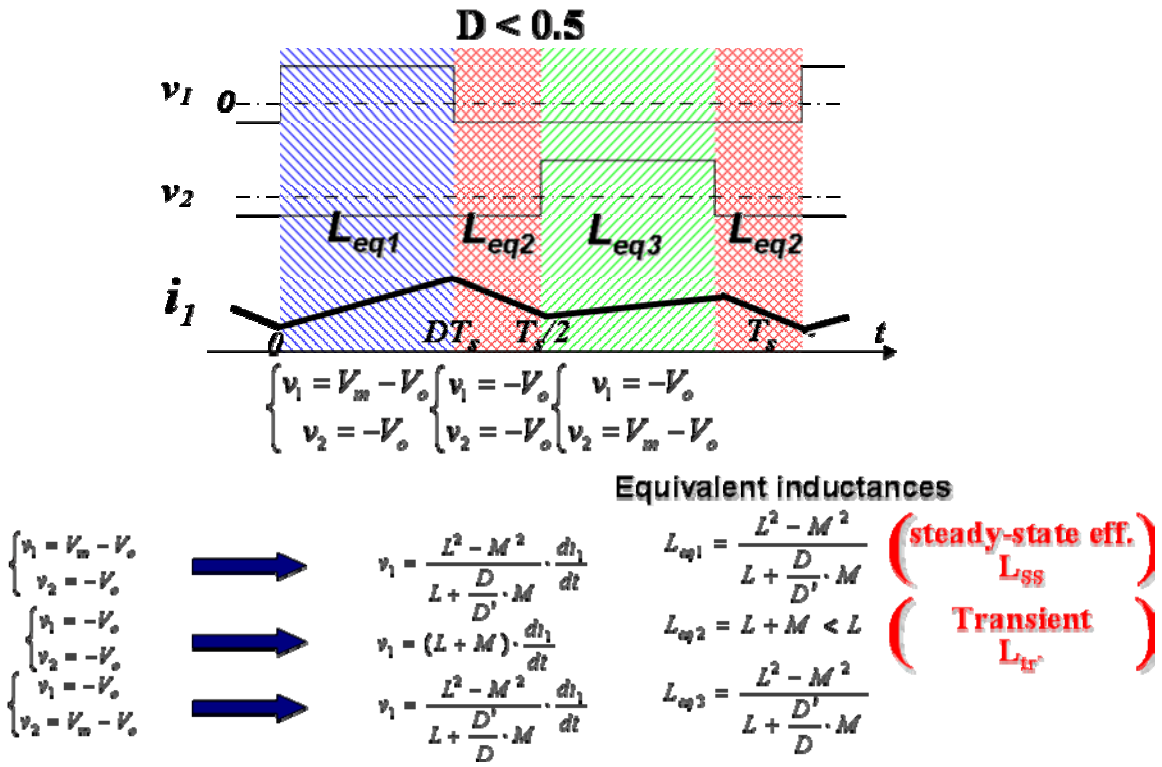

Figure 6.4. Waveforms and derivation of the inductance equations.

The advantage of using a coupled inductor rather than 2 single-phase non-coupled inductors is based on ripple reduction and transient improvement. For 2-phase non-coupled inductor (L_{nc}) buck and 2-phase coupled inductor buck, if $L_{tr}=L_{nc}$, then the phase current ripple reduction of coupled inductor buck compared to the non-coupled inductor buck is determined by the L_{tr}/L_{ss}. If $L_{ss}=L_{nc}$, then the transient response improvement is also determined by the L_{tr}/L_{ss}. [88]. Figure 6.5 shows the graph of phase current ripple reduction vs. the coupling effect and the duty cycle for the 2-phase

coupled inductor buck [119]. The stronger the coupling is and the nearer to 0.5 the D is, the better the performance.

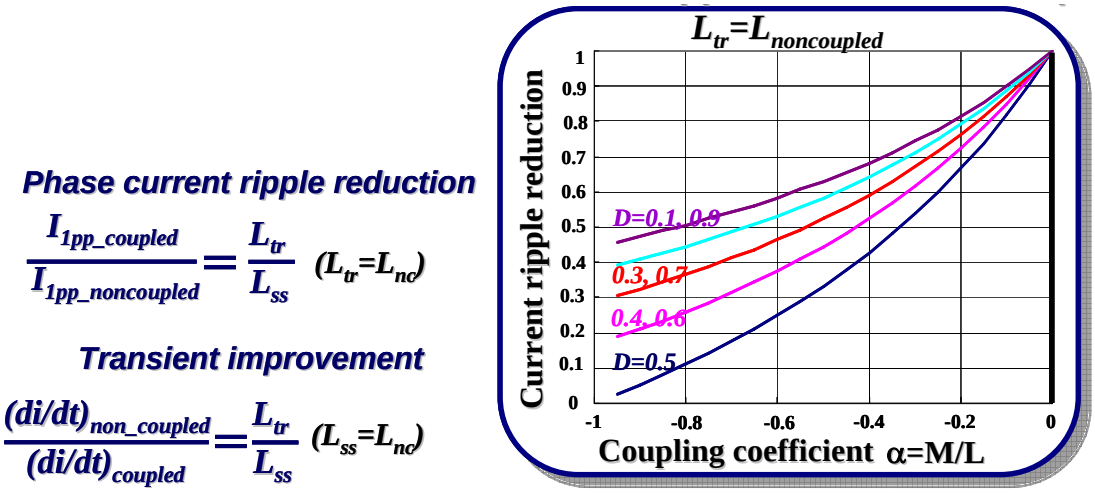

Figure 6.5. Advantages of coupled inductors over equivalent non-coupled inductors.

So the trade off in designing a coupled-inductor using LTCC material is that since there is no air gap due to the pressing and sintering processes, there isn't the flexibility to increase the self-inductance of coupled inductors *by reducing the air-gap*, the way it is done with a ferrite core. This yields weaker coupling. The relationship between L_{ss} and L_{self} is shown by Equation 40.

$$L_{ss} = \frac{L_{self}^{\,2} - M^2}{L_{self} + \frac{D}{D'} \cdot M} = \frac{1-\alpha^2}{1+\frac{D}{D'} \cdot \alpha} L_{self} \tag{40}$$

Stronger coupling leads to greater levels of flux cancellation, which in turn reduces flux density. This reduction, due to the nonlinearity of the B-H curve of the LTCC material will, will lead to stronger coupling. As the coupling approaches -0.9, the

steady-state inductance will drop to 25% of the self inductance. But with weak coupling, say the -0.6 we are targeting for this coupled LTCC inductor, then steady-state inductance will be 80% of self inductance (keeping L_{self} constant).

6.3 Active and Passive Layers for 2-Phase Design

Using LTCC material to make a coupled-inductor is trickier than a ferrite core because of the lack of air gap. The missing air gap makes designing for a very specific coupling coefficient more difficult. However this also presents a couple advantages as was outlined in Chapter 3 for the non-coupled LTCC inductor design: non-linear inductance to boost light load efficiency with high inductance, and achieve fast transient speed with low inductance at heavy load. Add this to what was just discussed and we have a great opportunity to reduce the size of the inductor and still maintain high performance. This work was completed with the help of Yan Dong from CPES [118, 119].

For this 2-phase structure, we looked at the smallest size we could make the active layer. It consists of 2 drivers, 4 switches, and necessary capacitors. In order to reduce the footprint, the drivers are stacked on top of the embedded devices. The arrangement of the components was chosen such that the converter would be a long this rectangular shape. As was shown in Chapter 5, this shape offers higher cooling capability than a square one, so long as the airflow is down the length. Figure 6.6

shows the active layer AutoCAD layout that offers an impressively small footprint of 9 x 23mm.

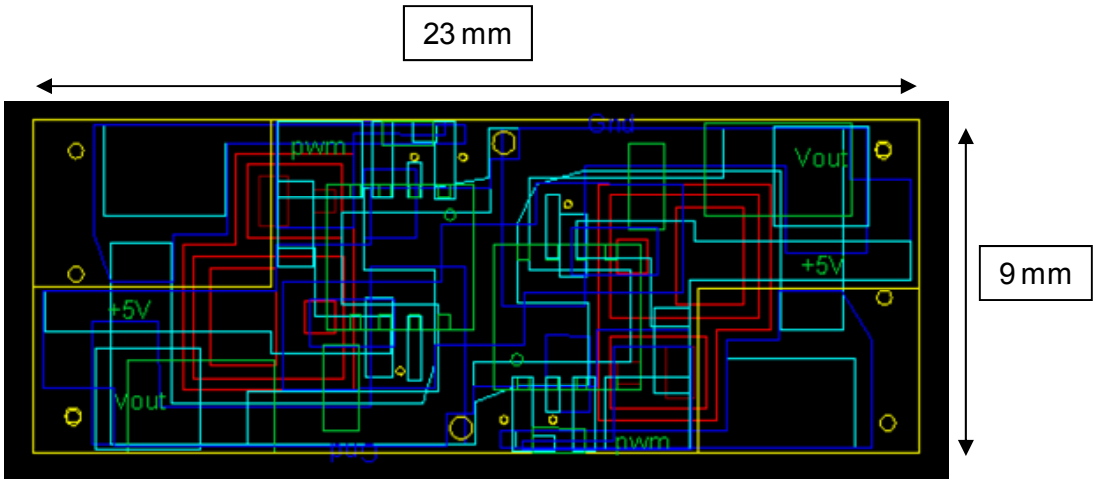

Figure 6.6. The two active layers on top of each other in wireframe

The top layer of the two is shown in Figure 6.9. It features an “S” shape because we must keep the input decoupling capacitor as close to the switches as possible. With the two opposite corners cut out, we can place the capacitor on the device layer, right next to the devices. This keeps the loop as short as possible to reduce parasitic inductance and improve efficiency as outlined in Chapter 3. Also, a closed ground path is located below the drivers and above the switches to short out the magnetic fields generated. This is shown by the loop superimposed on Figure 6.7.

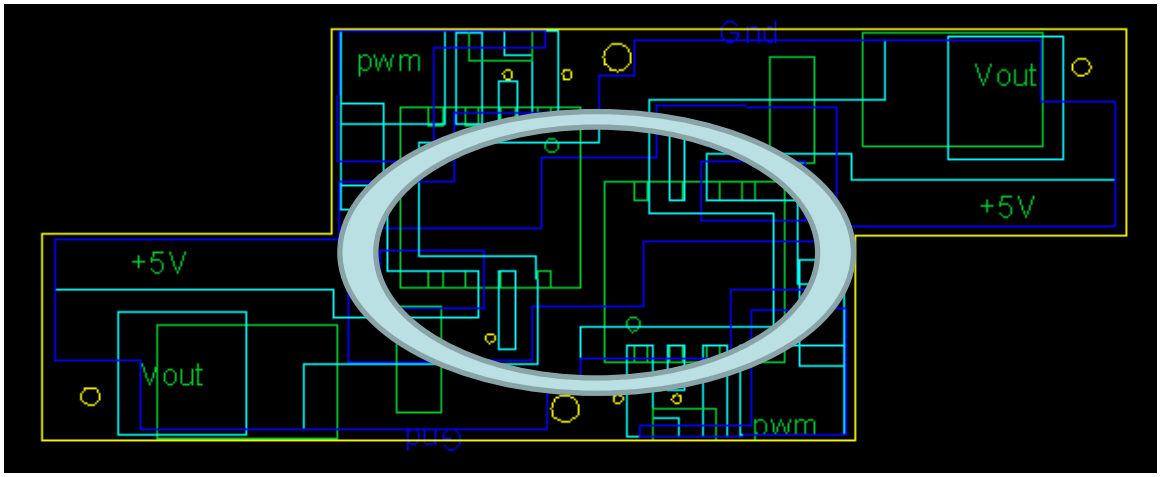

Figure 6.7. The top (driver) layer of the active stage showing ground loop

Another important step in the active layer design, as discussed in Chapter 3, is the connection of the driver and the switch node between the devices. This path must be kept short so that no opposing voltage reduces the top switch gate drive. This path is only 4 mm long in this design, as shown by the white arrows in Figure 6.9. The final active layer hardware is shown in Figure 6.10.

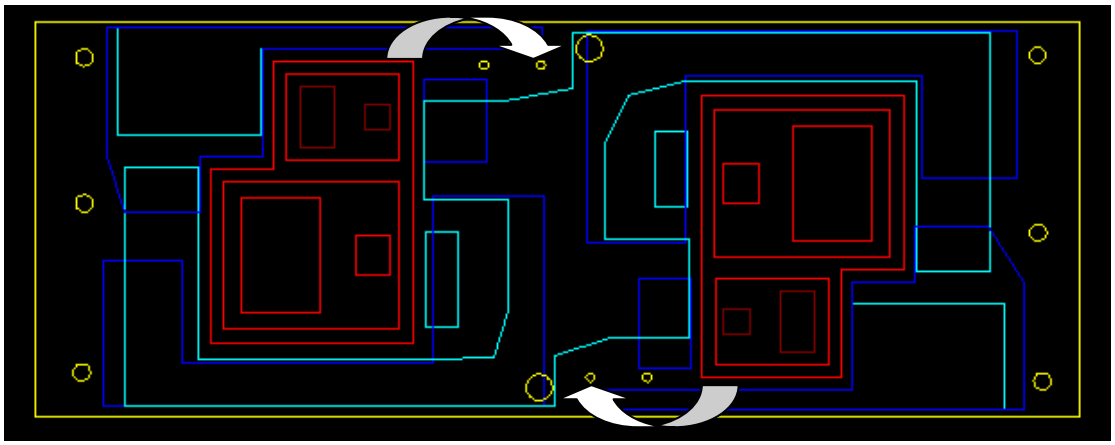

Figure 6.8. The bottom (device) layer showing short Vsw-driver paths

Figure 6.9. Active layer hardware shown next to a quarter coin for scale.

Based on the small 9 x 23mm footprint we were able to achieve with the silicon components, we embarked on a mission to get the highest possible steady-state inductance and smallest footprint for the inductor. We used the same methodology as outlined in section 4.2 for the single-phase module and modified it according to 6.2 to accommodate for the coupled-inductor differences [118, 119].

One of the modifications involves a process change due to the extra thickness of the LTCC material. The larger quantity of material led to cracking issues because of improper off-gassing of the silver paste during curing. The gas wanted to escape from between the LTCC layers and so caused cracks in the sides of the inductor. In order to improve the gas transfer without it having to crack the LTCC, small perforations were made in the top and bottom sides to reduce pressure build-up and allow for a more reliable fabrication method. In addition, the width of the LTCC material on either side of the traces was increased for better mechanical strength and to make the perforations more effective. Figure 6.10 shows three samples that were made, each one with width w1 wider than the previous, and with additional holes in to the top and bottom sides.

The value of w1 was varied from 1 mm, 1.5 mm, and finally 2 mm and the hole count went from 3, to 7, to 9 respectively. With w1 equal to 2 mm and with 9 holes in the top and bottom layers, there was no more cracking on the sides.

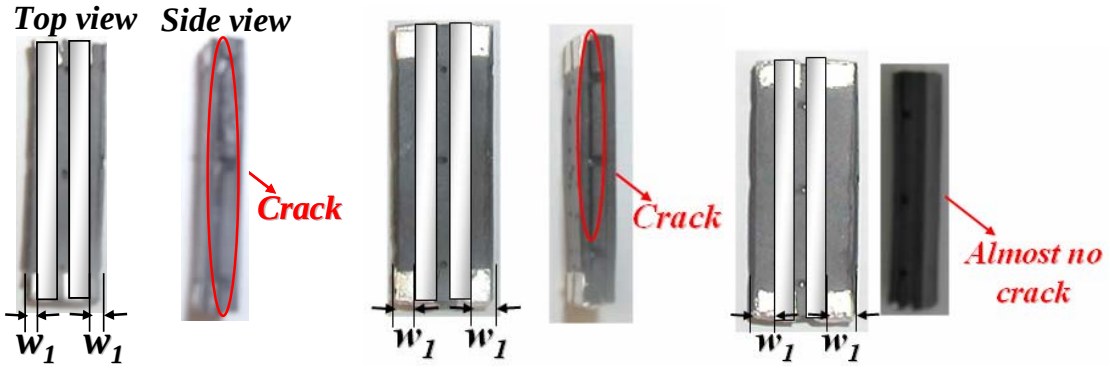

Figure 6.10. Three design iterations to eliminate cracking of LTCC [119].

Two coupled designs were made and are compared with the Generation 2 non-coupled inductor in Figure 6.11. The final structure 2 design is a long one for thermal reasons (discussed in Chapter 5) as well as the fact that we were able to get the same inductance and similar DCR as structure 1 but in a much smaller footprint of 8 x 20 mm. We achieved a 51% size reduction with the final coupled design when compared to the size of two non-coupled inductors together. Also, structure 2 has higher inductance than structure 1, as shown in Figure 6.12 [118]. This allows us to make a very small 2-phase module.

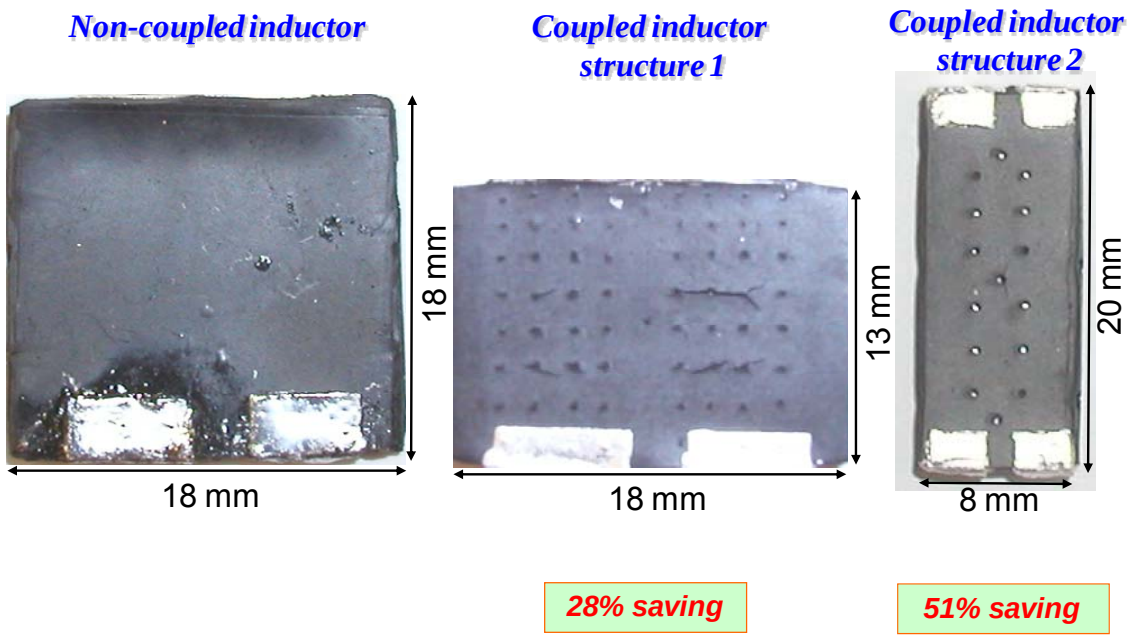

Figure 6.11. Inductor comparisons of two coupled designs and one non-coupled for reference.

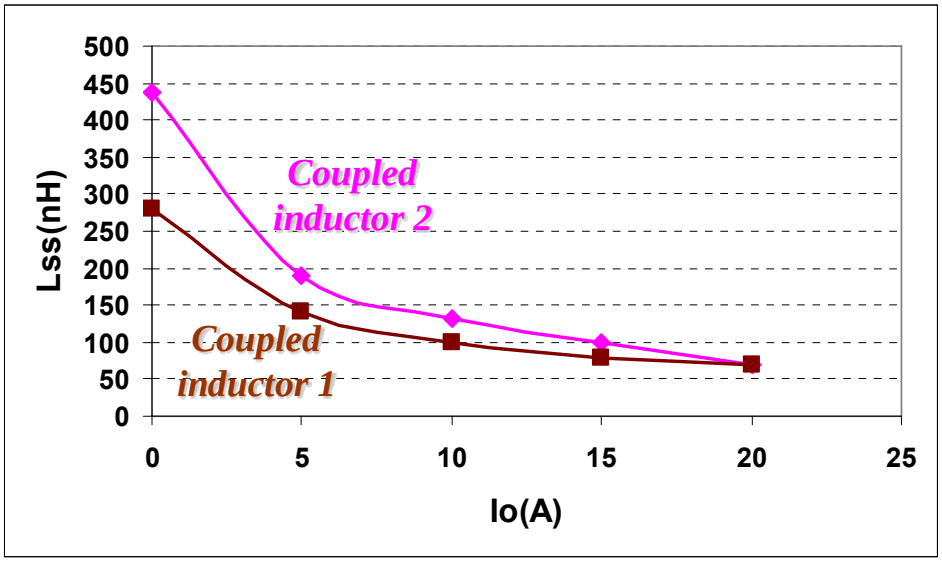

Figure 6.12. Inductance comparison between the 2 coupled structures [118].

The waveforms of chosen Structure 2 are shown in Figure 6.13. This was run on a PCB test board at 20A output with both phases operating. The coupling effect is seen

in the waveform where every other peak is reduced in amplitude. The peak-to-peak ripple is 10A and the steady-state inductance calculated from the graph is 70nH at 20A load, which is right where we wanted it.

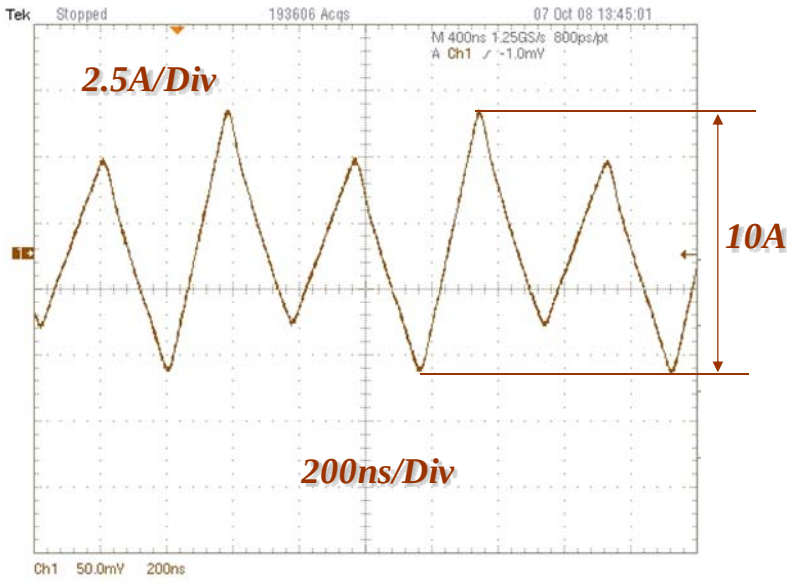

Figure 6.13. Coupled inductor ripple current waveform at 20A.

6.4 Stacked Power with Coupled Inductor Results

The Stacked Power coupled inductor was designed to have the same heavy-load inductance as the non-coupled case, and similar, if not lower, DCR, so it would be natural to achieve similar heavy-load efficiency. Testing the module at 5V input confirmed this is the case. Figure 6.14 shows a comparison of the 2-phase coupled-inductor POL and running only 1 single phase with the Generation 2 inductor but doubling its output current to make it apples to apples.

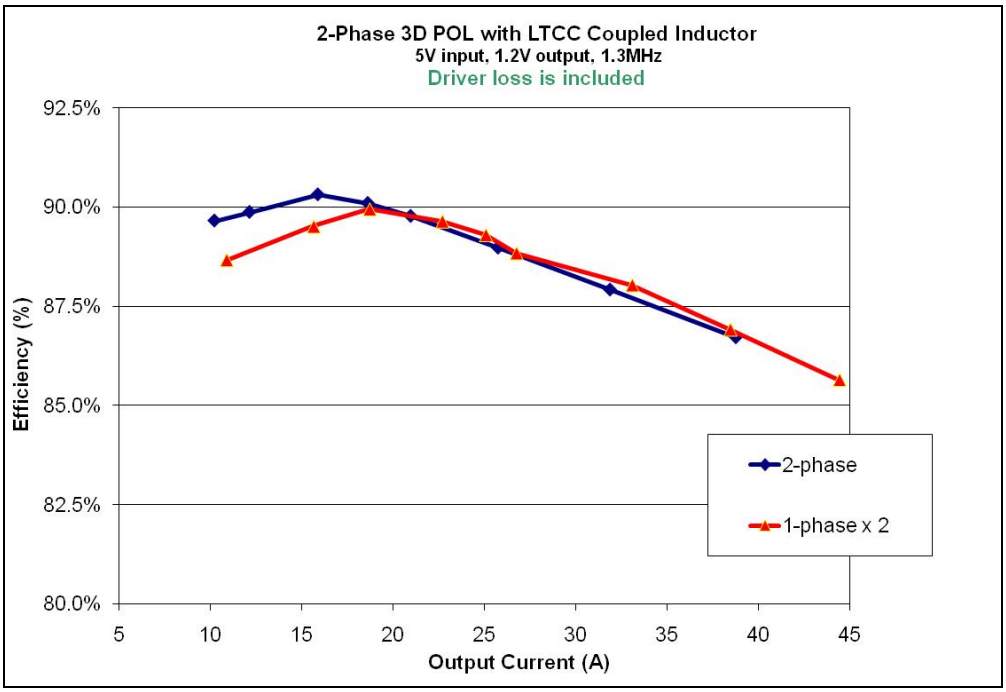

Figure 6.14. Efficiency comparison between 2-phase coupled and 1-phase non-coupled multiplied by 2.

It is clear in this efficiency graph that the coupled-inductor 2-phase converter works very well. Efficiency is on par with the single-phase inductor and is actually even better at light load, demonstrating a gain of 1% efficiency improvement thanks to the higher inductance of the coupled-inductor under this condition. However, the coupled module is 50% smaller than the Generation 2 converter and is nearly double the power density at 500 W/in^3! This is a huge improvement. The complete module with inductor substrate is shown in Figure 6.15 and has dimensions of 9 x 23 x 7 mm. It is shown next to an American 25¢ quarter coin for scale.

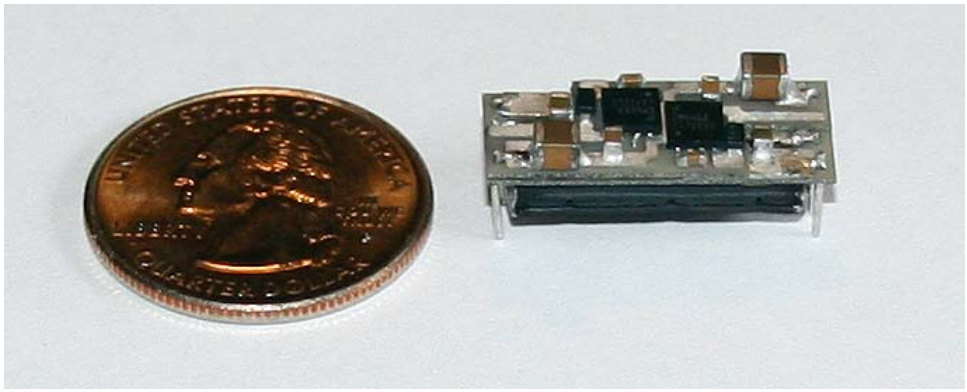

Figure 6.15. Complete Stacked Power 2-phase module next to a quarter

Thermally, the new 2-phase module is of course not of the level of the much-larger Generation 3 module, thereby requiring higher operating temperatures for similar output current. The Generation 3 can run at 20A in natural convection (no fan and no added heat sink), whereas the 2-phase module was at 90°C, tested with a calibrated thermocouple between the layers, at a level of about 13A in 25°C ambient and natural convection (both have the same power loss and 90% efficiency at this level). But by allowing the module to reach 130°C peak temperature, it achieved 25A output.

This was to be expected since at this level, both modules require similar heat transfer coefficient under these circumstances, as indicated by Equation (11). Table shows the values for each, and assumes that conducted heat is the same for both cases, which explains the slight differences in heat transfer coefficients.

Table 6.1.
Calculation values for natural convection capability for 2 Stacked Power POL modules

POL	Efficiency	Current	Loss	Surface area	Delta T	HTC
1-phase	87.5%	20A	3.4W	324mm^2	65°C	161W/m^{2}K
2-phase	89.0 %	25 A	3.7 W	207 mm^2	105°C	170W/m^{2}K

A thermal picture of the 2-phase module was not possible due to the number of wires used to interconnect it covering the board. However, to really be able to assess this performance objectively, Figure 6.16 shows a commercial example 40A non-isolated POL: The Power One ZY2140. It operates at 500 kHz and is 1.8 x 0.55 x 1.1 inches in size. The high profile of 1.1" is due to the fact a heat sink is mandatory for full current operation. Our module with no heat sink, 25°C ambient and with an airflow of 200 LFM (1 m/s) is able to reach its output of 40A with no heat sink in order to maintain a profile of only 0.3". This airflow is about 6 times greater than natural convection but it allows the module to dissipate its 7.5W of heat (86.5% at 40A, including driver loss) and still keep the delta T at 70°C. This is very impressive performance for such a small module.

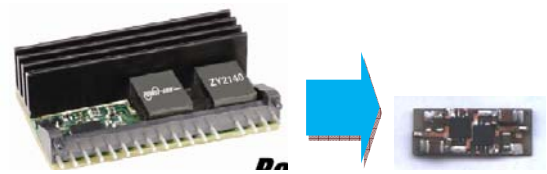

Figure 6.16. Power One ZY2140 with CPES module, both same relative scale.
Used with permission from Power One

It must be noted, however, that the ZY2140 has many more features than the Stacked Power module, thanks to its sophisticated control architecture. Here is a list of these features:

- Wide input voltage range: 8V–14V
- High continuous output current: 40A
- Wide programmable output voltage range: 0.5V– 3.65V
- Active digital current share
- Output voltage margining
- Overcurrent and overtemperature protections
- Overvoltage and undervoltage protections, and Power Good signal tracking the output voltage setpoint
- Tracking during turn-on and turn-off with guaranteed slew rates
- Sequenced and cascaded modes of operation
- Single-wire line for frequency synchronization between multiple POLs
- Programmable feedback loop compensation
- Differential output voltage sense
- Enable control
- Flexible fault management and propagation
- Start-up into the load pre-biased up to 100%
- Current sink capability

Nevertheless, the issue with the heat sink as used is that its fins are close together and thus airflow has trouble being forced between them unless a bulky shroud is used to force the air between the fins. However this force is directly translated into pumping losses which reduces the airflow rate accordingly – or requires to be offset by a larger and noisier fan. So without knowing the airflow rate between the fins (rather than the ambient free stream air flow rate), it is hard to use this derating curve for any more detailed analysis than was just presented. These are all issues to keep in mind when conducting thermal comparisons. The derating curve for the CPES 2-phase POL is shown in Figure 6.17. They are similar to the Power One module but keep in mind the fact that its surface area is much smaller since it doesn't have a finned heat sink – but the surface area it does have is connected to the dies with extremely low thermal

resistances (as outlined in Section 4.4). In addition, the Power One module's heat sink effectiveness is partially diminished by the thermal pad interface (again, as described in Section 4.4).

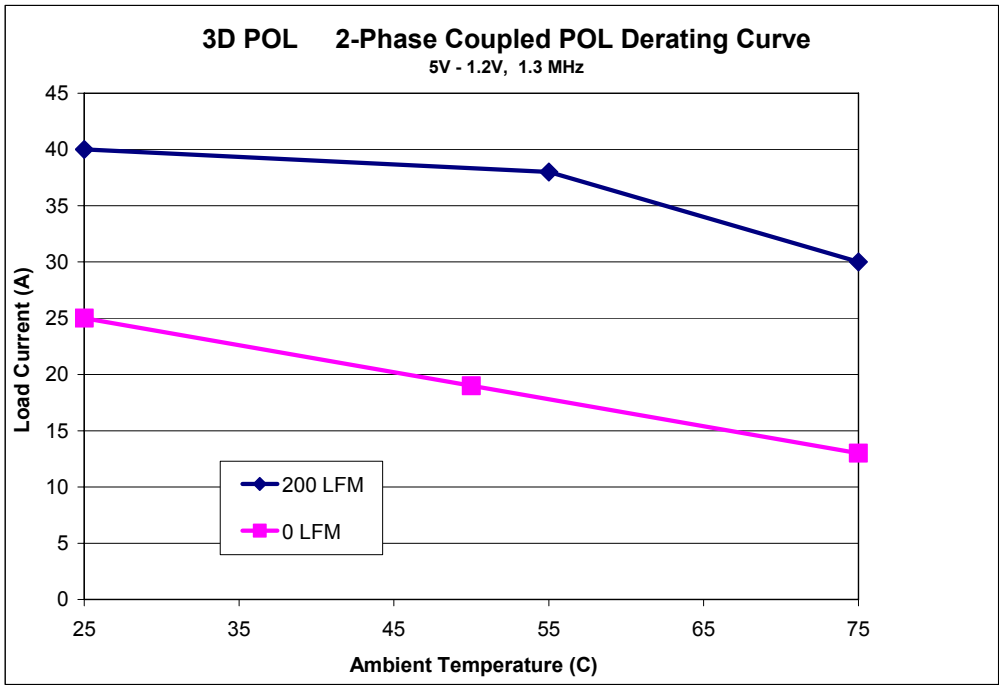

Figure 6.17. Derating curve for the CPES 2-phase module

Another example of 40A operation conducive to VRM applications is that of the iPowir IP2005A from International rectifier. However, this module does not include the inductor on-board but it does offer excellent performance and is very small at 7.7mm x 7.7mm x 1.7mm LGA. It is capable of 40A output at a switching frequency of 1MHz. Cooling is via thermal pads into the motherboard and it is claimed that no derating is needed up to a PCB temperature of 95°C. Further reading in the application notes shows that this condition is for a very large PCB area as shown in Figure 6.18.

The IRDCiP2003A-C is a 160A demo board using 4 40A modules. As you can see, each inductor is larger than each module, footprint usage (including the capacitors

and enough large traces to conduct heat) is very high, on the order of 20mm by 80mm just for the modules, caps, and inductors (so as to not unfairly include the test-and-mounting portions of the demo board). Such large surface area does indeed allow the thermal pad modules to cool sufficiently for such high power output but the overall size is still larger than CPES-designed 1MHz VRM discussed in Chapter 2. Thus, the only real advantage is the reduction of parts count to make a working multiphase board. Once again, we see that the thermal performance of pad modules with large plastic casings limits the final solution size reduction compared to what can be achieved with an open-frame, thermally-enhanced substrate design such as Stacked Power. On balance, the Stacked Power module would most likely cost more than the IR part due to the increased material costs.

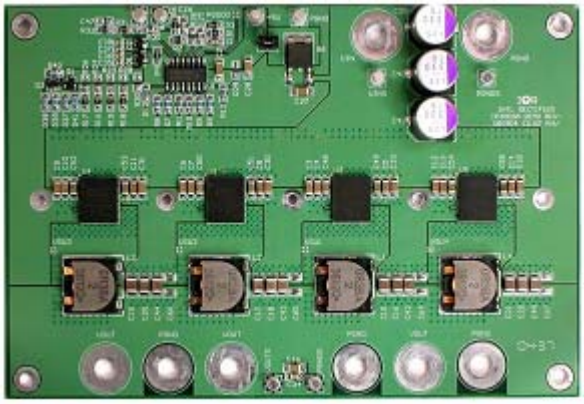

Figure 6.18. IPOWIR modules as building blocks for VRM type applications. The modules are visible as a row of four black squares.
Used with permission from International Rectifier

But for the joint conduction (through connecting wires) and convection (off the surfaces) cooling that our 2-phase module was subjected to during our bench testing, 200 LFM cooling before derating for a 40A, 500 W/in^3 (with inductor and enough

capacitors to keep ripple below 60%) module is very impressive nevertheless, even if we assume there may be less conduction occurring through the pins to the motherboard than through the large test leads we used. But having the device hotspot temperatures spread out over an area that is about 3.5x larger on two layers (the DBC surface area relative to die area), with no crowded heat sink fins to require additional fan force, allows us to achieve this unprecedented performance in the tiniest of packages.

Chapter 7. Conclusions and Future Work

It was shown in this dissertation that integration technology in the 3W output power range has been flourishing of late, with many attempts made at making an integrated POL on silicon. However, very little attention has been paid to integration in a higher power range such as 25W-50W, which cannot be implemented in CMOS today due to the large board area passives this power level would occupy. There is no doubt that a push to higher power integration will be forced upon everyone in the near future so this power range should not remain overlooked.

The main limiting factor in higher-power POL integration is thermal. Dealing with the heat that is developed at this power level, and at typical efficiencies, requires the use of big, bulky and expensive heat sinks. These heat sinks are detrimental to the package size, profile and power density, and obviously render any integration efforts moot. There are two main courses of action when dealing with this heat: First, reduce it with extremely high efficiency and, second, conduct heat away from the devices quickly and have a large surface area for convection.

Both of these steps have been implemented in the examples of the 3D POL designs using the novel “Stacked Power” process. This POL was not made to suggest that this is the only way to make it, but rather to outline the benefits, barriers, and limitations of integration in the 25-50W power level range and offer an example of how to successfully push these barriers farther away.

The AlN DBC layers extract heat from the devices and distribute it uniformly over the entire converter surface for higher convective efficiency, both due to the material's high thermal conductivity. This allows for a smaller size, which reduces the trace parasitic inductance and resistances since the interconnections are shorter. The electrical layout has been designed to take advantage of these parasitic reductions, allowing for high efficiency so that less waste heat is generated, mitigating the thermal issues at the source. These are the main advantages of Stacked Power at getting around the thermal limitation and allowing unprecedented size reduction.

When comparing other commercial and research products to the Stacked Power 3D POL Converter, it becomes clear that the added benefits of layering the circuit in the vertical Z-axis by using a substrate that maximizes the hotspot-to-package ratio, a thin inductor layer to reduce profile, high frequency to minimize the passives' size and shorten signal paths, and low die-to-board thermal resistances, pays big dividends. Switching frequency is currently at 1.3MHz and with the same 5V to 1.2V conversion, it outputs 24A in single-phase form and 40A in two-phase with coupled inductor! Efficiency of the single-phase power stage is a very high 88% thanks to careful electrical layout design, which pushes power density to 260 W/in^3, including all thermal management and all necessary parts onboard! These collective specifications better any industry product released so far.

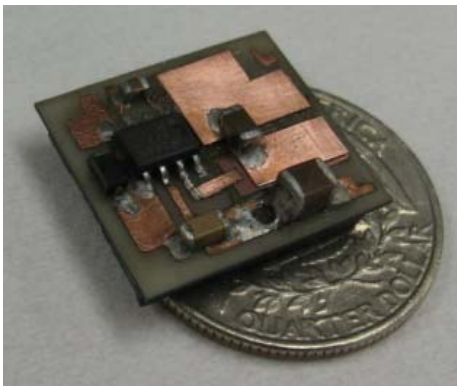

Figure 7.1. CPES Stacked Power 3D Integrated POL with inductor substrate.

So in the big picture summary, we have research on the monolithic buck, using thin-film inductor, at the low end of the power spectrum, followed by the Enpirion module which represents a great leap forward by using a more conventional inductor with its silicon, followed by the Linear Tech module which moves to the next power level by use of co-packaged active components, and at the top we have the new CPES Stacked Power module which combines co-packaging of active components with advanced thermal capability on top of a thin, layered inductor substrate as shown in Figure 7.2.

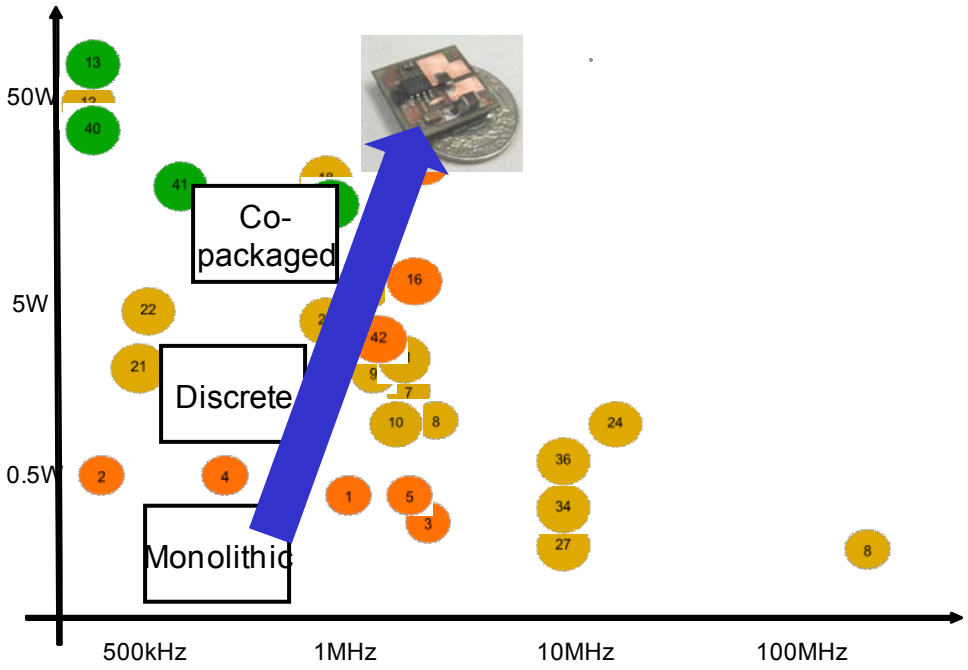

Figure 7.2. Power output improvements afforded by each of the 4 design concepts

In order to determine the impact that Stacked Power has overall on the packaging performance – both electrical and thermal, Figure 7.3 shows some comparisons with typical industry products of a similar type. It shows a graph of power density versus output current for a wide variety of low-voltage non-isolated DC/DC converters. All of them require output capacitors, except for the CPES Stacked Power module, but this is not figured into their power density figure – so they are shown higher than they should be. Also, these power densities are all calculated from datasheet specifications for 1.2V output. As we can see, the CPES Stacked Power module is the leader of the pack when both power density and output current are taken into account.

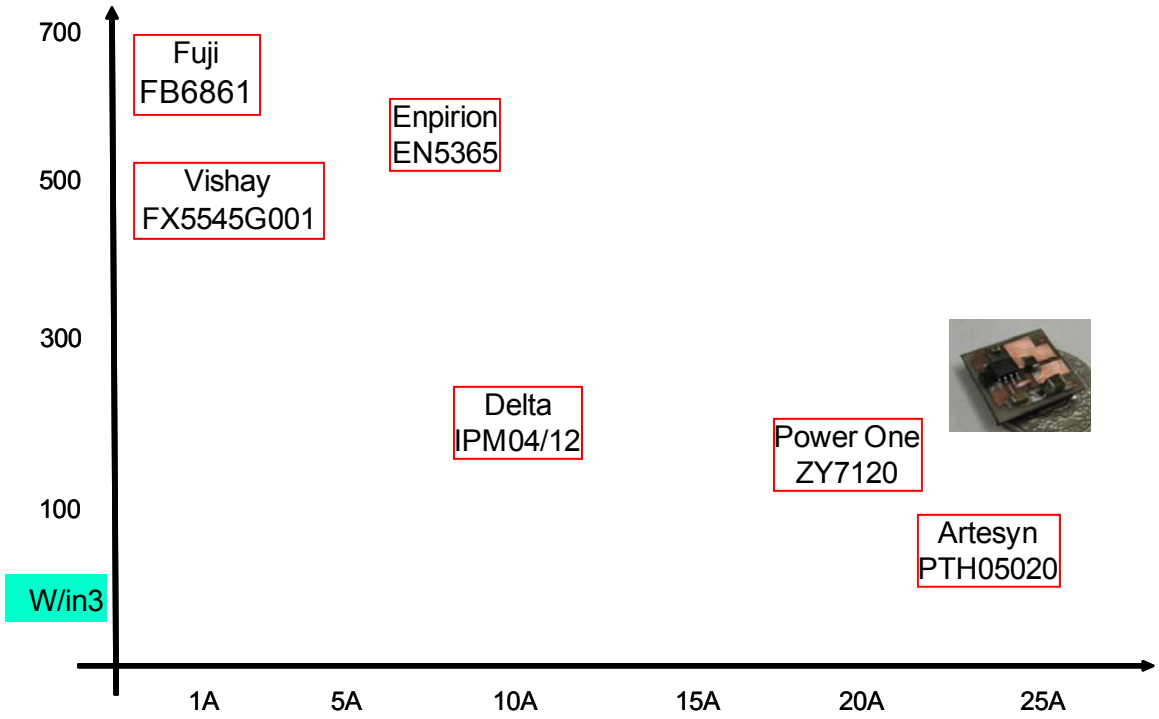

Figure 7.3. Power density versus output current for a range of low-voltage DC/DC converters

Figure 7.4 thermally compares a few low-voltage DC/DC industry products that are of similar size and power to Stacked Power POL. This graph shows the ambient temperature at which converter derating starts under natural convection conditions and at what load current that happens to be. So this is the maximum output current possible for the given ambient temperature. As we can see, the Stacked Power module is ahead of the rest thanks to it superlative thermal design.

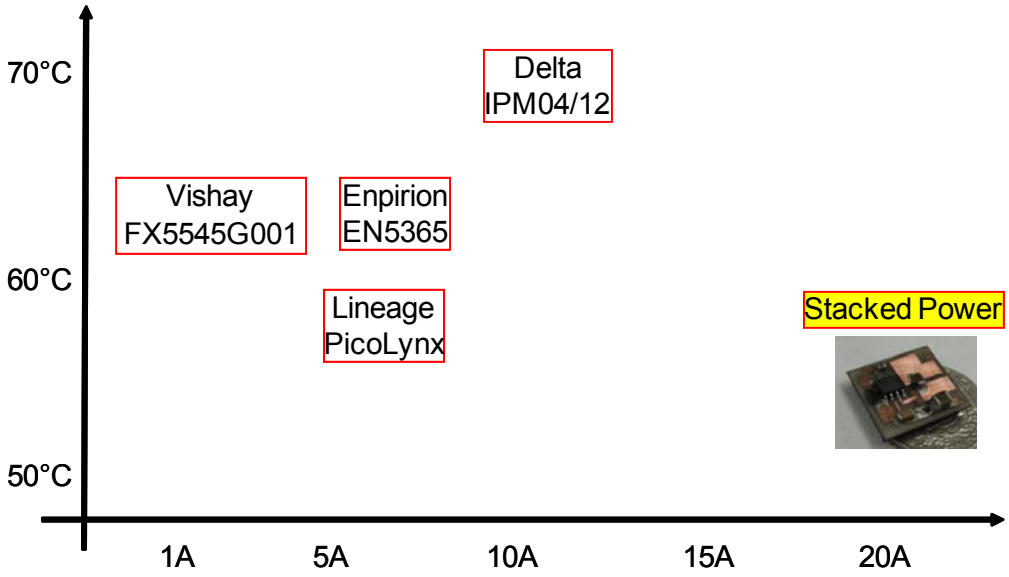

Figure 7.4. Maximum output current at the ambient temperature where derating begins for similarly-sized modules.

The single-phase efficiency graphs for the Stacked Power 3D POL converter, including driver loss, are repeated in Figure 7.5. As we can see, we get 12V-1.2V efficiency of 87% at 14A and 85% at 23A and for the low voltage input case, 5V-1.2V efficiency of 91% at 10A and 88% at 24A. These efficiencies are higher than what is currently found in industry, thanks to a thermally-enhanced design with ultra-low parasitics and thoughtful layout.

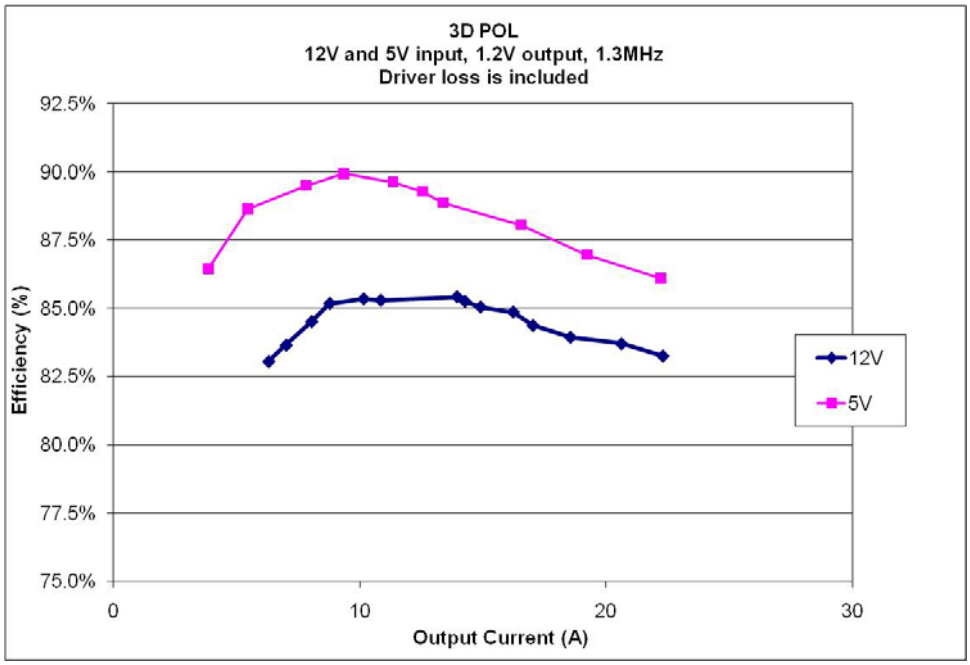

Figure 7.5. Efficiency curves for single-phase 12V and 5V input, including driver loss.

The 2-phase Stacked Power module with coupled inductor represents a huge leap forward in minituarization of high-current POL converters. Using the thermal benefits of the AlN DBC substrate, the module could be made tiny without requirements for huge airflow rates to reach its 40A output. Chapter 6 showed a CPES module next to a commercial 40A module, both pictures being scaled relatively well to each other. The tiny size of the CPES 3D POL is hard to believe and yet we have performance results to prove its merits.

The 2-phase efficiency is on par with the single-phase inductor and is actually even better at light load, demonstrating a gain of 1% efficiency improvement thanks to the higher inductance of the coupled-inductor under this condition. However, the coupled module is 50% smaller than the Generation 2 converter and is nearly double the power density at 500 W/in^3! This is a huge improvement. The 2-phase efficiency curve for 5V input is shown in Figure 7.6.

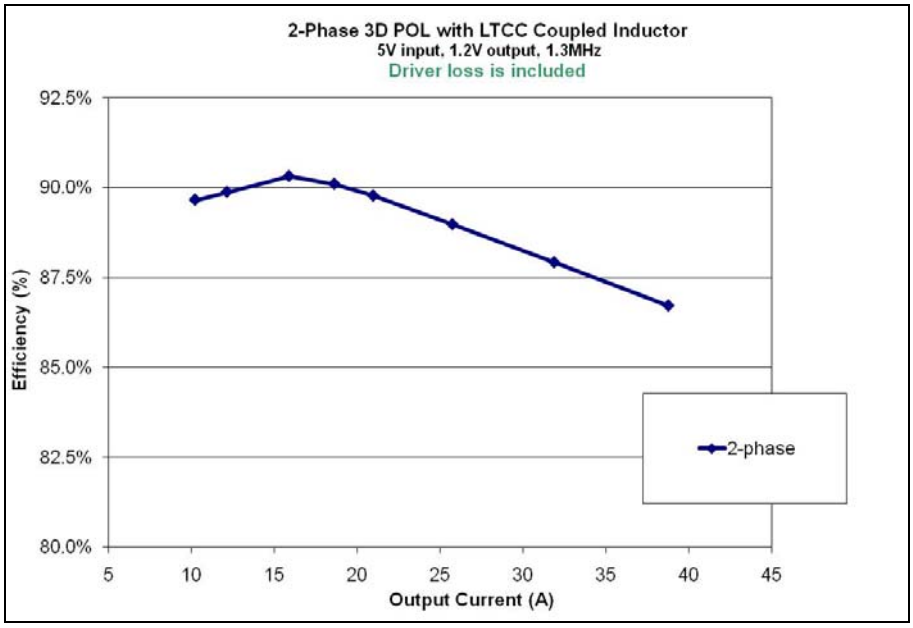

Figure 7.6. Efficiency of 2-phase Stacked Power with LTCC Coupled-Inductor

The work done thus far is already giving clear insight into the need for an industry-wide paradigm shift away from traditional PCB material if high performance in a small size is desired. PCB worked perfectly well when power densities were low but we are now rapidly approaching its thermal limit. To remove this barrier, different substrate materials are needed that are multi-functional rather than simply there to support parts and traces. An example of how to make a single-phase Stacked Power POL converter with a very-high power density of 260 W/in^3 (at 1.2V output and including thermal management for 20A with no heat sink and all capacitors onboard) and a 40A 2-phase module with a single coupled-inductor for both interleaved phases has been shown in this dissertation.

The summary of key points for Stacked Power are as follows:

1. Integrated thermal management ***replaces***:
   - ❖ PCB with high-performance ceramic substrate.
   - ❖ Bulky and expensive external heat sink.
   - ❖ Cooling fan up to 55°C ambient and 20A!
2. Active devices and inductor layer are integrated together
   - ❖ Integration allows for small size, low profile and short signal paths
3. Parasitics are reduced for less loss at high frequencies
   - ❖ High frequency reduces size of inductor and output capacitance
4. At 1.2 Vout,
   - ❖ For single-phase, power density of 260 W/in^3 achieved
   - ❖ For 2-phase, power density of 500W W/in^3 achieved
   - ❖ This includes ***all*** necessary thermal management!
5. High efficiency, with driver loss and an ultra-low profile of < 7 mm.
   - ❖ 1-phase 12V-1.2V efficiency of 86% at 10A and 83% at 22A
   - ❖ 1-phase 5V-1.2V efficiency of 90% at 10A and 86% at 22A
   - ❖ 2-phase 5V-1.2V efficiency of 90% at 10A and 86% at 40A

The difficult task to coming up with a module capable of this level of performance in the 25-40A range has been accomplished. The limitations of today's conventional

approach using PCB has been clearly laid bare. Fabrication techniques must change to obtain a quantum leap in performance at small sizes. The status quo of PCB use will have to come to an end in the very near future for high performance modules. This work lays the groundwork for accomplishing this seemingly-extraordinary challenge. However, mass manufacturing of Stacked Power is not that far from reality since ceramic-based hybrid modules are already being made in aerospace applications (such as products from VPT and Crane Aerospace). The work shown in this dissertation leads the way to making even better use of the ceramic in applications where such a notion hasn't even been considered fully yet. This is the most important part of this work.

Future work will include using this technique at lower current levels and thus, for even smaller modules. Novel inductor work being done by Mr. Qiang Li in CPES will be used in conjunction with Stacked Power to make tiny modules of relatively very-high power using a lateral flux inductor.

The potential of using Stacked Power in VRM applications has been outlined here but actual hardware testing for this structure is also future work. The drive to reduce size and increase power is no more demonstrated than in this challenging application. Stacked Power can step up to the plate and make this a reality.

Multi-layered modules are a very real possibility for future work. As has been shown, active device layers can be multiple and stacked together to make a multiphase module rather easily and of very low profile. Extending integration in the vertical axis is easy by this method and allows for the much-desired small footprint. No other integration method makes such good use of the vertical axis today. However, to assure

good reliability, an intensive thermo-mechanical analysis needs to be implemented to determine cycling strength, voiding acceptability, and fatigue stresses on AlN DBC inner layers.

References

[1] M. Mino, K. Tsukamoto, K.Yanagisawa,A. Tago, and T.Yachi, “A compact buck-converter using a thin-film inductor,” in Proc. IEEE Applied Power Electronics Conf. Expo. (APEC), vol. 1, Mar. 1996, pp. 422–426. (10x6x3mm, CoZrTa thinfilm L on Si, 1.6MHz, 85%, 3.3Vo, 0.12Ao)

[2] S. Sakiyama, J. Kajiwara, M. Kinoshita, K. Satomi, K. Ohtani, and A. Matsuzawa, “An on-chip high-efficiency and low-noise DC/DC converter using divided switches with current control technique,” in IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers, Feb. 1999, pp.156–157. (2Vo, 80mAo, 92%, 0.25umCMOS,10uH off-chip)

[3] Y. Katayama, S. Sugahara, H. Nakazawa, and M. Edo, “High-power density MHz-switching monolithic DC-DC converter with thin-film inductor,” in Proc. IEEE Power Electronics Specialists Conf., vol. 3, June 2000, pp. 1485–1490. (1Wo, 3Vo, 3MHz, 83%, thin film CoHfTaPd L with planar coil, 10x10x1.8mm

[4] J. Kim and M. A. Horowitz, “An efficient digital sliding controller for adaptive power-supply regulation,” IEEE J. Solid-State Circuits, vol. 37, no. 5, pp. 639–647, May 2002. (0.25umCMOS, 280mW, 2.3Vo)

[5] K. H. Kim, J. Kim, H. J. Kim, S. K. Han, and H. J. Kim, “A megahertz switching DC/DC converter using FeBN thin film inductor,” IEEE Trans. Magn., vol. 38, no. 5, pp. 3162–3164, Sep. 2002. (micromachined L, 2.7Vo, 1.8MHz, 1.5Wo, 300mAo)

[6] P. Hazucha, G. Shrom, J. Hahn, B. Bloechel, et al., “A 233-MHz 80%–87% Efficient Four-Phase DC–DC Converter Utilizing Air-Core Inductors on Package,” IEEE J. Solid-

State Circuits, vol. 40, no. 4, pp. 838–845, April 2005. (90nmCMOS, 0.9Vo, 0.3Ao, 87%pk)

[7] Product: TI – TPS62300 - 2.7W @ 3MHz, 2mm2, 500mAmax

[8] Product: Maxim – MAX8640Z – 1.25W @ 4MHz, 13.4mm2, 500mAmax

[9] Product: Vishay – FX5545G201 – 3W @ 2MHz, 0.3mm2, 1Amax

[10] Product: National Semi – LM3677 – 2.16W @ 3.5MHz, 5.9mm2,m 600mA

[11] Product: ON Semi – 1.96W @ 3MHz, 4mm2, 600mA

[12] Product: Artesyn SIL40C – 40A, 1V, 300k

[13] Product: Astec POLA PTH series – 100W, 350k,30Amax

[14] Mikura, T.; Nakahara, K.; Ikeda, K.; Furukuwa, K.; Onitsuka, K., "New substrate for micro DC-DC converter," *Electronic Components and Technology Conference, 2006. Proceedings. 56th* , vol., no., pp. 5 pp.-, 30 May-2 June 2006

[15] Product: Vishay FunctionPAK FX5545G001,15W, 600mAmax, 2MHz, 15x12mm

[16] Product: Enpirion EN5360 – 6A, 1.2V at 5MHz

[17] Product: Linear Tech LTM4601 – 10A, 1.5V at 850k

[18] Product: Renesas Dr. MOS – 25A, 1.3V at 1MHz

[19] Ball, Arthur; Lim, Michele; Gilham, David; Lee, Fred C., "System design of a 3D integrated non-isolated Point Of Load converter," *Applied Power Electronics Conference and Exposition, 2008. APEC 2008. Twenty-Third Annual IEEE* , vol., no., pp.181-186, 24-28 Feb. 2008

[20] Product: Monolithic Power Systems, MP2106 – 1.5W, 800k, 1.5Amax

[21] Product: ISL8540, 2A, 1.3V, 600k

[22] Product: LT1370: A 500kHz, 6A Monolithic Boost, 2Vout

[23] Product: MPS MP2360, 2A, 3.3Vo; integrated top sw., driver, control

[24] James R. Warren, III, Kathryn Anne Rosowski, and David J. Perreault, “Transistor Selection and Design of a VHF DC-DC Power Converter,” IEEE Transactions on Power Electronics, Vol. 23, No. 1, January 2008, pp. 27-37. (2W resonant boost, 30MHz, 7Vo)

[25] Juan M. Rivas, David Jackson, Olivia Leitermann, Anthony D. Sagneri, Yehui Han, David J. Perreault, “Design Considerations for Very High Frequency dc-dc Converters,” Power Electronics Specialists Conference, June 2006, pp. 1-11. (11.5W resonant rectifier, 12Vo, 100MHz)

[26] Robert C. N. Pilawa-Podgurski, Anthony D. Sagneri, Juan M. Rivas, David I. Anderson, David J. Perreault, “Very High Frequency Resonant Boost Converters,” Power Electronics Specialists Conference, June 2007, pp. 2718-2724. (integrated one is 50MHz, 17W, 32Vo)

[27] Gary Hanington, P.F. Chen, V. Radisic, T. Itoh, P.M. Asbeck, “Microwave power amplifier efficiency improvement with a 10 MHz HBTDC-DC converter,” Microwave Symposium Digest, On page(s): 589-592 vol.2, Jun 1998 (300mW, GaAs boost, 9Vo)

[28] Wataru Saito, Tomokazu Domon, Ichiro Omura, “Demonstration of 13.56-MHz Class-E Amplifier Using a High-Voltage GaN Power-HEMT,” IEEE Electron Device Letters, Vol. 27, No. 5, May 2006 (1.9A, 13.4W)

[29] K. L. R. Mertens and M. S. J. Steyaert, “Fully Differential CMOS Class-E Power Amplifier,” IEEE Journal of Solid-State Circuits, Vol. 37, No. 2, Feb. 2002 (700MHz, 1W into 50ohms, 62%, 0.35umCMOS, 2.64mm2)

[30] V. G. Krizhanovski; V. A. Printsovskii, "800 MHz and 1035 MHz UHF Class E Power Amplifiers on CLY5 MESFET," Microwave and Telecommunication Technology, 2006. CriMiCO '06. 16th International Crimean Conference , vol.1, no., pp.175-176, Sept. 2006 (500mW, 0.063A)

[31] Frey, R., "500 W, class E 27.12 MHz amplifier using a single plastic MOSFET ," Microwave Symposium Digest, 1999 IEEE MTT-S International , vol.1, no., pp.359-362 vol.1, 1999

[32] El-Hamamsy, S.-A., "Design of high-efficiency RF Class-D power amplifier," Power Electronics, IEEE Transactions on , vol.9, no.3, pp.297-308, May 1994, 300W, 13.56MHz

[33] Power One ZY7120, 25W max at 1MHz based on eff. Curve

[34] Titus, J.L.; Gehlhausen, M.A.; Desko, J.C., Jr.; Nguyen, T.T.; Shibib, M.A.; Hollenbach, K.E.; Roberts, D.J., "Characterization of a fully resonant, 1-MHz, 25-Watt, DC/DC converter fabricated in a rad-hard BiCMOS/high-voltage process," Nuclear Science, IEEE Transactions on , vol.42, no.6, pp.2143-2148, Dec 1995, AT&T Bell Labs

[35] Hanington, G.; Pin-Fan Chen; Asbeck, P.M.; Larson, L.E., "High-efficiency power amplifier using dynamic power-supply voltage for CDMA applications," Microwave Theory and Techniques, IEEE Transactions on , vol.47, no.8, pp.1471-1476, Aug 1999 (500mW, 10MHz, 3.5Vo ?)

[36] Casey, L.F.; Schlecht, M.F., "A high-frequency, low volume, point-of-load power supply for distributed power systems," Power Electronics, IEEE Transactions on , vol.3, no.1, pp.72-82, Jan 1988, 42V, 3.6MHz, 50W

[37] Hanington, G.; Metzger, A.; Asbeck, P.; Finlay, H., "Integrated DC-DC converter using GaAs HBT technology," Electronics Letters , vol.35, no.24, pp.2110-2112, 25 Nov 1999 (1.5W, 10MHz, 3.5Vo)

[38] Tabisz, W.A.; Lee, F.C.Y., "Zero-voltage-switching multi-resonant technique-a novel approach to improve performance of high-frequency quasi-resonant converters," Power Electronics, IEEE Transactions on, vol.4, no.4, pp.450-458, Oct 1989, 5W, 13MHz, 3.3Vo

[39] Hannon J., O' Sullivan D., McCarthy K.G., Egan M.G., "Design and Optimization of a High Current High Frequency Monolithic Buck Converter", IEEE APEC 2008, 20MHz, 2A max, 600mA tested, 70% eff., 3.3V-1.2V

[40] Product: Delta Delphia IPM04S0A0S10FA, 40W max, 10A max, 400kHz, 18x16

[41] Product: Lineage Power PicoTLynx, 30W max, 6A max, 600kHz, 12.2x12.2mm

[42] Product: National Semiconductor LM3218, 650mA, 2MHz, 3x2.5x1.2mm, 2.2W

[43] Zhenxian Liang, Yu Meng, Ming Xu, J.D. van Wyk, F.C. Lee, "Integrated Packaging for a High Frequency Voltage Regulator Application," CPES Conference and Annual Meeting, Proceedings of the 2005 CPES Conference.

[44] M. H. F. Lim, J. D. van Wyk and Z. Liang, "LTCC Technology, an Integration Solution for Power Electronics", Ceramics Interconnect and Ceramics Microsystems Conf., 25-27April 2006, session WP3.

[45] M. H. Lim, Z. Liang, J. D. van Wyk, "An improved low profile integratable LTCC power inductor for microprocessor power delivery applications", 4th Int'l Conf. on Integrated Power Electronics Systems, 7-9 June 2006, Naples, Italy, pp. 299-304.

[46] M. H. F. Lim, Z. Liang, and J. D. van Wyk, “Low Profile Integratable Inductor Fabricated Based On LTCC Technology For Microprocessor Power Delivery Applications”, to be published in IEEE Trans. on Comp. and Packaging Tech.

[47] M. H. F. Lim, J. D. van Wyk and Z. Liang, “Impact of Conductor Cross-sectional Shape on Component Losses and System Performance”, Ceramics Interconnect and Ceramics Microsystems Conf., 25-27 April 2006, session WA3.

[48] M. H. F. Lim, J. D. van Wyk and Z. Liang, “Effect of geometry variation of LTCC distributed air-gap filter inductor on light load efficiency of DC-DC converters”, Ind. Appl. Soc. Conf., IAS, Tampa, Florida, 8-12 Oct 2006.

[49] S. Mathuna, T. O’Donnel, N. Wang, K. Rinne, “Magnetics on Silicon: An Enabling Technology for Power Supply on Chip,” *IEEE Trans. Power Electronics*, vol. 20, no. 3, pp. 585–592, May 2005.

[50] G. Chen, M. S. Dresselhaus, J.-P. Fleurial, C. Caillat, “Recent Developments in Thermoelectric Materials,” International Materials Reviews 2003, vol. 48 no. 1

[51] R. Hahn, A. Ginolas, M. Schmidt, J. Wolf, V. Glaw, M. Töpper, O. Ehrmann, H. Reichl, “High Power Multichip Modules Employing the Planar Embedding Technique and Microchannel Water Heat Sinks,” IEEE Semi-Therm Symposium, 1997, pp. 48-56

[52] R. Ponnappan, “A Novel Micro-capillary Groove-wick Miniature Heat Pipe,” AIAA 2000-2947, pp. 818-826

[53] A. Mallik, G.P. Peterson, M. H. Weichold, “Fabrication of Vapor-Deposited Micro Heat Pipe Arrays as an Integral Part of Semiconductor Devices,” IEEE 1995, pp. 119-131

[54] D.A. Benson, R.T. Mitchell, M.R. Tuck, D.R. Adkins, D.W. Palmer, “Micro-Machined Heat Pipes in Silicon MCM Substrates,” IEEE 1996, pp. 127-130

[55] I. Yusuf, A. Watwe, H. Ekhlassi, “Integrated Heat Sink-Heat Pipe Thermal Cooling Device,” Inter Society Conference on Thermal Phenomena 2000, pp. 27-30

[56] M. Lee, M. Wong, Y. Zohar, “Integrated Micro-Heat-Pipe Fabrication Technology,” IEEE 2003, pp. 138-147

[57] M. Lee, M. Wong, Y. Zohar, “Design, Fabrication, and Characterization of an Integrated Micro Heat Pipe,” IEEE 2002, pp. 85-88

[58] S. Kalahasti, Y.K. Joshi, “Performance Characterization of a Novel Flat Plate Micro Heat Pipe Spreader,” IEEE 2002, pp. 554-560

[59] P. Rodgers, “Prediction of Microelectronics Thermal Behavior in Electronic Equipement: Status, Challenges, and Future Requirements,” Journal of Microelectronics and Electronic Packaging, vol. 1, no.1 1st Quarter 2004, pp. 16-29

[60] K. Exel and J. Schulz-Harder, Water Cooled DBC Direct Bonded Copper Substrates. Eschenbach, Germany: Curamik Electronics GmbH, 2005.

[61] J. Schulz-Harder, Direct Bonded Copper Substrates for Semiconductor Power Devices. Eschenbach, Germany: Curamik Electronics GmbH, 2005.

[62] Mariya Ivanova, Yvan Avenas, Christian Schaeffer, Jean-Bernard Dezord, and Juergen Schulz-Harder, Heat Pipe Integrated in Direct Bonded Copper (DBC) Technology for Cooling of Power Electronics Packaging, IEEE Transactions on Power Electronics, Vol. 21, No. 6, November 2006, pp. 1541-1546

[63] Schulz-Harder J., Dezord J., Schaeffer, Avenas, Puig O., Rogg. A., DBC (Directed Bonded Copper) substrate with integrated flat heat pipe, EMPC 2005, June 12-15, Brugge, Belgium

[64] P. Cadre, A. Petitbon, N. Changey, E. Ranchiy, New Encapsulation Technique Using Standard DBC Substrate Technology, Alcatel Alsthom Recherche, IEEE 1998, pp. 2103-2107

[65] Dr. Jurgen Schulz-Harder, DBC Substrates as a Base for Power MCMs, 2000 Electronics Packaging Technology Conference, 2000, pp. 315-320

[66] James M. Fusaro, Guillermo L. Romero, Pablo Rodrigucz and Joe L. Martincz, Jr., Thermal Characterization of DBC and MMC Stacks for Power Modules, IEEE 1996, pp. 1411-1417

[67] M. März, P. Nance, “Thermal Modeling of Power-Electronic Systems,” Fraunhofer Institute and Infineon Technologies white paper, April 2000

[68] A. Hefner, D. Blackburn, “Simulating the Dynamic Electrothermal Behavior of Power Electronic Circuits and Systems,” IEEE TPE 1993, pp. 376-385

[69] G. Crowe, “Thermal Analysis and Optimization of a Small, High Density DC Power System by Finite Element Analysis,” IEEE 1996, pp. 718-722

[70] A. Walker, D. Williams, “Thermal Design Consideration in the Design and Application of DC_DC Converters,” IEEE 1996, pp. 990-996

[71] R. Ridely, M. Reynell, S. Kern, “Thermal Considerations for Distributed Power Converter Systems,” IEEE 1999, pp. 886-872

[72] P. Cheasty, J. Flannery, M. Meinhardt, A. Alderman, S. C. O’Mathúna “Benchmark of Power Packaging for DC/DC and AC/DC Converters,” IEEE 2002, pp. 141-148

[73] F. Liang, K. Wolf, V.J. Tottuvelil, F. Alameel, “An Evaluation of Board-Mounted Power Module Packages,” IEEE 2002, pp. 112-118

[74] W. Odendaal, J. Ferreira, “A Thermal Model for High-Frequency Magnetic Components,” IEEE Transactions on Industry Applications, IEEE 1999, pp. 924-931

[75] D. van der Linde, C. Boon, J. Klaassens, “Design of a High-Frequency Planar Power Transformer in Multilayer Technology,” IEEE 1991, pp. 135-141

[76] M. Rascon, J. Ara, R. Madsen, J. Navas, M. Perez, F. San Miguel, “Thermal Analysis and Modeling of Planar Magnetic Components,” IEEE 2001, pp. 97-101

[77] M.M. Yovanovich, “General Expression for Laminar Forced and Natural Convection Heat Transfer from Isothermal Flat Plates for all Prandtl Numbers,” AIAA 1993, #2750

[78] S. Song, S. Lee, V. Au, “Closed-form Equation for Thermal Constriction/Spreading Resistance with Variable Resistance Boundary Condition,” IEPS Conference 1999, pp. 111-121

[79] J. Sturgess, T. Preston “An Economic Solution for 3-D Coupled Electromagnetic and Thermal Eddy Current Problems,” IEEE 1992, pgs. 1267-1270

[80] P. Vong, D. Rodger, P. Coles, H. Lai, “On modelling weakly coupled Electromagnetic-Thermal Problems with Prescribed Non-linear Surface Heat Transfer,” IEEE 2002, pgs. 260-264

[81] J. Brauer, P. Wallen, “Coupled 3D Electromagnetic, Structural, and Thermal Finite Element Analysis as Integral Components of Electronic Product Design,” ISBN 0-7803-3274-1

[82] J. Zhou, M. Xu, J. Sun, K. Yao and F.C Lee, “1 MHz self-driven DC-DC converter for non-isolated 12V VRM,” to be published in IEEE Proc. APEC, 2004.

[83] J. Culham, M.M. Yovanovich, “Factors affecting the calculation of effective conductivity in printed circuit boards,” IEEE ISCTP, 1998, pp. 460-467

[84] J. Lohan, P. Tiilikka, C-M. Fager, J. Rantala, “Effect of both PCB Thermal Conductivity and Nature of Forced Convection Environment on Component Operating Temperature: Experimental Measurement and Numerical Prediction,” IEEE Semi-Therm, 2000, pp. 128-139

[85] L. Ha, G. Matijasevic, P. Gandhi, C. Gallagher, “Thermal study of additive multilayer circuitry on polymer and metal substrates,” 14th IEEE Semi-Therm, 1998, pp. 104-110

[86] Zhenxian Liang; Lee, F.C.; Lu, G.Q.; Borojevic, D., "Embedded Power- A Multilayer Integration Technology for Packaging of IPEMs and PEBBs," *Integrated Power Packaging, 2000. IWIPP 2000. International Workshop on* , vol., no., pp.41-45, 2000

[87] J. Bai, Z. Zhang, J.N. Calata, and G-Q. Lu , “Characterization of Low-Temperature Sintered Nanoscale Silver Paste for Attaching Semiconductor Devices,” *Proceedings of the Seventh IEEE CPMT Conference on High-density Microsystem Design and Packaging and Failure Analysis (HDP'05),* pp.272-276, Shanghai, PR China, June 27 th to June 30 th , 2005.

[88] Pit-Leong Wong’s dissertation, Virginia Tech ETD online library

[89] F. Kreith, M. S. Bohn, *Principles of Heat Transfer*, 5th Edition, PWS Publishing Company, 1997

[90] Ying-Feng Pang’s dissertation, Virginia Tech ETD online library

[91] Ha, L.; Matijasevic, G.; Gandhi, P.; Gallagher, C.; Lee, C.C., "Thermal study of additive multilayer circuitry on polymer and metal substrates," Semiconductor Thermal

Measurement and Management Symposium, 1998. SEMI-THERM Proceedings 1998., Fourteenth Annual IEEE, vol., no., pp.104-110, 10-12 Mar 1998

[92] D. Sterk, Ball A., Xu M. Lee, F.C., “Novel Integrated Transfromer Winding Scheme for Self-Driven ZVS Interleaved Asymmetrical Half-Bridge for Telecommunications Quarter Brick,” *CPES Conference*, *2004*, CPES Proceedings, April 18 2004

[93] Ball, A.; Jinghai Zhou; Lee, F.C., "Thermal design considerations for voltage regulator modules," *CPES Annual Seminar, 2005.* CPES Proceedings, March 6 2005

[94] Ball, A.; Jinghai Zhou; Lee, F.C., "Thermal design considerations for voltage regulator modules," *Applied Power Electronics Conference and Exposition, 2005. APEC 2005. Twentieth Annual IEEE* , vol.1, no., pp.134-140 Vol. 1, 6-10 March 2005

[95] Ball, A.; Lee, F.C., "The Combination of Thermal and Electrical Improvements in a 1U 100A VRM," *CPES Seminar and Annual Meeting, 2006.* CPES Proceedings April 23, 2006

[96] Ball, A.; Sterk, D.; Lee, F.C., "The Combination of Thermal and Electrical Improvements in a 1U 100A VRM," *Power Electronics Specialists Conference, 2007. PESC 2007. IEEE* , vol., no., pp.15-20, 17-21 June 17, 2007

[97] Pinjala, D.; Iyer, M.K.; Chow Seng Guan; Rasiah, I.J., "Thermal characterization of vias using compact models," Electronics Packaging Technology Conference, 2000. (EPTC 2000). Proceedings of 3rd, vol., no., pp.144-147, 2000

[98] Lee, S.; Lemczyk, T.F.; Yovanovich, M.M., "Analysis of thermal vias in high density interconnect technology ," Semiconductor Thermal Measurement and Management Symposium, 1992. SEMI-THERM VIII.,Eighth Annual IEEE , vol., no., pp.55-61, 3-5 Feb 1992

[99] Zach Zenjie, J. Bai, S.-Y. Lee, CPES presentation on Interface Thermal Resistance

[100] ITRS 2003 Roadmap of Handheld Electronics

http://www.itrs.net/Links/2003ITRS/Home2003.htm

[101] Report On The 2006 PSMA Power Technology Roadmap

http://www.psma.com/index.php?directory=fpubl&area=pubs&page=fpublAab#pub4

[102] Dr. Eric Bogatin, “Looking Under the Sheets: Inductance,” article for Printed Circuit Design and Fabrication, November 2007 issue.

[103] Ball, A.; Lim, M.; Gilham, D.; Lee, F.C., "System design of a 3D integrated non-isolated Point Of Load converter," *CPES Power Electronics Conference, 2008*, CPES Proceedings, February 24, 2008

[104] Ball, A.; Lim, M.; Gilham, D.; Lee, F.C., "System design of a 3D integrated non-isolated Point Of Load converter," *Applied Power Electronics Conference and Exposition, 2008. APEC 2008. Twenty-Third Annual IEEE* , vol., no., pp.181-186, 24-28 Feb. 2008

[105] Linear Technology Application Note 103

http://www.linear.com/pc/downloadDocument.do?navId=H0,C1,C1003,C1424,P14856,D10319

[106] NetPower ERS4033x015xxx datasheet

http://www.netpowercorp.com/Products/SaveAttachment.aspx?path=../UserImages/product_datasheet/adminERS1033x015_Specifications.pdf

[107] International Rectifier product information page on voids

http://www.irf.com/product-info/directfet/dfvoiding.html

[108] EFD reflow profiling guide

http://www.efd-inc.com/NR/rdonlyres/30EAB1D2-ED61-4042-849C-E5EADD898599/0/EFDFirstPrinciplesSolderReflow.pdf

[109] EFD solder datasheet

http://www.efd-inc.com/NR/rdonlyres/B0BEBE89-BB04-4AEE-9161-E48C68C3ABB8/0/EFD__Reflow_Profiling_Guide.pdf

[110] Qiang Li, presentation: "Survey of Trends for Integrated Point-of-Load Converters," Power Supply on Chip Symposium (PSoC), Cork Ireland, September 24, 2008.

[111] Web article written by International Rectifier, Jason Zhang, http://eepn.com/locator/products/ArticleID/29270/29270.html

[112] CPES industry presentation by Yan Dong "*Integrated Coupled Inductor POL converters,"* available on CPES website http://www.cpes.vt.edu

[113] Julu Sun; Ming Xu; Lee, F.C.; Yucheng Ying, "High Power Density Voltage Divider and Its Application in Two-Stage Server VR," *Power Electronics Specialists Conference, 2007. PESC 2007. IEEE* , vol., no., pp.1872-1877, 17-21 June 2007

[114] J. Claassens, J.N. Calata, I. Hofsajer, G.Q. Lu, K.D.T. Ngo, "Mechanical Design and Material Selection of a High Temperature Diode Package," CPES Annual Conference, CPES 2007 Proceedings, pp. 127-132, April 2007

[115] M. H. F. Lim, J. D. van Wyk, "Applying the Steinmetz Model to Small Cores in LTCC Ferrite for Integration Applications", in Proc. Appl. Power Electronics Conf., Washington, DC, pp. 1027-1033, Feb 20

[116] C. Harper, *Electronic Packaging and Interconnection Handbook*, 3rd edition, McGraw-Hill Handbooks, 2000

[117] Hashimoto, T.; Kawashima, T.; Uno, T.; Satou, Y.; Matsuura, N., "System in Package with Mounted Capacitor for Reduced Parasitic Inductance in Voltage Regulators," Power Semiconductor Devices and ICapos, 2008. 20th International Symposium, 18-22 May 2008, pp. 315 – 318

[118] Yan Dong, Arthur Ball, David Gilham, Qiang Li, Michele Lim, Fred C. Lee and Ming Xu, "High Power Density Integration of the Point-of-Load Converter," CPES Conference Proceedings, April 2009, pp 1-5.

[119] Yan Dong, David Gilham, Qiang Li, Arthur Ball, Michele Lim, Fred C. Lee and Ming Xu, "Design of LTCC coupled inductors for the integrated POL converters," CPES Conference Proceedings, April 2009, pp. 1-13

[120] Reusch, D.; Lee, F.C.; Ming Xu; Sterk, D., "Improved transformer design for high frequency VRM applications," Applied Power Electronics Conference and Exposition, 2008. Twenty-Third Annual IEEE, 24-28 Feb. 2008 Page(s):1483 – 1489

[121] Liang, Z.; van Wyk, J.D.; Lee, F.C.; Boroyevich, D.; " An integrated power switching stage with multichip planar interconnection construction", 4th Int'l, Volume 1, pp. 364 – 369.

[122] Zhenxian Liang; van Wyk, J.D.; Lee, F.C.; Boroyevich, D.; Scott, E.P.; Zhou Chen; Yingfeng Pang; "Integrated packaging of a 1 kW switching module using a novel planar integration technology", IEEE Trans. Power electronics, Volume 19, Issue 1, Jan. 2004 Page(s):242 – 250.

[123] Y. F. Pang, J. Z. Chen, E. Scott, K. Thole, and D. Boroyevich, “Electrical and Thermal Layout Design and Optimization Considerations for DPS Active IPEM”, Proceedings Center for Power Electronics Systems Seminar, pp. 77–82, 2002.

[124] Pang, Y.; Scott, E.P.; Zhu, N.; van Wyk, J.D.; Liang, Z.; “Assessment of thermo-mechanics for an integrated power electronics switching stage” 39th IAS Annual Meeting, Industry Applications Conference, 2004, Volume 4, 3-7 Oct. 2004, pp. 2309 – 2314.

[125] Bai, J.G.; Guo-Quan Lu; Xingsheng Liu, “Flip-chip on flex integrated power electronics modules for high-density power integration,” IEEE Transactions on Advanced Packaging, Volume 26, Issue 1, Feb. 2003 Page(s): 54 – 59

[126] Xu, Jing, van Wyk, J. D., Ngo, Khai, “Package Alternatives for Integrated Power Electronic Modules (IPEM) with Improved Voltage Rating,” Applied Power Electronics Conference, APEC 2007 - Twenty Second Annual IEEE, Feb. 25 2007-March 1 2007, pp. 416-421

[127] M. H. F. Lim, D. Gilham, J. D. van Wyk, F. C. Lee, K. D. T. Ngo, “Performance Comparison of a Buck Converter Using Shielded-Substrate and Co-Packaged Planar Inductors,” CIPS 2008. Page(s): 59-64

[128] J. M. Lopera, M. J. Prieto, A. M. Pernia, F. Nuno, M. J. M. De Graaf, J. W. Waanders, L. A. Barcia, “Design of integrated magnetic elements using thick-film technology”, IEEE Trans. Power Electronics, vol. 14, no. 3, May 1999, pp. 408 - 414

[129] Lim, M.H.; Dong, J.; van Wyk, J.D.; Lee, T.C.; Ngo, K.D.T., “Shielded LTCC Inductor as Substrate for Power Converter,” Power Electronics Specialists Conference, 2007. PESC 2007. IEEE Volume , Issue , 17-21 June 2007 Page(s):1605 – 1611

[130] Electronic System Cooling, Reference Manual, MAYA Thermal Solutions. Published by SDRC, 1997

[131] Michael Gdula, Alexander Yerman, Vikram Krishnamurthy and Raymond Fillion, "The GE High Density Overlay MCM Interconnect Method Solves High Power Needs of GaAs Designs," IEEE Symposium 1993, pp.339-345

[132] Maxwell 3D User Guide, available from Ansoft's help files or from this website: http://www.scribd.com/doc/13084184/Ansoft-Maxwell-3D-v11-Userguide.

[133] EIA/JEDEC Standard Reference Manual, available from http://www.jedec.org/download/search/jesd51-7.pdf, or from JEDEC directly.

[134] Eberhard Waffenschmidt, Bernd Ackermann, and J. A. Ferreira, " Design Method and Material Technologies for Passives in Printed Circuit Board Embedded Circuits", IEEE Transactions on Power Electronics, vol. 20, no. 3, May 2005, pp. 576-584.

Appendix A

IR DirectFET Synopsis SABER Models used in simulations.

Obtained from International Rectifier website and unmodifed thereof

IRF6633

```
template irf6633 d g s
# Saber Model with Thermal RC Network
#************************************
#      Model Generated by MODPEX     *
#Copyright(c) Symmetry Design Systems*
#         All Rights Reserved        *
#    UNPUBLISHED LICENSED SOFTWARE   *
#   Contains Proprietary Information *
#      Which is The Property of      *
#     SYMMETRY OR ITS LICENSORS      *
#Commercial Use or Resale Restricted *
#   by Symmetry License Agreement    *
#************************************
# Model generated on Feb 21, 05
# MODEL FORMAT: Saber
# Symmetry POWER MOS Model (Version 1.0)
# External Node Designations
# Node d -> Drain
# Node g -> Gate
# Node s -> Source
electrical d,g,s
{
# BODY_BEGIN
spm..model mm=(type=_n,
level=1,is=1e-32,rd=1e-6,
vto=2.39383,lambda=0,kp=68.6566,rs=0.00106046,
cgso=1.05599e-05,cgdo=9.84564e-07)
spd..model md=(is=4.41067e-09,rs=0.00164681,n=1.41634,bv=20,
ibv=0.00025,eg=1,xti=1,tt=1e-07,
cjo=1.23214e-09,vj=0.647736,m=0.377695,FC=0.5)
# Default values used in MD1:
#   RS=0 EG=1.11 XTI=3.0 TT=0
#   BV=infinite IBV=1mA
spd..model md1=(is=1e-32,n=50,
cjo=1.82166e-10,vj=1.86105,m=0.3,fc=1e-08)
# Default values used in MD2:
#   EG=1.11 XTI=3.0 TT=0 CJO=0
```

```
#   BV=infinite IBV=1mA
spd..model md2=(is=1e-10,n=0.4,rs=3e-06)
# Default values used in MD3:
#   EG=1.11 XTI=3.0 TT=0 CJO=0
#   RS=0 BV=infinite IBV=1mA
spd..model md3=(is=1e-10,n=0.4)
spm.M1 n9 n7 s s =model=mm,l=100u,w=100u
spd.d1 s d =model=md
spr.rds s d =3e+07
spr.rd n9 d =0.0001
spr.rg g n7 =3.19203
spd.d2 n4 n5 =model=md1
spd.d3 0 n5 =model=md2
spr.rl n5 n10 =1
spf.fi2 n7 n9 i(spv.vfi2) =-1
spv.vfi2 n4 0 =0
spe.ev16 n10 0 n9 n7 =1
spc.cap n11 n10 =7.1214e-10
spf.fi1 n7 n9 i(spv.vfi1) =-1
spv.vfi1 n11 n6 =0
spr.rcap n6 n10 =1
spd.d4 0 n6 =model=md3
}
# Saber 5-Layer Thermal Model
# The customer will have to have the Optional Template Library License
# to run this part of the library

template irf6633t tj tc
{
rtherm.r1    tj  n4  = 0.667695
ctherm.c1    tj  tc  = 0.000098847
#rtherm.r2    n4  n3  = 1.046285
#ctherm.c2    n4  tc  = 0.00085636
rtherm.r2    n3  n2  = 1.561168
ctherm.c2    n3  tc  = 0.002809435
rtherm.r3    n2  n1  = 29.28222
ctherm.c3    n2  tc  = 0.023433331
rtherm.r4    n1  tc  = 25.45502
ctherm.c4    n1  tc  = 1.257119412
}
```

IRF6691

```
template irf6691 d g s
# Saber Model with Thermal RC Network
#************************************
#      Model Generated by MODPEX     *
#Copyright(c) Symmetry Design Systems*
#         All Rights Reserved        *
#    UNPUBLISHED LICENSED SOFTWARE   *
#   Contains Proprietary Information *
#      Which is The Property of      *
#     SYMMETRY OR ITS LICENSORS      *
#Commercial Use or Resale Restricted *
#   by Symmetry License Agreement    *
#************************************
# Model generated on Mar 23, 04
# MODEL FORMAT: Saber
# Symmetry POWER MOS Model (Version 1.0)
# External Node Designations
# Node d -> Drain
# Node g -> Gate
# Node s -> Source


electrical d,g,s
{
# BODY_BEGIN
spm..model mm=(type=_n,
level=1,is=1e-32,rd=1e-6,
vto=2.6309,lambda=0,kp=445.592,rs=0.000647069,
cgso=6.07538e-05,cgdo=6.07106e-08)
spd..model md=(is=7.29907e-09,rs=0.00390276,n=0.800767,bv=20,
ibv=0.001,eg=1,xti=1,tt=1e-07,
cjo=5.06928e-09,vj=1.11502,m=0.589727,FC=0.5)
# Default values used in MD1:
#   RS=0 EG=1.11 XTI=3.0 TT=0
#   BV=infinite IBV=1mA
spd..model md1=(is=1e-32,n=50,
cjo=1.0326e-09,vj=8.0469,m=0.3,fc=1e-08)
# Default values used in MD2:
#   EG=1.11 XTI=3.0 TT=0 CJO=0
#   BV=infinite IBV=1mA
spd..model md2=(is=1e-10,n=0.400032,rs=3e-06)
# Default values used in MD3:
#   EG=1.11 XTI=3.0 TT=0 CJO=0
```

```
#   RS=0 BV=infinite IBV=1mA
spd..model md3=(is=1e-10,n=0.400032)
spm.M1 n9 n7 s s =model=mm,l=100u,w=100u
spd.d1 s d =model=md
spr.rds s d =1e+06
spr.rd n9 d =0.0001
spr.rg g n7 =1.78576
spd.d2 n4 n5 =model=md1
spd.d3 0 n5 =model=md2
spr.rl n5 n10 =1
spf.fi2 n7 n9 i(spv.vfi2) =-1
spv.vfi2 n4 0 =0
spe.ev16 n10 0 n9 n7 =1
spc.cap n11 n10 =2.63006e-09
spf.fi1 n7 n9 i(spv.vfi1) =-1
spv.vfi1 n11 n6 =0
spr.rcap n6 n10 =1
spd.d4 0 n6 =model=md3
}


# Saber Thermal Model
# The customer will have to have the Optional Template Library License
# to run this part of the library
template irf6691t tj tc
{
rtherm.r1    tj  n3  = 0.678448546
rtherm.r2    n3  n2  = 17.29903034
rtherm.r3    n2  n1  = 17.56646856
rtherm.r4    n1  tc  = 9.470128245
ctherm.c1    tj  tc  = 0.001267598
ctherm.c2    n3  tc  = 0.033386842
ctherm.c3    n2  tc  = 0.508924145
ctherm.c4    n1  tc  = 11.19309024
}
```