International Rectifier

VRM Design Optimization for Varying System Requirements

George Schuellein, Director Computing Design & Applications Center
International Rectifier – 233 Kansas St., El Segundo CA 90245

As presented at Powersystems World, Nov 03

Introduction

Intel, AMD, and other microprocessor vendors publish specifications for VRMs (Voltage Regulator Module) intended to provide power for a particular processor family. In addition to establishing electrical performance, these specifications often define the mechanical dimensions, VRM to motherboard connector, capacitors mounted on the motherboard, and environmental factors such as airflow and ambient temperature.

In practice computer OEMs rarely follow the microprocessor vendor's VRM guidelines. Although the electrical specifications are generally adhered to the overall design of a particular system must be considered resulting in deviation from the standard specification. The computer OEM commonly specifies the mechanical, environmental, and capacitor requirements that are appropriate for a given system.

This paper will review the power requirements of current and future generation microprocessors, various system level requirements, and a number of different VRM designs. The designs all meet the latest Intel VRM10.x electrical requirements and deliver approximately 80A to 150A at less than 1.3V. Each design targets different system requirements and was optimized to achieve different size, efficiency, capacitor, thermal, and cost goals.

Microprocessor Power Trends

Over the past 10 years power requirements for state of the art microprocessors used in high end desktops and servers has followed the trend shown in Figure 1. Increases in operating frequency and introduction of new instruction sets leads to increases in supply current requirements while the transition to improved semiconductor processes with smaller process geometries reduces operating voltages and supply currents.

Prior to 2000 reductions in voltage kept pace with increases in current resulting in total power dissipation of less than 25W. With the introduction of processors such as Intel's Pentium 4 and AMD's Athlon currents began to outpace the drop in voltage and power began to significantly increase. This trend will accelerate as the ability to reduce supply voltage is hitting some fundamental barriers. Voltage drops induced by the impedance between the die and its power source threaten signal integrity and ultimately the promise of Moore's law to continuously improve processor performance.

Starting around 2004, processor voltage will stabilize around 1V with further reductions requiring significant technical advancements. To maintain increasing performance supply currents will rapidly increase from 100A in 2004 to 150A+ in 2006. This represents a significant challenge to the processor power supply.

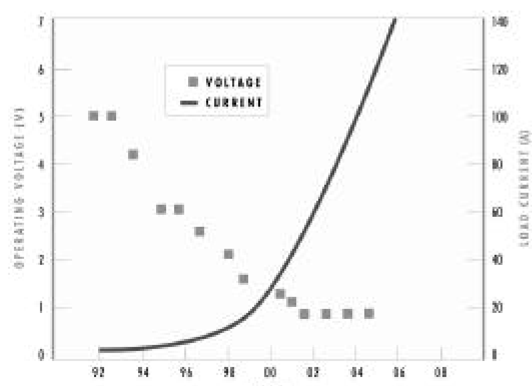

Figure 1 – Microprocessor Power Trends

VRM10.x Power Requirements

Intel will introduce its next generation desktop processor codenamed Prescott in the 4th quarter of 2003. A new voltage regulator specification, VRD10.0 has been developed [1] to support it. VRD10.0 introduces some significant new requirements compared to the previous VR9.x specification;

- ?? New "Dual Range" 6 Bit VID with 0.8375-1.5V range and 12.5mV LSB
- ?? Dynamic changes in VID code
- ?? One 12.5 mV step every 5 µs, up to 36 steps (450mV) in 180 µs.
- ?? 1.3 mohm load line out to 90A a with 50mV tolerance band (see Figure 2)
- ?? Differential Remote Sensing at CPU pins
- ?? Supply Voltage may overshoot 50mV for 25us during load step down

Subsequent versions of VR10.x will be introduced to support the 2004-2006 generation of Intel processors. While the VRM feature set is expected to remain the same, currents will increase to the 150A range and load line impedance and tolerance will see significant reductions to keep processor die voltage around 1V.

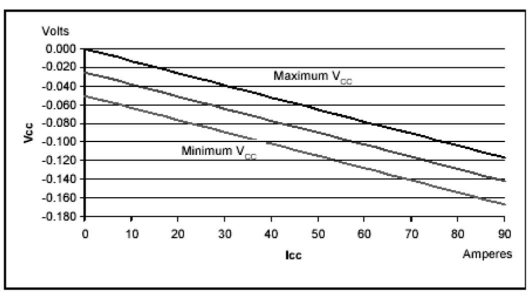

Figure 2 – VRD10.0 Load Line

Varying System Requirements

Although microprocessor voltage regulator specifications are generally adhered to by large computing OEMs, short cuts are routinely taken by commodity "white box" desktop motherboard OEMs in order to reduce cost. Common examples include not designing for the maximum specified current or skimping on the number of output caps required to maintain regulation during load transients. This can lead to "blue screen" lock-up of the computer in some circumstances but this has generally been tolerated by the end customer as part of the experience of buying PCs priced under $1000. It remains to be seen whether this practice will continue as signal integrity and thermal margins decrease.

Another varying requirement is the desired useful life of a particular motherboard or system. Microprocessor vendors commonly provide guidelines for the VR (voltage regulator) requirements of the highest speed grade expected to be used in a particular processor/chipset platform. Motherboard and computing OEMs can support this upgrade path or something less to reduce cost. Due to the relatively long life of most servers compared to desktops, OEMs often wish to provide support for speed grades beyond even the most distant processor specification. Thus we see VR10.x server designs being specified today to handle the 150A currents expected out in the 2006 timeframe.

Mechanical dimensions impose additional system constraints that effect VR requirements. Power supply specifications from Intel and other processor vendors generally assume a traditional tower mount desktop or server system. However, smaller form factors are emerging and gaining market share such as the 1U rack mount server where the total height of the server is only 1.75 inches. This precludes the use of low cost electrolytic bulk capacitors on the motherboard forcing the use of more expensive alternatives. The height of the VRM itself is constrained to only 1 inch instead of the standard 2.5 inches. The small physical size results in higher thermal impedance requiring a much lower power loss in order to avoid component overheating. All these factors tend to drive up cost; in general reducing voltage regulator size will increase cost.

System airflow is another important factor. In most desktops there is no airflow dedicated to VR cooling. The VR usually receives some residual airflow in the range of 150-200 LFM (Linear Feet per Minute) from the CPU fan. Airflows greater than 200 LFM can produce acoustic noise levels that are unacceptable for the desktop environment. Servers generally have more airflow available as they typically reside in a data center where higher noise levels are acceptable. Airflow has a huge impact on keeping component temperatures under control. Increasing airflow allows higher VR power loss and lower cost.

Finally, reliability requirements can have an impact on VR design. High temperature is the single most important factor leading to reduced component lifetime. PCB (Printed Circuit Board) temperatures are generally the limiting factor and should be kept to less than 105C for high reliability when using low cost FR4 material. This practice is often violated in low cost desktop and server solutions where PCB temperatures exceeding 125 deg C are commonly observed if the VR is tested to the processor thermal design current.

The reliability of low cost electrolytic capacitors is generally only several thousand hours if operated at their rated ripple current. Use of higher cost alternatives such as ceramics or Tantalum is required if high reliability is desired. To minimize the cost impact, higher number of phases or increased switching frequency can be used to improve the transient response and reduce the number of output capacitors.

Higher switching frequency also allows the use of physically smaller inductors. The penalty of increasing switching frequency is a corresponding increase in power loss. Figure 3 provides a normalized graph of power loss versus switching frequency for a synchronous buck converter optimized for low switching losses. In most cases the benefit of smaller component size is negated by the problem of dissipating the additional power loss associated with higher switching frequency. With components presently available, the optimal switching frequency will be in the range of 250-500kHz.

2 Phase VRD Converter

The first approach reviewed in this paper targets initial VRD10.0 processors and is designed to provide a 80A thermal design current. It uses a basic two phase VRD (voltage regulator down) configuration to eliminate the connector, additional PCB, and assembly costs associated with a VRM. Figure 4 provides a photo of the design. The PCB area required for the VRD is approximately 5 by 1.5 inches. Consistent with standard desktop practices 4 layers are used with half once copper on the external layers and one ounce copper on internal.

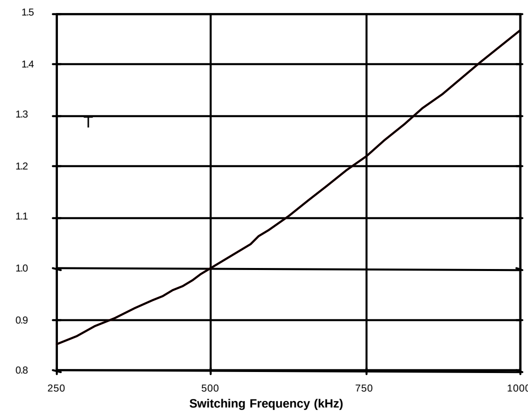

Figure 3 – Normalized Power Loss versus switching frequency

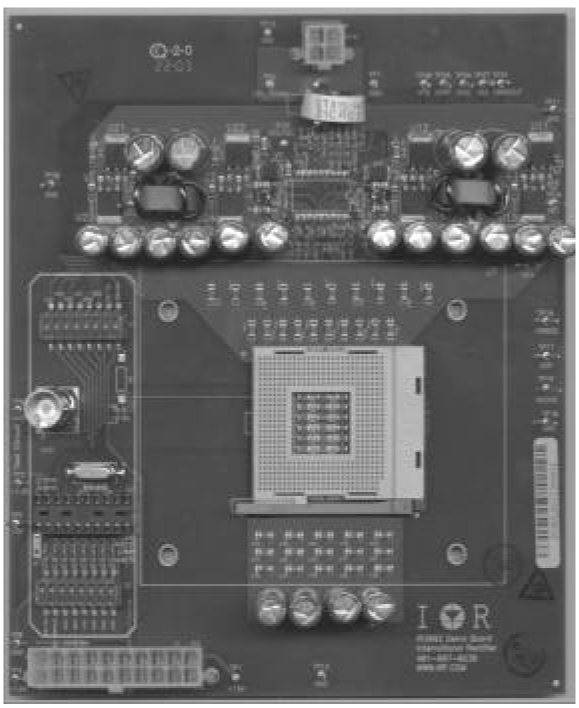

Figure 4 – 2 Phase VRD

A control IC with integrated gate drivers is used to minimize cost. A total of eight MOSFETs in Dpak are used with 2 paralleled for the control and synchronous rectifier function per phase. Different FET types are used for the control and synch positions to optimize performance. Low RDSon is required for the sync FETs to reduce conduction losses as they will be on during most of the switching cycle. Low gate charge is required in the control FETs to reduce switching losses. The sync FETs have a 4.5 mohm (typ) RDSon @ 10V Vgs and 23nC (typ) total gate charge while the control FETs have 8.4 mohm (typ) RDSon @ 10V Vgs and 8.3nC (typ) total Gate Charge.

The output filter consists of low cost 0.9uH Toroid output inductors (0.8 mohm ESR) and 16 low cost 560uF Al-Poly bulk output caps (8mohm ESR) along with 28 10uF ceramics in the 1206 package. To minimize parasitic impedance associated with the PCB 16 of the ceramic capacitors are located in CPU socket cavity. A switching frequency of 250kHz is selected to maximize efficiency.

Efficiency and power loss of this design are provided in figures 5 and 6. Peak efficiency of around 85% occurs at only 20 amps and drops to around 75% at 90 amps. Measurements made at 25C ambient and with 200 LFM airflow show component temps to be in the 110 to 120 deg C range. Although consistent with common desktop design practices this approach will clearly not scale to the higher currents that will be required in the future. The 0.5 watts per Amp power loss of the 2 phase VRD would result in an unacceptable 75W power loss at an output current of 150A.

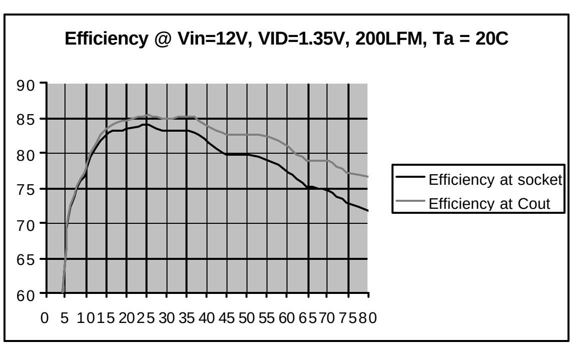

Figure 5 – 2 Phase VRD efficiency

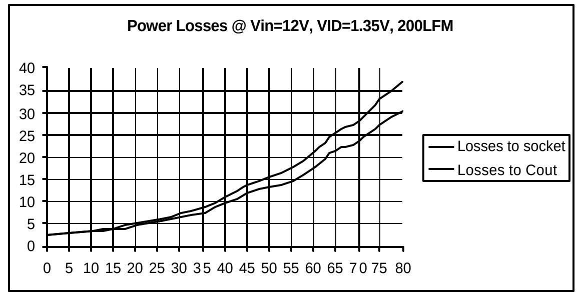

Figure 6 – 2 Phase VRD Power Loss

The load transient response of the 2 phase VRD is shown in figures 7, 8, and 9. Although the PWM control IC quickly changes duty cycle it takes about 8us to increase the inductor currents and 12us to decrease them. The bulk capacitors supply the load current while the inductor current ramps.

3-6 Phase standard VRM

The next design approach to consider is a VRM in Intel standard format. Dimensions are 3.8 inches by 2.5 inches with a "gold finger" connector to the motherboard. There are several key advantages of the VRM over the VRD, one of which is that construction does not need to be compatible with motherboard practices. A 6 layer PCB is used with 4oz copper on internal layers and 3oz for external layers. This provides much lower losses in the PCB copper and also enhances thermal dissipation. Using a VRM frees up valuable motherboard real estate, improves cooling by getting the VR up and into whatever airflow is available, and reduces noise and heat transfer into the mother board. The gold finger connector allows easy field upgrade and replacement of the VRM.

A flexible design is chosen allowing 3 to 6 phases to be implemented with the same PCB layout. A scalable IC architecture is used with a single control IC providing support for 3 to 6 phase ICs each of which includes gate drivers and circuitry required to monitor and control a single phase. To improve performance an enhanced MOSFET package is selected to reduce package losses and provide a means to directly attach a heat sink for additional cooling. Surface mount inductors with lower inductance and ESR (0.36uH and 1.1 mohm) are used to improve efficiency and transient response. Input and output capacitors remain on the motherboard.

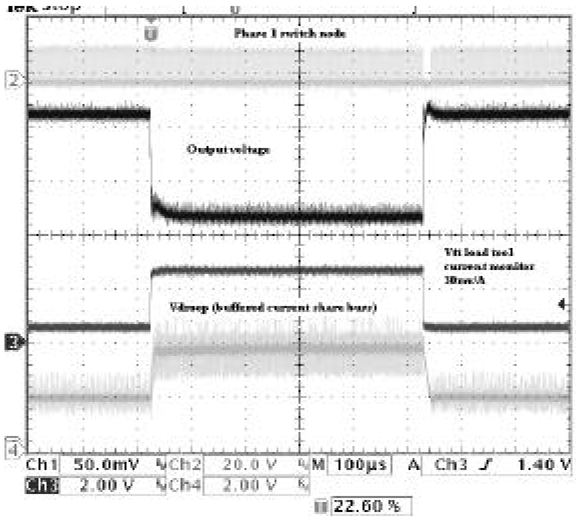

Figure 7 – Load Transient turn-on-off-

o

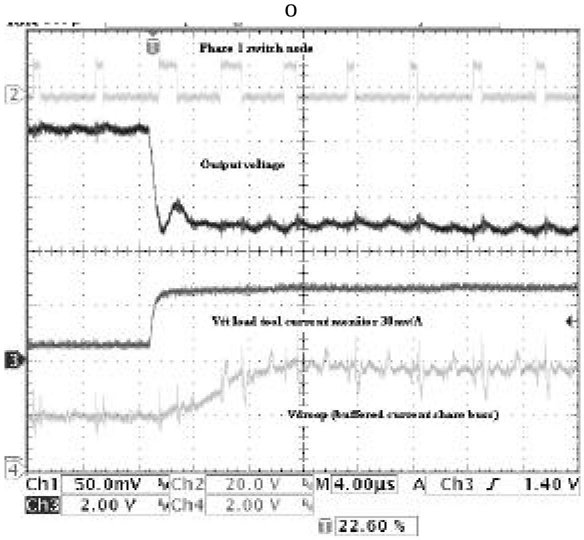

Figure 8 – Load Transient turn-on

Pictures of the standard VRM with and without heat sink attached are shown in figures 10, 11, and 12. A switching frequency of 300kHz was chosen to maximize efficiency, slightly higher than the 2 phase VRD due to the lower value output inductors.

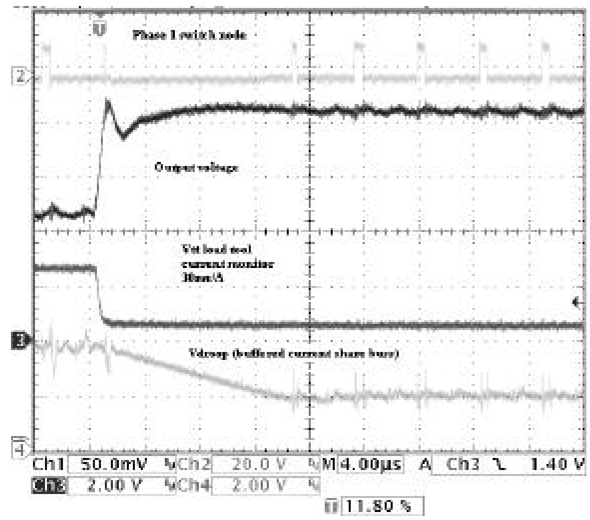

Figure 9 – Load Transient turn-off.

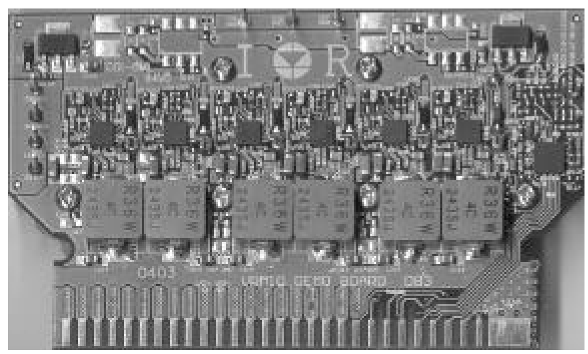

Figure 10 – Standard VRM (front view)

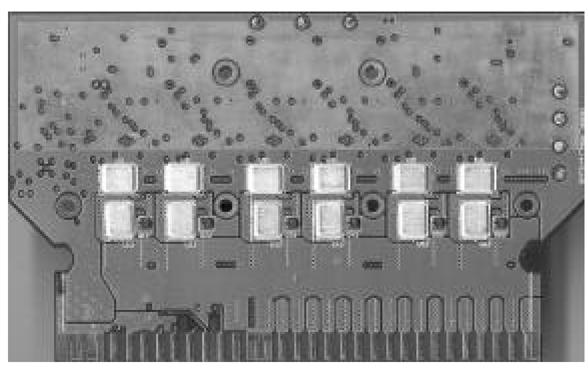

Figure 11 – Standard VRM (back view)

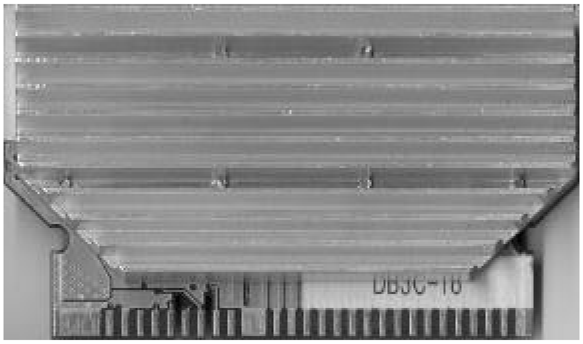

Figure 12 – Standard VRM with Heat Sink

Efficiency and power loss of the standard VRM with 3, 4, 5, and 6 phases is shown in figures 13 and 14. Peak efficiency for the 3 and 4 phase is around 40A and 60A respectively. The 5 and 6 phase designs both peak around 70A but the 6 phase efficiency drop off with increasing current is more gradual. At an output current of 120 amps the 6 phase power loss is around 24 watts while the 3 phase is about 50% higher at over 36 watts.

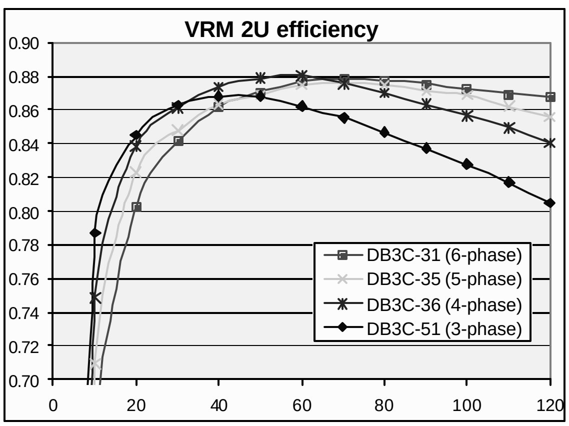

Figure 13 – Standard VRM 3-6 phase efficiency

Further investigation of the thermal performance was pursued using the 6 phase standard VRM with the results shown in figure 16. Airflow was adjusted to 200, 300, 400, and 600 LFM and PCB temperature recorded. At 200 LFM maximum load current reduced to about 90A while at 600 LFM an output current of up to 150A may be possible. Thermal resistance versus airflow was extrapolated and is shown in figure 17. The effect of airflow was a somewhat linear reduction in thermal resistance resulting in a considerable increase in output current capability.

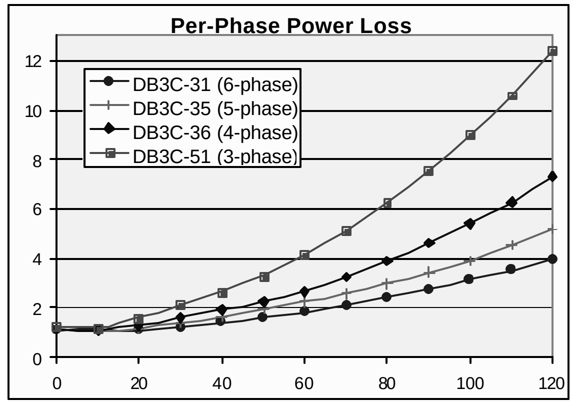

Figure 14 – Standard VRM 3-6 phase power loss

Wind tunnel testing of the standard VRM with 4, 5, and 6 phases was performed to investigate thermal performance and the results are provided in figure 15. With a 45C ambient temperature and airflow of 400 LFM output current was increased until PCB temperature exceeded 105 deg C. Under these conditions the 4, 5, and 6 phase designs can provide approximately 100, 110, and 120A of output current respectively.

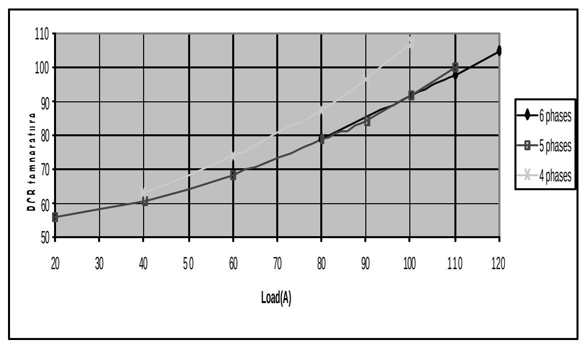

Figure 15 – Thermal evaluation of standard VRM

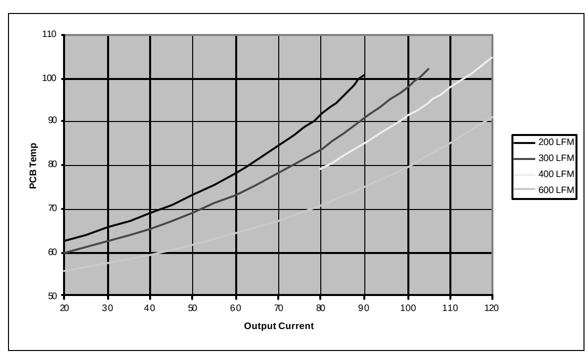

Figure 16 – PCB temperatures with varying airflow

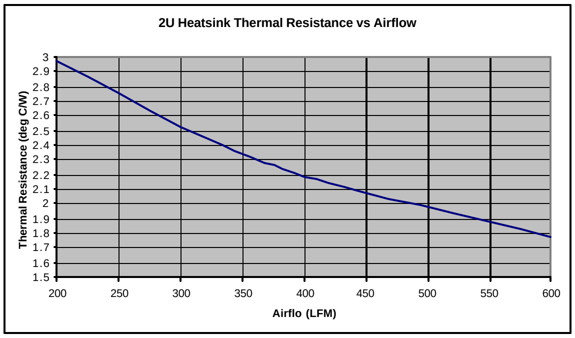

Figure 17 - Thermal resistance versus airflow

1U 5 Phase VRM

The next design to be reviewed is a VRM that meets the size constraints of a 1U server using the Intel defined gold finger connector. PCB dimension were reduced to 3.8 inches by 0.985 inches. A five phase design was selected using improved components compared to the standard VRM. To reduce size and power loss very small 0.22uH surface mount inductors with a reduced ESR of only 0.44 mohm were selected. Improved MOSFETs were used with lower RDSon and gate charge than had been used in the standard VRM.

A 400kHz switching frequency was determined to provide the optimal efficiency. Although ripple current is greater than in the standard VRM the reduction in inductor losses more than compensated for increased conduction losses in other components. Pictures of the 1U VRM are provided in figures 18, 19, and 20.

Efficiency, power loss, and PCB temperature were measured at 25C ambient temperature with an air flow of 400 LFM. The results are provided in figures 21, 22, and 23. Peak Efficiency increased to nearly 90% and was still over 88% at an output current of 105A. Power loss was under 18W at 105A or 0.17W/A, a vast improvement over the 2 phase VRD.

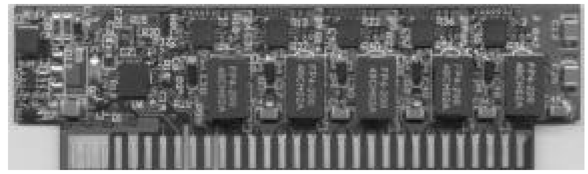

Figure 18 – 1U VRM (front side)

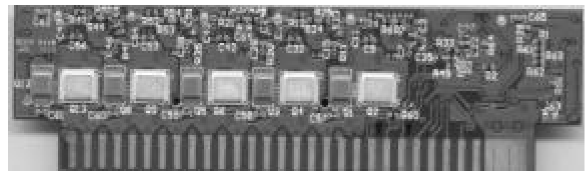

Figure 19 – 1U VRM (back side)

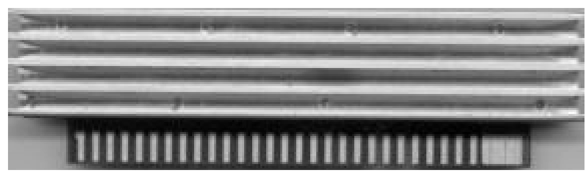

Figure 20 – 1U VRM with heat sink

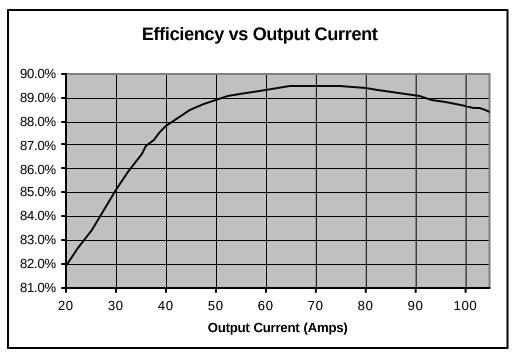

Figure 21 – 1U VRM Efficiency

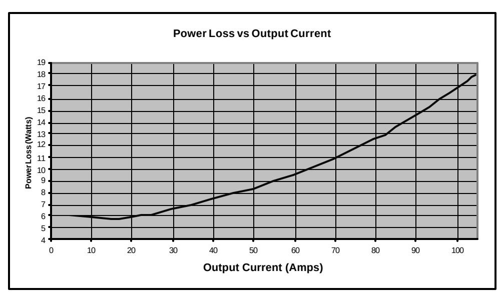

Figure 22 – 1U VRM Power Loss

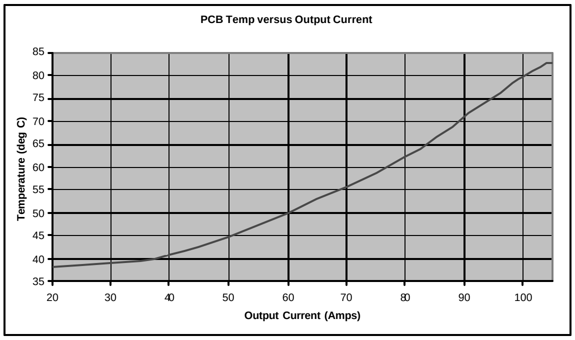

Figure 23 – 1U PCB Temperature

Solder-in VRM

The final design to be reviewed uses the same components as the 1U VRM, but with a RDI header. This approach has the advantage of a lower cost connector than the standard gold finger and does not require gold plating of the VRM PCB. Other advantages include a lower profile since the VRM is very close to the motherboard. Connector impedance is decreased resulting in reduced power losses and improved transient response. The main disadvantage of the solder-in VRM is that it can't be easily replaced in the field. A picture of the solder-in VRM is provided in figures 24 and 25.

Conclusions

There is no "one-size–fits-all" solution for VRM design. Every design should be approached with the system requirements in mind. A clear understanding of the load, mechanical dimensions, available airflow, and expected lifetime and reliability are essential. The approach selected for a particular VRM should attempt to find a solution to the requirements while minimizing cost.

This paper has provided a number of design examples and provided data on the performance that can be expected. A clear path has been outlined to migrate from the 60-80 amps needed today to the 150 amp that will be needed within the next several years. It can be expected that continuous improvements in component performance will ease the transition.

Reference

[1] Intel Corporation, "Voltage Regulator-Down (VRD) 10.0 design guide", April 2003, http://www.intel.com/design /Pentium4/guides/25288501.pdf

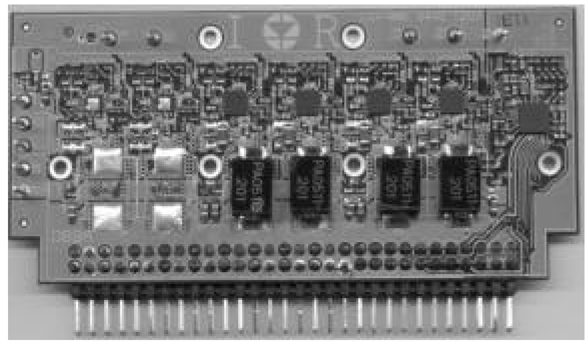

Figure 24 – Solder-in VRM (front side)

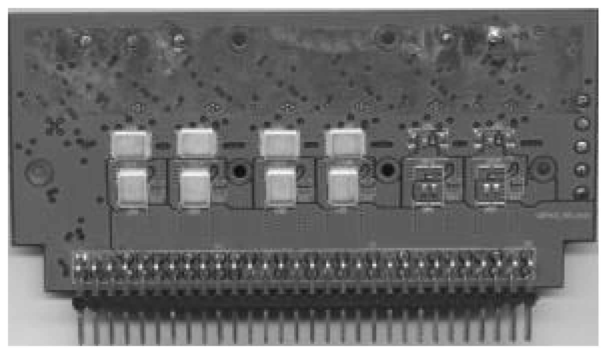

Figure 25 – Solder-in VRM (back side)