

may contain hidden “time bombs.” These designs meet neither the CPU power-supply demands nor the expected life-span characteristics and will eventually cause computer shutdown. The main causes of failure are poor electrical designs and underspecified discrete components.

A good design can satisfy all of the requirements of the microprocessor-core voltage supply and provide a product with a long life.

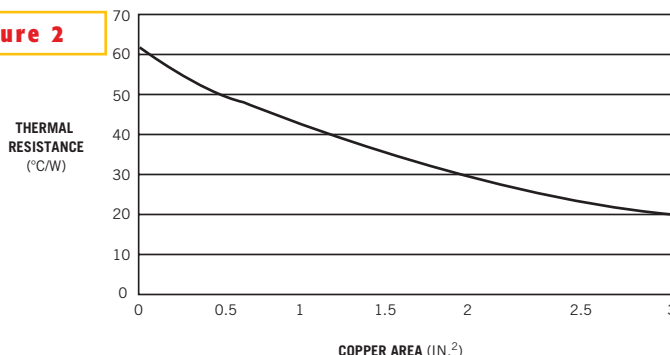
For example, you can design a 60A CPU-core voltage regulator using the ADP-3160 synchronous buck controller, which has a flexible, multiphase architecture (Figure 1). The circuit’s compensation scheme helps the circuit achieve the optimal response to meet the VRM9.0 dc/dc static and transient specifications. Using a four-phase, 120A controller is also possible. Such designs allow building low-gain, optimal-response power supplies that fully comply with the latest power-supply specifications.

Table 1 shows the main voltage and current specifications for 1.6V VRM9.0-compliant CPUs. Of particular interest are the tight tolerances of the core voltage at the processor pins and the high slew rate of the output current. The supply current for a single CPU can jump from a negligible value to 60A with a slew rate of 50A/μsec. During this transient event, the core voltage must stay within –125 to 0 mV of V_{ID} (voltage-identification code).

A CASE OF OVERSPECIFICATION?

VRM8.4, the CPU core’s voltage specification for Pentium III and Celeron processors with supplies of 1.5 to 2.05V, specifies a typical static variation of ±3.5% and a dynamic variation of ±7% with a slew rate of 20A/μsec at full-load excursions. The core-voltage power sup-

Figure 2



The junction-to-air thermal resistance, T_{JA}, of a D2PAK case in still air is approximately 37°C/W for a 2-oz copper area of 1 in.².

ply is either a socketed VRM (voltage-regulator module) or is part of the motherboard’s one- or two-phase buck regulator.

The newer VRM9.0 specification is much more demanding. The transient-voltage-regulation specification is 0/–7% with slew rates as high as 50A/μsec. Achieving these transient specifications is difficult, and most power-supply designs meet only the static-regulation limit. Most designs usually violate the same limit for output-voltage transients, and these designs are thus not fully compliant.

You should pay special attention to the output-current slew rate of 50A/μsec. Such fast transients make it practically impossible to meet the output-voltage dynamic tolerances without special design tricks. Consider a VRM on a socket next to a processor. Because the CPU and the VRM power supply are not integrated into one unit, there is some physical distance between them. Even with the minimum distance between the CPU and the VRM power supply, the pc-board traces between the two would have a total length of approximately 1 to 1.5 in. The inductance of the stand-alone, not-sandwiched pc-board traces would have a typical range of 10 to 25 nH, depending upon their geometry.

Assuming a favorable case of only 10

nH/in. and a total trace length of 1 in., the total wire’s parasitic inductance between the CPU and VRM hardware would be 10 nH. The voltage step at the CPU, therefore, is: $V = -Ldi/dt$, so $\Delta V = 10 \text{ nH} \times 50 \text{ A}/\mu\text{sec}$ or $\Delta V = 500 \text{ mV}$.

These 500-mV deviations should cause the CPU to make many errors because the VRM9.0 specification tolerances are only 0/–125 mV. Interestingly enough, the

VRM/CPU combination still works well, even with CPUs that operate at 900 MHz or more. In this case, the specification has fallen prey to overspecification, one of the seven sins of power-supply design. But the message is clear: The interconnection parasitics will soon limit the CPU power-supply quality and thus its speed and reliable operation.

To minimize trace inductance, the solution is to interleave the V_{CC} and ground planes in a layered-sandwich construction and to interleave layers when there are more than two. If L is the length, W is the width, H is the distance between layers of the trace islands, and N is the total number of copper planes in mils, then the leakage inductance for both negative and positive voltage planes is:

$$L = 31.9 \times L \times H / [W \times (N-1)]$$

in picohenries.

Two 1×1-in. copper traces with a spacing of 20 mils has a total trace inductance of approximately 640 pH; four such interleaved planes has 210 pH. These values are much lower than the inductance of single plane, which is approximately 10 nH.

MAKE DESIGN TRADE-OFFS

To meet motherboard cost targets, designers often use low-cost, low-ESR (equivalent-series-resistance) electrolytic capacitors. Reducing the capacitor value and size too much, however, reduces the expected life span of the product. The current flowing through the ESR of the capacitors dissipates power and causes temperature to rise. Capacitor life is inverse-

TABLE 1—INTEL VRM9.0 CPU-VOLTAGE AND -CURRENT SPECIFICATIONS

Symbol	Parameter	Minimum	Maximum	Units
V _{OUT VRM}	VRM output voltage at VRM connector	1.53	1.600 (=V _{ID})	V
V _{OUT CPU}	V _{CC} core voltage at processor socket	1.475	1.600 (=V _{ID})	V
V _{MAX}	Maximum nonoperating voltage		2.1	V
I _{OUT MAX}	Maximal static current		60	A
di _{OUT/dt MAX}	Output-current slew rate		50	A/μsec

ly proportional to temperature, size, and ESR.

Another trade-off is the thermal design versus the cost of the power supply. The motherboard-manufacturing process calls for 100% SMT (surface-mount technology) components, and most motherboards on the market use D2PAK- or DPAK-packaged power FETs in their power regulators. The achievable thermal impedance with convection cooling when a power transistor is soldered to the copper area of the pc board depends on the area and thickness of the copper pad and the velocity of the cooling air. **Figure 2** shows the junction-to-air thermal resistance of a D2PAK case in still air.

With an output power of 192W (1.6V/120A) and 84% efficiency, the MOSFET losses in each of the four phases are approximately 4.8W per pair, or 2.4W per transistor for a properly designed regulator. Assuming a maximum temperature rise of 60°C, each transistor would require a minimum pc-board area of 2.5 in.². This area is too large for practical regulator designs.

Consequently, you must use forced air to reduce the area to an acceptable size, and multilayer, 2- to 6-oz copper would also help. Successful pc-board thermal design for currents higher than 50 to 60A with no external heat sinks is a challenging task and requires a variety of analytical and experimental work. The ADP-3160 1.6V, 120A demo board is an example of such a combination of thermal and electrical design. The maximum temperature rise per MOSFET is 60°C at full load without heat sinks, with forced-air cooling of 400 lfm. This air velocity is typical for a workstation application.

LAYOUT CONSIDERATIONS FOR 120A

In less than two years, the peak output current that a VRM must supply to a CPU has risen from 14 to 120A, which is approximately a ninefold increase. The important issues related to the layout of the VRM/CPU combination on the

motherboard are static regulation and dynamic regulation.

Standard copper thickness per layer of a CPU motherboard is 1.4 mils for a 1-oz board. **Table 2** gives some information about the current-carrying capacity of 1-oz copper. For 0.5- or 2-oz copper, multiply or divide the 1-oz data by 2. You should expect errors as great as 20% when calculating maximum current.

A better solution is 2-oz copper or 4- to 10-oz copper if the number of layers is four or fewer. Eight-layer motherboards give good performance for high-current applications. However, if all layers have 1-oz thicknesses, you should place four or more layers in parallel for every output terminal to achieve reasonably low total trace resistance. (Some boards have different copper thicknesses for each layer.) To achieve tight voltage regulation and good efficiency, your goal should be no more than a 20°C temperature rise of copper trace due to resistive losses, which requires a minimum of four layers with 1-in.-wide, 1-oz copper for each output trace.

Minimizing the effect of trace length is also critical for dynamic performance. Every tenth of an inch counts, especially after the output-capacitor bank. One inch of 1-in.-wide, 1-oz copper gives a resistance of 0.92 mΩ and an inductance of 10 to 20 nH—values high enough to spoil dreams of compliance with static and dynamic-specification limits. To reduce the inductance by a factor of 20 to 50, you should use a sandwich construction of power planes. For good results, four-layer pc board is a minimum.

AVOID COMMON LAYOUT SINS

For a successful design, avoid layout overlaps between the switching nodes and the signal nodes. These overlaps are ideal ways to inject switching noise into the converter-control circuit. These violations can lead to poor regulator performance or even complete malfunction of the design.

Overlap between switching nodes and

input and output traces is also a bad idea. Switching nodes can easily inject large amounts of noise to nearby traces and loops via unavoidable parasitic capacitance and inductance. You can minimize but not eliminate parasitic elements through proper schematic and layout design. The best way to reduce the effect of parasitics is to minimize not only their values, but also the noise that the switching node generates. You can either minimize the layout area of the node or bypass the high-frequency noise of the node by placing a high-quality capacitor across the node.

Routing critical signal lines, such as voltage- and current-sense lines, through power circuitry also causes problems. It is best if you interpose a signal-ground plane between the signal lines and the traces of the power circuitry. The output capacitors should be as close as possible to the microprocessor. If the load is distributed, as in the case of multiprocessor systems, the capacitors also should be distributed in proportion to where the load is more dynamic.

A small ferrite-bead inductor placed in series with the drain of the lower MOSFET is the best way to prevent this source of switching-power loss. The routing of the output-power path, although less critical than the switching-power path, should also encompass a small area. The output-power path is the current through the inductor, through the output capacitors, and back to the input capacitors. □

AUTHOR'S BIOGRAPHY

Angel Getchev, PhD, is staff application engineer at Analog Devices Inc (Santa Clara, CA), where he has worked for more than two years. In this current position, he provides research, development, and application support in power management. He received his PhD in Space Power Electronics from the Bulgarian Academy of Sciences in 1986. He enjoys acoustic and rock guitar, music, and windsurfing.

REFERENCES

1. Liswood, Woody, "Whose time bomb is it anyway?" *Microtimes*, Oct 13, 1997.
2. Analog Devices Inc ADP3160, ADP-3412 data sheets
3. Hogan, Jerry, *et al*, "The seven sins of power-supply design," PCIM Conference, July 1998.

TABLE 2—CURRENT-CARRYING CAPACITY OF 1 OZ COPPER

Line width, (in.)	Resistance per foot	Cross section area (mils ²)	maximum current per 20°C temp rise	Maximum current per 40°C temperature rise
0.25	0.048	340	11	16
0.5	0.024	680	19	26
0.75	0.016	1010	25	34
1.0	0.011	1350	30	44