

f e a t u r e

Multiphase Voltage Regulator Tackles 150-A Applications

Ralph Monteiro
International Rectifier

George Schuellein
International Rectifier

Power demands of the latest microprocessors from AMD, IBM, Intel, and other vendors continue to soar.

Indeed, CPUs currently under development for next-generation servers require 150 A or more, while desktop processors require 120 A. Along with these stringent requirements come additional constraints such as cost, board area (power density), and thermal environment, which further inflame the design challenge.

To explore the various options for voltage-regulator design, this article will review several multiphase buck voltage-regulator-down (VRD) solutions that target the latest CPUs in server applications. A thermal design current (TDC) of 130 A, with a peak current of 150 A, will be used as a design target. The evaluation will be based on the use of an eight-layer pc board or PCB, which is commonly implemented in server motherboards. The power and ground planes use 2-oz copper (middle layers), while the rest of the layers are 1-oz copper.

The thermal challenge

When delivering continuous currents of 130 A at 1.35 V_{OUT}, an 85% efficient VRD design would need to dissipate about 31 W to the ambient—a requirement that makes the thermal design a

formidable task. In rack-mounted server environments, ambient temperatures inside the chassis are as high as 45°C, further raising the performance bar. In such a system, saving every watt of power loss is important because the heat-saturated PCB has a high junction-to-ambient thermal resistance. PCB temperatures generally must be kept below 105°C to maintain reliability. In most cases, this is the factor that restricts the amount of voltage-regulator power loss, rather than the component temperatures that generally can be much higher.

The thermal design problem can be solved either by increasing the voltage regulator's efficiency, by finding a more effective way to remove heat from the PCB, or by doing both simultaneously. For a given set of controllers, drivers, and passives, efficiency improvements can be achieved quite simply by increasing the number of MOSFETs paralleled or by upping the number of phases in a design.

Having more paralleled MOSFETs helps reduce MOSFET conduction losses, but it doesn't help reduce the stress on other components, such as the inductors. Paralleling more MOSFETs per phase also increases layout complexity and switching losses, limiting the design's switching frequency. Furthermore, paralleling MOSFETs increases the area of the switch node (the connection point of the MOSFETs and inductor), which causes an increase in radiated noise. Because high-speed digital lines usually must be run under

ANALOG & POWER



Don Tuite

Forecast Kudos, APEC

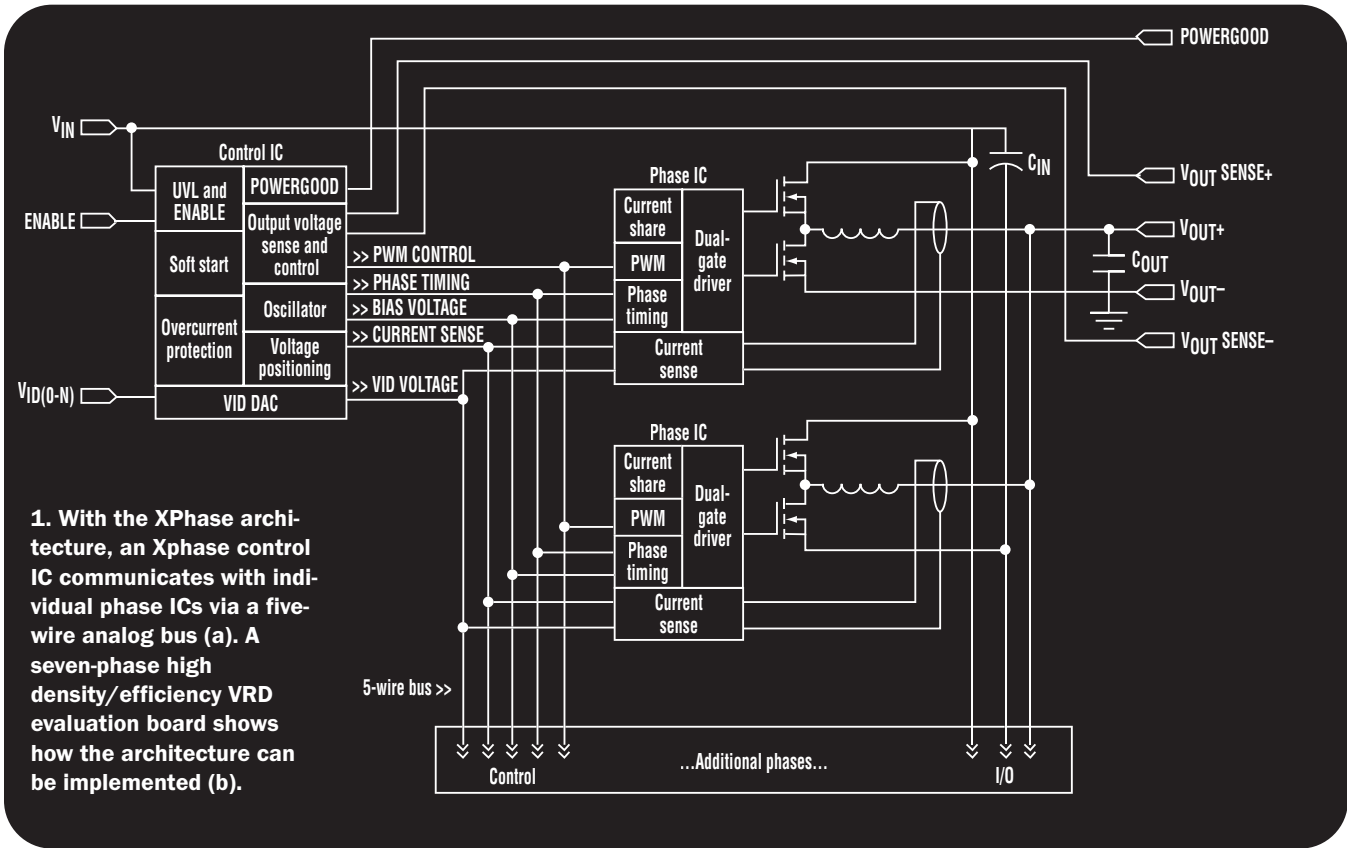
Our January 13 Technology Forecast was a monumental effort. Along with the great stuff in print, look for commentaries by our staff and industry pros online at www.elecdesign.com. I'd also like to thank the folks who spent hours briefing me: Linear Technology's Robert Dobkin; Analogic Tech's Richard Williams; ADI's Lew Counts, Steve Sockolov, and Eric Nolan; Intersil's Simon Dutton; Artesyn's Todd Hendrix; True Circuit's John Maneatis; IR's Steve Clemente; ON Semi's Jeff Olsen; Power-One's Dave Hage; Vicor's Patrizio Vinciarelli; and Virtual Silicon's John Ford.

Meanwhile, the annual Applied Power Electronics Conference and Exhibition (www.apec-conf.org) will hit Austin, Texas, March 6-10. Two highlights are worth noting.

A new, full-registration package lets you see everything—seminars, sessions, social events, and exhibit hall—and gets you a CD-ROM of the proceedings plus a copy of the seminar workbook. It costs \$50 less than the separate costs of individual passes to the seminars and sessions. Register by Feb. 5, and you'll save \$75 or \$100 on any of the packages. (IEEE lifers and students pay half the early registration rate at any time.)

APEC also will present the ever-popular MicroMouse contest, which premiered at the 1977 IEEE show in New York. The grand prize was presented by Claude Shannon, who amazed us all with the original MicroMouse, built at Bell Labs in 1949. It explored and remembered its way through a reconfigurable maze using a rack full of telephone switching relays.


ED Online 9479



or near the switch nodes, this can cause interference and prevent proper motherboard operation. For these reasons, the VRD designs considered in this article will avoid the use of paralleled MOSFETs.

Boosting the number of phases helps reduce stress in every component. Moreover, it's possible to use smaller and less-expensive components, reduce ripple current in I/O filters, and improve transient response. However, with thermal design currents typically exceeding 100 A, designs are already utilizing all four phases provided by commonly used multiphase control ICs possessing four MOSFETs per phase. Recently, multiphase controller manufacturers have introduced controllers capable of four-, five-, and six-phase operation. But these are based on existing architectures with a fixed number of phases that require point-to-point wiring between the phase circuitry and the controller.

The problem could be solved with a multiphase architecture that lets five, six, or even more phases be implemented while minimizing layout complexity. The IR3081 XPhase control IC



J3: BNC for output waveform

Bias inputs: 3.3 V 5 V 12 V GND

ATX12V input

GND Output connectors Vcc_core

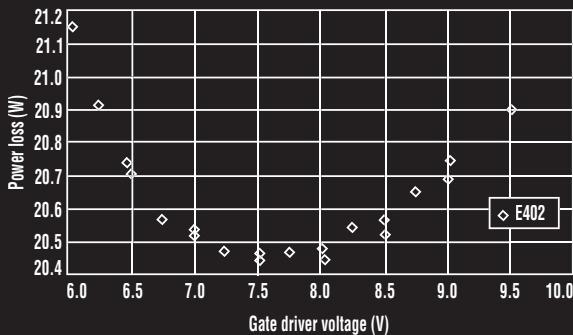
Vcc_core GND S1: VID PS_ON OutEnable Controller for dynamic VID S2: Dynamic VID Gate driver bias supply Test points: PowerGood OutEnable IShare bus VDAC PWM_Ramp EA_out

and IR3086 phase IC permit such an architecture by including all of the “one-per-converter” circuitry in the control IC and the “one-per-phase” circuitry in the phase IC. The control IC communicates with phase ICs via a novel five-wire analog bus that consists of bias voltage, phase timing, average current, error-amplifier output, and voltage-identification (VID) voltage. By eliminating the need for point-to-point wiring between the controller and the phase ICs, the five-wire bus shortens the interconnects to cut parasitic inductance and noise.

This architecture can support any number of phases, and it facilitates rapid design tradeoffs. Thus, designers can upgrade their designs if a higher-current CPU is introduced, or depopulate a phase or two if the current requirements drop. As will be discussed later, this can be done without changing the fundamental design.

High-density, low-profile, server VRD design

Many high-density server designs with two or more proces-



2. Here, a gate-drive voltage of 7.6 V produces the lowest power loss. Load current is 120 A.

sors and closely spaced VRDs require efficiencies in excess of 85% to reduce the generated heat. Height restrictions or the need for unrestricted airflow to the processors might prevent the use of a heatsink. To solve this problem, a seven-phase VRD demo board exhibiting extremely high density and efficiency was created.

The architecture used for this demo board allows for flexibility in designing multiphase, interleaved buck converters with 1 to X phases, while providing the ability to facilitate rapid design tradeoffs. It consists of an IR3081 control IC and a scalable array of phase blocks, each using an IR3086 single-phase (Fig. 1a). The control IC communicates with the phase ICs through the five-wire analog bus.

The control IC contains the entire one-per-converter circuitry, which includes VID, pulse-width modulation (PWM) ramp oscillator, error amplifier, bias voltage, fault detection, and other necessary functions to meet the load line, as well as the transient response requirements of present and future microprocessors. Each phase IC comprises all one-per-phase circuitry that includes gate drivers, PWM comparator and latch, overvoltage

protection, and current sensing and sharing.


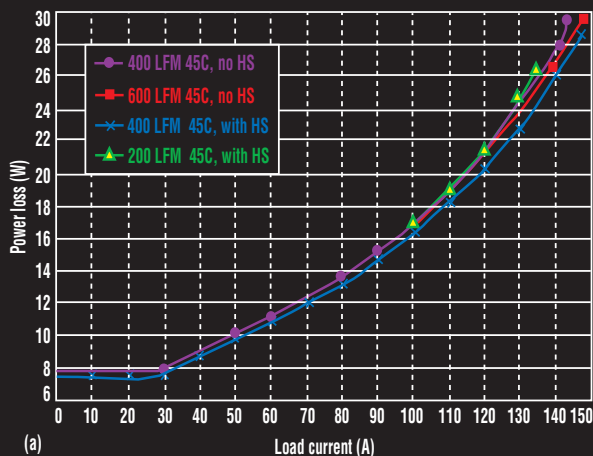
Figure 1b shows the demo board developed using the XPhase architecture and the IRF6617/IRF6691 DirectFET MOSFET pair. The IRF6617 features low gate charge (11 nC typical) for fast switching, while the IRF6691 provides an extremely low $R_{DS(ON)}$ (1.2 m Ω typical). To minimize switching losses, size, and cost, only one IRF6617/6691 pair was used per phase. The control IC is located to the right bottom corner of the board, far away from the heat and noise of the power train. One DirectFET pair per phase and small (10- by 7-mm) Pulse PA0511 220-nH inductors help shrink the power train to 2.5 by 0.95 in. This is considerably smaller than what's possible in four-phase designs that use four MOSFETs per phase and much larger inductors.

A constant 400-kHz switching frequency per phase achieves an equivalent ripple frequency of 2.8 MHz. This frequency was chosen as a tradeoff between overall converter size, transient performance, and efficiency (power loss). The Pulse inductors have a small footprint and low power loss due to their low dc resistance of 0.36 m Ω typical. But they also have a low 32.5-A saturation current that prevents using a lower switching frequency.


A lower switching frequency could yield higher efficiency. But it would require using a higher-value output inductor with a larger footprint, and it could possibly expand the size and space needed for input and output decoupling capacitors. A higher switching frequency could make for a more compact design, but it would increase switching losses and lower the efficiency.

Using seven phases with small inductor values improves the load transient performance, eliminates the need for bulky electrolytic capacitors, and greatly reduces the required PCB area. This design meets the transient response requirements for a load step of 100 A using both ceramic (60- by 10- μ F 1206 MLCC) and low-profile POSCAPs (8 by 2R5TPE470M9 470- μ F POSCAP with 9-m Ω equivalent series resistance).

A simple linear regulator with an npn pass transistor is used



3. These curves illustrate the seven-phase VRD's power loss (a) and efficiency (b) under various thermal environments.



4. At a temperature of 45°C and an air flow of 400 LFM, the seven-phase VRD's infrared thermal images show PCB temperatures both with (a) and without (b) a heatsink.

to supply a gate-drive bias of about 7.6 V from the 12-V input rail. Selecting a gate-drive voltage between the traditional 5- and 12-V supplies lets the designer optimize the gate drive for the lowest power losses (Fig. 2).

The seven-phase VRD was evaluated in a wind tunnel at 45°C and experiencing varying amounts of airflow, with and without the heatsink (Fig. 3). Power loss for the design was approximately 24 W, and efficiency is about 87% at 130 A. With the heatsink, the highest temperature on the VRD was about 83°C (Fig. 4). Alternatively, less airflow is required to cool the seven-phase design or a higher ambient temperature is possible. Without the heatsink, more airflow is required to maintain safe temperatures (Fig. 5). But the VRD height is reduced to a maximum of just 5 mm, which can be an important factor in space-constrained applications such as blade server designs, where height restrictions limit the use of heatsinks.

Five-phase design

It's possible to reduce the number of phases in applications that have space for additional bulk capacitors. However, power losses increase and a heatsink becomes necessary to remove heat from the board. Due to the flexibility of the multiphase architecture, the seven-phase design was easily converted to a five-phase design by simply depopulating components from two phases.

A finned heatsink, measuring 3.25 in. long, 0.5 in. wide, and 0.5 in. high, was attached to the PCB by four screws extending up from the backside of the board. An adhesive thermal tape (gap pad) was used between the heatsink and DirectFET devices as a thermal interface and an electrical insulator. Other attachment methods and alternate thermal interface materials can easily be accommodated, depending upon the designer's preference.

In this design, the gap pad material used was Berquist A3000 with a thermal conductivity of 2 W/mK. Larger inductors (Pulse PA0515, 225 µH, dc resistance = 0.63 mΩ) with a higher saturation current of 55 A had to be employed because of the higher per-phase current. Switching frequency was 400 kHz.


The five-phase design was evaluated in a wind tunnel. Performance data from the testing is presented in Figure 6. Efficiency of the five-phase design is a couple of points lower at 130 A, and it dissipates 27 W of power loss. This 12.5% higher power loss compared to the seven-phase design is all it takes to drive board temperatures up another 10 degrees. Because most server motherboards have at least two CPUs per board, this greater power loss necessitates a heatsink to help cool the PCB.

Besides the reduction in thermal performance, transient performance is also affected. This design would need 12 to 16 560-µF electrolytic capacitors, which would consume board space and increase costs. **ED**

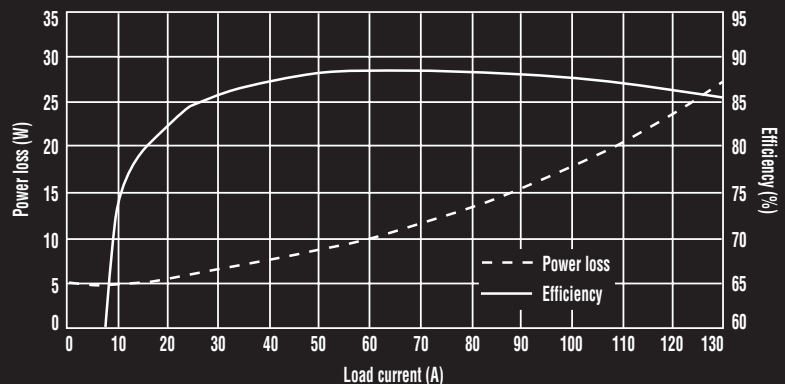
ED Online 9496

Ralph Monteiro is marketing manager, desktop & servers for International Rectifier (www.irf.com), El Segundo, Calif. He holds a [need degree received] from [need university, city, state]. Monteiro can be reached at [need e-mail].

George Schuellein is director, computing applications engineering for International Rectifier. He holds a [need degree received] from [need university, city, state]. Schuellein can be reached at [need e-mail].



5. This curve shows the maximum PCB temperature versus airflow. No heatsink is used. VID = 1.36 V, load current = 130 A, and the ambient temperature is 45°C.



6. A five-phase design exhibits these power-loss and efficiency curves. Air flow = 400 LFM, VID = 1.36 V, and the ambient temperature is 45°C.

Just One Microcontroller Pin Sets Programmable Timer's Interval

Ricardo Jimenez, Virgilia Moreno, and Mario Maciel

ricardo.jimenez@imperial.edu

Instituto Tecnológico de Mexicali (ITM), Mexicali, B.C. Mexico

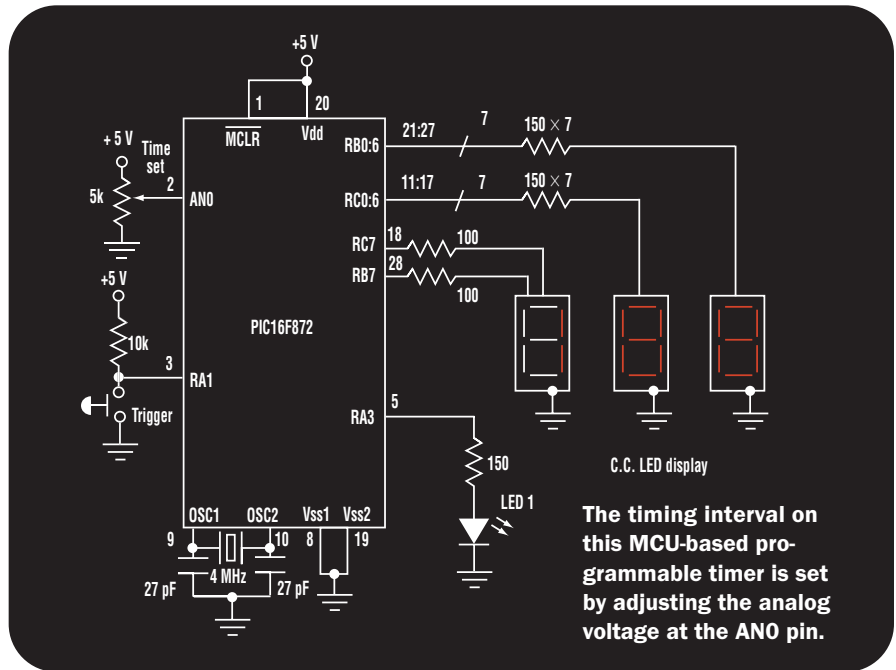
most digital timers require a set of dip switches or rotary BCD encoders to preset their interval. But those techniques consume eight or more inputs from an MCU's I/O lines. To alleviate the problem, this circuit forms a digital timer that only requires one input pin of the PIC16F872 MCU to adjust its interval (see the figure). The MCU performs all of the functions, analog-to-digital converting, timing, and decoding to drive the displays.

This circuit's timing interval ranges from 1 to 199 seconds with a resolution of 1 second. It can be adjusted with potentiometer R1, which generates an analog voltage to the PIC's AN0 input. The ADC reads the voltage and then generates the respective binary reading. Using lookup tables in the software routine, such binary readings are converted to BCD code and then translated to a seven-segment code. This code is required to drive the 2.5-digit LED digital display. The digital reading displayed now represents the timing interval in seconds.

Once the trigger is actuated by a negative-edge pulse generated by the pushbutton connected to RA1, output RA3 goes to a logic one, and the digital reading starts decrementing each second. When the digital reading reaches zero seconds, RA3 goes low and the display presents the original preset number. This way, the timer can be triggered again with the same interval if desired.

A trigger pulse on RA1 typically has a 50- μ s duration. The MCU reads this input every 73 μ s, which is sufficient to catch such a pulse, and updates the display every 73 μ s. Debouncing the pushbutton is performed through software because only the first high-to-low transition is read to start the timing routine.

A new timing interval can be obtained by adjusting R1 before pulsing the trigger input. Once the cycle



The timing interval on this MCU-based programmable timer is set by adjusting the analog voltage at the AN0 pin.

starts, the interval can't be changed. The MCU is constantly reading the AN0 voltage. Changing this voltage increases or decreases the reading.

This MCU's software routine, DB2090CODE.asm.doc, is online at www.elecdesign.com, Drill Deeper 9499. By changing the constants in the subroutine "DELAY," the MCU can be adjusted to supply pulses from 1 to 199 ms, minutes, or even hours.

When the input voltage is changed, it's first compared against the ADC reference voltages V_{REF+} and V_{REF-} , which are con-

nected to V_{DD} and V_{SS} , respectively. Register ADCON0 configures these reference voltages. Once the a-d conversion is processed in 20 μ s, the result is stored in register ADRESH. Then, the BCD to Binary conversion routine "BCDBIN" is executed. Finally, this result is converted to a seven-segment format to drive the displays. The table shows nine typical cases. If the user tries to set a timing interval that's out of range, the display will flash "Err" to indicate an error.

The output LED connected to RA3 can be substituted by an optocoupler or a transistor to drive a relay. For battery-powered circuits,

the output pulse from RA3 can be replaced by a 1-kHz frequency with a duty cycle of 50% to drive a coil's relay and save energy. Also, the C.C. LED displays can be substituted by an alphanumeric LCD readout, but an additional subroutine is required to handle such a device. Finally, for industrial applications, it's recommended to use a 4.00-MHz crystal with an extended temperature range. **ED Online 9498**

Ricardo Jimenez, electronics professor at Imperial Valley College, received a BSEE from the Technological Institute of Durango, Mexico.

Virgilia Moreno and **Mario Maciel** are in their fourth year of electronics engineering at ITM.

TIMING CONSTANT VS. INPUT VOLTAGE

ANO (V)	ADRESH	T _{SET} (seconds)
0.00	00000000	0
0.03	00000001	1
0.71	00100111	39
1.45	01001110	78
2.22	01110101	117
2.77	10011100	156
3.46	11000111	199
3.80 or higher	11001000	Error
5.00	—	Error