Design Considerations for 12-V/1.5-V, 50-A Voltage Regulator Modules

Yuri Panov and Milan M. Jovanovic

Delta Products Corporation
Power Electronics Laboratory
P.O. Box 12173
5101 Davis Drive
Research Triangle Park, NC 27709

Abstract - **The paper presents design considerations for a 12-V/1.5-V, 50-A voltage regulator module (VRM) for the next generation of microprocessors. The module has stringent power-density and transient-response specifications, which are hard to meet with traditional design techniques. The proposed design solutions increase the VRM efficiency, as well as achieve the desired transient response with a minimum amount of the output capacitance.**

I. INTRODUCTION

To decrease power consumption and increase the speed, the next generation of computer microprocessors will operate at significantly lower voltages and higher currents than today's generation. At the same time, these microprocessors will require a highly accurate supply voltage regulation which cannot be achieved by a centralized power system. A specified regulation accuracy can be accomplished with the distributed power system where a high-quality power is delivered to the microprocessor by a voltage regulator module (VRM), which is located on the motherboard next to the load. Generally, the VRM is required to have a high power density and to operate with a high efficiency. To meet these requirements and to provide a fast transient response, the power conversion must be performed at a high switching frequency, which presents a serious design challenge.

This paper deals with the design of a VRM supplying power from a 12-V tightly regulated bus to a 0.9-1.5 V, 50-A load which exhibits current transients with a slew rate of 50A/μs. For the present VRMs with the load current in the 15-A to 20-A range, the conventional buck topology with synchronous rectifier (SR), shown in Fig. 1, has been proven to represent a good performance/cost trade-off. However, if a single buck-converter topology were employed in the 12-V/1.5-V, 50-A VRM, then, to achieve the specified load transient response, a large amount of the output-filter and on-board decoupling capacitance would be required [1]-[4]. The size of the VRM would increase as well as the required space on the motherboard, making the conventional single-module buck-converter topology not practical.

The amount of required output-filter and decoupling capacitances can be minimized by employing the interleaving technique, as demonstrated in [2]. Generally, the interleaving technique is implemented by paralleling a number of converter modules, and by phase-shifting (interleaving) their drive signals. The main benefit of the interleaving is the increased output ripple frequency which is equal to the product of the single-module switching frequency and the number of the interleaved modules. Specifically, the increased output-ripple frequency makes possible to reduce the output-filter capacitance, as well as to increase the control-loop bandwidth to improve the transient response. Additional benefits of interleaving include better thermal management and packaging flexibility.

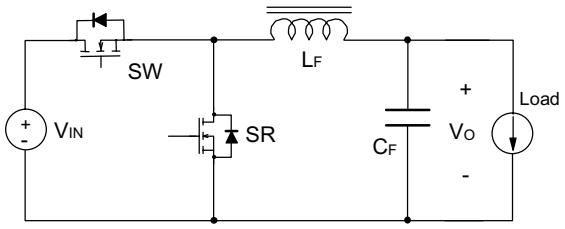

Fig. 1. Buck converter with synchronous rectifier.

Recognizing that the interleaving approach is currently the only viable approach in low-voltage, high-current applications with highly dynamic loads, a number of IC manufacturers have introduced dedicated controllers for interleaved VRMs. Generally, these multiphase controllers offer different number of phases and switching frequency ranges, as well as different integration levels of control, drive, and power components. Also, to achieve a uniform current distribution among the interleaved modules, some of the controllers employ active current-sharing techniques, whereas the others try to provide identical duty cycles for each module, and rely on the identical layout of the modules, as well as on the tolerances of the components and circuit delays.

Due to extremely challenging requirements, the design of the next generation of VRMs requires a thorough undertanding of the performance and design trade-offs. The objective of this paper is to discuss these trade-offs, and to propose some design solutions for optimizing the performance of the 12-V/1.5-V, 50-A VRMs.

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

II. Power Stage Design Considerations

B. Trade-Off Between Efficiency and Transient Response

One of the most important issues of the VRM power-stage design is the selection of the output LC-filter parameters. Generally, for VRMs this selection is based not on the output-voltage ripple spec, but on the trade-off between the specified VRM efficiency and transient response.

The minimum capacitance which is required to keep the transient output voltage V_O within the regulation limits, can be estimated using the approach presented in [1]. Assuming that the VRM control responds immediately to the load change, i.e., assuming that the control-loop bandwidth is infinite, the buck converter equivalent circuits during the load step-up and step-down transients are shown in Figs. 2(a) and 2(b), respectively. From Figs. 2(a) and 2(b), the rate of the inductor current change is

$$\frac{di_L}{dt} = \frac{V_{IN} - V_O}{L_F} \text{ during step-up transient,} \qquad (1(a))$$

and

$$\frac{di_L}{dt} = -\frac{V_O}{L_F} \text{ during step-down transient.} \qquad (1(b))$$

According to Eqs. (1(a)) and (1(b)), for a 12-V/1.5-V VRM, the rate of inductor current change is much higher during a step-up than during a step-down transient because input voltage V_{IN} is much higher than output voltage V_O. Therefore, the output-voltage overshoot during a load step-down transient sets the limit on the VRM transient performance.

To keep VRM output voltage V_O within regulation spec ΔV_O during a load transient of magnitude ΔI_{OMAX}, the minimum required output-filter capacitance is

$$C_{FMIN} = \frac{1}{2} \cdot \frac{\Delta I_{O\,MAX}^2}{\Delta V_O} \cdot \left(\frac{L_F}{V_O} - \frac{1}{di_O/dt} \right), \qquad (2)$$

where di_O/dt is the load-current slew rate.

According to Eq. (2), output-filter inductance L_F has to be minimized to achieve a fast transient response with the minimum output capacitance. However, a low inductance value increases the inductor current ripple, which has a detrimental effect on the VRM efficiency. Not only increased inductor-current ripple increases conduction losses due to the increased rms currents, but more importantly it dramatically increases the buck switch turn-off loss due to the increased peak value of the inductor current. The detrimental effect of a high inductor-current ripple on the VRM efficiency is illustrated in Fig. 3 which shows the measured efficiency of a 12-V/1.5-V, 20-A single-module VRM as a function of the switching frequency for different values of the output-filter inductance. As can be seen in Fig. 3, at any switching frequency the VRM efficiency decreases as the output-filter inductance decreases from 470 nH to 160 nH. Generally, the efficiency drop is more pronounced at lower switching frequencies, i.e., below 300 kHz. Specifically, at f_S = 200 kHz the efficiency drop is 5.5% when the inductance is reduced from 470 nH to 250 nH, whereas the efficiency drop when inductance is reduced from 250 nH to 160 nH is around 10%. However, at f_S = 700 kHz, for example, the corresponding efficiency drops are only 4% and 2%.

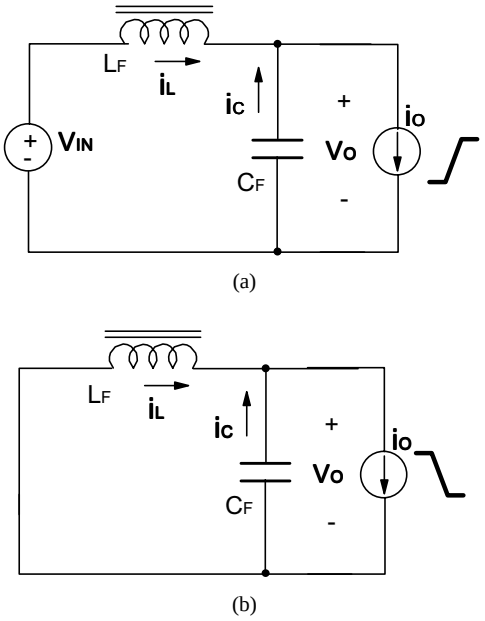

Fig. 2. Equivalent circuit of VRM with ideal control (i.e., with infinite control loop bandwidth) during load transients: (a) load step-up transient; (b) load step-down transient.

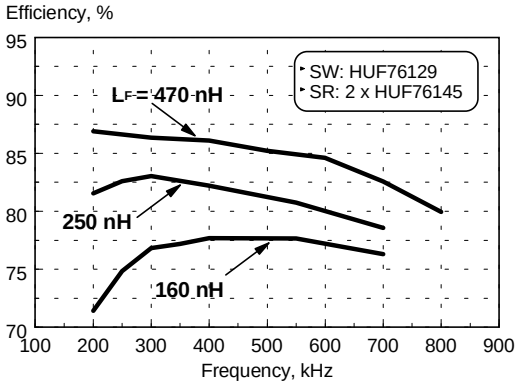

Fig. 3. Measured 12-V/1.5-V, 20-A VRM efficiency as function of switching frequency for different values of output-filter inductance.

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

Figure 3 also shows that for a given output-filter inductance value there is an optimal switching frequency at which the VRM efficiency is maximized. For a large output-filter inductance, e.g., L_F = 470 nH, the efficiency monotonically increases as the switching frequency decreases. The improvement of the efficiency at lower frequency is caused by reduced switching losses, in particular, the turn-off switching loss of the buck switch. However, for lower values of the output-filter inductance, i.e., for L_F = 250 nH and L_F = 160 nH, the maximum efficiency does not occur at the minimum switching frequency. In fact, for L_F = 250 nH, the maximum efficiency occurs at f_S=300 kHz, whereas for L_F=160 nH, the optimal switching frequency is in the 400-550 kHz range. The efficiency decrease at low frequencies for low output-inductance values is caused by the increased turn-off switching loss of the buck switch because of the increased peak inductor current.

The reduction of the output-filter inductance without penalizing the conversion efficiency can be achieved by employing the interleaving approach. Since for an interleaved converter the output-filter inductors of the individual modules are effectively connected in parallel, the transient response of the interleaved converter is governed by effective inductance

$$L_{F(EFF)} = L_F/N, \quad (3)$$

where N is a number of interleaved modules. Consequently, in an interleaved converter, the desired transient response can be achieved with a smaller output-filter capacitance than in a single-module converter [2].

Optimization of the VRM efficiency and transient performance requires careful selection of the switching devices, switching frequency, output-filter components, and number of interleaved modules. Selection of the switch and SR devices is driven by their operating conditions. Since SR conducts the inductor current for the most of the switching cycle, its conduction loss is considerably higher than that of the switch. At the same time, the SR switching loss is minimal because, by proper selection of delays between the switch and SR gating signals, SR can be turned on and off with zero voltage across it. Therefore, it is desirable to select the device with the lowest on-resistance for SR. However, for the buck switch, it is crucial to select the device with the lowest turn-off loss which is determined by many parameters such as the device fall time, gate charge, internal gate resistance, and package parasitic inductance [5].

After the switching devices are chosen, the next design step is to select inductance L_F and the module switching frequency which correspond to the specified VRM efficiency. This design optimization can be most efficiently performed empirically, using a single-module (one-phase) prototype circuit. The data similar to that shown in Fig. 3 is helpful in determination of the optimal switching frequency.

The final design step is to estimate the minimum amount of the output capacitance which satisfies the transient spec. If the estimated C_F value is unacceptable, the effective inductance $L_{F(EFF)}$ has to be decreased by increasing the number of interleaved modules.

B. Driving Loss Optimization

As it was mentioned in the previous section, the SR device in the 12-V/1.5-V, 50-A VRM must have a very small on-resistance. When the required on-resistance cannot be obtained by selecting a proper device, SR is implemented by connecting several devices in parallel. In any case, the SR gate capacitance is large, and it causes a significant driving loss at high switching frequencies. As an illustration, Fig. 4 shows the measured efficiencies of the 12-V/1.5-V, 20-A single-module VRM without and with the control losses included. As can be seen from Fig. 4, at 20-A load current the efficiency drops by 6 % because of the control loss which for 70-80 % consists of the buck switch and SR driving loss.

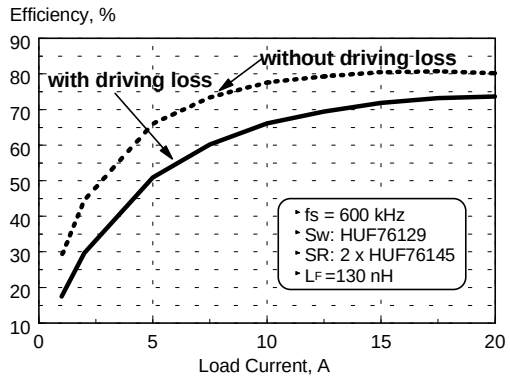

Fig. 4. Measured 12-V/1.5-V, 20-A VRM efficiency as function of load current with and without control loss included.

Generally, the conventional, "hard-switched" MOSFET driver, shown in Fig. 5, dissipates each switching cycle twice the energy necessary to charge the MOSFET gate capacitance [6]. The driving loss can be reduced by replacing the conventional drive with a resonant drive, which recycles the energy stored in the device gate-source capacitance. The maximum possible efficiency of the resonant drive is limited by the MOSFET internal gate resistance. For an ideal MOSFET with zero gate resistance, the resonant drive can theoretically operate with 100% efficiency. However, for the majority of today's low on-resistance MOSFET devices, the internal gate resistance limits the maximum efficiency of the resonant drive to 50-85% [6].

Figure 6 shows the circuit diagram of the buck converter with a resonant drive. According to Fig. 6(a), when switch SW is turned on, capacitor C1 is charged to voltage V_{IN}, as shown in Fig. 6(b). When switch SW turns off at $t = t_0$, switch S1 is turned on, and capacitor C1 resonantly dischar-

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

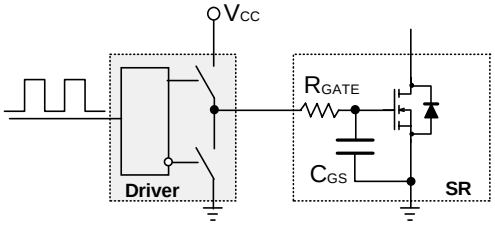

Fig. 5. Conventional MOSFET driver.

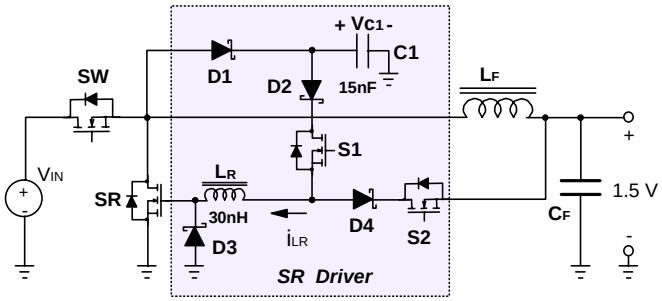

(a)

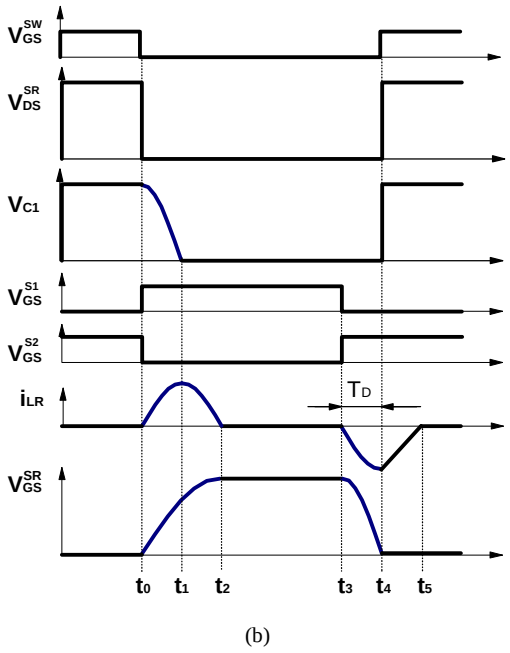

(b)

Fig. 6. Resonant SR driver: (a) circuit diagram; (b) key waveforms.

ges into the SR gate capacitance through diode D2, switch S1, and resonant inductor L_R, turning on SR. After voltage V_{C1} across C1 reaches zero at $t = t_1$, diode D1 starts conducting. During $[t_1\text{-}t_2]$ interval, inductor current i_{LR} decreases to zero, whereas the SR gate capacitance continues to charge. At $t = t_3$, which occurs before the turn-on of switch SW at $t = t_4$, switch S1 is turned off and switch S2 is turned on. As a result, the SR gate capacitance discharges in a resonant fashion into the output through inductor L_R, diode D4, and switch S2, turning off SR. The discharge of the SR gate capacitance ends at $t = t_4$, when diode D3 starts conducting. During the $[t_4\text{-}t_5]$ interval, current i_{LR} decays linearly to zero. As can be seen from waveforms in Fig. 6(b), both driver switches S1 and S2 turn off at zero current, that greatly reduces their switching losses. In addition, capacitor C1 serves as a lossless snubber which limits the voltage overshoot across switch SW and SR after the turn-on of switch SW.

Experimental waveforms of the resonant drive are shown in Fig. 7. It should be noted that the measured V_{C1} waveform differs from the ideal waveform in Fig. 6(b). Namely, because of the power-stage parasitics, capacitor C1 charges to 20 V, instead of to the input voltage, whereas, because of the SR internal gate resistance and losses in the driver components, C1 discharges to 6 V, instead of to zero.

Measured efficiencies of the VRM with the conventional and proposed resonant SR driver are shown in Fig. 8. As can be seen from Fig. 8, the resonant drive provides an efficiency gain of 1.5% at the full load of 20 A. This VRM efficiency gain corresponds approximately to the 40% reduction of the driving loss. The further driving loss reduction is limited by the loss in the internal gate resistance of the SR, as well as by the conduction losses in the driver components.

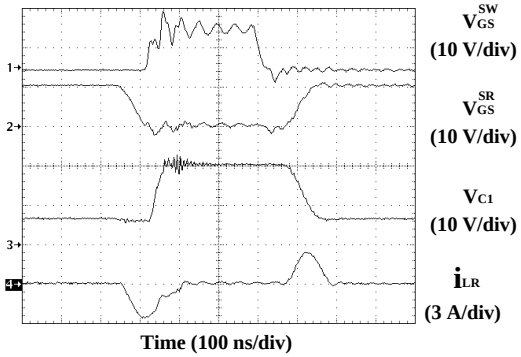

Fig. 7. Experimental waveforms of resonant SR driver.

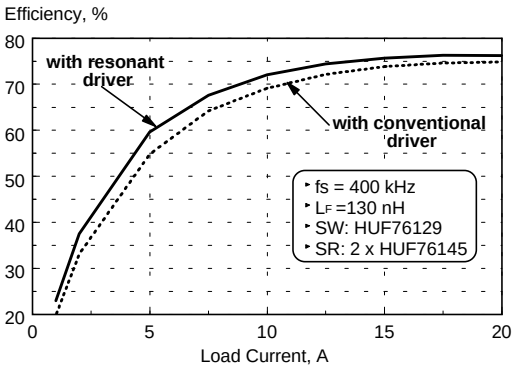

Fig. 8. Measured 12-V/1.5-V, 20-A VRM efficiency with conventional and resonant SR drive.

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

III. CONTROL DESIGN CONSIDERATIONS

A. Limitations of Conventional VRM Control

Generally, interleaved VRMs require a high-performance feedback control that can provide a low overshoot of the output voltage during load transients, as well as even current sharing among the interleaved modules. A general block diagram of the voltage-feedback interleaved VRM control is given in Fig. 9. The block diagram of the conventional implementation of the PWM and phase-shift circuitry along with the key waveforms is shown in Fig. 10 for two interleaved converters, but it can be easily extended to a larger number of modules.

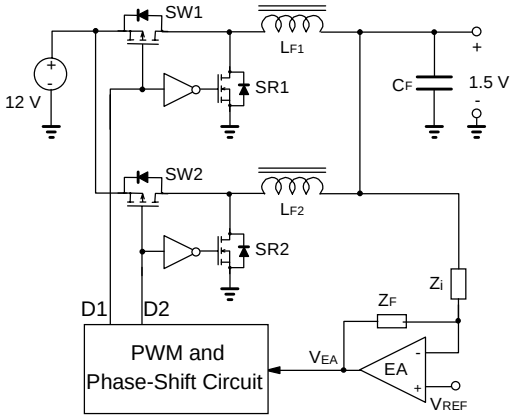

Fig. 9. Block diagram of voltage-feedback interleaved VRM control.

Generally, the major drawback of the conventional control, shown in Fig. 10, is related to the load-current distribution among interleaved modules. If the interleaved modules had identical layouts and their duty cycles were tightly matched, an acceptable current sharing among the modules would be accomplished without an active current-sharing control. The tight matching of the duty cycles requires phase-shifted ramp signals V_{RAMP1} and V_{RAMP2}, shown in Fig. 10(b), be tightly matched. However, the accurate matching of the ramps is very difficult to accomplish. With today's integrated-circuit technology, the duty-cycle matching within ≈ 1 % can be accomplished. Any further improvement in the matching accuracy would require additional design and manufacturing steps which substantially increase the controller cost. With the 1% duty-cycle matching accuracy, the current-sharing error cannot be reduced below 10-20%. Furthermore, the current-sharing accuracy degrades as the switches with a lower on-resistance are used.

Another drawback of the conventional control is associated with the employment of RS latches, shown in Fig. 10(a). As can be seen from Fig. 10, clock signal C1 sets the RS Latch #1 and turns on switch SW1 at instant t = 0.

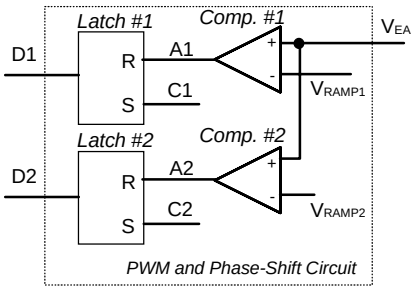

(a)

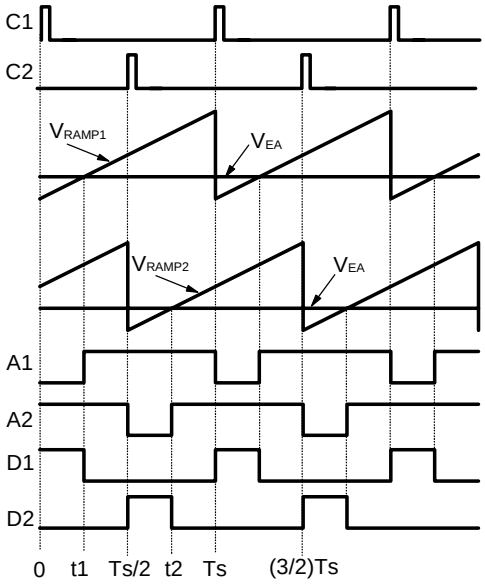

(b)

Fig. 10. Conventional implementation of interleaved VRM control: (a) simplified block diagram of PWM and phase-shift circuit; (b) key waveforms.

Latch #1 is reset and switch SW1 is turned off at instant $t = t_1$, when ramp voltage V_{RAMP1} becomes larger than output voltage V_{EA} of error amplifier EA. After $t = t_1$, RS Latch #1 prevents the turn-on of SW1 until the next clock signal at $t = T_S$. Similarly, after $t = t_2$, RS Latch #2 prevents the turn-on of SW2 until the next clock signal at $t = (3/2)\cdot T_S$. Because of the presence of the latches, which delay the turn-on of the switches until the next clock signal, the output voltage may experience a large negative overshoot in the case of the load step-up, as illustrated in Fig. 11. For example, if the load current is increased at $t = t_{1+}$, i.e., immediately after the turn-off of switch SW1, the desired control response is to immediately turn on both switches SW1 and SW2. However, due to the RS latches, switch SW1 cannot be turned on earlier than at $t = T_S$, and switch SW2 cannot be turned on earlier than at $t = T_S/2$. As a result, the output capacitor C_F needs to support the increased load current for a longer time than in the case when controller is implemented without the latches.

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

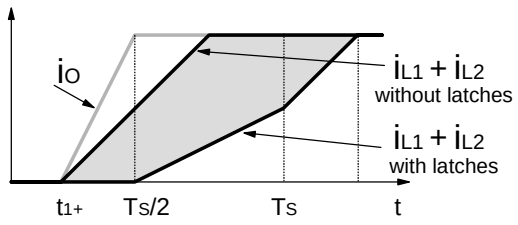

Fig. 11. Effect of RS latches on VRM transient response to load step-up. Shaded area is proportional to excessive charge drawn from capacitor C_F.

Therefore, due to the presence of the latches, a larger output capacitor is required. The detrimental effect of the RS latches is illustrated in Fig. 11 by the shaded area which is proportional to the excessive charge drawn from capacitor C_F.

The fast response to a load disturbance requires a wide bandwidth of the feedback loop. However, as the control bandwidth increases, the task of maintaining VRM stability under all operating conditions becomes progressively harder, and the noise immunity of the control suffers as well.

These drawbacks of the conventional VRM control can be mitigated with the control scheme which is presented in the next section.

B. Proposed Interleaved VRM Control

The VRM transient response can be improved without sacrificing stability by keeping the control loop gain low during steady-state operation, and by increasing the gain during load transients. As the PWM gain is inversely proportional to the slope of the ramp signal, the variable-gain approach can be implemented by replacing the conventional constant-slope ramp signal with the variable-slope ramp signal, as shown in Fig. 12 for the case of a single module. In Fig. 12(b), the PWM gain is low during steady-state operation to maintain the VRM stability. When the load step-up occurs at $t = T_S$, duty ratio for the next cycle jumps to unity. As can be seen from Fig. 12(a), with the constant-slope ramp, it takes more than one switching cycle for the output-filter inductor current to reach the new level of the load current. With the variable-gain modulator, the inductor current reaches the new level in one cycle, as illustrated in Fig. 12(b), thus, reducing the output-voltage overshoot.

The variable-slope ramp considerably changes the PWM input/output characteristic. In the constant-gain conventional control, as V_{EA} increases, duty ratio d increases linearly until it reaches unity, as shown in Fig. 13(a). In the variable-gain control, the duty ratio change is identical to that of the constant-gain control for $V_{EA} < V_{TH}$. However, the duty ratio changes abruptly to unity at $V_{EA} = V_{TH}$, as shown in Fig. 13(b). As can be seen from Fig. 13(b), at $d=D_{TH}$, the modulator incremental gain becomes infinite, which helps to improve the transient response.

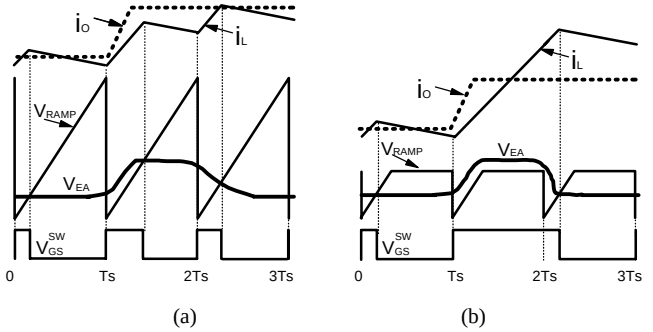

Fig. 12. VRM control transient response: (a) with constant-slope ramp; (b) with variable-slope ramp.

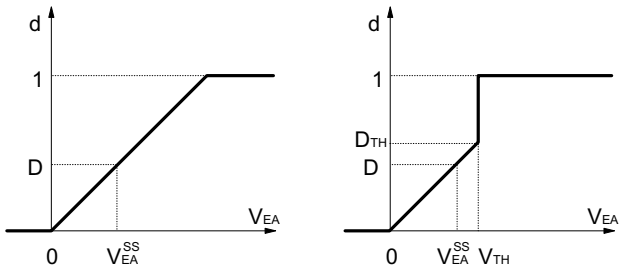

Fig. 13. PWM input/output characteristic: (a) with constant-slope ramp; (b) with variable-slope ramp.

The variable-slope ramp approach for interleaved modules is implemented as shown in Fig. 14. Signals A1 and A2 in Fig. 14 are the output signals of the phase-shift circuit.
During steady state, pulse-width modulation is performed by Comparator #1. Pulse train B at the output of Comparator #1 is then distributed between switches SW1 and SW2 by gates G1, G2. If during a load step-up EA output voltage V_{EA} exceeds threshold level V_{TH}, Comparator #2 turns on the switches of all modules to accelerate the response. Although shown for two modules, the proposed circuitry can be easily modified for a larger number of modules. Since a single ramp is used to generate gate-drive signals for several modules, the maximum number of interleaved modules is limited by the relationship $D < D_{TH} < 1/N$. Therefore, the proposed control scheme limits the number of 12-V/1.5-V interleaved modules to approximately five modules that is sufficient for most applications.

The important feature of the control circuit in Fig. 14 is the absence of the RS latches. Generally, RS latches enhance noise immunity of control by preventing multiple switchings during steady-state operation. Without the latches, multiple switchings may occur for a few cycles during severe load transients. This is usually acceptable as far as the VRM steady-state operation is not affected.

Finally, it should be noted that the control in Fig. 14 also improves current sharing among the interleaved modules since it uses the same ramp for all modules and, therefore,

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

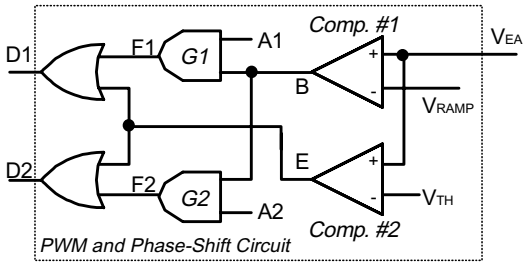

(a)

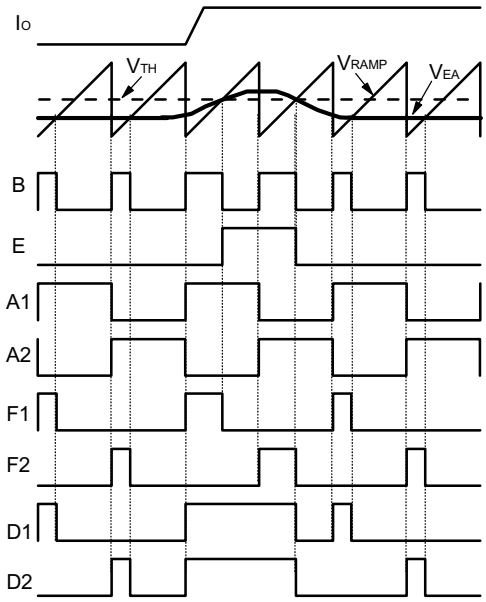

(b)

Fig. 14. Proposed implementation of interleaved VRM control: (a) simplified block diagram of PWM and phase-shift circuit; (b) key waveforms.

eliminates the main source of duty-ratio mismatch in the conventional control. Nevertheless, to preserve a good current sharing, the layout and drive-circuitry delays must also be matched.

III. Design Evaluation

An experimental, 12-V/1.5-V, 50-A VRM prototype was built with the described variable-gain control. The prototype was implemented with three interleaved modules, each operating at 400 kHz. The following major components of the VRM power stage were selected: switch SW - 2 x IRF7811, SR - 4 x IRF7811, and L_F = 500 nH. The core of inductor L_F is a combination of E14/3.5/5-3F3 E-core with PLT14/5/1.5 plate. The inductor winding has three turns of 175/40 Litz wire. The output capacitor bank consists of twenty one 220 μF-2.5 V POSCAP capacitors, twenty seven 33 μF-25 V ceramic capacitors, and nine 1.5 mF-4 V tantalum capacitors.

The measured VRM efficiency as a function of the load current is shown in Fig. 15. At the 50-A load the VRM efficiency is 81 %, whereas at the 55-A load it drops to 80.2 %. It was found that in the experimental prototype, at these high current levels, the layout-related conduction loss contributes significantly to the total VRM loss. Namely, the VRM prototype was built on a PCB with 2-mil thick copper. With a thicker copper layer, the VRM efficiency is expected to increase by ≈2 %.

The accuracy of current sharing among interleaved modules is shown in Fig. 16, where the relative current-sharing error is plotted as a function of the load current. The relative current-sharing error for a k-th module is defined as $I_{Lk}/\left[\left(I_{L1}+I_{L2}+I_{L3}\right)/3\right]-1$, where k = 1,2,3, and I_{L1}, I_{L2}, and I_{L3} are average inductor currents of individual modules. As can be seen from Fig. 16, the current sharing error is within ±5 % for load currents above 20 A.

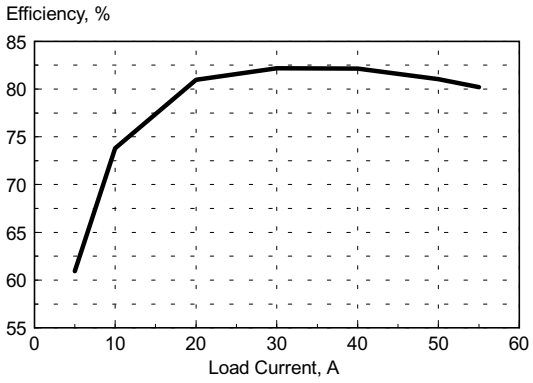

Fig. 15. Measured efficiency of three-module interleaved 12-V/1.5-V VRM prototype.

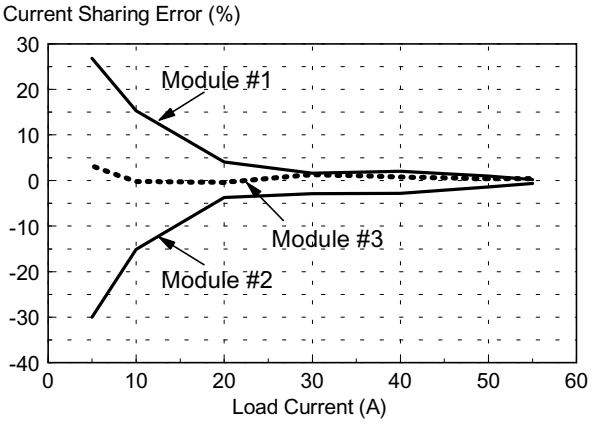

Fig. 16. Measured current distribution among three interleaved modules of 12-V/1.5-V VRM prototype.

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

The measured VRM transient responses to a 50-A load step-up for the constant-slope ramp and variable-slope ramp are shown in Figs. 17(a) and 17(b), respectively. In both cases the bandwidth of the control loop was 30 kHz, and it was limited by stability considerations. As can be seen from Figs. 17(a) and 17(b), for both constant-ramp and variable-ramp controls, it takes 1-2 switching cycles for inductor currents to start rising in response to the load step. For the VRM with the variable-slope ramp, the initial output-voltage drop causes voltage V_{EA} to exceed the threshold level V_{TH}. As a result, the switches of all modules turn on simultaneously at instant t_1, and the sum of modules' inductor currents increases at the maximum rate, thus, reducing the output-voltage deviation from the steady-state value. However, for the VRM with the constant-slope ramp, the switch of only one module is ″on″ during any given time, resulting in a slower transient response. As can be seen from Figs. 17(a) and 17(b), the employment of the variable-slope ramp reduces the output-voltage overshoot by 16-17 %. It also should be noted that in Fig. 17(b) the i_{L1} waveform, shown within the encircled area, indicates multiple switching during one switching period due to the absence of the RS latches in the controller. This switching had no detrimental effect on the overall VRM performance.

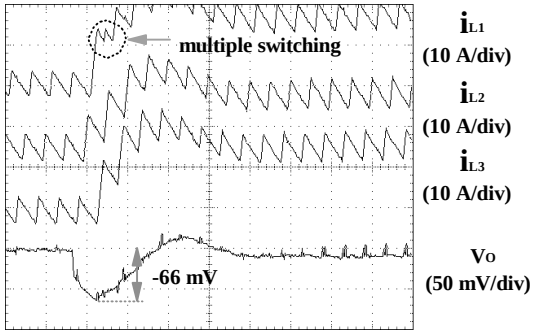

(a)

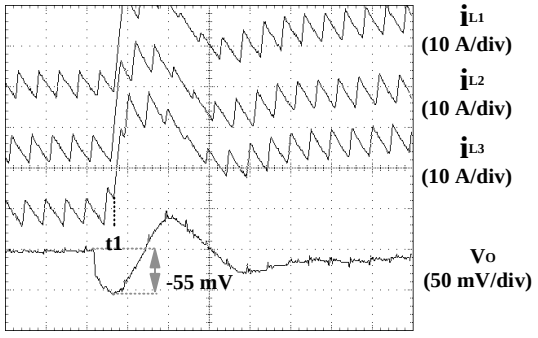

(b)

Fig. 17. Measured VRM transient response to 50-A load step-up: (a) with constant-slope ramp; (b) with variable-slope ramp.

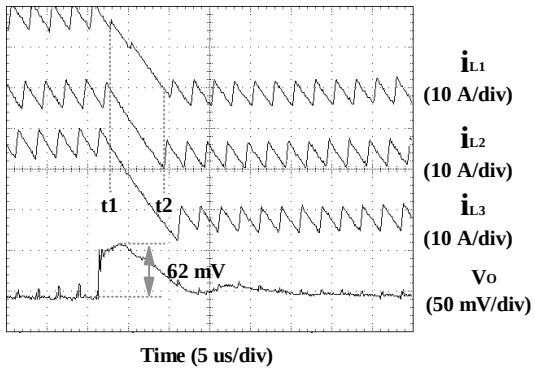

Fig. 18. Measured VRM transient response to 50-A load step-down.

Finally, the VRM response to a 50-A load step-down is shown in Fig. 18. The decay of inductor currents starts within one switching period after the load step-down. During $[t_1, t_2]$ time interval, all three modules operate with zero or minimum duty ratio that helps to reduce the output-voltage overshoot. As can be seen from Fig. 18, the V_O overshoot has approximately the same magnitude as the overshoot during the load step-up transient.

IV. SUMMARY

Design considerations for the interleaved 12-V/1.5-V, 50-A VRM were presented. The VRM power-stage design which can meet the specified efficiency and transient requirements was discussed. The limits of the conventional voltage-mode control in interleaved VRM applications were demonstrated. In order to overcome those limits, a simple control scheme, which improves the transient performance as well as the current distribution among the interleaved modules, was proposed.

REFERENCES

[1] M. Zhang, M. Jovanovic, F. C. Lee, "Design Considerations for Low-Voltage On-Board Dc/Dc Modules for Next Generations of Data Processing Circuits," *IEEE Trans. on Power Electronics*, vol. 11, no. 2, Mar. 1996, pp. 328-337.

[2] X. Zhou, et al., "Investigation of Candidate VRM Topologies for Future Microprocessors," *IEEE Applied Power Electronics Conf. Proc.*, pp. 145-150, Feb. 1998.

[3] P. L. Wong, et al., "VRM Transient Study and Output Filter Design for Future Processors," *Virginia Power Electronics Center Seminar Proc.*, pp. 1-7, Sep. 1998.

[4] A. Rozman, K. Fellhoelter, "Circuit Considerations for Fast, Sensitive, Low-Voltage Loads in a Distributed Power Systems," *IEEE Applied Power Electronics Conf. Proc.,* pp. 33-42, Mar. 1995.

[5] L. Spaziani, "A Study of MOSFET Performance in Processor-Targeted Buck and Synchronous-Rectifier Buck Converters," *High Frequency Power Conversion Conf. Proc*, pp. 123-137, Sep. 1996.

[6] "Investigation of Power Management Issues for Next Generation Microprocessors," *VRM Consortium Quarterly Progress Report*, Center for Power Electronics Systems, Virginia Tech, Sep. 1999.

0-7803-5867-8/00/$10.00 (c) 2000 IEEE

View publication stats