POWER MANAGEMENT SOLUTIONS FOR DESKTOP AND MOBILE CPUS

Richard Redl
Analog Devices, Inc.
Montévaux 14
1726 Farvagny, Switzerland
richard.redl@analog.com

and

Gabor Reizik
Analog Devices, Inc.
3550 North First Street
San Jose, CA 95134, U.S.A.
gabor.reizik@analog.com

Abstract

This paper focuses on the power management solutions for desktop and mobile CPUs, with special emphasis on the control aspects and the related control IC requirements and architectures for VRM applications. The paper discusses the following topics: power converter topologies, load transient response, the ADOPT™ voltage positioning technology, maintaining high efficiency over a wide load-current range, current and thermal balancing, current-sense solutions, protection functions, functional block schematics of VRM control ICs, and challenges for the IC designers.

1. Introduction

Designing a power supply (often called VRM for Voltage Regulator Module) for CPUs (Central Processor Units, or microprocessors) for desktop and mobile applications is a challenging task. The microprocessors are highly dynamic loads, switching between the minimum and maximum currents with a di/dt that can exceed 450A/μs at the processor socket [1]. The supply voltages (at the time of writing this article) typically fall in the 0.8 to 1.6-V range, determined by a five-to-eight bit VID (Voltage Identification) code [2] coming from the microprocessor. (Note: This means that the VRM controller has to include a five-to-eight bit DAC.) Dynamically programmed lower voltages down to about 0.5 V also begin to appear, especially in battery-powered mobile applications, where energy conservation is extremely important. Depending on the processor type, the current consumption in high-current active mode varies over a wide range, and can exceed 100 A for high-performance microprocessors. Further design difficulties arise from the requirement that the CPU supply voltage must be a tightly controlled linear function of the load current. This means that the power supply must have a specified resistive output impedance, from dc to several hundred kHz. Additional difficult-to-meet requirements include high voltage step-down ratio, high efficiency down to a fraction of a percentage of the full load (especially in mobile applications), differential remote sense, processor power sequencing, fast on-the-fly VID programming, and protection against sustained overload, overtemperature, output overvoltage, and output reverse voltage. It also goes without saying that the cost and size of the power supply should be minimized.

J. H. Huijsing et al. (eds.), Analog Circuit Design, 201-224.

This paper presents an overview of the most widely employed power-converter topologies for CPU VRM applications, discusses the most important control and protection issues with relevance on the architecture of the PWM control ICs, and presents the functional block diagrams of two example control ICs for microprocessor VRMs.

2. Power-converter topologies

Most CPU power supplies are based on the synchronous buck converter (Fig. 1). That circuit is capable of efficiently converting the voltage of the primary power source to the low output voltage demanded by the microprocessor, with the simplest possible topology. (The primary power source is either the +12V or +5V output of the "silver box" – the line power supply in desktop computers –, or the battery in mobile applications. The battery voltage is typically between 8 and 19V including battery charger operation.)

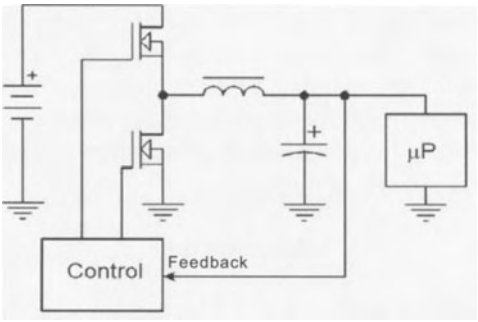

Fig. 1. The synchronous buck converter.

The basic synchronous buck converter has several limitations:

- It has a large input ripple current – this increases the cost of the input filter.
- Due to the limited rate of change of current in the buck inductor, the load transient response will be sluggish – this increases the cost of the output filter capacitor.
- The output current capability is limited by device availability.
- The converter has a high hot-spot temperature due to the concentrated heat dissipation.

The output current capability can be increased and the hot-spot temperature can be reduced either by paralleling power switches or by paralleling identical converters. The load transient response can be improved by reducing the inductance of the buck inductor, but that will lead to increased conduction losses and increased turn-off losses of the upper (control) FET. The conduction-loss increase caused by reduced inductance can be prevented by increasing the switching frequency, but that will lead to an increase in the switching losses. There is, however, a technique that can overcome all above limitations of the basic synchronous buck converter. That

technique combines the idea of paralleling converters with an even distribution of the switching instants of the individual converters over the switching period. This concept is called multiphase interleaved conversion, and it has been around since at least the early 70s [3]. Fig. 2 shows the topology of a three-phase interleaved synchronous buck converter.

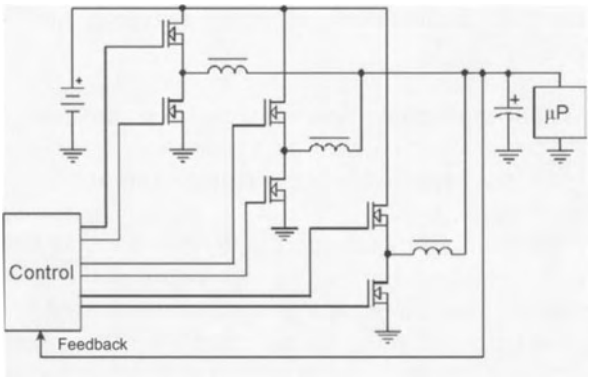

Fig. 2. The three-phase interleaved synchronous buck converter.

Perhaps the most important advantage of the multiphase converter is that the equivalent switching frequency increases in proportion to the number of phases, without increasing the switching losses. The higher equivalent switching frequency means that the equivalent inductance can be decreased without an efficiency penalty. Since the inductance is the major limiting factor for a fast load transient response, a multiphase converter is expected to produce a better response than its single-phase counterpart operating with the same switching frequency.

Unfortunately, using a multiphase topology is not without its drawbacks. Those drawbacks are:

- More FETs, drivers, and output inductors are needed. (Note, however, that the total FET size and inductor volume remain about the same.)
- The control circuitry becomes more complex.
- With voltage-mode control the current sharing among the phases is not guaranteed.
- With current-mode control, more than one current sensor might be required.

Although at present the single-phase or multiphase synchronous buck converter is by far the most popular choice for desktop and notebook computers, in multiprocessor server applications, where the total processor current can exceed several hundred amperes, often transformer-coupled converters are used. The main motivation for transformer coupling is to reduce distribution losses and parasitic voltage drops. The transformer coupling allows the input voltage for the converters to be raised to a higher level (typically 48V). Stepping down such a high voltage to the low voltage

required by the microprocessor with a buck converter is impractical due to the extremely narrow duty ratio and the related high switching and conduction losses. Ref. [4] presents a critical evaluation of six transformer-coupled topologies for such applications. Other recommendations for implementing high step-down ratios include the usage of high-frequency ac power distribution [5], multiphase tapped-inductor buck converters with active clamp [6], and cascade connection of the transformer-coupled active-clamped inductorless forward converter and the multiphase synchronous buck converter [7].

3. The multiphase synchronous buck converter

3.1. Input and output ripple currents

A major benefit of the multiphase converters is the fact that both the input and output ripple currents are reduced. Fig. 3 shows the rms values of the ac input currents, for single-phase, two-phase, three-phase, and four-phase interleaved buck converters vs. the duty ratio of the high-side FET, for the case when the ripple currents in the inductors are negligible. (The base of normalization is the output current.) It can be seen that the rms current in the input filter capacitor of the multiphase converter is much smaller than that of the single-phase converter. Also, as the plots in Fig. 4 indicate, the ripple current flowing in the output filter capacitor is substantially reduced. That figure shows the normalized ripple currents of the two-phase, three-phase and four-phase buck converter, with the ripple current of the single-phase converter used as the basis of normalization. It is interesting to note that both the input and output ripple currents become zero at the duty ratios M/N, where M is an integer number between 1 and N-1 and N is the number of phases.

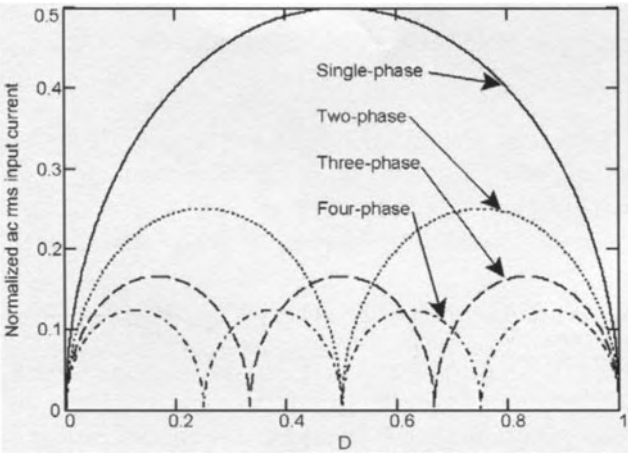

Fig. 3. Normalized rms ac input current of the multiphase buck converter vs. the duty ratio (base of normalization: output current).

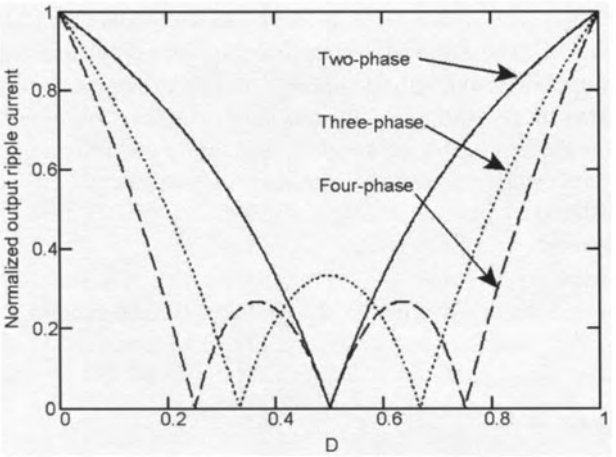

Fig. 4. Normalized output ripple current of the multiphase buck converter vs. the duty ratio (base of normalization: output ripple current of the single-phase converter).

3.2. Theoretical load transient response limits

The main dynamic requirement for a converter powering a microprocessor is that it responds rapidly to step changes in the load current. In this respect, a multiphase converter is essentially equivalent to its single-phase version operating at N times higher switching frequency and having an inductor with an inductance that is 1/N times smaller.

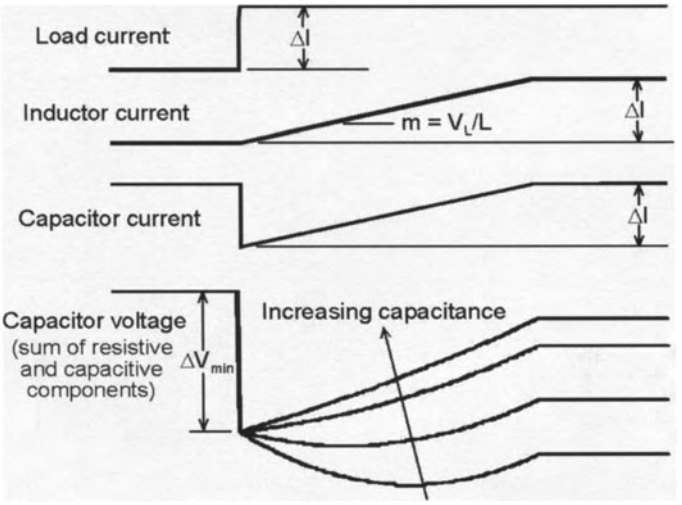

Fig. 5. Waveforms for determining the theoretical load transient response limits. (V_L is the maximum voltage available to change the current in the inductor.)

The theoretical load transient response limit can be understood by considering the waveforms in Fig. 5. The assumptions are that (1) the only parasitic component to consider is the equivalent series resistance (ESR) of the output capacitor, and (2) the ripple component of the inductor current is negligible. (Note: While the second assumption is not always true for a single-phase converter, it is a good approximation for multiphase converters, where the individual inductor ripple currents tend to cancel each other.)

It is straightforward to calculate the voltage deviation vs. time, for a step change in the load current and for a control that forces the inductor current to change towards the new steady-state value as fast as possible [7]. The result is:

$$\Delta v_o(t) = t^2 \frac{m}{2C} + t\left(mR_e - \frac{\Delta I}{C}\right) - R_e \Delta I \tag{1}$$

where C is the capacitance of the output capacitor, R_e is the ESR of the output capacitor, ΔI is the amplitude of the load step change, and m is the slope of the inductor current. In the case of an upward step in the load current

$$m = \frac{V_{in} - V_{out}}{L} \tag{2}$$

and in the case of a downward step in the load current,

$$m = -\frac{V_{out}}{L} \tag{3}$$

where L is the equivalent inductance of the converter, as defined at the beginning of this section.

From (1) the peak deviation is either

$$\Delta V = R_e \Delta I \tag{4}$$

or

$$\Delta V = \frac{\Delta I^2}{2mC} + \frac{mCR_e^2}{2} \tag{5}$$

depending on the time constant R_eC of the output capacitor. When that time constant is larger than $\tau_{crit} = \Delta I/m$, (4) gives the peak deviation and when the time constant is smaller than τ_{crit}, (5) gives the peak deviation. Due to the fact that the rising slope tends to be much smaller than the falling slope, it is possible that the peak deviation

for a upward load step is determined by (4) and for the downward load step it is determined by (5).

The significance of the capacitor time constant is that when it is relatively large (as in the case of electrolytic capacitors), the peak deviation is determined only by the ESR, and does not depend on the inductor current slope. This means that compliance with the load transient specifications can be achieved even when the slopes are small, i.e. when the switching frequency is relatively low (and the inductance is large), using an output bulk capacitor whose size and cost are essentially independent from the frequency. On the other hand, when the capacitor time constant is small (as in the case of multilayer ceramic capacitors, or MLCCs), the peak deviation depends on the slope, and so there is a strong inverse correlation between the capacitor size/cost and the switching frequency. This means that MLCCs as bulk output capacitors can only be used cost-effectively at high switching frequencies.

3.3. Duty-ratio considerations for multiphase converters

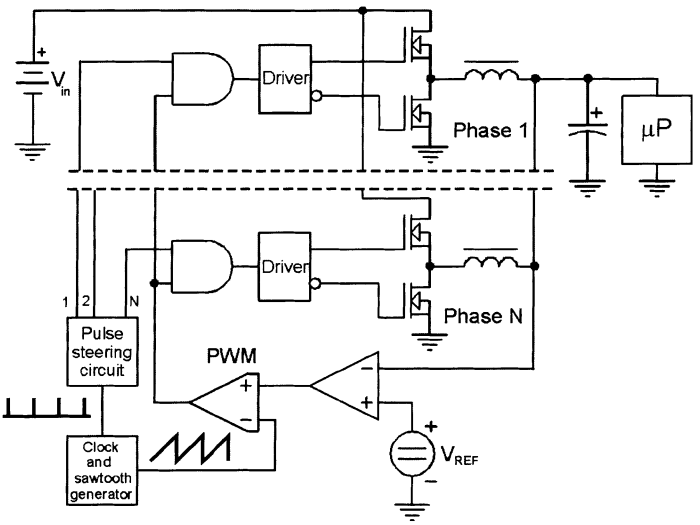

Fig. 6. N-phase buck converter with a single pulse-width modulator (for non-overlapping conduction of the control FETs).

Fig. 6 shows a typical control architecture for a multiphase interleaved buck converter. In that architecture there is a single pulse-width modulator (PWM), whose output pulses are sequentially sent to the control FETs of the individual phases. Fig. 7 shows another architecture, where there are as many PWMs as phases, and each PWM output is sent only to one phase. In this case the PWM carrier waveforms (usually sawtooth waves) are shifted from each other by 360°/N. The controller for the architecture in Fig. 6 is the less expensive one of the two versions due to its

simplicity (only one PWM and one PWM carrier waveform generator). Another advantage of that architecture is that in steady state the duty ratios of the control pulses for all phases are equal, while in the architecture of Fig. 7 the inevitable small differences in the PWM carrier waveforms lead to duty-ratio differences. The duty-ratio differences, in turn, can lead to unequal current sharing among the phases. The drawbacks of the architecture of Fig. 6 are: (1) It does not allow the maximum duty ratio to exceed 1/N, so in applications with low input voltage the number of usable phases is limited. For example, with 5V input and 1.5V output, where the steady-state duty ratio is close to 0.3, the maximum number of phases is limited to three. (2) Since overlapping conduction of the control FETs is not allowed, the maximum available slope of the inductor current for upward load steps is reduced, and actually becomes less than the available slope for downward load steps in the input voltage range from NV_{out} to $2NV_{out}$. In practice, this means that if the input voltage drops below $2NV_{out}$ (e.g. below 9V for a three-phase configuration with 1.5V output), the cost of the output capacitors to meet the load transient specifications will be higher than for the architecture of Fig. 7.

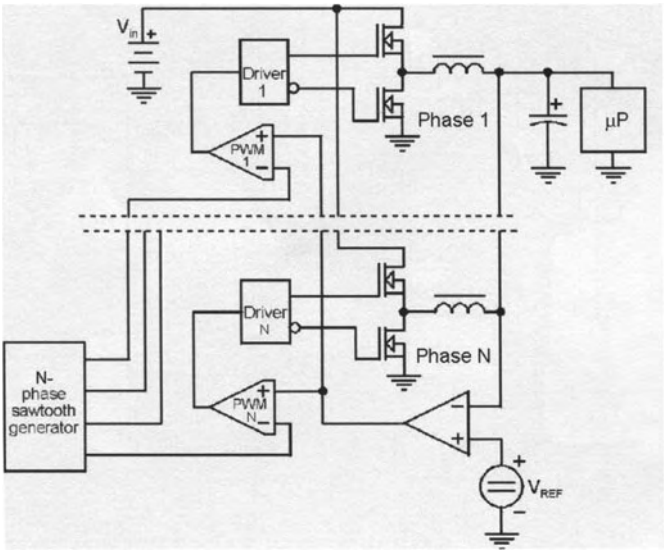

Fig. 7. N-phase buck converter with N pulse-width modulators (for overlapping conduction of the control FETs).

4. Control and protection considerations

4.1. PWM control techniques

The fundamental task of the controller for a CPU VRM is to ensure the regulation of the output voltage. Voltage regulation in a buck converter is implemented with duty-ratio control. There are many different types of duty-ratio modulators (often called

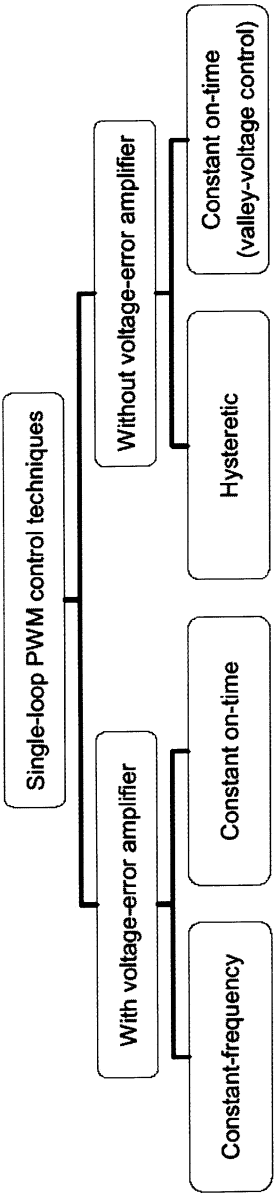

Fig. 8. Hierarchical chart of commonly used single-loop PWM control techniques.

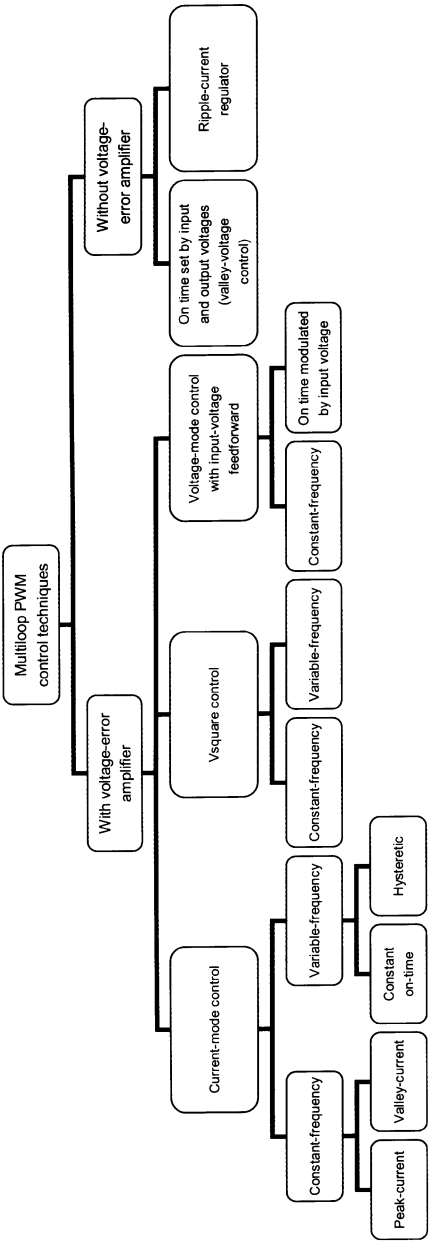

Fig. 9. Hierarchical chart of commonly used multi-loop PWM control techniques.

pulse-width modulator or PWM), including constant-frequency and variable-frequency types, with single-loop or multiple-loop feedback/feedforward control, and with or without voltage-error amplifier. Figs. 8 and 9 show the hierarchical charts of the various types used in, or considered for, CPU buck converters. Many of those control techniques, including the single-loop control techniques with voltage-error amplifier, the hysteretic controller without error amplifier, and the various current-mode controllers, are described in [9]. The Vsquare control is described in [10].

Recently, the idea of digital control of VRMs has received increased attention [11]-[14], and at least one vendor released a chipset with digital control for such applications [15]. Whether digital control of VRMs is a more cost-effective way of achieving compliance with the required specifications than analog control still remains to be seen.

4.2. Implementing the load line or voltage positioning

The microprocessor specifications require that the output voltage be a well-defined linearly decreasing function of the load current (see, e.g. Figs. 1 and 2 of [1], or Section 2.2 of [2]). This requirement is equivalent to having a resistive output impedance over a wide frequency range. The resistive output impedance allows the output voltage to decrease, or droop, in proportion to the instantaneous processor current. The droop is often called "voltage positioning," due to the fact that heavy load positions the output voltage near the lowest voltage where the CPU is still operational, and light load positions the output voltage near the maximum voltage acceptable for the CPU. The optimized voltage positioning (the one where the bulk capacitor is selected to have an ESR to be equal to the resistance represented by the load line) provides two major advantages: (1) It allows the minimization of the bulk output capacitor by essentially reducing the maximum peak-to-peak voltage deviation to be equal to the product of the maximum load step magnitude and the ESR of the bulk capacitor, at any pulse width of the load current step. (2) By allowing the output voltage to sag to its specified minimum value at full load reduces the average dissipation in the microprocessor.

There are several ways of implementing voltage positioning. The first, widely used, solution was adding a "droop" resistor with a resistance equal to the value represented by the required load line between the inductor and the bulk capacitor, and taking the voltage-regulating feedback from the junction of the inductor and the resistor (Fig. 10). Unfortunately this, traditional, approach has several drawbacks, including the cost of the high-current precision resistor, reduced efficiency, and most importantly the fact that the transient response to a step load change shows a spike-back, which, in case of narrow load pulses, prevents the output voltage to remain between the upper and lower tolerance limits of the load line. (This issue will be discussed in more detail in the next subsection.)

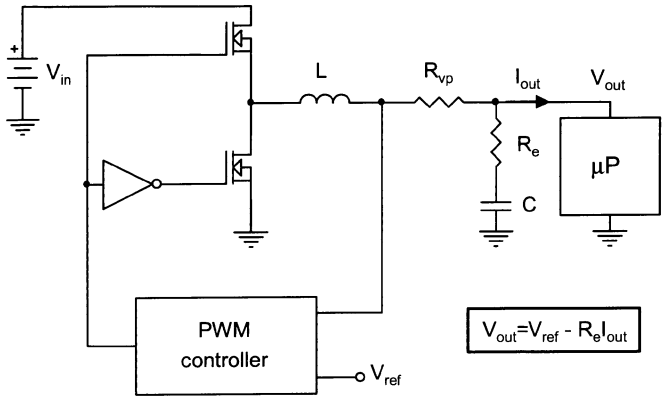

Fig. 10. Traditional implementation of the voltage positioning using a resistor between the inductor and the bulk capacitor.

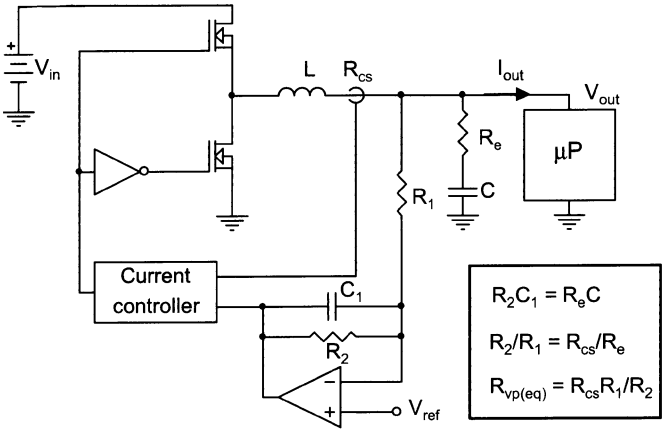

Fig. 11. Implementing ADOPT™ voltage positioning in a current-mode controlled converter.

The next three methods (collectively called ADOPT™, for Analog Devices' Optimal Positioning Technology) are capable of alleviating or completely avoiding the drawbacks of the droop resistor technique. They are implemented by using either (1) a current-mode-controlled converter with an optimized compensation of the voltage-error amplifier (Fig. 11, [8], [16], [18]), or (2) a converter with fast voltage regulation (e.g. a hysteretic regulator) complemented with a current sensor and an external RC network between the output, the reference voltage and the noninverting input of the PWM controller (Fig. 12, [16], [19], [25]), or (3) a voltage-mode-controlled converter, in which the total inductor current signal is subtracted from the reference voltage and the voltage-error amplifier is optimally compensated (Fig. 13, [17], [20]). All three ADOPT™ methods allow setting the output impedance of the converter to

be equal to the ESR of the bulk capacitor over a wide frequency range, thereby avoid the spike-back and permit the minimization of the size and cost of that capacitor. These methods can also be used in the cases where the ESR is smaller than the specified load-line resistance or where a substantial amount of low-ESR capacitors (typically multilayer ceramic capacitors for local bypassing around the microprocessor) is placed in parallel with the bulk capacitor.

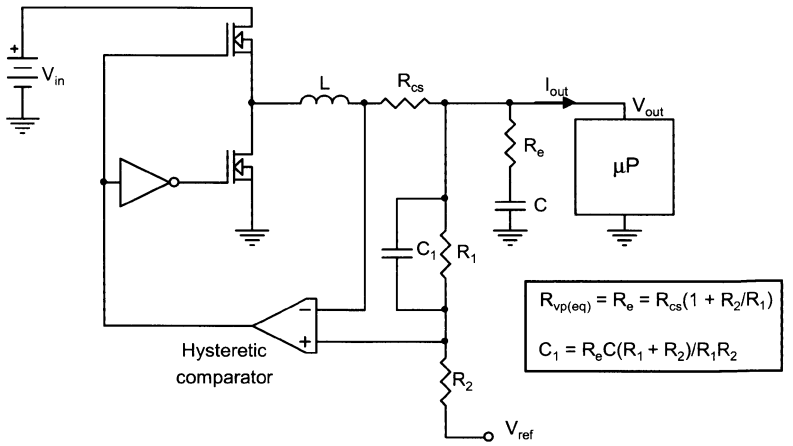

Fig. 12. Implementing ADOPT™ voltage positioning in a converter with hysteretic regulation.

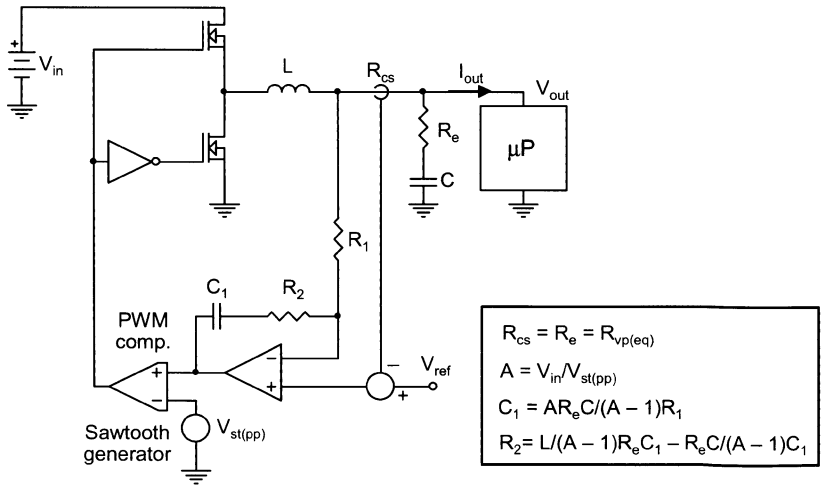

Fig. 13. Implementing ADOPT™ voltage positioning in a voltage-mode controlled converter.

4.3. Load transient response

The reason why the traditional voltage positioning technique discussed at the beginning of the previous subsection does not work well with narrow load pulses can be understood with the help of Figs. 14 and 15. Fig. 14 shows the details of the spike-back for the case of load-current step-up. Initially the voltage drop is equal to the product of the current step ΔI_{out} and the ESR of the bulk capacitor, but then the voltage-regulating feedback loop attempts to bring back the voltage rapidly to the value corresponding to the product of ΔI_{out} and the parallel equivalent of the ESR, R_e, and the droop resistance, R_{vp}. The speed of the voltage rise is limited by the available rate of inductor current change (di/dt), so the peak of the spike-back voltage is delayed by a time interval that depends on di/dt. Eventually the output voltage will settle at the new value determined by the product of ΔI_{out} and R_{vp}.

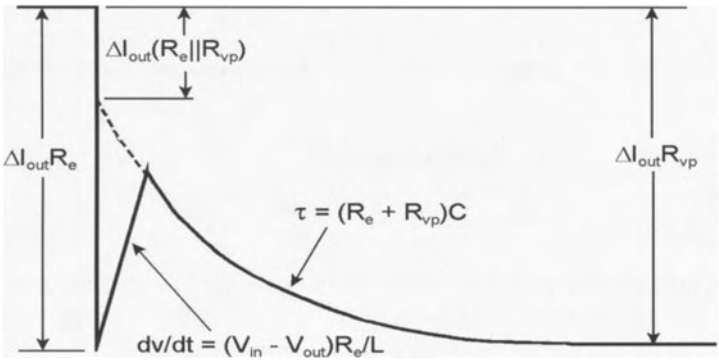

Fig. 14. Details of the spike-back waveform.

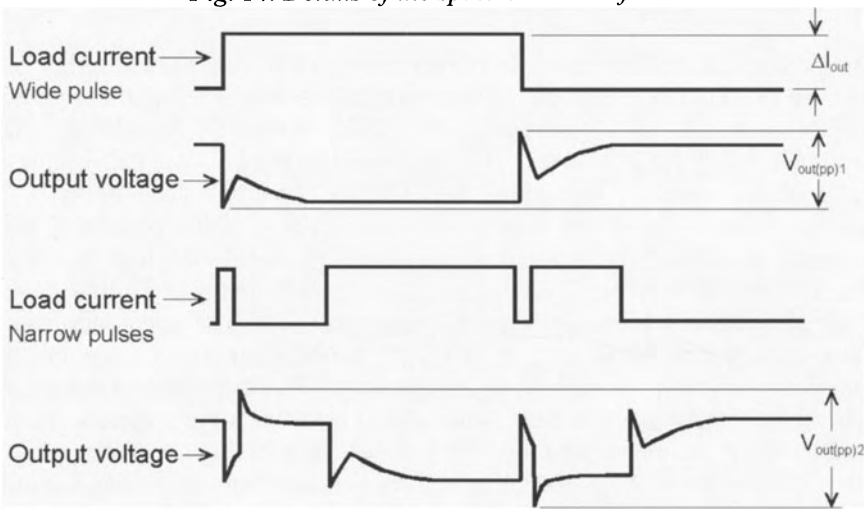

Fig. 15. Waveforms with traditional voltage positioning.

Fig. 15 shows what happens when the load current pulse is narrow and the peak or valley of the spike-back coincides with the termination of the pulse. Clearly, overshoot and/or undershoot appear, which increase the peak-to-peak deviation. The worst-case deviation can be almost twice the product of the current step amplitude and the voltage positioning resistance.

In the case of ADOPT™, the load transient response is an almost ideal step, as can be seen in Fig. 16, which shows the measured load-current step response of a CPU converter on two different time scales.

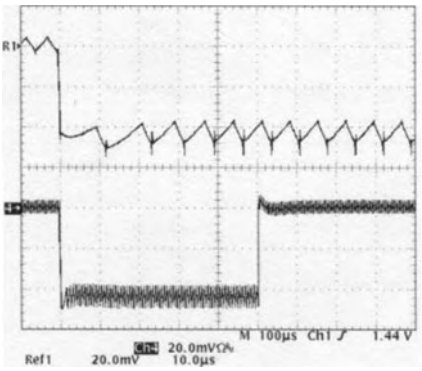

Fig. 16. Measured output voltage response of an ADOPT™ converter to a 14A step change in the load current. Top time scale: 10 µs/div, bottom time scale: 100 µs/div. Voltage scale: 20 mV/div.

4.4. Maintaining high efficiency over a wide range of load-current variation

In mobile applications it is important to maximize the battery lifetime. This effort is helped by the improved power management of the newer generations of mobile microprocessors (IMVP, QuickStart, DeepSleep, Enhanced SpeedStep, Thermal Throttling, see, e.g. [21]-[23]). It is also important that the losses in the various power supplies of the computer (including the VRM for the CPU) be minimized. Those losses are around 10% of the total power consumption of the computer. Since in battery-powered operation the microprocessor spends most of the time in low-power mode, it is important for the VRM to maintain high efficiency down to a fraction of one percent of the full load. In addition to optimized design of the power section of the converter, the high efficiency at light load can be achieved by a combination of essentially two types of power-saving measures, (1) decreasing the switching losses, which tend to be dominant at light loads, and (2) eliminating the circulating current caused by the reverse conduction of the low-side, synchronous rectifier, FET when the load current drops below the value of critical conduction (i.e., when the inductor valley current reaches zero).

Above the load current corresponding to critical conduction, the switching-loss reduction can be achieved by phase shedding. Phase shedding means that out of the N phases of the multiphase converter one or more phases are turned off whenever the load current drops below a certain level. The drawback of phase shedding is that the output ripple voltage increases.

If the load current drops below the critical conduction, the losses caused by the circulating current can be reduced by turning off the synchronous rectifier in the instant when the inductor current reaches zero. In this case the synchronous rectifier essentially emulates a true diode, i.e. it does not allow the current in it to reverse (see Fig. 17). With diode emulation below the load current corresponding to critical conduction, the converter operates in discontinuous inductor-current mode (DCM), and the rms currents in the inductor and also in the power FETs are reduced, and so the conduction losses are cut back.

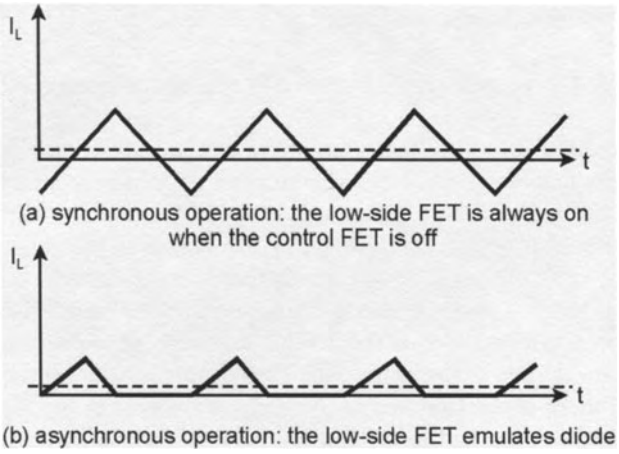

Fig. 17. Inductor current waveforms at light load.

In DCM the switching losses can be reduced by frequency modulation, without the penalty of increasing output ripple. By making the switching frequency proportional to the load current, the efficiency can be made essentially independent from the load current, at least until the bias power consumptions of the controller and driver ICs become comparable to the output power.

Fig. 18 shows the calculated efficiency of a typical 300kHz four-phase converter with diode emulation, in four-phase operation and in single-phase operation (i.e. after three phases were shed). Fig. 19 shows the calculated efficiency of a typical 300kHz single-phase converter in three cases, (1) without diode emulation, (2) with diode emulation but without frequency modulation in DCM, and (3) with both diode emulation and frequency modulation in DCM. In the last case the switching frequency is proportional to the load current.

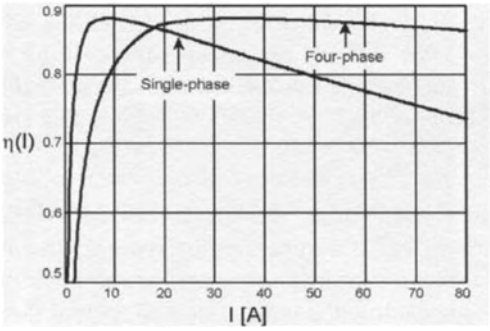

Fig. 18. Calculated efficiency of a 300kHz four-phase converter with diode emulation, in four-phase operation and after three phases were shed, (i.e. in single-phase operation).

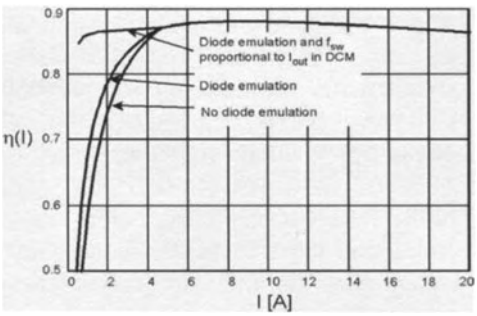

Fig. 19. Calculated efficiency of a 300kHz single-phase converter without diode emulation, with diode emulation, and with diode emulation and frequency modulation in DCM ($f_{sw} \sim I_{out}$).

4.5. Current/thermal balance in multiphase converters

Imbalanced operation in multiphase converters can manifest itself in uneven current sharing and/or in uneven junction temperatures of the power FETs. The causes of imbalances can be (a) duty-ratio differences caused by either unequal PWM sawtooth amplitudes or uneven propagation delays, (b) differences in the channel resistances of the power FETs, or (c) thermal-resistance differences. Duty-ratio differences lead to unequal current sharing among the phases. Employing current-mode control or adding the current information to the PWM sawtooth in voltage-mode control are effective ways to balance the currents. The problem with these approaches is, however, that if there is a difference in the $R_{ds(on)}$ values or in the thermal resistances among the phases, then current balancing will exacerbate the temperature differences. The temperature differences caused by the $R_{ds(on)}$ or thermal-resistance differences can be reduced by allowing the positive temperature coefficients of the FET channel resistances (about 0.4%/°C) equalize them. The higher the temperature of a FET is, the higher its resistance will be, which will reduce the current flowing in it. The optimal solution seems to combine the two balancing techniques, i.e. to trade some current balance for thermal balance by using voltage-mode control with a small amount of current signal added to the PWM sawtooth. The best thermal balance is obtained if the voltage drops across the conducting FETs are used as current signals. The drawback of a good thermal balance is, however, that it forces uneven phase currents, therefore the inductor sizes must be increased in order to be able to handle the highest possible current without entering in saturation.

4.6. Current-sense solutions

The load-line specification requires the controller to receive accurate information about the total inductor current. A straightforward way of generating the current

signal is to use a current-sense resistor between the common junction of the inductors and the output bulk capacitor. The drawbacks of that solution are cost, extra dissipation, and layout difficulties. A more practical approach is to use the winding resistances of the inductors. The current signal information can be extracted by using an RC filter across the inductor with a time constant that matches the $L/R_{winding}$ time constant of the inductor [24]. Fig. 20 shows the implementation of the idea in a two-phase converter. Unfortunately, sensing the voltage drop across the winding resistances introduces a positive temperature coefficient, and this makes it difficult to meet the accuracy requirements. Fig. 21 shows an economical way of canceling the temperature coefficient of the winding resistances with a single NTC resistor in the feedback network of the current-sense amplifier [17], [20].

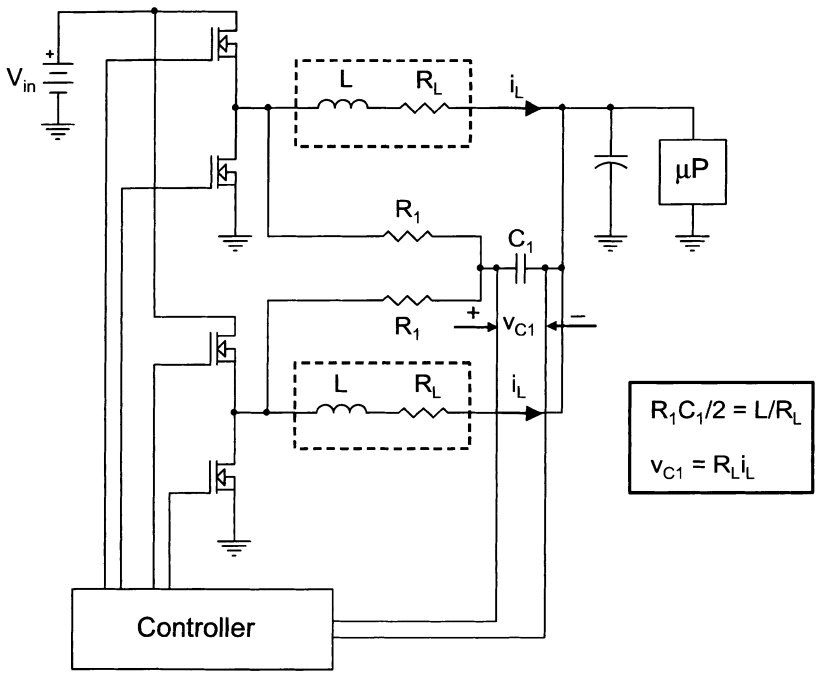

Fig. 20. Sensing the total inductor current in a two-phase converter, using the winding resistances of the inductors. Two RC circuits with a shared C recover the total inductor current signal.

4.7. Optimal tri-loop control for multiphase converters

Fig. 22 shows a tri-loop control architecture for multiphase voltage-mode converters [17]. The three control loops are (1) a common feedback loop for output voltage regulation, (2) individual phase current loops for current/thermal balancing, and (3) a common impedance loop to implement the required load line.

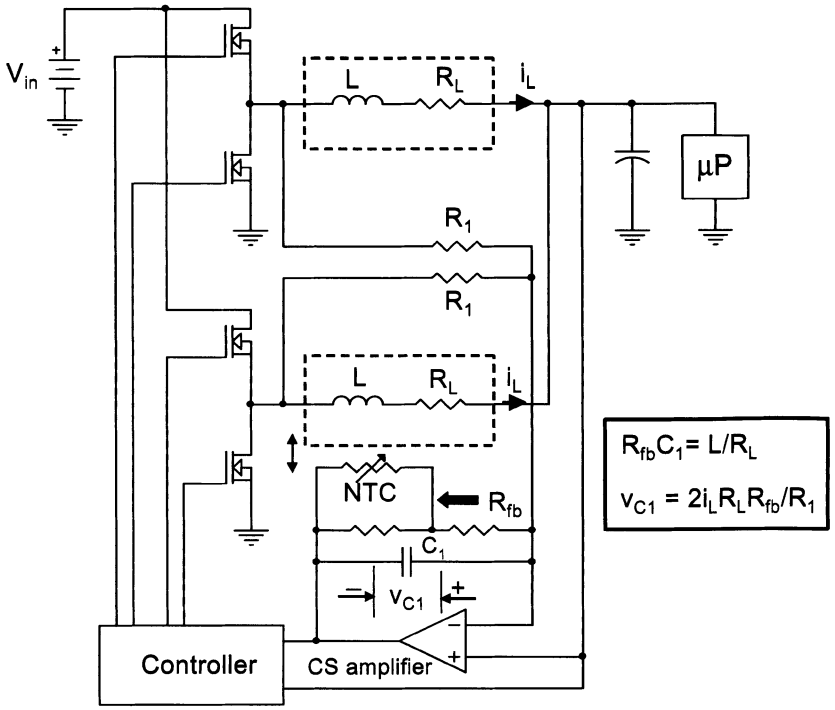

Fig. 21. Sensing the total inductor current in a two-phase converter, using the winding resistances of the inductors. A single NTC resistor cancels the temperature dependence of the winding resistances.

The voltage-regulating feedback loop has high dc gain for accurate tracking of the difference between the reference voltage and the total inductor current signal, and is compensated for ADOPT™ load transient response according to the equations in Fig. 13. The phase current loops are using the voltage drops across the synchronous rectifier FETs, thus providing the best trade-off between current and thermal balance. The common impedance loop is implemented by subtracting the total inductor current signal from the reference voltage. That current signal can be generated either by a sense resistor between the junction of the inductors and the bulk capacitor or by using the winding resistances of the inductors, as per Fig. 21. The tri-loop control of Fig. 22 represents the optimal solution for multiphase voltage-mode converters for VRM applications, and is implemented in the ADP3168 IC from Analog Devices.

4.8. Protection considerations

It is expected that the VRM is (a) protected against sustained overload at the output, and (b) designed such that under no circumstances does an overvoltage or reverse voltage appear on the output.

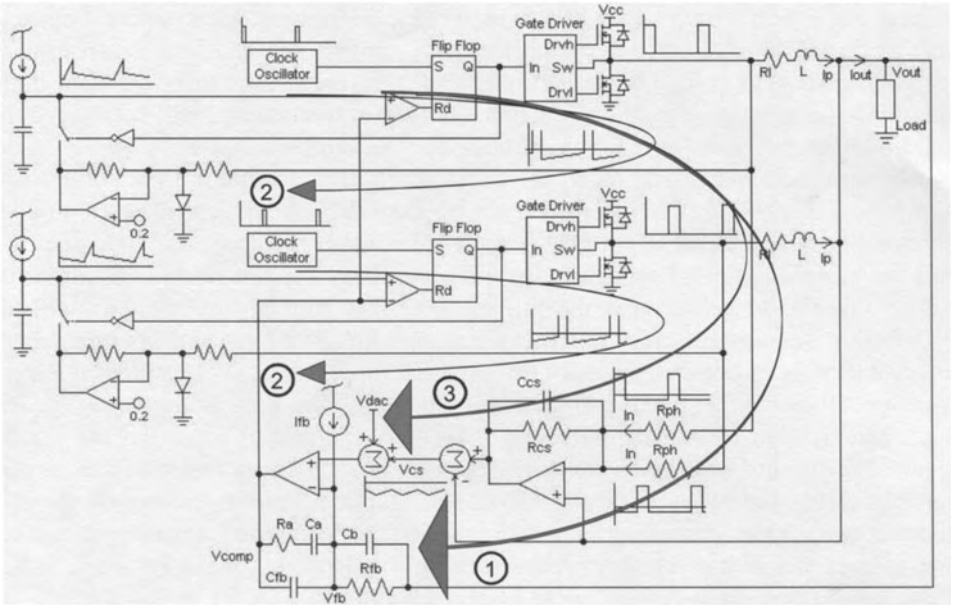

Fig. 22. Optimal tri-loop control of multiphase converters in VRM applications.

The overload protection can be continuous (cycle-by-cycle) type, hiccup-type, or latch-off type. Each type has its set of advantages and disadvantages. The continuous protection provides safe start-up with any output capacitance, and if the overload is removed the output recovers without the need for recycling the input supply. On the other hand, the continuous operation during sustained overload leads to hot spots and a complete battery discharge. Also, the VRM can overheat if the output current limit is not well controlled. (Note: Foldback current limit helps preventing the overheating.) The hiccup protection reduces the average current by the hiccup on/off ratio, so hot spots and overheating are avoided. Also, if the overload is removed the output recovers without recycling the input supply. The drawbacks of the hiccup protection include the danger of start-up difficulty with large output capacitances and the additional complexity. The latch-off type protection completely eliminates the power consumption after the latch-off, but can fail to start with a large output capacitance; also, it does not provide automatic recovery after the overload is removed.

The most likely causes of an output overvoltage are sudden load removal (note: this should never happen in a VRM that complies with the voltage regulation specifications) or a drain-source short of the high-side FET. The VRM specifications expect that if the output is above the maximum VID level by a certain margin (usually 200 mV), or above a set absolute voltage specified for the CPU, the VRM shuts off the supply to the processor. This should be done by an output overvoltage

protection (OVP) circuit with a path separate from the voltage sense path used for voltage regulation. The concern here is to avoid false tripping due to external noise or an inaccurate threshold. Also, the protection circuit should enable a low-resistance path (a crowbar) to ground with a sufficiently fast response time such that if the high-side FET shorts to the input power source then the output voltage will not exceed the maximum specified voltage for the processor. The crowbar is usually the low-side, synchronous rectifier, FET. When the overvoltage happens due to the failure of the high-side FET, the crowbar is turned on by the OVP circuit and blows a fuse. A potential problem is that if the fuse is replaced without replacing the high-side FET and the system is turned on again, the supply voltage for the control IC comes up slowly and the processor sees the full input voltage through the shorted high-side FET before the crowbar could activate again, and this destroys the processor. A fail-safe solution is to use a dedicated crowbar FET directly across the output that runs from an "always-on" rail and is turned on by default when the system comes up. The dedicated crowbar FET is also effective against reverse voltage caused by the current reversal in the inductor when the low-side FET is shorted. In that case the reverse current leads to a voltage reversal across the output capacitor and possibly to the destruction of the processor. Fig. 23 shows the combined implementation of overvoltage and reverse-voltage protection with a dedicated crowbar FET in a single-phase converter with hysteretic ripple-current control and ADOPT™.

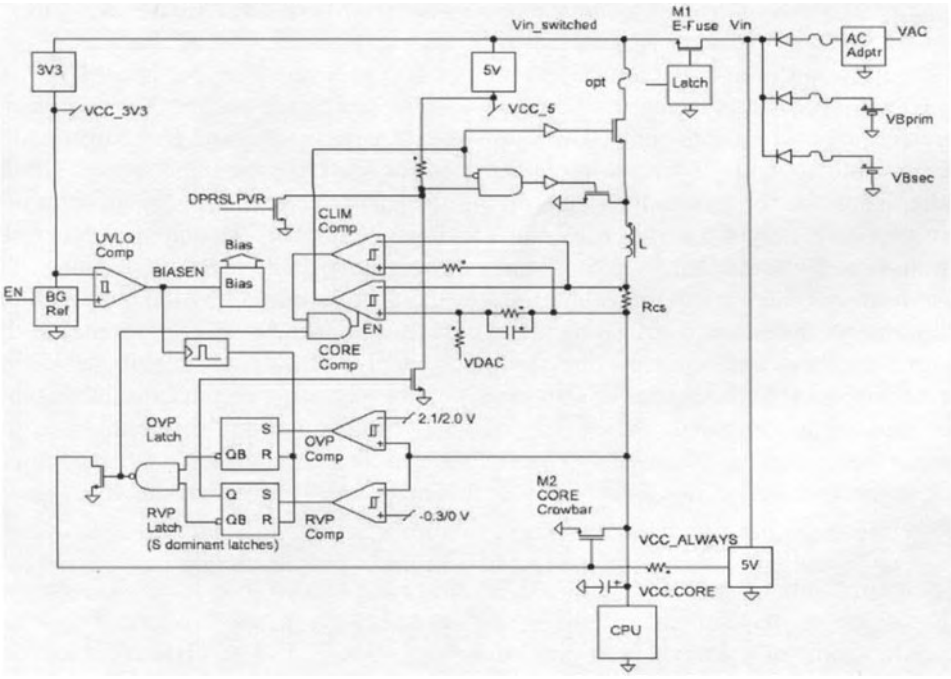

Fig. 23. Combined implementation of overvoltage and reverse-voltage protection with a dedicated output crowbar FET.

5. Example control IC architectures

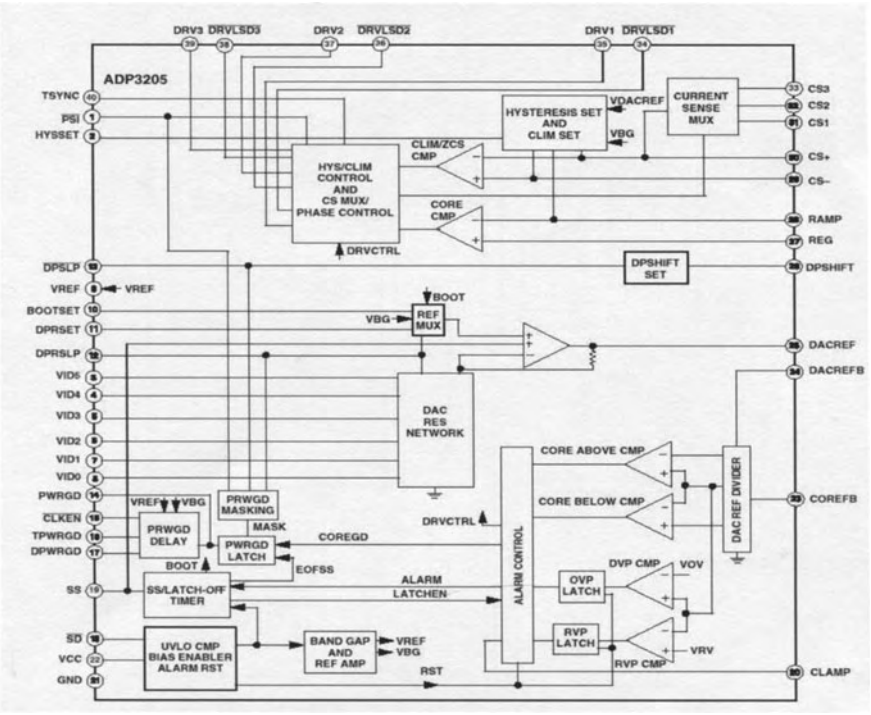

Fig. 24. Functional block diagram of the ADP3205 IC.

Fig. 24 shows the functional block schematic of a dc-dc controller IC, the ADP3205 from Analog Devices [25] developed for notebook computer applications. The device utilizes hysteretic ripple current control, which is a combination of the hysteretic regulator principle with ADOPT™ voltage positioning. Its features include the following:

- Pin-programmable one-, two-, or three-phase operation
- Compliance with Intel IMVP-IV specifications
- Inherent static and dynamic current sharing
- ADOPT™ voltage positioning
- Noise blanking after a leading edge to avoid false switching due to noise pickup
- Diode-emulation control of the low-side FET for optimized light-load efficiency
- Six-bit VID DAC
- Soft DAC output voltage transition for dynamic VID change
- Masked Power Good during output voltage transients
- Cycle-by-cycle current limiting with latched or hiccup overload protection
- Soft start-up
- Two-level overvoltage and reverse-voltage protection

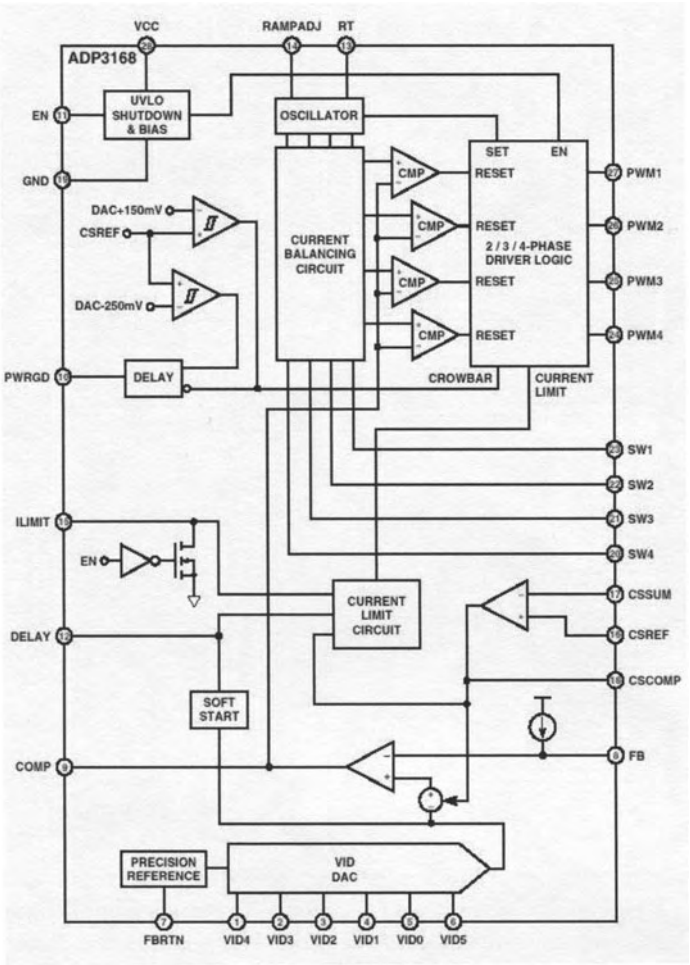

Fig. 25. Functional block diagram of the ADP3168 IC.

Fig. 25 shows the functional block schematic of the ADP3168, also from Analog Devices [20]. The device was developed for desktop dc-dc converter applications, and utilizes the optimal tri-loop control architecture discussed in Section 4.7. Its features include the following:

- Pin-programmable two-, three-, or four-phase operation
- Compliance with Intel VRD-10 specifications
- 1MHz per phase switching frequency
- Active current/temperature balancing
- Temperature-compensated current sensing, using the inductor winding resistances as sense resistors

- ADOPT™ voltage positioning
- Six-bit VID DAC
- Built-in Power-Good and crowbar blanking to support on-the-fly VID code changes
- Programmable overload protection with programmable latch-off delay

6. Challenges for the designers of VRM control ICs

The design of control ICs for VRM applications is a challenging task. The challenges include, but are not restricted to, the following:

- Low cost and extremely short development schedule
- High-accuracy band-gap reference [less than ±1% error over wide temperature range (0 to 100 C°) and over process variations], with good rejection of the input ripple voltage (40 dB or better) at the effective switching frequency, and with floating ground arrangement to comply with the remote sensing requirements
- VID DAC with a resolution up to eight bits, with monotonicity and good linearity, and with controlled slew rate to comply with VID on-the-fly specification
- Voltage-error and current-sense amplifiers with low offset (less than ±1-2 mV)
- In mobile applications: Low supply voltage (3.3V nominal), with full functionality down to around 2.7 V
- Complex system interface (PowerGood signal, power-saving control features, VID on-the-fly control, soft transition between active and power-saving modes for low acoustic noise)

7. Summary

This paper discussed the converters used for notebook and desktop CPU VRM applications, with emphasis on the control aspects having relevance on the architecture of PWM control ICs. The discussion included the review of power converter topologies, some of the important characteristics of the multiphase buck converter, the various PWM control techniques, the implementation of the load-line specification, the load transient response and the related ADOPT™ voltage positioning technology, the techniques to maintain high efficiency over a wide range of load-current variation, current and thermal balancing, current-sense solutions, and protection against overload, output overvoltage and output reverse voltage events. The paper concluded with the presentation of the functional block schematics of two control ICs for microprocessor VRMs and with a summary of control IC design challenges.

References

[1] Intel document, "VRM 9.1 DC-DC Converter Design Guidelines," January 2002, Order Number: 298646-001.

[2] Intel document, "Voltage Regulator-Down (VRD) 10.0 Design Guide," February 2004, Document Number: 252885-003.

[3] D. R. Garth, W. J. Muldoon, G. C. Benson, and E. N. Costague, "Multi-phase, 2 kilowatt, high voltage, regulated power supply," *IEEE Power Conditioning Specialists Conference Record*, pp. 110-116, 1971.

[4] M. Ye, P. Xu, B. Yang, and F. C. Lee, "Investigation of topology candidates for 48V VRM," *Proc. APEC '02*, Paper no. S15P7.

[5] M. Qiu, P. K. Jain, and H. Zhang, "An AC VRM topology for high Frequency AC power distribution systems," *Proc. APEC '02*, Paper no. S5P1.

[6] P. Xu, J. Wei, K. Yao, Y. Meng and F. C. Lee, "Investigation of candidate topologies for 12V VRM," *Proc. APEC '02*, Paper no. S15P5.

[7] Y. Ren, M. Xu, K. Yao and F. C. Lee, "Two-stage 48V power pod exploration for 64-bit microprocessor," *Proc. APEC '03*, Paper no. S10P3.

[8] R. Redl, B. P. Erisman, and Z. Zansky. "Optimizing the load transient response of the buck converter," *Proc. APEC '98*, pp. 170-176.

[9] A. S. Kislovski, R. Redl, and N. O. Sokal, *Dynamic Analysis of switching-mode dc/dc converters*, Design Automation, Inc., Third printing, June 2003.

[10] U.S. Patent, 5,770,940, "Switching regulator," June 23, 1998, by D. Goder.

[11] A. Wu, J. Xiao, D. Markovic, and S. R. Sanders, "Digital PWM control: application in voltage regulation modules," *PESC '99 Record*, Paper no. 02_6.

[12] J. Xiao, A. V. Peterchev, and S. R. Sanders, "Architecture and IC implementation of a digital VRM controller," *PESC '01 Record*, Paper no. 10637.

[13] J. A. Abu-Qahouq, N. Pongratananukul, I. Batarseh, and T. Kasparis, "Multiphase voltage-mode hysteretic controlled VRM with DSP control and novel current sharing," *Proc. APEC '02*, Paper no. S15P1.

[14] S. Saggini, M. Ghioni, and A. Geraci, "A low-complexity high-performance digital control architecture for Voltage Regulator Modules," *PESC '03 Record*, Paper no. 07_01.

[15] ISL6580 and ISL6590 data sheets from Intersil Corp.

[16] U.S. Patent 6,229,292, "Voltage regulator compensation circuit and method," May 8, 2001, by R. Redl, B. P. Erisman, J. M. Audy, and G. Reizik.

[17] U.S. Patent 6,683,441, "Multi-phase switching regulator," Jan 27, 2004, by T. F. Schiff, J. C. Buxton, and R. Redl.

[18] ADP3164 data sheet from Analog Devices, Inc.

[19] ADP3422 data sheet from Analog Devices, Inc.

[20] ADP3168 data sheet from Analog Devices, Inc.

[21] Intel document "Intel Low Power Technologies for Mobile PCs".

[22] M. Gochman, et al, "The Intel® Pentium® M processor: microarchitecture and performance," *Intel Technology Journal*, May 21, 2003, pp. 21-36.

[23] D. Genossar and N. Shamir, "Intel® Pentium® M processor power estimation, budgeting, optimization, and validation," *Intel Technology Journal*, May 21, 2003, pp. 44-49.

[24] B. Carsten, "Designing high frequency current shunts and current transformers," *HFPC - April 1994 Proceedings*, pp. 539-549.

[25] ADP3205 data sheet from Analog Devices, Inc.