Achieving a Stable Power Supply with Fast Transient Response Through Digital Control

Throughout the entire history of power system design there has been one requirement that has been constant: the need for a stable power supply. With the evolution of digital control, the elusive problem of achieving power supply stability has been conquered. With the introduction of Intersil's ChargeMode™ control loop technology, the ZL8800 implements a digital control topology that provides a compensation free design environment. Designed for fast transient response, the ZL8800 is able to react to a transient event in a single cycle, reducing the amount of output capacitance needed in a system. The end result is a savings in cost and board space. **Chance Dunlap, Sr. Marketing Manager Infrastructure Power, Intersil, USA**

New trends and changing requirements in the power management industry influence how we design power supplies. Some, like the need for telemetry and system information are more recent due to advancements in technology. Others, like the demand for smaller solutions, higher levels of integration, faster transient response and high switching frequencies have been around for decades. This article details the common problems plaguing analog voltage mode control loops and shows how a digital control loop is able to provide bandwidth that was previously dismissed as unobtainable. By walking through the variables associated with stability and comparing it to an analog control system, it can be seen that a digital control loop is able to achieve faster transient response and improved performance while maintaining a stable loop.

Designing for stability

When designing point of load (POL) power supplies from a distributed rail such as +12 V, numerous options exist. The common architecture is an analog fixed frequency, voltage mode control loop with type 3 external compensation. This is popular due to the ability to adjust poles and zeroes to extend the converter bandwidth while compensating for a range of output capacitors. With a constant switching frequency the inductor and capacitor currents are predictable allowing for optimal component choice. The

downside of this topology however, is the complexity of the compensation. Five external components are required which can take significant time to design. Given all the variations in the power train components it is often difficult to take full advantage of the bandwidth, the end result being a power supply with mediocre transient response. Rules of thumb suggest that a power supply bandwidth should be set in the range of 1/10 of the switching frequency (F_{sw}). But the concern is the higher that the bandwidth is pushed, the more likely the system is to suffer from pre-mature field failure. This can be seen by walking through the process of power supply design and device compensation.

When first approaching a design the common starting point is to design the power stage per the system requirements, then attempt to compensate it using resistors and capacitors in the feedback loop. The inductor and output capacitors are always sized to meet the performance specifications of the power supply, such as output current, voltage ripple, and transient performance. These components are not modified for the purposes of overall compensation as this is the responsibility of the controller or the small signal compensation components, otherwise the cost of the system could be impacted. To achieve a well compensated power supply, the expectation is that all the variables in the output filter remain constant. Unfortunately, if an analysis is performed the associated distribution of each variable

can have a significant effect. For a worse case model, a Monte Carlo simulation could be run combining all the variants of L and C to create the range of transfer functions from the power stage components. Based on these results, the only potential for a stable system would be to dramatically reduce bandwidth to avoid the effect from the double pole filter.

For instance, when determining the value of L, documentation usually calls out using the nominal value that is specified on the inductor datasheet. If a worse case analysis was to be performed, this would not be adequate. Inductors have variances in inductance based on multiple factors such as inductor current, temperature, frequency and aging effects. A good example of this is when looking at non-ferrite based inductors that are popular for switch mode power supplies. Across the rated current range, the inductance can vary significantly, dropping to less than 50 % of their initial value.

Output capacitors also have wide variations in both capacitance and ESR based on operating conditions that need to be accounted for. DC bias voltages (especially with ceramic capacitors), initial tolerances, temperature, and aging effects will all change the double pole location of the output filter in the frequency domain from one board to another. With changes in DC bias and temperature, ceramic capacitors can easily drop to less than 40 % of their stated value.

From just a cursory view it can be noted

that the only reliable way to compensate the device is to start decreasing the bandwidth, relying on a low frequency dominant pole to determine the bandwidth. This is in effect how internally compensated power devices are usually configured. They provide such a low bandwidth that any variation in inductance is unlikely to cause a stability problem. The downside is the restriction to use a large output capacitor to overcome the performance limitations.

This is the classic trade-off that has plagued power design engineers for years. You can optimize bandwidth at the expense of additional time calculating components, running the risk of instability due to component variation and aging. Or, utilize internal compensated devices that provide faster design time with the downside of higher BOM (bill of material) cost due to the extra capacitors required to compensate for the poor performance. The turning point to this dilemma has been the advent of digital control. By sampling the output voltage and converting it to the digital domain, advantages can be achieved through use of signal processing that would be impossible in the analog domain. The benefit of this digital signal processing is evident with the introduction of the ZL8800, a fourth generation digital controller that uses a proprietary ChargeMode control architecture to provide a compensation free solution without sacrificing performance.

ChargeMode controller

As can be seen in Figure 1, the ZL8800 is a dual channel PWM controller capable of converting a +12 V or +5 V to lower output voltages ideal for powering point of load applications. With a full digital control loop and internal non-volatile memory, the device eliminates the multitude of external components normally found on analog
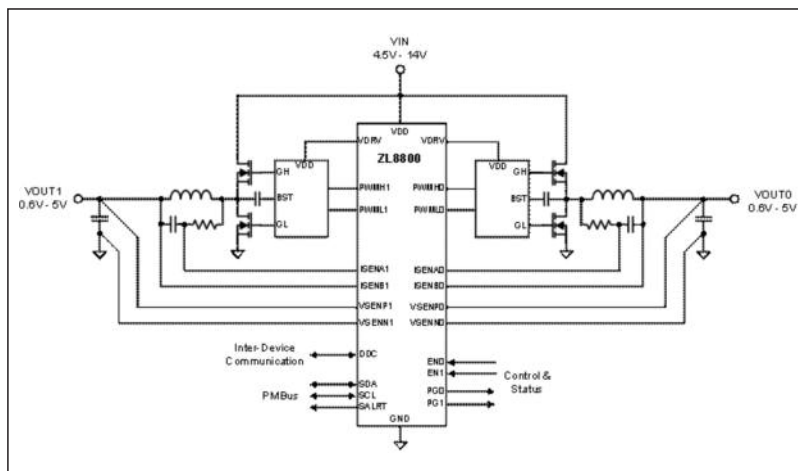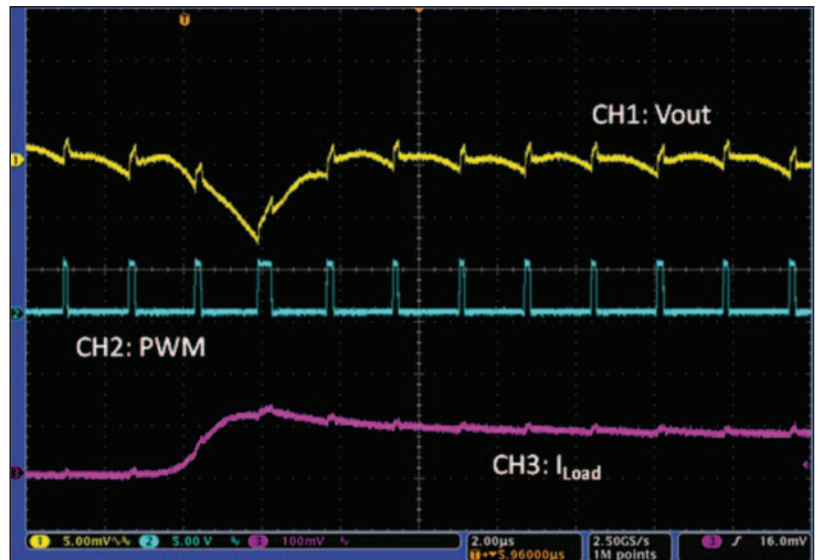

Figure 2: **ZL8800 single cycle response (1 V output with 10 A load step in 1 µs)**

controllers. With the ChargeMode control loop technology, the burden of selecting compensation components completely disappears.

The simplification of ChargeMode control scheme is that it responds immediately to any transient event by precisely modifying the duty cycle so that the amount of charge lost on the output capacitor can be replaced in a single cycle. This is achieved by over sampling the output voltage so that corrections can be made without having to know the actual capacitor value. The benefit of the non-linear response is that a load transient can be responded to and possibly corrected in one cycle with minimal ringing or overshoot. Comparatively in analog systems the ability to react quickly usually results in an under-damped system with low phase margin risking instability and excessive ringing.

An example of the ZL8800's ability to respond to a transient event is shown in Figure 2. In this situation a 10 A load step

is applied with a slew rate of 10 A/µs.

The ability to respond to a transient response at this speed sets the ZL8800 apart from other power controllers. But the real ability lies in the fact that this is achieved without compensation. In addition to the high speed loop, the ZL8800 ChargeMode control architecture is inherently stable. This means that any combination of inductance and output capacitance can be applied and the loop will remain stable. By designing a digital architecture that enforces system stability, any effect on the outside circuit such as removing capacitors dynamically or aging/thermal effects will not cause a problem.

One of the other benefits of the control loop is the ability to extend the bandwidth up to and beyond $F_{sw}/4$. The ability to achieve this is due to several approaches taken in the device. The first is a double edge modulation technique allowing the PWM signal to have a fixed frequency, but modulate both edges providing double the sampling rate in the system. Combined with a high speed ADC over-sampling the output, a wide bandwidth can be achieved without the high frequency phase roll-off common in other converters. As a result, when the bandwidth is increased the phase margin will remain stable.

Adjustable feedback gain

To allow the designer a degree of freedom in selecting the bandwidth, the ZL8800 incorporates a feedback gain term that can be adjusted to increase the response. The default gain of 256 provides a stable setting that will provide a bandwidth equivalent to an analog product that has been nominally compensated. Increasing the gain (typical range of 100 to 1200) allows a designer to dial in a faster
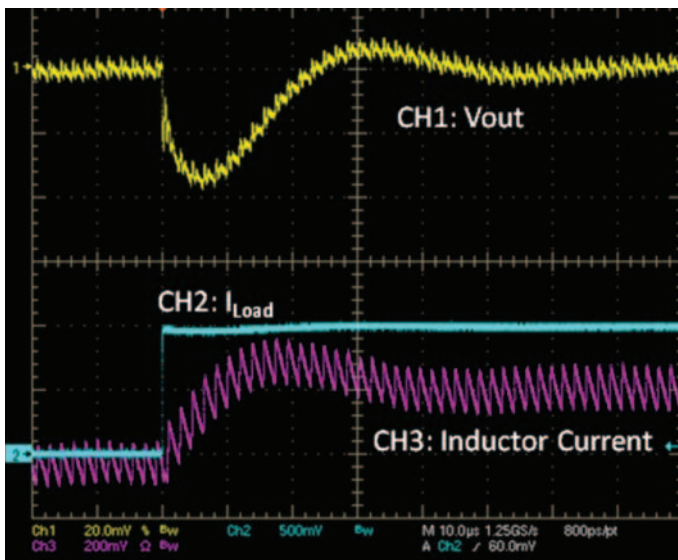

Figure 1: **Typical application diagram of ZL8800**

Figure 3:
Transient response with gain setting = 250 (bandwith = 26 kHz)

Figure 4:
Transient response with gain setting = 650 (bandwith = 70 kHz)

Figure 5:
Transient response with gain setting = 1050 (bandwith = 140 kHz)

response with an associated trade-off of higher jitter in the PWM signal. With typical systems employing low ESR capacitors, any increase in jitter has minimal effect, hence the recommendation is always to optimize for transient response.

The scope shots shown in Figures 3, 4 and 5 illustrate how increasing the gain of the ZL8800 can result in higher bandwidth and improved transient response. In all cases, the same control loop algorithm is applied, maintaining a stable system without the need for compensation. The scope shots were taken with a board under the following conditions:

- V_{in} = 12 V
- V_{out} = 1V
- C_{out} = 1500 µF
- Inductor = 300 nH
- F_{sw} = 550 kHz
- 10 A load transient applied at 100 A/µs.

Figure 3 shows the ZL8800 transient response while running with the feedback gain set to 250. When the loop gain was measured, this equated to a loop bandwidth of 26 kHz. Increasing the gain to 650 resulted in a faster transient response with improved settling time. This can be seen in Figure 4. With this setting, the equivalent loop bandwidth was 70 kHz. Increasing the gain further to 1,050 provided even faster response, allowing the loop to compensate in almost a single cycle. As can be seen in Figure 5 the inductor current slews up to the optimal point allowing the output voltage to immediately recover without overshoot or ringing. In this situation the loop bandwidth was 140 kHz which corresponds to just over 25 % of the switching frequency.
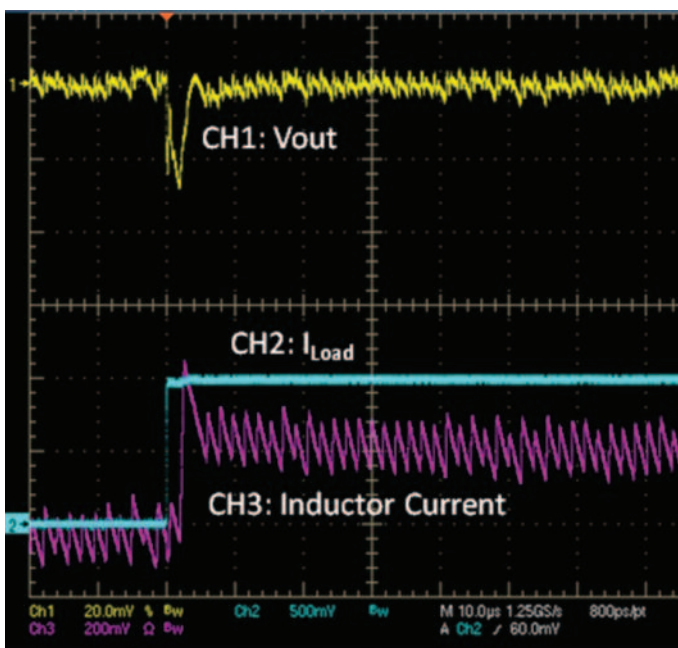
Conclusions

In conclusion, it can be seen that power converters relying on analog compensation techniques have limitations in providing high bandwidth designs without compromising stability or long term reliability due to component variation and aging. With the ZL8800 digital control architecture it is possible to design a power converter that is compensation free with high bandwidth. These capabilities allows for a power supply to be developed in a short period of time while saving on output capacitance.

ZL8800 Digital DC/DC Power Controller

Base-stations, routers, and similar infrastructure designs require fully control and monitoring of every power rail in order to maximize reliability, and typically use monitors and sequencers to provide added functionality around the POL converters. The ZL8800 integrates all of these functions, so a designer can monitor and control every aspect of the power supply through a PMBus™ interface, without the need for additional parts. Telemetric data is available on the power supply, including temperature, input current, output current, input voltage, output voltage and fault status.

Each ZL8800 output can operate independently or together in a dual phase configuration for high current applications. The device supports a wide range of output voltages (0.54 V to 5.5 V) operating from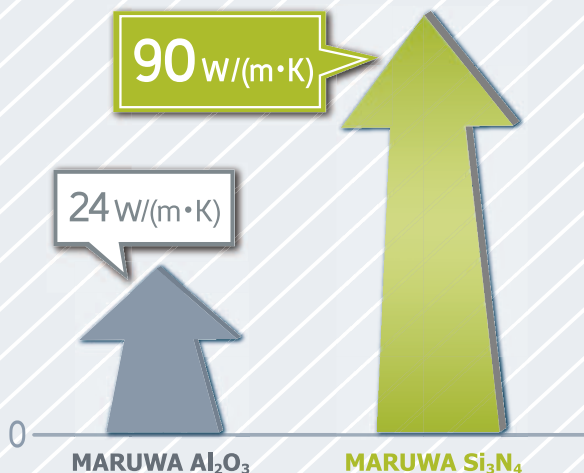 input voltages from 4.5 V up to 14 V. The ZL8800 is available in a 7 x 7 44-pin QFN package with prices starting at $4.90 each, in 1,000-piece units. There are also two evaluation kits available, the ZL8800-2CH-DEMO1Z 2-output demo board, and the ZL8800-2PH-DEMO1Z, a 60 A 2-phase demo board. Both are available for $150.00 each.

Intersil also announced a new version of its PowerNavigator™ graphical user interface, enabling designers to leverage all the capabilities of the ZL8800 digital controller without writing a line of code. Designers can use the drag-and-drop utility to set up and control any power supply architectur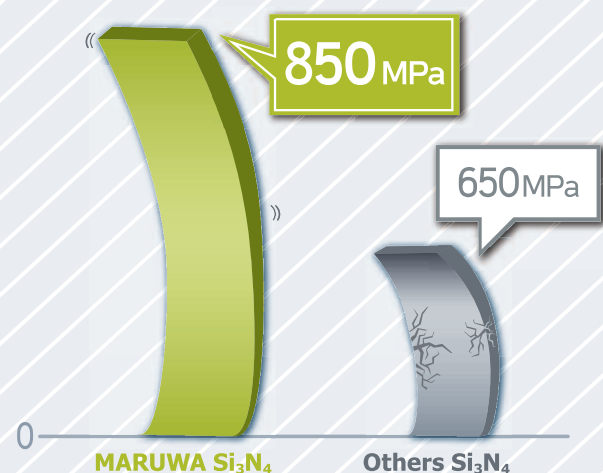e, further simplifying the development effort.