

Predictive Energy Balancing Enhances Control of Power Converters

By Tom Lawson, CogniPower, Malvern, Penn.

Conventional pulse width modulation (PWM) control for switched-mode power supplies (SMPSs) must compromise two conflicting goals; stability versus agility. Stability is obtained by employing a substantial output filter capacitor. However, the time lag introduced by such filtering tends to cause overshoot and undershoot in the feedback loop. That tendency toward under- and overcorrection makes oscillation (i.e. instability) a constant concern.

Bill Morong, the principal inventor of Predictive Energy Balancing, likes to use a driving analogy to describe the problem. Delayed feedback is like steering while looking out the rear window. That can work reasonably well on a straight section of road, but on a twisty road you are bound to weave back and forth in drunken fashion, no matter how attentive you are. If you limit yourself to small, slow corrections, you won't weave as much, but you have a better chance of driving entirely off the road. Compensation schemes are strategies to adjust the degree and speed of correction for the circumstances.

By reducing the gain and adding a compensation network, control loops can obtain reasonably fast response, with reasonable stability. This type of compensation can only be optimal under one set of conditions, so the compensation must be compromised to work reasonably well over a range of circumstances. The process can seem more like art than science. Even under best-case conditions, overshoot and undershoot are characteristic when SMPSs respond to rapidly changing conditions.

Various schemes for improving the outcome of the stability/agility compromise fill the literature, but do not dispatch the issue. If the filter capacitance is increased relative to the switched inductance, agility suffers, and the voltage feedback amplitude is reduced, requiring extra gain. Higher gain comes with its own set of problems. Often, these control loops rely on the equivalent series resistance (ESR) of the output filter to obtain a feedback signal, making them sensitive to component substitutions. When the switched inductance is relatively large, the inability to instantaneously change the inductive current confounds control. All of this is a long way of saying that there is no generalized control solution for PWM.

CogniPower has developed and patented new ways to manage switched-mode power supplies that side step the PWM control problem entirely. Through energy prediction, the phase lag of the output filter is removed from the feedback path. This is analogous to looking out the front of the car, instead of out the back. Without the delay in the feedback, you can drive straight down the center of your lane. The underlying principle is that the voltage on the filter capacitor after the inductive energy from the switched inductor has been transferred can be calculated in advance. Given that information, the decision to switch from energizing the inductor to transferring inductive energy to the output can be made on the basis of the energy outcome at the end of the control cycle. That simple concept removes constraints that have long limited the performance of power converters. It constitutes the basis for a control technique known as Predictive Energy Balancing.

Underlying Mathematics

The mathematics for energy balancing in any type of SMPS can be derived from a few fundamental formulas. The kinetic energy held in an inductor, L, is

$$KEL = (I^2 \times L) / 2$$

where KEL is inductive energy in joules, I is current in amps, and L is switched inductance in henries.

The kinetic energy held in a capacitor, C, is:

$$KEC = (V^2 \times C) / 2$$

where KEC is capacitive energy in joules, V is voltage in volts, and C is filter capacitance in farads.

In general, the inductive energy term represents the supply and the capacitive energy term represents the demand. The demand is the difference between the instantaneous capacitive energy and the desired capacitive energy. At the regulation voltage, Reg, the energy held in the output filter capacitor would be:

$$K_{EReg} = (Reg^2 \times C) / 2$$

where K_{EReg} is desired capacitive energy in joules, Reg is Reference point in volts, and C is filter capacitance in farads.

The capacitive energy deficit is then $K_{EReg} - K_{EC}$.

For a simple, discontinuous flyback converter, the energy balance point is the moment in time when the inductive energy is equal to the capacitive energy deficit:

$$K_{EL} = K_{EReg} - K_{EC}$$

In an energy balancing power converter, when the balance point is reached, inductive energizing stops, and the inductive energy is transferred to the output filter capacitor. With the balance properly scaled, after the inductive energy has been transferred, the voltage on the filter capacitor will equal the regulation voltage.

Note that a load current may be discharging the filter capacitor while the inductive energy is being transferred. That effect can cause a slight undercorrection, but does not tend to destabilize the control mechanism. In critical applications, the predicted load energy required during the remaining time before the end of the control cycle can be added to the demand term.

Another possible enhancement involves continuous conduction, where not all the inductive energy is transferred during a single cycle. In that case, the supply term becomes the difference between the instantaneous inductive energy, and the predicted inductive energy at the end of the control cycle. For clarity, here we will limit the discussion to discontinuous mode (DCM) without an explicit load correction.

A Practical Implementation

It might appear that the above discussion is nice enough in theory, but impractical in practice. In fact, the extra circuitry required to implement a predictive converter with near-ideal behavior is entirely manageable. Fig. 1 shows a simple flyback converter in block form. Inputs to the control circuitry are a representation of inductor current from resistor i_{Sense} , voltage reference V_{Ref} , output voltage V_{out} , and a Clock signal to pace synchronous operation.

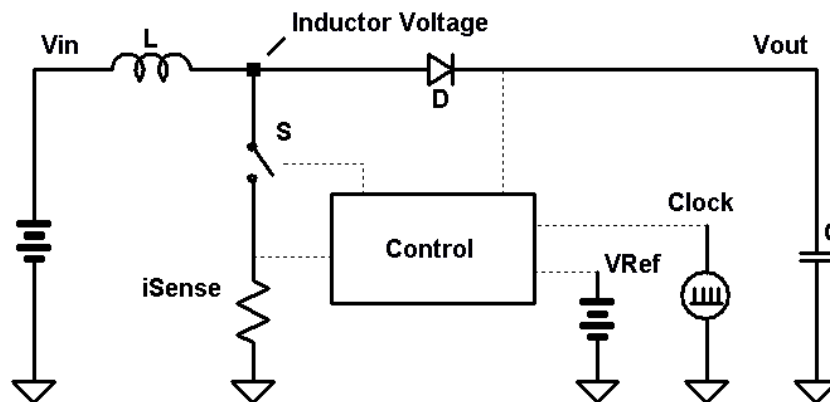


Fig. 1. Simplified block diagram of a flyback converter.

Fig. 2 shows a SPICE representation of the inductive and capacitive energy terms described above, plus the voltage on the switched end of the inductor. (The evaluation hardware represents the energy terms as currents, not voltages, so a screen shot would not be practical here.)

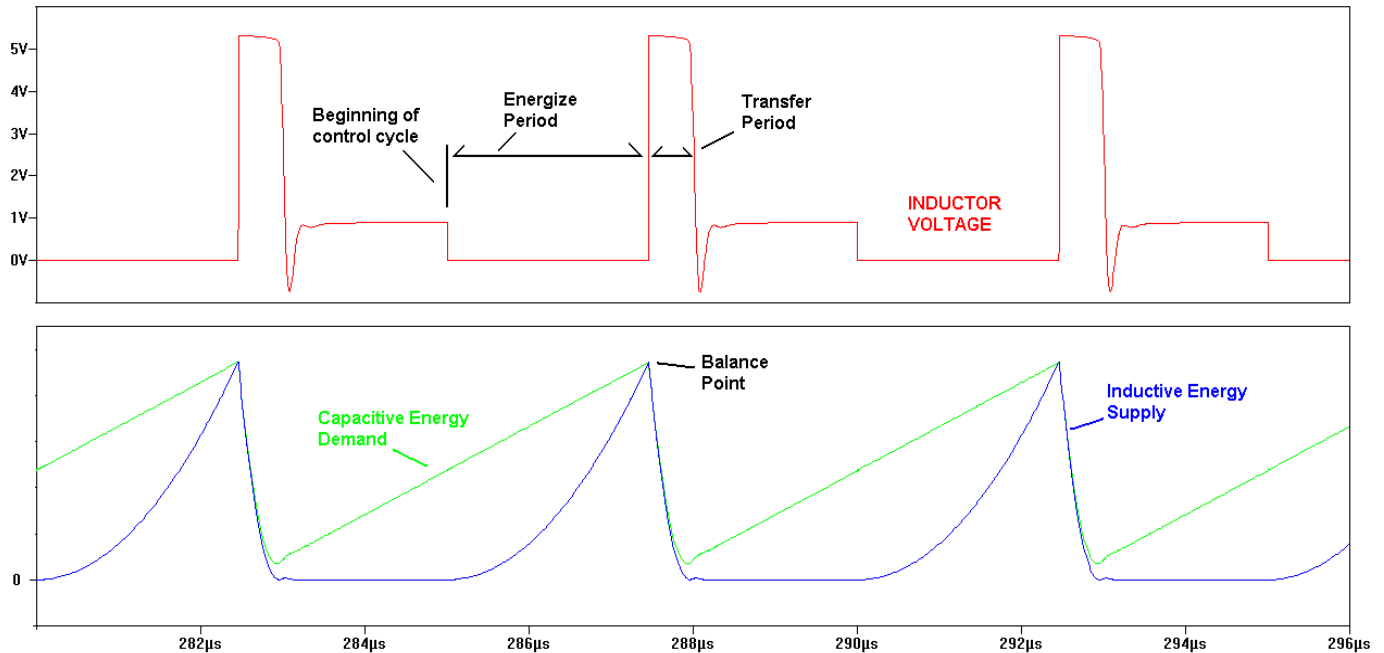


Fig. 2. SPICE waveforms for energy terms associated with flyback circuit in Fig. 1.

You can see the three distinct portions of the control cycle in the Inductor Voltage waveform. At the start of the cycle, as signaled by Clock, switch S closes so that the voltage at the switched end of the inductor falls to very near zero. As long as switch S remains closed, the inductive energy increases. When switch S opens, the Inductor Voltage flies to a diode drop above the output voltage, beginning the energy transfer period. Once the inductive energy has been transferred, no current flows in the diode or in the inductor, so the Inductor Voltage returns to the battery voltage. The situation remains unchanged until the arrival of the next Clock edge.

The Inductive Energy Supply term rises during the energize period, overtaking the more slowly rising Capacitive Energy Demand term. Balance occurs when the traces intersect. Then, the switch opens and flyback commences. The Capacitive Energy Demand term falls to almost exactly zero by the end of the transfer period, as does the Inductive Energy term. In this case, the demand begins to increase immediately after the transfer period because the load is draining the filter capacitor.

The behavior for each subsequent cycle repeats in the same fashion. It is key to note that there is no dependence on the behavior of the previous cycle in this control scheme. Each decision is based on instantaneous conditions. The absence of sensitivity to recent history allows the loop gain to be maximized without introducing a tendency toward oscillation. By providing enough agility to correct for changing conditions in a single cycle, the characteristic alternation between two output levels can be eliminated. The tractability obtainable using these techniques will be a pleasant surprise for those familiar with the usual methods.

Fig. 3 shows a generalized, practical Control Block in more detail. The supply term is derived by squaring the voltage at the inductor current sensing resistor. The size of the sense resistor should be chosen based on the maximum allowed inductor current. The demand term is the difference between the output energy and the regulation point energy, as determined by a differential amplifier.

The amplifier output is scaled by VSCALE to match the sensitivity of the supply term. As long as the inductance and filter capacitance are relatively stable, the L and C terms do not need to be included in the real-time calculation. Here, the L/C ratio is implicit in the scaling factor, VSCALE. If the L/C ratio is actually changing with time or temperature, a slow adaptive control loop can be added to bring the dc voltage to the regulation point without harming the dynamic performance made possible by the predictive control loop.

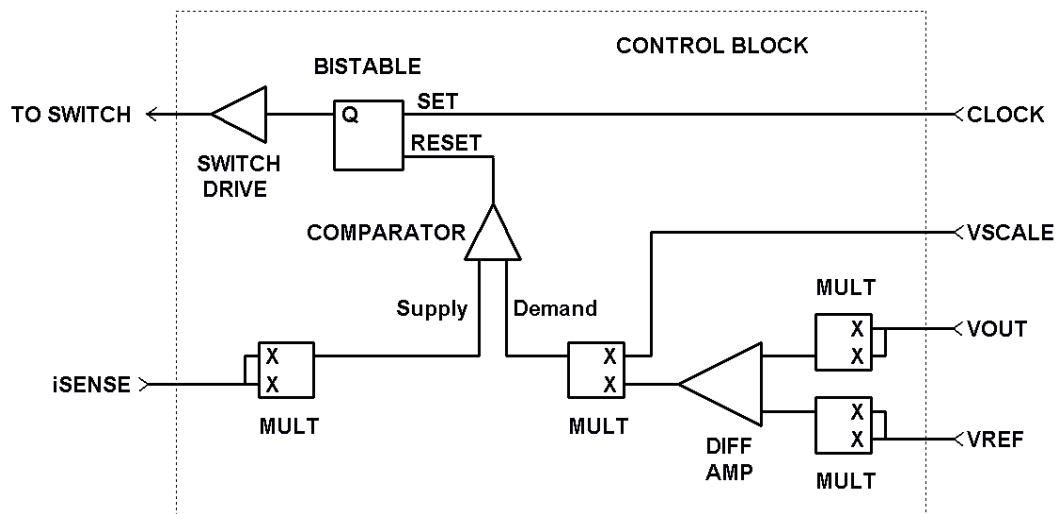


Fig. 3. Generalized control block for Predictive Energy Balancing.

When the supply term matches the demand term, a comparator responds so that the attainment of energy balance resets a flip flop, stopping the energize period. A small amount of additional logic can enforce maximum and minimum energize periods or can provide soft start or overcurrent protection. A switch drive buffer completes the control block.

Experimental Results

There are a number of circuit simplifications that make little or no difference in performance under most circumstances. In an SMPS with a fixed dc output voltage, K_{Reg} is constant, so it can be represented by a fixed voltage or current. Also, when only small voltage excursions appear at the output capacitor, a linear approximation matches the geometric demand term well enough for many purposes.

The CogniPower flyback converter evaluation boards now available use that simplification. Without the multipliers for squaring, the demand term circuitry is essentially the same as in a conventional power converter. If the SMPS was intended to respond to major disruptions with best dynamics, or to follow a changing reference input, the squaring circuitry should be replicated for the demand calculation, as shown here.

To aid observation of dynamic response, the evaluation boards are fitted with a sawtooth load generator which may be connected to the converter output. This sawtooth load repetitively rises from a minimum of about 2 mA to a maximum of about 20 mA over a period of approximately 5 ms. Then, in only 25 μ s, the load falls back to 2 mA. The board is also fitted with a sync test point to facilitate synchronizing an oscilloscope with the load sawtooth. The load waveform is not synchronized with the dc-dc conversion frequency.

In Fig. 4, a screen shot reveals measurements taken on a CogniPower flyback converter evaluation board. The top waveform is the load, while the bottom waveform is the output voltage. Note that the ripple is at the minimum theoretical value for the load, filter capacitance, and clock frequency. That means the smallest possible amount of filter capacitance is needed for a given amount of ripple. Changing the filter capacitance would dictate a change in the scaling setting, but the dynamics will be equally good once the gain is reset to match the new capacitance.

In situations where switched inductance or filter capacitance can change on the fly, adaptive gain can be used to quickly readjust for accurate prediction. In Fig. 4, each converter cycle lasts 15 μ s. Only minimal changes would be required for faster or slower operation. The absence of subharmonic behavior, undershoot, or overshoot, clearly differentiates this control strategy from all other methods.

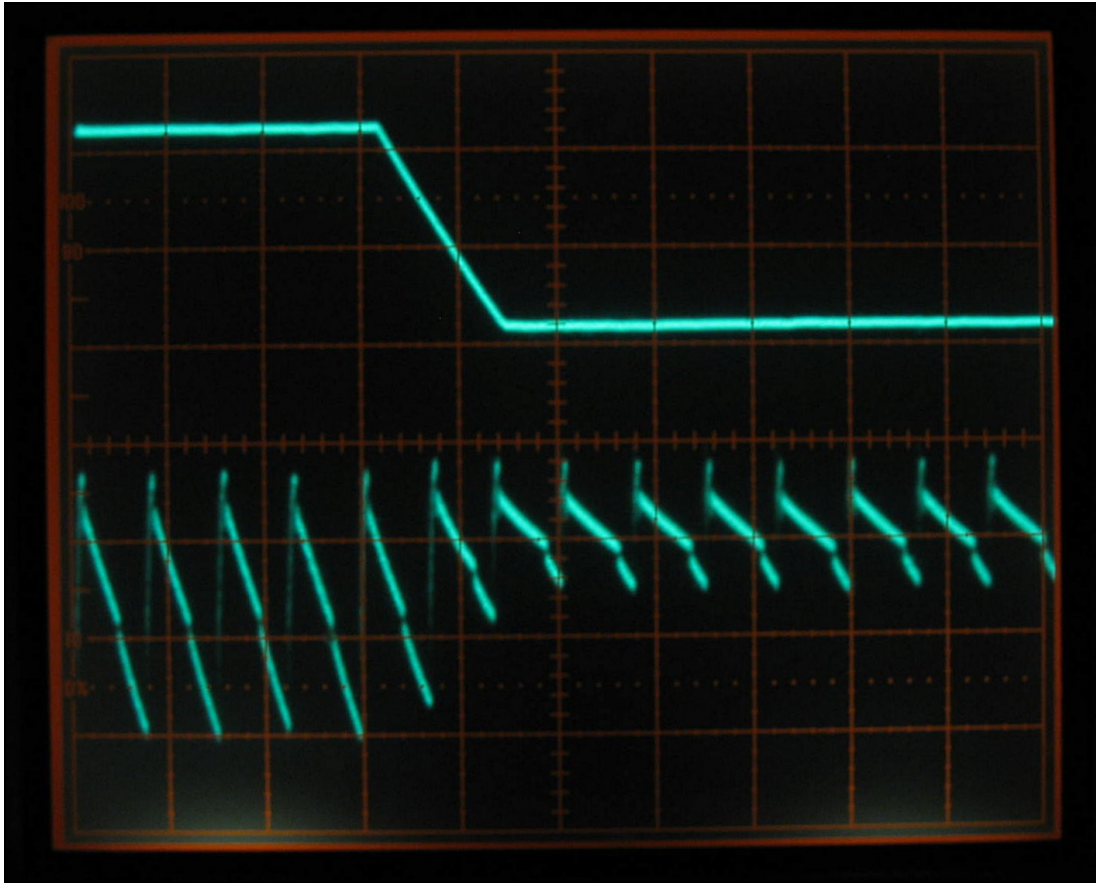


Fig. 4. Load current (upper trace, 10 mA/div.) and output voltage (lower trace, 10 mV/div.) as measured on a Predictive Energy Balancing flyback converter evaluation board. Time scale = 20 μ s/division.

Summary

These controls are scalable and are adaptable over a wide range of applications, from handheld devices to high energy physics. The computational circuitry for energy balancing is very similar for buck or forward converters. By rearranging the circuit blocks, with a few variations, the same principles can be applied to other converter topologies, with similar benefits.

CogniPower converters can respond almost instantly to load changes or to digital control, while most power converters require many cycles to recover from a discontinuity. Digital controls must respect the stability limits of the power converters they oversee. With predictive, single-cycle response, more aggressive power management becomes practical.

Because CogniPower control methods follow the theory and don't require delicate compensation, the same control block can be relied on to maintain its intrinsic stability, dynamic response, and efficiency across a family of applications. That flexibility enables these power converters to be generalized, and more closely integrated with the devices being powered.

In some cases it will make economic sense to integrate the power converter controls with the circuit under power. That close-coupling will bring further improvements in transient response and even greater efficiencies. Overall size reductions become possible. Further, higher frequencies of operation reduce magnetics size and the size and cost of the filter capacitance required. System-on-chip (SoC) devices and chipsets with integrated power conversion functions can be less dependent on external circuit layout. The resulting ease of application is always a plus.

The evaluation board is pictured in Fig. 5. In the flyback converter implemented on this board, the only additional components required for predictive energy balancing are several dual-transistor packages and associated resistors seen here near the center of the board. Even built from discrete components, the total

extra area involved is less than that required for a single screw terminal. There are a limited number of evaluation systems now available.




Fig. 5. Predictive Energy Balancing flyback converter evaluation board.

About The Author



Tom Lawson has been involved with instrumentation since 1968. During the 1970s he worked in medical electronics with Bill Morong, the principal inventor of predictive energy balancing. During the 1980s and 90s he built his own instrumentation company. Since rejoining with Bill Morong, the focus has been power conversion. Lawson started [CogniPower](#) in 2009 to begin the commercialization process. Lawson is named on eight issued patents and five patents pending, spanning four decades.

For more on power supply control methods, see the [How2Power Design Guide](#), select the Advanced Search option, go to Search by Design Guide Category, and select "Control Methods" in the Design Area category.