ISSUE: July 2019

Surprisingly Simple Control Method Slashes Vampire Power

by Tom Lawson, CogniPower, Malvern, PA

Wall-mounted ac-dc converters have become so common they are largely taken for granted. Efficiency and reliability are good enough. They don't use that much electricity because most chargers are 65 W or below to avoid mandated requirements for power factor correction. These small devices may consume a few kilowatt-hours each per year. Even if you could substitute a charger with 100% efficiency, it could not pay for itself out of the power savings of pennies per year.

Still, there are many billions of these power converters in daily use, which means that collectively they consume many gigawatts of power annually. For perspective, think of that as about ten coal-fired power plants operating continuously just to power wall adapters.

The worst of it is that most of those chargers are doing nothing useful most of the time. They stay plugged in when not in use because it is hard to reach them where they tend to be found, down along the baseboard. Further, if they are left plugged in to the device they charge, say a remote speaker, the chargers are running at a very low percent of their rated output just to keep the batteries topped up. Charger efficiency is reduced at lower loads, so they run most of the time at very low efficiency.

What is the best way to approach the problem of that wasted energy, often referred to as vampire power? Mandates are difficult to put in place and are inconvenient. Market forces alone are unlikely to move the needle when the contribution of each charger is so small. CogniPower's answer is make a more efficient charger with much lower standby power that is smaller and cheaper than the best of the current crop. Smaller and cheaper will succeed in the market. Then, increased efficiency and lower vampire power will follow as painless benefits. A new control method known as Demand Pulse Regulation (DPR) makes these objectives possible.

Instead of adding components to improve efficiency, using DPR we remove components and complexity to shrink the size, save power, and lower the cost. The fundamental power converter structure does not change. You still need to rectify the ac input, switch it into a transformer primary, and rectify and filter the output at the transformer secondary. The design challenge in such a power converter revolves around the feedback mechanism and the control circuitry. Because the control circuitry must run all the time, it is mostly responsible for the vampire power, and will be our focus here. This article explains the principles of operation under DPR, its efficiency benefits, experimental results, opportunities for further efficiency gains and details of implementation.

DPR Versus Conventional Feedback

The universal power input range is 85 to 265 Vac, and the dc output range may be 2 to 20 V or more, so the ratio of input to output can vary by a factor of over thirty. The flyback form of power converter is the only form that offers anywhere near that degree of flexibility. Accordingly, these wall adapters are flyback converters.

Conventional feedback is achieved using optocouplers or by reflecting a representation of the output back through the transformer to the primary side. Optocouplers are physically large, power hungry, temperature sensitive, and they change with age and use. If you save power by under-driving them, they get very slow, which introduces more delay into the feedback loop, adding to existing control challenges. Reflected voltages are transient, noisy, subject to loading variations, and dependent on non-idealities, which vary with the details of transformer construction and operation.

Conventional feedback is based on the output voltage, which is filtered by the LC filter formed by the transformer secondary and the output filter capacitor. That filter induces delay, and control loops with delayed feedback tend to oscillate. Most of the large body of literature written about switched mode power converters addresses this control issue. The controller must constantly compromise between speed and stability. There is no correct answer for such a controller, only good enough, or not.

The key to Demand Pulse Regulation (DPR) is to make all control decisions in the simplest fashion and at the optimal location. As will be explained, the feedback becomes an instant in time, which is easily transmitted across an isolation barrier from the secondary side to the primary side. A tiny pulse transformer is our preferred means of coupling that instant of time to the primary side. The key functional elements of a DPR power converter are shown in the simplified schematic in Fig. 1.

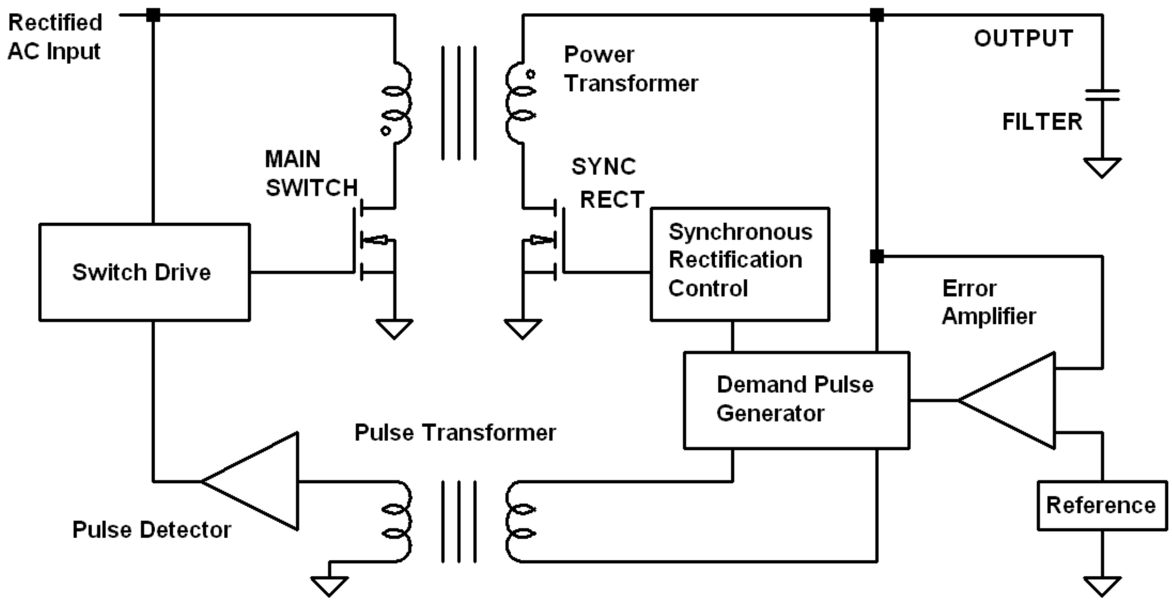

Fig. 1. Simplified block diagram of a DPR power converter.

An error amplifier on the secondary side produces a signal proportional to the difference between the desired output, represented by the reference, and the actual output. That error signal becomes the input for a Demand Pulse Generator circuit. When the output falls below the reference, a demand pulse is sent through the pulse transformer to the primary side of the isolation barrier. That pulse is detected and causes the primary-side switch driver to turn on the main switch.

The decision regarding when to turn off the switch is made on the primary side. That decision is based on switch current and/or time. Typically, the main switch turns on until a particular switch current is reached, or for a maximum length of time. The switch current threshold for turning off the switch may be dynamically adjusted on the primary side to achieve a switching frequency in a desired range.

In summary, the decision of when to turn on the switch is made on the secondary side, by a simple comparison of reference and output, while the decision of when to turn off the switch is made on the primary side, based on local conditions at the main switch. That partitioning of the control circuitry enables the lower standby power and higher efficiency of DPR.

Because the feedback in DPR is based on the instantaneous output on a cycle-by-cycle basis, the filter pole does not contribute delay to the feedback loop, so there is no tendency to oscillate during discontinuous operation. To be clear, in what is called discontinuous mode, the current in the transformer falls to zero once each switching cycle. In continuous mode, the transformer current rises and falls each cycle, but never falls to zero. More about continuous mode will follow later.

Cycle-by-cycle control provides the fastest possible response to changes in load. In discontinuous operation, overshoot is limited to any excess energy contained in a single cycle, and undershoot results in the converter operating at its maximum allowed switching rate. Because there is no compensated feedback loop, there is no compromise needed to accommodate changes in output voltage, current, or capacitive loading. Cycle-by-cycle control with regulation decisions made on the secondary side offers another big advantage.

USB chargers, for example, need to interact with their load to change charging voltage or current. The simplest and easiest way to add that flexibility is to have the decision of when to turn on the main switch made on the secondary side. Simple and easy translates to less circuitry, which means smaller and lower power. USB-PD and Quick Charge standards are two examples. Any present or future digital protocol can be implemented with a minimum of added complexity given secondary-side control. Switching waveforms for a DPR power converter are shown in Fig. 2.

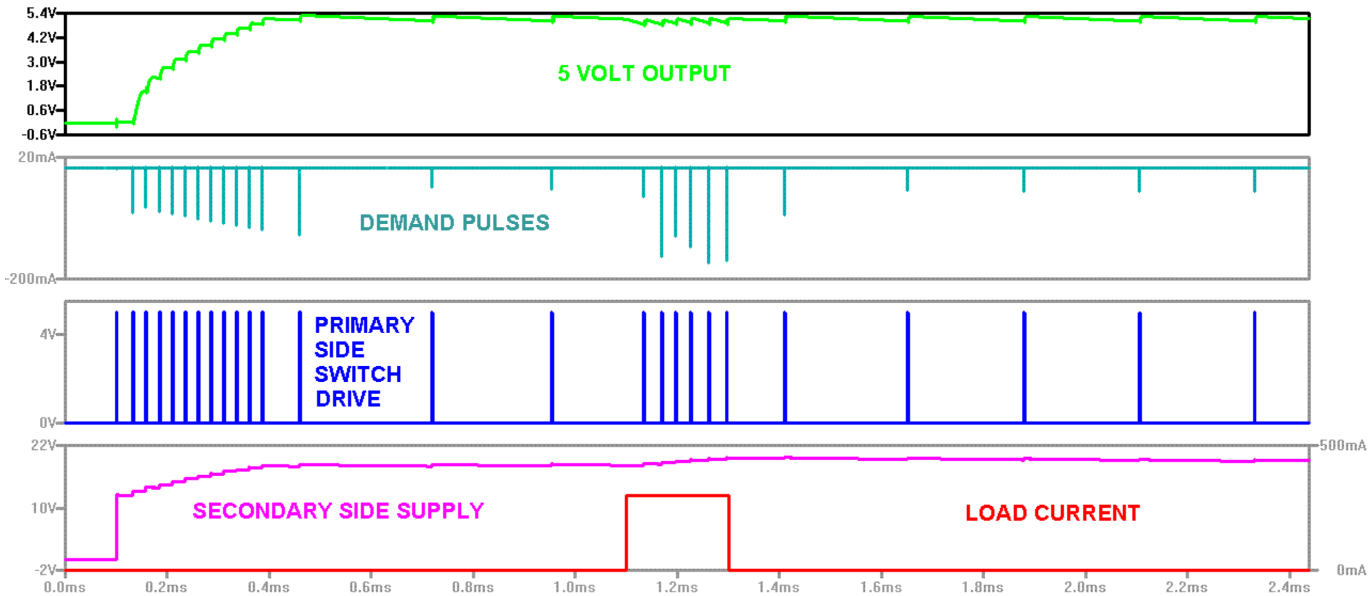

Fig. 2. Simulated waveforms of a DPR power converter during start up and operation.

When the regulation is controlled on the secondary side, a chicken-and-egg situation can be created at startup. Care is required to assure that all possible combinations of ac and dc power and load are handled gracefully. In Fig. 2, the primary side sends one start up pulse to begin operation. That pulse powers the secondary side which then takes over the task of deciding when to turn on the main switch.

In Fig. 2, note that only the first main switch drive cycle does not correspond to a demand pulse sent from the secondary side. That startup pulse has little effect on the output voltage in green. Instead, the energy from the startup cycle goes to bring up the secondary side supply shown in pink. Once that supply is up, the error amplifier, reference, and demand pulse generator begin to operate. Subsequent cycles of the main switch are seen to be triggered by demand pulses.

The extra power required to charge the output filter capacitor at start up needs to be limited. Here, the demand pulse generator runs at its maximum allowed rate while the output voltage ramps to the regulation point. If softer starting is desired, the reference can be brought up more slowly, allowing the converter to stay in regulation during startup.

In Fig. 2, the converter is lightly loaded to begin, so after the output reaches the regulation point, demand pulses slow to a low rate. When the load is applied (red trace) the demand pulses immediately increase in frequency, causing the output voltage to be maintained at the desired voltage. When the load is removed, demand pulses fall back to the slower rate.

Note that the secondary side supply is not regulated. Regulation is not needed here because the error amplifier, reference, and demand pulse generator can operate over a wide range of secondary-side supply voltage. If digital circuitry were required on the secondary side, a 3.3- or 5-V regulator would be added.

Experimental Results And Performance Potential

In Fig. 3, a prototype of a DPR power converter is presented. This converter was designed for under 1 W at full load to illustrate that high efficiencies could be obtained even on the smallest of power supplies running at a fraction of full load. To achieve that efficiency, standby power must be reduced to a minimum.

On the bench, this power converter delivered 95.6% efficiency at 12.5% of full load. That was six years ago. Since then, we have made further improvements and can deliver efficiencies over 96% even at just 1% of full load. Standby power can now be under 400 µW, almost 10 times lower than the prototype pictured.

If needed, with additional care, that standby power can be reduced even further. You might observe that the power transformer is rather large for such a small supply. The limitation is ease of hand-winding. Much smaller transformer cores would be used in an automated production environment.

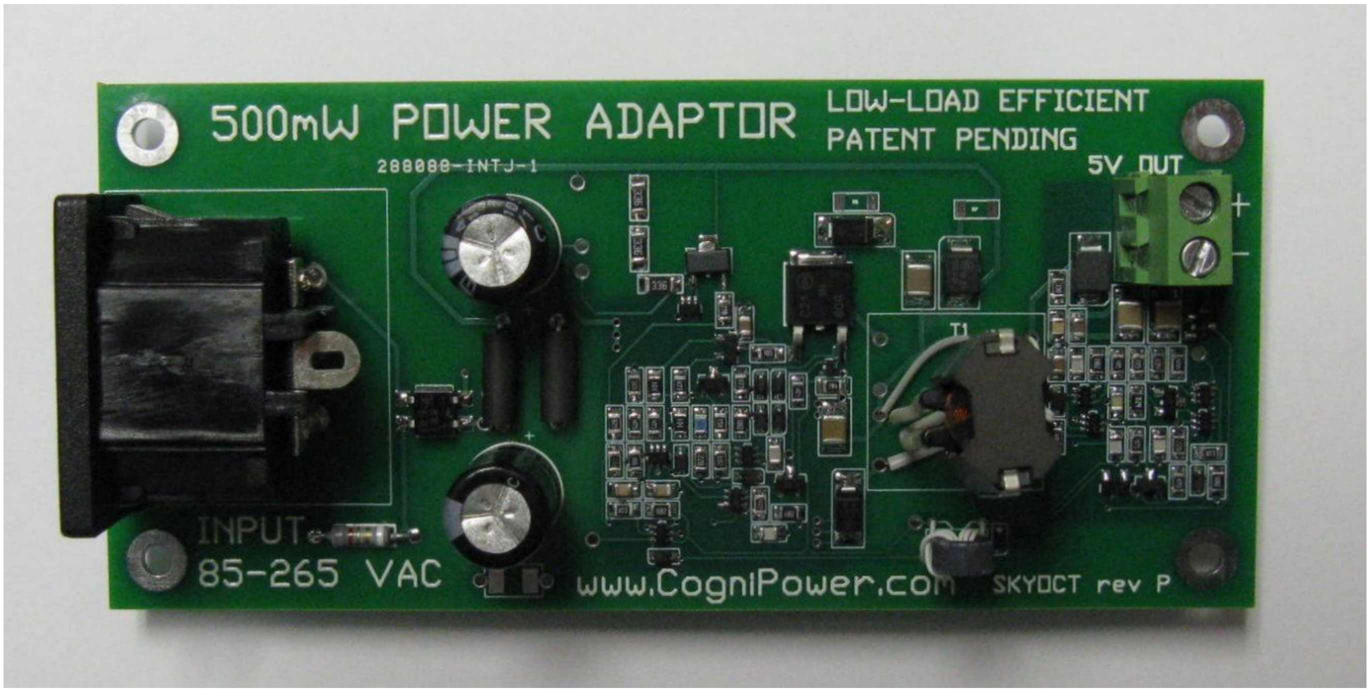

Fig. 3. Prototype DPR power converter measured at over 95% efficiency.

Putting the standby power aside for the moment, there are two other ways that DPR helps to reduce losses for best efficiency. First is power transformer construction. Because DPR does not need a reflection of the output voltage on the primary side, the power transformer design can be optimized for good coupling between primary and secondary. That good coupling reduces the need for snubbing, which reduces what can be a significant power loss.

Second, by generating the signal that turns on the main switch from the secondary side, the synchronous rectification control has the advantage of knowing when a cycle is about to begin, instead of needing to detect that cycle and responding reactively. Avoiding just a few nanoseconds delay in driving the synchronous rectifier switch can prevent losses that would otherwise show plainly as decreased efficiency, particularly at lower output voltages.

The signal transformer for transmitting the demand pulses across to the primary side is seen below the power transformer in Fig. 3. Though small, and with few turns of wire, this signal transformer still occupied volume and involved some cost. We have since improved the demand pulse generator so that the smallest practical toroid with a single stitch of wire each for primary and secondary is sufficient to transmit demand pulses.

One of the windings uses wire with high-voltage insulation, and the spacing between the primary and secondary terminations can be set to obtain the desired isolation voltage. Such a signal transformer is faster, lower power, smaller, cheaper and more reliable than an optocoupler.

DPR does not much alter the requirements for a primary-side controller except to remove the circuitry responsible for regulation. The primary-side controller has three tasks to perform: It must efficiently turn on the main switch in response to a demand pulse. It must turn off the main switch after a certain current develops in the switch or after a certain interval, and it must send a pulse, or pulses, of energy to power the secondary side if a set period of time passes without a demand pulse. The primary-side controller can also keep track of the recent repetition rate of demand pulses and make occasional adjustments of the main switch on-time to steer the converter toward an advantageous frequency of operation.

The secondary-side controller of a DPR power converter handles most of the job. It is worth going to some lengths to reduce the power required on the secondary side because the advantages of reduced power multiply. Reducing secondary power brings easier startup, lowers the minimum cycle rate for no-load conditions, and reduces the size of the secondary-side local power supply filter capacitor. The secondary side need consume only enough power to preserve the function of the error amplifier, and it must store enough capacitive energy to generate a demand pulse when the error amplifier indicates a need.

For ultra-low standby power, a hibernate mode could be added where the error amp and reference are starved for power. However, if a few hundred micro watts of standby power is acceptable, there is no need for the complication of this extra mode of operation.

Secondary-Side Circuitry

We prototype the secondary side in bipolar transistors because they are easily available in pairs that match. For a commercial implementation, presumably MOS transistors would be used. Matching would not be an issue in an integrated circuit. Also a MOSFET version of the circuit could be even lower power than the bipolar version shown here.

Fig. 4 shows a DPR secondary side in actualized form. Starting with the upper half of the drawing, M is the secondary of the power transformer, Rect is the rectifier, and F is the filter capacitor. The rectifier is shown here as just a diode for clarity. In fact, at higher output voltages, say 20 V, 95% efficiencies can be obtained without the need for synchronous rectification. D in Fig. 4 (for differential amp) is the error amplifier comparing the output to a reference (Ref), shown here as a reference diode. A clamp diode for overvoltage protection is also shown, but it is not essential.

The output of the error amplifier is the input signal for the demand pulse generator, labeled DPG. The feedback signal is labeled C for control. P is the primary of the signal transformer. The demand pulse generator is a proprietary circuit that may be somewhat nonintuitive. That circuit provides multiple advantages, and is the key to simplicity and low standby power. L is the unregulated local power supply that was shown in pink in Fig. 2, filtered by the capacitor on the left. The rightmost two capacitors provide positive feedback for the demand pulse generator. The capacitor on the right is the timing capacitor that sets the maximum frequency of operation.

Much below the regulation point the DPG runs as an oscillator at the maximum frequency. When the timing capacitor charges enough to begin to turn on the upper transistor, positive feedback through the capacitor on the left causes both transistors to turn on hard and fast. Both capacitors in the DPG then discharge rapidly through the pulse transformer primary producing an extremely rapid edge which propagates easily through the signal transformer forming the demand pulse seen by the primary-side detector. After that discharge, the timing capacitor again starts to charge, beginning the next cycle.

If the output approaches the regulation point, the control voltage rises, causing the transistor at point C to divert charging current from the timing capacitor. That slows down the charging rate which lowers the frequency of demand pulse generation. At, and above, the regulation point, all the charging current is diverted so the timing capacitor never charges and the DPG never triggers.

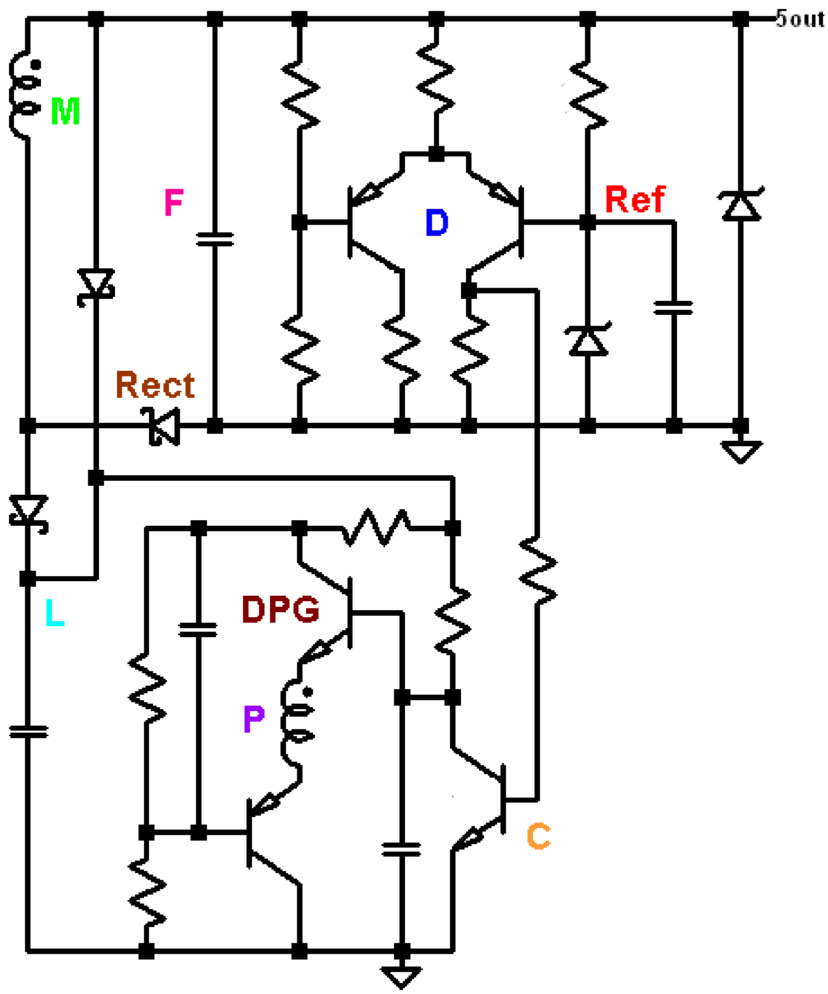

Fig. 4. Detail of actualized secondary-side circuitry.

This demand pulse generator circuit provides more function than is immediately evident. It sets the maximum switching frequency. It operates in a linear fashion near the set point for smooth regulation. It obviates the need for a regulated local power supply. It can run on as little as 2 V and as much as the voltage ratings for the transistors. It produces a very high edge rate pulse which allows the use of a minimal signal transformer. And, it does all that while consuming a vanishingly small amount of power.

As shown in Fig. 5, running at 500 Hz, the demand pulse generator consumes 1.25 μW of power. The demand pulses in Fig. 5 peak at 10 mA and persist for about 40 ns. The blue trace in the upper axis of Fig. 5 is the voltage at the top of the pair of transistors in the DPG. That voltage can be used as a logic input to optimize the timing for a synchronous rectifier.

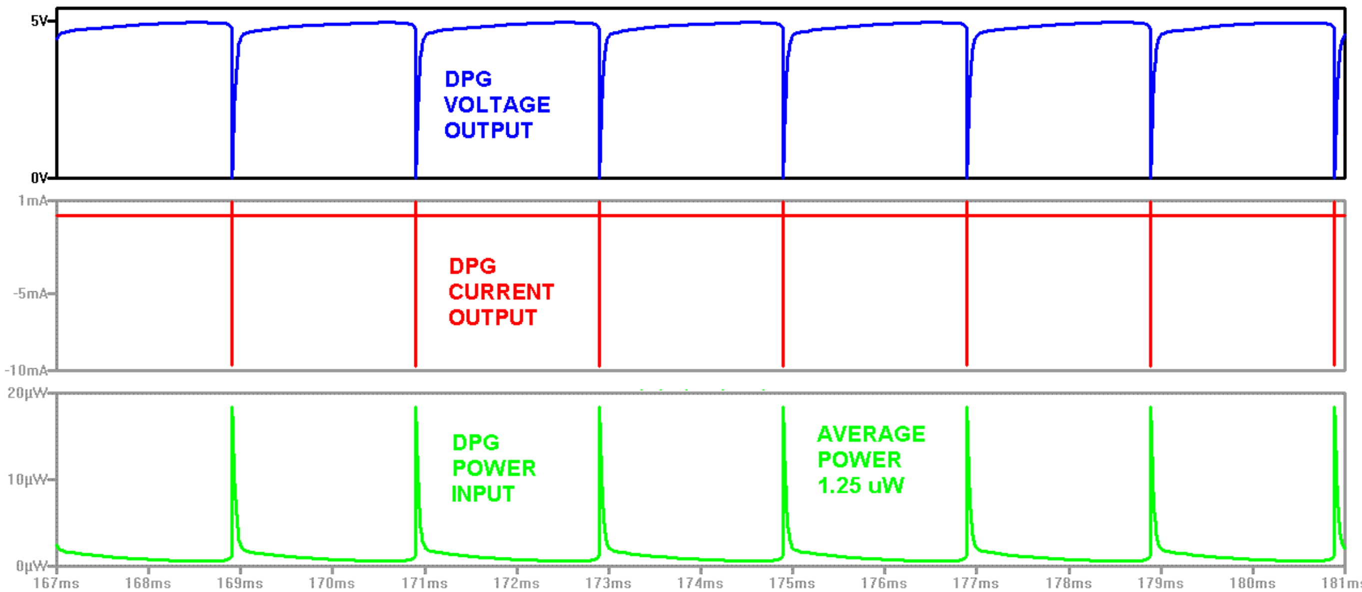

Fig. 5. Demand Pulse Generator waveforms illustrating 1.25 µW of power consumption.

Solving The Standby Power Problem

The DPG solves a problem that is easy to overlook, but one that is responsible for the relatively high standby power of many circuits which combine analog and digital functions. Normally you can think of a logic input to a MOS gate as a high-impedance point. The power consumption of the MOS is modeled as a capacitor and is calculated based on the frequency of operation.

With the input voltage near the plus or minus rail, a MOS logic input draws almost no current. But, at a point when the logic input is between a valid high and a valid low voltage, both FETs in the MOS input structure are on at once, and much more than the usual current is drawn. A 74AS logic gate that draws a few microwatts under static conditions can draw 25 mW at the transition point. (See reference [1].)

If slowly changing analog signals are used as logic inputs or clock inputs, the standby power will go up by orders of magnitude. The DPG circuit shown in Fig. 4 can accept a very slowly changing analog input signal and respond with a very fast digital edge when a threshold is reached without drawing a heavy current near the transition point. That circuit can be used to make a low-power hysteretic buffer with general utility. It can be used to accelerate a slow clock signal or to condition a sensor output to generate a digital signal. Type D flip flops have a maximum clock transition specification. Placing this buffer at the clock input would remove that limitation.

Any efficient ac-dc converter requires switched-mode power to supply the primary-side controller. That switched-mode converter requires a regulator. Such a regulator would consist of an error amplifier gating a power switch, very much like the secondary-side circuitry described above, but without the need for isolation. A conventional regulator's efficiency suffers from the extra power required by its logic when inputs pass slowly through the transition zone. That extra power may be ignored in a regulator drawing many milliwatts of standby power, but not if microwatt performance is needed. So, a variation of the Demand Pulse Generator of Fig. 4 can be used to improve efficiency in other power converters, as well.

In order to move more power while avoiding high peak currents, continuous mode is essential. In continuous mode, the transformer's inductance acting with the output filter capacitance sustains a current that persists, even though the main switch is still turning on and off. When the power converter shifts into continuous mode, the relationship between the on-time of the main switch and the power transferred to the output changes.

Current in the primary winding goes up linearly with on-time, but the energy stored in the power transformer goes up by the square of the current.

If the on-time is sufficient to increase the primary current by 1 A, in continuous mode, the increase might be from 1 A to 2 A. That change places four times more energy in the power transformer than the discontinuous case where the primary current increases from 0 A to 1 A. For that reason, conventional controls change gain more than a little when going into or out of continuous mode. Gain changes on the fly can be challenging to implement.

DPR doesn't need to change gain in continuous mode because of the cycle-by-cycle control, but it does need to slow down to match the actual response of the converter. The continuous current has the effect of straddling cycles, and that effect must be handled properly to prevent subharmonic behavior. If the converter in continuous mode overshoots, demand pulses stop, but the continuous current goes on flowing into the filter capacitor. That means overshoot may be prolonged while continuous current is flowing.

When the output does fall below the regulation point, the continuous current may be gone, so it will take multiple cycles to re-establish continuous operation. That will tend to cause undershoot. The result of that delayed response is a tendency toward oscillating around the set point, a pattern that is all too familiar for users of conventional compensated controls.

The answer is to slow the application of the feedback so that the converter approaches the regulation point exponentially. The DPG will do the exact correct thing if the action of the control voltage is limited. The DPG then responds in a linear fashion as the control point is approached. Demand pulses are spread out or bunched closer instead of just turned off and on. Simply increasing the value of the base resistor on the control transistor adjusts for proper behavior when going into or coming out of continuous mode.

More Details Of Operation And Implementation

At very low loads, the operating frequency can be at the low end of the audible frequency range, but so little power is moving that it cannot be heard. At higher power, the converter will run above the audible range. There may be intermediate regions where audible noise is an issue. The primary side controller can take care of that potential problem by retaining a short history of the operating frequency. That history can be kept in analog fashion as a voltage on a capacitor, or digitally, as data. If the operating frequency is in an undesirable range, the amount of energy loaded into the transformer with each cycle can be changed. The result will be a higher or lower operating frequency.

Because the regulation scheme is simple and uncompensated and driven from the secondary side, digital control is as simple as changing the reference voltage to change the output voltage. Regulating output current in place of voltage is no different for DPR than for standard practice. Most of the usual things that can be done to get incremental efficiency improvements in conventional power converters can be done with DPR. These include active clamping, quasi-resonance, the use of HEMT switches, etc.

Our 95.6% bench-measured efficiency was achieved without going to those lengths, and, was achieved using MOSFETs dating from 15 years ago. Modern MOSFETS have lower on-resistance and reduced gate charge and reduced leakage and parasitics. All of that leads to higher efficiencies with newer parts. Higher efficiencies result in lower operating temperatures, providing a reliability boost. With less heat to remove, physical size can be reduced, as well.

In order to achieve the efficiencies we describe, more than passing attention must be paid to circuit layout and transformer construction. Stray inductance is determined mostly by the geometry of the layout and should be minimized. When currents are switched on and off, stray inductance stores energy which, when released, appears as voltage spikes. If spikes must be snubbed, losses are incurred.

Newer, higher-voltage FETS can be used to avoid the need for snubbing. Uncoupled inductance in the transformer is a result of the details of transformer construction. Minimizing stray and uncoupled inductance adds design effort, but has little or no effect on unit cost. The added design effort pays off many times over in a product produced in high volumes.

Summary

The problem of vampire power extends beyond wall chargers. Literally billions of small ac-dc power converters are put into service each year. These are not only battery chargers, but also IoT power, LED lighting, and power for communications, computer standby supplies and consumer goods. Any line-connected product that can be run from a remote control needs to have a power supply that functions when the device is nominally "off". Any appliance with remote access features needs an always-on power supply.

An increasing variety of other equipment has always-on USB charging ports built in, including printers, plug strips, and even the wall sockets themselves. A sobering comparison tells us that in the U.S., we now waste more electricity as vampire power than the total amount of electricity generated here in 1950.[2]

DPR offers several efficiency advantages including ultra-low vampire power and that ultra-low standby power enables unmatched low-load efficiency. Another benefit is that advanced notice from secondary-side control for when the primary-side switch is about to turn on enables proactive, instead of reactive, synchronous rectification for higher efficiency at higher loads. Secondary-side control also minimizes the need for extra circuitry for implementation of present and future digital protocols.

By not relying on the power transformer to sense the output voltage as reflected to the primary side, power transformer design is uncompromised. This enables better coupling, and therefore less need for dissipative snubbing. Overall, the simple and efficient design achieved by DPR means less waste heat, smaller size, higher reliability, and crucially, lower cost. That last point, lower cost, makes this proprietary new technology particularly compelling.

References

1. "Solving CMOS Transition Rate Issues Using Schmitt Triggers" by Shreyas Rao, TI White Paper SLLA364A, April 2017 revised May 2017.
2. U.S. Energy Information Administration, Monthly Energy Review, Table 7.2A, March 2019.

About The Author

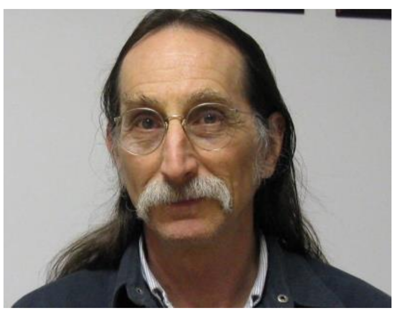

Tom Lawson has been involved with electronics since 1968. During the 1970s he worked in medical electronics with Bill Morong, the principal inventor of Demand Pulse Regulation. During the 1980s and 90s he built his own instrumentation company. Since rejoining with Bill Morong, Lawson's focus has been on power conversion. Lawson started CogniPower in 2009 to begin the commercialization process. He is named on eighteen issued patents with more pending, spanning four decades. Lawson can be reached at tlawson@cognipower.com.

For further reading on power supply control methods, see the How2Power Design Guide, locate the Design Area category and select "Control Methods".