Power Factor Correction: Part I

by Dale Eagar

Introduction

The term "power-factor correction" (PFC) refers to the reduction of the harmonic content, and/or the aligning of the phase angle of incoming current so that it is in phase with the line voltage required to operate an electronic device. PFC is considered very beneficial to the environment because it makes more efficient use of existing power plants. PFC is subject to legislation and policy making throughout the industrialized world.

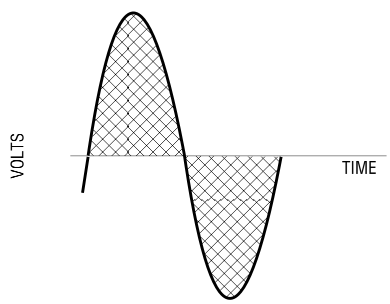

Figure 1a. Input line voltage

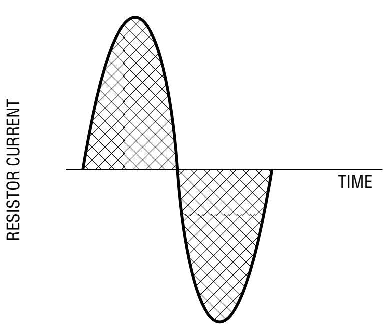

Figure 1b. Current drawn by a pure resistive load

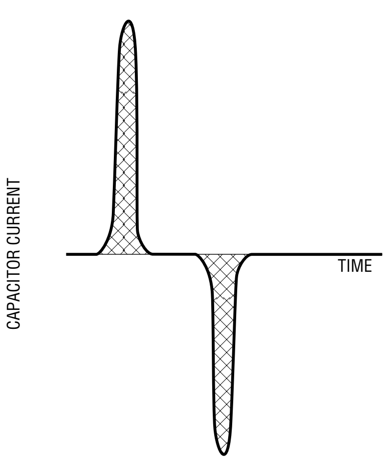

Figure 1c. Current drawn by a capacitive-input power supply

The European standard (IEC 555) sets maximum permissible values for the harmonics of the input line current that may be produced by equipment meeting these standards. By 1996, TV sets and other consumer equipment will be required to incorporate PFC. The benefit of PFC is realized as energy savings seen throughout the power distribution system. With PFC implemented throughout the industry, fewer new power plants will need to be built to meet projected energy demands. Consumers will pay more for "poor power factor" power at the power meter; it is hoped that, as a result, they will choose the beneficial "green" or PFC-equipped devices to save energy dollars.

Equipment that uses DC voltages derived from the AC line generally have a poor power factor because of the capacitive input to the DC power section. The waveforms in Figure 1 show the "evils" of capacitor input power supplies. Figure 1a represents the input line voltage; Figure 1b represents a "nice" waveform of current as drawn by a resistive load; Figure 1c represents the harmonic-rich current waveform drawn by a capacitive-input power supply.

How PFC Performs Its Magic

The trick behind PFC is simple: make the input look as much like a resistor as possible. Resistors have the perfect power factor (unity). From the power utility company's viewpoint, unity power factor is the load of choice, a load that allows their power distribution system to operate at its maximum efficiency.

Emulating a Resistor

A resistor is emulated at the input port of a PFC by loading the incoming power line with a programmable current sink that is programmed with a voltage proportional to the instantaneous line voltage (Figure 2a.) The programmable current sink is the input characteristic of a "lossless energy converter" (detailed in Figure 2b). The energy converter intercepts instantaneous power from the power line, which is the product of the instantaneous voltage and the instantaneous current entering the energy converter. All energy intercepted by the energy converter is delivered to the load device.

Although the devices detailed in Figures 2a and 2b emulate resistors, they provide no means of controlling the overall level of power intercepted from the power line. The circuit in Figure 2c allows for variation in both line voltage and load power. The "load device" detailed in Figures 2b and 2c is invariably a low AC impedance device; such devices include, but are not limited to, capacitors, batteries, and voltage sources.

The overall goal of PFC is to transfer power from a "wiggly" source such as an AC power line to a relatively benign DC voltage. This task must be performed without stuffing a bunch of harmonic junk back on the AC power line.

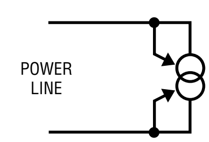

Figure 2a. Programmable current sink

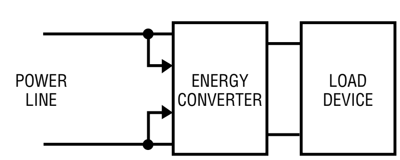

Figure 2b. Non-programmable energy converter (aka PFC)

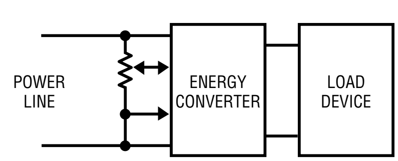

Figure 2c. Programmable energy converter (aka PFC)

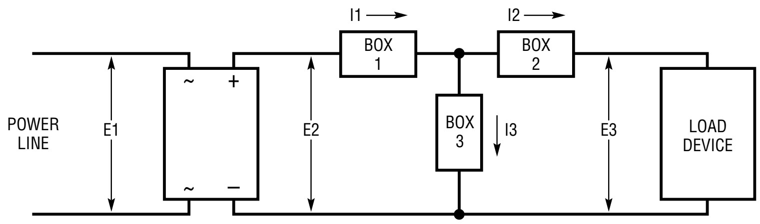

Figure 3. Detailed block diagram of energy converter

The Energy Converter Box

How it Works

The energy converter shown in Figure 2b obeys the laws of conservation of energy (as we all must). As the energy intercepted at the input is transformed from one voltage to another, the current is also transformed from one value to another. The energy stays the same.

The "guts" inside the energy converter block of Figure 2b are further detailed in Figure 3. Regardless of the circuitry in boxes 1–3, we can be sure that Kirkhoff will have his way: I1 + I2 + I3 = 0. Further, we shall assume that whatever occupies the three boxes is lossless (a pretty good assumption in a PFC with better than 95% efficiency). Since the energy intercepted from the input power line E1 cannot be dissipated in the lossless contents of Boxes 1, 2, and 3, it will be losslessly transferred to the output E3.

Figure 4 illustrates the waveforms of a 300W, 120V-to-382VDC power conditioner (refer also to Figure 3). Figure 4a shows E1, the input power line voltage of $120V_{RMS}$. Figure 4b details E2, the full-wave-rectified sine wave.

The energy converter does magic things in the three boxes to cause the waveshape of the input current I1 (Figure 4c) to be a replica of the input voltage E2, with only the magnitude being different. The power intercepted from the input is P1:

P1 = I1 × E2 (see Figure 4d).

Note that the input power is a sinusoidal waveshape, is always positive, and is at twice the line voltage frequency. This is exactly what the waveshape and frequency of the power delivered from a sine wave source to a resistive load looks like.

All of the power intercepted by the energy converter circuit is transferred

continued on page 21

PFC: (Power Factor Correction)— The process used to make capacitors look like resistors. PFC became popular in the early 1990's when the earthlings realized that about 10% of the power they harnessed on their planet was being converted to heat. This heat, which was dissipated through their power distribution network, become a contributing factor in their global warming trend. (see *History of the Sol System*, Vol. 17, pp. 137,657–137,698.)

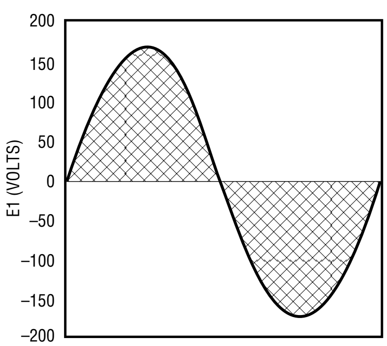

Figure 4a. E1, input voltage waveform: 300W idealized PFC

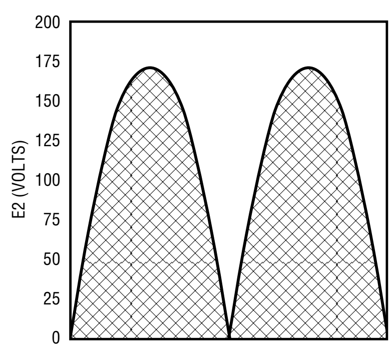

Figure 4b. E2, full-wave-rectified sinewave: 300W idealized PFC

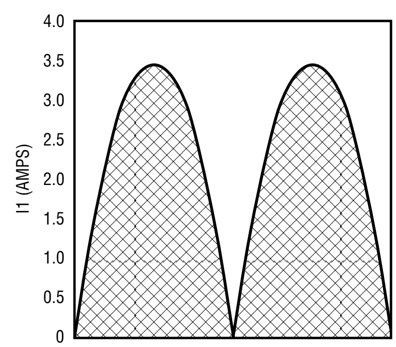

Figure 4c. I1, input current: 300W idealized PFC

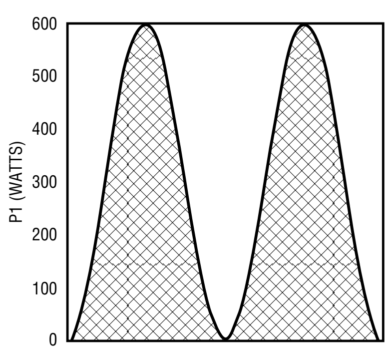

Figure 4d. P1, power intercepted from the input (P1 = I1 × E2): 300W idealized PFC

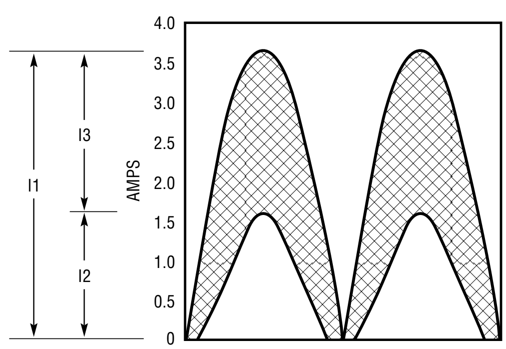

Figure 4e. Input current, I1, and output current, I2. I3 is obtained by subtraction. 300W idealized PFC

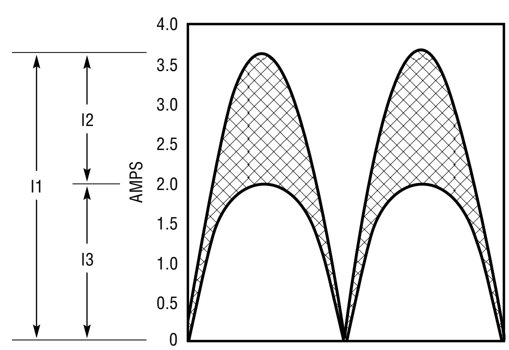

Figure 4f. Input current, I1, and current through Box 3, I3. I2 is obtained by subtraction. 300W idealized PFC

Circuit-Board Layout Considerations

All the capacitors in the decoupling network should be installed on power- and ground-plane areas on the top side of the board. An absolute minimum of one feedthrough per end for each capacitor into the internal power and ground plane should be used. It is preferable to use two feedthroughs per capacitor end (64 total). Any more than that proves to be of no benefit, but at 30 total, expect about a 2mV increase in transient droop. This is about a 5% degradation in performance. Decoupling capacitors should be connected with planes rather than traces, since the traces will be far too inductive. The total network ESR must be less than 6.5 milliohms and the total ESL less than 0.07 nanohenry for the P54C-VR.

Input Capacitance

Another important consideration is the amount of capacitance on the power supply input. The ripple-current rating must be high enough to handle the regulator input ripple. In addition, this capacitance will decouple the load transients from the 5V supply. If insufficient capacitance is used, the disturbance on the 5V supply will exceed the 5% specification for the TTL logic powered by this voltage. Because the magnitude of this disturbance is quite dependent upon the nature of the 5V power supply and because the performance of these supplies varies widely, it is difficult to say just how much capacitance is needed. In general, however, if enough capacitance is present to handle the ripple current, the disturbance on the 5V supply will be acceptable. Good transient response on the 5V supply translates to a need for less input capacitance. If sufficient bulk capacitance is present on the motherboard for the 5V supply, less additional capacitance will be required on the processor supply input. As a minimum, there should be at least one low-ESR capacitor within an inch of the regulator. Be careful to look at the level of disturbance on the 5V supply to make sure it remains within specifications.

Powering the P54C

The same basic circuit is used for both the 5 Amp and the 10 Amp designs. The necessary substitutions are shown on the schematic (Figure 1). If 12V is available to power the LTC1266, the bootstrap capacitors and diodes can be eliminated. The 12V solution is preferred, as it is simpler and somewhat more efficient. If no 12V is available, use the bootstrap circuit. Note also that different MOSFETs are specified for the 5 Amp and 10 Amp circuits. The Si4410 is a new part from Siliconix, which offers less than half the on-resistance of the Si9410 used in the 5 Amp circuit.

High-Accuracy Solution—Basics of Operation

The solution for the P54C-VR and -VRE shown in Figure 2 relies on the accuracy of the LT1431. The internal reference is specified at 2.5V ±0.4% (worst case) at 25°C. The LT1431 consists of a precision reference and a wide-bandwidth amplifier with an open-collector output. The feedback divider is set to place the reference input pin at 2.5V with the desired output present. The 2.5V is further divided to 1.15V to drive the LTC1266s VFB pin. In a normal application, this pin will servo to 1.25V. Hence, the LTC1266 sees the output as being too low and forces its internal error amplifier to the positive rail, which, in this case, is 2.0V. This output shows up as a current out of the I_{TH} pin. The open collector of the LT1431 draws enough current from this pin to set the output of the supply at the desired voltage. Since this constitutes a high-gain servo loop, the output is regulated very accurately. Loop compensation is accomplished by R5, C7, and C8. The internal error amplifier of the LTC1266 will act as an overvoltage protection loop should the LT1431 ever fail.

Conclusion

The Pentium microprocessor offers some interesting challenges to the power system designer. To operate the microprocessor at higher clock speeds requires very stringent supply voltage specifications. Stop-clock power saving modes have introduced severe load transients not present in older generations of processors. However, with careful attention to detail both in component selection and mechanical layout, the required performance can be obtained. Also, the need for high efficiency can be met while providing the required dynamic performance.

PFC, continued from page 18

to the output. Therefore, the input power, P1, flows into the load device, E3, as output current I2. Further, since E3 is a constant voltage, I2 will have a waveshape that is identical to P1. Figure 4e details the input current I1 and the output current I2. By subtracting I2 from I1 we can see what I3 looks like. Figure 4f details the input current I1 and the current through box 3 (I3).

This is the first in a series of articles explaining power-factor correction. The next article will present more component level circuitry using the LT1248 and LT1249 PFC devices. In the meantime, if you require more information contact the LTC factory.

Power Factor Correction, Part Two — Filling in the Boxes

by Dale Eagar

In part one of this article, we investigated power factor correction (PFC) by looking at its line frequency voltage, current, and power waveforms. The device that performs PFC is called a power factor correction conditioner (PFCC).

The waveforms shown in part one are ideal, in that they reflect an average of what is happening inside the three boxes, ignoring higher frequency effects. We showed that PFC could be performed if the appropriate components were implemented. We further developed the concept of an instantaneously adjustable DC variac as an equivalent circuit for the power handling part of the PFCC. In part two, we will develop the circuitry for the implementation of the DC variac by introducing the boost converter.

Why the Boost Converter?

Even though several circuit configurations can perform PFC conditioning, the boost topology is by far the most popular, because of the topology's inherently low input ripple current. This ripple current is at the switching frequency of the boost converter, and must be filtered by an EMI filter at the input terminals of the PFCC. Unfiltered switching-frequency ripple may be conducted down the power line as EMI.

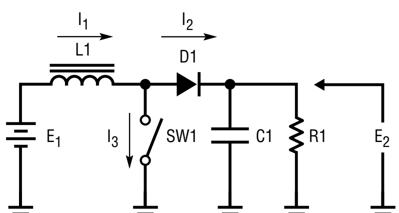

Figure 1. Simple boost converter

An Ideal Boost Converter

A simple boost converter is shown in Figure 1. The boost converter has two modes of operation, each with its own characteristics. The two modes are known as *discontinuous mode* and *continuous mode*. A boost converter functioning as a PFCC will operate in both modes. The criterion for determining in which of these two modes a switcher is running at any given time is whether the inductor is left unloaded for any part of a switching cycle (transitions don't count). If the inductor is unloaded (SW1 and D1 both off) during a switch cycle, the switcher is operating in the discontinuous mode. With the circuit shown in Figure 1, operation becomes discontinuous when the inductor current decays to zero. This happens during the part of the switching cycle when SW1 is open.

To understand the workings of a switching regulator, it is necessary to have at least a mild immunity to "inductorphobia."

The global outbreak of Inductorphobia of the late 20th century threatened to wipe out all analog circuit design. Fortunately, the requirement for power factor correction mandated the use of the inductor. This use is largely responsible for the continuation of the practice of the art through the black age of digital design. [See "The Black Age." History of the Sol System, *Volume 17, p. 12,947]*

Immunization against Inductorphobia involves exposing oneself to inductors, and is highly recommended.

When used in a boost-mode converter, the inductor is placed across the input line and allowed to intercept and store some energy. The inductor is then placed between the input and the output to dump its energy, (along with some additional intercepted line power) into the load.

An Introduction to the Ideal Inductor

1. An ideal inductor will act to prevent DC voltage across its terminals. The inductor will steal energy from any source that attempts to impose a voltage across its terminals.
2. An ideal inductor will store the minimum possible energy.

   The inductor will attempt to dump any energy that it has stolen at the first possible moment.
3. An ideal inductor, having an inductance of L, will stretch and store L volt seconds of charge for each ampere of current flowing through it. The inductor will relax and return L volt seconds of charge to the circuit upon withdrawal of each ampere of current flow.

If, during a switch cycle, the inductor can successfully unload all of its energy, the operating mode is said to be discontinuous. Otherwise, the inductor is forced to store some amount of energy through multiple switch cycles. The presence of such stored energy indicates the continuous mode of operation.

When used to implement a DC variac, a boost-mode switcher controls the duty cycle of the switch SW1 so that the volt seconds across the inductor always add up to zero over any complete switching cycle.

The Continuous-Mode Boost Converter

In the steady state, the continuous-mode boost converter implements the function of the DC variac. The duty factor is set by the constraint on volt seconds, namely that the total volt seconds imposed on the inductor will be zero when looked at over a complete switch cycle.

The voltage transformation ratio of the DC variac is:

$E_2/E_1 = 1 + DF/(1 - DF)$

where DF = Duty Factor

Some interesting properties of this DC variac are:

- ❑ DF = 0, $E_2 = E_1$
- ❑ DF = 0.5, $E_2 = 2 \times E_1$
- ❑ DF = 1, E_2 = Infinity

It's easy to see that as duty factor approaches unity, things get interesting, and can, in fact, become quite a problem.

The problem is not merely that you get infinite voltage, but that you get infinite voltage and infinite current at the same instant—that means infinite power, which tends to rearrange galaxies.

We know that this is a problem, because after a past catastrophic galactic self destruction, we sent scout ships to the estimated center of the galaxy, only to find a breadboard of the circuit shown in Figure 1 floating in space with switch SW1 open. Evidently, ideal components don't vaporize.

One interesting property of a boost converter with a DF of unity is that the switch SW1 never opens, which can be good or bad. In the non-ideal boost converter, the switch simply blows up, limiting current. In the ideal circuit, the power stored in L1 increases with the square of the duration of the time SW1 is closed. This leads us to one final statement about the boost converter with a unity DF:

If, in your engineering adventures you happen across the circuit shown in Figure 1, implemented with ideal components, and with the switch SW1 closed, do not open SW1! Call 911, and, for added safety, take the first hyper-light shuttle out of the galaxy!

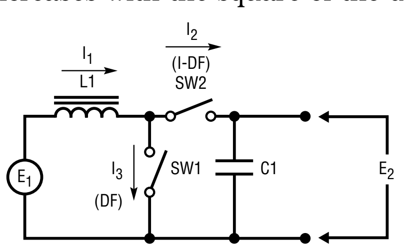

Figure 2a. Continuous-mode boost converter

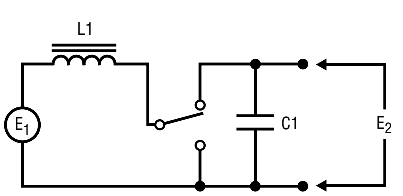

Figure 2b. Simpler version of Figure 2a

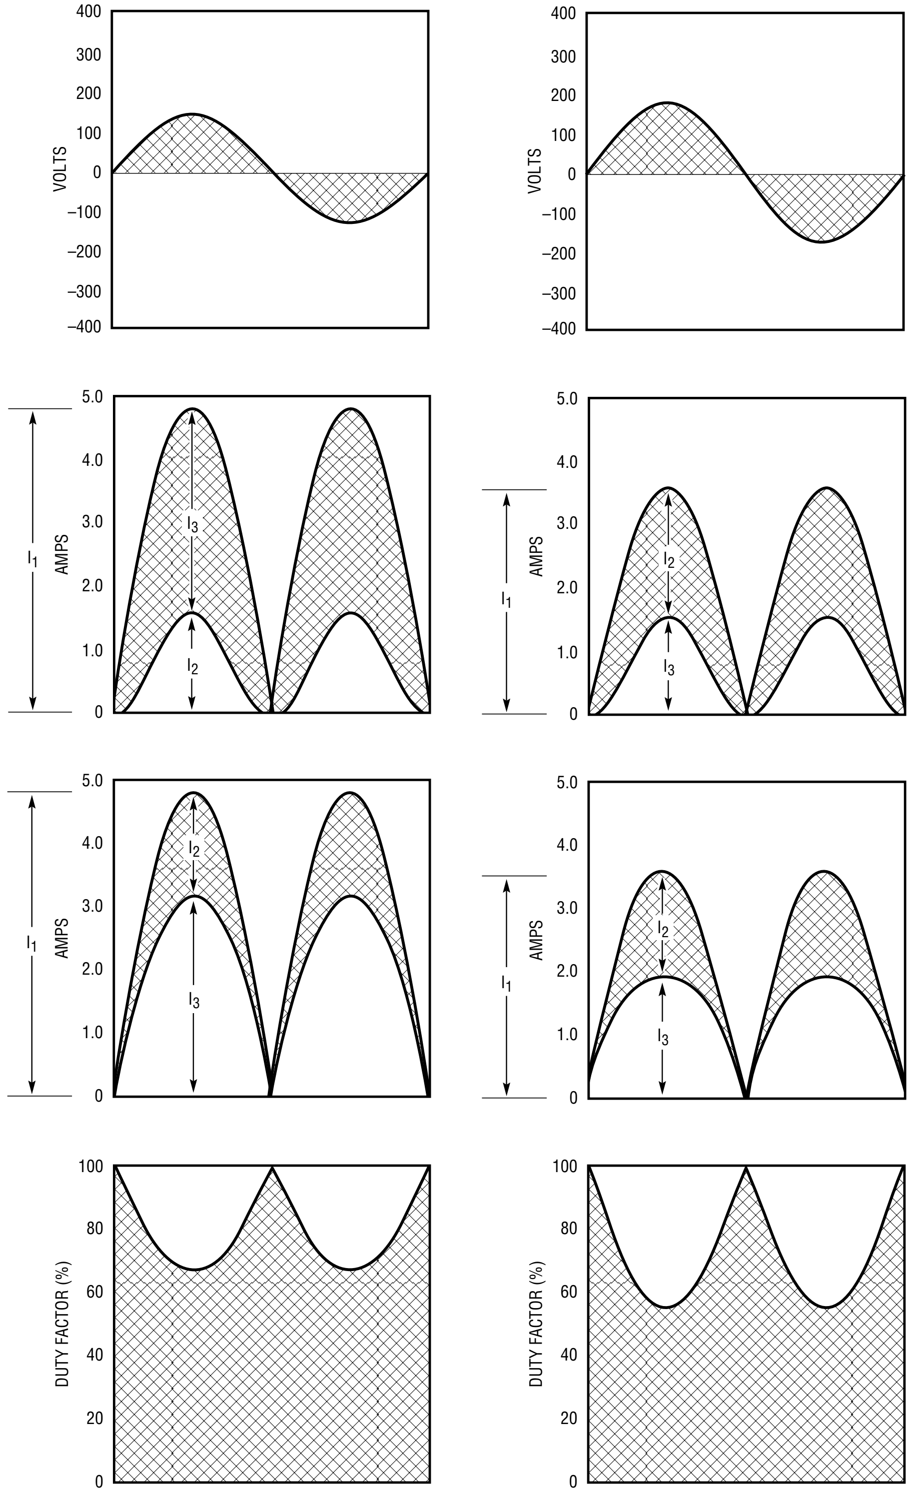

Figure 3. Waveform gallery

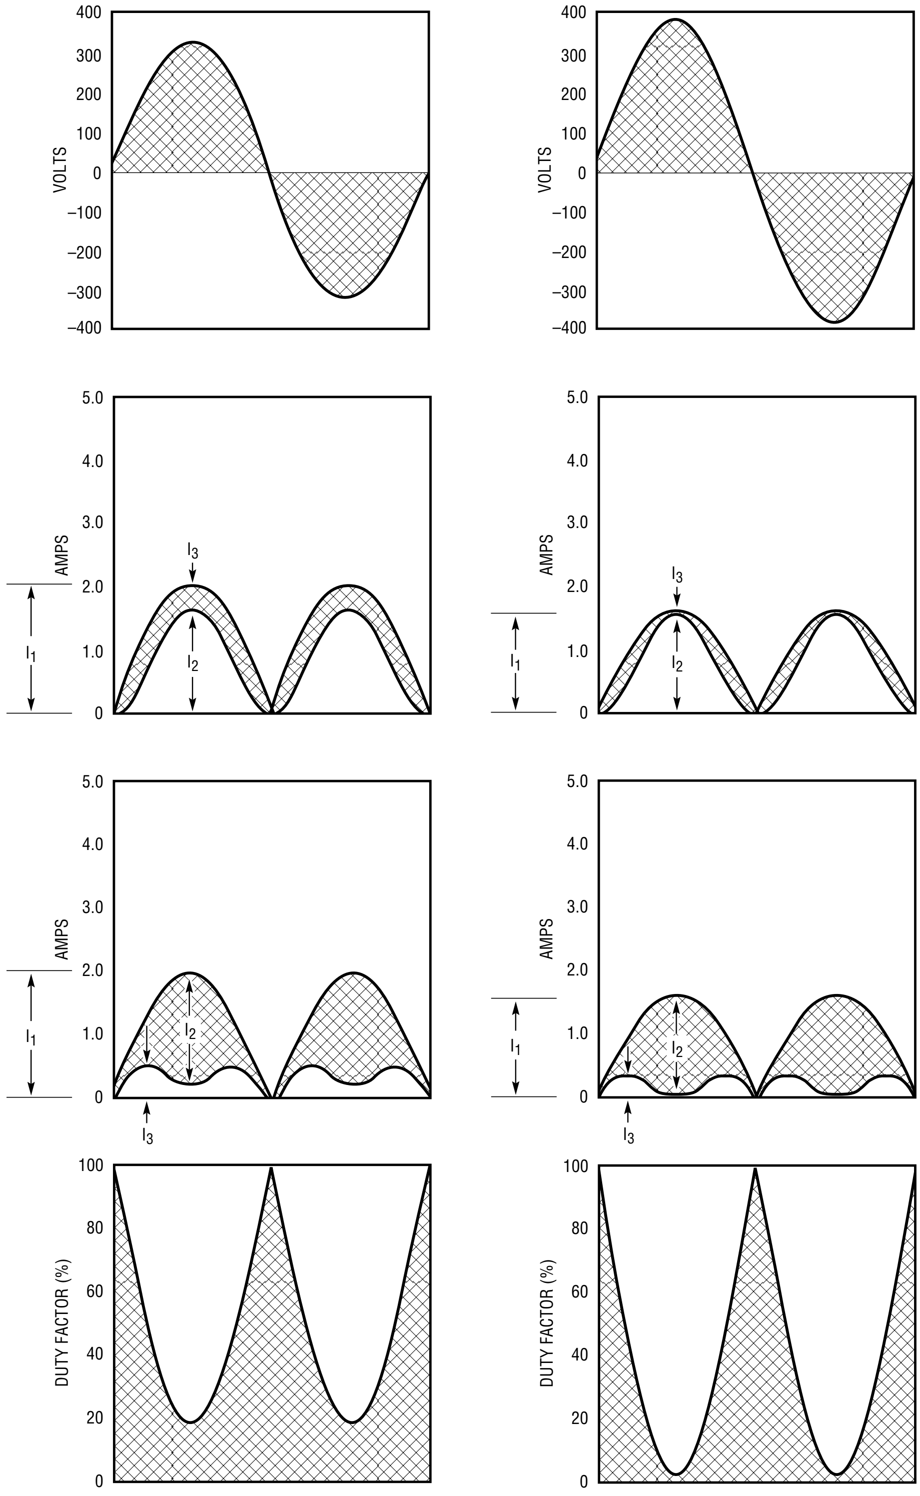

Figure 3. Waveform gallery (continued)

The Continuous Boost Converter as a PFCC

By setting the switching frequency of the boost converter to many hundreds to several thousand times the line frequency, we get the freedom to analyze the boost converter in terms of average values over multiple cycles. Because the system detailed in Figure 1 behaves as a linear system in both states of SW1, the average values of I_1, I_2, and I_3 will obey the same laws obeyed by the instantaneous values of I_1, I_2, and I_3. Thus, Kirchhoff's laws apply to the averaged steady-state values of I_1, I_2, and I_3 just as it applies to the instantaneous values. Once past all of that linear system stuff, we can get to much more interesting things like circuits and waveforms.

To implement a continuous-mode PFCC with a boost converter, a slight modification needs to be made to the circuit in Figure 1. The modification involves the substitution of a second switch SW2 for the output diode D1, as shown in Figure 2a. The opening and closing of the switch SW2 is out of phase with the opening and closing of SW1. Thus, the action of SW1 and SW2 constitutes a single-pole, double-throw switch, as shown in Figure 2b. This modification causes the boost converter to *always* operate in the continuous mode.

Figure 2b details this continuous-mode boost converter implementation of the PFCC. Figure 3 shows the waveforms of the PFCC shown in Figure 2b. One can see that the duty factor is changing continuously, and is directly related to the input voltage. An interesting property of the continuous-mode boost converter is that the duty factor does not change significantly with load current. (This is to be expected for a collection of things whose purpose is imitating a DC variac.)

The circuit of Figure 2b is not trivial to implement in the real world. It not only requires a switch to be subsisted for D1 (Figure 1), but also requires four additional switches to implement the input rectifier bridge (so conveniently missing from Figure 2b).

The Nonsteady-State Boost Converter Operating in Continuous Mode

One of the problems of an ideal approach to a problem like PFC is oversimplification. Here we have developed a model for a DC variac that works wonderfully well in all steady-state conditions. If the load current changes, does the duty factor need to change?

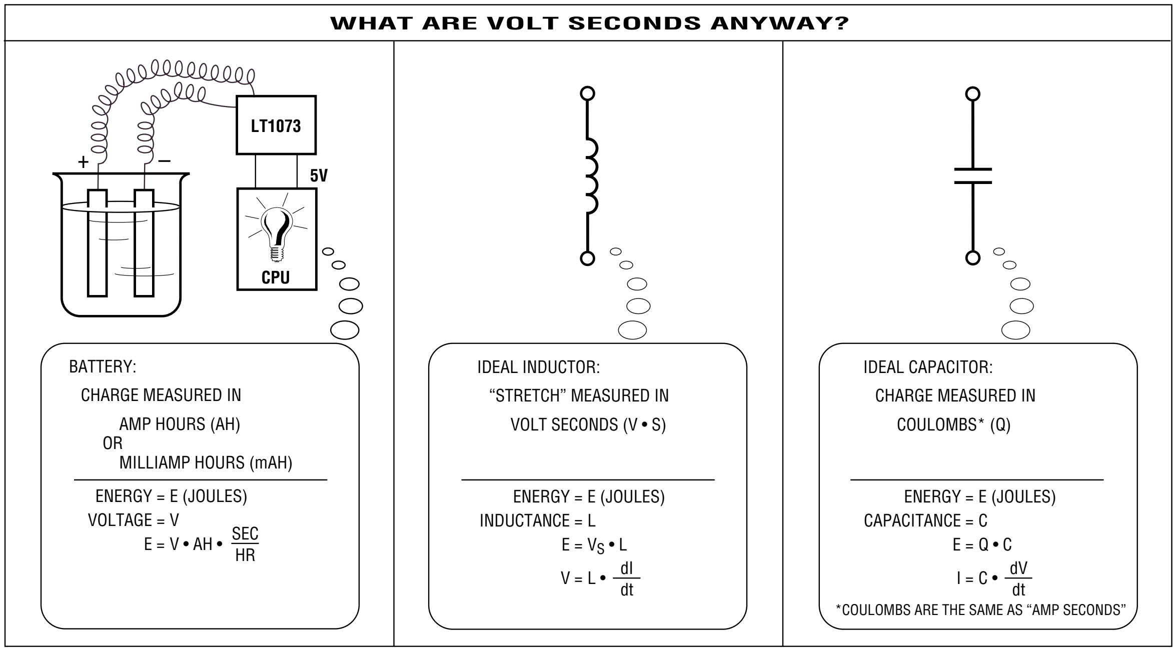

The answer is both yes and no. The model of the DC variac is not quite accurate for a continuous-mode boost converter. The real model of the DC variac needs to include inductance L1 in series with the input. The net effect of L1 in the model of the DC variac can be seen when the system responds to steps in current at fixed input and output voltages. To allow a change in current, L1 will capture and store L1 volt seconds per amp of current change. The duty factor will have to change momentarily from the steady-state value to allow the choke to capture the needed volt seconds. This is illustrated in Figure 4.

In the working PFCC, the effect of capturing and releasing volt seconds is seen in a slight shift in phase of the duty factor waveform. The amount of phase shift is determined by the value of L1, and is negligible for all practical purposes.

Conclusion

The boost converter can be used to implement the DC variac function required to perform PFC, but this requires the duty factor to be well controlled.

In part three of this series, we will investigate the discontinuous-mode boost converter and how it differs from the continuous-mode boost converter.

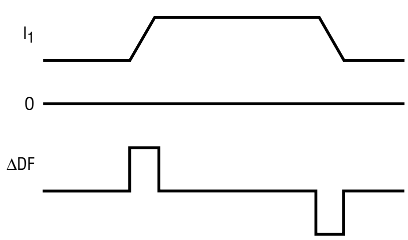

Figure 4. Current in the ideal choke and duty factor

Volt Seconds — the measure of the area under the curve of voltage when plotted against time in the Cartesian coordinate plane. The volt second, the unit of measure of stretch, was popularized in the late 20th century with the advent of the switching power supply. Later in the 21st century, the volt second took on its present sinister meaning when Sumlioux Midge was sentenced to twelve million volt seconds for making the absurd statement "Digital electronics is a mere subset of analog electronics." Midge made the now infamous statement in 2027 near the peak of the "Era of Digital Decadence." [*Solclopedia, 2120.* Volume V, p. 324.]

Stretch — The measure of lattice deformation due to magnetostriction in a ferromagnetic material. Stretch is also used to describe the presence of volt seconds in an inductor. This usage is not strictly canonical in that it is used irrespective of the medium that actually contains the magnetic lines of force. —*Solclopedia, 2120* Volume V, p. 285.

Power Factor Correction — Part Three: Discontinuity

by Dale Eagar

In Part One of this series (in *Linear Technology* V:1, pp. 17–18, 21), we investigated power factor correction (PFC) by looking at its line-frequency voltage, current, and power waveforms. We developed the concept of a DC variac and demonstrated its utility for performing PFC.

In Part Two (*LT* V:2, pp. 3–7), we developed the boost converter and investigated its properties when operating in the continuous-mode. We showed that a continuous-mode boost converter does a pretty good job of emulating the DC variac.

Part Three continues our investigation of the boost converter by focusing on its properties when operating in the discontinuous-mode. We will see that a discontinuous-mode boost converter doesn't look anything like a DC variac.

There is elegance and simplicity in the concept of the DC variac. There is beauty in its application to PFC, a beauty that is manifest in the portions of line cycle where the boost converter is operating in the continuous-mode.

Insubordination Wreaks Havoc

In analog engineering, we like to view things as continuous and smooth, with all nonlinearities well defined. We look at an AC sine wave and see the voltage go from positive, through zero, to negative with no discontinuity. This explains a good portion of the genuine perplexity experienced by the author when the bank practices alchemy on perfectly legitimate checks—changing the paper into rubber. One would think the bank would know that it is only two days until payday and the check book is still half full.

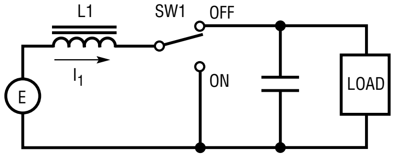

Figure 1a. The boost converter from heaven

Research into the character of banking employees reveals the startling truth that the local bank is staffed by the genetic precursor of the digital engineer, a human so vile, so base, as to actually enjoy inserting discontinuities into an otherwise continuous system. In the boost converter it is the output diode that acts as the banker.

The boost converter shown in Figure 1a is the boost converter from heaven. Its glory is directly attributed to its having a switch, rather than a vile diode, in its output. The switch allows current flow in both directions, and thus inserts no additional discontinuities. The boost converter from heaven always operates in the continuous-mode. Back here on earth we must live with discontinuities.

In a boost converter operating in continuous-mode, the volt-seconds, and thus the current in the inductor, is entirely controlled by:

1. The input voltage
2. The output voltage
3. The duty factor.

The boost converter of Figure 1b is the mortal version. The inductor current (I_1) in the boost converter, driven by the three factors above, goes to zero when the input line voltage swings through zero. Figure 2 shows that the input line current is actually the average of the inductor current over a full switching-frequency cycle. Figure 2a is the inductor current when the input line voltage has reached its peak with the average value representing the line current at the plug. (The ripple is filtered out by the input EMI filter.) Similarly, Figure 2b shows the inductor current and input current when the instantaneous input line voltage is at 40% of its peak value. Finally, Figure 2c shows the inductor current and input line current as the instantaneous input line voltage is at 10% of its peak value. Notice that on each switching-frequency cycle, the inductor current goes to zero and then tries to reverse. When the inductor current attempts to reverse, the diode commits the disgraceful act of insubordination, destroying the beauty, simplicity, and continuity of the control function of the continuous-mode boost converter. Figure 2c also shows the area gained due to the clipping of the negative portion of the inductor current. The net effect is that the average current is higher than the current that the continuous-mode's duty factor control of this DC variac would have requested. The very act of preventing the inductor current from reversing, a hideous act indeed, necessitates a new set of models to describe the discontinuous-mode operation of the boost converter. Further, a new control law is required to control the discontinuous-mode boost converter.

The Boundary is a Moving Target

The boost converter could be likened to a fish tank, with the surface of the water representing the boundary of continuity, above which the boost converter operates in the continuous-mode and below which it operates in the discontinuous-mode. Folks living outside the fish tank tend to obtain their oxygen from the air, whereas folks living inside the fish tank tend to get their oxygen from the water. Air folk and water folk have great trouble communicating, as

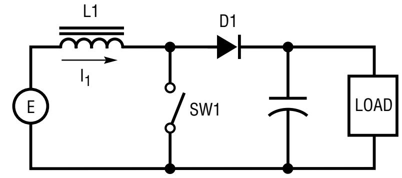

Figure 1b. Mortal boost converter

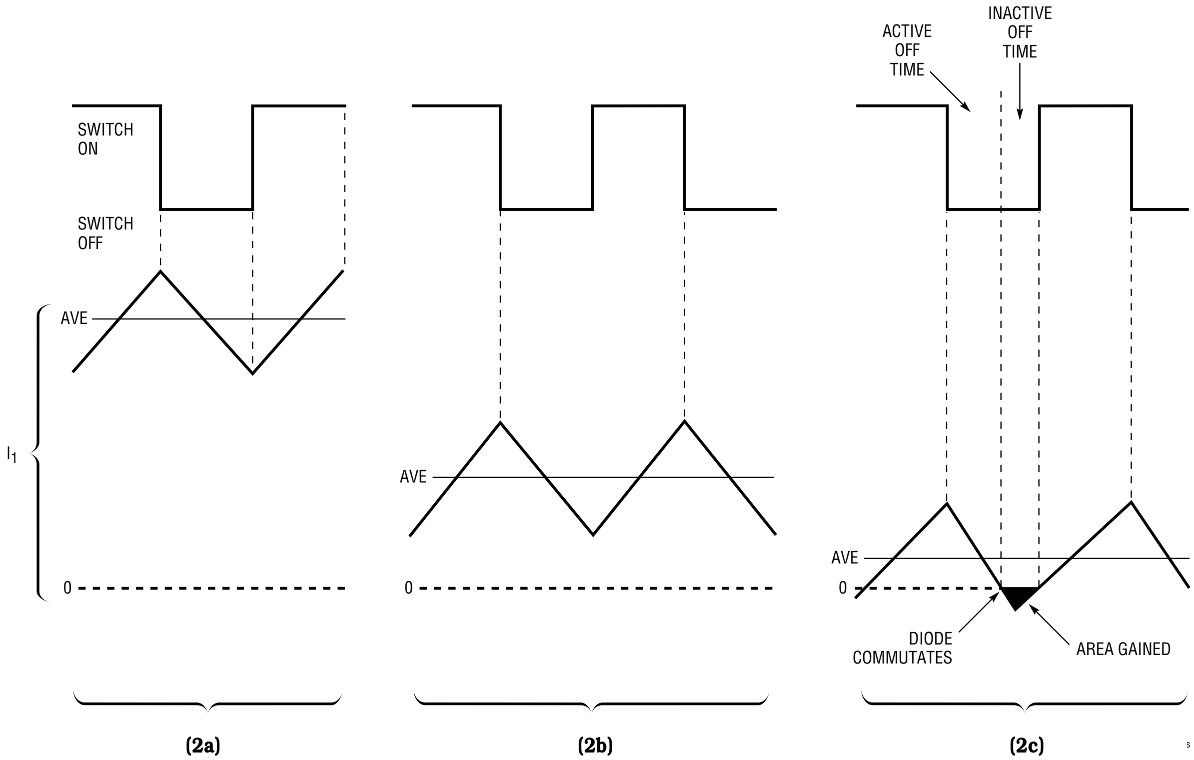

Figure 2. Choke-current waveforms

neither is equipped to speak, let alone breathe, in the other's medium. This lack of communication has had a significant effect on the way the two groups think.

When air folk look at controlling their boost converters, they speak of duty-factor, the all important controlling entity. Water folk control their boost converters with on-time. We, looking at the whole system from the outside, know that the switching-frequency is constant, and thus on-time and duty-factor are describing the same thing.

When air folk plot duty-factor, they see a smooth curve extending down to the surface of the water, and then being refracted at the surface to a different slope as it goes further down to the bottom of the tank. The air folk mathematicians have developed equations that predict the duty-factor plot, equations that exactly describe everything it does in the air. The air folk equations are incapable of predicting what happens when the plot is below the water. The water folk mathematicians have the exact same problem in reverse. We the onlookers can inspect the equations, and when we do, we see that the respective equations have no hope of being consolidated into one. The air folk equation states that the duty-factor is a function of instantaneous input and output voltage and the rate of change in the current flow, but is altogether independent of the steady-state value of the current flow, the operating frequency, and the inductor's inductance value.

The water folk equation states that the on-time is a function of the steady-state current flow, the inductance of the inductor, the switching frequency, and the input voltage, and is independent of the output voltage and the rate of change in current flow. Fortunately for all of us, both equations converge on the same point at the water's surface.

Air folks use volt-seconds to describe the stretch in their inductors, and water folk use joules to describe the energy stored in their inductors. As onlookers, we can convert between Joules and volt-seconds by using the following equations:

$$J = \frac{1}{2L}(V \bullet S)^2$$

$$V \bullet S = \sqrt{2LJ}$$

J = Energy in Joules
L = Inductance in Henries
V • S = Stretch in Volt Seconds

To further complicate matters, the index of refraction of the water changes and is directly related to the water level. The water level changes due to input voltage, output voltage, switching frequency, output current, and inductance of the inductor.

Memory Lost

In continuous-mode, the choke current, and thus power intercepted from the input line, depends on what the choke current was during the previous cycle. Evaluation of continuous-mode boost is best performed by looking at the overall action of an integral number of switching cycles of a running converter.

In discontinuous-mode, the choke current goes to zero during each

continued on page 15

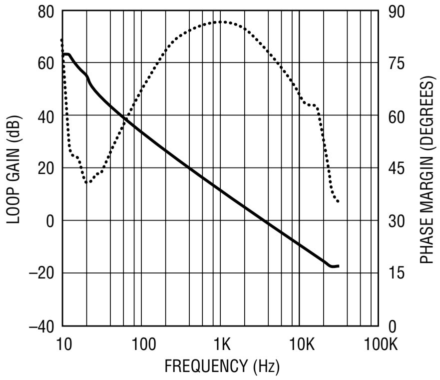

Figure 4. Bode plot of the circuit shown in Figure 3

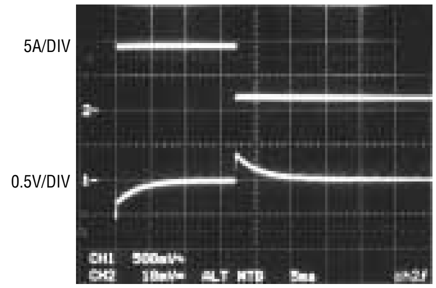

Figure 5. Figure 3's response to a 2A to approximately 10A load step

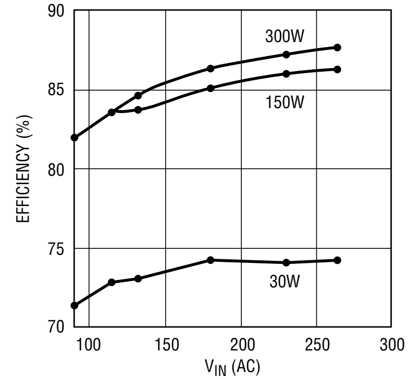

Figure 6. Efficiency curves for Figure 3's circuit

secondary added to the 70-turn primary of T1, bootstraps V_{CC} to about 15VDC, supplying the chip's 13mA requirement as well as about 25mA to cover the gate current of the three FETs and high-side transformer losses. A 0.15Ω sense resistor senses input current and compares it to a reference current (I_M) created by the outer voltage loop and multiplier. Thus the input current follows the input line voltage and changes, as necessary, in order to maintain a constant bank voltage. The forward converter sees a voltage input of 382VDC unless the line voltage drops out, in which case the 180µF main capacitor discharges to 250VDC before the PWM stage is shut down. Compared to a typical off-line converter, the effective input voltage range of the forward converter is smaller, simplifying the design. Additionally, the higher bus voltage provides greater hold-up times for a given capacitor size. The high-side transformer effectively delays the turn-on spike to the end of the built-in blanking time, necessitating the external blanking transistor.

Conclusion

The LT1508 and LT1509 offer complete solutions for low-to-high power, isolated, off-line power-factor-corrected supplies. The many built-in features ensure a simplified, low-parts-count design for a variety of applications.

1118, continued from page 6

The enable pin provides many other useful capabilities to the resourceful designer. Figure 5 shows two LT1118-5 regulators configured to provide sequential turn-on and turn-off from logic control. A logic high level turns on the first regulator; the second regulator turns on only after the first output voltage is stable. Turn-off is in the reverse order, due to the rapid discharge of capacitor C3 through the diode.

These and other applications will benefit from the LT1118 family's unique performance attributes. Any application that requires fast transient load response, unconditional stability under capacitive loading, or logic control of the power output is a potential application for the LT1118 family.

PFC, continued from page 10

off-time. This allows us to model the discontinuous-mode boost converter as a sum of discrete events. Each event consists of an on-time and an off-time, with the off-time being longer than the time it takes to empty all stored energy from the inductor.

Figure 2c shows that the off-time consists of two separate times:

- ❑ the time when the diode is conducting, the active off-time.
- ❑ the time after the diode has committed, yet before the next on-time, the inactive-off-time.

The Energy Bucket

We look at a discontinuous-mode boost converter as an energy bucket that gets filled and dumped with each cycle. During each cycle the bucket takes a certain amount of energy (J). Over a given number of cycles the bucket will intercept many bucketsful of energy and deliver them to the load. A boost converter operating at a switching frequency of f and intercepting an amount of energy of J during each cycle, will be intercepting $f \times J$ watts of power from the input power line. Because this converter is very efficient, essentially all of the intercepted power is delivered to the load.

This concludes part three of our series introducing the power factor correction conditioner.

Power Factor Correction, Part Four: The Brains Behind the Brawn

by Dale Eagar

The Control Room

In the three previously published parts of our series on power factor correction, we investigated the workings of the "power transfer mechanism" as used in the ever so popular "boost-mode power factor correction conditioner." This power transfer mechanism is shown in Figure 1 as the "engine room." Also shown is the "control room," where the decisions are made and actions taken to make the PFC power conditioner behave like a resistor (the ultimate disguise for a switching power supply). Having had a brief tour of the engine room in an earlier article, we can summarize its workings in a few rules:

1. The amount of power intercepted from the input line is related to the duty factor of the boost engine.
2. When the duty factor is at 0%, no power is intercepted from the input line.
3. Under all conditions, increasing the duty factor will increase the amount of power intercepted from the input line.
4. The boost engine is essentially lossless, so all power intercepted from the input line ends up in the load.
5. The load voltage is to be held constant.
6. The load voltage shall be set above the peak input voltage, where the peak input voltage is 1.414 times the maximum RMS input power-line voltage; this is where the power factor is to be corrected.

If the peak input voltage exceeds the output voltage at any time, the output diode (internal to the boost engine) takes it upon itself to forward conduct, transferring power from the input to the output without any authorization from the control room. Pure Insubordination! The only way to prevent this disgraceful act by the output diode is to keep the output voltage out of reach of the peak input voltage.

At the Console

In the control room is a control console, a panel with meters and a duty factor control. Sitting in the control seat is the controller, an intelligent entity. The controller is given the task of watching the three meters (Figure 1) and controlling the duty factor to keep things in check.

The controller is given two tasks to perform, namely:

1. Make the input look resistive.
2. Keep the output voltage constant at 386V, regardless of the load current or input line voltage.

Careful inspection of the above two tasks may surprise you. Let's look at each task separately.

Task One: Making the Input Look Resistive

It would not be too difficult to teach the controller to make the input look resistive. One would need to teach the controller to watch the input volts and amps while adjusting the duty factor. The duty factor is constantly adjusted to make the input amps equal to some constant (K) times the input volts. Presto resistor! Looking into the input terminals, one would see a pure resistor. The input would look resistive at a value of $Z_{IN} = 1/K$. This impedance would be fairly constant for frequencies from DC to the fastest frequency at which the controller could accurately keep up with things. In the real world there is this guy named Nyquist who says, among other things, that the controller would make serious blunders if he were to try to control the duty factor faster than 1/2 of the switching frequency. With the switching frequency set to somewhere around 100kHz, the input looks resistive from DC to many kiloHertz, which is good enough for the 60Hz line and many of its nasty harmonics.

Careful inspection of the output would reveal an output impedance of V_{OUT}^2/K, a big number. The whole system would look like a constant power source, with the power level set to K times V_{IN}^2. These observations are entered in the first row of Table 1.

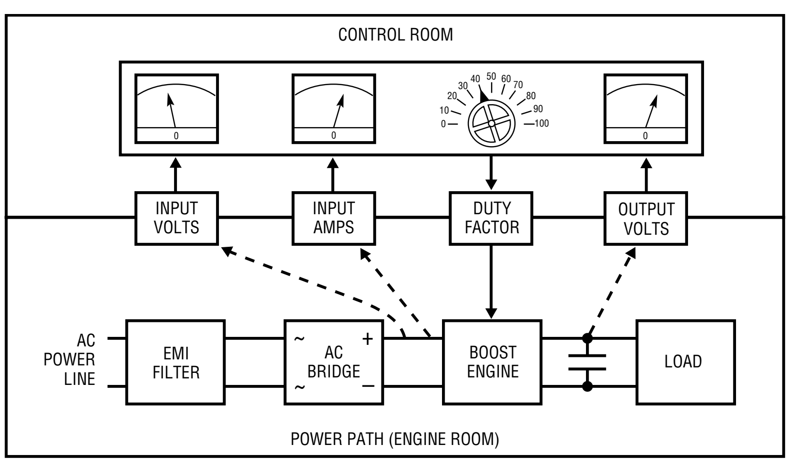

Figure 1. Boost PFC block diagram

Task-2: Keeping the Output Voltage Constant at 386V

Looking at task-2 in the absence of any task-1 constraints, we teach our controller to adjust the duty factor to keep the output voltage constant, regardless of the load current or input line voltage. Having visualized such a system, we can look at its characteristics. Looking into the input port at any constant load current, we see a negative input impedance. (As the input line voltage goes up, the input current goes down.) Looking into the output port we see an impedance of zero. (Any finite increase in output current causes no change in output voltage.) Finally, looking at the power intercepted from the input line we have a constant at $V_{OUT} \times I_{OUT}$ regardless of input voltage. These observations occupy row two of Table 1.

Putting Them Together

It should be apparent that no device can exhibit positive impedance and negative impedance at the same time. However, we know how to cheat—we take the reciprocal of time to get frequency.

This trick doesn't need smoke and mirrors, unless, of course, you hook it up wrong, in which case you get seven years of bad luck. This is a trick in the frequency domain.

The first step is to loosen up the output specification a little. We can see that the power delivered from a sine wave into a resistive load is anything but constant (see Figure 2). We would have to have an infinitely large output capacitor to realize zero output ripple when there is input ripple current. Infinite capacitors went out of vogue back in the Cretaceous era. First we loosen up the output voltage specification to allow a few volts of ripple, bringing the output capacitor down to a realistic size. Second, we further loosen the output specification to allow the output voltage to swing tens of volts during heavy load steps.

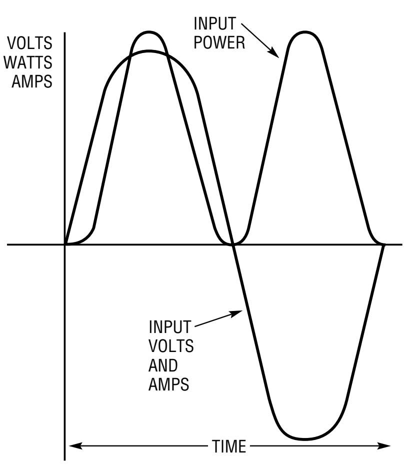

Figure 2. Input power waveform

Table 1. Task-1 and task-2 characteristics

Task	Input impedance	Output impedance	Power transferred
Make Input Resistive	1/K	V^2_{OUT}/K	$K\ V_{IN}$
Constant V_{OUT}	Negative	Zero	$V_{OUT} \times I_{OUT}$

We do this by teaching the controller to keep the output voltage at 386V when averaged over any 3-second interval. This is in effect telling the controller to make the input impedance negative for frequencies from DC up to several Hertz. We are instructing the controller to adjust the K value in the task-1 operation slowly to keep the output voltage at an average of 386V. See Figure 3 for the input impedance versus frequency characteristics.

Summary/Hindsight

Power factor correction is incorporated into a product design for one of several reasons:

A. Government standards, such as IEC555, compliance with which is required to sell product in several worldwide markets.

B. The marketing boys upstairs feel that they can sell it as a feature.

C. Management needs to keep the power supply design group out of trouble until the next major design.

The boost converter is the best way to perform PFC at any power greater than 50W. The crowning advantage of the boost topology is its continuous input current when running in the continuous mode. The loss of continuity of input current in the discontinuous mode is unimportant, as discontinuous mode only happens at low input currents where input ripple is less critical.

The boost converter PFC provides no isolation from the line, has high inrush current, has a very high output voltage, has 120Hz ripple on the output and exhibits slow load-step response. All of these difficulties of the boost converter are small compared to the problems in the design of an input EMI filter if other technologies were used.

Isolation, hold-up time and regulation are performed by the switching power supply that follows the PFC conditioner. Inrush protection is provided by a negative tempco thermistor or other external circuitry.

It is advantageous to synchronize the PFC conditioner and the following switching power supply 180 degrees out of phase to reduce the ripple current on the output capacitor and reduce overall system noise. The Linear Technology PFC family of parts all have synchronizing ability, and the LT1508 and LT1509 controllers have both a PFC conditioner and a switching power converter on one chip, fully synchronized at 180 degrees.

The input impedance of a functioning PFC will be negative from DC to about 10Hz; from there the impedance goes reactive, then to positive resistance at frequencies above 10Hz. The output impedance of the PFC conditioner is near zero ohms from DC to about 10Hz; above 10Hz the output impedance looks like a variable resistor in parallel with the output capacitor's reactance.

continued on page 25

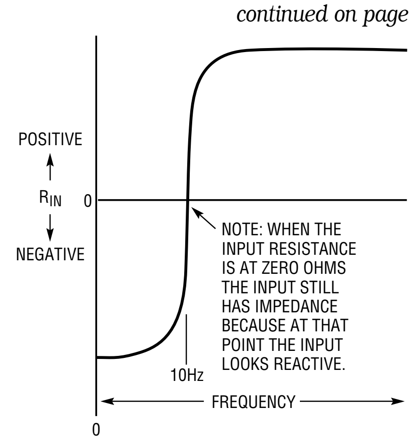

Figure 3. Input impedance (resistive component) versus frequency

C-load, continued from page 22

span is set to 5V. Input signals with magnitudes that cover this span will use all of the ADC's codes. However, an input signal that only spans 2V will use only 2/5 of the codes, resulting in a loss of resolution. When the reference voltage is reduced to 2V, the 2V signal will now use all of the ADC's codes. But with the 2V reference, an input signal having a 4V span is likely to exceed the input span, resulting in clipped signals and missed information. In cases such as this, an adjustable reference is needed to adjust the ADC's full-scale gain.

The circuit in Figure 7 uses the LTC8043 12-bit CMOS 4-quadrant DAC driving the LT1220 C-Load amplifier to generate variable ADC reference voltages that allow the user to adjust the ADC's gain. An LT1004-2.5 micropower 2.5V reference is used as the LTC8043's reference source, setting its output span to 0V–2.5V. The LT1220's output is applied to the LTC1410's V_{REF} input. The LT1220 C-Load amplifier ensures that, when changing the LTC1410's reference voltage, the amplifier does not oscillate even while driving the reference input bypass capacitor C5. The LTC1410 can easily handle a reference voltage as low as 1V (±$2V_{P-P}$ input range) with only a 0.5 LSB change in linearity error.

Conclusion

Using advanced process technology and innovative circuit design, LTC has created a series of C-Load operational amplifiers that not only tolerate capacitive loading, but remain stable driving any capacitive load. C-Load amplifiers meet the challenging and difficult capacitive loads by remaining stable and settling quickly.

PFC, continued from page 11 (Design Features)

Conclusion

PFC is fun. It offers technical challenges not because it is complicated but because there are so many little weirdnesses, all happening together in concert to provide a beautiful function. One of the biggest challenges is to find a place to hook your scope probe to see the PFC in action. Scope probes don't read duty factor very well, and probing on the FET drain reveals a waveform that doesn't look like anything you have ever seen before, and may awaken you in the middle of the night. It's not that the waveform isn't pretty, it's that the waveform is very busy, and it is hard to see anything like a sine wave in it. The input current waveform is far easier to enjoy, it's just harder to get to without a good current probe.

Authors can be contacted at (408) 432-1900

Index of Circuits in *Linear Technology*, Issues I:1 Through V:4

Included with this issue is a complete index of the circuits featured in *Linear Technology* Magazine, issues I:1 through V:4. This index evolved as a result of a request from one of our Field Applications Engineers, who was unable to locate a circuit he believed had appeared in *Linear Technology* Magazine. We set out to create an index in which the reader could "browse" for a circuit that fit his or her design needs or search for a specific Linear Technology part; we believe that the table design we used here thoroughly fulfills this mission. The reader can page to the desired circuit category, scan the columns for the circuit type and/or part number, and easily find the appropriate article, issue, page number, and figure number(s).

Over 250 circuits from the first five years of *Linear Technology* magazine are included in this index. Hereafter, we will update the index yearly, so new circuits won't get lost. Pleasant Reading!