FEBRUARY 2005
WWW.POWERELECTRONICS.COM

POWER ELECTRONICS® TECHNOLOGY

FORMERLY PCIM MAGAZINE

FOR ENGINEERS, DESIGNERS AND INTEGRATORS

Boosting Power Supply Efficiency for Desktop Computers

- Buck Converter Losses Under the Microscope
- Sensorless Control for PMDC Motors
- Easy Calculation Yields Load Transient Response
- Liquid Cooling for High-Power Electronics

A PRIMEDIA PUBLICATION

Boosting Power Supply Efficiency for Desktop Computers

By Dhaval Dalal, Technical Marketing Director, Power Supplies, ON Semiconductor, Phoenix

Adopting active PFC for harmonic-reduction while making other architectural- and component-level improvements can cost-effectively reduce losses in ATX power supplies.

ATX power supplies—the generic name for Intel-specified internal power supplies for desktop computers—are designed for the lowest cost possible and offer poor active-mode efficiency. As utilities and regulators seek to save energy, manufacturers are finding enhancements to ATX power supplies that are technologically superior and cost effective.

For years, the ATX power supplies used in personal computers (PCs) have been designed using traditional forward topologies, and there has been little motivation for innovation. As a result, their typical efficiency ranges from 65% to 70%, and they require heavy heatsinking and fan cooling for operation. Now, with various degrees of energy concern around the world, utilities and regulatory agencies are looking for any possible way to reduce energy consumption. Consequently, PC power supplies are a prime target.

One innovative program will award rebates to manufacturers of computers with power supplies having efficiency greater than 80% (see www.80plus.org). Going from 67% to 80% efficiency on a typical 300-W ATX power supply can save 75 W at full load. With more than 130 million PCs sold every year, the potential for energy savings is obviously tremendous.

This article identifies the scope for efficiency improvements and the projected impact on global energy requirements; provides an overview of existing topological and component approaches to ATX power supplies; and suggests paths for efficiency enhance-

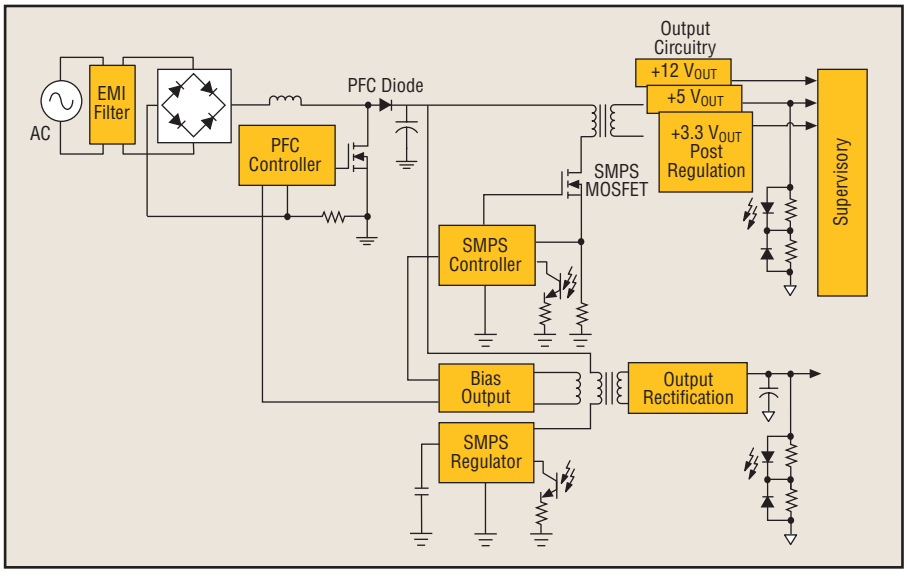

Fig. 1. *Typical block diagram/structure of an ATX power supply.*

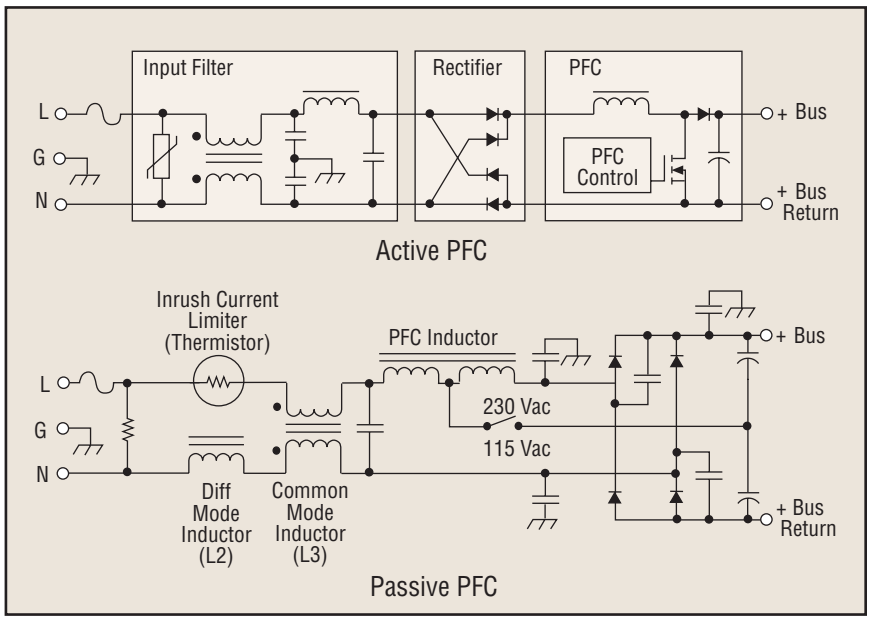

Fig. 2. *Passive and active PFC implementations.*

ments that are technologically superior and cost-effective.

Fig. 1 shows a generic block diagram/structure of an ATX power supply. The power supply output ranges from 250 W to 350 W, supplied primarily at three major output voltages—12 Vdc, 5 Vdc and 3.3 Vdc. Any ATX power supply sold in Europe, Japan and some other regions of the world is required to comply with the harmonic current requirements specified in IEC1000-3-2. As a result, a harmonic reduction front-end in the form of a passive choke or an active power stage (i.e., power factor correction boost) is necessary.

The main power conversion stage for this type of supply is usually a single- or two-switch forward converter with a wide range of reset schemes. While the cross-regulation between the 5-V and 12-V outputs is managed by use of coupled inductors, the 3.3-V output requires dedicated postregulation that is achieved using classic mag-amp circuitry. In addition to the main power stage, most ATX power supplies have an auxiliary (standby) power supply that supplies the power (usually 10 W to 15 W) during the standby mode. This arrangement eases the burden of meeting the low-power standby requirements and also allows easy biasing of the main power stage components.

While the arrangement shown in **Fig. 1** has its merits—including its proven approach, low cost, easy component availability—it does not distinguish itself when it comes to delivering an efficient power supply solution. Easily identifiable areas of improvement can be classified in three major categories:

- A more appropriate harmonic reduction front-end for improved system efficiency
- Architectural-level improvements to eliminate the losses
- Component-level improvements and upgrades to reduce losses.

Each of these areas alone can help improve the efficiency of the power supply. But when judiciously combined, the result is an even more optimized solution. For example, component-level improvements alone may lead to a much higher cost to achieve a required level of efficiency than combining architectural improvements and optimized front end.

Impact of the PFC Front-end

Active power factor correction (PFC) techniques have evolved significantly over the past decade and a half in terms of knowledge base, component availability, design optimization and cost reduction. However, a large portion of the ATX power supplies required to meet the IEC 1000-3-2 requirements still use a passive PFC approach.

Despite the bulk and "clunkiness" of the passive inductor required, economic rationalizations are available for this choice. However, in the broader context of system optimization and end-customer preferences, these rationalizations may soon lose validity. End-user requirements for higher functionality and higher power levels are accompanied by the need for more compact solutions. In the active PFC

Attribute	Active PFC	Passive PFC	Comments
Complexity	Medium (full-power stage needed)	Low (choke, range switch, extra bulk cap)	Reduced complexity for active PFC with newer components
Input rms current	3.0 A	3.69 A	Higher current leads to larger filter size
Output voltage range	300 V to 415 V	200 V to 375 V	Impact on SMPS stage operation
Bulk capacitance	220 µF, 420 V	2 x 1000 µF, 200 V	Passive value often traded off against Vo range
Protection features	Incorporated	Not available	Added circuit costs
Reliability	Foolproof (no range switch)	Potential failure due to range switch	
Efficiency at 115 Vac	~94%	~96%	Active efficiency can improve with better semiconductors

Table 1. *Comparison between passive and active PFC solutions for a 250-W ATX application.*

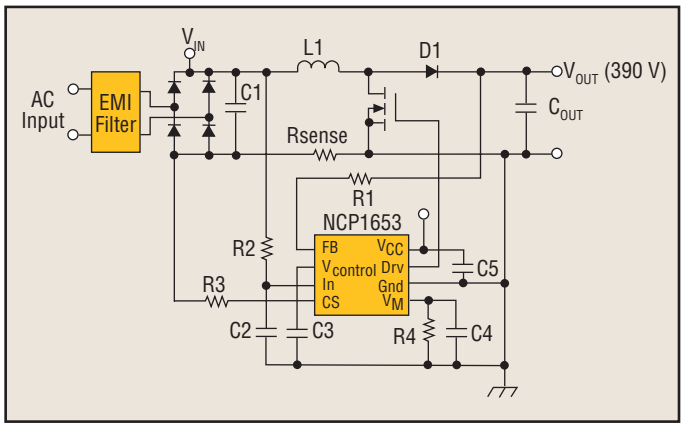

Fig. 3. *Active PFC implementation is simplified using the NCP1653.*

realm, these demands are easier to satisfy than with the passive approaches.

A key issue to consider when choosing a harmonic reduction circuit is its impact on the rest of the system. We can quantify this effect by examining a 250-W ATX example. **Fig. 2** shows the representative circuits for the active and passive approaches. The passive PFC solution requires a range switch and a voltage doubler to keep the bulk voltage relatively constant. The differences between the two approaches are quantified in **Table 1**.

One of the most significant benefits of the active PFC is realized in the design of the downstream SMPS converter. The minimum input voltage for the SMPS stage is increased (as shown in **Table 1**, 300 V versus 200 V for 250 W) for the same holdup time and for much smaller capacitor values.

With the narrower input voltage range, the SMPS design has a flexibility that can be exploited for lower system cost, better efficiency or both. For example, the rms current in the SMPS stage power MOSFET and the transformer is reduced by 33% (from 2.05 A to 1.36 A) with the inclusion of PFC stage, which leads to a 56% reduction in conduction losses in these elements if the same $R_{DS(on)}$ and wire gauge are used. Alternatively, lower-cost MOSFETs may be used to achieve the same conduction loss.

Another area where the active PFC solution helps is in meeting the holdup time (or line dropout) requirements—usually one line cycle at full load and low line. The active PFC boost naturally provides great holdup performance without requiring a large capacitance value since the energy is stored in boost output capacitor at a much higher voltage (~ 400 V).

Any non-PFC or passive-PFC solution stores the energy at the peak of the line voltage, which is about 140 V for low line (100 V). With the voltage doubler, this voltage becomes twice that value. The impact of this is reflected in the size of the holdup capacitor shown in **Table 1**. Remember that the values shown in **Table 1** are for different minimum voltages. If the same minimum voltage were used, the difference in capacitance value would be much more pronounced.

Architectural-level Improvements to Eliminate Losses

Under the new high-efficiency paradigm, the entire architecture of the ATX power supply has to be reviewed in order to optimize its performance. The new architecture must present the potential to meet the long-term performance objectives—including harmonic reduction, improved efficiency, compact size, robust operation, better acoustics and design flexibility—in a cost-effective manner. Most of the present-generation designs have not used the holistic approach, instead approaching the design improvements in a piecemeal fashion. As a result, opportunities for improving performance are missed.

Given that harmonic reduction is an increasing requirement, the natural question is whether this requirement can be met by using a single power stage. Many single-stage solutions have been proposed and utilized in lower-power applications. However, in the ATX application, the holdup time requirements rule out some of these approaches. Also, the higher power level makes it difficult to optimize some of these single-stage approaches for both efficiency and cost. Future topological innovations and/or component innovations may allow a more feasible single-stage approach for the ATX power applications. In the present context, however, the two-stage approach is more relevant.

The boost topology is the ideal and universal choice for the active PFC front-end for several reasons. First, it is a good choice for low-input voltage coupled with high-output voltage, because it has continuous input current and the discontinuous output current is low. It also absorbs input transients well. In addition, it contains the energy-storage capacitor that stores the energy at minimum size, since the energy is proportional to V^2 and the size is proportional to V. Finally, recent emergence of controllers and other components has reduced the total component count and reduced the complexity significantly. For example, a complete PFC front-end using ON's recently introduced NCP1653 is shown in **Fig. 3**.

The choice of the optimal second-stage topology is less clear. Traditional usage of single-transistor forward-type topology is under increased scrutiny due to the need for better efficiency solutions and increasing power demands. Major constraints of forward converters include their inability to handle a wide input-voltage range effectively. The voltage stress is maximum at high line and current stress is maximum at low line, requiring components to handle severe stresses on both operating conditions.

Furthermore, there is the requirement of a very high-voltage switching device (above 800 V) for any system that has active PFC front-end. Other limitations include the inability to extend the duty cycle beyond 50%, and achieve lower current stresses in the primary and lower voltage stresses in the secondary. Finally, it is not possible to push switching frequency higher because of hard-switching waveforms. All of these constraints become more pronounced in the new paradigm of higher power re-

Attribute	Half-Bridge	Active Clamp	Active Clamp Comments
Dmax	0.5 (0.45) per FET	0.61 (0.6)	Higher D leads to several advantages
V_{ds}/rating	425/500 V	687/800 V	Two 500 V vs. one 800 V
$R_{DS(on)}$	0.27 Ω	0.8 Ω	Higher $R_{DS(on)}$ for same loss
Iprim (rms/pk)	1.01/1.5 A	0.86/1.11 A	Lower currents
Inductor	X µH	1.5*X µH	50% larger inductor for same frequency
Sec. peak voltage	42 V	32 V	Allows low-voltage synch MOSFET—high efficiency
Transformer	CT secondary	No CT on secondary	Simpler construction

Table 2. *Comparison of active clamp and half-bridge topologies for 12-V, 180-W output*

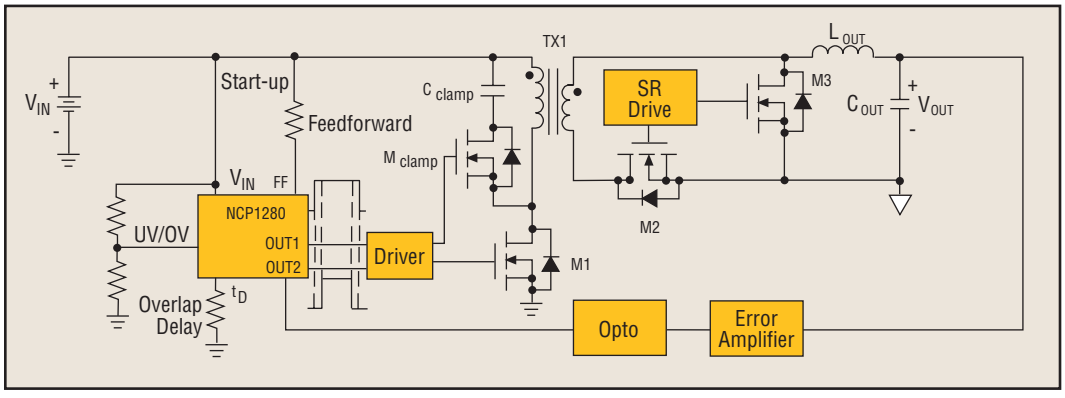

Fig. 4. *Active clamp forward topology (simplified) using the NCP1280.*

quirements, active PFC front-ends and higher system efficiencies.

A viable alternative is the use of the active-clamp forward topology. The active clamp forward topology (**Fig. 4**) allows a more efficient solution for this application at a competitive cost. In this topology, the transformer reset is accomplished with an active clamp circuit (consisting of a FET and a capacitor) that replaces the reset winding used in traditional single-switch forward converters.

This arrangement offers many benefits. For example, the duty cycle can go higher than 50%, resulting in higher turns ratio, lower primary currents and secondary voltages, and smaller output inductor. Also, the voltage stress on the primary FET remains relatively constant over the full input voltage range, leading to better overall efficiency. In addition, soft-switching is possible with this approach, which can lead to further size reduction by increasing the switching frequency.

To demonstrate the effectiveness of this approach, a prototype circuit was built using the NCP1280 as the controller for the active clamp converter and delivering 265 W for the ATX outputs. The resultant efficiency was measured at 93% at low line (325 V) and 92.34% at nominal line (400 V). This represents a gain of 60 W in power losses compared with the traditional forward approach.

It is obvious that the overall added complexity of the active clamp approach is minimal compared with the forward converter and with significant performance advantages. With some better component usage and coupled with a more efficient PFC stage, this approach can lead to the ultimate goal of a compact, fanless power supply for the desktop computer.

While the active clamp approach offers significant advantages, other candidate topologies can be considered as alternatives for this function. One is the half-bridge converter, a traditional topology that offers some benefits with the addition of another power switch (better core utilization, smaller inductor, lower voltage stresses, etc.). A quick comparison like the one in **Table 2** shows that the active clamp approach is simpler in terms of silicon use and complexity.

Meanwhile, the half-bridge resonant converter (LLC resonance) topology is gaining popularity in many performance-driven applications because it offers high efficiency and low EMI. It also eliminates the output inductor requirement. However, the control is more complex and this topology does not handle input-voltage variations well. The implication for the ATX application is that if the holdup time requirements are managed without a significant drop in SMPS minimum input voltage, then this topology may be attractive.

Having addressed the front-end and SMPS requirements, the final piece of the puzzle is the post regulation for the 3.3-V output. Traditional methods for achieving this include linear regulation, which is both inefficient and cheap, or mag-amp post regulation. The mag-amp method is limited to lower switching frequencies. Also, the ability to go to synchronous rectification is lacking. An approach where the mag-amp is replaced by a solid-state switch is made possible by integrated controllers such as NCP4330 and offers significant advantages, including better overall efficiency, the ability to go to higher switching frequency without increasing losses and additional protection functions. **Fig. 5** depicts this approach.

Component-level Improvements

Once a given architectural approach is chosen, there is always room to improve the efficiency of the power supply by using components that provide better performance. As in most designs, this usually is a cost/performance tradeoff. Careful component selection can help achieve better performance without adding cost if the right parameter tradeoffs

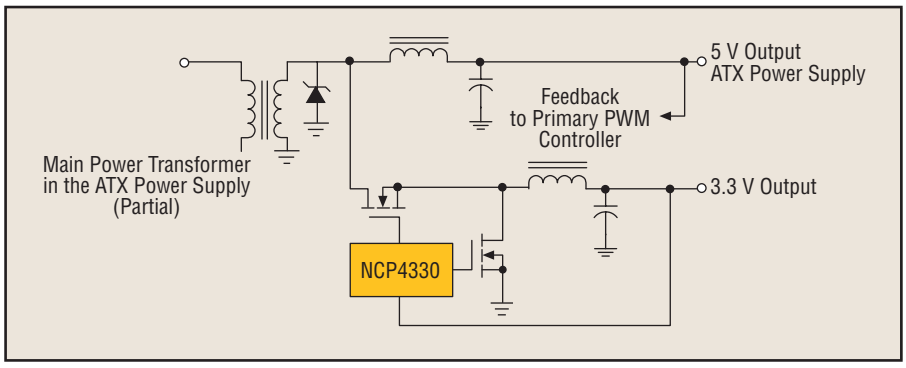

Fig. 5. *A switching post regulator approach for 3.3 V output using the NCP4330.*

are made. Also, the trend in semiconductor technology in recent years has led to tremendous cost reductions, and that trend is likely to continue. As a result, designers have more flexibility in improving their system efficiency by choosing circuits with better semiconductor content.

For example, consider the output rectification in ATX power supplies. The rectification is traditionally accomplished by Schottky rectifiers. For an output current of 20 A on the 12-V output, the conduction loss for a standard Schottky is computed to 17 W. With a 200-mV improvement in V_f, that can go down to 13 W. Even further loss reduction can be achieved if the Schottky rectifier is replaced by a MOSFET acting as a synchronous rectifier. Use of MOSFETs with an $R_{DS(on)}$ of 10 mΩ each can result in conduction losses getting down to 4 W.

Similar loss-reduction opportunities are available in high-voltage FETs, PFC rectifiers and input-bridge rectifiers. Magnetic components also are major contributors to the losses, but in recent years, the rate of innovation in magnetic materials has been slower compared with semiconductors. Lower losses can be achieved with better winding techniques, and these should be applied whenever feasible.

The Path Forward

This article has outlined the increasing challenges power supply designers face from end customers, regulatory agencies and competitive pressures. Against the backdrop of these challenges, the natural, more defensive response would be to keep evolving from the existing designs, not rocking the boat too much. Often, that type of response yields the needed cost reductions. However, there is a better more sophisticated way to meet the challenges by taking a renewed look at the whole system and identifying new architectures.

As demonstrated in this article, it is easy to achieve an efficiency of >85% (>93% PFC and >92% SMPS) with the right system architecture. The goal of the power system designers should be the same as the ultimate need for the end consumer—a desktop computer that is light, compact, quiet and affordable. Currently, the power supply for computers is affordable, but it is far from being light, small and quiet enough. Wouldn't it be nice to have a power supply that does not stand out like a sore thumb in your computer system? Not to mention the fact that each year such a transition would save >10 billion kWh of energy globally.[1] **PETech**

References

1. *Estimate based on: annual kWh reduction from 300 kWh (PIER data) to 240 kWh per PC due to efficiency improvement applied to one-fourth of the global PC usage (~680 M units).*