

BEAD INDUCTOR MEETS CPU POWER CHALLENGES

Introduction

Wound-toroid inductors have been the mainstay for desktop core voltage regulators (V_{CORE}) for years. Historically, the bulky size, sloppy tolerances and high power losses of these inductors were not a concern as they were the lowest-cost solution, making them the solution of choice.

However, three evolving V_{CORE} requirements have highlighted the negative aspects of the wound toroid (size, tolerances, efficiency). Efforts to modify the toroid design to address these issues along with broader global economic conditions have continually increased the price of the once low-cost wound toroid. As a result, the wound toroid is now an ineffective and ultimately higher cost choice for V_{CORE} regulators.

An alternative approach, the through-hole technology (THT) power bead, overcomes the drawbacks of wound toroids while addressing the evolving V_{CORE} requirements. The power bead's performance in the V_{CORE} application is demonstrated and compared with the toroid's through calculated and measured results.

Ferrite-based power beads, an alternative to wound-toroid inductors, offer greater efficiency, tighter tolerances and smaller size in V_{CORE} voltage regulators.




Fig. 1. Historically, wound-toroid inductors used in microprocessor core voltage regulators employed low-cost, high-permeability powdered-iron cores.

Fig. 2. The THT power bead inductor structure consists of a single-turn winding on a gapped-ferrite core.

Regulator Requirements

Each subsequent generation of processor demands faster transient response times. Faster transient response requires the ability to change the current through the V_{CORE} inductor quickly. However, the magnetic field within an inductor resists change ($di/dt = V_{\text{OUT}} / L$); therefore, with a fixed output voltage (V_{OUT}), the only way to increase the di/dt is to reduce the value of inductance (L).

Unfortunately, as inductance values drop, the ripple current through the inductor increases, dramatically raising the switching losses in the inductor. Over the past several years, inductance values have dropped twice and are expected to decrease by similar amounts over the next two years.

As inductance values have dropped, the parasitic inductances in the circuit, such as the inductance of pc-board trace lengths from the inductor to the processor, have become more evident. To eliminate these unwanted parasitic inductances, it is necessary to locate the V_{CORE} inductor closer to the processor. However, to get close, the inductors must be able to fit underneath the overhanging heatsink.

In addition, the requirement to get more components close to the processor necessitates a reduction in the footprint of each component in order to fit them in the now-reduced space. And so, the requirement for faster transient response has led to the requirement for lower-inductance, lower-height and smaller-footprint inductors.

It is now standard practice to use the distributed dc resistance (DCR) of the inductor winding as the current-sense element to control overcurrent protection and the output voltage droop. This is done by measuring the voltage drop across the inductor and then filtering out the portion of this drop that is attributable to the inductance.

Consequently, the tolerance of the inductor DCR and inductance, which was historically never considered, now directly affects the accuracy of the current sensing. Over the past several years, there has been an increased focus on this tolerance and increasing pressure for improvements.

Historically, the efficiency of V_{CORE} solutions was only looked at if there was a thermal issue with the design. However, there is now a new focus on product efficiency due to the overall energy consumption costs to the end user, new Energy Star efficiency mandates and the increased thermal management difficulties of removing more power from a smaller physical space. The ability to meet these efficiency requirements necessitates a reduction in the inductor switching losses that have been increasing due to faster transient response needs as well as a reduction in the dc losses created by the inductor's DCR.

The DCR current-sensing scheme requires some minimum nominal value of DCR to maintain a good current sense signal to noise ratio. However, recent improvements in the pulse-width modulated circuits now enable the use of lower nominal DCR values. To optimize efficiency, it is necessary to design inductors at these new lower nominal DCR values.

Despite the more stringent requirements of lower inductance, lower profile, smaller footprint, tighter DCR and inductance tolerances, and improved efficiency, there is still no tolerance for any price increase in the final product. Therefore, all these improvements must be achieved without impacting the cost of the inductor components.

The Problem with Wound Toroids

In the past, wound-toroid inductors (Fig. 1) were the least expensive V_{CORE} solution. Initially, they used high-permeability (75 perm), low-cost powdered-iron cores. As inductance values decreased, it was necessary to drop the number of turns ($L \approx \text{Turns}^2 \times \text{perm}$) to achieve the lower values.

However, reducing the number of turns increases the operating flux density inside the core ($\Delta B \approx 1 / N$), and higher flux density increases power loss in the core ($P_{CORE}(W) \approx (\Delta B)2x$).

As inductance values dropped further, simply lowering the number of turns was no longer possible without excessive core losses.

Instead, it was necessary to switch to lower-permeability (55 perm) and higher-cost cores with better core-loss characteristics. These cores require more turns for a given inductance, but the higher number of turns lowers the flux density. This cycle of using more turns or the same number of turns on lower-perm cores to achieve lower values of inductance has continued to even lower-permeability cores (35 perm and 14 perm).




Fig. 3. When two 325-nH inductors are tested in a VR11-type, 250-kHz, three-phase voltage regulator, the efficiency of the bead design is several percentage points higher than the toroid.




Fig. 4. A 220-nH bead inductor achieves greater efficiency than a comparable toroid at higher loads when both are measured in the VR11 test setup (250-kHz, three-phase voltage regulator).




Fig. 5. Tests taken on two 160-nH toroid designs and one 160-nH bead design in the VR11 setup reveal a slight gain in efficiency for the bead inductor versus the better of the two toroids.

Although lower-perm powder cores have lower losses, they are still relatively high when compared to other core materials. In the end, this trend results in a higher-priced solution with still relatively high core losses and high dc copper losses due to the number of turns.

The tolerance of the DCR of a wound-toroid inductor is affected by the tolerance of the actual core dimensions and the tightness of the winding around the core as well as the variation in the resistance of the wire used. The resistance of round magnetic wire is tightly controlled to a $\pm 2\%$ tolerance. However, toroid core dimensions vary widely and, because toroids are typically wound by hand or, at best, using a machine-assist process, the tightness of the winding can also vary. It is not possible, without screening, which would dramatically increase the price of the part, to achieve a DCR tolerance of less than $\pm 10\%$ using a wound toroid.

In addition, toroid core inductances typically are controlled only to $\pm 10\%$, and added to this number should be variations in leakage inductance as a result of variations in winding placement. Although tighter tolerances may be quoted, there is no data to support such numbers. If the industry requires tighter tolerances than $\pm 10\%$, a new inductor solution is needed.

Toroids, by their nature, do not make efficient use of space. There is typically a large area inside the center of the toroid that is empty, and the windings bulge out beyond the core dimensions with gaps of empty space between them, increasing the overall footprint. Variations in winding tightness further increase the size of the wound toroid. As such, it is not possible to make a space-efficient wound-toroid design.

Finally, it should be remembered that multi-turn toroids take time to build and use relatively large amounts of copper. Labor costs globally are increasing and the price of copper and other raw materials remains high. These cost drivers coupled with the more expensive low-perm powder cores will continue to drive up the cost of wound-toroid inductors. As such, it seems impossible to meet the evolving requirements of smaller size, higher efficiency, tighter tolerances and lower cost using the wound-toroid solution.

THT Power Bead Inductors

The goals of tightening tolerances, reducing size, improving efficiency and decreasing cost have led to the development of several alternative solutions. One alternative that meets all of these requirements is the THT power bead inductor (Fig. 2). The power bead uses a single-turn winding on a gapped-ferrite core structure.

Ferrite cores have roughly 20 times lower core losses than powdered iron for a given flux density and frequency. This property allows for a dramatic reduction in core losses. In addition, ferrite, unlike powdered iron, is not susceptible to thermal aging, which is a process by which the binder in the powdered cores breaks down at elevated temperatures, causing an increase in core losses as well as more heat and a thermal runaway condition.

The power beads' single-turn winding means that the nominal DCR value can be designed to whatever minimum value is acceptable for inductor DCR current sensing. The combination of lowest-possible DCR and low core losses makes the THT power bead a highly efficient inductor solution.

in the DCR tolerance. The lead can be preformed to tight tolerances and, because the lead dimensions do not rely on a hand-assembly winding process or the core tolerances, DCR tolerances of $\pm 4\%$ have been achieved. This is a dramatic improvement over the existing $\pm 10\%$ tolerances available with wound toroids.

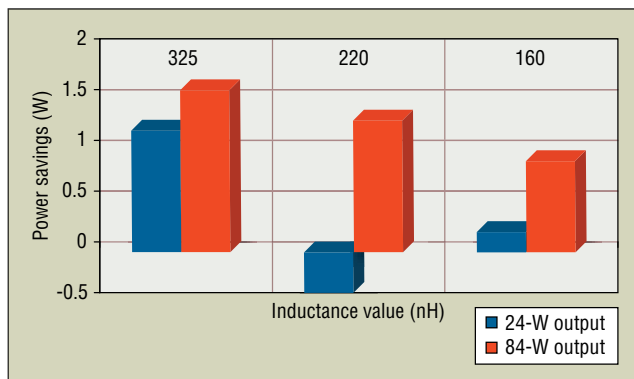


Fig. 6. The measured power savings for a bead inductor versus a toroid inductor in a VR11 voltage regulator vary depending on the inductance value and the load condition.

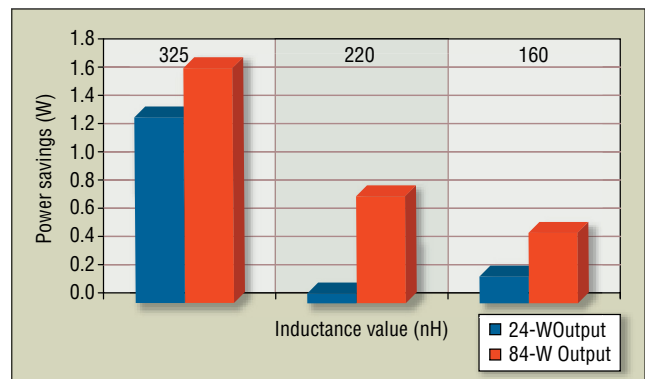


Fig. 7. The predicted values of power saved by a bead inductor are greater than the measured savings at light load (24 W) and less than the measured savings at heavy load (84 W).

220-nH Inductor Solution

To make a low-inductance (220-nH) toroid solution without excessive core losses, it is necessary to move to a lower-perm, higher-cost powder core. The toroid solution for a 220-nH V_{CORE} inductor (Pulse PA2164NL) uses a 14-perm, 0.44-OD powder core with six turns of 2 × 18-GA wire. This design has a footprint of 14.5 mm × 14 mm, a DCR of 1.1 mΩ nominal and a calculated core loss of 136 mW.

The THT power bead equivalent (Pulse PA1894NL) has a 10-mm × 10-mm footprint, a DCR of 0.51 mΩ nominal and a calculated core loss of 130 mW. As seen in Fig. 4, the efficiency at full load is better for the power bead, but at light load, the low-perm toroid actually performs slightly better.

The light-load (24-W) and heavy-load (84-W) power savings are -0.4 W and 1.3 W, which means there is less savings than calculated at light load and more savings than calculated at heavy load (Fig. 7). The discrepancy in power loss could be a result of component placement or variation in actual core loss from that calculated. In any case, the 1.3-W savings at heavy load, coupled with the footprint reduction of more than 50% (Fig. 8), makes the power bead a better solution.

160-nH Inductor Solution

The toroid solution for a 160-nH V_{CORE} inductor (Pulse power bead an optimized solution for low-inductance CORE PA2142NL) again uses a 14-perm, 0.44-OD powder core with five turns of 17-GA wire. This design has a footprint of 14.5 mm × 14 mm, a nominal DCR of 0.7 mΩ and a calculated core loss of 202 mW.

The THT power bead equivalent (Pulse PA2080NL) has a footprint of 10 mm × 7.5 mm, a DCR of 0.5 mΩ nominal and a calculated core loss of 150 mW. As seen in Fig. 5, the efficiency of the power bead is marginally better at light load and better still at heavy load.

As a reference, a toroid design using a high-perm core (35 perm) has also been included on the efficiency curve. It is clear that even though this design has a lower DCR, the excessive core losses make it a poor toroid solution. The light-load and heavyload power savings, respectively, are 0.2 W and 0.9 W (Fig. 6), which compares nicely to the calculated values of 0.2 W and 0.5 W (Fig. 7). In addition to the power savings, the footprint has been reduced by more than 60% (Fig. 8).

It is clear from this analysis that using lower-perm powder cores can offset some of the efficiency losses when using toroids. But, the resultant increase in DCR still makes the THT power bead a more efficient solution. The efficiency gains, coupled with the dramatic size reduction and DCR tolerance improvements, make the THT power bead an optimized solution for low-inductance V_{CORE} applications.

Written By: John Gallagher & Fuxue Jin