

Pick the right inductor construction for a desktop-CPU voltage regulator

CHOOSING THE BEST INDUCTOR CONSTRUCTION FOR DESKTOP-CPU VOLTAGE REGULATORS REQUIRES A GOOD UNDERSTANDING OF INDUCTORS AND REGULATORS.

Until recently, cost and the ability of a part to function in an application have driven the choice of inductors in desktop-computing systems' processor-voltage regulators. Designers have not been greatly concerned with component size, parameter tolerance, or performance requirements. As a result, inductor design for these systems lagged behind the state of the art. Over the past three years, however, the industry has begun to awaken, and vendors have proposed and evaluated several new inductor platforms to overcome new challenges resulting from changing regulator requirements.

First, overall power requirements for desktop processors continue to increase. Moreover, as processor-voltage levels drop, current levels increase dramatically. Increased current causes increased thermal issues and inductor-copper losses: $P_{CU} = I_{RMS}^2 \times R_{DC}$, where R_{DC} = dc resistance.

The second challenge is that processor transient-response times continue to decrease, which means the power supply must be able to respond much faster to changes in load conditions. One of the limitations to this response is the inductance value. Inductors store energy and slow down current changes (di/dt , that is, rate of current change = output voltage/inductance). Ideally, you could simply decrease this inductance to the value required to meet the transient-response criteria.

However, decreasing only the inductance would keep the regulator from meeting the third technical challenge: decreased output-voltage ripple. Output-voltage ripple is a function of the output capacitors' ESR (equivalent series resistance) and the ripple current from the inductor: $V_{RIPPLE} = ESR_{CAP} \times I_{RIPPLE}$. The only way to minimize voltage ripple is by using the expensive option of reducing the capacitor ESR by paralleling more capacitors or decreasing the ripple current through the inductor. However, to reduce the ripple current, either the operating frequency must increase—that is, dt must decrease—or the inductance must increase: $I_{RIPPLE} = Vdt/L$.




Figure 1 New constraints have forced designers to look at new types of inductors in addition to traditional desktop voltage regulators (a). These alternatives include horizontally mounted toroids (b), power-cube inductors (c), rod cores (d), power beads (e), and ferrite versions of the power cube (f).

To some extent, frequency increases have occurred, but the resultant increase in switching losses has limited these increases. The approach commonly used over the past five years to reduce transient-response time and ripple current has been to add parallel power trains running out of phase with one another and then sum the phase outputs at the regulator output. The result of this multiphase scheme is that each phase can have a lower inductance for faster transient response. But, because of cancellation, the summed output-ripple current does not increase, and, therefore, output-ripple voltage can be minimal. Over the past five years, the inductance that most desktop-system applications require has decreased from approximately 600 nH to as little as 160 nH. Because of the decreased inductance per phase, the inductor in each phase sees increased peak and ripple currents, which cause additional power losses in the inductor core— $P_{\text{CORE}} \propto \Delta I^2$ —and associated thermal issues.

The fourth technical challenge is that the regulator must be as close to the processor as possible. This requirement limits stray inductances and losses in regulation that result from having long traces between the power regulator and the processor input. To meet this challenge, the inductors must fit underneath the processor's overhanging heat sink. This requirement limits the inductor height to less than 10 mm, making impractical the use of tall, vertically mounted toroids.

Finally, for a regulator to regulate the current it delivers, it must accurately measure this current through some circuit element and feed this information back through the control loop. In the past, it was common to use current-sense resistors, which had tightly controlled resistance tolerances and minimized stray inductance and capacitance. It was relatively simple to measure the voltage drop across these elements: $I_{\text{DC}} = V_{\text{DROP}}/R$. However, use of the current-sense resistor increases power losses— $P = I_{\text{DC}}^2 \times R$ —and adds cost.

ALTERNATIVE CURRENT-SENSING SCHEMES

Designers have tried several alternative current-sensing schemes, but the standard approach in desktop computing is to use inductor sensing. In this scheme, the inductor's winding resistance replaces the current-sense resistor to form a so-called lossless current sensor. Unlike a current-sense resistor, however, the inductor has both a dc voltage drop— $V_{\text{DROP_DC}} = I_{\text{DC}} \times R_{\text{DC}}$ —and an ac voltage drop associated with the component's inductance. As a result, to accurately determine the dc current through the inductor, you must accurately know the inductor's resistance, which implies a tight tolerance on the inductor's R_{DC} , and you must use an RC filter to remove the voltage's ac component. Tuning the RC filter to the inductance




Figure 2 A distributed-gap powder core exhibits a soft-saturation characteristic.




Figure 3 Reducing the number of turns greatly increases the core losses. The increased core loss and the variability of the inductance make the use of high-permeability powder cores impractical.

requires tight control over the inductance's value and tolerance. The need for tight tolerances and known nominal values complicates the inductor design. In addition, the use of inductor sensing requires that you must maintain the inductor resistance at some minimum value; otherwise, the voltage-drop— $V_{\text{DROP}} = I_{\text{DC}} \times R_{\text{DC}}$ —signal will disappear within the measurement noise and offset.

These technical challenges—increased current, resulting in additional losses; faster transient response; and tighter output-voltage ripple—require adopting multiphase architectures, which further stress the inductors. The use of inductor sensing requires tighter inductance and resistance tolerances, and the need for more critical component placement dictates a maximum inductor height. These new constraints have forced designers to look at new types of inductors (Figure 1).

DESKTOP VOLTAGE REGULATORS

Historically, desktop voltage regulators employed high-permeability, low-cost iron-powder cores wound with a single strand of magnetic wire on a vertically mounted toroid (Figure 1a). These inductors are cost-effective, occupy limited board space, and, because of the soft-saturation characteris-

tic of the distributed-gap powder core (Figure 2), tolerate unexpectedly high transient or peak currents. These parts have relatively high inductance of 0.8 to 1.4 μH , whose value varies greatly from light load to full load. As transient requirements increased and designers employed the multiphase architecture, these high-permeability cores were not suited to the lower inductance requirements. Essentially, to gain a low inductance from a high-permeability core, designers need to reduce the number of turns, $L \propto N^2$. However, reducing the number of turns greatly increases the core losses. The increased core loss and the variability of the inductance make the use of high-permeability cores impractical (Figure 3).

The next approach designers employed was the use of a low-permeability powder core in the same vertically mounted package. Such cores exhibit much less inductance swing with varying load and reduce core losses by approximately 33%, but their cost is approximately 80% greater than that of the high-permeability cores.

The introduction of restricted component height—to place the inductors close to the processor and under the overhanging heat sink—made vertically mounted toroids infeasible. The most obvious way to reduce the component height is simply to turn the vertically mounted inductor on its side and make it a horizontally mounted toroid (Figure 1b). The horizontal toroid's 55% increase in footprint (Table 1 and Figure 4) initially made this approach unpalatable. Instead, designers proposed a power-cube inductor (Figure 1c). Depending on the vendor, the industry refers to such components as either green or black cubes. This inductor comprises a soft-saturating, low-permeability powdered-iron shaped core with properties similar to those of the toroid core but with a different shape. This type of inductor meets the reduced-height requirements, increasing the footprint by only 22% over that of the vertically mounted toroid. Moreover, because the inductor can be machine-wound on a tightly dimensioned mandrel, the approach can significantly reduce the resistance tolerance. This inductor's drawback is that its overall losses are only 12% lower than those of the low-permeability vertical toroid, yet it costs 39% more and almost 2.5 times as much as the original high-permeability vertical toroid.

Whereas the power cube was adequate, the industry was determined to find a lower cost approach. In addition to the main switching inductor, most desktop-processor applications use an input filter. Because this filter sees little ac-ripple current, a low-cost and effective configuration is a rod-core inductor (Figure 1d). A rod core comprises a cylindrical rod of ferrite with a coil placed over it. Although ferrite cores

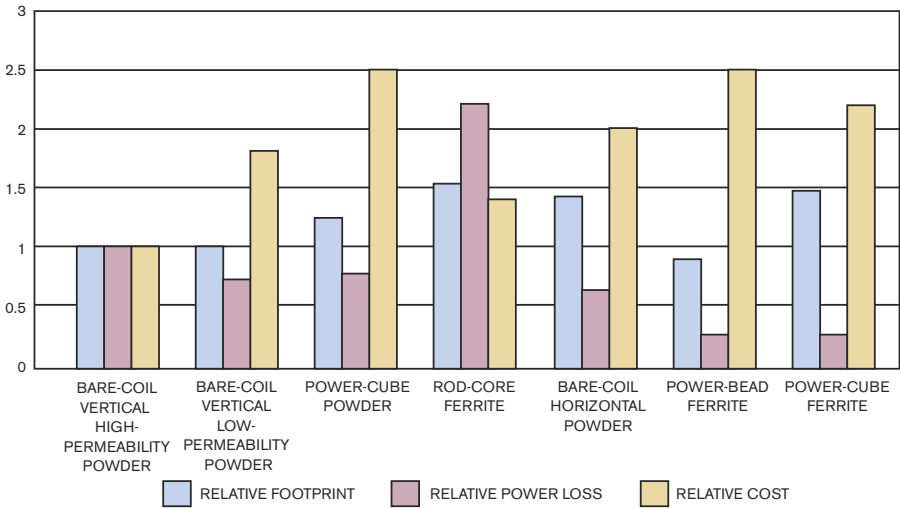


Figure 4 This chart compares the footprint, efficiency, and cost of seven inductor constructions using the older, vertical bare coil as the baseline.

exhibit much lower core losses than do powdered cores, they require an air gap within the flux path to store energy. The air gap for a rod core comprises the air around the component. Consequently, the component does not contain the magnetic path. On the surface, this approach appears acceptable. The core loss and dc-winding loss are low, and the cost is 56% less than that of the power cube, but the footprint increases by 25%. However, unlike an input filter, the main switching inductor can see more than 15A of ac current. Any ac current produces an ac magnetic field. In the case of a rod core, the core does not contain this magnetic field. Any uncontained or stray ac magnetic field induces eddy currents in nearby metal, such as the inductor winding, pc-board traces, and capacitor bodies. These eddy currents create unpredictable power losses throughout the circuit. In customer testing, the rod core, though less expensive, produces 2 to 3% lower overall circuit efficiency than the power cube. (That is, there was as much as 5W of additional loss that the inductors cause.) Moreover, the induced eddy currents created noise on nearby signal traces, thus complicating system control. Although you can use rod cores for processor power, the reduction in efficiency and the unpredictable effect of the magnetic fields on control signals typically make the cost savings not worth the effort.

Technical issues with the rod-core inductor and the high cost of the power-cube inductor necessitate the evaluation of alternative inductor constructions. One possibility is to re-evaluate the low-permeability-powder horizontal-bare-coil inductor. Although the industry initially rejected this approach as occupying too much board space compared with vertical toroids, the horizontal toroid is actually smaller than the rod core and only 18% larger than the power cube. From a cost standpoint, the horizontal toroid is midway between the rod core and the power cube and is therefore a reasonable choice after all. However, for current sensing, vertical and horizon-

tal toroids present the same inductance- and resistance-tolerance problems. Unlike the power cubes and rod cores, toroids are essentially hand-wound around the core, and designers cannot use precision mandrels or torque gauges. As a result, the core's dimensional tolerance and the winding tightness affect the R_{DC} tolerance. Core mechanical tolerances are typically ± 0.010 in., which alone results in a $\pm 5\%$ deviation in R_{DC} tolerance. Coupled with the variation in winding tightness, it's difficult to maintain better than a $\pm 10\%$ R_{DC} variation. Although some vendors specify $\pm 5.5\%$, it is extremely questionable whether this tolerance is achievable in mass production.

TOROID-INDUCTANCE TOLERANCE

Toroids have two main issues with respect to inductance tolerance. First, because the winding typically does not cover the entire core, leakage inductance can be quite high. This inductance essentially adds in series with the core's magnetizing inductance, increasing the inductor's nominal value. If the inductor windings are not always in the same space, this variation adds to the already-high $\pm 10\%$ core-magnetizing-inductance variation. If the application can compensate for the wide inductance and resistance tolerance and the additional footprint and relatively high power losses are unimportant, a horizontal toroid is a good choice because of its low cost.

Two other options, the power bead (Figure 1e) and a ferrite version of the power cube (Figure 1f), allow for tighter inductance and resistance tolerances, much lower power loss in the inductors, and decreased footprint. Both use ferrite cores, but, unlike in the rod core, the gap is small, and the field is contained, so no issues exist with stray magnetic fields. Ferrite offers the advantages of minimizing leakage inductance and minimizing core losses but has the disadvantage of hav-

MORE AT EDN.COM

Go to www.edn.com/ms4181 and click on Feedback Loop to post a comment on this article.

ing a hard saturation characteristic. As a result, you must know an application's absolute peak currents so that the inductor does not saturate.

The power bead enables the highest inductor efficiency (61% higher than a horizontal toroid's), smallest component footprint (38% smaller than a horizontal toroid's), tightest inductance tolerance ($\pm 8\%$), and tightest R_{DC} tolerance ($\pm 6.5\%$).

Unfortunately, you pay for these improvements, because power beads cost as much as power cubes and 25% more than horizontal toroids. The ferrite version of the power cube is less expensive than the power bead and costs only 10% more than a horizontal toroid. The ferrite power cube allows for a tight inductance tolerance ($\pm 5.5\%$), a tight R_{DC} tolerance ($\pm 8\%$), and high efficiency (57% higher than a horizontal toroid's). The ferrite power cube's footprint is larger than a power bead's and is essentially the same size as a horizontal toroid's.

Depending on a desktop-system application's key design drivers, three options appear worth investigating. For systems whose cost is key, the horizontal-toroid inductor appears to be the right choice. If, however, system performance, tolerance control, or footprint reduction are worth a 10 to 25% increase in inductor cost, the ferrite-power-cube or power-bead approaches are the best. **EDN**

AUTHOR'S BIOGRAPHY



John Gallagher is the North American field-applications engineer for the Power Division of Pulse (San Diego), where he is responsible for magnetics design and technical support for key accounts and power-supply-IC partners. He has a bachelor's degree in engineering from Dalhousie University (Halifax, NS, Canada). In his spare time, he enjoys surfing, kayaking, and hiking.

TABLE 1 PLATFORM COMPARISON

	Bare-coil vertical, high-permeability powder	Bare-coil vertical low-permeability powder	Power-cube powder	Rod-core ferrite	Bare-coil power horizontal powder	Bead ferrite	Power-cube ferrite
Dimensions (length×width×height, mm)	17×18.3×15	17×18.3×15	15.5×11×8	16.5×13×10	14.5×14×8	11.2×11.2×9	14.5×14.5×9
Nominal dc resistance (milliohms)	0.84	1.1	0.7	0.83	0.76	0.58	0.57
Peak current (A)	NA	NA	NA	45	NA	45	50
DC-resistance tolerance (%)	10	10	5.5	10	10	6.5	5.5
Inductance tolerance (%)	10	10	10	10	10	8	8
Core loss (W)	1.4	0.45	0.99	0.06	0.71	0.04	0.08
Copper loss (W)	0.76	0.99	0.63	0.75	0.68	0.52	0.51
Total loss/phase (W)	2.16	1.44	1.62	0.81	1.39	0.56	0.59
System loss (W)	8.6	5.8	6.5	3.2	5.6	2.2	2.4
System footprint (mm ²)	141	141	171	215	203	125	210
Relative cost	1	1.8	2.5	1.4	2	2.5	2.2

Note: Based on a processor-voltage regulator with 12V_{IN}, 1.2V_{OUT}, 300-kHz, four-phase, 120A-dc output.