Basic Power Inductor Design

Charles R. Sullivan, Professor, Thayer School of Engineering at Dartmouth
Copyright 2010—2021

This article explains a basic procedure for designing a power inductor. Many papers and textbook chapters on this topic offer more sophisticated methods, but it is harder to find a clear outline of a basic design process. Studying and practicing a basic design process is important for beginners to understand the fundamental tradeoffs in design, and to build intuition. For more advanced work, the basic design process is useful it avoids relying on assumptions that might not be valid with, for example, high-frequency loss effects that are ignored in the development of some more sophisticated methods, or that constrain other methods to narrow, specific cases.

The method here is not a new invention—rather it is an attempt to clearly articulate what many people have practiced and taught for many years.

The outline of the article is as follows:

- **Introduction and Design Starting with the Saturation Constraint** (p.1)
 - Fundamentals: Effect of a gap (p. 1), Constrained designs (p. 2), Design equation $LI = NAB$ (p. 3).
 - Design procedure starting with an saturation constraint (p. 3)
- **Optimizing the Number of Turns for Core/Winding Loss Balance** (p.6)
 - Core loss model (p. 6)
 - Deriving the optimal number of turns (p. 6)
 - Applicability (p. 7)
 - Procedure for Inductor Design (p. 8)
- **Appendix: Fitting the Steinmetz Equation** (p. 8)

Introduction and Design Starting with the Saturation Constraint

Effect of a gap

The weird thing about designing a power inductor is that one ordinarily puts an air-gap in the magnetic path. (Any non-magnetic, non-conductive material can be used, as it's functionally equivalent to air for the purpose of magnetic fields.) That adds reluctance and decreases inductance:

$$L = N^2 / (\mathcal{R}_{core} + \mathcal{R}_{gap}) \quad (1)$$

where N is the number of turns and $\mathcal{R}_{core}$ and $\mathcal{R}_{gap}$ are the reluctances of the core and gap. From the point of view of creating inductance, that's a bad thing; or more from a design point of view, if we know the inductance we need, increasing reluctance will require a larger number of turns to achieve that inductance, and thus will increase winding resistance and loss. So it's initially a mystery as to why anyone would do that.

The reason for including a gap is to reduce the flux and flux density in the core, which can be calculated as:

$$\Phi(t) = Ni(t) / (\mathcal{R}_{core} + \mathcal{R}_{gap}) \quad (2)$$

$$B(t) = \Phi(t) / A_c = Ni(t) / [A_c (\mathcal{R}_{core} + \mathcal{R}_{gap})] \quad (3)$$

where A_c is the cross-sectional area of the core and $i(t)$ is the winding current. Having B too high is a concern for two reasons—core loss and saturation.

So the gap size needs to be chosen avoid the two hazards of too high a flux density in the case of a small gap (3) and too low inductance in the case of a big gap (1). (The latter problem would, in practice, become the problem of needing too many turns to get the inductance we need in a given design problem.)

It starts to seem that we have a complicated design process in which there are multiple constraints and considerations. But we can simplify the design process by assuming that we are only interested in the designs that meet an inductance spec, L_{sp} . In any design, we are going to adjust something to make it meet that spec. Let's assume that we adjust the gap (and thus the gap reluctance) to meet that spec. The next section examines the implications of that constraint.

Constrained designs: reluctance chosen to meet inductance spec

If gap reluctance is chosen to meet the inductance design spec, for any number of turns, all the designs we are considering satisfy (1) and thus $\mathcal{R}_{core} + \mathcal{R}_{gap} = N^2 / L_{sp}$. We can substitute that into (3) to get

$$B(t) = Ni(t) L_{sp} / (A_c N^2) = i(t) L_{sp} / (A_c N)$$

Or, as it is more commonly written,

$$L_{sp} i(t) = N A_c B(t) \quad (4)$$

This is a funny equation. It is something that becomes true if we choose the right gap to satisfy the inductance spec, not something that describes a physical way to get the inductance we are looking for. But it's very useful, because it can show us the design tradeoffs more directly than considering (1) and (3) simultaneously. In particular, it shows that, for a given core area, inductance spec, and current spec, the real tradeoff is between number of turns N and flux density B . We can decrease B by increasing N , at the expense of higher winding loss from the larger number of turns. This conclusion is only true when we maintain constant inductance—examining (3) would lead us to the opposite conclusion about the effect of N on B . But that opposite conclusion would be for the case of constant reluctance, whereas our practical interest is in designing for a particular inductance spec, in which case we use a different reluctance with each N . The beauty of (4) is that it implicitly includes this adjustment of reluctance so we don't need to fuss with it manually.

Note that (4) is written in terms of $i(t)$ and $B(t)$, it applies equally to I_{dc} and B_{dc} , I_{peak} and B_{peak} , etc. Whatever parameter of the time varying waveform $i(t)$ you use on the left-hand side, the corresponding parameter for $B(t)$ is the correct thing to use on the right-hand side. If someone asks you, "in the equation $LI = NAB$, is I the RMS value, the peak value, or what?", the answer is that it can be whichever of those (and more) you want, as long as you use the corresponding version of B on the other side of the equation.

Applying $LI = NAB$ for Design: N vs. B Tradeoff.

In the previous section, we reduced two equations, based on inductance and flux in the magnetic circuit to one equation (4), based on the assumption that we will always adjust the reluctance value to get the inductance we want. (That is accomplished in practice by varying the gap length.) Now we proceed to use this for inductor design.

Based on $LI = NAB$ we see that we have a tradeoff is between number of turns N and flux density B . (You might ask what I mean by B . Peak? Amplitude of the AC component? Dc value? The answer is that all of

these are inversely proportional to the number of turns, so I don't need to specify yet.) Increasing N is bad because it leads to high winding losses. On the other hand decreasing N and thus increasing B is bad because it leads to high core loss and/or magnetic saturation. Thus, we'd generally like to decrease N until we run into trouble with either saturation or core losses. In some situations, saturation is the limiting factor, and we decrease N as far as possible without getting too close to saturation ("saturation limited designs"). In other cases, core loss becomes too high before we get close to saturation, and we increase N beyond what's needed to avoid saturation in order to control loss (loss limited designs). Thus there are two different design criteria. One way to deal with this is to assume that one or the other (core loss or saturation) is going to be the primary consideration, and proceed with a design on that basis, and then go back and check that the other issue with the size of B (saturation or core loss) has not become a problem.

Design procedure starting with a saturation limit

Figure 1 shows a design procedure starting by assuming a saturation-limited design. This is often easier than the alternative of considering the trade-off between core loss and winding loss as the initial step, because the saturation calculation is simpler than optimizing the loss trade-off. Methods of performing that optimization are addressed in a separate document. The steps in this procedure are described as follows:

1. *Choosing a core.* Typically one picks a standard core from a catalog. Although some magnetics design methods use formulas to help you guess the right size core to start with, those formulas tend to be rough approximations, and if you are going to iterate anyway, you can also simply guess—there is no need to guess right. (If you do want to guess right, think about other designs you've seen, and what size they are, such as the two different inductors we used in lab. Very roughly speaking, the size of an inductor should be proportional to its peak energy storage, $\frac{1}{2}LI_{peak}^2$).
2. *Finding the minimum number of turns to avoid saturation.* The core material data sheet will have a value for the saturation flux density. That is the point at which the core is completely saturated, and the inductance is the same as you get with no core at all.

That is way past the point at which you would want to operate to avoid any saturation effect.

Choosing to design for flux density equal to that value would be like looking at the datasheet for a polymer material, and deciding it is okay to operate at a temperature where it has not only started to soften, but has been incinerated, leaving only ashes. Unfortunately, core data sheets do not

Figure 1: design procedure starting with the assumption that the design will be saturation limited, and then checking this assumption

typically list a maximum recommended operating flux density. A rule of thumb might be to use 75% of the saturation flux density. As the saturation flux density is a function of temperature, one should use a value for the highest expected operating temperature, which might be 100 to 130° C. One can also calculate more precisely what happens to the inductance as one approaches saturation [1], but that's rarely done rigorously in practice.

After selecting the maximum safe value of flux density to use (let's call that B_{MS}), one simply must make sure that the operating flux density never exceeds that. In other words, the maximum value that $B(t)$ ever gets to (let's call that B_{max}), should be less than or equal to the maximum safe value: $B_{max} \leq B_{MS}$. The maximum value that $B(t)$ ever gets to will occur when $i(t)$ is at its maximum value; so we can find B_{max} maximum from $LI = NAB$, using in this case the version in which B and I are both the maximum values the functions of time ever hit $LI_{max} = NAcB_{max}$. Note that I_{max} may occur during some transient or other abnormal condition, and may be higher than the peak value during normal study-state operation.

3. *Finding winding loss.* Winding loss would be simply $I_{rms}^2 R$, except that many windings have frequency-dependent resistance, and many waveforms in power electronics applications include many frequencies, such as harmonics of the switching frequency, and often also include dc. The frequency dependent resistance effects are an advanced topic not covered in this document, but it should be noted that they are more complex and severe than would be predicted by considering only the skin effect. For the purpose of understanding basic inductor design, it is useful to first ignore these effects. This could be valid in cases in which the dc component of the current dominates I_{rms} , or in cases when special techniques, such as the use of litz wire, are used to control high-frequency loss effects in windings.

Resistance can be calculated for a given wire size and number of turns by looking up the resistance per unit length on a wire table, looking up the average length of a turn on a given core on the datasheet of a bobbin for that core, and multiplying by the number of turns to get the total length of wire. Particularly if one is implementing this procedure in a spreadsheet or computer program, it can be useful to calculate resistance for a given diameter from $R = \ell \rho / A_t$ where ρ is the resistivity of the conductor (Cu or Al) at the expected operating temperature, A_t is the cross-sectional area of the wire (the area of a circle with diameter equal to the diameter of the wire), and ℓ is the total length of wire, and equal to the average length per turn times the number of turns.

The wire size is often chosen to be the maximum that can fit with the specified number of turns. This results in the lowest dc resistance, although it does not necessarily result in the lowest ac resistance. The number of turns that can fit can be estimated based on the available area for the winding, called the *window area*. This value can typically be found on the datasheet of the *bobbin* to be used. The area available for an individual turn of wire would be the window area divided by the number of turns, except that when wire is packed into the window, there is still lots of air space between turns, and the insulation also takes up space. These factors can be estimated by using a "packing factor" F_p such that $A_t = F_p A_w / N$, where A_w is the window area available for the whole winding. A safe value for packing factor for a hand-wound coil is 0.5. With care, 0.6 can usually be achieved. Professionals using coil winding equipment can achieve higher values, but without special techniques or specially shaped the wire, the theoretical limit is ideal square packing, which would achieve $F_p = \pi/4$.

4. *Finding core loss.* Usually this is a matter of looking this up on a chart on a core-material data sheet. Unfortunately, core-material data sheets assume purely sinusoidal excitation, and most power electronics uses other waveforms. Techniques exist to approximately correct for this [2], but that is another advanced topic beyond the scope of this introduction, and so we assume here that we can simply use the data-sheet chart. To use the chart, we need to do the frequency and a value of B . But there's a problem: we are assuming that the flux density is a sinusoidal function of time (perhaps with a dc offset), but we need a single numerical value to apply the chart. There are many different parameters of a sinusoidal that we might use. We could use the peak value, or the RMS value, or the peak-to-peak value, or the dc value, and very few core material data sheets specify which of these to use. There is no way to figure out which to use from fundamental physics, because it is merely a convention for how these data sheets are expressed. So that leads to the one piece of information in this whole document that is something that you couldn't derive yourself – the one truly valuable piece of information here. Ready? **The value used on the datasheet is one half of the peak-to-peak amplitude.** That's the same as the zero-to-peak amplitude of the ac component of the waveform, and is often written as $\hat{B}$.

To find $\hat{B}$, recall that $LI = NAB$ works for any parameter of B as long as we use the same parameter of I on the other side of the equation. So $\hat{B} = \frac{L\hat{I}}{NA_c}$, where $\hat{I}$ is one half the peak-to-peak amplitude of the current (a.k.a. The zero-to-peak amplitude of the ac component of the current).

The core loss from the chart is provided in loss per unit volume, and then you need to multiply by the volume of the core. You could calculate the volume from the dimensions, but usually the volume is given directly, which makes this much easier. But you didn't need to read this to figure that out. The only special knowledge you need is what parameter of B to use. That's why that information is in bold above.

5. *Comparing core and winding loss.* If these are similar, the design is close to optimal for the trade-off between core loss and winding loss. If the core loss is much smaller than the winding loss, it might seem that there would be an opportunity to re-balance losses, but if you got to this point following this design procedure, there is no such opportunity, because you are already at the saturation limit, and cannot increase the flux density or decrease the number of turns. However, if the core loss is much larger than the winding loss, this means that there is an opportunity to reduce total loss by increasing the number of turns. A larger number of turns could reduce the core loss more than it increases the winding loss. One could figure out how much to increase the number of turns by trial and error, or by using the optimization method described in the second half of this article.
6. *Checking total loss.* If there is a loss specification, one may simply compare the sum of core and winding loss to the specification to determine whether the design meets the spec. The loss also determines the temperature rise, and with a value for thermal resistance, one may determine whether the loss results in an overheating problem. The maximum temperature would be limited either by the range where the core performs well, or the maximum temperature beyond which the insulation materials are damaged. In many cases, there is no definitive loss spec, and one may wish to evaluate the performance and cost of several designs before making a final decision. Note that if the loss is well below the specification, it may be worth considering using a smaller core to reduce cost (and the size and weight if these are important).

7. *Gapping.* The use of $LI = NAB$ was predicated on the assumption that we would use whatever gap length was necessary to get the desired inductance. So the final step is to calculate or measure the gap length needed. If you do this by calculation, the first step is to calculate the necessary reluctance from (1), and then calculate the gap length needed to achieve that reluctance. The standard reluctance formula for a rectangular block gives a reasonable approximation for gaps that are short compared to their lateral dimensions, but as the gap length gets up to a significant fraction of the lateral dimensions, corrections for "fringing" effects are necessary for an accurate calculation. A simple and crude approximation is to use an effective area for the gap reluctance calculated by adding the gap length to the lateral dimensions of the gap. More precise corrections are described in [3], but in many cases, a crude approximation is the only calculation needed, and the actual length needed can be found experimentally.

Optimizing the Number of Turns for Core/Winding Loss Balance

A core-loss model for optimization

Increasing the number of turns in a winding will increase winding loss (proportional to N^2), but decrease core loss (assuming the gap is adjusted for each new number of turns to maintain inductance at the specified value). That indicates that there is likely to be a number of turns that minimizes total loss. Let's figure out what that is. To do that analytically, rather than by iteratively trying numbers of turns, N , we need an expression for the core loss. The manufacturer's data shows core loss at a particular frequency, as a function of flux density, to be a straight line on a log-log plot, indicating that a power law will fit the data well:

$$P_V = k_f B_a^\beta \quad (1)$$

where B_a is half the peak-to-peak flux density. (This is *not* the same as the peak flux density, B_{pk} .) This approximation was invented by Steinmetz in 1891 and is called the *Steinmetz equation*. The constant k_f is valid for only one frequency—a different value is needed for each frequency. The best-fit value of β may also be different at each frequency. Typical values of β for ferrites are in the range of 2 to 3.

The term *Steinmetz equation* is also often applied to a power law as a function of frequency *and* flux density:

$$P_V = k B_a^\beta f^\alpha \quad (2)$$

which is also a widely useful approximation. However, plots of loss vs. frequency are not as close to straight lines, so (2) is not as good an approximation (1).

The Appendix describes a simple approach to finding the Steinmetz model parameters from plots or tables of loss data, but first we'll discuss applying such a model to finding the optimal number of turns.

Deriving the optimal number of turns

From $LI_a = NAB_a$ we have $B_a = LI_a / (NA)$, and so

$$P_V = k_f (LI_a / A)^\beta N^{-\beta} \quad (3)$$

Thus we see that the dependence on N is $N^{-\beta}$.

Knowing how core loss and winding loss vary with N , we could write total loss as

$$P = C_c N^{-\beta} + C_w N^2 \quad (4)$$

We could find the two constants, C_c and C_w , and then proceed to optimize N . Or we can start with an initial design, and consider scaling N by a factor M , i.e., $N_{new} = MN_0$. We then have, in terms of the core and winding loss of the initial design, $P_{C,0}$ and $P_{W,0}$, respectively,

$$P = P_{C,0} M^{-\beta} + P_{W,0} M^2 \quad (5)$$

To find the optimum value of M , we take the derivative and set it equal to zero:

$$0 = -\beta P_{C,0} M^{-\beta-1} + 2 P_{W,0} M \quad (6)$$

Solving for M :

$$\beta P_{C,0} M^{-\beta-1} = 2 P_{W,0} M \quad (7)$$

$$\frac{\beta}{2} P_{C,0}/P_{W,0} = M^{\beta+2}$$

$$M_{opt} = \left(\frac{\beta}{2} \cdot \frac{P_{C,0}}{P_{W,0}} \right)^{\frac{1}{\beta+2}} \quad (8)$$

$$\text{and so } N_{opt} = M_{opt} N_0 = \left(\frac{\beta}{2} \cdot \frac{P_{C,0}}{P_{W,0}} \right)^{\frac{1}{\beta+2}} N_0.$$

We can also multiply (7) through by M to obtain

$$P_{C,0} M^{-\beta} = \frac{2}{\beta} P_{W,0} M^2 \quad (9)$$

or

$$P_{C,opt} = \frac{2}{\beta} P_{W,opt} \quad (10)$$

A common rule of thumb is to make core losses equal to winding losses. This shows that that's only exactly correct in the case that $\beta = 2$. In a more typical case of $\beta \approx 2.5$, the core losses should be 4/5 the winding losses to minimize total loss.

Applicability

The results derived in the previous section for the optimal ratio of winding loss to core loss and the optimal number of turns apply as long as the resulting design does not result in saturation at the maximum current the inductor is required to handle. If a check of saturation fails, one must increase the number of turns to meet that requirement. This will increase the winding loss and decrease the core loss, resulting in an overall increase in loss.

The above analysis is based on several other assumptions:

- High-frequency winding loss effects (skin effect and proximity effect) have been neglected. The results also hold if the ac resistance factor (ratio of ac resistance to dc resistance) is constant, independent of the number of turns, which would be true with some kinds of winding designs optimized for ac resistance after the number of turns is chosen.
- Manufacturers' core loss data and the Steinmetz power law model are based on sinusoidal waveforms, but we are often interested in other waveforms. A power-law model still works well, but the correct constants for that power law will be different.

Procedure for Inductor Design

The result above can be applied in an inductor design procedure as shown in this flowchart. It is essential to check for saturation, because in some cases, particularly in an inductor with a large dc current component, or at low frequencies where core loss is small, the value of N calculated above would lead to saturation.

Another option is to start with the assumption that the design is limited by saturation, find the minimum number of turns to avoid saturation, and only optimize N if necessary, following the flowchart in the first section of this article.

Appendix: Fitting the Steinmetz Equation

I. Single Frequency

To find parameters for a single-frequency Steinmetz model (1), k_f and β , All that is needed is a pair of loss data points at different flux densities (but the same frequency). If you have a plot, look for points at the intersections of grid lines so it's easy to read them precisely. (If you do need to estimate values between grid lines, note that on a log scale, it's easier to do this accurately with values in the range of 7 to 10 than with values in the range of 1 to 3.)

The ratio of the power loss density at these two points is,

$$\frac{P_V(B_{a1})}{P_V(B_{a2})} = \left(\frac{B_{a1}}{B_{a2}} \right)^\beta, \quad (11)$$

so we have

$$\beta = \frac{\log\left(\frac{P_V(B_{a1})}{P_V(B_{a2})}\right)}{\log\left(\frac{B_{a1}}{B_{a2}}\right)}. \quad (12)$$

Then the value of k_f can be found simply by plugging in the data for one point.

If the data follows the equation exactly, this approach is adequate. If the data is noisy, or deviates from this law, fitting the data to only two points will be less accurate than using more data points. Two approaches to this are:

1. Plot the data on a log-log scale, and graphically choose lines that fit the portion of the data that you are most interested in. The record a two points from the line, not from the original data, to use in the above equations.
2. Take the log of the data, and perform a linear fit to this data in log space.

II. Incorporating frequency dependence

To derive a power law for frequency dependence as well as amplitude dependence (2), we need, in addition to a pair of loss data points at different flux densities but the same frequency, a pair of points at different frequencies but the same flux density. The value of β is found from the first pair of points as above (12). Then similarly, for the two points at different frequencies,

$$\frac{P_V(f_1)}{P_V(f_2)} = \left(\frac{f_1}{f_2} \right)^\alpha, \quad (13)$$

so we have

$$\alpha = \frac{\log\left(\frac{P_V(f_1)}{P_V(f_2)}\right)}{\log\left(\frac{f_1}{f_2}\right)}. \quad (14)$$

Again, the value of k can be found simply by plugging in the data for one point.

Although the assumption that P_V is proportional to B^β is very good, the proportionality to f^α is less accurate. To get a good fit in the region of interest it is important to choose data points in the frequency range of interest. Examining a loglog plot of P_V vs. f is useful for determining how wide a frequency range can be modeled using one set of parameters.