

THE FAMILIAR THREE-ELEMENT MODEL OF AN INDUCTOR HAS SOME SERIOUS SHORTCOMINGS, EVEN AT FREQUENCIES WELL BELOW THE INDUCTOR'S SELF-RESONANT FREQUENCY. A MORE USEFUL MODEL CORRECTLY PREDICTS AN INDUCTOR'S BEHAVIOR OVER A RANGE OF FREQUENCIES.

RF-inductor modeling for the 21st century

OF THE THREE BASIC passive components—resistors, capacitors, and inductors—inductors are the least ideal and, therefore, the most difficult to model. The importance of a good model is that it not only tells you which inductor is best for your application, it also tells you how that inductor will behave in your application.

Since early in the last century, engineers have recognized the effect of distributed capacitance within an inductor and the resulting self-resonant frequency (Reference 1). The resulting simple model of a single shunt capacitance across a lossy inductor is therefore nearly 90 years old. This model needs to be replaced because it does not, and never has, correctly modeled an inductor.

STUDY THE OLD MODEL

The old model is easy to understand (Figure 1). You measure the inductance at some frequency well below the SRF (self-resonant frequency). You calculate the capacitance from the SRF using the following formula:

$$C = \frac{1}{(2\pi \bullet \text{SRF})^2 \times L}$$

An important note to make here is that SRF means *self*-resonant frequency; it does not mean *series*-resonant frequency. The confusion comes about because a capacitor is series-resonant at its *self*-resonant frequency. An inductor is parallel-resonant at its self-resonant frequency. The resistance in the model is not the dc resistance of the inductor but, rather, some different value that more accurately simulates the response around the operating frequency.

This model correctly predicts that the measured inductance will increase as the frequency ap-

proaches the SRF, hence the earlier statement about measuring the inductance well below the SRF. The usual formula for the inductance is as follows, where HF is high frequency and LF is low frequency:

$$L_{\text{HF}} \approx \frac{L_{\text{LF}}}{1 - (f/\text{SRF})^2}$$

Obviously this formula “blows up” at the SRF, but its main use is to estimate how low a frequency you must use to correctly measure the “low-frequency inductance.” For example, at $f = \text{SRF}/10$, the formula predicts an error of only 1% in the measured value.

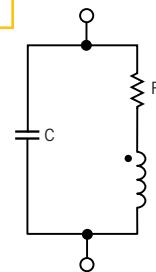
Notice that the model has the capacitor across the series RL circuit. It is unacceptable to put the capacitor directly across the inductor, although you may occasionally see such models. A moment's thought reveals that putting the capacitor directly across the inductor gives infinite impedance at self-resonance, which is not a useful model.

So far, this old model seems to be workable, particularly with the often-stated rule that it is applicable up to only some arbitrary frequency, such as $\text{SRF}/5$.

OLD MODEL RUNS INTO DIFFICULTIES

It is impossible to cite one overriding parameter of an inductor that takes precedence over all others. For example, if you are designing an oscil-

Figure 1



The 90-year-old, three-element model of an inductor is a single shunt capacitance in parallel with a lossy inductor.

lator or a tuned circuit, Q is the most important parameter. If you are using the inductor as an RF-blocking component, however, you are interested only in its impedance magnitude versus frequency. If you are using it in a filter, you want the inductance to remain constant with frequency.

Thus, there are three key facets of the inductor to consider: its SRF, its Q as a function of frequency, and its impedance as a function of frequency. A useful model must correctly simulate all three facets over several decades of frequency. This model then provides all the necessary information as the operating frequency changes.

In the old model, it is a simple matter to calculate the peak-impedance frequency; it occurs at self-resonance, where, perhaps less obviously, the Q falls to zero. It takes some more serious manipulation to discover that the peak Q occurs at $SRF/\sqrt{3} \approx SRF/1.7$.

In practice, manufacturers' typical data states that Q peaks at a point much lower than this $\sqrt{3}$ factor. The range is at least 2 to 50, and a range of 3 to 10 is quite normal. This range is, of course, from inductor family to family and does not represent the variation of a particular part. It is impossible to deduce the variation from part to part because the data is unavailable.

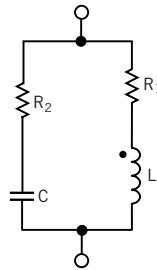
So, the old model does not correctly predict the position of the peak Q. Also, the model fails miserably above self-resonance, where it predicts that the Q rises monotonically with frequency to infinity. This failure is not very useful either.

Putting a resistor in series with the capacitor handles the high Q above self-resonance (Figure 2). Interestingly, this approach also allows for variation of the peak Q frequency, so that the model can fit a Q peak frequency up to around one-fifth the SRF. However, the model is not good in the sense that the resistor in series with the capacitor rapidly gets up into the 1-kΩ region, and the curve fit between the predicted Q and the measured Q is not particularly close.

CONSTRUCT A BETTER MODEL

In contrast with the three-element model of Figure 1, which is easy to deal

Figure 2



Simply adding a resistor in series with the capacitor improves the model in Figure 1.

with mathematically, the four-element model of Figure 2 is somewhat unpleasant, to say the least. A more interesting model uses both series and shunt resistors (Figure 3). The resistor in series with the inductor gives a Q that increases with frequency. The resistor in parallel with the inductor gives a Q that decreases with frequency. The combination is therefore widely adjustable to give a peak Q at any arbitrary frequency. If this frequency is a factor of 10 or more away from the required SRF, then the self-capacitance will have little effect on the established peak Q.

The beauty of this model comes from the simplicity of the equations, where Q_{MAX} is the maximum value of Q:

$$Q_{MAX} = \frac{1}{2} \sqrt{\frac{R_2}{R_1}}$$

and

$$f_{Q(MAX)} = \frac{\sqrt{R_1 R_2}}{2\pi L}$$

You can rearrange the formulas to give the resistor values directly from the Q_{MAX} value and its frequency position:

$$R_1 = \frac{2\pi f_{Q(MAX)} L}{2Q_{MAX}}$$

and

$$R_2 = 2Q_{MAX} \times 2\pi f_{Q(MAX)} L$$

As the Q_{MAX} frequency approaches the SRF, there is a strong interaction with the self-capacitance. You should, therefore, expect some degree of iterative tuning of values. However, working out the math for

this new five-element model is not something for the fainthearted to contemplate.

A PHYSICAL MODEL GETS REAL

So far, the presented model has been a mathematical exercise, but you need to answer a question about the physical inductor: What components, in the correct place, give the right sort of response? A more scientific basis for the model results when you use the known characteristics of the constituent parts of a real inductor. Of course, the construction of an inductor is generally so simple that it is hard to imagine the specific constituent parts. However, you can highlight two specific parts: the core and the conductive path.

The conductive path is usually copper, either in wire form or as a thin layer of metal. The century-old “skin effect” then comes into play, and at some radio frequency, the resistance of this path will increase with the square root of frequency. You could say, for argument’s sake, that this relationship between resistance and frequency happens when the skin depth is one-third the wire diameter or, alternatively, one-third the conductor thickness.

For isolated copper conductors, the skin depth at various frequencies is simple to calculate (Table 1). The skin depth becomes more complicated when turns in a winding are close to each other. This “proximity effect,” however, concerns details that are beyond the scope of this article.

The core loss is another matter. For an air-cored inductor, there is no core loss. For a UHF inductor wound on a ceramic or other nonmagnetic core, the core loss should also be minimal. On the other hand, for iron or ferrite cores, the loss will be at least proportional to frequency. Analysis of ferrites used in switched-mode power supplies, for example, suggest that for the core loss, the power law of frequency is 1.3 to 1.5.

Theoretical analysis of single-layer

TABLE 1—FREQUENCY VERSUS SKIN DEPTH

Frequency (kHz)	Skin depth (μm)
0.050	9300
5	930
500	93
50,000	9.3

coils suggests that the self-capacitance is a function of the size and shape of the coil and is unrelated to the number of turns (Reference 2). This approximation is good when there are approximately 20 or more turns equally spaced along the length but becomes increasingly less valid as the number of turns decreases.

A good example of this constant self-capacitance rule is in the published data on Ohmite (www.ohmite.com) KM-series molded RF inductors, which have phenolic cores (Table 2). This self-capacitance rule also applies to iron-cored inductors, such as the 9230 series of molded RF chokes from JW Miller Magnetics (www.jwmiller.com). The chokes in the range of 1.2 to 10 μH all have a self-capacitance of approximately 1 pF.

When the number of turns is very small, the self-capacitance within a family of inductors increases with the number of turns, as shown by the SRF data on Coilcraft (www.coilcraft.com) Mini Spring air-cored inductors (Table 3).

The proof of any model is whether or not it agrees with the real-world measured response. Unfortunately, most standard Spice-based simulators cannot directly simulate resistors whose values are a function of frequency. On such simulators it would therefore be necessary to write a simulation script to step the frequency, change the resistor values, and plot the results. This process is not necessarily complicated, but it can certainly be time-consuming. At least three Spice simulators can handle resistors that are a function of frequency: PSpice (www.orcad.com), ICAP (www.intusoft.com), and SpiceAge (www.thoseengineers.co.uk).

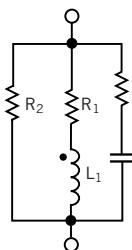
CREATE A MATHEMATICAL MODEL

The first thing you need to do for simulation purposes is decide how to plot Q against frequency. You may remember the definition of Q for a series resonant circuit:

$$Q = \frac{\omega_0 L}{R}$$

You can represent a general impedance as $Z=R \pm jX$. Thus, you can write the Q of any general component as

Figure 3



The combination of the resistor in series with the inductor, which gives a Q that increases with frequency, and the resistor in parallel with the inductor, which gives a Q that decreases with frequency, results in a widely adjustable combination that gives a peak Q at any arbitrary frequency.

$$Q = \frac{X}{R} = \frac{|\text{Im}(Z)|}{\text{Re}(Z)},$$

where Re(Z) and Im(Z) are the real and imaginary parts of Z, respectively. This formula allows you to write a script to plot Q.

An alternative approach is to simulate the inductor mathematically using a general math solver, such as Mathcad.

Although the math is complicated, it is not complicated to just write down the equation:

$$Z = \frac{1}{\frac{1}{K\sqrt{f} + j2\pi fL} + \frac{1}{R_{HF} + 1/j2\pi fC}}$$

This equation represents just two series circuits, wired in parallel. The resistor, R_{HF} , in series with the capacitor, modifies the frequency position of the maximum Q point. It is difficult to use R_{HF} to tune the response of the model above self-resonance because manufacturers' data seldom even goes up to the SRF, let alone beyond it.

This model of an inductor is essentially the same one that Coilcraft uses to model its inductors. The only difference is that Coilcraft adds in a small extra fixed resistor in series with the model given above. Unfortunately, the Coilcraft model does not fit the Coilcraft-published Q data very well. However, the published typical data is ac-

curately plotted and does provide useful data against which you can compare an alternate model.

You can achieve a better model by using a different variation with frequency for the series loss resistance, which the impedance equation represents using the $K\sqrt{f}$ term. This approach then causes another problem with the model; the dimensions of K become bizarre. A good approach to this problem is to normalize the loss resistor to the self-resonant frequency. For any inductor worth modeling, it is a good approximation to say that the self-resonant frequency is simply:

$$\text{SRF} = \frac{1}{2\pi\sqrt{LC}}$$

You can then conveniently write the impedance equation as follows, where $0.5 \leq \eta < 1$, and R_S is the series loss resistance:

$$Z = \frac{1}{\frac{1}{R_S(f/\text{SRF})^\eta + j2\pi fL} + \frac{1}{R_{HF} + 1/j2\pi fC}}$$

You can adjust this model to accurately fit the manufacturer's data, but it also has a couple of extra qualities. For example, you can change the value of η slightly without completely destroying the previous response. If you used η on the un-normalized frequency, the change in K necessary to bring the model back close to its previous setting would be large. In other words, there is a huge "gearing" effect on the parameters in the un-normalized model. Another benefit is that as the self-capacitance changes, the effective loss resistance changes appropriately, and the model seems more realistic relative to an actual production spread.

AN APPLICATION EXAMPLE

The model of a 0805HQ16N inductor (Coilcraft) is a good example because there is a big curve of its typical

TABLE 2—CONSTANT SELF-CAPACITANCE RULE

Inductance (μH)	Minimum SRF (MHz)	Calculated maximum self-capacitance (pF)
0.15	525	0.61
0.47	310	0.56
1.5	160	0.66
4.7	90	0.67

Q in the manufacturer's data sheet. The model values are: $R_{HF}=6\Omega$, $L=16.8\text{ nH}$, $C=75\text{ pF}$, $R_s=1.95\Omega$, and $\eta=0.54$. This new model follows the published typical data to within a Q of approximately ± 2 but also realistically extrapolates Q beyond the range given in the data sheet (Figure 4).

In the final application, you are really interested in the Q of an actual resonant circuit. From the graph of the inductor's Q, you might conclude that because the Q is 0 at its SRF, you can't use the inductor at this frequency. But this conclusion is wrong. You need to understand that there is a world of difference between the Q of a component and the Q of a resonant circuit. The Q of the inductor is 0 at its self-resonant frequency because it appears entirely resistive. However, if you were to plot the impedance curve given by the model, you would find a sharp peak at the SRF (Figure 5).

Ultimately, you want to know the selectivity of the resonant circuit in terms of the 3-dB bandwidth. The relationship between this bandwidth and the Q is:

$$3\text{-dB BANDWIDTH} = \frac{\text{CENTER FREQUENCY}}{Q}$$

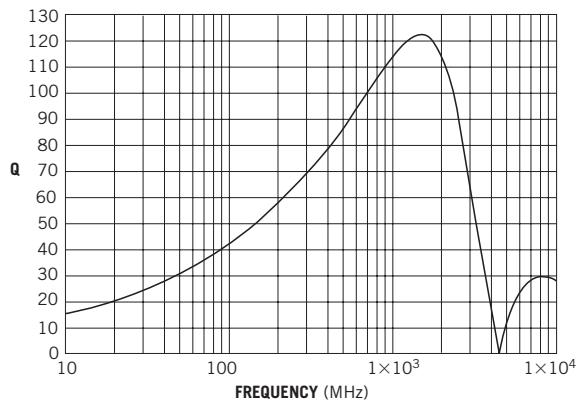
A self-resonant inductor can have a reasonable Q as a resonant circuit. Unfortunately, you can't use an inductor as a resonant circuit because you never know where this self-resonant frequency will be. Manufacturers never tell you the spread of self-resonant frequencies,

TABLE 3—SELF-CAPACITANCE VERSUS NUMBER OF TURNS

Number of turns	Inductance (nH)	Minimum SRF (GHz)	Calculated maximum self-capacitance (pF)
one	2.5	12.5	0.06
two	5.0	6.5	0.12
three	8.0	5.0	0.13
four	12.5	3.3	0.19
five	18.5	2.5	0.22

yet this figure is important to design engineers. Without this information, you can't estimate the tolerance on the self-capacitance of the inductor. Thus, you can't determine the necessary trim range for the capacitor that you use to

Figure 4



The new model follows the published data for Q to within approximately ± 2 but also extrapolates the Q beyond the range in the data sheet.

resonate with the inductor.

Now that you have a representative model of the inductor as a component, you can predict the overall Q of an actual circuit. For this representation, assume that the external capacitor has an arbitrarily high Q. In other words, assume that the capacitor is lossless.

To plot the circuit Q, you now want to calculate the circuit in terms of its parallel equivalent. You then use the low-frequency value of the inductance to evaluate the circuit Q, where R_p is the parallel equivalent resistance:

$$Q = \frac{R_p}{2\pi f_0 L} = \frac{1/\text{Re}\{1/Z\}}{2\pi f_0 L} = \frac{1}{2\pi f_0 L \times \text{Re}\{1/Z\}}$$

A comparison of the overall circuit Q to the inductor Q reveals an interesting result: The circuit Q is higher than the inductor Q (Figure 6). This result is probably contrary to what you have been taught. The point is that the inductor Q is worse because of the lossy self-capacitance. The external capacitance is considered lossless and, therefore, increases the overall Q. Note that the circuit Q above self-resonance is meaningless; you would need to resonate the inductor with another induc-

tor above its self-resonant frequency, which is a pointless exercise.

MODELING IN PRACTICE

This mathematical model is not easy to use because the values are interrelated. The iterative procedure to arrive at the values is as follows:

- Set the inductor to the low-frequency inductance value.
- Set the capacitance according to the typical self-resonant frequency. Guess at 15% higher than the manufacturers' minimum SRF in the absence of other data, where

$$C = \frac{1}{(2\pi \cdot \text{SRF})^2 \times L}$$

- Set $R_{HF}=1\Omega$ and $\eta=0.50$; then adjust R_s until the peak Q is approximately correct.
- If the peak Q is in the wrong place, you shift it by moving the SRF (adjusting C) or by changing R_{HF} . Making R_{HF} bigger moves the peak Q to a lower frequency.
- Repeat moving the peak Q frequency and then bringing the peak Q amplitude back to the correct value until you are happy with the peak.
- Now adjust the loss-correction term, η , to modify the Q values leading up to the peak. A larger value of η will make the low-frequency Q readings higher (after the peak Q has been set back to its correct value using R_s). Increasing η also slightly lowers the peak Q frequency.

Note that if the peak Q occurs at much less than SRF/5, then the model will require an additional parallel resistor term. Expect this additional resistor to be necessary for ferrite- and iron-based core materials, where the Q peak can be as many as 50 times lower than the SRF.

The actual simulation time in this process is essentially zero. The overall time necessary depends on how good you are at changing values to produce the desired result. This approach is not elegant, but it does produce a workable result. You ordinarily will not need to model that many inductors.

DRAW SOME CONCLUSIONS

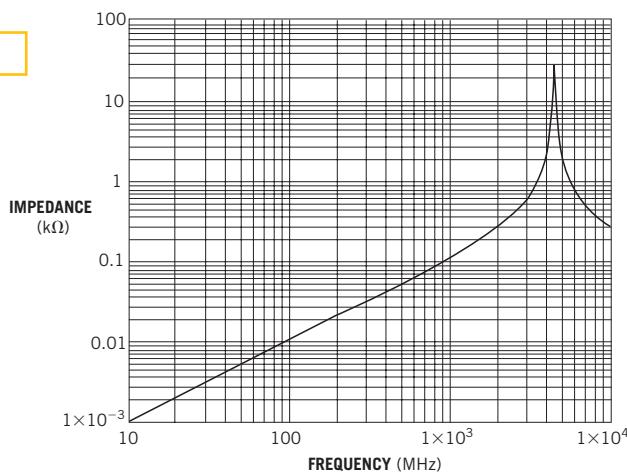
It is impossible to theoretically predict where the peak Q of an inductor occurs relative to its SRF. Manufacturers' data suggest that a figure of $SRF/f_{Q(MAX)}$ from 2 to 9 exists even for the air- or ceramic-cored inductors. Thus, the typical data on the peak Q frequency of an inductor is a vital piece of data to obtain from the manufacturer.

To get the best performance from a resonant circuit, the inductor Q needs to be as high as possible. To achieve this high Q, use an appropriate family of inductors and a value such that the Q peaks near the operating frequency.

Although inductors are available with tolerances as low as $\pm 2\%$, low tolerance does not help greatly when the inductor is running near SRF/2. The implied variation of SRF from its minimum specified value to infinity gives rise to a greater than $\pm 10\%$ uncertainty in the desired tuning frequency. It is, therefore, important for manufacturers to give both the upper and lower limits of the SRF. Without manufacturers' data, users end up inventing their own tolerances for the self-capacitance. In such cases, a figure of approximately $\pm 15\%$ is not unreasonable.

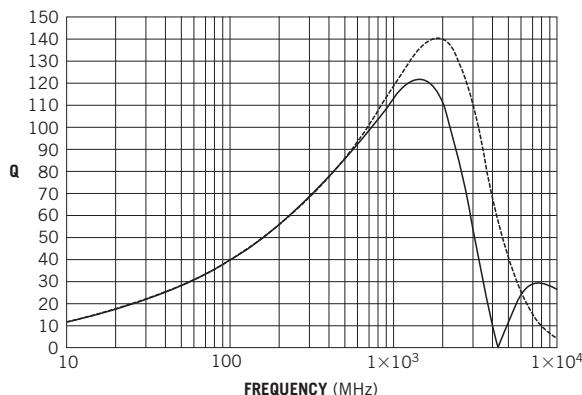
For use as an ac-blocking component, the model of an inductor need not be complicated. All you need to know is the minimum low-frequency inductance and the SRF. You know that the blocking impedance will always be greater than or equal to $2\pi fL$ to the minimum SRF. If you go beyond the SRF, you will eventually hit a series-res-

Figure 5



A plot of the model's predicted impedance versus frequency shows a sharp peak at the inductor's self-resonant frequency.

Figure 6



The circuit Q, in a dotted line, is higher than the inductor Q.

onant point at which a local minimum impedance exists. It is hard enough getting data up to self-resonance, let alone up to the first series resonance. However, you should be fairly safe if you don't subject the inductor to frequencies greater than approximately 1.5 times the quoted minimum SRF.

The key thing to remember is that if the data is not in the data sheet, you need to contact the supplier to get it. If enough engineers pester them, manufacturers will provide the data you need. □

AUTHOR'S BIOGRAPHY

Leslie Green has a bachelor's degree in electrical and electronic engineering from Imperial College (London). He has 20

years of experience designing test-and-measurement equipment and has spent the last 15 years designing digital-storage oscilloscopes for Gould-Nicolet Technologies in the United Kingdom.

REFERENCES

1. Kolster, Frederick K, "The Effects of Distributed Capacity of Coils Used in Radio-Telegraphic Circuits," Proceedings of the Institute of Radio Engineers, 1 (1913), pg 19 to 34.
2. Medhurst, RG, "HF Resistance and Self-Capacitance of Single-Layer Solenoids," *Wireless Engineer*, 24 (1947), pg 35 to 43 and 80 to 92.