

Capacitors Demystified: Practical information on capacitor usage in EE133

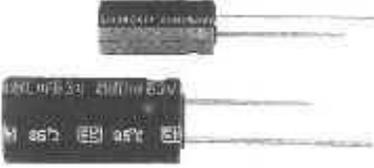
Why do I need to read this?

Despite all we know about capacitors from previous exposure (including everyone's favorite capacitor fact – 'Well, I know the voltage across a capacitor cannot change instantaneously!'), there is a quite a bit about the (good) properties and (bad) non-idealities of these devices that effects the way they are actually used in RF circuits. So, with the practical predisposition of EE133, let's take another look at our old friend...

The Starting Line-up

Like fine cheeses or candy bars, capacitors come in wide variety of sizes, shapes, and prices- and they can be made from a host of different materials (except for nougat). However, all of these characteristics are interrelated because, in classic Murphy fashion, a capacitor that ranks well in one or two of these properties is usually terrible in the other categories.

The physical size and shape are only of marginal importance to us since the former is pretty much irrelevant at the size of our solder board and the latter is only for aesthetics (which is a foreign term to EE's anyway). But, the composite material is an important factor in determining the behavior of the capacitor, so for the purpose of this class we will identify four major types (Refer to page 21 of "The Art of Electronics" by Horowitz & Hill for a more complete listing of capacitor families):

Material/Comments	Picture
Ceramics (10 pF – 22μF): Inexpensive, but not good at high frequencies	
Tantalum (0.1μF – 500μF): Polarized, low inductance	
Electrolytic (0.1μF – 75μF): Polarized, not good at high frequencies (NOTE: longer lead is usually the positive terminal)	
Silver Mica (1 pF – 4.7 nF): Expensive with only small capacitive values, but great for RF	

Some Commonly Asked Questions about Capacitors

Resistors are color-coded. How are capacitors labeled?

I have attached a sheet that describes the various notations for labeling capacitors with their value. But, the system is not nearly as well standardized as the codes for resistors, so often it is a good idea to just measure them. Especially since students often return capacitors to the wrong bins.

What was that loud popping sound and why is there goo on my circuit?

Congratulations, you have just blown up your first capacitor! ‘Popping a cap’ as it is called by EE’s (though it has certain other connotations in the rap community) is almost always the result of ignoring its polarization. Polarized capacitors are designed to only function if one lead (generally labeled by a plus ‘+’ sign) is always at a higher DC potential than the other lead (generally labeled by a negative ‘-’ sign, or an arrow with a negative sign pointing to it). Only use polarized capacitors if you can actually assure, based on the operation of the circuit, that the two nodes between which the capacitor is placed will be of the same signed potential difference for all instants in time (i.e. AC fluctuations are OK if both sides are DC biased). Otherwise, watch out for flying goo!

When is a capacitor not a capacitor?

The short answer is, when it’s an inductor! Although we like to model a capacitor as an ideal reactive element with a -90° phase shift ($Z = 1 / sC$), they do not have such pure behavior. Instead, on account of the resistance and inductance of the conductors from which it is constructed (i.e. the wires, contacts, and sometimes casing) every capacitor is better modeled as a series RLC resonant circuit (do you know why it’s not a parallel RLC?). Using this series RLC model, it becomes evident that the capacitor has a *self-resonance*.

?? Self-Resonance

Recalling our knowledge of RLC networks, we note that the impedance of the series RLC tank is dominated by the capacitor at low frequencies (approximately $1/sC$) and is dominated by the inductor at high frequencies (sL). But, at some frequency in between, typically denoted as the resonant or natural frequency of the circuit, the impedance transitions between these extremes and is purely resistive.

For capacitors, the natural frequency is called the self-resonant frequency and essentially tells you the frequency above which the capacitor starts behaving like an inductor. Why is this so important to us? Imagine taking your hard-earned design and replacing all the capacitors with inductors of random values. Then, sweep the frequency of the input and watch how it responds. I don’t think you’ll like the result.

?? Don’t believe the hype -- Size Matters

How does self-resonance vary as a function of the capacitor properties mentioned above? In general, the larger the capacitor’s value, the lower its self-resonance. For example, a 1- μ F capacitor will cross over into self-resonance at a much lower frequency than a 5-pF cap. However, the size is not the only consideration because depending on the material from which the capacitor is made, the frequency behavior

can be improved. Using our example, we would find that a 1- μ F ceramic capacitor has a much lower self-resonance than a 1- μ F made from silver mica.

For our RF work in EE133, high frequency is the name of the game and so we need to be sure our components are not inherently constrained below 24 MHz. Speaking of EE133, let's take a look at the specific ways in which we will employ capacitors in this class and how the self-resonance phenomenon effects these usages:

Some Common Applications of Capacitors

Application #1: Power Supply Bypass

As you learned in Lab 1, RF circuits can be extremely sensitive to noise on the power and ground rails. This noise (which can be caused from a variety of sources including but not limited to the circuit itself) is transmitted throughout the circuit via the rails and can wreck havoc with transistors as well as chips. Here's why: if one rail is noisy relative to the other, then instead of having, say, a 9V difference between them, it is possible for that difference (and thus the relative voltages seen by the circuit) to vary randomly with time.

Solution for Application #1:

To combat this problem, we need to install something between the power and ground rails that will not disrupt their DC differential but at the same time will effectively ensure that the noise is common mode. The obvious choice is (surprise!) a capacitor. Over a wide range of non-zero frequencies (seeing as how noise can occur at almost any frequency) a capacitor will short the two rails together and ensure that deviations in one are mimicked by identical variations in the other. This will result in a constant voltage difference between the two while reducing the noise from the circuits' perspective.

Not so fat...

While simple conceptually, this solution is tricky to execute. Ideally, we would use an infinitely large capacitor for power supply bypass. After all, we want it to appear as a short circuit at all frequencies (except DC), so to make $Z = 1/sC$ zero, we would need $C = \infty$. But, in case you haven't been to Halted lately, I'll fill you in on a little secret – they don't carry ∞ -F caps (at least not when I was there last). As engineers, this doesn't really bother us because we can just use something “big” that will approximate infinity. Typically, this is referred to as a BFC, or Big Fat Capacitor (to keep this a G-rated handout I won't mention how the “F” is sometimes translated differently). OK, so basically all I've said is that we need a BFC for power supply bypass. Sounds easy enough. Except for one small problem...

Self-resonance (again)

Most BFC's have very low self-resonant frequencies since they are typically either electrolytic (which are used because polarity is fixed in bypassing), or tantalum. This means, at 24 MHz, the BFC will look like a small inductor and its impedance will be $Z = sL$. Since Z is proportional to frequency then at 24 MHz this impedance will certainly not be zero, and that means trouble in terms of solving our noise problem. So, if we can't use a BFC because it will be a SSI (Stupid Small Inductor) at 24 MHz, what do we do?

The solution is to use two capacitors in parallel when bypassing the power supplies. One should be a BFC that will provide the desired short circuit for low frequencies while the other should be a smaller cap that will provide the short circuit at higher frequencies and will essentially overwhelm the $Z = sL$ of the BFC. Since the smaller capacitor should have a higher self-resonance, we can safely model it as a capacitor at 24 MHz. While the BFC should be electrolytic or tantalum, the smaller shunt capacitor should be either ceramic or silver mica. This brings up an important question:

The Age-old, Time-worn, Oft-repeated, Classic Question

People always ask: “What do you mean by ‘big’ and ‘small’?” The answer is that it depends. In the EE133 lab, “big” refers to caps in the $0.1\mu\text{F} - 47\mu\text{F}$ range while “small” is something in the $0.1\text{ nF} - 33\text{ nF}$ range. The best way to decide what you need is to pick a cap in these ranges and then evaluate its impedance at 24 MHz. Regardless of the BFC, the “small” cap should be in the tens of ohms because this is the ‘active’ capacitor at this frequency.

Application #2: AC coupling / DC blocking

As our signal travels through our system, passing from one stage to the next, it is often desirable to eliminate the DC component of that signal between each stage. This process, referred to equivalently as *AC coupling* or *DC blocking*, is necessary so that the DC value of the signal, which we often cannot control (or even predict), does not disturb the bias points of the next stage. This concept may be familiar from EE113 where establishing the DC bias points of your transistors was fundamental to the validity of the small signal model. If the bias points drift, the small signal parameters will no longer be valid, destroying the design.

Solution for Application #2:

So, we need an element that will block the DC component of the signal, while allowing the AC variations to pass through unaffected. If this sounds familiar, it should. But there is one major difference between AC coupling and power supply bypass. In AC coupling, we are concerned about blocking all DC signals, but we really only need to pass certain (rather than all) AC signals. In fact, we often would like to filter out AC signals at all frequencies other than those the next stage is designed to handle. So, if we can use an element that will block DC but only pass the frequencies of interest we will all set.

Before we jump into the obvious (read capacitive) solution, I should address the obvious alternative. Someone is likely to be saying “Well, it sounds like what we need is an AC bandpass filter between stages.” In fact, this is precisely the right idea, and ideally is what we would do. However, in lieu of the added complexity and number of components this would require (as well as the low probability that you could actually tune all the bandpass filters between stages to operate precisely the same frequency), we opt to leave the filtering as a function of each individual stage and simply provide the DC blocking (with a minimal amount of frequency selectivity) in between stages.

To do this, we use AC coupling (or sometimes just ‘coupling’) capacitors. Again, these should be chosen to appear as short-circuits at the frequency of interest so that the desired signal can pass through undisturbed (of course, we have the DC blocking covered regardless of what size we use). Note that we do not need to use two in parallel because we are really only concerned about passing a certain range of frequencies. Without the need for wide-band performance, we can get away with a narrow-band, single capacitor solution.

The Age-old, Time-worn, Oft-repeated, Classic Question (again)

Now, we fire the deadly question: “What size should we use for coupling?” And, using our bullet-proof answer, respond: “It depends”. Here, the dependence is even more severe, though, because instead of trying to remove noise, we are trying to capture our signal. While the presence of the former is annoying, the absence of the latter is unacceptable. So, at the frequency of interest, we not only need to be sure that our capacitors are below their self-resonance (i.e. still look like capacitors), we also must be sure they have the value we expect (so that we can accurately predict their impedance).

To illustrate, let’s explore the differences between high and low frequency use:

?? High-Frequency Example

Use your imagination (if you don’t have any, Pre-Lab #1 will do) to examine the case of an amplifier designed to work at 24 MHz. If the signal we care about is at 24 MHz, we will need to use silver mica capacitors for coupling because they have a higher self-resonance and are very accurate in terms of their value at such a frequency. The exact size can again be determined by using $Z = 1/sC$ to keep the impedance of the capacitor to 1% - 5% of the total impedance looking into/out of a given stage. If we have an amplifier whose input impedance is 200 Ω , then we can tolerate no more than 10 Ω (5%) across the capacitor. So 1-nF might be a nice choice ($Z = 6.6\Omega$ at 24 MHz). Obviously, for lower frequency signals, the impedance will be larger and thus the signals will be attenuated in voltage as they pass through the capacitor (at 100 kHz, a 1-nF cap will have an impedance of 1.6 k Ω !). Hey, lookie there -- free high-pass filtering with your paid subscription to ‘AC Coupling: The Magazine’ (sure beats a fleece pull-over)!

?? Low-Frequency Example

An example of the opposite extreme is the microphone circuit. Here, we will be dealing with signals in the audio range (roughly 0.3 Hz – 20 kHz). Clearly, if we use the same coupling caps as in the previous case, we will be in bad shape. The impedance of those smaller coupling caps will be so large at audio frequencies (at 20 kHz, a 1-nF cap will have an impedance of about 8 k Ω !) that no signal will get through. Instead, we must select larger (and, therefore, not silver-mica) capacitors to couple at low frequency. Sticking to our 1%-5% rule for the tolerable impedance, we would need a 1- μ F cap at 20 kHz (impedance of 7.9 Ω). However, 1- μ F turns out to be a very expensive capacitor, so we don’t have many in the lab. Just look for something in that range (err on the side of being too large). Ceramics and tantalum

(beware of polarity) are acceptable, and electrolytics will pass if you can't find the size you need in any other flavor.

Application #3: Resistor Bypass

Though not as common as the previous ones, this application arises when there is a need to short out a resistor at a particular frequency. Again, this situation should be reminiscent of Lab #1 and the emitter-degeneration resistor in the common-emitter amplifier. Though the resistor (R_e) is helpful for biasing, we want it to look like a short circuit in the small signal model so that it won't degenerate our transconductance. To do so, we typically place a capacitor across it. And, you may be asking yourself: "What kind of capacitor should I use?" (Or, at this point, you may just be asking when this handout will ever end.)

Solution for Application #3:

Despite its name, resistor bypassing is an application for capacitors where our solution looks more like that for application #2 than for application #1. Since the capacitor is again in the signal path, we can achieve some degree of filtering by insisting on a narrow-band as opposed to wide-band solution. A single capacitor should be chosen with constraints similar to those describe in the 'Solution to Application #2' (Typically, resistor bypassing is done to attain some benefit at higher frequency so it is likely you will need to look at the 'High Frequency Example').

Application #4: General Usage

Last, but most common, is the use of capacitors in oscillators, filters, timing circuits, impedance matches, rectifiers, and other RF circuits. Here, the function of the capacitor can vary drastically, but the general principle is always the same: you are designing under the assumption that you will have a capacitor that will be of some value at some particular frequency or range of frequencies. So, although there is no standard solution in this case, you should always select a cap that meets those expectations.

The Last Word

Keep in mind the differing self-resonant frequencies of the different sizes and types of capacitors. Attention to this detail will ensure that the capacitor is of the value you think (and, more importantly, that you don't have an inductor) at the frequencies of interest. Also, beware of polarized caps and be certain, even if they are acceptable for a given situation in which you can guarantee the polarity of the nodes, that you install them properly. Last, but not least, always measure your capacitors! Happy capping!