STUDENT MANUAL FOR

THE ART OF ELECTRONICS

THOMAS C. HAYES
PAUL HOROWITZ

Class 2: Capacitors and *RC* Circuits

Topics:

- *new:*
  - Capacitors: dynamic description
  - *RC* circuits
    - Time domain:
      - step response
      - integrator, differentiator (approximate)
    - Frequency domain: filters

Capacitors

Today things get a little more complicated, and more interesting, as we meet frequency-dependent circuits, relying on the *capacitor* to implement this new trick. Capacitors let us build circuits that "remember" their recent history. That ability allows us to make timing circuits (circuits that let *this* happen a predetermined time after *that* occurs); the most important of such circuits are *oscillators*—circuits that do this timing operation over and over, endlessly, in order to set the frequency of an output waveform. The capacitor's memory also lets us make circuits that respond mostly to changes (*differentiators*) or mostly to averages (*integrators*), and—by far the most important—circuits that favor one frequency range over another (*filters*).

All of these circuit fragments will be useful within later, more complicated circuits. The filters, above all others, will be with us constantly as we meet other analog circuits. They are nearly as ubiquitous as the (resistive-) *voltage divider* that we met in the first class.

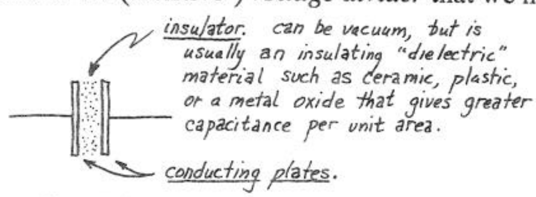

Figure N2.1: The simplest capacitor configuration: sandwich

This capacitor is drawn to look like a ham sandwich: metal plates are the bread, some dielectric is the ham (*ceramic* capacitors really are about as simple as this). More often, capacitors achieve large area (thus large capacitance) by doing something tricky, such as putting the dielectric between two thin layers of metal foil, then rolling the whole thing up like a roll of paper towel (*mylar* capacitors are built this way).

Text sec. 1.12

A *static* description of the way a capacitor behaves would say

$$Q = CV$$

where Q is total charge, C is the measure of how big the cap is (how much charge it can store at a given voltage: $C = Q/V$), and V is the voltage across the cap.

This statement just defines the notion of capacitance. It is a Physicist's way of describing how a cap behaves, and rarely will we use it again. Instead, we use a dynamic description—a statement of how things change with time:

$$I = C\, dV/dt$$

This is just the time derivative of the "static" description. C is constant with time; I is defined as the rate at which charge flows. This equation isn't hard to grasp: it says 'The bigger the current, the faster the cap's voltage changes.'

Again, flowing water helps intuition: think of the cap (with one end grounded) as a tub that can hold charge:

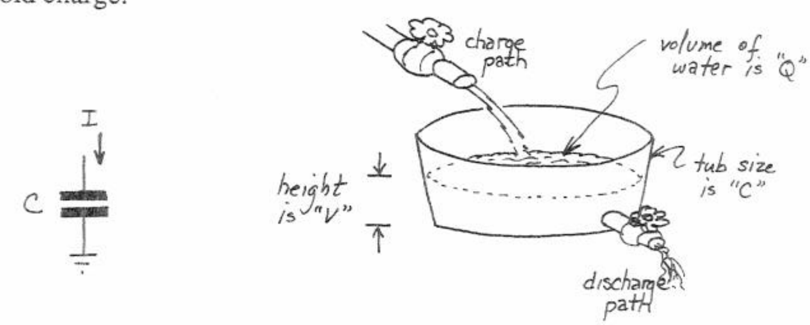

Figure N2.2: A cap with one end grounded works a lot like a tub of water

A tub of large diameter (cap) holds a lot of water (charge), for a given height (V). If you fill the tub through a thin straw (small I), the water level—V—will rise slowly; if you fill or drain through a fire hose (big I) the tub will fill ("charge") or drain ("discharge") quickly. A tub of large diameter (large capacitor) takes longer to fill or drain than a small tub. Self-evident, isn't it?

Time-domain Description

Text sec. 1.13

Now let's leave tubs of water, and anticipate what we will see when we watch the voltage on a cap change with time: when we look on a scope screen, as you will do in Lab 2.

*An easy case: constant **I***

Text sec. 1.15; see Fig. 1.43

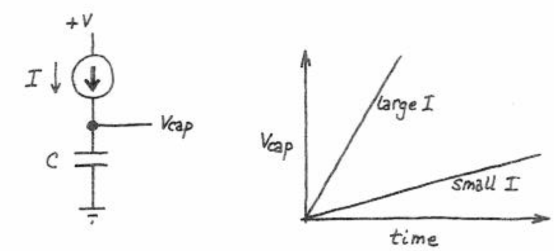

Figure N2.3: Easy case: constant I —> constant dV/dt

This tidy waveform, called a *ramp*, is useful, and you will come to recognize it as the signature of this circuit fragment: capacitor driven by constant current (or "current source").

This arrangement is used to generate a triangle waveform, for example:

Compare Text sec.1.15, Fig. 1.42: ramp generator

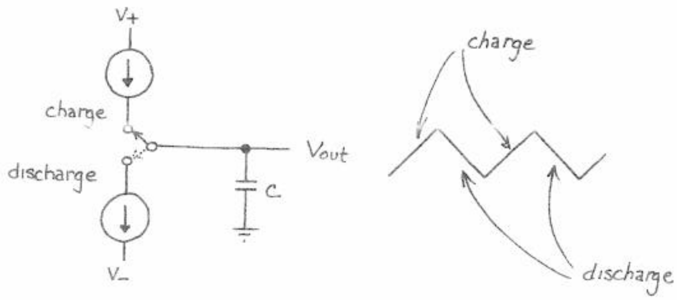

Figure N2.4: How to use a cap to generate a triangle waveform: ramp up, ramp down

But the ramp waveform is relatively rare, because current sources are relatively rare. Much more common is the next case.

A harder case but more common: **constant voltage** source in series with a resistor

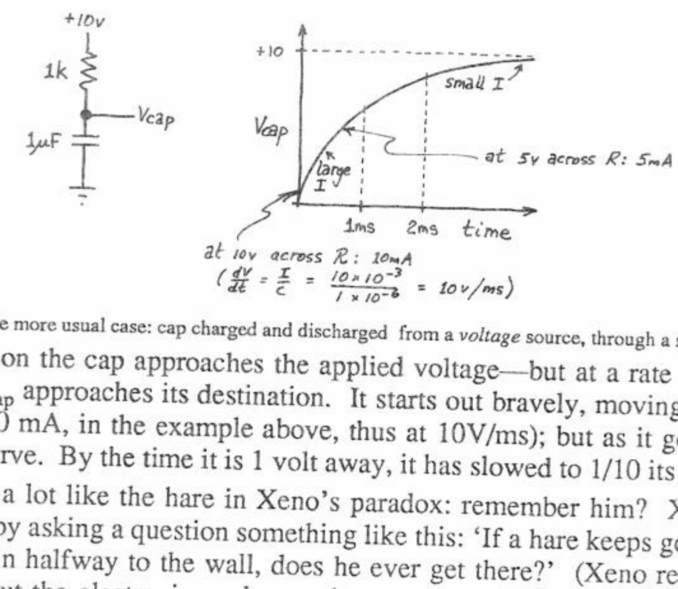

Figure N2.5: The more usual case: cap charged and discharged from a *voltage* source, through a series resistor

Here, the voltage on the cap approaches the applied voltage—but at a rate that diminishes toward zero as V_{cap} approaches its destination. It starts out bravely, moving fast toward its V_{in} (charging at 10 mA, in the example above, thus at 10V/ms); but as it gets nearer to its goal, it loses its nerve. By the time it is 1 volt away, it has slowed to 1/10 its starting rate.

(The cap behaves a lot like the hare in Xeno's paradox: remember him? Xeno teased his fellow-Athenians by asking a question something like this: 'If a hare keeps going halfway to the wall, then again halfway to the wall, does he ever get there?' (Xeno really had a hare chase a tortoise; but the electronic analog to the tortoise escapes us, so we'll simplify his problem.) Hares do bump their noses; capacitors don't: V_{cap} never does reach $V_{applied}$, in an *RC* circuit. But it will come as close as you want.)

Here's a fuller account:

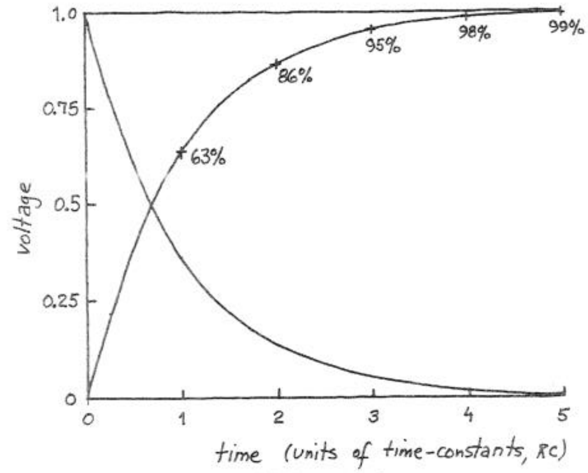

Figure N2.6: *RC* charge, discharge curves

Don't try to memorize these numbers, except two:

- in *one RC* (called "one time-constant") 63% of the way
- in *five RCs*, 99% of the way

If you need an exact solution to such a timing problem:

$$V_{cap} = V_{applied}\,(1 - e^{(-t/RC)})$$

In case you can't see at a glance what this equation is trying to tell you, look at

$$e^{(-t/RC)}$$

by itself:

- when t = *RC*, this expression is 1/e, or 0.37.
- when t = very large (» *RC*), this expression is tiny, and $V_{cap} \approx V_{applied}$

> *A tip to help you calculate time-constants:*
> MΩ and μF give time-constant in seconds
> kΩ and μF give time-constant in milliseconds

In the case above, *RC* is the product of 1k and 1μF: 1 ms.

Integrators and Differentiators

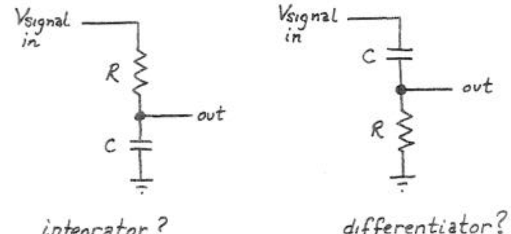

Figure N2.7: Can we exploit cap's $I = C\,dV/dt$ to make differentiator & integrator?

The very useful formula, $I = C dV/dt$ will let us figure out when these two circuits perform pretty well as differentiator and integrator, respectively.

Let's, first, consider this simpler circuit:

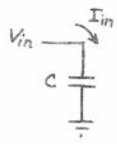

Figure N2.8: Useless "differentiator"?

The current that flows in the cap is proportional to dV_{in}/dt: the circuit differentiates the input signal. But the circuit is pretty evidently useless. It gives us no way to measure that current. If we could devise a way to measure the current, we would have a differentiator.

Here's our earlier proposal, again. Does it work?

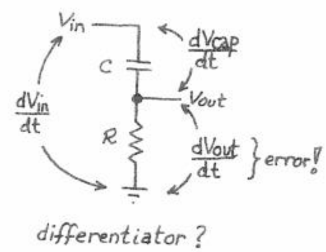

Figure N2.9: Differentiator? —again

Answer: Yes and No: Yes, to the extent that $V_{cap} = V_{in}$ (and thus $dV_{cap}/dt = dV_{in}/dt$), because the circuit responds to dv/dt across the cap, whereas what interests us is dV_{in}/dt—that is, relative to ground.

So, the circuit errs to the extent that the output moves away from ground; but of course it must move away from ground to give us an output. This differentiator is compromised. So is the *RC* integrator, it turns out. When we meet operational amplifiers, we will manage to make nearly-ideal integrators, and pretty good differentiators.

The text puts this point this way:

Text sec. 1.14

> for the differentiator:
> "…choose R and C small enough so that $dV/dt \ll dV_{in}/dt$…."

Text sec. 1.15

> for the integrator:
> "…[make sure that] $V_{out} \ll V_{in}$….$\omega RC \gg 1$."

We can put this simply—perhaps crudely: assume a sine wave input. Then,

> the *RC* differentiator (and integrator, too) works pretty well **if** it is **murdering** the signal (that is, attenuates it severely), so that V_{out} (and dV_{out}) is tiny: hardly moves away from ground.

It follows, along the way, that differentiator and integrator will impose a 90° phase shift on a sinusoidal input. This result, obvious here, should help you anticipate how *RC* circuits viewed as "filters" (below) will impose phase shifts.

Integrator

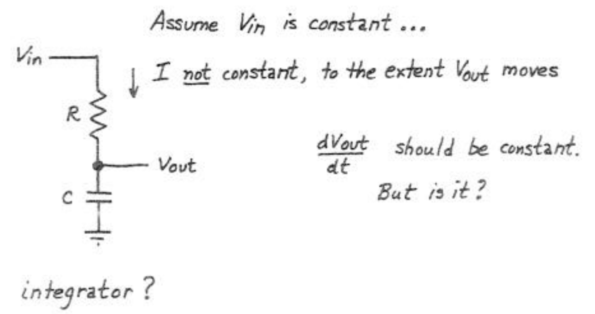

Figure N2.10: Integrator?—again

One can make a similar argument to explain the limitations of the *RC* integrator. To keep things simple, imagine that you apply a *step* input; ask what waveform out you would like to see out, and what, therefore, you would like the current to do.

RC Filters

These are the most important application of capacitors. These circuits are just voltage dividers, essentially like the resistive dividers you have met already. The resistive dividers treated DC and time-varying signals alike. To some of you that seemed obvious and inevitable (and maybe you felt insulted by the exercise at the end of Lab 1 that asked you to confirm that AC was like DC to the divider). It happens because resistors can't remember what happened even an instant ago. They're little existentialists, living in the present. (We're talking about ideal R's, of course.)

The impedance or *reactance* of a cap

A cap's impedance varies with frequency. ("Impedance" is the generalized form of what we called "resistance" for "resistors;" "reactance" is the term reserved for capacitors and inductors (the latter usually are coils of wire, often wound around an iron core)).

Compare Text sec. 1.12

It's obvious that a cap cannot conduct a DC current: just recall what the cap's insides look like: an insulator separating two plates. That takes care of the cap's "impedance" at DC: clearly it's infinite (or *huge*, anyway).

It is not obvious that a rapidly-varying voltage can pass "through" a capacitor, but that does happen. The Text explains this difficult notion at sec. 1.12, speaking of the *current* that passes through the cap. Here's a second attempt to explain how a *voltage* signal passes through a cap, in the high-pass configuration. If you're already happy with the result, skip this paragraph.

When we say the AC signal passes through, all we mean is that a wiggle on the left causes a wiggle of similar size on the right:

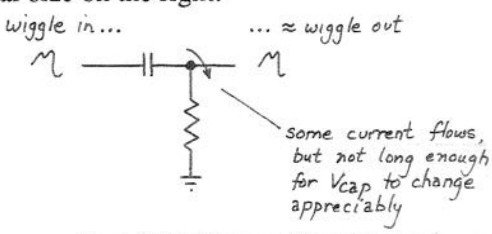

Figure N2.11: How a cap "passes" a signal

The wiggle makes it "across" the cap so long as there isn't time for the voltage on the cap to change much before the wiggle has ended—before the voltage heads in the other direction. In other words, quick wiggles pass; slow wiggles don't.

We can stop worrying about our intuition and state the expression for the cap's *reactance*:

$$X_C = -j/\omega C = -j/2\pi fC$$

And once we have an expression for the impedance of the cap—an expression that shows it varying continuously with frequency—we can see how capacitors will perform in voltage dividers.

RC Voltage dividers

Text sec. 1.18

You know how a resistive divider works on a sine. How would you expect a divider made of capacitors to treat a time-varying signal?

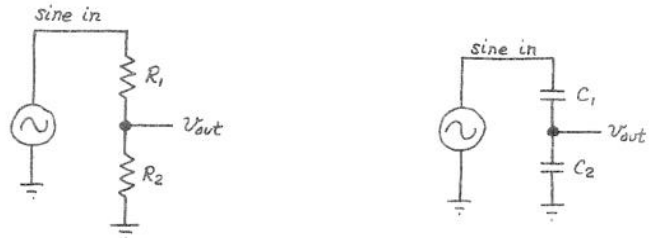

Figure N2.12: Two dividers that deliver 1/2 of V_{in}

If this case worries you, good: you're probably worrying about phase shifts. Turns out they cause no trouble here: output is in phase with input. (If you can handle the complex notation, write $X_C = -j / \omega C$, and you'll see the j's wash out.)

But what happens in the combined case, where the divider is made up of both R and C? This turns out to be an extremely important case.

Text sec. 1.19

This problem is harder, but still fits the voltage-divider model. Let's generalize that model a bit:

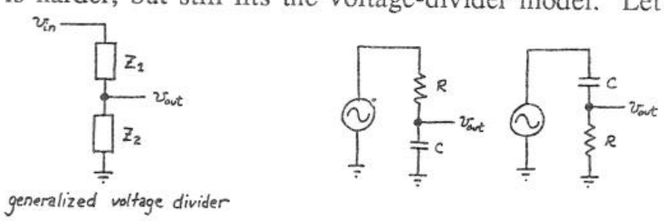

Figure N2.13: Generalized voltage divider; and voltage dividers made up of R paired with C

The behavior of these voltage dividers—which we call ***filters*** when we speak of them in frequency terms, because each favors either high or low frequencies— is easy to describe:

1. See what the filter does at the two frequency extremes. (This looking at extremes is a useful trick; you will soon use it again to find the filters' worst-case Z_{in} and Z_{out}.)

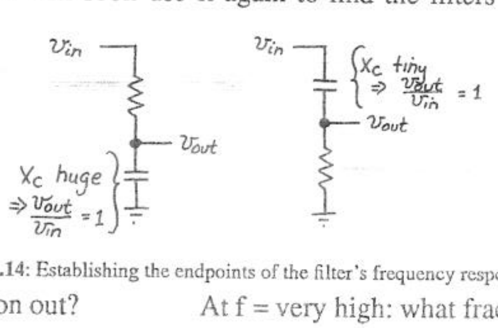

Figure N2.14: Establishing the endpoints of the filter's frequency response curve

At f = 0: what fraction out? At f = very high: what fraction out?

2. Determine where the output "turns the corner" (corner defined arbitrarily[1]) as the frequency where the output is 3dB less than the input (always called just "the 3 dB point"; "minus" understood).

Knowing the endpoints, which tell us whether the filter is *high-pass* or *low-pass*, and knowing the 3dB point, we can draw the full frequency-response curve:

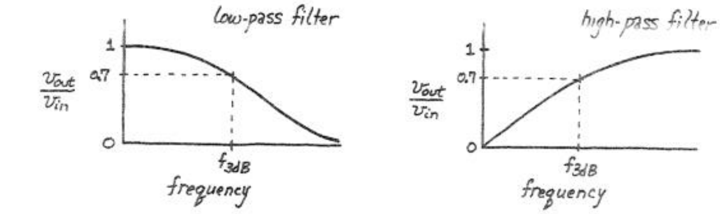

Figure N2.15: *RC* filter's frequency response curve

The "3dB point," the frequency where the filter "turns the corner" is

$$f_{3dB} = 1/(2\pi RC)$$

Beware the more elegant formulation that uses ω:

$$\omega_{3dB} = 1/RC.$$

That is tidy, but is very likely to give you answers off by about a factor of 6, since you will be measuring period and its inverse in the lab: frequency in hertz (or "cycles-per-second," as it used to be called), *not* in radians.

Two asides:

Caution!

Do not confuse these *frequency-domain* pictures with the earlier *RC* step-response picture, (which speaks in the *time-domain*).

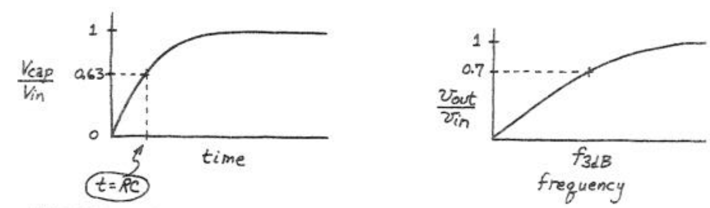

Figure N2.16: Deceptive similarity between shapes of time- and frequency- plots of *RC* circuits

Not only do the curves look vaguely similar. To make things worse, details here seem tailor-made to deceive you:

- *Step response*: in the *time RC* (time-constant), V_{cap} moves to about 0.6 of the applied step voltage (this is $1 - 1/e$).
- *Frequency domain*: at f_{3dB}, a frequency determined by *RC*, the filter's V_{out}/V_{in} is about 0.7 (this is $1/\sqrt{2}$)

Don't fall into this trap.

A note re Log Plots

You may wonder why the curves we have drawn, the curve in Fig. 1.59, and those you see on the scope screen (when you "sweep" the input frequency) don't look like the tidier curves shown in most books that treat frequency response, or like the curves in Chapter 5 of

1. Well, not quite arbitrarily: a signal reduced by 3dB delivers half its original power.

the Text. Our curves trickle off toward zero, in the low-pass configuration, whereas these other curves seem to fall smoothly and promptly to zero. This is an effect of the logarithmic compression of the axes on the usual graph. Our plots are linear; the usual plot ("Bode plot") is log-log:

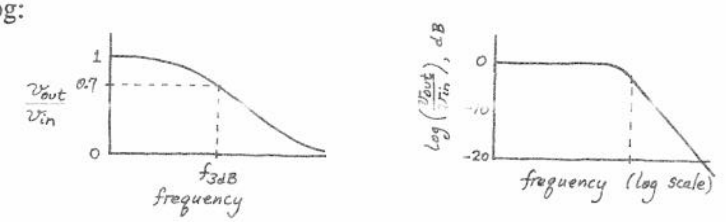

Figure N2.17: Linear versus log-log frequency-response plots contrasted

---

Input and output impedance of an *RC* Circuit

If filter A is to drive filter B—or anything else, in fact—it must conform to our *10X* rule of thumb, which we discussed last time, when we were considering only resistive circuits. The same reasons prevail, but here they are more urgent: if we don't follow this rule, not only will signals get attenuated; frequency response also is likely to be altered.

But to enforce our rule of thumb, we need to know Z_{in} and Z_{out} for the filters. At first glance, the problem looks nasty. What is Z_{out} for the low-pass filter, for example? A novice would give you an answer that's much more complicated than necessary. He might say,

$$Z_{out} = X_C \text{ parallel } R = -j/\omega C \cdot R \,/\, (-j/\omega C + R)$$

Yow! And then this expression doesn't really give an answer: it tells us that the answer is frequency-dependent.

We cheerfully sidestep this hard work, by considering only *worst case* values. We ask, 'How bad can things get?

We ask, 'How bad can Z_{in} get?' And that means, 'How *low* can it get?'

We ask, 'How bad can Z_{out} get?' And that means, 'How *high* can it get?'

This trick delivers a stunningly-easy answer: the answer is always just *R*! Here's the argument for a low-pass, for example:

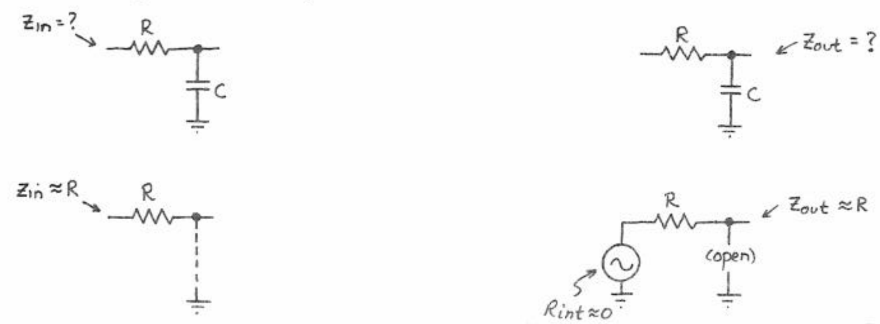

worst Zin: cap looks like a short: Zin = *R* (this happens at highest frequencies)

worst Zout: cap doesn't help at all; we look through to the source, and see only R: Zout = *R* (this happens at lowest frequencies)

Figure N2.18: Worst-case Z_{in} and Z_{out} for *RC* filter reduces to just *R*

Now you can string together *RC* circuits just as you could string together voltage dividers, without worrying about interaction among them.

Phase Shift

You already know roughly what to expect: the differentiator and integrator showed you phase shifts of 90°, and did this when they were severely attenuating a sine-wave. You need to *beware* the misconception that because a circuit has a cap in it, you should expect to see a 90° shift (or even just noticeable shift). That is *not so*. You need an intuitive sense of when phase shifting occurs, and of roughly its magnitude. You rarely will need to calculate the amount of shift.

Here is a start: a rough account of phase shift in *RC* circuits:

> If the amplitude *out* is close to amplitude *in*, you will see little or no phase shift. If the output is much attenuated, you will see considerable shift (90° is maximum)

And here are curves saying the same thing:

Text sec. 1.20, fig. 1.60, p. 38

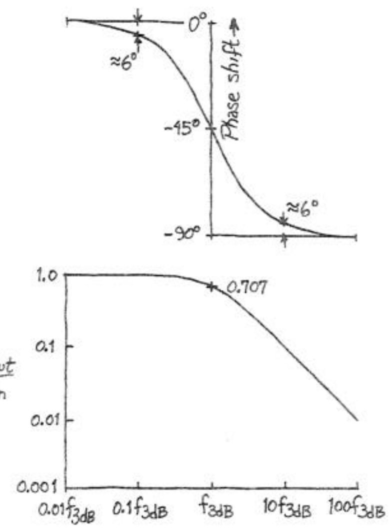

Figure N2.19: Attenuation and phase shift (log-log plot)

Why does this happen? Here's an attempt to rationalize the result:

- voltages in R and C are 90° out of phase, as you know.
- the sum of the voltages across R and C must equal, at every instant, V_{in}.
- as frequency changes, R and C share the total V_{in} differently, and this alters the phase of V_{out} relative to V_{in}:

Consider a low-pass, for example: if a lot of the total voltage, V_{in}, appears across the cap, then the phase of the input voltage (which appears across the RC series combination) will be close to the phase of the output voltage, which is taken across the cap alone. In other words, R plays a small part: V_{out} is about the same as V_{in}, in both amplitude and phase. Have we merely restated our earlier proposition? It almost seems so.

But let's try a drawing:

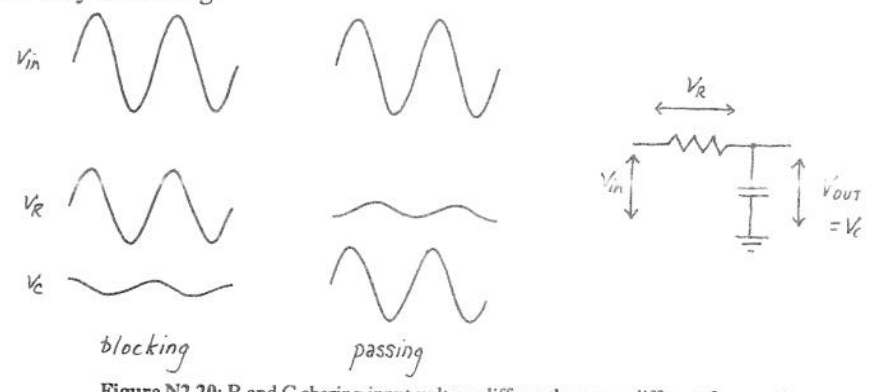

Figure N2.20: R and C sharing input voltage differently at two different frequencies

Now let's try another aid to an intuitive understanding of phase shift: phasors.

Phasor Diagrams

These diagrams let you compare phase and amplitude of input and output of circuits that shift phases (circuits including C's and L's). They make the performance graphic, and allow you to get approximate results by use of geometry rather than by explicit manipulation of complex quantities.

The diagram uses axes that represent resistor-like ("real") impedances on the horizontal axis, and capacitor or inductor-like impedances ("imaginary"—but don't let that strange name scare you; for our purposes it only means that voltages across such elements are 90° out of phase with voltages across the resistors). This plot is known by the extra-frightening name, "complex plane" (with nasty overtones, to the untrained ear, of 'too-complicated-for-you plane'!). But don't lose heart. It's all very easy to understand and use. Even better, *you don't need to understand phasors*, if you don't want to. We use them rarely in the course, and always could use direct manipulation of the complex quantities instead. Phasors are meant to make you feel better. If they don't, forget them.

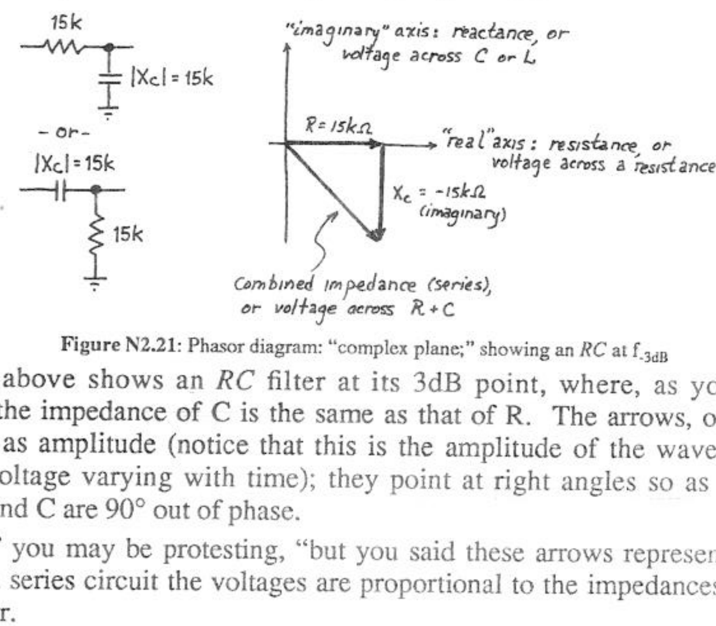

Figure N2.21: Phasor diagram: "complex plane;" showing an *RC* at f_{-3dB}

The diagram above shows an *RC* filter at its 3dB point, where, as you can see, the *magnitude* of the impedance of C is the same as that of R. The arrows, or vectors, show phase as well as amplitude (notice that this is the amplitude of the waveform: the peak value, not a voltage varying with time); they point at right angles so as to say that the voltages in R and C are 90° out of phase.

"Voltages?," you may be protesting, "but you said these arrows represent impedances." True. But in a series circuit the voltages are proportional to the impedances, so this use of the figure is fair.

The total impedance that R and C present to the signal source is *not* 2R, but is the vector sum: it's the length of the hypotenuse, $R\sqrt{2}$. And from this diagram we now can read two familiar truths about how an *RC* filter behaves at its 3dB point:

- the amplitude of the output is down 3dB: down to $1/\sqrt{2}$: the length of either the R or the C vector, relative to the hypotenuse.
- the output is shifted 45° relative to the input: R or C vectors form an angle of 45° with the hypotenuse, which represents the phase of the input voltage.

So far, we're only restating what you know. But to get some additional information out of the diagram, try doubling the input frequency several times in succession, and watch what happens:

each time, the length of the X_C vector is cut to half what it was.

First doubling of frequency:

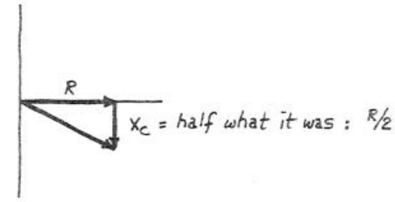

Figure N2.22: *RC* after a doubling of frequency, relative to the previous diagram

The first doubling also affects the length of the hypotenuse substantially, too, however; so the amplitude relative to input is not cut quite so much as 50% (6dB). You can see that the output is a good deal more attenuated, however, and also that phase shift has increased a good deal.

Second doubling of frequency

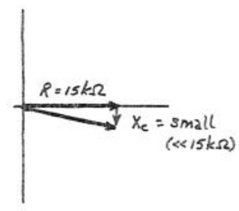

Figure N2.23: *RC* after a doubling of frequency, relative to the previous diagram

This time, the length of the hypotenuse is changed less, so the output shrinks nearly as the X_C vector shrinks: nearly 50%. Here, we are getting close to the *–6dB/octave* slope of the filter's rolloff curve. Meanwhile, the phase shift between output and input is increasing, too—approaching the limit of 90°.

We've been assuming a *low-pass*. If you switch assumptions, and ask what these diagrams show happening to the output of a *high-pass*, you find all the information is there for you to extract. No surprise, there; but perhaps satisfying to notice.

LC circuit on phasor diagram

Finally, let's look at an LC *trap* circuit on a phasor diagram.

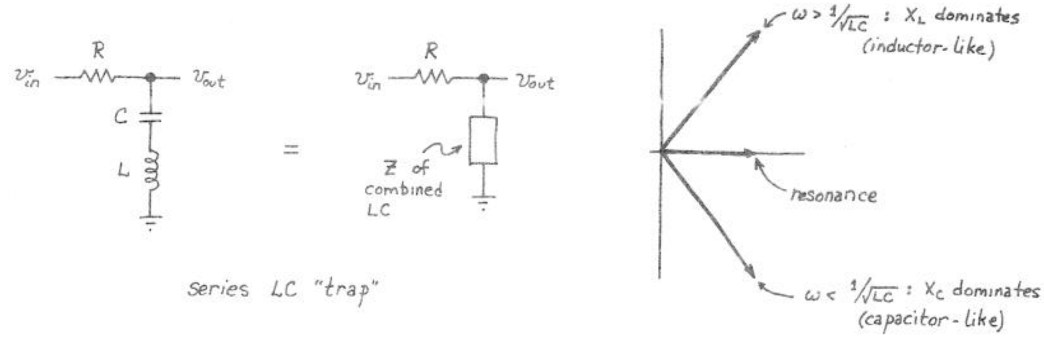

Figure N2.24: LC trap circuit, and its phasor diagram

This is less familiar, but pleasing because it reveals the curious fact—which you will see in Lab 3 when you watch a similar (parallel) LC circuit—that the LC combination looks sometimes like L, sometimes C, showing its characteristic phase shift—and at resonance, shows no phase shift at all. We'll talk about LC's next time; but for the moment, see if you can enjoy how concisely this phasor diagram describes this behavior of the circuit (actually a *trio* of diagrams appears here, representing what happens at *three* sample frequencies).

To check that these LC diagrams make sense, you may want to take a look at what the old voltage-divider equation tells you ought to happen:

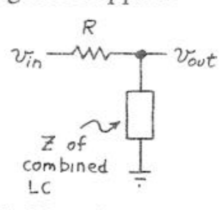

Figure N2.25: LC trap: just another voltage divider

Here's the expression for the output voltage as a fraction of input:

$$V_{out}/V_{in} = Z_{combination} / (Z_{combination} + R)$$

But

$$Z_{combination} = -j/\omega C + j\omega L.$$

And at some frequency—where the magnitudes of the expressions on the right side of that last equation are equal—the sum is zero, because of the opposite signs. Away from this magic frequency (the "resonant frequency"), either cap or inductor dominates. Can you see all this on the phasor diagram?

Better Filters

Having looked hard at *RC* filters, maybe we should remind of the point that the last exercise in Lab 2 means to make: *RC*'s make extremely useful filters, but if you need a better filter, you can make one, either with an LC combination, as in that circuit, or with operational amplifiers cleverly mimicking such an LC circuit (this topic is treated in Chapter 5; we will not build such a circuit), or with a clever circuit called a 'switched-capacitor' filter, a circuit that you will get a chance to try, in Lab 11, and again in Lab 21. Here is a sketch (based on a scope photo) comparing the output of an ordinary *RC* low-pass against a 5-pole Butterworth low-pass, like the one you will build at the end of Lab 2.

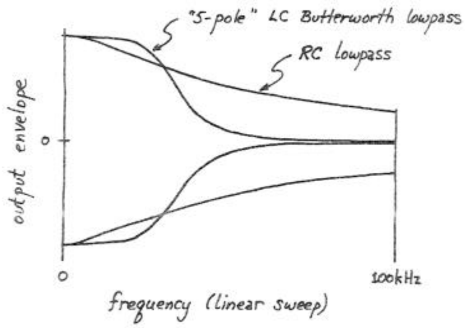

Figure N2.26: Simple *RC* low-pass contrasted with 5-pole Butterworth low-pass

("5-pole" is a fancy way to say that it something like, 'It rolls off the way 5 simple *RC*'s in series would roll off'—it's five times as steep as the plain *RC*. But this nice filter works a whole lot better than an actual string of 5 *RC*'s.)

Not only is the roll-off of the Butterworth much more abrupt than the simple *RC*'s, but also the "passband" looks much flatter: the fancy filter does a better job of passing. The poor old *RC* looks sickly next to the Butterworth, doesn't it?

Nevertheless, we will use *RC*'s nearly always. Nearly always, they are good enough. It would not be in the spirit of this course to pine after a more beautiful transfer function. We want circuits that work, and in most applications the plain old *RC* passes that test.

Chapter 1: Worked Examples: RC Circuits

Two worked exercises:

1. filter to keep "signal" and reject "noise"
2. bandpass filter

1. Filter to keep "signal" and reject "noise'

Problem:

Filter to remove fuzz:

Suppose you are faced with a signal that looks like this: a signal of moderate frequency, polluted with some fuzz

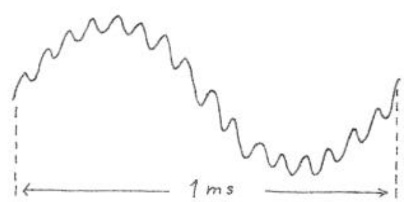

Figure X2.1: Signal With Fuzz Added

1. Draw a *skeleton* circuit (no parts values, yet) that will keep most of the good signal, clearing away the fuzz.
2. Now choose some values:
    a. If the *load* has value 100k, choose R for your circuit.
    b. Choose f_{3dB}, explaining your choice briefly.
    c. Choose C to achieve the f_{3dB} that you chose.
    d. By about how much does your filter attenuate the noise "fuzz"?

    What is the circuit's input impedance—

    a. at very low frequencies?
    b. at very high frequencies?
    c. at f_{3dB}?
3. What happens to the circuit output if the load has resistance *10k* rather than 100k?

A Solution:

1. Skeleton Circuit

You need to decide whether you want a low-pass or high-pass, since the signal and noise are distinguishable by their frequencies (and are far enough apart so that you can hope to get one without the other, using the simple filters we have just learned about). Since we have called the lower frequency "good" or "signal," we need a *low-pass*:

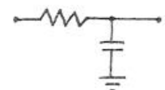

Figure X2.2: Skeleton: just a low-pass filter

2. Choose R, given the load

This dependence of R upon load follows from the observation that R of an RC filter defines the *worst-case* input and output impedance of the filter (see Class 2 notes). We want that output impedance low relative to the load's impedance; our rule of thumb says that 'low' means low by a factor of 10. So, we want $R \leq R_{load}/10$. In this case, that means R should be ≤ 10k. Let's use 10k.

3. Choose f_{3dB}

This is the only part of the problem that is not entirely mechanical. We know we want to pass the low and attenuate the high, but does that mean put f_{3dB} halfway between good and bad? Does it mean put it close to good? …Close to bad? Should both good and bad be on a steeply-falling slope of the filter's response curve?

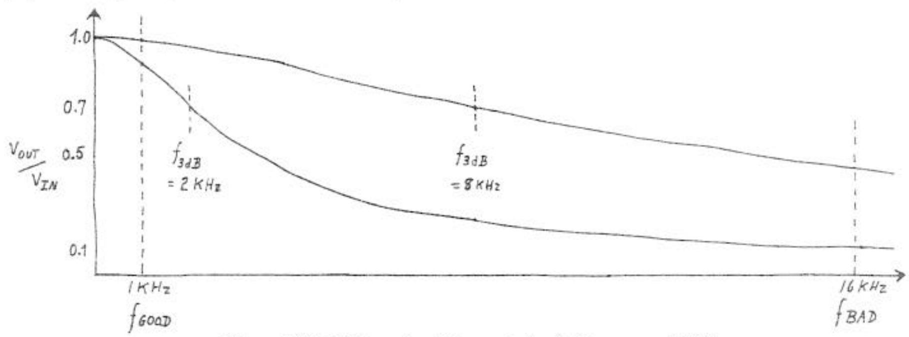

Figure X2.3: Where should we put f_{3dB} ? Some possibilities

Assuming that our goal is to achieve a large ratio of good to bad signal, then we should not put f_{3dB} close to the noise: if we did we would not do a good job of attenuating the bad. Halfway between is only a little better. Close to signal is the best idea: we will then attenuate the bad as much as possible while keeping the good, almost untouched.

An alert person might notice that the greatest relative preference for good over bad comes when both are on the steepest part of the curve showing frequency response: in other words, put f_{3dB} so low that *both* good and bad are attenuated. This is a clever answer—but wrong, in most settings.

The trouble with that answer is that it assumes that the signal is a single frequency. Ordinarily the signal includes a range of frequencies, and it would be very bad to choose f_{3dB} somewhere *within* that range: the filter would distort signals.

So, let's put f_{3dB} at $2 \cdot f_{signal}$: around 2kHz. This gives us 89% of the original signal amplitude (as you can confirm if you like with a phasor diagram or direct calculation). At the same time we should be able to attenuate the 16kHz noise a good deal (we'll see in a moment *how* much).

4...Choose C to achieve the f_{3dB} that you want

This is entirely mechanical: use the formula for the 3dB point:

$f_{3dB} = 1/(2\pi RC) \Longrightarrow C = 1/(2\pi f_{3dB}R) \approx 1/(6 \cdot 2 \cdot 10^3 \cdot 10 \cdot 10^3) = 1/(120 \cdot 10^6) \approx 0.008 \cdot 10^{-6}F$

We might as well use a 0.01 μF cap ("cap" ≡ capacitor). It will put our f_{3dB} about 25% low—1.6Khz; but our choice was a rough estimate anyway.

5. By about how much does your filter attenuate the noise ("fuzz")?

The quick way to get this answer is to count octaves or decades between f_{3dB} and the noise. f_{3dB} is 2 kHz; the fuzz is around 16kHz: $8 \cdot f_{3dB}$.

Count octaves: we could also say that the frequency is doubled three times ($= 2^3$) between f_{3dB} and the noise frequency. *Roughly,* that means that the fuzz amplitude is cut in half the same number of times: down to $1/(2^3)$: 1/8.

This is only an approximate answer because a) near f_{3dB} the curve has not yet reached its terminal steepness of –6dB/octave; b) on the other hand, even *at* f_{3dB}, some attenuation occurs. But let's take our rough answer. (It happens that our rough answer is better than it deserves to be: we called V_{out}/V_{in} 0.125; the more exact answer is 0.124.)

6. What happens to the circuit output if the load has resistance 10k rather than 100k?

Here's a picture of such loading.

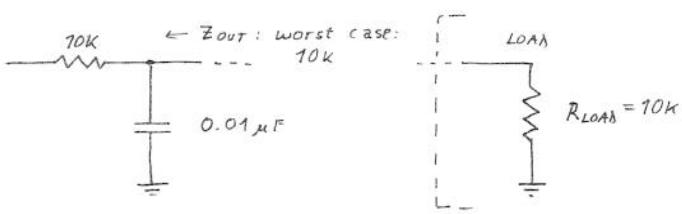

Figure X2.4: Overloaded filter

If you have gotten used to Thevenin models, then you can see how to make this circuit look more familiar:

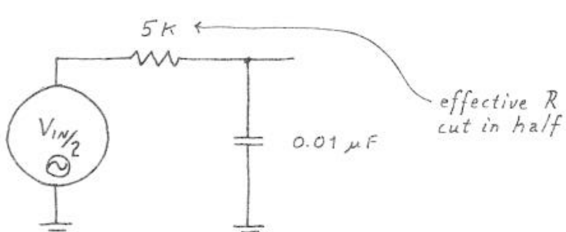

Figure X2.5: Loaded circuit, redrawn

The amplitude is down; but, worse, *f3dB* has changed: it has doubled. You will find a plot showing this effect at the start of the class notes for next time: Class 3.

7. What is the circuit's input impedance—

1. at very low frequencies?
   Answer: *very large*: the cap shows a high impedance; the signal source sees only the load—which is assumed very high impedance (high enough so we can neglect it as we think about the filter's performance)

2. at very high frequencies?
   Answer: *R*: The cap impedance falls toward zero—but *R* puts a lower limit on the input impedance.

3. at f_{3dB}?

This is easy if you are willing to use phasors, a nuisance to calculate, otherwise. If you recall that the magnitude of $X_C = R$ at f_{3dB}, and if you accept the notion that the voltages across R and C are 90° out of phase, so that they can be drawn at right angles to each other on a phasor diagram, then you get the phasor diagram you saw in the class notes (fig. N2.21):

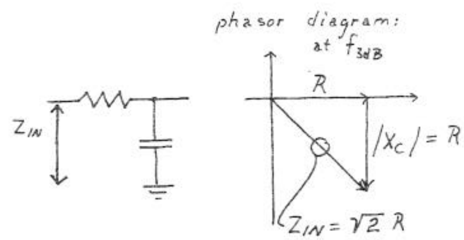

Figure X2.6: Phasor diagram showing an RC filter at its 3dB point

—then you can use a geometric argument to show that the hypotenuse—proportional to Z_{in}—is $R\sqrt{2}$.

2. Bandpass

Text exercise AE-6 (Chapter 1)

Problem:

> ***Bandpass filter***
>
> Design a bandpass filter to pass signals between about 1.6 kHz and 8kHz (you may use these as 3dB frequencies)..
>
> Assume that the *next* stage, which your bandpass filter is to drive, has an input impedance of 1 M ohms.

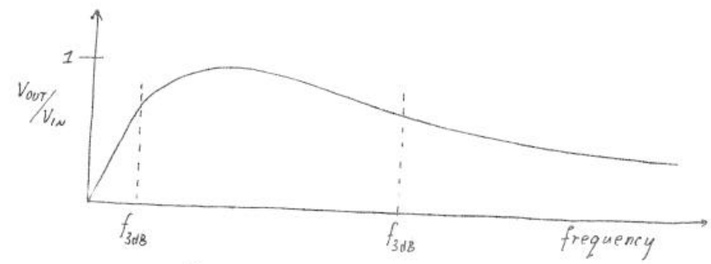

Figure X2.7: Bandpass frequency response

Once you recognize that to get this frequency response from the filter you need to put *high-pass* and *low-pass* in series, the task is mechanical. You can put the two filters in either order. Now we need to choose R values, because these will determine worst-case impedances for the two filter stages. The later filter must show Z_{out} low relative to the load, which is 1MΩ; the earlier filter must show Z_{out} low relative to Z_{in} of the second filter stage.

So, let the second-stage R = 100k; the first-stage R = 10k:

Choosing C's

Here the only hard part is to get the filter's right: it's hard to say to oneself, 'The high-pass filter has the lower f_{3dB};' but that is correct. Here are the calculations: notice that we try to keep things in *engineering notation*—writing '$10{\bullet}10^3$' rather than '10^4.' This form looks clumsy, but rewards you by delivering answers in standard units. It also helps you scan for nonsense in your formulation of the problem: it is easier to see that '$10{\bullet}10^3$' is a good translation for '10k' than it is to see that, say, '10^5' is *not* a good translation.

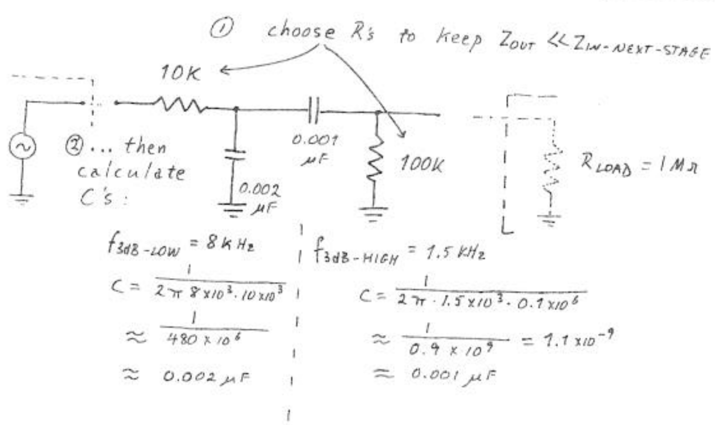

Figure X2.8: Calculating C values: a plug for engineering notation

A Note on Reading Capacitor Values

Why you may need this note

Most students learn pretty fast to read resistor values. They tend to have more trouble finding, say, a 100 pF capacitor.

That's not their fault. They have trouble, as you will agree when you have finished reading this note, because the cap manufacturers don't want them to be able to read cap values. ("Cap" is shorthand for "capacitor," as you probably know.) The cap markings have been designed by an international committee to be nearly unintelligible. With a few hints, however, you can learn to read cap markings, despite the manufacturers' efforts. Here are our hints:

Big Caps: electrolytics

These are easy to read, because there is room to write the value on the cap, including units. You need only have the common sense to assume that

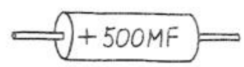

Figure CAP.1: A big cap is labeled intelligibly

means 500 micro farads, not what it would mean if you took the capital M seriously.

All of these big caps are *polarized*, incidentally. That means the capacitor's innards are not symmetrical, and that you may destroy the cap if you apply the wrong polarity to the terminals: the terminal marked + must be at least as positive as the other terminal. (Sometimes, violating this rule will generate gas that makes the cap blow up; more often, you will find the cap internally shorted, after a while. Often, you could get away with violating this rule, at low voltages. But don't try.)

Smaller Caps

As the caps get smaller, the difficulty in reading their markings gets steadily worse.

Tantalum

These are the silver colored cylinders. They are polarized: a + mark and a metal nipple mark the positive end. Their markings may say something like

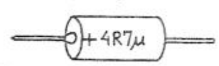

Figure CAP.2: Tantalum cap

That means pretty much what it says, if you know that the "R" marks the decimal place: it's a 4.7 μF cap.

The same cap is also marked,

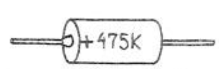

Figure CAP.3: Tantalum cap: second marking scheme on same part

Here you meet your first challenge, but also the first appearance of an orderly scheme for labeling caps, a scheme that would be helpful if it were used more widely.

The challenge is to resist the plausible assumption that "k" means "kilo." It does not; it is not a unit marking, but a *tolerance* notation (it means ± 10%). (Wasn't that nasty of the labelers to choose "K?" Guess what's another favorite letter for tolerance. That's right: M. Pretty mean!)

The orderly labeling here mimics the resistor codes: 475 means

$47 \times$ ten to the fifth.

What units?

10^5 what? 10^5 of something small. You will meet this question repeatedly, and you must resolve it by relying on a few observations:

1. The only units commonly used in this country are
   - microfarads: 10^{-6} Farad
   - picofarads: 10^{-12} Farad

   (You should, therefore, avoid using "mF" and "nF," yourself.)
   A Farad is a huge unit. The biggest cap you will use in this course is 500 μF. That cap is physically large. (We do keep a 1F cap around, but only for our freak show.) So, if you find a little cap labeled "680," you know it's 680 **pF**.
2. A picofarad is a tiny unit. You will not see a cap as small as 1 pF in this course. So, if you find a cap claiming that it is a fraction of some unstated unit—say,".01"—the unit is μF's: ".01" means 0.01 μF.
3. *Beware* the wrong assumption that a *picofarad* is only a bit smaller than a microfarad. A *pF* is *not* 10^{-9} F (10^{-3} μF); instead, it is 10^{-12}F: a *million* times smaller than a microfarad.

So, we conclude, this cap labeled "475" must be 4.7×10^6 ***picofarads***. That, you will recognize, is a roundabout way to say

$$4.7 \times 10^{-6}\ \text{F}$$

We knew that was the answer, before we started this last decoding effort. This way of labeling is indeed roundabout, but at least it is unambiguous. It would be nice to see it used more widely. You will see another example of this *exponential* labeling in the case of the CK05 ceramics, below.

Mylar

These are yellow cylinders, pretty clearly marked. .01M is just 0.01 μF, of course; and .1MFD is *not* a tenth of a megafarad. These caps are not polarized; the black band marks the outer end of the foil winding. We don't worry about that fine point. Orient them at random in your circuits.

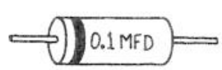

Figure CAP.4: Mylar capacitor

Because they are long *coils* of metal foil (separated by a thing dielectric—the "mylar" that gives them their name), mylar caps must betray their coil-like construction at very high frequencies: that is, they begin to fail as capacitors, behaving instead like inductors, blocking the very high frequencies they ought to pass. Ceramics (below) do better in this respect, though they are poor in other characteristics.

Ceramic

These are little orange pancakes. Because of this shape (in contrast to the *coil* format hidden within the tubular shape of mylars) they act like capacitors even at high frequencies. The trick, in reading these, is to reject the markings that can't be units:

Disc

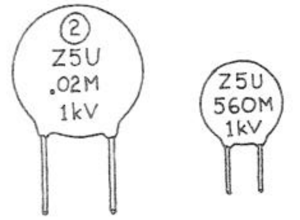

Figure CAP.5: Disc capacitor markings

Z5U: Not a unit marking: cap type

.02M, 560M That's it: the M is a tolerance marking, as you know (± 20%); not a unit

Common sense tells you units:

".02?" microfarads. "560?" picofarads.

1kV Not a unit marking. Instead, this means—as you would guess—that the cap can stand 1000 volts.

CK05

These are little boxes, with their leads 0.2" apart. They are handy, therefore, for insertion into a printed circuit.

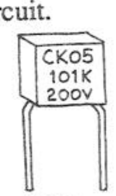

Figure CAP.6: CK05 capacitor markings

101k: This is the neat resistor-like marking. This one is 100 pF.

Tolerance Codes

Just to be thorough—and because this information is hard to come by—let's list all the tolerance codes. These apply to both capacitors and resistors; the tight tolerances are relevant only to resistors; the strangely-asymmetric tolerance is used only for capacitors.

Tolerance Code	Meaning
Z	+80%, –20% (for big filter capacitors, where you are assumed to have asymmetric worries: too small a cap allows excessive "ripple;" more on this in Lab 3 and Notes 3)
M	± 20%
K	± 10%
J	± 5%
G	± 2%
F	± 1%
D	± 0.5%
C	± 0.25%
B	± 0.1%
A	± 0.05%
Z	± 0.025 (precision resistors; context will show the asymmetric cap tolerance "Z" makes no sense here)
N	± 0.02%

Figure CAP.7: Tolerance codes