

S2014, BME 101L
Applied Circuits Lab 3b: Sampling and Aliasing

Kevin Karplus

April 7, 2014

1 Design Goal

This lab is more of a demo than a design lab. The main point is to look at sampling frequency and how it affects the representation of signals stored on computers.

You will have to design a couple of RC filters in order to do the demo: a high-pass filter to block the DC component of the signal generator and shift the voltage up to the mid-voltage of the analog-to-digital converter, and a low-pass filter to look at the effects of low-pass filtering on signals.

2 Background

2.1 Sampling

When we record signals on the computer, we cannot store infinitely many data points to represent continuous functions of time. Instead we store data at fixed time intervals (the *sampling period*). Because it is often easier to deal with frequencies rather than times, we refer to the inverse of the sampling period as the *sampling frequency*, f_s .

If our signals are low frequency compared to our sampling frequency, the computer representation of them is quite good, and little, if any, information is lost in storing them at discrete times. There is a theorem that signals that contain no energy above $0.5f_s$ can be accurately reconstructed from the sampled data—the frequency $0.5f_s$ is referred to as the *Nyquist frequency* or the *Nyquist limit*. (The math is beyond the scope of this course, but should be covered in the “Signals and Systems” course.)

In this lab we’ll look at what happens to the recordings of signals as they get close to or exceed the Nyquist limit.

One interesting effect is *aliasing* in which a frequency higher than the Nyquist frequency appears indistinguishable from a frequency lower than the Nyquist frequency. This effect is commonly observed with stroboscopes (which give you visual samples at fixed time intervals) and Moiré patterns, which sample images at a fixed spatial frequency.

To avoid artifacts in data recording, it is usually necessary to filter the input signal to remove all components that have frequencies above the Nyquist frequency. You may have already experienced anti-aliasing filters on your computer—characters displayed on the screen are often deliberately blurred (low-pass filtered) to remove the “jaggies” that are discretization artifacts (in space rather than in time). We’ll look at the effect of a very crude RC low-pass filter in this lab. More sophisticated filter design is left for the courses “Signals and Systems” and “Digital Signal Processing”.

2.2 RC filters

An RC filter is simply a voltage divider which uses a resistor for one of the two impedances and a capacitor for the other (see Figure 1). We looked at one RC filter very informally in the microphone lab, where you looked at DC-blocking capacitors with the oscilloscope input serving as the other impedance. In this lab you will need to design some filters with specific properties, so you need a more formal model of how they work.




Figure 1: Voltage dividers that act as high-pass and low-pass filters. There is no need to memorize which is which, as it is easy to derive from the behavior at ∞ Hz and 0Hz: at very high frequency, the capacitors are almost short circuits, and at very low frequency they are almost open circuits.

The basic idea is to combine two formulas: the voltage divider formula, which we use in almost every design, and the formula for the impedance of a capacitor. Let’s derive the formula for impedance of capacitors from basic physics concepts.

2.3 Impedance of capacitor

First, impedance is the ratio of voltage to current (for DC circuits impedance is just resistance), so we’ll want to look at voltages and currents for sinusoids.

Second, current is movement of charge. We know that the charge on a capacitor is the product of the capacitance and the voltage across the capacitor: $Q = CV$, and any current into one plate of the capacitor is going to increase the charge there. That means that the current is the derivative of the charge on the capacitor: $i(t) = \frac{dQ(t)}{dt}$, just as velocity is the derivative of position.

If we are dealing with a constant capacitance (a good approximate model for most of the components we’ll deal with—except capacitive or electret microphones), then the current is proportional to the derivative of voltage: $i(t) = C \frac{dv(t)}{dt}$.

Now let’s look at sinusoids. Since trigonometry gets a bit confusing, I like to work always with complex numbers and the unit circle instead, using Euler’s formula: $e^{j\theta} = \cos(\theta) + j \sin(\theta)$, where $j = \sqrt{-1}$ (electrical engineers use i for current and so need a non-standard symbol for $\sqrt{-1}$). We can represent a sinusoidal voltage $v(t) = \cos(\omega t) = \Re(e^{j\omega t})$, where $\Re(z)$ means the real part of z . The parameter ω is the *angular frequency* in radians per second, but we usually prefer to talk about the frequency f in cycles per second (Hertz), with $\omega = 2\pi f$. It turns out to be very convenient not to take the “real” part of the exponential formula, but to leave the voltage as a complex number: $v(t) = e^{j\omega t}$.

You can think of any complex number as a amplitude A (how far from 0) and an angle from the real axis ϕ : $Ae^{j\phi}$. We can equate these points with different sinusoids (all with the same angular frequency): $A \cos(\omega t + \phi) = \Re(Ae^{j(\omega t + \phi)}) = \Re(Ae^{j\phi}e^{j\omega t})$. The number $Ae^{j\phi}$ is referred to as a *phasor* when it is representing a sinusoid like this. (Read the Wikipedia article <http://en.wikipedia.org/wiki/Phasor>.) The real number A is the *amplitude*, and the angle ϕ is the *phase* of the sinusoid.

Exponentials are very nice for differentiation: if $v(t) = e^{j\omega t}$, then $dv(t)/dt = j\omega e^{j\omega t}$, and

$$\begin{aligned}
 i(t) &= C \frac{dv(t)}{dt} \\
 &= C j \omega e^{j\omega t} \\
 &= C j \omega v(t)
 \end{aligned}$$

That means that the impedance for a capacitor is

$$Z_C(\omega) = \frac{v(t)}{i(t)} = \frac{1}{j\omega C} .$$

Note that at DC (0Hz), $\omega = 0$ and $Z_C(0) = \infty$, so the capacitor is an open circuit. At very high frequency $Z_C(\infty) = 0$, and the capacitor is effectively a short circuit.

Notice that the impedance of a capacitor is not a real number, but an imaginary one. What does that mean? Basically, it means that the phase of the sinusoid is changed between the current and the voltage. Because the current is proportional to the derivative of the voltage across a capacitor, it is 0 when the voltage is at its maximum and minimum (the derivative of a cosine is the negative of the sine). The 90° shift in the phase of the signal corresponds to the 90° angle between the real and imaginary axes of the complex plane.

This class will look mainly at filters for their effect on the amplitude of sinusoids. We'll leave the effect on phase to more advanced electronics or signal processing classes.

2.4 RC voltage divider

Now that we have formula for impedance of a capacitor, we can get formulas for the behavior of an RC filter by using our voltage divider formula.

For the high-pass filter in Figure 1 we have

$$\begin{aligned} \frac{V_{out}}{V_{in}} &= \frac{R_1}{R_1 + Z_{C_1}(\omega)} \\ &= \frac{R_1}{R_1 + 1/(j\omega C_1)} \\ &= \frac{j\omega R_1 C_1}{1 + j\omega R_1 C_1} \end{aligned}$$

For the low-pass filter in Figure 1 we have

$$\begin{aligned} \frac{V_{out}}{V_{in}} &= \frac{Z_{C_2}(\omega)}{R_2 + Z_{C_2}(\omega)} \\ &= \frac{1/(j\omega C_2)}{R_2 + 1/(j\omega C_2)} \\ &= \frac{1}{1 + j\omega R_2 C_2} \end{aligned}$$

Note that both filters depend not on R and C separately, but only on the product RC , which is measured in seconds (an ohm-Farad is a second). You can see the high-pass and low-pass nature of the filters by looking at what the voltage divider ratio would be at $\omega = 0$ and $\omega = \infty$. Even better is to plot the magnitude of the voltage-divider ratio on a log-log scale as in Figure 2.

A straight-line approximation to the log-log plot is referred to as a *Bode plot*, after Hendrik Wade Bode, who developed quick ways to plot amplitude and phase response of filters without extensive calculation. The Wikipedia article http://en.wikipedia.org/wiki/Bode_plot will allow you to explore the concept further, but it goes well beyond what you need for this course.

Note that for frequencies well below the corner frequency $f_{corner} = \frac{1}{2\pi RC}$, the low-pass filter has a voltage divider ratio of essentially constant 1, and the high-pass filter has a magnitude of f/f_{corner} .

For frequencies well above the corner frequency, the high-pass filter's voltage divider ratio is essentially 1, and the low-pass filter has a magnitude of f_{corner}/f .

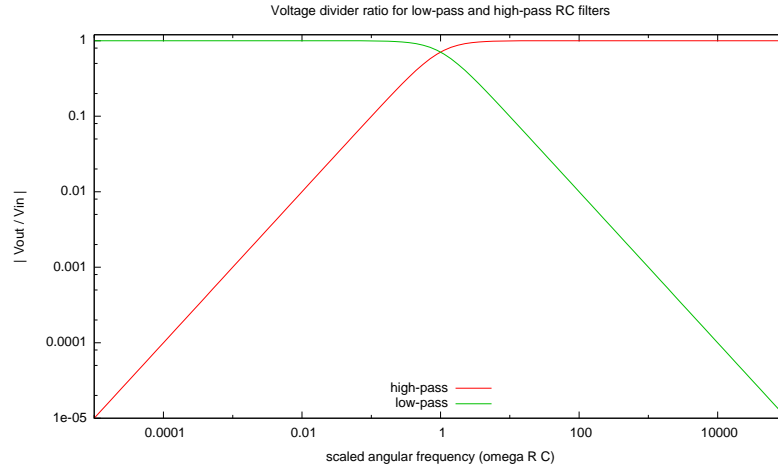


Figure 2: Plot of the effect on amplitude of sinusoids of simple RC filters, such as those in Figure 1. The *corner frequency* where $\omega = \frac{1}{RC}$ is an important point. (This is often expressed as $f = \frac{1}{2\pi RC}$.) The log-log plot allows the behavior of the filter to be approximated by straight lines, which is handy for quick designs.

2.5 Electrolytic Capacitors

The low frequencies in this lab require large RC time constants, which in turn requires large capacitance.

To get a large capacitance, you need a large plate area, a very thin insulator, and a high dielectric constant for the insulator. One standard way to achieve this is with an *electrolytic capacitor*, which has metal foil for one electrode (usually aluminum, but sometimes tantalum), coated with a thin oxide layer as the insulator. The other electrode is a wet electrolyte connected to another sheet of aluminum. The whole sandwich is coiled up and stuffed into a cylindrical can.

The most common electrolytic capacitors have the oxide coating on only one of the electrodes, making them *polarized*. The idea is that if tiny holes occur in the oxide layer, any current flow through them will cause more oxidation of the aluminum, repairing the hole. If the current flowed the *wrong direction*, then you would get reduction of the oxide layer and an increase in the size of the holes, eventually leading to a large current through the capacitor and heating of the electrolyte. This can happen quite quickly—fast enough to turn the wet electrolyte to steam and explode the capacitor.

Please read the Wikipedia article http://en.wikipedia.org/wiki/Electrolytic_capacitor.

2.6 Level shifting

A lot of the high-pass filters you will need to design are primarily for changing the DC voltage level that a signal has. For example, in the microphone lab, you were trying to center the output at 0V, rather than at 1v. For amplifier labs coming up, you’ll sometimes want to center the AC voltage inputs at the mid-point between the two power rails—you’ll sometimes hear me refer to this point as a “virtual ground”, because amplifier design traditionally used a symmetric power supply with a positive and a negative voltage source around ground.

Figure 3 shows two simple circuits for shifting a signal with an unknown DC bias to be centered at 1.65v. The first one uses a power supply set to 1.65v, which is not generally available (though we’ll see how to fake one in a later lab). The second one takes advantage of the notion of *Thévenin equivalence*.

Thévenin’s theorem says that that any circuit made up of passive components (resistors, capacitors, and inductors) and sources (voltage sources and current sources) that has only 2 ports (connections to




Figure 3: These two circuits show how to shift the DC offset of a signal to be centered at 1.65v. In both circuits, the capacitor blocks any DC current. With a constant V_{in} , the capacitor maintains a constant charge, and the output is a constant 1.65v, as long as no current is taken through the V_{out} port. Both circuits act as high-pass filters, with corner frequency $1/(2\pi RC)$. Note that the resistances in the second circuit are effectively in parallel, so that we need to make them twice as big to get the same RC time constant.




Figure 4: The Thévenin equivalent of a voltage source and voltage divider, with the two ports V_{out} and ground, is a voltage source whose voltage is set by the voltage-divider formula and a resistance that is the two resistors in parallel.

the outside circuitry) can be replaced by a single voltage source in series with a single impedance. There is even an easy way to determine what the voltage source and impedance are: The voltage source has whatever voltage you would see across the ports if they are not connected to any external circuitry, and the impedance is what you would get if all the voltage sources were replaced by short circuits and all the current sources by open circuits. (Note: Thévenin's theorem doesn't say anything about how much power is converted to heat internally in the circuit—only about how the circuit behaves electrically as viewed from the two ports.)

Read http://en.wikipedia.org/wiki/Thevenins_theorem for more explanation of Thévenin equivalence.

Figure 4 shows how Thévenin's theorem applies to a voltage divider, regarded as a 2-port device (V_{out} and ground). Note that if the two resistors have identical values, then the equivalent circuit has half the voltage of the source in the original circuit, and a series resistance half the size of the resistors in the original circuit. The level shifters in Figure 3 consist of one of the circuits of Figure 4 with a capacitor connecting V_{in} to V_{out} .

3 Pre-lab assignment

You will use a signal generator in the lab to generate waveforms of known frequency and look at them using the PteroDAQ code. (You'll need a recent version, 2014 April 6 or newer, to have the downsampling feature implemented.)

The first thing to do is to get a sine wave to display nicely on analog input PTB0. The function generator produces sine waves that are centered at 0V, but the KL25Z analog-to-digital converter wants signals between 0V and 3.3V.

- Design a high-pass filter that blocks the DC component of the function generator and has a 1.65v DC output when there is no input.

We're going to be looking at very low frequency signals, so the corner frequency of the high-pass filter should be around 30mHz (0.03Hz)

- Design another low-pass filter that has a corner frequency around 4Hz.

4 Parts, tools, and equipment needed

Parts for this lab from kit:

- breadboard
- resistors
- capacitors (particularly electrolytic)
- KL25Z
- USB cable

Tools for this lab:

- Wire strippers

Equipment in lab:

- Signal generator (0.1Hz to 100Hz desired).
- Oscilloscope for checking voltage range before connecting to KL25Z.

5 Procedures

Wire up your high-pass filter, using the KL25Z board's 0v and 3.3V lines for power. Set the function generator to have about a 1.6V peak-to-peak sine wave at 10Hz and look at the output of your high-pass filter on the oscilloscope (DC-coupled). Adjust the function generator output voltage (and, possibly, the design of your high-pass filter) until the scope shows that your output remains between 0 and 3.3 volts, but gets below 0.5V and above 2.8V. Note that the very large RC time constant for your high-pass filter means that there is a slow power-on transition when you first apply power to the filter, as the electrolytic capacitor slowly charges up. Wait for that to settle before you start trying to adjust the design.

Check that the signal remains in the desired range (low peak between 0V and 0.5V, high peak between 2.8V and 3.3V) at 0.1Hz, 1Hz, 10Hz, and 100Hz.

Connect the output of the high-pass filter to an analog pin on the KL25Z board. In PteroDAQ, set the timed trigger to 20msec, so that you are sampling at 50Hz. Set up two channels, both looking at the same pin, and use the “downsampling” option for one of the channels to downsample by a factor of 5, so that it is sampling at $50\text{Hz}/5 = 10\text{Hz}$. Note: the downsampling option is a rarely used feature, so it is hidden under the `>>` button on the channel.

Set the function generator for a very low frequency (say 0.1Hz) and observe the waveforms with PteroDAQ. Make notes in your lab notebook of what you observe. Make sure your notes include the frequency of the sine wave and the sampling frequency of each channel, so that you can reproduce the effect later. Save the recording to files, so that you can plot them in your report.

Increase the frequency gradually up to 10Hz. What happens as you change the frequency? What happens at 5Hz? at 10Hz? What if you are just a little off 5Hz (4.9Hz or 5.1Hz)? What happens if you go up to 25Hz?

Find interesting behavior at various frequencies. Describe the patterns you see and record them for inclusion in your report. Again, save any data you wish to talk about in files.

What happens to the interesting waveforms if you put your 4Hz low-pass filter between your 0.01Hz high-pass filter and the PteroDAQ? (Use another analog pin for the filter output, and add more channels.) Does the filter do what you expect? Is its corner frequency close to 4Hz?

Again, record some waveforms, keeping careful records in both your lab notebook and the meta-data “Notes:” field of the PteroDAQ files.

6 Demo and writeup

For the demo, show me some of the interesting waveforms, with and without the low-pass filter.

Your report should have schematics for your entire test setup. The capacitors can be drawn as polarized (with a little plus sign) in SchemeIt. Make sure your schematics show the polarity the right way around.

Your report should contain several of the interesting waveforms, together with explanations of why they occur and the effect of low-pass filtering on them.

Superimposing the higher sampling frequency plot from one channel and the lower sampling frequency plot from the other channel on the same graph helps to explain what is going on. You may need to specify a small xrange (a small time interval) to show what is going on clearly.

7 Design Hints

The low-pass filter does not do any level shifting, so can be designed as an RC voltage divider with no special tricks.

Remember that electrolytic capacitors have a polarity (a positive and negative side). Make sure that you connect them up the right way around, or they might blow up.

Remember to save some of your PteroDAQ recordings as files, so that you can use gnuplot to make figures for your lab report.