S2014, BME 101L: Applied Circuits Lab 9: class-D power amplifier

Kevin Karplus

May 19, 2014

1 Design Goal

In Lab 6, you made an audio amplifier for a microphone using an op amp, but it was not powerful enough to provide loud sounds from a loudspeaker, because the op amp had limited output current, and the loudspeaker had a low impedance so the power to the loudspeaker, I^2R, was small. (With a 23mA limit and 8Ω the maximum power was about $I^2R = 4$mW.)

In this lab, you will design an amplifier that is capable of delivering 2W–4W of power to an 8Ωloudspeaker, about 500–1000 times more than the one-op-amp audio amplifier of Lab 6.

Because you will be using power field-effect transistors (FETs), a class-D amplifier design is probably the easiest to get working. [3]

Design constraints: In addition to the parts in your kit, you should use the triangle wave output from the function generator and both the 6v power supply and the $\pm 25V$ dual power supply. The loudspeaker should be powered from the dual supply (though not at 25V!), but the preamplifier from the 6V supply.

2 Background

2.1 Setting the power supply

If you want to deliver 10W (the maximum RMS power) to your 8Ω loudspeaker, what is the maximum voltage and current you should apply? You should look at most powerful signals you can deliver to the speaker, which would be square waves that alternate between the positive and negative voltage rails. If you are using the symmetric bench supply for driving the loudspeaker, with one end of the loudspeaker on the 0V ground in the middle, the maximum power you can deliver is $(\pm V)^2/R$ or I^2R, so your power-supply voltage should be limited to about 9V, and the current limited to 1.1A, assuming an 8Ω resistive load. The symmetric power supply only provides up to 1A, though, so the largest power you can deliver to the loudspeaker is 8W.

The comparators we plan to use this year (TLC 3072) have a 16v power-supply limitation, so you are further limited to ±8v, which would also be an 8W limitation on an 8Ω load.

Prelab: If you want to deliver 4W, what voltage and current would you need from the power supply?

Note that the MCP6004 op-amp chips you are using have a 6v power supply limit, and probably should have a separate 6V supply to power any op-amp chips you use. Having a separate power supply for the preamp helps prevent feedback through the power-supply lines.

When setting your power-supply voltages—one warning: be sure that the output of the preamp remains between the power rails for the next stage (the comparator). Even after removing any DC offset, the preamp output has a peak-to-peak voltage that can be as big as the power supply to the preamp. If you use a separate power supply for the pre-amp, make sure its voltage is no higher than the range of voltages in the final-stage power supply, so that large input signals don't exceed the power supply limits.

2.2 Pulse-width modulation (PWM)

The idea behind class D amplifiers is that you don't provide the output as a voltage that is a fixed multiple of your input voltage (the classic idea of a linear amplifier). Instead, you provide a pulse train that rapidly

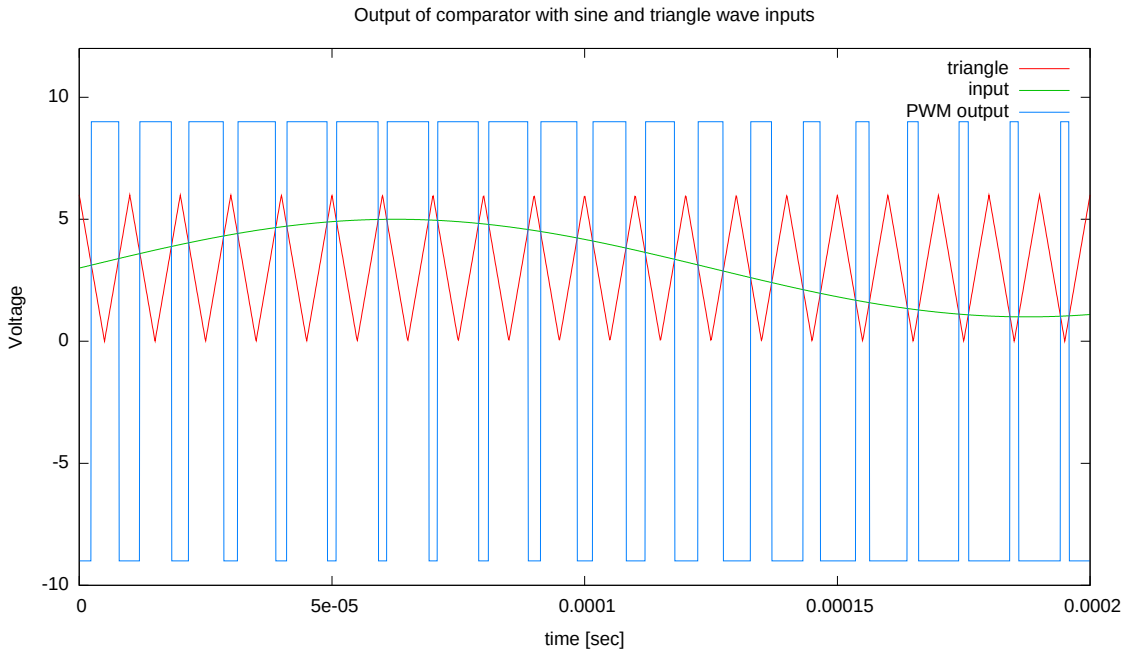

Figure 1: The output of a comparator for different input voltages with a triangle wave on the negative input of the comparator. The comparator power inputs are set to the positive and negative rails of the power-amp stage (here, ±9v). Note that the DC offset and amplitude of the triangle wave affect the behavior of the comparator. Here I've used a $\pm 3V$ triangle wave centered at 3V.

switches between the positive power rail and the negative power rail. What you adjust is how much time you spend at the positive rail and how much time you spend at the negative rail. This is referred to as *pulse-width modulation*, abbreviated PWM, because you are adjusting the duration of the high pulses. [4] A PWM signal is usually specified with two parameters: its *frequency* and its *duty cycle* (what percentage of the period the output is high).

You can get a pulse-width modulated signal from a voltage input by using a comparator and a triangle wave generator. If the voltage of your input is higher than the triangle wave, then your output is high. If the voltage of your input is lower than the triangle wave, then your output is low. (See Figure 1.) The amount of time the output spends in the high state corresponds to how much of the time the triangle wave is lower than your input. Since the triangle wave spends an equal amount of time at each voltage, the fraction of the time spent high (the duty cycle) is proportional to the input voltage, and the frequency of the PWM signal is the same as the frequency of the triangle wave.

If you make the PWM frequency very high, it is difficult to make the rise and fall times for the gate voltages of the output field-effect transistors (FETs) fast enough to follow the pulses (see Section 2.4). If you make the PWM frequency too low, it is difficult to filter out the PWM pulses adequately without also filtering out some of your desired audio frequencies. Normally, a PWM frequency in the range 40kHz to 100kHz is chosen to balance these conflicting constraints for audio work, but lower frequencies are often chosen for motors.

2.3 Inductive loads for PWM

Pulse-width modulation is often used for driving inductive loads like motors and loudspeakers. Both motors and loudspeakers are more dependent on the current through their coils than on the voltage across them,

because the magnetic field strength of a coil is proportional to the current.

PWM works well with motors and loudspeakers because the rapid changes in voltage result in much slower changes in current. The current is the integral of the voltage, so ramps up when the voltage is positive and ramps down when the voltage is negative. Thus one gets sawtooth-like fluctuations in the current around the average value, with the average value determined by the duty cycle of the PWM signal.

Furthermore, if the PWM signal has a very high frequency, the mass of the loudspeaker or motor keeps it from moving much in response to the high-frequency part of the waveform, so you only observe the sound waves or motor movement due to the lower frequency components.

In this class we won't rely on this mechanical filtering to remove the PWM signal, but use an external filter.

2.4 FETs as output stage

In this class, you use FETs mainly as non-linear, digital devices—that is, as switches that are either "on" or "off". One reason for using the power FETs as switches is that they are very energy efficient if used that way (they dissipate little power as heat). Another reason is that their characteristics are very temperature-dependent, so designing linear circuits with power FETs is quite difficult (beyond the scope of this class).

Why are FETs power efficient as switches? Let's look at them both when they are turned off and when they are turned on:

When an FET is off, it has a very high resistance between the source and drain and passes little current. I measured about 10nA–50nA for the AOI518 nFET transistor with $V_{GS} = 1V$, though its data sheet allows up to 5μA with for $V_{GS} = 0V$ and $V_{DS} = 30V$, if the transistor warms up to 55°C. Note that even with the data sheet worst-case conditions, the transistor would only dissipate 30V 5μA = 0.15mW of heat when it is off. This amount of power dissipation is unimportant in a power amplifier—too small to be worth bothering with.

When an FET is on, it has a very low resistance, and so the voltage drop across it is small. The data sheet for the AOI518 nFET specified an on-resistance of at most 11.9mΩ when $V_{GS} = 4.5V$ and 6.6mΩ when $V_{GS} = 10V$. If you use the FET to drive an 8Ω loudspeaker, the current will be limited by the power supply driving the speaker (a 9V power supply would be delivering 1.125A), so the power dissipated in a turned-on FET would be $I^2R = 15$mW, while the power to the loudspeaker would be about 10W.

The transistors are designed to handle large currents (up to 46A if kept cool, 36A if allowed to get up to 100°C)—much larger currents than you'll be dealing with. The main limitations of the FETs when they are fully on, though, is with self-heating: if you try to put 40A through them they'll dissipate $(40A)^2\, 11.9m\Omega = 19W$ and quickly heat up past their maximum temperatures unless good heat sinks are used to reduce the thermal resistance.

Although the nFETs you are using are supposedly capable of dissipating 50W of power without damage if kept at 25°C, you don't want to go there, as you can't keep them that cool. The junction-to-ambient thermal resistance with a small heat sink (a square inch of 2oz copper on a PC board) is 50°C/W, so keeping the junction below the 150°C maximum junction temperature limits us to about 2.5W dissipated by the nFET (about 14A with a resistance of 11.9mΩ). Without a heat sink, as you are using them, you need to keep the power dissipation in the FETs even lower—luckily we will be, since we're only going up to about an amp (dissipating about 12mW).

Prelab: Look up the data sheet for the IRFU9024NPbF pFET and figure out the maximum power and current you could safely have it handle without a big heat sink.

In summary, FETs dissipate little power because they have extremely low current when off and low voltage drops when on. When FETs are switching between the off and on states they can have intermediate resistance values with large currents and large voltages and so dissipate large amounts of power. For

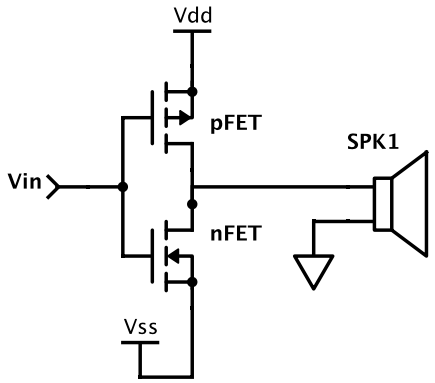

Figure 2: A cMOS output stage, consisting of one pFET with its source connected to the positive voltage rail and one nFET with its source connected to the negative voltage rail, and the output from the two connected drains. Here one side of the loudspeaker is connected directly to the drains of the FETs, but see Section 2.8 for information about adding an output filter.

example, if you were driving an 8Ω load and used an FET to get the voltage only half way to the power rail, then the FET would have the same resistance as the load, and each would be dissipating the same amount of power (2.5W in the example of a 9V power supply). To maintain efficiency and prevent overheating the FETs, we try to switch between the on and off states as quickly as possible.

The standard configuration for an FET output stage in an amplifier (either a class-D amplifier like you're designing or in a cMOS digital circuit) is to have one nFET and one pFET, with their drains connected to the output and their sources connected to the power rails, as shown in Figure 2. The acronym "cMOS" stands for "complementary Metal Oxide Semiconductor" and refers to this sort of pairing of an nFET with a pFET.

Note: because nFETs have somewhat better electrical characteristics than pFETs, high-power circuits sometimes use nFETs for both legs, but the circuitry needed to control an nFET that is connected to the positive supply gets a little complicated, and we don't need that much power, so we'll stick with the simpler cMOS design.

The cMOS FET output stage works great if one of the two FETs is turned on and the other off, providing a low-resistance connection to the desired voltage rail and no connection to the other rail. If both FETs are turned on, though, it can pass a large current from one power rail to the other (called *shoot-through current*), heating up the FETs to no useful effect.

Various circuits can be used to ensure that you don't have both FETs on at once. The simplest of them is just connecting the two gates together and switching the gate voltage very rapidly from a voltage near the top power rail to a voltage near the bottom power rail. This design is used in logic gates called *inverters*, whose name comes from their function of changing a high input to a low output and vice versa. The Schmitt-trigger inverters you used in the hysteresis oscillator probably have this sort of design in their output stage.

Note that when the gate voltages are at the top power rail for a ±6v supply, the pFET has $V_{GS} = 0V$ and the nFET has $V_{GS} = 12V$ so the pFET is off and the nFET is on. When the gate voltages are at the bottom power rail, the pFET has $V_{GS} = -12V$ and the nFET has $V_{GS} = 0V$, so the pFET is on and the nFET is off. There are roughly three voltage ranges for the tied-together gates: below the nFET threshold (only the pFET is on), between the thresholds (both FETs on), and above the pFET threshold (only the nFET is on).

The simple inverter design has problems as the power supply voltages on the nFET and pFET sources

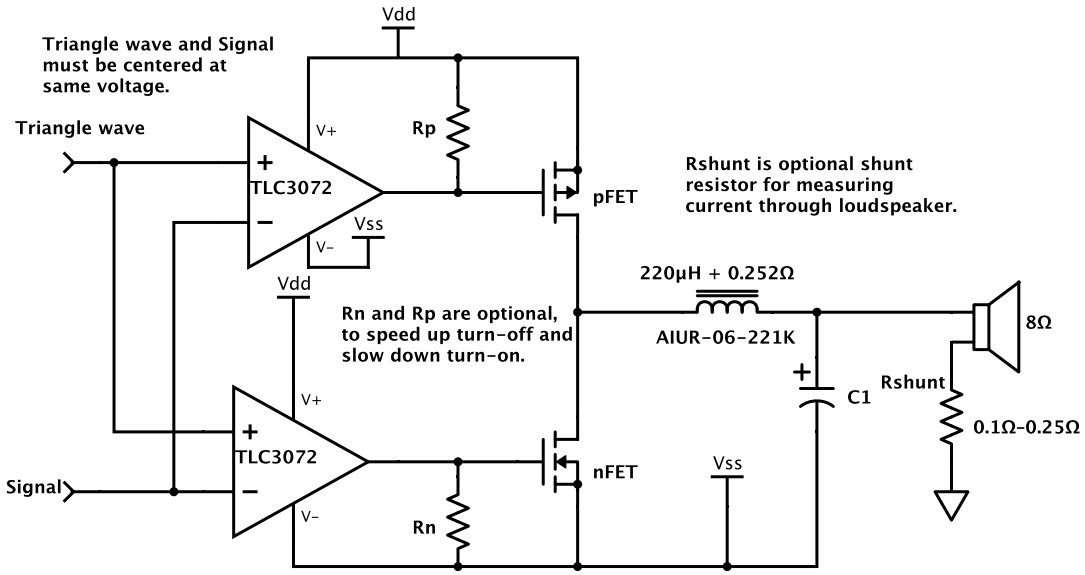

Figure 3: By having separate control signals for the gates of the nFET and pFET, we can arrange the signals to the gates so that the FETs are not simultaneously on. The basic idea is to make sure that the FETs are turned off quickly, even at the cost of turning on more slowly.

increase, for two reasons:

- The nFET threshold is referenced to the nFET source (the lower power rail) and the pFET threshold is referenced to the pFET source (the upper power rail). Increasing the voltage difference between the power rails means that there is a wider range of gate voltages between the thresholds, where both FETs to be on.
- The comparators that drive the gates (see Section 2.6) have limited current capability, but the gate capacitance is roughly constant, so as the power supply voltage increases, it takes longer to charge or discharge the capacitance of the gate. That means that more time is spent in the region where both FETs are on.

Unfortunately, while the input is switching, the power consumption is quite high, and the FETs get hot, so we won't use this circuit. Instead of tying the gates of the nFET and pFET together to make a single-input cMOS inverter, you'll need to create separate control voltages for the two gates.

By controlling the two gates separately, we can try to turn off one FET before turning on the other (or, at least, keeping the time when both are on a very small fraction of the time), reducing the wasted power that heats the FETs when both are on. Figure 3 shows one approach for doing this.

2.5 Switching speeds of FETs

Estimating how long it takes for an FET to switch can be tricky—simple RC charging models for the gate don't work very well, because the switching signal on the output is capacitively coupled back to the gate.

The simplest model that works fairly well is the "gate-charge" model, which has 3 parameters: the *gate source charge* Q_{gs}, the *gate drain charge* Q_{gd}, and the total gate charge Q_g, which is larger than the sum of the other two.

When turning on an nFET, the gate voltage ramps up until it reaches the threshold voltage, at which point the drain current I_{ds} ramps up while the drain voltage V_{ds} stays high, then the drain voltage ramps

down while the current stays high. The initial ramping up of current lasts until a charge of Q_{gs} has been added to the gate, and the ramping down of voltage until another charge of Q_{gd} has been added. The total gate charge is an estimate of how much charge has to be added to the gate to get it all the way to its final voltage. The time each phase takes can be estimated from the charge divided by the current that is charging the gate.

When turning off an nFET, the gate voltage ramps down until it is near the threshold voltage, at which point the drain voltage ramps up while drain current stays high, then the drain voltage V_{ds} stays high, while the current ramps down. The ramping of the drain voltage takes the time to transfer Q_{gd} off the gate and the ramping of the drain current the time to transfer Q_{gs}. The extra charge in Q_g mainly determines how much delay there is before the drain voltage starts to change.

So Q_{gd} is associated with the ramping of the drain voltage for both rising and falling edges, and Q_{gs} with the ramping of the drain current. For both rising and falling edges, the time when there is power being dissipated in the FET is roughly the time it takes to move $Q_{gs} + Q_{gd}$ onto or off of the gate.

A pFET behaves the same way, except that the voltages and currents are all negated.

Note: I've neglected some important stuff here, in the interest of keeping the model simple enough for back-of-the-envelope calculations. There are more detailed explanations in application notes from FET manufacturers [1].

2.6 TLC3072 comparators

Your FET output stage needs to have gate voltages that make a large voltage change rapidly, but our MCP6004 op amps are not capable of doing that. Not only are they limited to a maximum voltage swing of about 6V, they have a designed-in speed limitation: a slew rate of 0.6V/μs, so it would take them 10μs to do even their full 6V swing.

Because of these limitations of the op amp, we'll be using a different chip to drive the FETs, a dual comparator chip. The parts kit has a LM2903 chip, but we won't be using the LM2903 chips this year, as I realized that the nFET we are using this year will be very difficult to turn off rapidly enough to avoid shoot-through current using the LM2903. Instead, I have ordered TLC3072 comparators, which have a larger current capability than the LM2903 and can provide both pull-down and pull-up current (the LM2903 has an "open-collector" output which can provide pull-down current but needs an external resistor to provide pull-up current). The TLC3072 comparators have a maximum supply voltage of only 16v (limiting the power supply we can use for the amplifier to ±8v).

The resistors R_n and R_p in Figure 3 may not be needed. They are optional additions that you can use to speed up turning off one of the FETs, by providing extra current to pull its gate towards its source. For example, R_n would pull the gate of the nFET towards V_{ss}, turning off the nFET more rapidly. The problem with adding the resistors is that they slow down turning on the FET. When the nFET is being turned on, it can only get up to a voltage of $V_{GS} = I_{OH}R_n$, since at that voltage the resistor is stealing all the current and no further charging can take place. If R_n is too small, V_{GS} will always be below the threshold voltage and the nFET won't turn on.

Prelab: look up the maximum currents that the TLC3702 can supply at the power supply voltages you plan to use. The I_{OH} current is for pulling up and I_{OL} current for pulling down the output. Look up the threshold voltages as which the nFET and pFET are guaranteed to turn on. Based on these numbers, what are the smallest R_n and R_p values you could possibly consider?

2.7 Real power

To compute the energy dissipated in an impedance you need to multiply the instantaneous voltage by the current, then integrate over time. To get the average power, you divide the energy by the time you integrated over.

For example, if the voltage across a device is $V(t) = A\cos(\omega t + \phi)$ and the current into the device is $I(t) = B\cos(\omega t + \psi)$, then their product is

$$AB\cos(\omega t + \phi)\cos(\omega t + \psi)$$

or, using the trig identity $\cos(a)\cos(b) = 0.5(\cos(a+b) + \cos(a-b))$,

$$0.5AB\left(\cos(2\omega t + \phi + \psi) + \cos(\phi - \psi)\right) \ .$$

Note that there is a periodic part, with frequency twice as high as the waveform we started with and a constant part. If you average the power over a full period, the periodic part integrates to 0, and you have just the constant part: $0.5AB\cos(\phi - \psi)$.

Note: for a resistive load, current and voltage are in phase ($\phi = \psi$), the cosine of the difference is 1, and the real power is $0.5AB$, just as we computed when we talked about RMS voltage earlier this quarter.

It turns out to be easier to analyze sinusoids using *RMS phasors*, representing

$$V(t) = A\cos(\omega t + \phi) \text{ as } \bar{V} = Ae^{j\phi}/\sqrt{2} \ ,$$

and

$$I(t) = B\cos(\omega t + \psi) \text{ as } \bar{I} = Be^{j\psi}/\sqrt{2} \ .$$

So a voltage amplitude of A results in an RMS voltage of $A/\sqrt{2}$ and a current amplitude of B results in an RMS current of $B/\sqrt{2}$. The main point of RMS voltage and RMS current is that you can multiply them to compute power as $\bar{V}\bar{I}^*$, where the asterisk represents the complex conjugate (negating the imaginary part, which has the same effect as negating the phase). This gives us the right result for voltage and current being in phase ($\phi = \psi$), no matter what phase that is.

If $\bar{V}$ and $\bar{I}$ are RMS phasors, then $\bar{V}\bar{I}^*$ is called the *complex power*, $|\bar{V}\bar{I}^*|$ is called the *apparent power*, $\text{real}(\bar{V}\bar{I}^*)$ is called the *real power*, and $\text{imag}(\bar{V}\bar{I}^*)$ is called the *reactive power*. Power engineers like to change the names of the units (watts for real power, volt-amps for apparent power, and vars for reactive power). They're all the same unit, but changing the name reminds you that they are not computed the same way and are not interchangeable. You will not be held responsible for those names in this course, and may use watts for all forms of power, but we will expect you to mean "real power" unless you explicitly modify the noun.

The real power is what you computed by doing the integral of the cosine waves—it represents the power actually dissipated in the device. Note that a sine wave generator dissipates negative power (the current into the device and the voltage have opposite signs).

The reactive power represents power that is being temporarily stored in the device during part of the cycle, and released in a different part of the cycle. If you have a purely reactive circuit (just capacitors and inductors with no resistors, so the impedance is purely imaginary), then the voltage and current are 90° out of phase, and the real power is $0.5AB\cos(\pi/2) = 0$. No energy is dissipated in a purely reactive circuit.

Note that if you use *amplitude phasors* (representing $A\cos(\omega t + \phi)$ as $\bar{V} = Ae^{j\phi}$ without the $\sqrt{2}$ correction) rather than *RMS phasors*, the power computation $\text{real}(\bar{V}\bar{I}^*)$ gives us double the actual power. That is why our loudspeakers are rated as 10W RMS power (the actual power they can dissipate safely) or 20W peak power (the number you get if you multiply the amplitude phasors).

If you know the RMS voltage phasor $\bar{V}$ for a sinusoid to a complex impedance Z, the RMS current phasor is just $\bar{V}/Z$, and the real power is $\bar{V}(\bar{V}/Z)^*$.

You'll want to compute the real power that your amplifier delivers to the loudspeaker, to make sure you don't exceed 10W RMS. You'll also want to compute how much power is delivered at frequencies that

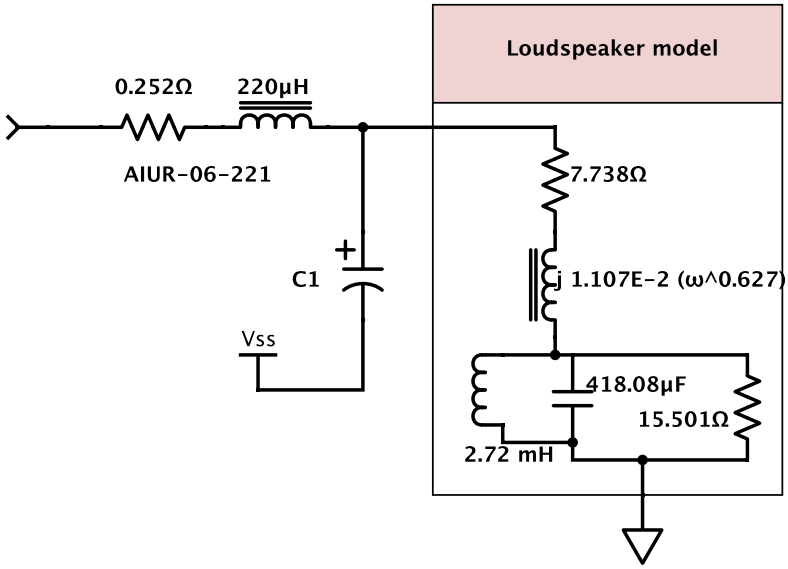

Figure 4: An LC low-pass filter combined with the loudspeaker model, without a Zobel network for compensating the speaker inductance. The 0.252Ω resistance is not a resistor that you add, but a model for the resistance of the wire in the inductor.

do not produce useful sound (like the PWM frequency discussed in Section 2.2), to make sure that you are not wasting power heating up the loudspeaker to no useful effect.

Manually computing the complex impedance of a loudspeaker or a complicated network of resistors, capacitors, and inductors is tedious and unnecessary. Use gnuplot with models for the loudspeaker that you developed in Lab 5A to plot the impedance or power (or use the "print" statement to do a calculation of a single value).

2.8 Output filter overview

In professional class-D amplifiers for audio work, the output of the cMOS stage is filtered by an LC low-pass filter (see Figure 4) [2]. The corner frequency of the filter needs to be between the highest input frequency to be amplified and the frequency of the triangle wave, which is a much higher frequency.

Motor drivers often don't bother with the LC filters, though they can reduce heating in the motor and make electromagnetic interference easier to control. One reason is that motors require a lot of current, and power inductors capable of handling large currents are expensive. The little ones I bought for the class are only about 60 cents, but they have a saturation current of 1.9A—plenty for our 10W loudspeakers but not a lot for a motor.

You could skip the output filter in your design, and connect your loudspeaker directly from the drains of the pFET and nFET to the ground of the dual power supply. The amplifier would work, but you would have high-frequency ripple in the current to the loudspeaker, some of which may cause radio-frequency interference with AM radios.

One trick that speaker designers sometimes use is to add an R+C impedance in parallel with the speaker, to try to flatten out the impedance at high frequencies (a *Zobel* network). This makes the speaker look more like a constant resistive load to the amplifier, so that the amplifier design can be more independent of the loudspeaker.

In the handout for Winter 2013, I spent a lot of space analyzing a Zobel network, concluding that

the current through the resistor and capacitor gets quite large, wasting power and requiring a high-power resistor. We don't have high-power resistors in the parts kit (only 1/4W resistors), and we don't want to waste the power from the amplifier.

Can you design the low-pass filter directly, without a Zobel network to make the speaker look like a simple resistive load (like in Figure 4)?

We need to choose the inductor and the capacitor values. Initially, I was going to have students wind their own inductors around cardboard tubes, picking the number of turns of wire using an inductance calculator like the ones at
`http://www.pronine.ca/multind.htm` and
`http://www.crystalradio.net/professorcoyle/professorcoylecyl.shtml`
But after figuring out how much wire would be needed and making an attempt to wind one myself, I decided it would be better and cheaper if I just bought some inductors for the class. I ended up buying a 220μH inductor with a 0.252Ω resistance and a maximum current of 1.9A (the AIUR-06-221K made by Abracon Corporation). These are high-power ferrite-core inductors, but they have one problem: their leads are too fat for most breadboards, so you'll probably have to attach them to your circuit with clip leads.

To make design easier for you, I've also written a gnuplot script (Figure 5) that plots the real power into the loudspeaker if you use this inductor with various standard capacitor sizes. It even allows you to set the value for a small resistor in series with the loudspeaker for viewing current waveforms on the oscilloscope. (Such a resistor does change the LC design somewhat.)

You can also use the gnuplot script to help you choose a power-supply voltage that will deliver between 2W and 10W of power to the loudspeaker, using the detailed model of the loudspeaker and LC filter, rather than just an 8Ω estimate. You should replace the loudspeaker model in the script (Zloud) with the model that you developed in Lab 5A. I show the results for one setting with the parameters from my loudspeaker model in Figure 6.

It is important to remember that (for a fixed load) current will go up along with voltage, and so the real power will go up with the square of voltage. Also, the power computations in the gnuplot script are for sine waves, but the real power to the loud-speaker may be higher with other input signals (with square waves maximizing the power for a given power supply—twice the power of a sine wave).

```
set logscale xy
set xlabel "frequency [Hz]"
set xrange[10:1e6]
set mxtics 10 # to get standard logarithmic tics at 10,20,30,...,90

set key bottom left Left
set samples 1000
set grid

j=sqrt(-1)
conjugate(a) = real(a) - j*imag(a)
phase(v) = imag(log(v))
real_power(V,Z)= real(V*conjugate(V/Z)) # V is RMS voltage phasor

divider(Zup,Zdown) = Zdown/(Zup+Zdown)
Zc(f,C) = 1/(j*2*pi*f*C)  # f= frequency [Hz], C=cap [farads]
Zl(f,L) = j*2*pi*f*L
Zlsemi(f,L,power) = j*(2*pi*f)**power*L
Zpar(z1,z2) = z1*z2/(z1+z2)

#model for impedance of loudspeaker  USE YOUR MODEL
#impedance of loudspeaker
Zloud(f) = Rnom+Zlsemi(f,Lhi,Lhipow)+Zpar(Rs,Zpar(Zl(f,Ls),Zc(f,Cs)))

Rnom=7.738
Lhi= 1.107e-2    ;    Lhipow = 0.627
Ls=2.72e-3      ;    Cs=418.08e-6   ;  Rs=15.501

# CHOOSE A SERIES RESISTOR IF YOU PLAN TO MEASURE LOUDSPEAKER CURRENT
Rseries = 0.0 # extra resistor for displaying current on scope

# real power into loudspeaker with voltage from amplifier V at frequency f
# if no filter is used
nofilter_power(V,f)=real_power(V*divider(Rseries,Zloud(f)), Zloud(f))

# LC FILTER MODEL

# INDUCTOR VALUES FROM SPEC SHEET
Lfilt=220e-6;   Rfilt=0.252     # parameters for AIUR-06-221 inductor
# Lfilt=270e-6;   Rfilt=0.270     # parameters for AIUR-06-271 inductor
# Lfilt=330e-6;   Rfilt=0.394     # parameters for AIUR-06-331 inductor

# voltage divider with Lfilt+Rfilt and (Cfilt || loudspeaker+resistor)
gain(f,Lfilt,Rfilt,Cfilt) = \
divider(Zl(f,Lfilt)+Rfilt, Zpar(Zloud(f)+Rseries,Zc(f,Cfilt)))

# total impedance, in case we want to plot that also
Ztotal(f,Lfilt,Rfilt,Cfilt) = Zl(f,Lfilt)+Rfilt + Zpar(Zloud(f)+Rseries,Zc(f,Cfilt))

# real power into loudspeaker with voltage from amplifier V at frequency f
# using the LC filter
loud_power(V,f,Lfilt,Rfilt,Cfilt) = \
real_power(V*gain(f,Lfilt,Rfilt,Cfilt)*divider(Rseries,Zloud(f)), \
  Zloud(f))

 # CHOOSE YOUR SUPPLY VOLTAGE
V_supply=4    # +/- power supply voltage

V_RMS = V_supply/sqrt(2)        # RMS voltage for largest sine wave from supply

half_power=0.5*loud_power(V_RMS,300,Lfilt,Rfilt,3.3e-6) # half-power point

set title sprintf("Low-pass LC filter for loudspeaker with %.3gohm series resistor",Rseries)
set yrange [1e-4:40]
set ylabel sprintf('Real power to speaker with +- %.3gV supply [W]',V_supply)

plot half_power title sprintf("half-power=%.3gW",half_power), \
        nofilter_power(V_RMS,x) title "no filter", \
        loud_power(V_RMS,x,Lfilt,Rfilt,0.1e-6) title sprintf("L=%.3guH C=%.3guF", Lfilt*1e6, 0.1), \
        loud_power(V_RMS,x,Lfilt,Rfilt,1e-6) title sprintf("L=%.3guH C=%.3guF", Lfilt*1e6, 1), \
        loud_power(V_RMS,x,Lfilt,Rfilt,2.2e-6) title sprintf("L=%.3guH C=%.3guF", Lfilt*1e6, 2.2), \
        loud_power(V_RMS,x,Lfilt,Rfilt,3.3e-6) title sprintf("L=%.3guH C=%.3guF", Lfilt*1e6, 3.3), \
        loud_power(V_RMS,x,Lfilt,Rfilt,4.7e-6) title sprintf("L=%.3guH C=%.3guF", Lfilt*1e6, 4.7), \
        loud_power(V_RMS,x,Lfilt,Rfilt,10e-6) title sprintf("L=%.3guH C=%.3guF", Lfilt*1e6, 10)
```

Figure 5: Gnuplot script to choose capacitor values to provide low-pass filtering to the speaker. You'll have to modify the script to incorporate your loudspeaker model, and adjust the power supply voltage and the series resistor. This script will be made available at `http://users.soe.ucsc.edu/~karplus/bme101/s14/loudspeaker-low-pass-choose-C-power-law.gnuplot`

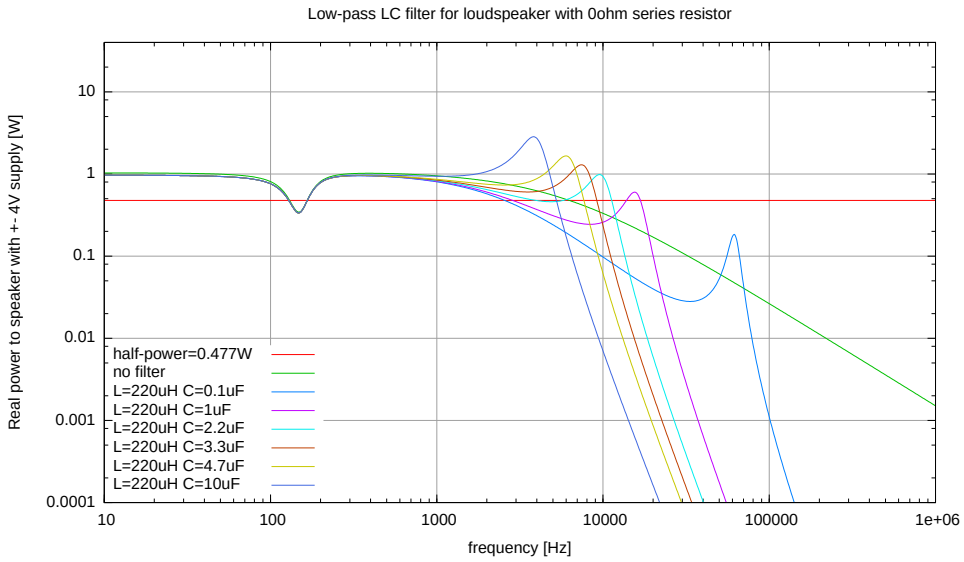

Figure 6: The result of running the script in Figure 5 for modeling power to the loudspeaker. The power to the loudspeaker should be very low at the PWM frequency, fairly constant in the range of frequencies you want to amplify, and below the 10W limit for the speaker at all frequencies.

3 Pre-lab assignment

Select a partner at least 2 days before the lab and start working with them.

Make some design specifications for your amplifier: what is the frequency range you wish to amplify? what is the power output of the amplifier? What AC voltage or current does you input have?

What frequency do you want your PWM signal to be at? It should be at least 3 times the frequency of the highest frequency you want to amplify and above the range of human hearing. If it is too high, the FETs will be spending a greater portion of their time in the intermediate state where they are neither on nor off, and the efficiency of the amplifier will be reduced (and the FETs might get hot).

As you did for the one-op-amp audio amplifier lab, make a block diagram of the whole amplifier, from microphone to loudspeaker, giving the function of each block and any constraints (voltage, current, frequency, ...) on the signals between the blocks. What are the voltage (or current) levels (DC and AC) for every signal that goes between blocks? Remember to include blocks for the power supplies and for the function generator.

Your preamplifier using the op-amp chip should be powered from the 6v single supply, while the comparator, cMOS output stage, and loudspeaker should be powered from the dual power supply. Think about how these separate voltage sources relate to one another, and make sure you include them in your schematics.

If you need to change DC voltage levels from one block to another, remember that you can use high-pass filters to do the conversion (as we did between the microphone and the op amp in the one-op-amp amplifier).

Think about how you can control the gain of the amplifier, since it might be too loud and no one wants to listen to feedback squeal. If your pre-amp stage has too much gain, you may get clipping and distortion, even if the overall loudness is reasonable. There are at least three different ways to control the output volume—so try to think up more than one and figure out what the trade-offs are between them

Design each block.

Put a 0.25Ω resistor in series with your loudspeaker (you may have to make it out of larger resistors in parallel), so that you can use the oscilloscope to look at the current through the speaker as well as the voltage across it. Larger resistor values could cause problems with the amount of power dissipated in the resistor—if you have 1A of current, the power from a 1Ω resistor would be 1W, and our resistors can only handle 0.25W.

Draw a complete schematic of the entire amplifier, showing every component value and every pin number. You will be building and debugging from this schematic, so it must be as complete and accurate as you can make it.

4 Parts to look up specs for

Parts for this lab from kit:

- CUI inc CMA-4544PF-W electret microphone
- MCP6004 quad op amp
- AOI518 nFET
- IRFU9024NPbF pFET
- 10W 8Ω loudspeaker `http://www.parts-express.com/pe/pshowdetl.cfm?&Partnumber=299-919`
- breadboard
- various resistors and capacitors

Loaned by instructor:

- AIUR-06-221K 220μH inductor
- TLC3072 dual comparator

5 Procedures

Build and test your circuit a block at a time, starting from the microphone. It is very hard to debug a large circuit when the problem could be anywhere.

It might be a good idea to build all the 6V circuitry (the preamp) on one breadboard and all the dual-supply circuitry (the power amp) on a different breadboard. There should then be only 2 wires between the boards: the input to the comparator and a reference voltage to make sure that the input remains within the power rails of the comparator.

Check the function generator voltage with an oscilloscope before hooking it up to the comparator input.

It is a good idea to build the preamp starting forward from the microphone, testing each block as you add it. The LC filter and loudspeaker should be thought of as a single block, since the LC filter does something quite different without the loudspeaker as a load.

It is also a good idea to start out the power amp section with a lower voltage and current limit on the power supply than you intend to finally use it, to avoid overheating components before the wiring has been debugged.

6 Demo and writeup

Demonstrate your amplifier working, producing undistorted amplified speech. Show the output waveforms (current and voltage) on the oscilloscope. It is worthwhile to show the waveforms with current and voltage as the two channels, so the phase relationship between the two can be seen. It is then useful to observe it at at two different scales: one where the PWM waveform is clearly visible, maybe 2μs/division, and one where the audio output is clearly visible, maybe 1ms/division.

Document both the system-level design and the details of each block in your report.

References

[1] Jess Brown. Power MOSFET basics: Understanding gate charge and using it to assess switching performance. Technical Report Application Note 608, Vishay Siliconix. `www.vishay.com/docs/73217/73217.pdf`

[2] W. Marshall Leach. *Introduction to Electroacoustics and Audio Amplifier Design*, chapter The Class-D Amplifier. Kendall/Hunt, second edition—revised printing edition, 2001. `http://users.ece.gatech.edu/mleach/ece4435/f01/ClassD2.pdf`

[3] Wikipedia. Class-d amplifier. read 17 Feb 2013. `http://en.wikipedia.org/wiki/Class_d_amplifier`

[4] Wikipedia. Pulse-width modulation. read 17 Feb 2013. `http://en.wikipedia.org/wiki/Pulse-width_modulation`