Department of Engineering

Oscillators

Matthew Spencer

Harvey Mudd College

E151 – Analog Circuit Design

1

In this video series we're going to talk about oscillators. We've focused a bit on how to make our amplifiers stable, so it's fun to flip the script and ask how to make them unstable on purpose. This is our last video series, so thanks for watching and I hope you enjoyed the class!

Department of Engineering

Phase Shift Oscillators

Matthew Spencer
Harvey Mudd College
E151 – Analog Circuit Design

2

In this video we're going to talk about phase shift oscillators, which are oscillators that use a few poles to set their loop gains to -1.

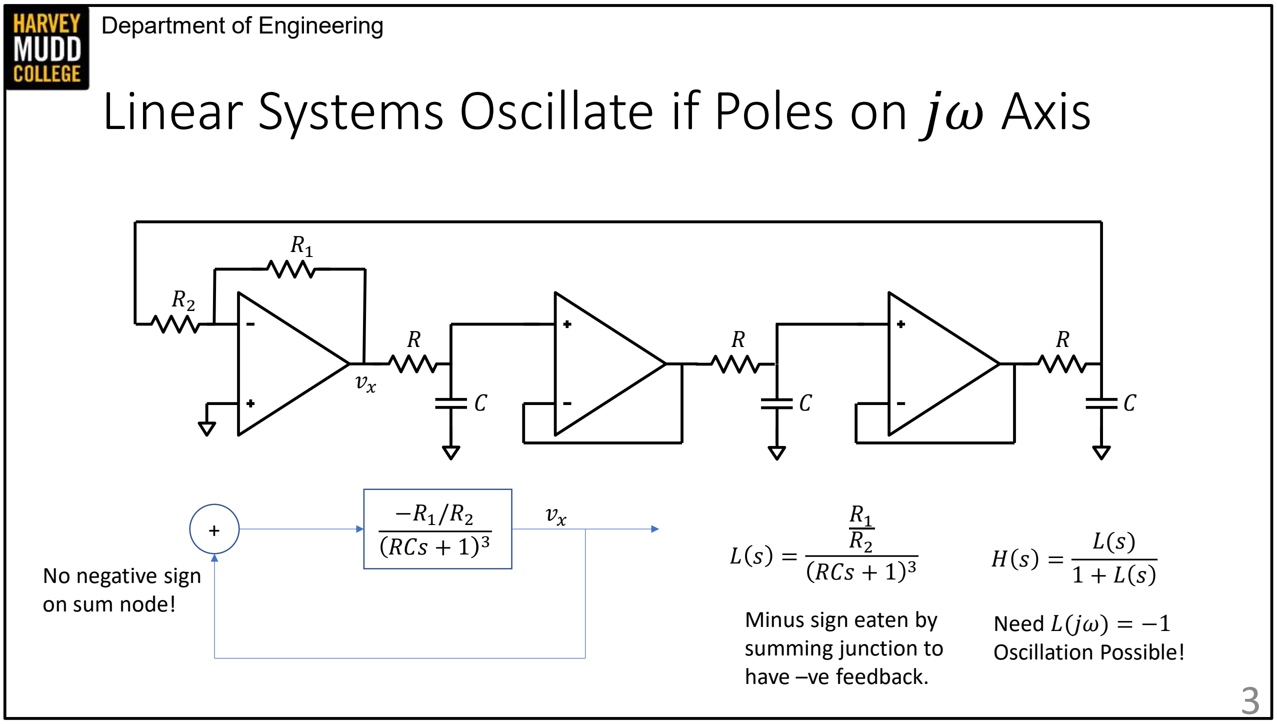

We're going to start with a brain teaser. Pause the video and try to analyze whether it is possible at all for this system to oscillate, and assume you can pick whatever R, C, R1 and R2 you want in order to achieve that oscillation. This is asking you to solve a problem you haven't seen yet, so think about what you know about loop gain to try to answer the question.

CLICK The answer is, yes, it is possible for this system to oscillate. Getting to that answer relies on a bit of analysis, and we'll kick that off with this block diagram. We found the loop gain by multiplying the gain of each op-amp – 1, 1, and –R1/R2 – and the transfer function of each of the dividers at the op-amp outputs – 1/(RCs+1) three times, so it's cubed. The system is in unity gain feedback because we just have a wire from the final stage back to the first one. However, this block diagram looks a bit funny because the summing junction doesn't have a negative sign. (It also doesn't have an input, but that's normal for an oscillator.) That means this looks like positive feedback around an inverting gain.

CLICK So that means we should factor out the factor of -1 from the block in our diagram to find our loop gain. The factor of -1 gets eaten up by our summing junction to make this a normal negative feedback loop.

CLICK In negative feedback, the closed loop transfer function is L(s)/(1+L(s)) as we've seen before, so we want L(jw)=-1 to have oscillation. We have three poles, so we can get a phase of up to -3pi/2, and that means we can pick some frequency where the phase is –pi, then pick R1/R2 to achieve a magnitude of 1. It's worth noting that this analysis hinges on the phase of the loop gain, while the magnitude is something we feel like we can correct as an afterthought. That's an insight that should guide the analysis of most oscillators; the phase of the loop gain is hard to change, so it determines whether a circuit will oscillate, while the magnitude of the loop gain is relatively easy to tune, so it doesn't often prevent oscillation.

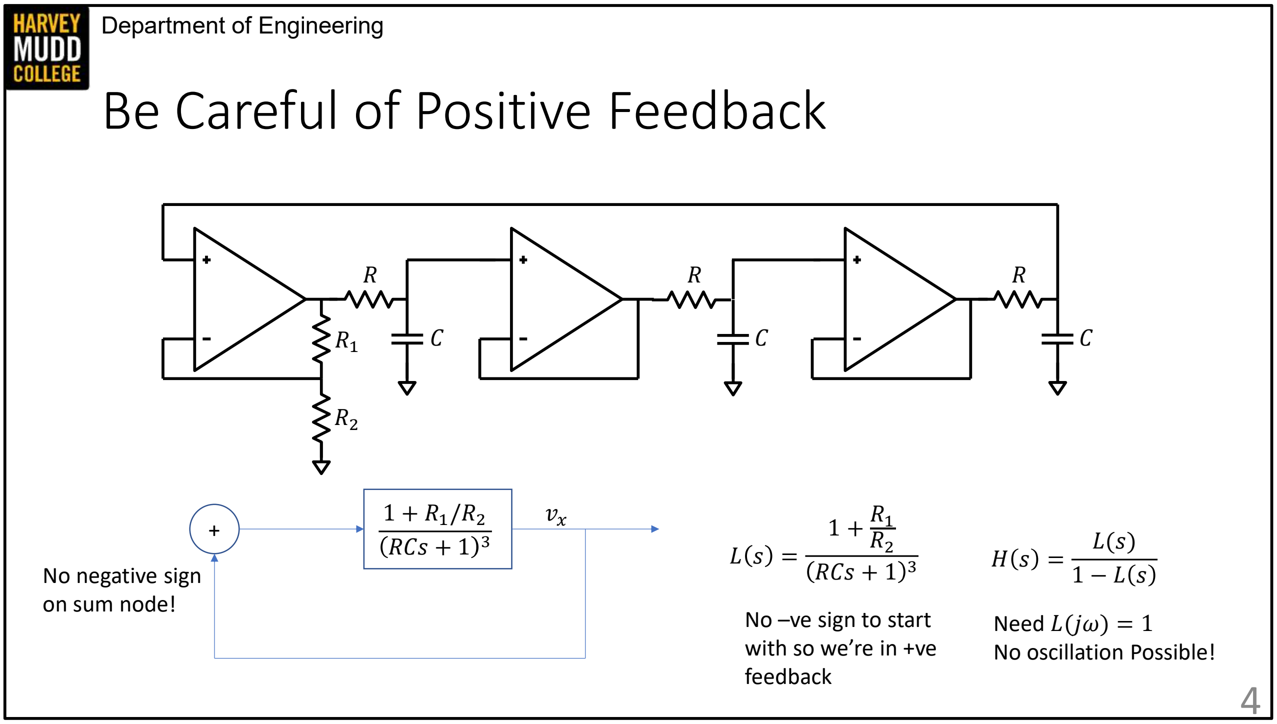

OK, we're going to do this again with a new amplifier. Pause the video and figure out whether it's possible for this circuit to oscillate. Note that the only thing that has changed relative to the last slide is that the first op-amp is a non-inverting amplifier instead of an inverting amplifier.

CLICK This circuit can't oscillate, which you might find surprising because it's in positive feedback. You've probably been taught that positive feedback is a recipe for instability, but it's not a guarantee of oscillation. That said, we found this block diagram the same way as the block diagram on the previous slide. The inverting gain of –R1/R2 has been replaced by the non-inverting gain of 1+R1/R2 in the numerator, and we have the same three collocated poles. Like before, this is in positive feedback, we don't have a negative sign at our summing junction.

CLICK That means our loop gain is just the contents of our block, but that we're in net positive feedback.

CLICK The closed loop gain of a system in positive feedback is L(s)/(1-L(s)), so we need L(jw)=1 to make the system oscillate. We have up to 270 degrees of phase shift in this circuit, but we'd need 360 to achieve a condition sufficient for oscillation.

Department of Engineering

Can Find Gain Per Stage to Achieve Oscillation

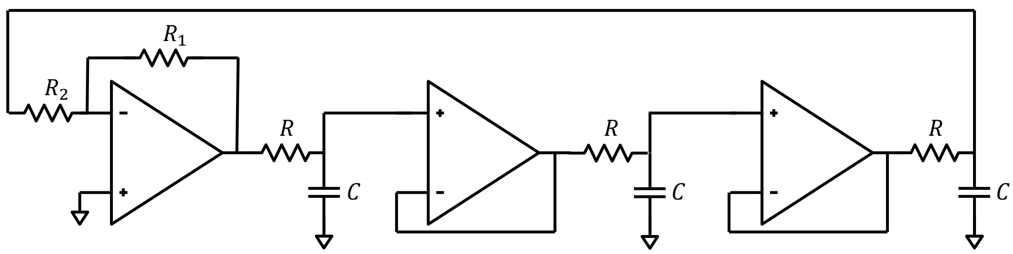

$$L(s) = \frac{\frac{R_1}{R_2}}{(RCs+1)^3}$$

Need $L(j\omega) = -1$

All poles co-located, so each gives 60 degrees of phase

$$H_{1pole}(j\omega_{0.33\pi}) = \frac{1}{j\omega_{60}RC+1}$$

$$\angle H_{1pole}(j\omega_{0.33\pi}) = -60° = -\operatorname{atan}\omega_{0.33\pi}RC$$

$$\rightarrow \omega_{0.33\pi} = \sqrt{3}/RC$$

$$\left|L\left(\frac{j\sqrt{3}}{RC}\right)\right| = \frac{R_1}{R_2}\left|\frac{1}{\sqrt{\left(RC\cdot\frac{\sqrt{3}}{RC}\right)^2+1^2}}\right|^3 = \frac{R_1}{R_2}\cdot\frac{1}{8}$$

$\frac{R_1}{R_2} = 8$ for oscillation. Oscillate at $\omega_{osc} = \omega_{0.33\pi}$.

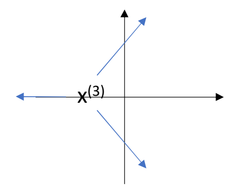

5

OK, so we’re going to go back to the circuit that we know can oscillate, which is called a phase shift oscillator, for reference. We’d like to figure out what values of R1 and R2 we should use to make this oscillate, and we’d also like to know what frequency it will oscillate at. I’ve copied our loop gain and oscillation condition on the right of the slide for reference.

CLICK We’re going to start by noting that because all of our poles are in the same location, they each have to contribute equally to the desired phase of -pi degrees. That means each pole will be contributing –pi/3 degrees of phase at our frequency of oscillation.

CLICK We can look at our transfer function for one pole, which is just given by an impedance divider between R and C. I introduced the variable omega_pi/3, which is the frequency at which the pole contributes –pi/3 radians of phase. By definition, we know the angle of H1pole at omega_pi/3 is –pi/3, but we also know that has to be equal to the arctangent of omega_pi/3*R*C. If we dig deep in our basic trigonometry, we can remember that pi/3 is 60 degrees and that 30-60-90 triangles have a height of sqrt(3), so the argument to our arctangent is sqrt(3), which in turn tells us that omega_pi/3 is sqrt(3)/RC.

CLICK We can substitute this frequency into our loop gain and take the magnitude of it. This will tell us how much gain we need to cancel loss from the poles. The RC values in the frequency cancel with the RC values in the transfer function, so we find that each pole has a magnitude of 1/2, and the combined effect of all the poles has a magnitude of 1/8. That means R1/R2 need to equal 8 so that we have a loop gain with magnitude 1 at omega_pi/3. Also, we know that we're going to oscillate at a frequency of omega_pi/3 because that's where our loop gain meets the oscillation criteria.

CLICK Finally, we can compare this loop gain analysis to a root locus plot. The root locus for three co-located poles shows them flying away at equal angles from one another. It's plausible that a gain of 8 is the exact value that results in the conjugate pair of poles landing on the y axis.

Summary

- Linear systems oscillate if poles are on the $j\omega$ axis.
  - Same condition as $L(j\omega) = -1$ if system is in –ve feedback.
- Can design gain of amplifiers in feedback to achieve $L(j\omega) = -1$
- Classic example is the phase shift oscillator:

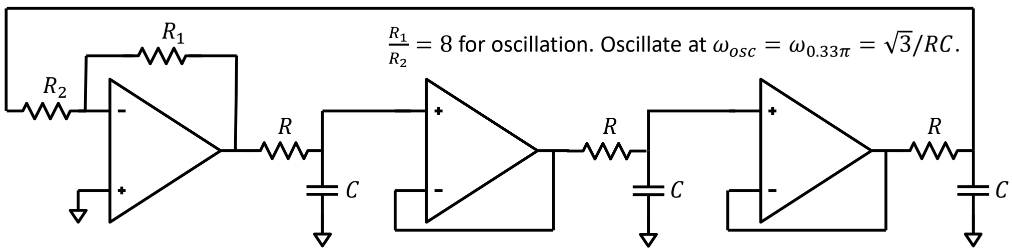

Department of Engineering

LC Oscillators

Matthew Spencer
Harvey Mudd College
E151 – Analog Circuit Design

7

In this video we're going to study LC oscillators, which we can also analyze using loop gains. Like phase shift oscillators, these circuits implement feedback around linear amplifiers to create oscillators. However, using second order feedback networks with both inductors and capacitors lets you make much more compact oscillators, which we'll see soon.

HARVEY MUDD COLLEGE Department of Engineering

LC Networks Have Useful Gain and Phase

Tapped Capacitor Resonator

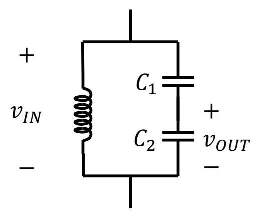

$$Z_{IN} = Ls||\frac{1}{C_{eff}s} = \frac{Ls}{LC_{eff}s^2 + 1}$$

Where $C_{eff} = C_1C_2/(C_1 + C_2)$

$$\frac{v_{OUT}}{v_{IN}} = \frac{C_1}{C_1 + C_2}$$

Tapped Inductor Resonator

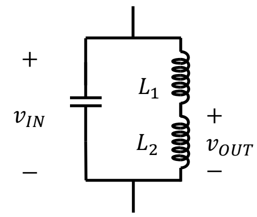

$$Z_{IN} = L_{eff}s||\frac{1}{Cs} = \frac{L_{eff}s}{L_{eff}Cs^2 + 1}$$

Where $L_{eff} = L_1 + L_2$

$$\frac{v_{OUT}}{v_{IN}} = \frac{L_2}{L_2 + L_1}$$

8

We're going to start analyzing this by looking at some interesting LC networks. The network pictured on this slide is called a tapped capacitor network.

CLICK The input impedance of this network is given by the parallel combination of the inductor impedance and the two series capacitor impedances. Because we are assuming there is no resistance in this network, we get a funny looking second order system in the denominator of this expression. If you substituted jw into the denominator, you'd find that it's always real, such that the overall impedance is always imaginary. You'd also find that the phase shifts from positive to negative abruptly at omega=1/sqrt(LCeff) when the sign of the denominator changes. These funny behaviors are because we haven't included any resistors in our circuit.

Those behaviors are distracting, but they paint the broad strokes of how a parallel LC resonator works. At low frequencies it looks purely inductive, so it has a purely imaginary and positive impedance. At high frequencies it looks purely capacitive, so it has a purely imaginary and negative impedance. At the resonant frequency, the imaginary parts of the parallel branches cancel out and you get a purely real impedance. Therefore, current and voltage are only in phase at resonance, and current leads or lags voltage at other frequencies. These behaviors hold even when there is resistance in the tank.

CLICK There’s also a gain in this network between vIN and vOUT. That ratio is given by a capacitive divider.

CLICK There’s a natural dual of these tapped capacitor networks called tapped inductor networks. These are particularly easy to build because you can just grab some wires in the middle of a wound inductor to make the vout port.

Both of these networks seem interesting in feedback because they have phase crossing zero at one frequency and a tunable gain. Maybe we can make them into oscillators!

… (Quartz crystal electromechanical models are similar to these.)

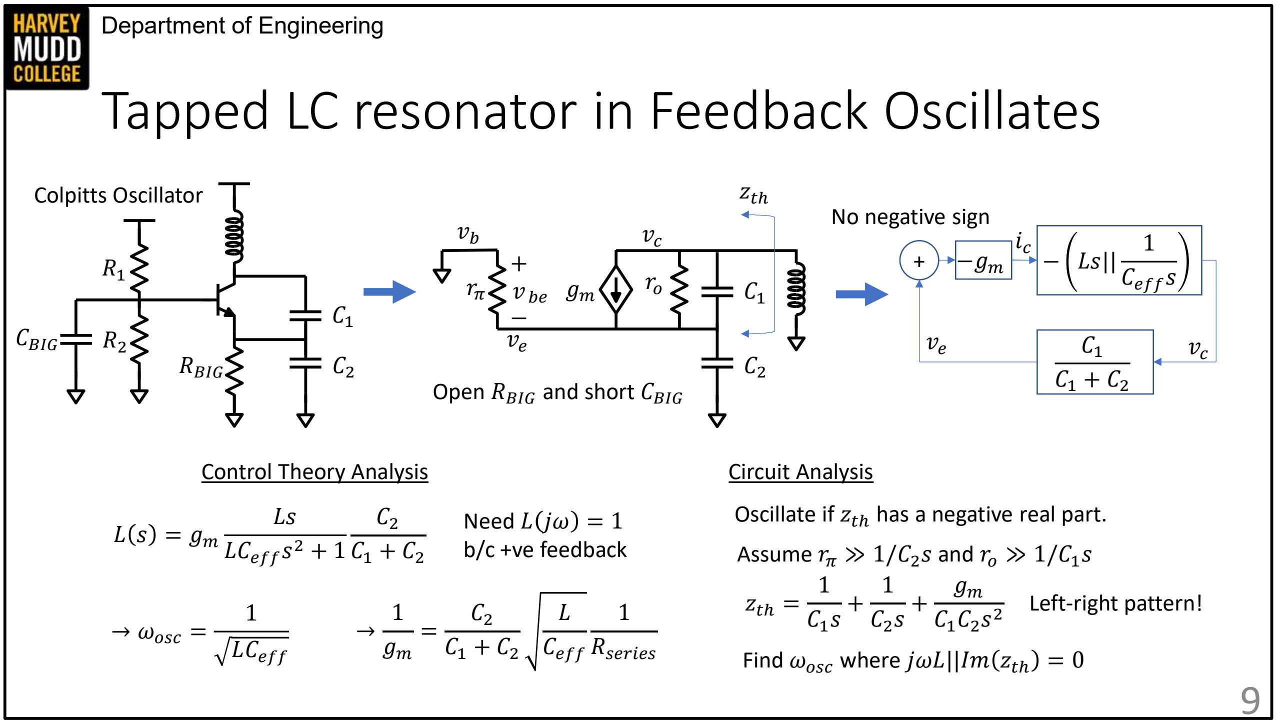

There are a lot of oscillators that use resonators in feedback, and we don't have time to analyze all of them. We're going to start with the Colpitts oscillator in common base configuration, which uses a tapped capacitor network and appears on the left. As the name of the circuit suggests, there are other variants of the Colpitts oscillator, most notably the emitter follower configuration, which is the kind of oscillator that student accidentally build in lab most often. There are also lots of other kinds of oscillators, like Hartley oscillators, which use tapped inductor feedback, Clapp oscillators, which add a few more tuning capacitors, and Pierce oscillators, which are often used with quartz crystal resonators. All of them work on similar principles to this Colpitts Oscillator.

CLICK We can make a small signal model for this oscillator, and we see that the inductor going to the power supply is connected to small signal ground, which makes this look much more like the tapped capacitor network we saw on the previous slide. We opened RBIG and shorted CBIG when we made this model, and we could pick values of RBIG and CBIG to make that simplification reasonably accurate. We'll also be assuming ro is much bigger than 1/C1s and rpi is much bigger than 1/C2s later in the analysis, which means this circuit kind of just looks like a gm generator driving into this resonator.

CLICK We can describe this circuit using a block diagram. Decreases in ve reduce gm current, and decreases in ic increase vc because less current gets pulled through the impedance of the resonator network. The voltage at the collector then sets the voltage at the emitter through the capacitive divider. The loop gain around this block diagram is net positive, so this is a system in positive feedback.

CLICK We can write our loop gain out using the ZIN of the resonator that we found on the last slide. We also note that we need L(jw) to be 1 in order to have poles on the jw axis in positive feedback.

CLICK This expression only has a phase of zero degrees exactly at the resonant frequency, so the oscillation frequency is 1/sqrt(LCeff). Calculating the magnitude of the loop gain at the oscillation frequency is tough because the denominator is zero. However, if there were a tiny bit of resistance, it would remain when the Lceff*s^2 term cancelled with 1. That means the gain will be equal to 1 if gm cancels out the attenuation of the tapped capacitor network and the ratio of sqrt(L/Ceff)/Rseries. That latter ratio – sqrt(L/Ceff)/Rseries – is interesting: it's a number called the quality factor of the resonator, and it's super important in oscillator design. Regardless, you can tune the biasing of the base and the capacitance ratio to achieve the magnitude constraint.

CLICK This control theory analysis answers the questions we care about: we know the frequency of oscillation and how much gain we need. However, there's another way to come to this result that is rooted in circuit analysis. It's worth seeing so that you can understand it when you see it in other references. The circuit analysis focuses on the Thevenin impedance seen from the inductor.

CLICK The oscillation condition for circuit analysis is that there is no resistance in the oscillator. An ideal inductor and cap will oscillator forever, and this analysis views the gm generator as a way to cancel out the real series impedance in the inductor to achieve that permanent oscillation. So if we find that zth has a negative real part that is bigger than the series resistance in the inductor (which is zero in this case), then we'll oscillate. We're going to assume that most of our current flows in the cap network and little flows in ro and rpi, which presumes we're at a reasonably high frequency.

CLICK When we do that, we find that zth can be calculated as a left-right pattern. C1 is in parallel with the gm generator, and the control voltage is in parallel with C2. The gm term of this impedance is going to be negative and real, so this feedback network has the potential to oscillate if the gm term is more negative than the real impedances it sees in the inductor and the BJT. This winds up to be the same condition as the gain criterion we found with control theory.

CLICK This technique can get us the oscillation frequency too. We need to have the imaginary parts of the inductance and zth cancel at resonance, so we look for the frequency were jwL || Im(zth) is zero.

As a final note, I want to emphasize how easy these are to make by accident. For instance, I've seen unintentional Colpitts oscillators where C2 is a bypass cap, C1 is breadboard capacitance, and L is a long wire to the power supply.

Summary

- Tapped LC resonators let us build oscillators because they provide attenuation and in-phase voltage and current at resonance.
- The Colpitts oscillator oscillates at $\omega_{osc} = 1/\sqrt{LC_{eff}}$.
- Analyze LC oscillators with either control theory or circuit analysis.
- It’s easy to make an LC oscillator by accident.

Department of Engineering

Describing Functions

Matthew Spencer
Harvey Mudd College
E151 – Analog Circuit Design

11

In this video we're going to look at a mathematical technique for analyzing non-linearities in our oscillators, which is called a describing function. This is just a glimpse of the world of non-linear control theory, which is full of really fun math and modeling. I'm picking this particular technique because it can give you some insight about amplitude of oscillation and about whether an oscillation is sustained.

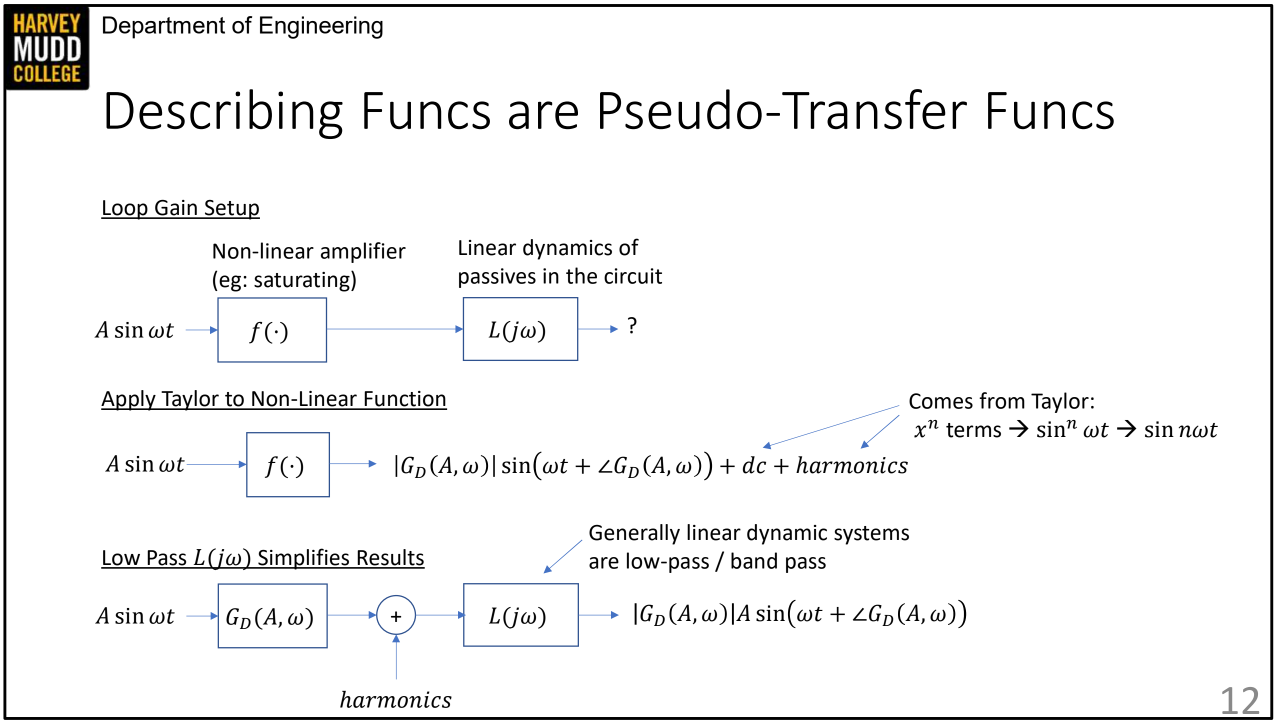

We need to do this describing function analysis because all Practical Oscillators are Non-Linear. A linear system with two poles on the jw axis would oscillate forever if you kicked it with an impulse, but a second impulse would create another oscillation that stacked on top of the first one, and there's no reason that it would start up of its own accord. We don't see that behavior in oscillators in lab, which seem to start up on their own and then self-limit their amplitude somehow.

That behavior can be traced back to the non-linearities in the amplification elements that we use. Any amplifier we build has its gain change when the output approaches the rails, and that looks like a non-linearity. If that non-linearity was just in feedback on its own, we'd be in trouble. Non-linear analysis has seemed hard from the very start of the class, and this would probably be no exception. However, our circuit has passive, linear elements that will filter the output of the non-linearity somehow. Maybe those elements will save our analysis.

CLICK With trepidation, we start our analysis by Taylor expanding our non-linearity using Taylor expansion. When we put a sine wave into a Taylor expansion, we find a sine wave at the output that is phase-shifted and scaled relative to the input wave, plus a DC component, plus harmonics of the sine wave, which are sine waves at integer multiples of the frequency omega. The frequency omega is called the fundamental frequency because all the harmonics depend on it. The DC and harmonic terms, and even the amplitude dependence of the fundamental component, come from the way that high order terms in the Taylor series exponentiate the input sine wave. Squaring sine results in a sinusoid at double the frequency and a DC offset, and similar things happen for higher order terms.

CLICK Finally, we assume that our loop dynamics are either low-pass or band-pass, which means we can mostly ignore the harmonics. All of the high frequency harmonics get filtered by the low pass behavior of L(jw), and the DC component is just an offset that we're OK to ignore in small signal analysis. The upshot of this is that we can plausibly claim that putting a sinusoid through a non-linearity and loop dynamics results in another scaled and phase-shifted sinusoid. The pseudo-transfer function that determines the scaling and phase-shifting, GD, is called a describing function, and the major difference between describing functions and transfer functions is that the describing function changes with the amplitude of the input sine wave. You can imagine that means the root locus of GD varies as the input amplitude changes, or you can imagine that this expression lets us change the gain of a system as the amplitude approaches the rail. Both of those capture the significance of having an amplitude dependence in our transfer function.

… Feed that back in and |GD| has to be self-consistent at all A values and frequencies
… Harmonic balance analysis, self-consistency in harmonics

HARVEY MUDD COLLEGE
Department of Engineering

Example of Finding a Describing Function

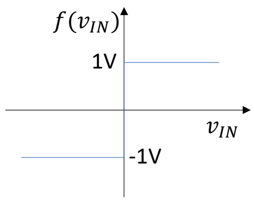

Input is $A \sin \omega t$. Output is $\pm 1\text{V}$ square wave w/ frequency ω: $\text{sgn}(A \sin \omega t)$

$$\text{sgn}(A \sin \omega t) \approx \frac{4}{\pi}\left(\sin \omega t + \frac{1}{3}\sin 3\omega t + \frac{1}{5}\sin 5\omega t + \cdots\right) \quad \text{by Fourier}$$

Low passing the harmonics and equating $G_D A$ w/ fundamental gain $\rightarrow G_D$

$$LPF(\text{sgn}(A \sin \omega t)) = \frac{4}{\pi}\sin \omega t = |G_D(A,\omega)| A \sin\left(\omega t + \angle G_D(A,\omega)\right)$$

$$G_D = \frac{4}{\pi A}$$ Describing function gain gets smaller as input amplitude gets bigger!

13

Lets look at an example of finding a describing function. We'll use this step non-linearity to represent a super high gain amplifier.

CLICK Per our setup on the last page, we drive this non-linearity with a sinusoid Asin(wt). When we do that, the output is a +/- 1V square wave because this non-linearity is +1 when the input sinusoid is positive and -1 when the input sinusoid is negative.

CLICK We can express the output square wave as a sum of sine waves using Fourier analysis. Note that because this function is odd and centered at 0V DC, we don't see a DC term or any cosines in this sum.

CLICK When we low-pass that expression, we find that only the fundamental term gets through. However, we know from the analysis on the previous slide that the coefficient of the fundamental has to be equal to the magnitude of the describing function times A. This non-linearity doesn't seem to introduce any phase changes relative to the sine wave, so we can safely sat the the angle of this describing function is zero.

CLICK So that means our describing function is 4/pi*A, which we got just by rearranging the coefficients of the previous line and cancelling out the sinusoids. This describing function does what we promised because it gets smaller as the input amplitude increases. That makes sense, no matter how big a sinusoid we put into this non-linearity, our output square wave is +/-1V. So our non-linearity scales up 1mV input sine waves, but scales down 10V input sine waves.

Department of Engineering

Describing Functions in Feedback

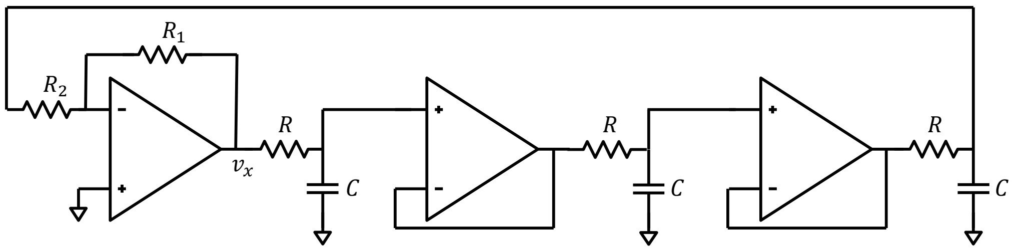

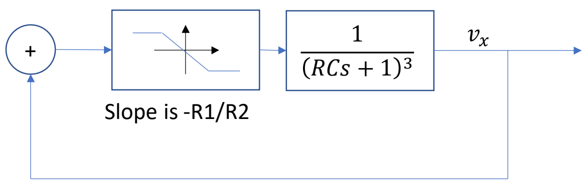

Slope is -R1/R2

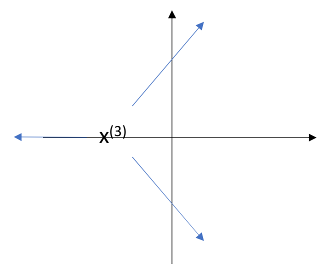

OK, let's see what this technique tells us about an oscillator in practice. We're looking at the phase shift oscillator from a few videos ago. We can rewrite the block diagram to show that the first stage is a saturating non-linearity: the linear gain stops working when we hit the rails, which flattens out the gain. The slope in the middle is the linear gain of the amplifier, -R1/R2. The remaining loop dynamics are a third order low-pass filter from three co-located poles. The other op-amps could introduce non-linearities when they rail out, but their gain is smaller than the first one, so it will saturate before the second or third one.

This is a new non-linearity, so it's going to have a new describing function. However, we can assume that describing function will get smaller as gain gets bigger just like the one on the previous page, and for the same reasons. No matter how big a sine wave we put in, our output is clipped to the rails, so the describing function has to decrease as A gets bigger.

That's interesting because the describing function looks like the gain in a root locus analysis. So the three poles of the loop gain will move out along the loop gain when the describing function is big, and shrink in toward the original location as the describing function gets small. This tells us about the startup behavior of the oscillator. When the first tiny perturbation starts a little sine wave, the describing function is small, so our poles move over onto the right half-plane. This causes the amplitude of the sine perturbation to get bigger until it starts compressing. Then the describing function shrinks until the poles rest on the j omega axis. This also tells us that we are safe using a gain bigger than eight in our feedback loop because the describing function will compress our gain if we overshoot the oscillation criterion. We can also guess that we'll get more harmonics as we increase our gain, which will shorten the linear range of our non-linearity. So even qualitatively, we can learn a bunch of things from the basics of Describing functions.

If slope is steep, then output is still basically a square wave and $G_D \approx 4/\pi A$

G_D is controlling gain for root locus plot.
Small A leads to RHP poles that increase A
Big A leads to reduced G_D keeping poles on $j\omega$ axis.

… gives you an idea of whether you'll oscillate and the amplitude of oscillation.

Department of Engineering

Summary

- Describing functions are transfer functions for non-linear systems that rely on low pass dynamics rejecting a harmonic.
- Describing functions change with amplitude, so their poles are functions of the amplitude of oscillation.
- Saturating non-linearities force poles back to the $j\omega$ axis in response to amplitude perturbations.

Department of Engineering

Relaxation Oscillators and Digital Oscillators

Matthew Spencer
Harvey Mudd College
E151 – Analog Circuit Design

16

In this video we're going to talk about a class of oscillators that really leans into non-linearity. All the oscillators we've talked about so far have relied on a linear system (or at least a system that's linear for part of its operation) to create a feedback loop with a loop gain of -1. Relaxation oscillators rely on digital systems, which are highly non-linear by design, to create oscillations. That's daunting, but digital systems are well-enough behaved that we can analyze these circuits.

Department of Engineering

Schmitt Trigger Oscillator

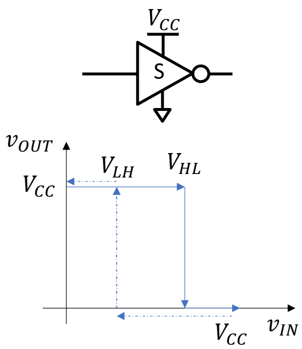

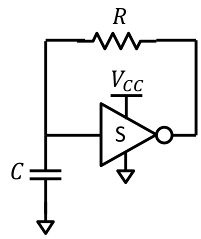

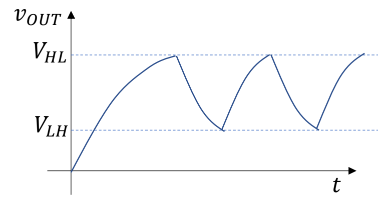

$$v_{RISE}(t) = V_{LH} + (V_{CC} - V_{LH})\left(1 - e^{-t/RC}\right)$$

$$v_{FALL}(t) = V_{HL}e^{-t/RC}$$

$$f_{osc} = \frac{1}{t_{rise} + t_{fall}}$$

17

The first oscillator we'll look is called a Schmitt trigger oscillator, which is a part of a family of oscillators called relaxation oscillators. Relaxation oscillators rely on charging and discharging an RC network to set the frequency.

Unsurprisingly, Schmitt trigger oscillators depend on a device called a Schmitt trigger, which is a digital gate with a hysteretic transfer function. Hysteresis is a type of memory that means your transfer function depends on past inputs. This sounds complicated, but it's not so bad if we look at an example. I'll fill out the Schmitt transfer function on the axes below.

CLICK When our input voltage is going from 0 to VCC, we see that our output looks like an inverter transfer function. One weird detail is that the inverter doesn't transition downward at VCC/2, but rather at a voltage VHL, which stands for the high-to-low voltage.

CLICK As you may have guessed from the arrows and VHL, the behavior of the Shmitt trigger is different when we decrease vIN from VCC to zero. This also looks like an inverter transfer function, but the output rises at a different voltage, VLH, which stands for the low-to-high voltage.

CLICK We can make an oscillator out of this pretty easily by attaching a simple RC network around the inverter. When the inverter output is high, the input will charge toward VHL.

Then when the input hits VHL, the output will switch low and the RC network will discharge towards VLH.

CLICK Each of these charge or discharge segments is governed by a standard first order differential equation, so we can quickly write the rising and falling voltage. Setting those expressions equal to VHL or VLH will give us the rise and fall times of the input waveform. We find the frequency of oscillation as one over the sum of the rise and fall time. Interestingly, the duty cycle isn't necessarily 50% in this circuit. You can find the duty cycle as the ratio of trise over tfall.

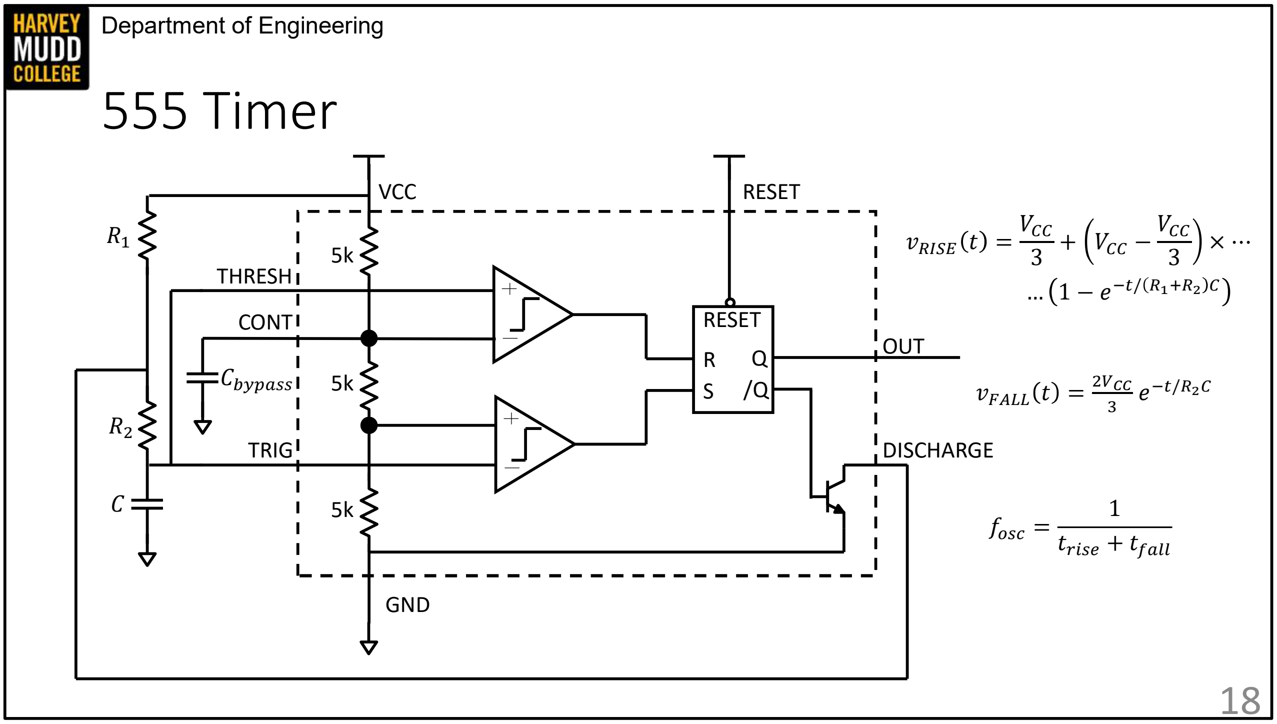

The most famous way of implementing relaxation oscillators is an integrated circuit from the 1970s called a 555 timer. These are remarkably commercially successful chips, they were released at the same time as the Intel 4004 and they're still on the market! They've become a bit of an obsession in some corners of the electrical engineering community, and you can find recipes to get them to do all sorts of tricks. We'll be focusing on making them into a Schmitt-like relaxation oscillator.

Before we hook up the relaxation network, let's look at the internals of the chip. The 555 timer gets its name from a trio of 5k resistors that provide some reference voltages to these comparators. That means the upper comparator puts out a high voltage when THRESH exceeds 2/3 of VCC and the lower comparator puts out a digital high voltage until TRIG exceeds 1/3 VCC. These thresholds can be configured by bypassing the top and middle 5k resistors using the CONT pin. The two comparators drive a SR latch, which sets its Q output high when S is triggered, and which sets its Q output low when R is triggered. /Q is the logical complement of Q, so it's low when Q is high and high when Q is low. /Q controls this handy BJT attached to the discharge pin. This type of BJT is called an open collector BJT, and it's a common type of interface to chips.

CLICK We short the THRESH AND TRIG pins to serve as the input to our relaxation oscillator, and we hook that input up to our state capacitor. The capacitor is pulled up to VCC through

R1 and R2, but the discharge pin is attached between those two resistors in a kind of feedback. The oscillation comes from charging and discharging this cap just like a Schmitt trigger oscillator. When Q is high, /Q will be low and discharge will be an open circuit. That means the capacitor will charge towards VCC through R1 and R2. When the capacitor voltage reaches 2VCC/3, a R signal will be sent to the SR latch forcing Q low. This will turn on the discharge BJT so that C will discharge through R2. When the voltage on C reaches VCC/3, a S signal is sent to the SR latch to set Q high again.

CLICK The rise time and fall time are first order RC behaviors again, but they have different time constants because the cap is pulled up through both R1 and R2, but only pulled down through R2. That gives rise to these equations. You can find the rising and falling time of the state capacitor by setting these equations equal to 2VCC/3 or VCC/3, and you might notice that means VCC will cancel from all of the terms in the equations. That means the charge time is independent of VCC, which is good for the frequency stability of the chip. The frequency of oscilation 1/(trise+tfall) again.

… Cbypass stabilizes CONT

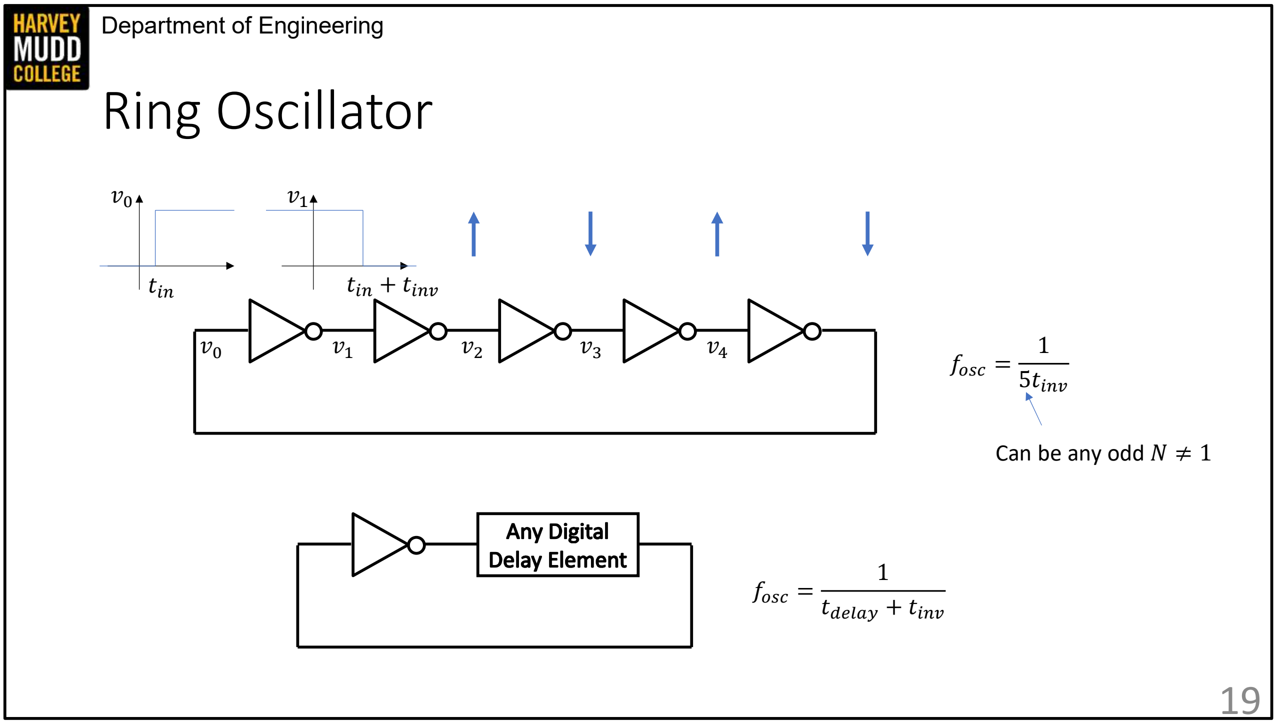

The last type of oscillator we'll talk about is a ring oscillator. Ring oscillators are hugely important calibration tools in digital chips, and they also get used in some clock generation circuits.

Understanding a ring oscillator depends on understanding inverters, but we can get away with the most basic model of an inverter, which suggests is that an inverter changes its output to the opposite of its input after a small delay. We can see this illustrated for the left-most inverter. We se a rising edge at time tin on node 0, and that results in a falling edge at time tin+tinv at node 1. tinv is referred to as the inverter delay.

CLICK This falling edge on node 1 will cause a rising edge on node 2 tinv later,
CLICK Which in turn will cause a falling edge on node 3 tinv after that,
CLICK Which causes a rising edge on node 4
CLICK And then a falling edge on node 0, which looks like a contradiction. After all, a rising edge on node zero started all these transitions. However, the falling transition on node 0 happens 5*tinv after the rising edge that started this transition, so it's fine for the v0 waveform to have a falling edge at this point. That falling edge reverses all of the transitions in the circuit, resulting in a transition from low to high at node 0 another 5 tinve later. This continues forever because an odd number of inversions will never reach a stable equilibrium.

CLICK The frequency of this oscillation is 5*tinv, but you can make ring oscillators of different frequencies by changing the number of inverters in the ring. Any odd number except for one will do.

CLICK In fact, you can generalize that concept to put other interesting logic gates into a ring oscillator loop. You could argue that all that the right four inverters are doing is sending a delayed copy of the left inverter's output back to its input, and you could accomplish the same goal with AND gates or OR gates that let you set the state of your oscillator. That's fun, but you might wonder what's so special about digital delays, could we use an analog delay instead? The answer is yes, though the math gets harder: you'd find a describing function for the inverter and calculate the loop gain for your delay.

… similarity to Schmitt trigger oscillator

Summary

- Relaxation oscillators digitally switch to charge/discharge a capacitor in feedback once the output voltage passes a threshold.
  - In Schmitt Trigger oscillators the threshold is set by the hysteresis of the Shmitt Trigger.
  - In 555 Timers the threshold is set by an internal 5k-5k-5k divider providing reference voltages to comparators.

$$f_{osc} = \frac{1}{t_{rise} + t_{fall}}$$

- Ring oscillators rely on an odd number of inversions in a loop to force digital gates to switch states.

$$f_{osc} = \frac{1}{N t_{inv}}$$