

Diodes in Circuits

Matthew Spencer

Harvey Mudd College

E151 – Analog Circuit Design

1

In this video series we're going to be putting diodes into circuits, and we're going to find out that it's hard to analyze circuits with non-linear elements in them. So we're going to come up with some workarounds. One of these workarounds is called small signal modeling and it's crucially important for the rest of the semester, so you've got something to look forward to!

The Resistor-Diode Circuit

Matthew Spencer


Harvey Mudd College

E151 – Analog Circuit Design

2

In this video we're going to combine a resistor and a diode in a circuit. I'm sure it's going to work great!

Resistor-Diode Circuits are Transcendental



$$i_D = I_s \left(\exp \frac{v_D}{n\phi_{th}} - 1 \right)$$

$$V - v_D = i_D R$$

$$V = v_D + R I_s \left(\exp \frac{v_D}{n\phi_{th}} - 1 \right)$$

Transcendental!

3

OK, I've made the simplest circuit including a diode that I can think of. It's just a voltage source, a resistor and a diode. I'd like you to take a minute and solve this circuit by finding i_D and v_D , which will definitely go smoothly and not reveal any complications.

CLICK OK, I'll show my solution. We start by noting that the diode voltage is just a resistor drop below the voltage source

CLICK Then we substitute the diode equation in for i_D and rearrange for a bit,

CLICK Then we throw our hands in the air because $x + e^x$ is a transcendental equation with no algebraic solution.

This means we're potentially in big trouble: we can't even solve the simplest non-linear circuit that we can build. If we can't solve the simplest one, then what are we going to do?

Take a minute to solve these very straightforward equations ...

How to Solve the Resistor-Diode Circuit

- Computers ← Don't go straight here because we want design insight
- Graphical Methods
- Approximations
- Linearization

4

I'm going to reflect that question back at you: I'd like you to write down some ways to solve the resistor diode circuit. I came up with four broad categories.

CLICK The first is using computers. Computers love doing the repetitive math to solve transcendental equations. So great, let's let them do that. We even have fancy simulators that make it simple to put circuit components into a computer.

CLICK However, we don't actually want computers to be our first choice. Though they are great at solving circuits, our value added as circuit designers is understanding what's going on. Fortunately for us, doing brute force searches over the circuit design space is implausible, so the insight that circuit designers bring is crucial. Never forget that your value as a circuit designer is understanding what's going on, and if you're just tweaking a simulation to get something to work, then you've left your value compared to a computer by the wayside.

CLICK Graphical solution methods, on the other hand, are great for humans to pick up intuition. They won't give us accurate answers, but drawing plots showing how your elements interact can lend considerable design insight.

CLICK We can also make simple approximations of the diode curve that make math much

easier.


CLICK Finally, we're going to learn about a really cool technique that lets us examine small deviations of the voltages and currents about known values of a circuit. This technique is called small signal modeling, and it relies on linearizing circuit elements using a Taylor series. We'll reuse it over and over again this semester.

Summary

- The resistor diode circuit is transcendental.
- We will learn about four alternative solution methods
 - Computers – which we're doing last
 - Graphical Methods
 - Approximations
 - Linearization

In this video we're going to draw some IV curves that capture the behavior of multiple circuit elements, allowing us to derive some design insight about their behavior.


IV Relationship of a Resistor Imposes Constraint on Diode IV



$$V - v_D = i_D R$$

Load line
$$i_D = \frac{V}{R} - \frac{1}{R} v_D$$

Device line
$$i_D = I_s \left(\exp \frac{v_D}{n\phi_{th}} - 1 \right)$$



We start by writing out the equations governing the operation of this circuit. The only interesting deviation from how we wrote them last time is that I've reformulated the equation showing that v_D is related to V by the voltage drop across a resistor. However, if we look close, this second equation gives us i_D as a function of v_D , just like the bottom equation for diode behavior. That means we could draw this constraint, which is imposed by the resistor, on the same axis as our diode curve.

I've done just that on the right.

CLICK The exponential curve describing the diode's behavior is sometimes called the device line.

CLICK and this line representing the resistor's behavior is called the load line. It's called that because resistors in amplifier circuits are often called loads. You can see that it starts at V/R on the y-axis, which would be true if v_D were zero, implying that the diode was just a short. It ends at V on the x-axis, which would imply no current flowed through R , or put another way, that the diode was an open.

CLICK Both of these lines need to be true when the circuit is operating: the equations for the diode and the resistor both have to be satisfied. That means the intersection of the two lines tells us the operating point of the circuit, giving us the correct values of v_D and i_D . The operating point of a circuit is often called the bias point.


Though I wouldn't trust my freehand sketch for calculating a voltage, I can often get within a few hundred millivolts if I use graph paper. Even better, this plot tells us what will happen if we change circuit parameters. For example, you could picture that a bigger value of R will make the load line flatter, moving the bias point down the exponential curve. This is a powerful technique for guessing at the sensitivity of your system to changes in your circuit.

Approximation with a Switch-Voltage-Source Model


Matthew Spencer
Harvey Mudd College
E151 – Analog Circuit Design

In this video we are going to devise a very simple approximation to the diode curve that lets us analyze circuits with decent accuracy.

Exponentials are Basically Vertical Lines



~0.7V, Increases by ~60mV/decade



Material	Von [V]	LED Color
Silicon	0.7	IR
InGaAs / InGaP	1.6	Red
Nd:YAg / InGaN	2.3	Green
Silicon Carbide	3.1	Blue

That approximation relies on the observation that exponentials change very quickly. Once an exponential gets going, it grows faster than any other relevant functions.

CLICK As a result, we can basically say it's a straight line. If we draw an IV curve that is dead flat at zero current, and then dead vertical, providing infinite current, at some voltage, we're pretty close to exponential behavior over a decade or so of current. So as long as you don't expect your current to swing around by more than a factor of 10-100, this approximation will give you a decent answer.

CLICK Of course, the choice of where to switch from horizontal to vertical is really important for this model. We name this parameter V_{ON} , and it's worth noting that V_{ON} increases by about 60 mV for every decade change in current. That's not a lot of voltage, 60 mV is only 3 times bigger than the noise limits of our oscilloscopes, and a decade change is a big change in current. So it's easy to pick a V_{ON} value that is valid over the range of currents that are usually used in PCB design. For the case of silicon diodes, that V_{ON} value is 0.7V for currents that are about 1A, and 0.6V for currents around 1mA. We usually use 0.7V as the assumed value in this class.

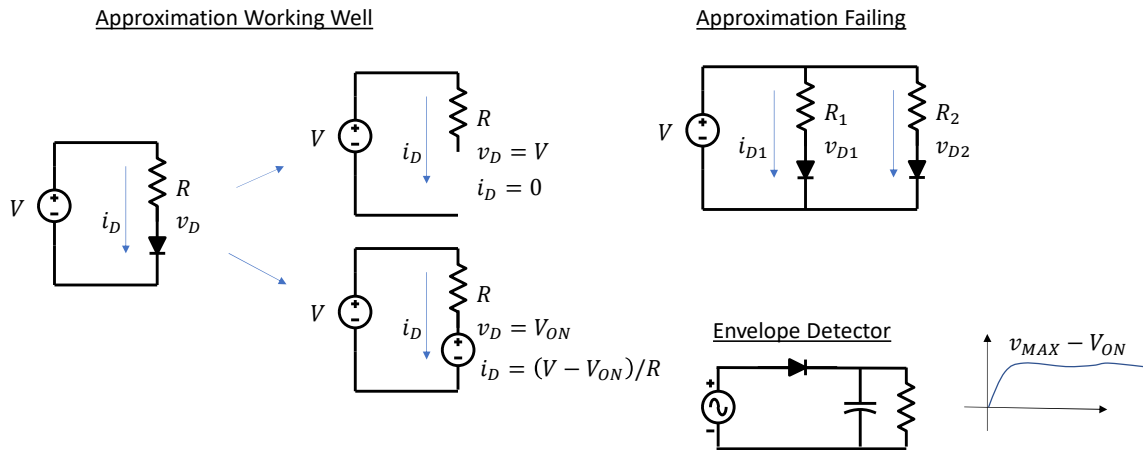
CLICK We represent this approximation in a circuit as a switch in series with a voltage source, which gives the model the creative name of "the switch-voltage source model".

The switch closes if the voltage across the diode is greater than V_{ON} , at which point the diode acts like a small voltage source. Voltage sources provide the same voltage regardless of how much current is pulled out of them, so the voltage source is responsible for the vertical line on the IV curve.

CLICK One aside that I love mentioning is that the value of V_{ON} depends on the material used to make diodes. Silicon diodes, which emit IR light if they are used as LEDs have an on-voltage of 0.7V. Indium-gallium-arsenide-phosphide alloys, which are used for red LEDs, have 1.6 V on voltages. Green LEDs are made of Neodymium doped yttrium-aluminum garnets or indium-gallium-nitride alloys, and they have a turn-on voltage of 2.3V. Finally, blue LEDs are often made of silicon carbide, and they have 3.1 V turn-on voltages. There's a pattern here where higher energy light seems to correspond to higher on voltages in diodes. That pattern arises from an energy barrier designed into the materials. Semiconductors with larger energy barriers release higher energy photons when carriers recombine, but the larger energy barrier results in a higher turn-on voltage.

While we're talking about energy barriers, it's important to note that V_{ON} is not that same as the built-in voltage in diodes, and trying to use VBI to find V_{ON} is a mistake. Don't get the two confused.

Powerful Model with Some Limitations



We can see two examples of the switch-voltage-source model in action on this page. On the left, the model works well. We check if the diode is conducting by comparing v_D to V_{ON} , and if we don't know v_D then we just solve the circuit both ways and pick the one that's correct (which is called the method of assumed states). The upper circuit shows the diode in its non-conducting state, and if no current is running through the diode then i_D is zero and v_D is V . The bottom circuit assumes the diode is conducting, so it's represented by a voltage source of V_{ON} that pins v_D to V_{ON} . i_D is calculated using the voltage difference across the resistor, and it is found to be $(V - V_{ON})/R$.

The right side of the slide shows a situation where this approximation works poorly. If you care about the difference between two diodes, or slight changes in the value of v_D , then this approximation won't tell you anything interesting. For this circuit, it would predict v_{D1} and v_{D2} are both about $0.7V$, which is true, but not helpful if you care about the difference between v_{D1} and v_{D2} . So be careful of using this model if you care about small differences between multiple diodes in a circuit.

However, this model works well for lots of situations, including the envelope detector in the lower right. The diode in the envelope detector allows the capacitor to be charged when the input sine wave is increasing, but prevents it from discharging when the input sine is decreasing. This means the voltage across the capacitor is almost constant, except

for a small discharge through the resistor in parallel with the cap. This discharge is corrected the next time the input sine wave reaches a peak, so it's often said that envelope detectors exhibit output ripple. The resistor might be put there deliberately to allow you to detect decreases in the amplitude of the input sine wave, or it might represent a load circuit.

Summary

- You can approximate exponentials with a piecewise-linear model of no-current, then infinite current: a switch + a voltage source
- You switch between no-current and infinite current at V_{ON}
- V_{ON} depends on exact current level but is $\approx 0.7V$ for silicon diodes
- This model fails when comparing diode voltages to one another

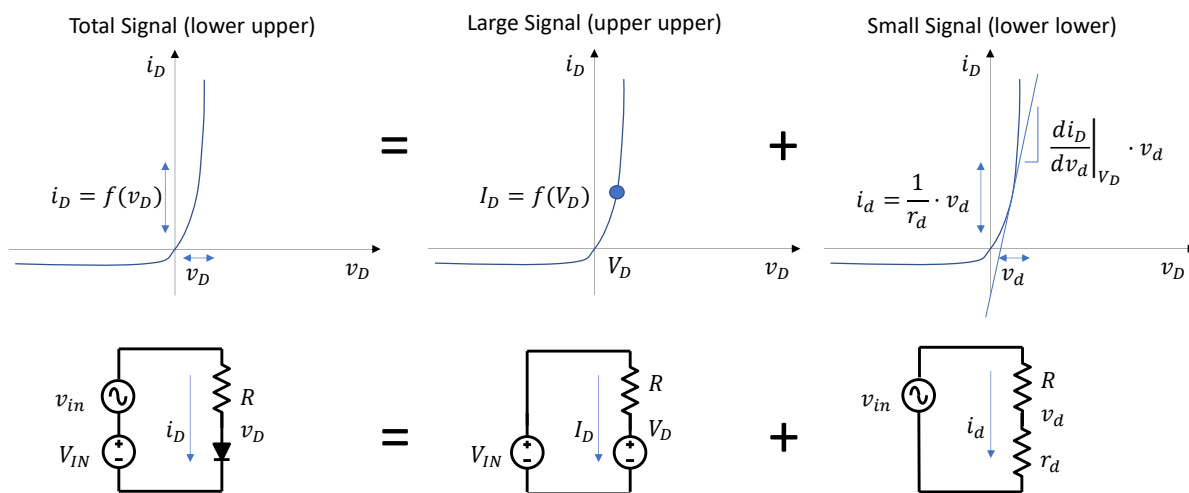
Linearization with a Small Signal Model

Matthew Spencer
Harvey Mudd College
E151 – Analog Circuit Design

13

In this video we're going to devise a way to analyze small deviations from the bias point that is called small signal modelling. We're going to use small signal models a lot this semester, so stay sharp!

Linearized Ckt Lets Us Analyze Small Changes



We start by assuming that our resistor diode circuit has an input that consists of a big fixed component plus a small variation. We call the combination of these two signals the total signal, we call the large fixed component the large signal, and we call the small variation the small signal. There's a notation convention that is associated with this type of modeling where total signal variables are denoted with a lower case letter and an upper case subscript, Large signal variables are denoted as upper case letters with upper case subscripts, and small signal variables are denoted with lower case letters and lower case subscripts. You can see this labeling in the circuit below the exponential curve. The big DC source is labeled upper case V, upper case IN, while the small varying source is labeled lower case v, lower case in.

It's common to draw small signals as AC sources, but the math we're going to pursue in a little bit doesn't require the small signal to be varying. You can have small DC variations that are effectively modeled by small signal modeling. I mention this because we're going to use some models later that do require sources to be time varying, and it's important to keep track of which assumptions are associated with which models.

A picture of the diode's experience of this model is shown in the curve in the upper left. Our input looks like a small variation with a DC offset, and that will cause the current to move up and down as the changing diode voltage indicates different parts of the diodes IV

curve. That means i_D is some function of v_D . Notice that I've used total signal labels for each of these expressions because we're talking about total signals here.

CLICK The big trick that we're using with small signal modeling is combining superposition and Taylor series expansions. This picture only shows the Taylor series part of that trick. Here we say that the output total signal i_D is equal to the sum of a large signal I_D , which is the first term of a Taylor series. That first term means that I_D is equal to the diode's IV function evaluated at the large signal voltage. This large signal I_D and V_D describe the bias point of the circuit.

This large signal representation is added to a linear approximation of the variation from the bias point, which is the small signal model. We find the small signal variation by calculating the slope of the IV curve at the bias point, and we multiply that by the small signal deviation to find the size of the current deviation. This is just like the second term of a Taylor series.

There are a few things to notice about this model. First, this model only deals with variations from the bias point, small signal models only tell you about small deviations, not about absolute values. Foreshadowing a bit, that means we're going to shut off sources in our small signal model circuits because sources never have small deviations in their value. Second, we calculated the slope by taking the derivative of the diode's IV curve, which means we may be able to come up with a tidy expression for di_D/dv_D by substituting the diode equation in for i_D . We'll do that on the next slide. Third, di_D/dv_D is a conductance because it has units of current over voltage, and we are mostly used to dealing with resistances in circuits. So we define di_D/dv_D as one over a resistance, r_d , which we call the differential resistance. Put another way, the differential resistance is the inverse of the slope at every point on the diode curve.

CLICK This mathematical representation can turn into a circuit representation, which using superposition to add a large signal circuit that captures our bias point to a small signal circuit that captures small deviations from the bias point. Our large signal circuit ignores the small signal input source because the small signal source is much smaller than V_{IN} . We often solve these large signal circuits using switch-voltage-source models. The small signal circuit zeroes out the input voltage source because the value of V_{IN} won't ever change. Put another way, there's zero small deviations from V_{IN} . The diode in the small signal circuit has been replaced by the differential resistance to capture the slope of the diode at the bias point.

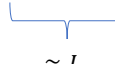
We could apply this model to figure out the variations in small signal v_d as a result of small signal v_{in} , and we'd see that they're just related by a voltage divider. If we had a value for r_d , this would be a very quick calculation.

Calculating Differential Resistance

$$\frac{1}{r_d} = \left. \frac{di_D}{dv_d} \right|_{V_D}$$

$$\frac{1}{r_d} = \frac{d}{dv_d} \left(I_s \left(\exp \frac{V_D}{n\phi_{th}} - 1 \right) \right)$$

$$\frac{1}{r_d} = \frac{1}{n\phi_{th}} I_s \exp \frac{V_D}{n\phi_{th}}$$



 $\approx I_D$

$$r_d = \frac{n\phi_{th}}{I_D}$$



15

There's one last step to completing our small signal model, which is finding the value of r_d . We've defined $1/r_d$ as the derivative of the diode's I-V curve at the bias point.

CLICK We can substitute the diode equation in for i_D

CLICK Then evaluate the derivative, which eliminates the -1 term and brings a $1/n\phi_{th}$ down from the exponential

CLICK This would be a usable expression, but it's clunky and hard to calculate with in lab. You have to take an exponential, which probably requires a calculator. However, the back half of this expression is almost the same as I_D , differing only in the -1 term we lost in the derivative. So we're going to substitute I_D for all of these expressions.

CLICK And if we do that and simplify, then we wind up with a very simple little expression, that r_d is equal to n times ϕ_{th} over I_D

I want to call attention to this substitution, because we most small signal models have a similar step. When you're making small signal models, look for this opportunity to substitute your large signal current in place of your large signal voltage. One classic test question is to give you a weird IV curve and ask you to make a small signal model for it, and this substitution is a place where many students have stumbled.

Finally, the art on the right side is a bit superfluous. It's just a visual reminder that we are approximating the diode's behavior as a combination of a large signal behavior, the V_{ON}

voltage source, plus a small signal behavior, the differential resistance.

Summary


- You can use superposition and Taylor Series to make linearized circuits that capture small deviations from the bias point.
- The circuit representing the first Taylor term is called a small signal model, and you find it by taking derivatives.
- The small signal model for a diode is a resistor called the differential resistance:

$$r_d = \frac{n\phi_{th}}{I_D}$$

In this video we're going to circle back to solving non-linear circuits with computers, which we usually do using programs called circuit simulators. Circuit simulators work really well, and they are staples of industrial practice. You should assume that you always have one available when you are working on your problems or labs, they are like calculators for this class.

DC Simulations Calculate Bias Points

- Newton's Method is an iterative approximation – can implement 1D version on TI-83+, Matlab, Python, etc.
- SPICE does N-dimensional Newton's Method on non-linear array of equations representing branch currents.



Output

V1=5.3V
V2=3.4V
I1=100mA
...

Or a graph if you sweep a variable

Simulators are capable of multiple types of simulations, and it's important to know how to use each of them. DC simulations are the first and, in some ways, the most important of these. DC simulations are used to find the bias point of circuits. You could imagine multiple ways of doing this, like guessing and checking a lot of solutions or using brute force to try every possible solution and seeing what fits, but circuit simulators mostly find bias points using an algorithm called Newton's method. This technique finds the roots of non-linear systems by making successively better approximations. You calculate your next approximation using the slope of the non-linear curve at your current guess, specifically by tracing the slope down to where it intersects with zero. You can see that pictured on this slide.

There are lots of things to say about Newton's method. For instance, deciding how many iterations to pursue before you stop requires a bit of thought. However, the most important thing for you to know as designers is that DC simulations can fail, and that you should be ready to adjust your models or simulations to get around it. The most common source of failure are highly non-monotonic or discontinuous models, both of which wreak havoc on the derivatives that are required for Newton's Method to function.


You can build Newton's method in most programming languages pretty easily, and one of my favorite implementations involves storing a function to the variable X on a graphing

calculator, then hitting enter a bunch of times to rerun the function for a better approximation. However, SPICE has done a lot of that implementation work for you because SPICE implements an N-dimensional Newton's method solve for systems of non-linear equations for its DC simulation. This is a useful trick, and SPICE was designed with that general purpose usefulness in mind. In fact, it was originally called the simulation program with integrated circuit emphasis, suggesting its use in other domains.

The output of a DC simulation is either a list of currents and voltages in different nodes and branches, which is the output if you perform an operating point or .op simulation, or a curve that shows how a voltage/current varies with some parameter, which happens if you carry out a .dc analysis.

AC Simulations Find the Small Signal Model

- Calculate a derivative at bias point, so depends on DC convergence
- Plots transfer function vs. frequency creating a Bode plot
- Need to include a small signal source, value doesn't matter b/c linear
- Get comfortable with decibels (or change your axis)



19


AC simulations effectively calculate the small signal transfer model of your circuit. AC simulations assume your input is sinusoidal and small, and they then calculate the transfer function of an input sinusoid to other variables by taking a N-dimensional derivative of your circuit's behavior at the bias point. Because you need the bias point to do an AC simulation, you need to make sure your DC simulation works. This means it needs to converge, and it also needs to be at the right spot. You should carry out a DC simulation before you do an AC simulation to check your operating point.

The output of an AC simulation will look like a Bode plot of some variable in your circuit. If you conduct a sweep with your AC simulation, then you will see a family of curves on your Bode plot, where each curve corresponds to one value of the swept variable. Many programs default to reporting AC outputs in decibels, so refamiliarize yourself with how decibels work or change the configuration of your y axes.

You need to remember to put a small signal source into your simulation. You often have the option to specify the amplitude of your small signal source, but it doesn't really matter because the AC model is linearized. As a result, your input source amplitude will just be multiplied by the gain from the input to your circuit variable regardless of how big the input source amplitude is. I like using input sources of 1V so that the Bode plot of a variable just shows me the gain rather than some measured value I'd expect with a smaller input source.

Transient Sims Capture Full Nonlinearity

- Full nonlinear simulation vs. time, captures non-linear effects
- Approximates system as ODE over short time step
- Initial condition is bias point, then previous time step
- Computationally expensive,
- Not always good for design insight
- Time steps can be deceptive



Multiple curves if you sweep a variable

Transient simulations capture all the non-linear behavior of a circuit, which often makes students think they are the best simulations. However, transient simulations are computationally expensive and they often don't grant as much design insight as AC simulations, so you can throw a lot of good time after bad with transient sims. Transient simulations also depend on the DC operating point of your circuit as an initial condition for the transient, which is calculated by finding the evolution of a series of ordinary differential equations that approximate your system. That means that you should carry out a DC simulation before a transient to make sure your circuit is at the correct operating point. I also recommend carrying out an AC simulation before a transient to glean design insight. The transient makes a great confirmation step to show that your system works, but it should be the last simulation you carry out in this class.

Finally, beware of the time step setting on transient simulations. Making it too short will cause your simulation to run forever, but letting it get too large can hide oscillations or other behaviors in your circuit. The time step will usually automatically adapt in SPICE simulators, but you have some control over it, and you should use that control to keep an eye out for suspicious simulation errors.

