INTRO | mki x es.edu

Hey there, thanks for buying this DIY kit! We – **Erica Synths** and **Moritz Klein** – have developed it with one specific goal in mind: teaching people with little to no prior experience how to design analog synthesizer circuits from scratch. So what you'll find in the box is not simply meant to be soldered together and then disappear in your rack.
Instead, we want to take you through the circuit design process step by step, explaining every choice we've made and how it impacts the finished module. For that, we strongly suggest you follow along on a **breadboard[1],** which is a non-permanent circuit prototyping tool that allows you to experiment and play around with your components. To help you with this, we've included suggested breadboard layouts in select chapters.
In addition to this, you can also play around with most of the chapter's circuits in a **circuit simulator** called CircuitJS. CircuitJS runs in your browser. You'll find web-links in the footnotes which will take you to an instance that already has example circuits set up for you. We strongly encourage you to fiddle with the component values and general structure of those circuits to get a better understanding of the concepts we're laying out.
Generally, this manual is intended to be read and worked through front to back, but there were a few things we felt should go into a dedicated appendix. These are general vignettes on electronic components & concepts, tools, and the process of putting the module together once you're done experimenting. Don't hesitate to check in there whenever you think you're missing an important piece of information. Most importantly though: have fun!

TABLE OF CONTENTS

[1] Note that there's no breadboard included in this kit! You will also need a pack of jumper wires and two 9 V batteries with clips. These things are cheap & easy to find in your local electronics shop.

THE mki x es.edu VCA

The VCA is one of the fundamental building blocks of any modular synthesizer. No rack is really complete without it – which is why I came up with this simple circuit for my own modular system. It's a fairly basic, yet super efficient little VCA made of just a couple transistors and op amps.[2]

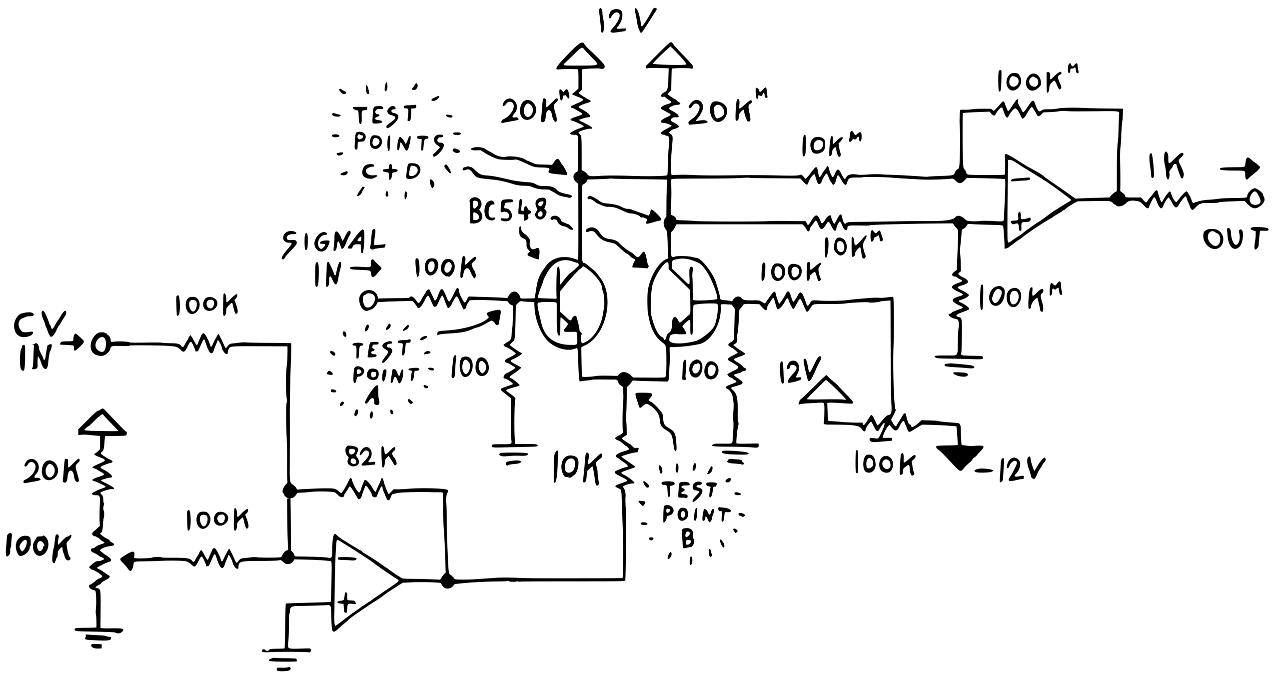

[2] Eagle-eyed readers might spot that some resistor values have changed compared to the version I showed off in my VCA video. Unfortunately, explaining that was out of scope for this (already super long) manual. But in effect, **those changes reduce the amount of noise on the output and improve the linearity of the VCA's CV response**.

BILL OF MATERIALS

Before we start, please check if your kit contains all of the necessary components. In the box, in addition to a PCB, panel and power cable, you should find:

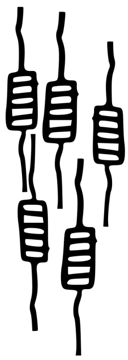

An array of resistors. The specific values (in ohms, which you should check for with a multimeter) are

100k	x 12
82k	x 2
20k	x 6
10k	x 6
1k	x 2
100	x 4
10	x 2

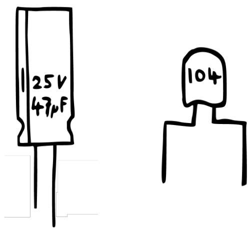

A few capacitors. The specific values (which are printed onto their bodies) are

47 µF (electrolytic)	x2
100 nF (104/ceramic)	x6

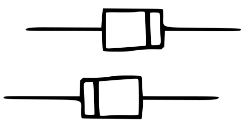

Some diodes. The specific model names (which are printed onto their bodies) are

SB140 (schottky) x2

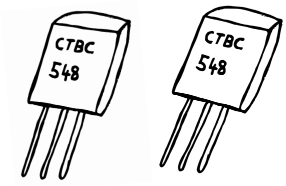

A couple transistors. The specific model names (which are printed onto their bodies) are

BC548 (NPN) x4

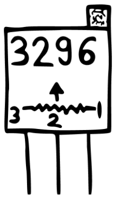

Some trimmer potentiometers. The specific values (which are encoded & printed on top) are

100k (W104) x2

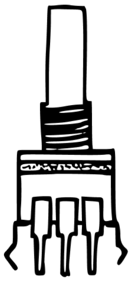

A few regular potentiometers. The specific values (which may be encoded & printed onto their bodies) are

100k (B104) x2

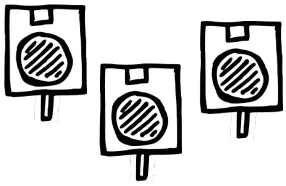

A bunch of jack sockets. The specific models (which you can identify by their color) are

Switched mono (black) x6

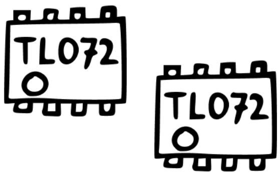

A couple chips. The specific models (which are printed onto their bodies) are

TL072 (dual op amp) x2

You will also find a few sockets that are only relevant when assembling the module in the end.

VOLUME CONTROL BASICS

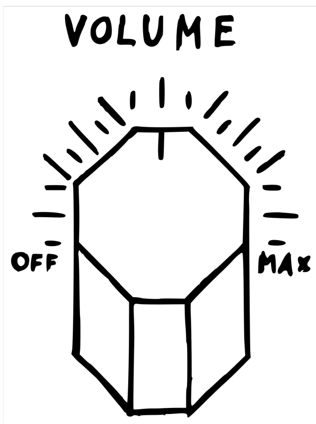

First off, let's answer the most essential question here: what does a VCA do, exactly? Well, the simplest way to visualize it is by looking at an ordinary volume knob. We probably all know how these work: turn it to the right to increase and to the left to decrease the volume. So far, so simple. But while they have their place even in a modular synth, they've usually got one major shortcoming: they're not automatable. You can only turn them by hand. In the modular world, this is a big no-no -– the whole point is being able to modulate any and all parameters using LFOs, envelopes and sequencers.

So what we're looking for is basically just a volume knob that can be controlled with a voltage. That's what a VCA is: a voltage controlled amplifier. Where a low level voltage corresponds to a quiet, while a high level voltage gives you a loud output signal. So if we control our VCA with an envelope, for example, we can give an oscillation a volume curve that mimics those of certain instruments. A slow attack sounds like a swelling violin. A super short attack, quick decay and long release sounds like a plucked guitar.

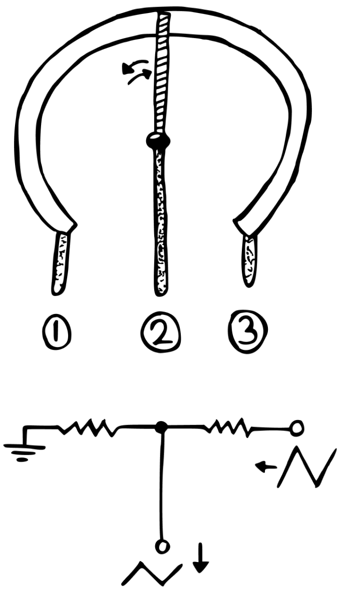

Now in order to be able to design such a VCA, we need to take a step back and talk about basic, boring volume knobs again. How do those work? Well, at their most basic, they're really just potentiometers set up as variable voltage dividers.[3] Because as you might know: in the world of electronics, sound is just a swinging voltage. And the bigger the swing, the louder the sound. Since voltage dividers do just what their name implies, they'll divide (or scale) down any voltage level you throw at them. So imagine we set our potentiometer to precisely 12 o'clock. This means that the two resistances inside the device are exactly equal, giving us a perfect 50% voltage divider that slashes our sound's volume in half. Turning it to the left will then gradually reduce that volume, until we get complete silence.

Why's that? **Because we're changing the relation between the two resistances that make up our voltage divider**. And that relation will determine the factor by which our voltages are divided down. To understand this on a more in-depth level, let's visualize it using the water analogy.

[3] Find out more about potentiometers in the components & concepts appendix (page 37).

VOLTAGE DIVIDER CLOSE-UP

In that analogy, we'd think of **electricity as water flowing through a system of pipes, with voltage being pressure and current being the flow amount per second**. Resistors would then be narrowings in the pipes, restricting the flow.[4]

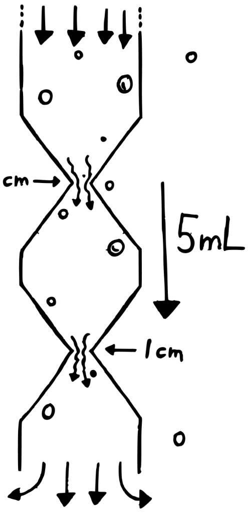

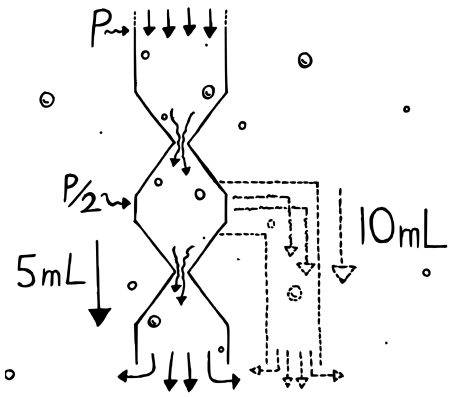

With this in mind, it should become clearer how our voltage divider works.[5] Because now imagine we apply a steady amount of pressure to the top here. Since our system is wide open at the bottom, we'll see water flowing downwards. How much? That directly depends on how narrow our pipes get, as this is quite literally the bottleneck in this scenario. To make our lives a little easier, let's quantify things. We'll say that both our narrowings are exactly equal in diameter: 1 cm. Individually, they'd let the same amount of water pass through in this setup: 10 ml per second. But placed in series, that amount would be slashed in half, giving us a flow of just 5 ml per second. **So it's fair to say that there is some unused potential here – the two narrowings are basically handicapping each other**.

I like to think of it like this. If we were to bypass the second narrowing, we'd suddenly see a flow of 10 ml. This is simply because we got rid of the additional bottleneck. Now the question is: where does this unused potential go when there's no bypass available? The answer is simple: nowhere. It just stays in the middle, between our two narrowings, as pressure. How much? **Well, since the full realized potential is a flow of 10 ml, but we see just half of that – the remaining pressure is exactly half the pressure we apply up top**.

Knowing this, we can derive what happens when we change the relation between these two narrowings. If we make the bottom one wider, we'll see two things happen: first, the total amount of flow will increase. Second, the amount of pressure in the middle will decrease. This is because we partly de-handicap the upper narrowing, which allows for

---

[4] Read more about resistors in the components & concepts appendix (page 33). Also, please note that the way I use water as a metaphor here might very well be nonsense, physically speaking. It's just a clarification crutch that has been quite effective for me.

[5] You can try this chapter's concepts in a circuit simulator. I've already set them up for you right here: https://tinyurl.com/y8dbcu2q – you can change all values by double clicking on components.

more potential to be used and converted into flow. On the flip-side, if we make the bottom narrowing tighter than the other, flow will de- and pressure in the middle will increase.

An important thing to keep in mind is that when we're just focusing on the pressure levels, **the absolute quantity of flow does not matter at all.** It really is only about relations and proportions. So regardless of wether we have a total flow of 5 ml or 5 l – as long as this is half of what could potentially flow with just one narrowing, we effectively halve the applied pressure.

Okay, but so far we've assumed the pressure at the top to be constant. An audio signal, as we said earlier, is a swinging voltage though, meaning that it changes over time. In our analogy, this would mean that our pressure is getting more or less intense – and sometimes even going negative, essentially turning into suction. Now obviously, a voltage divider works with audio signals. So it can handle big, small, negative voltages just fine, always dividing them by the same factor. But why is that? **Because resistors have two properties that make this possible: linearity and symmetry**. Linearity, in this context, means that there is a linear relation between the voltage applied to a resistor and the current that'll flow through it. So a doubling in voltage will always result in a doubling in current.

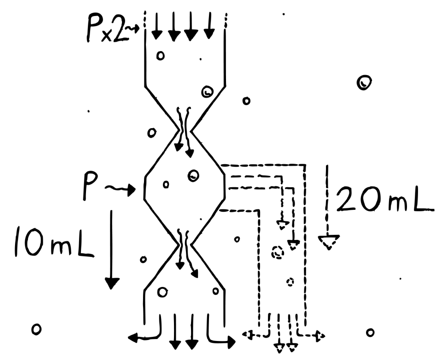

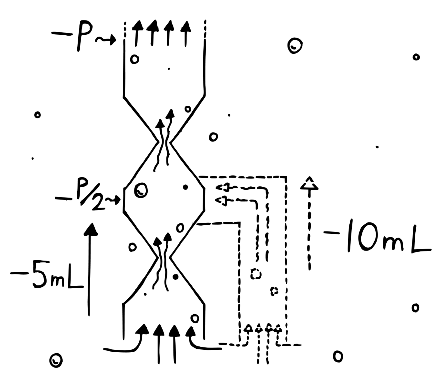

Same thing applies for our analogy: if we double the pressure, we get double the amount of water flowing through our narrowings. Since this is true not only for multiple narrowings in series, but also for each individual one in isolation, our potential/actual flow-model still works out the same way. Let's look at this in detail. For a doubling in pressure, we'd get a total flow of 10 ml instead of just 5. Add a bypass, though, and we'd see a whopping 20 ml flow through this narrowing. **Since 10 ml total is still half of these 20 ml that could potentially flow, we (again) get half the applied pressure between our two narrowings**.

Cool, but what about negative voltages? This is where the second property – symmetry – comes in handy. For resistors (and for our narrowings), the direction of flow does not matter. The restrictive effect will be the same in either direction. So if the pressure up top goes negative, we simply add a minus sign before all our values, indicating that the direction of flow has been inverted. **But since -5 ml is still half of -10 ml, everything still works out the same, and we get half the pressure in the middle.**

To try this in practice, grab your breadboard and set up a (very simple) circuit as shown below. It's just a 100k potentiometer[6] configured as a variable voltage divider, with an audio signal sent in via the right hand jack socket.[7] You should then connect a pair of headphones to the left socket. Since this is a passive circuit, you won't need a power supply or battery.

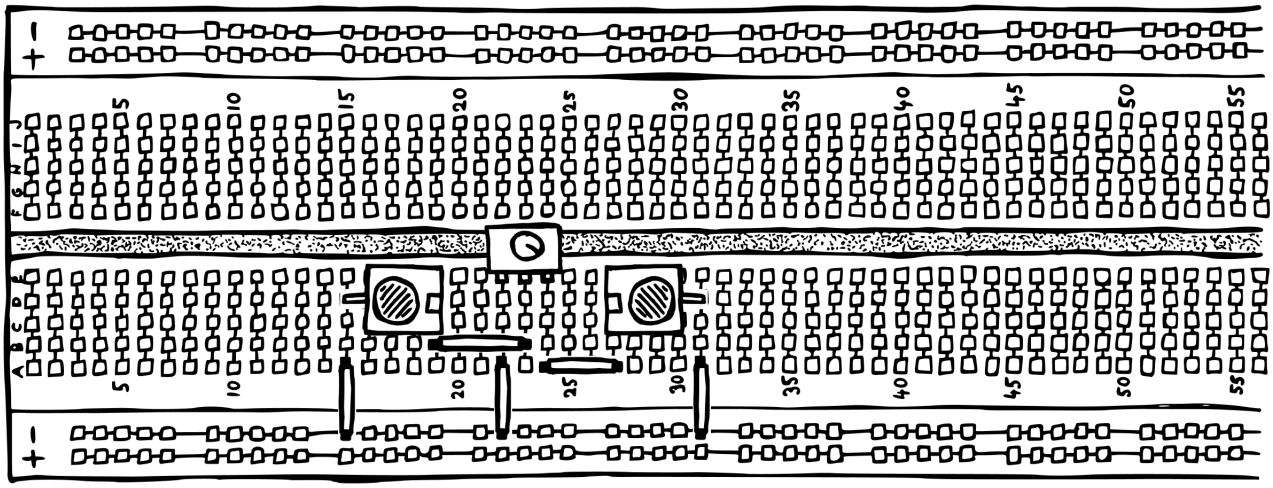

By turning the potentiometer, you should now be able to change the signal's volume.[8] With this, we know why voltage dividers are a good fit for manual volume control. **But the question is: how can we modify them to give us voltage-based volume control?** Well, if you've watched some of my earliest videos, you might already be thinking in a certain direction here.

[6] We're looking for a big component here, with a shaft on top. If the pot doesn't quite fit, try bending the legs upward a bit. No worries, this won't break it.

[7] The other elements connecting sockets and the potentiometer are called jumper wires. **Please note that these are not included in the kit**.

[8] Though depending on your headphones, the response could be a bit rocky. This is simply because our passive circuit is limiting the amount of current that'll drive your headphone's membranes.

RESISTORS VS. BIPOLAR JUNCTION TRANSISTORS

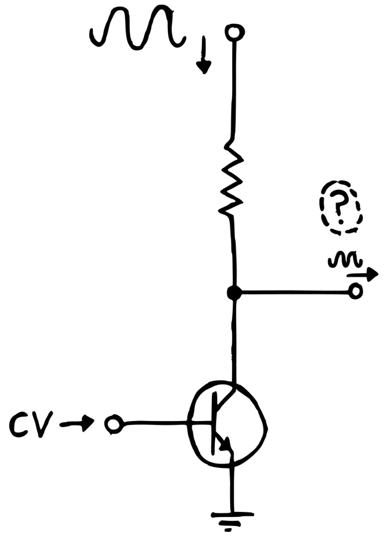

In my DIY VCO series, **I claimed that you can use NPN transistors as voltage controlled resistors.**[9] If that were true, then something like this should work, right? We could simply replace one of the two resistors in our voltage divider with a transistor. Then, by applying a voltage to the base, we should be able to control the amount of resistance between collector and emitter. Giving us a voltage controlled voltage divider in total. Unfortunately, if you build and try this, you'll be severely disappointed. And that's because what I said in that VCO video is at least misleading – if not flat out wrong. **A bipolar junction transistor is not a voltage controlled resistor.** There is no „resistance" (in that sense) between its collector and emitter.

To understand this better, let's talk about what resistors and transistors have in common – and what separates them. On the common side, both components allow us to restrict the amount of current flowing in a circuit. This is why superficially, they seem kind of interchangeable. Consider this example. We'll set up two separate, simple circuits.[10]

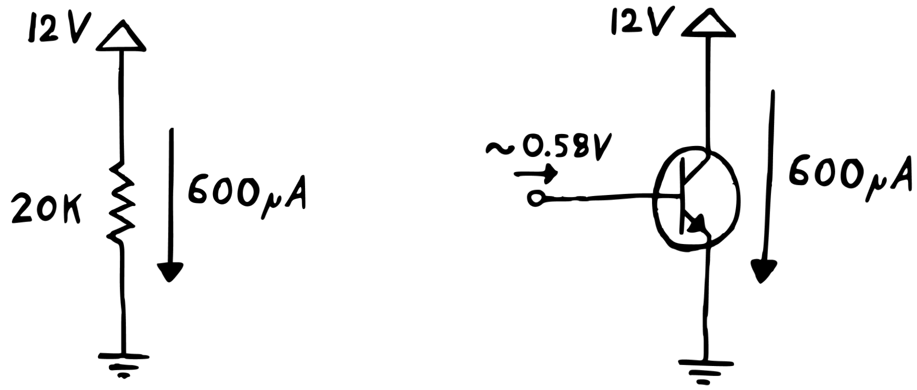

One is just a resistor between a voltage source and ground. The other is a transistor, with its collector connected to an identical voltage source, its base connected to another much weaker voltage source, and its emitter connected to ground. Now if we tune the base voltage carefully, we can get the exact same amount of current flowing through our

---

[9] Read more on NPN transistors in the components & concepts appendix (page 42).

[10] You can try this chapter's circuits in a simulator. I've already set them up for you right here: https://tinyurl.com/y9kqbqnj – you can change all values by double clicking on components.

resistor and our transistor – about 600 uA. **This might lead you to believe that for this base voltage, the path between collector and emitter acts just like a 20k resistor**.
If that were true, then it would need to exhibit the same two properties that resistors do: linearity and symmetry. Let's check for linearity first. Doubling the voltage applied to the resistor will give us twice the amount of current: 1.2 mA. But for the transistor, something odd happens: nothing.

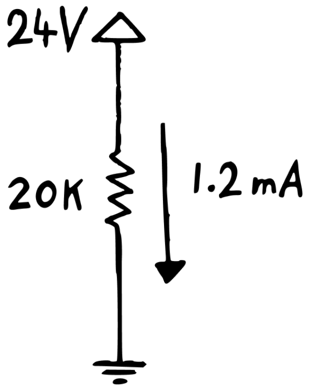

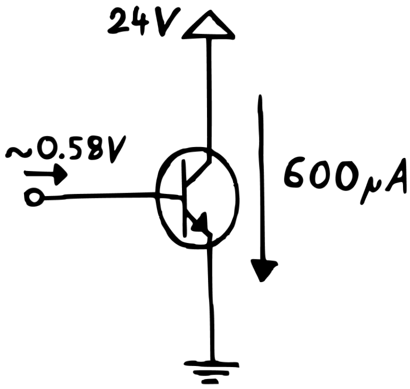

We'll still see the same 600 uA flowing through this circuit. Why is that? **Because the only thing determining the amount of current that can flow between collector and emitter is the amount of current flowing into the base.** And the latter only depends on the voltage at the base – not the voltage at the collector. Now while this would already disqualify the transistor, let's also check for the second property: symmetry. What happens when the collector voltage goes negative?

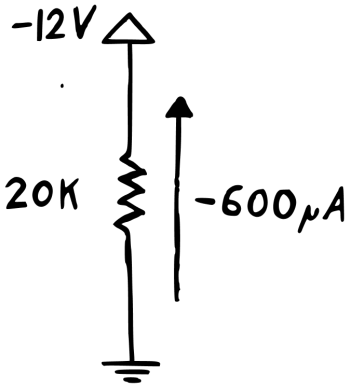

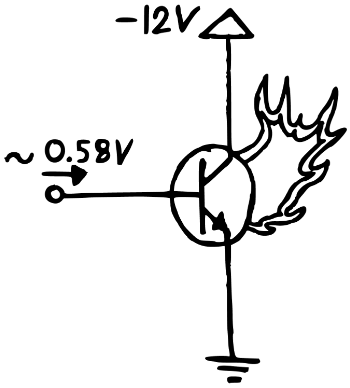

Interestingly, our transistor would theoretically still work roughly in the same way – we'll see a current flow between emitter and collector. But it won't be the same current, just inverted. Instead, it will be a ridiculously strong current that'll wreck our transistor. This is simply because the huge negative voltage between base and collector pulls a much too large amount of current into the base, disqualifying this use case without further alterations to the circuit. **So it's official: transistors are not voltage controlled resistors!**

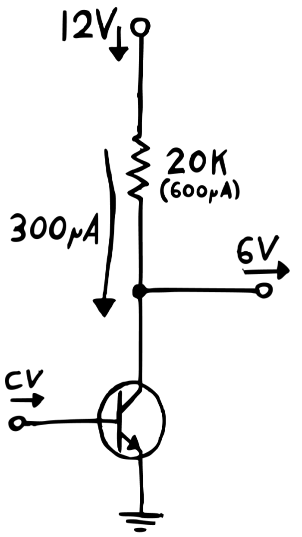

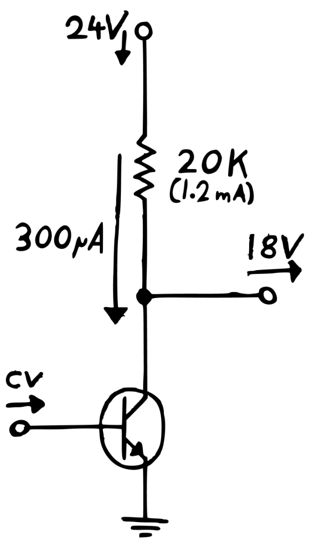

Which is precisely why the initial idea of a transistor-based voltage controlled voltage divider does not work. Since the current going through our transistor is fixed (if the base voltage does not change and the collector voltage stays positive), the total current flowing will be fixed as well. **So in- or decreasing the collector voltage will only change the potential, not the actual flow amount**.

This means that the ratio between potential and actual flow is not fixed anymore – and that's why this setup simply does not work as a voltage divider. Here's how this pans out in detail. Let's assume we've set the base voltage so that our collector-emitter-current is exactly 300 uA. Then the total current going through transistor and resistor is also 300 uA. Since the 20k resistor could potentially allow for 600 uA to flow – if we bypassed the transistor – the ratio between potential and actual flow is exactly 1/2. So we'd measure 6 V between resistor and transistor.

Now if we double the voltage up top to 24 V, the current flowing will stay the same: 300 uA. But since the resistor could now potentially allow for a flow of 1.2 mA, our ratio shifts from 1/2 to 1/4, meaning we get 3/4 the voltage at the transistor's collector: 18 V. **As we can see, the increased voltage simply gets added to the previous value – making this setup unusable for our purposes**.

TRANSISTOR AMPLIFICATION 101

Okay, but if we can't use transistors as voltage controlled resistors, then how can we use them? Well, one of their main applications is actually amplification. **Transistors can take a very small voltage swing and scale it up massively**. One way to do this is by setting them up in what's known as the common emitter amplifier configuration.[11] At its most basic, it could look like this:

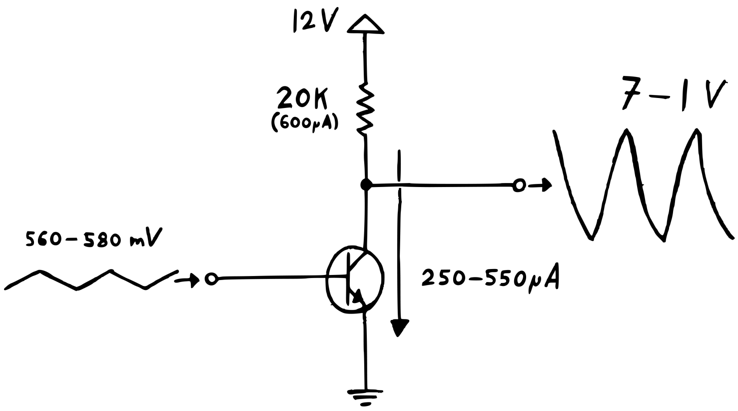

Here's how it works. We apply a very weak little 20 mV peak-to-peak signal with a comparably big 570 mV offset to the transistor's base, while monitoring the voltage at the transistor's emitter. And here we'll notice something strange. Compared to the oscillation we send into the base, there's a huge voltage swing at the collector that looks a lot like it, just slightly distorted.[12] What's up with that? To answer this, we'll first have to understand the way the base voltage relates to the amount of current allowed to flow between our transistor's collector and emitter.

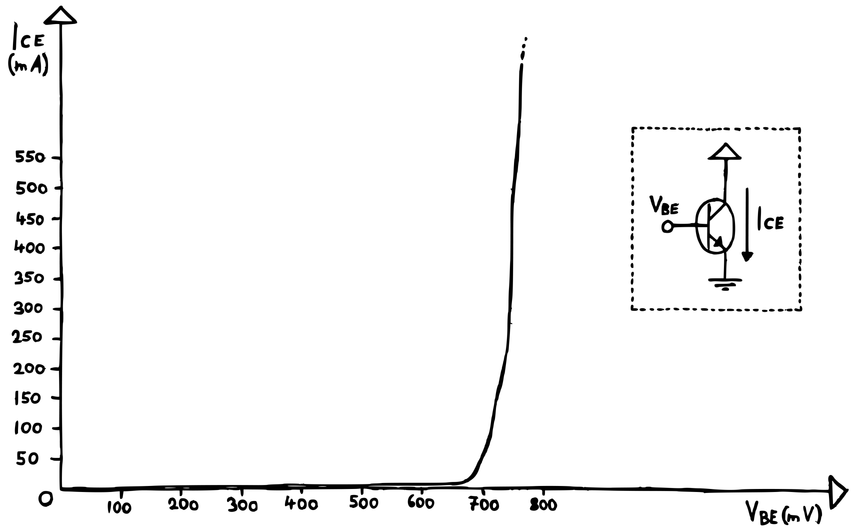

---

[11] You can try this chapter's circuits in a simulator. I've already set them up for you right here: https://tinyurl.com/ycdpkfnq – you can change all values by double clicking on components.

[12] Yup, it's also inverted. But this is not really important for our purposes here!

We'll simplify things by removing the resistor between voltage source and collector for now. Here's how this graph works, then: On the x-axis, we track the voltage applied to the base, while the y-axis shows us the measured collector-emitter-current. As you can probably tell, **the resulting plot looks suspiciously exponential. And that's because it is.** In the area between 0 mV and 600 mV base voltage, we only see a super slight increase in collector current. But then, suddenly, it starts going through the roof. And at around 775 mV, we already see a full amp of current flowing through our transistor. Or rather, we would see – if the transistor weren't already toast by then.
That's because most general purpose transistors can only handle a few hundred mA of collector-emitter-current. And so the actual usable range is somewhere between 0 mV and 700 mV base voltage. In our amplifier, we're using just a tiny sliver of this, since the scaled-down & offset signal swings between 560 mV and 580 mV.

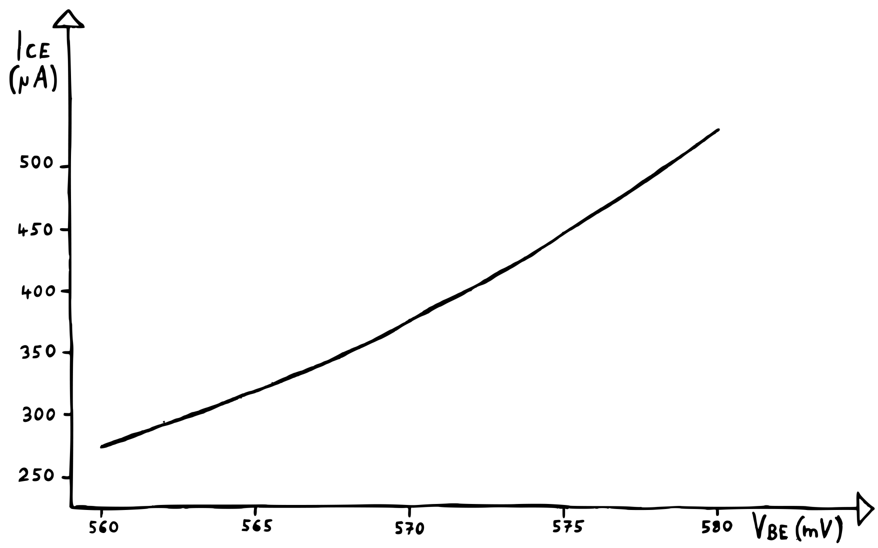

Here's what that sliver looks like when we zoom in on it. Now, since we are using a triangle oscillation as our to-be-amplified signal, we are basically moving our base voltage back and forth along the curve at a constant speed. So effectively, we are replicating the oscillation in the amount of collector current going through our transistor. And as we can see, for a rather small change in voltage we get a pretty significant change in current. Knowing this, it should should start to make sense why our setup above works as an amplifier. **Because if the current going through our transistor changes, the voltage at its collector will also change. That's because we are converting more – or less – of the resistor's potential- into actual current flow.**
Let's break this down in detail. So we know that our 20k resistor could potentially pass 600 uA in our setup. And on the graph above, we saw that our transistor's collector current is oscillating between around 250 uA and 550 uA. So the ratio between actual and potential current flow is also oscillating. At 250 uA collector current, we are using around 42% of the total potential. And at 550 uA, that number rises to about 91%. Which means that the collector voltage will swing between approximately 7 V and 1 V, giving us a 6 V

peak-to-peak output. This is an impressive amount of gain! Because remember, the signal we sent into the transistor's base was just a measly 20 mV peak-to-peak. But what about the distortion that our circuit introduces? Where's that coming from?

Simple. **If we take another close look at the relation between base voltage and collector current, we can clearly see that the plot is not linear**. **It's slightly bent**. That's because it is part of an exponential curve. And so as our voltage is oscillating at a constant pace – remember, we're using a triangle wave here – the current oscillation is slightly speeding up and slowing down. And since this change in current is directly translated into a voltage swing, that voltage swing will be slightly bent as well. Tough luck!

So what can we do about that? **There's just one solution here, really: use a smaller portion of our curve**. Why does that help? Well, the thing about exponential curves is that if you divide them into smaller and smaller chunks, these chunks begin to look more and more linear. The zoomed-in graph above is already a lot straighter than the full curve. But in order to minimize distortion, we'll have to scale our oscillation down even further.

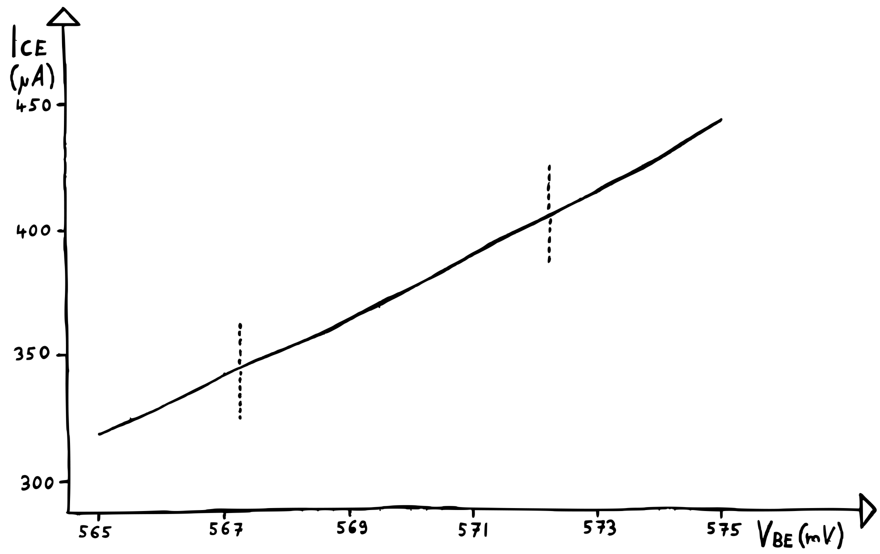

So here our signal is swinging only between 567.5 mV and 572.5 mV – netting us a tiny 5 mV peak-to-peak oscillation. **And as we can see, the resulting plot looks almost linear!** It's not perfect, mind you, but still close enough in my books.

THE EMITTER RESISTOR

Okay, so we'll just use a voltage divider to scale our input down even further, right? Actually, there's another way to achieve the same result, but with some added benefits: **adding a resistor between the transistor's emitter and ground.**[13]

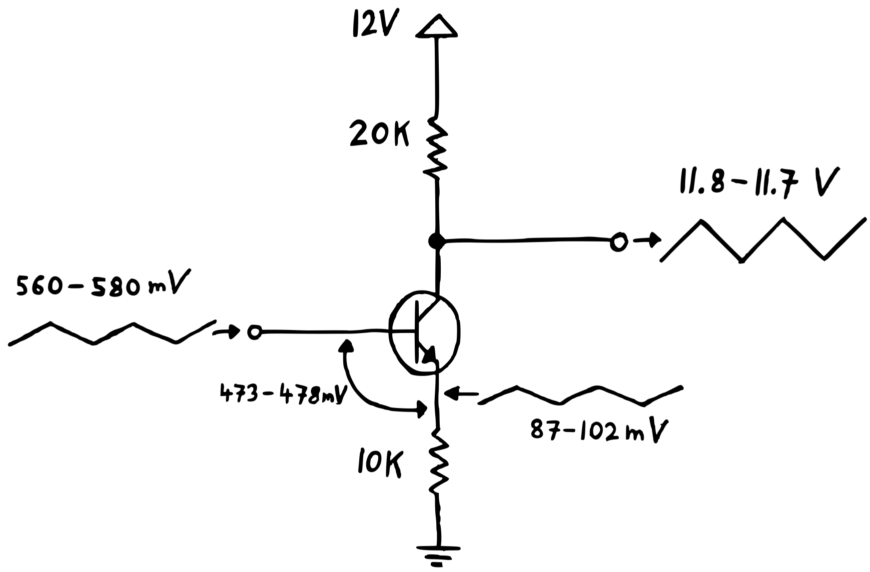

How does this work? Well, this is where it gets a bit tricky. **It works by introducing negative feedback.** As we know, our circuit operates by having a varying amount of current flowing through the transistor and to ground. Now if there's no resistance below the transistor, we have a direct, uncomplicated relation between the base voltage we apply and our collector current. But by adding in the resistor, we are restricting the current flow between emitter and ground. And this means that there's again some unused potential here – which manifests in a voltage at the emitter. How much? That depends on two factors: the amount of resistance – and the amount of current. Because the resistance is fixed, we don't have to worry about it too much. We'll only have to ask: what happens when the base voltage – and thereby the collector current – changes? Simple: if they increase, then the voltage at the emitter will increase as well. And if they decrease, it will be mirrored there, too.

Let's play this out with some actual numbers. We assume the emitter resistor to be a 10k. So our base voltage is still swinging between 560 mV and 580 mV. But if we probe the voltage at the transistor's emitter, we'll see that it is also oscillating between 87 mV and 102 mV. Now since for our transistor's operation, only the voltage difference between emitter and base voltage is what matters, we need to subtract the former from the latter. Leaving us with an effective base voltage swing between 473 mV and 478 mV. And voila:

[13] You can try this chapter's circuit in a simulator. I've already set it up for you right here: https://tinyurl.com/yd9lavjv – you can change all values by double clicking on components.

we’ve reduced it to just 5 mV peak-to-peak. But at the same time something weird happened to our amplified signal. Instead of 5 V we’re now only getting about 100 mV peak-to-peak there. Now of course, we’ve effectively scaled down our input to just 25% of what we fed into the base before. But 100 mV is not 25% of 5 V. Something else must’ve changed. And that something is our amplifier’s gain. To understand why, we just need to have another look at our base-emitter-voltage graph.

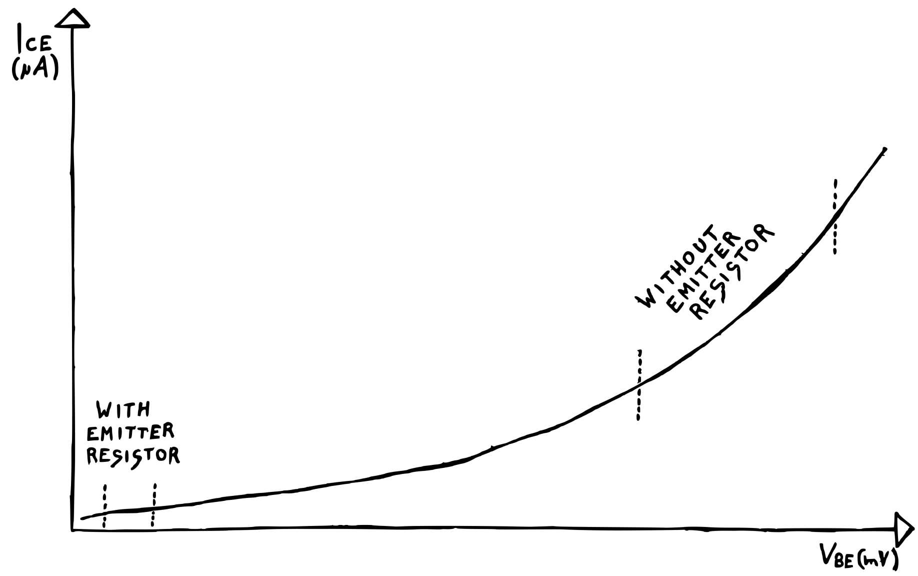

Before we added in the 10k emitter resistor, our signal was swinging between the two points up the curve. But now with the resistor in place, the active range is a much smaller slice – that’s also further down the curve. And because it is further down the curve, it is also noticeably flatter. Which is why our output signal got so much more quiet. **By moving the input oscillation down the curve, we reduced the amplifier’s gain.** We’ll get to the larger implications here in a second, but first, I want to talk about the added benefits I mentioned earlier. Because at first sight, this seems to be a rather bad deal: we needlessly cut down our amplifier’s ability to amplify. Right? Yes and no. While we did lose a significant amount of gain, we also got something important in return: stability and reliability. That’s because transistors have two severe flaws.

First, if you take two transistors – same type, same label, same manufacturer – and compare their base-voltage-to-collector-current-curves, you’ll probably notice that they don’t match up at all. That’s because of small variations in the silicon they’re made of. And second, those unequal curves are also heavily dependent on the transistor’s temperature. Heat it up by just a few degrees and your amplifier’s gain shoots up noticeably. Cool it down, and the output gets more quiet. Which is why all the „actual“ examples and numbers we’ve talked about so far are just rough approximations – your mileage trying this with a real transistor will vary quite a lot. That is, as long as you don’t use an emitter resistor!

Why’s that? Simple: because of the negative feedback it introduces. Let’s assume we build the initial circuit twice, with two different transistors. One of them naturally has a lot more gain, the other a lot less. So for the same amount of base voltage there would be very different amounts of current flowing through each respective circuit – if it weren't for the emitter resistors and the voltages they create.

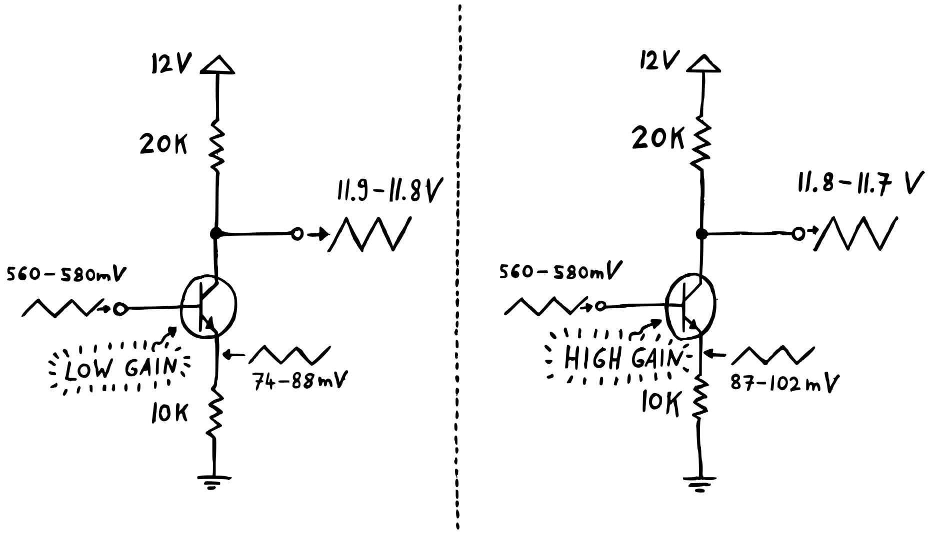

Now, since the potential amount of current on the right side is a lot higher, we'd expect to see a significantly larger voltage at the emitter there. In reality though, both emitters will just be a few mV apart. This is because the more the emitter voltage increases, the smaller the difference between base and emitter becomes – reducing the transistor's gain, and thereby the collector current, in the process. And as we've seen before, a very small reduction in base voltage can already give you a huge reduction in collector current. So **we are essentially turning the transistor's gain against itself** – the more gain the transistor has, the stronger the negative feedback and the more the gain is reduced. Conversely, if the transistor's gain is rather weak to begin with, then the negative feedback – and thereby the gain reduction – will be weaker as well. To illustrate this, I've drawn up four different plots here.

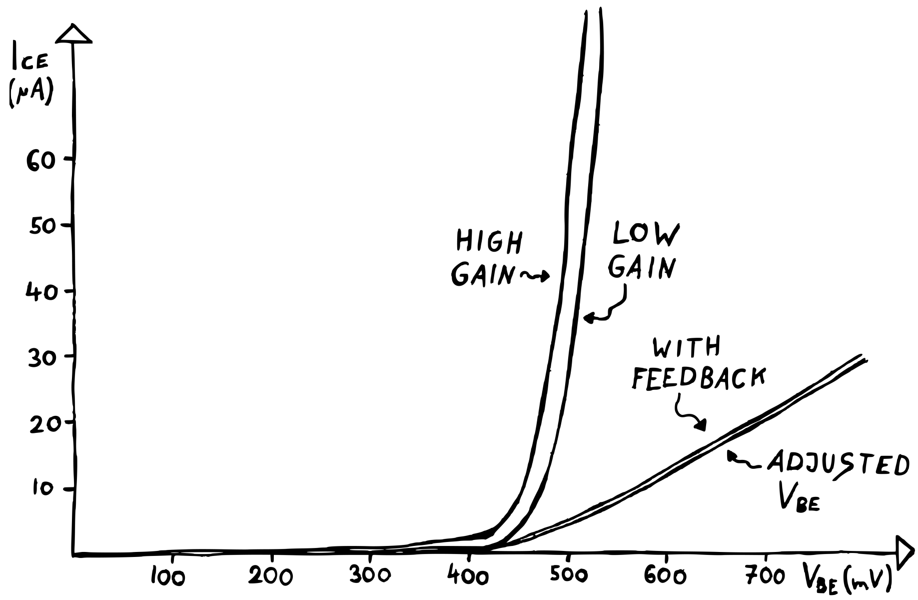

Two of them show us the base-voltage-to-collector-current-curves for our example transistors without negative feedback. The other two are what we'd measure with 10k emitter resistors in place. For better comparison, I used a little trick here: instead of measuring the actual voltage between base and emitter, I measured the voltage between base and ground. This way we can integrate the negative feedback directly into the gain curve. **And as you can see, the rather significant gain disparity between the two transistors is pretty much eliminated by the negative feedback**.

Which is why they'll perform in a very similar way in our example circuit. But that's not the only benefit. As I said before, a transistor's temperature also has an effect on its gain. Normally, that effect is quite strong – but the negative feedback again really mitigates it here. Because as the temperature rises, the gain attempts to increase, which in turn will raise the emitter voltage and that dampens the gain in response. Okay, but one final question remains: why did I pick a 10k resistor and not a bigger one? If the negative feedback is so beneficial, why not use a ton of it – to negate all differences between transistors and eliminate any temperature dependence? Simple: because we'd be completely crushing the gain – making our setup basically unusable as an amplifier. **So in the end, choosing a value for the emitter resistor is always a balancing act. How much gain are we willing to give up for reliability and stability?**

GAIN CHANGING TRICKS

That being said, the emitter resistor is not the only variable we can mess with to influence our amplifier's gain. We've got two more options here. **First, we can adjust the value of our collector resistor.**[14]

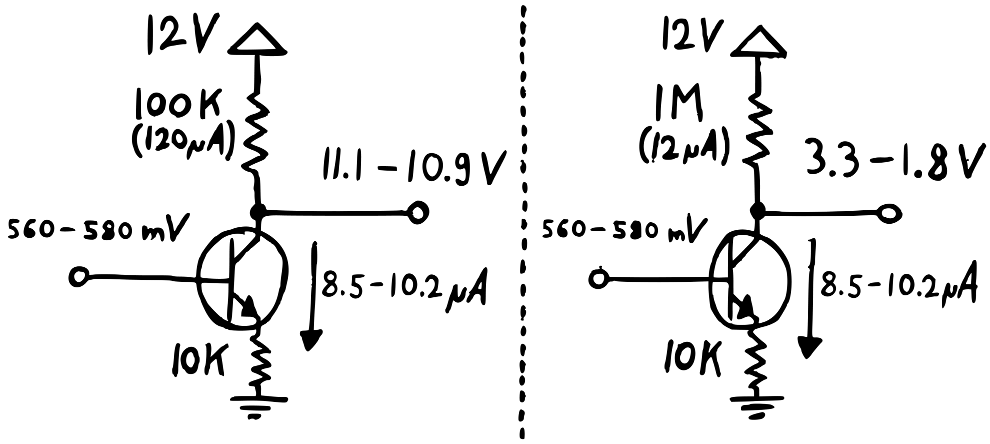

Because if we increase it, then we reduce the amount of current that could potentially flow through it. Now, since this does not affect the transistor's operation, the collector current swing will stay the same. Which means that essentially, it's swinging across a larger portion of the full potential, netting us a bigger voltage swing at the output. There is a hard limit here, though. **We can't just keep increasing the collector resistor's value in hopes of unlimited additional gain.**

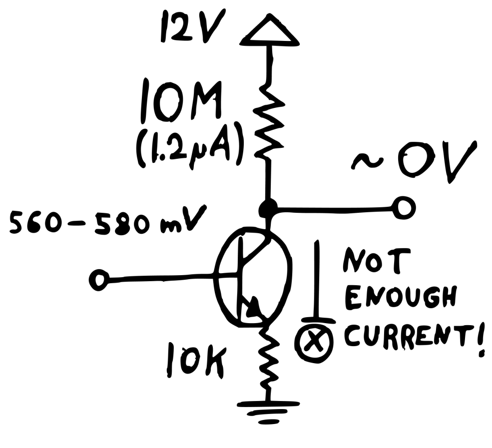

Because at some point the potential amount of current will drop below the set collector current, starving our transistor of power and distorting – or even killing – the oscillation at the output.

[14] You can try this chapter's circuits in a simulator. I've already set them up for you right here: https://tinyurl.com/y7ctexrd – you can change all values by double clicking on components.

Okay – and the second option? Well, this is where it actually gets interesting for us. Because the second option is adjusting the constant base voltage that our signal is riding on top of. Up until now, we've been using a 20 mV peak-to-peak triangle oscillation with a DC bias of 570 mV. So on our adjusted curve, we were moving back and forth across this tiny marked section. **Now what we could do is move that section down the curve – to where it is significantly less steep**.

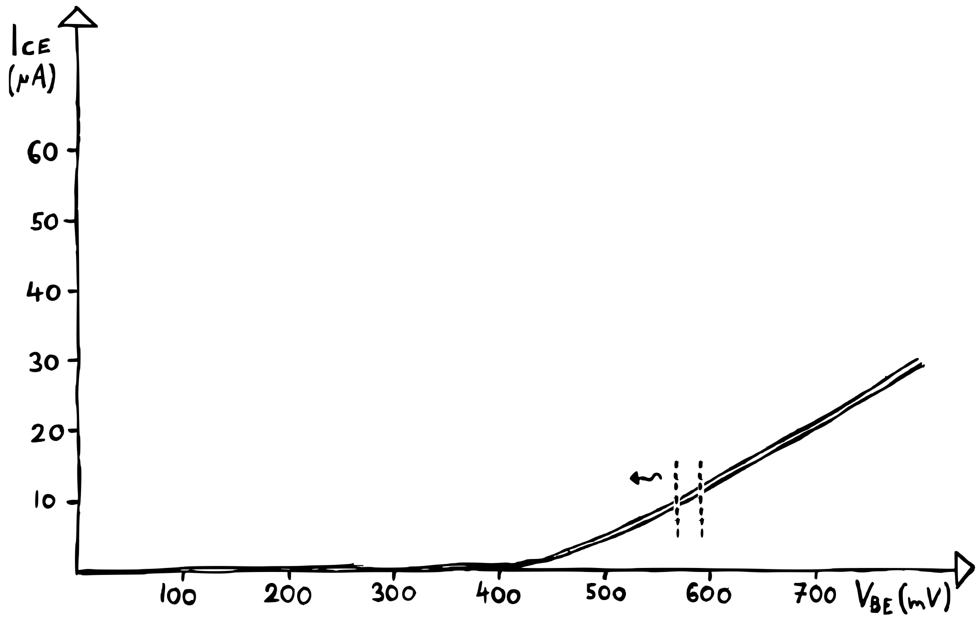

Since this will have a direct effect on our collector current swing, we should also see a reduction in gain. But while we could go through the trouble of adding our scaled-down input signal to a variable DC bias before applying the sum to the transistor's base, there is a simpler, more elegant solution.

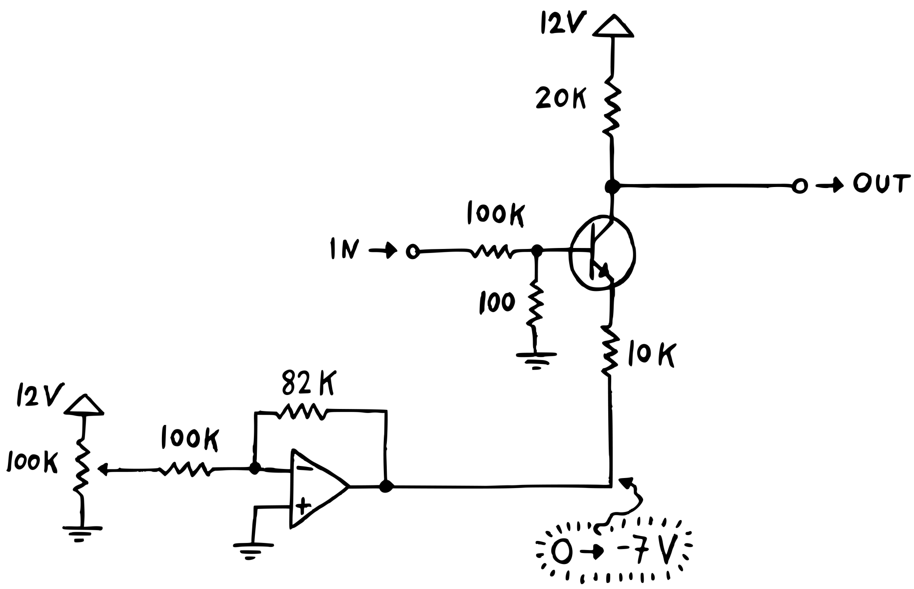

Instead of adding the bias at the base, we can just subtract it from the voltage below the emitter resistor. Because remember: the only thing that matters for the transistor's operation is the voltage difference between base and emitter. And by pulling the emitter resistor's south end below ground (so below 0 V), we can expect the voltage at the emitter to drop proportionally, too. This will increase the difference between base and emitter voltage – which in turn should push up our amplifier's gain. To try this out, we'll first set up a potentiometer as a variable voltage divider. This will be our control voltage source. Then, we'll invert and scale down that control voltage with an op amp.[15] By turning the potentiometer, we can now adjust the voltage at the op amp's output to any value between 0 V and -7 V. Going any more negative would increase the collector current too much and thereby starve our transistor – that's why I've chosen this range here.
Lastly, we'll need to scale our input signal down to the appropriate levels. Since the intention here is to pair this circuit with eurorack-compatible oscillators (which push out a 10 V peak-to-peak signal), I'm using a 100k/100 ohms voltage divider for this purpose. This way, we're effectively feeding our transistor a 10 mV peak-to-peak input. Be aware that this is an active circuit – so you'll need to hook your breadboard up to either a dual power supply or two 9 V batteries.[16] **Also, make sure that the TL072-chip (which houses two standard op amps) is set up exactly as shown here – if you reverse the power connections, it will heat up and die!** If you need help finding the components in your kit, check the components & concepts appendix.

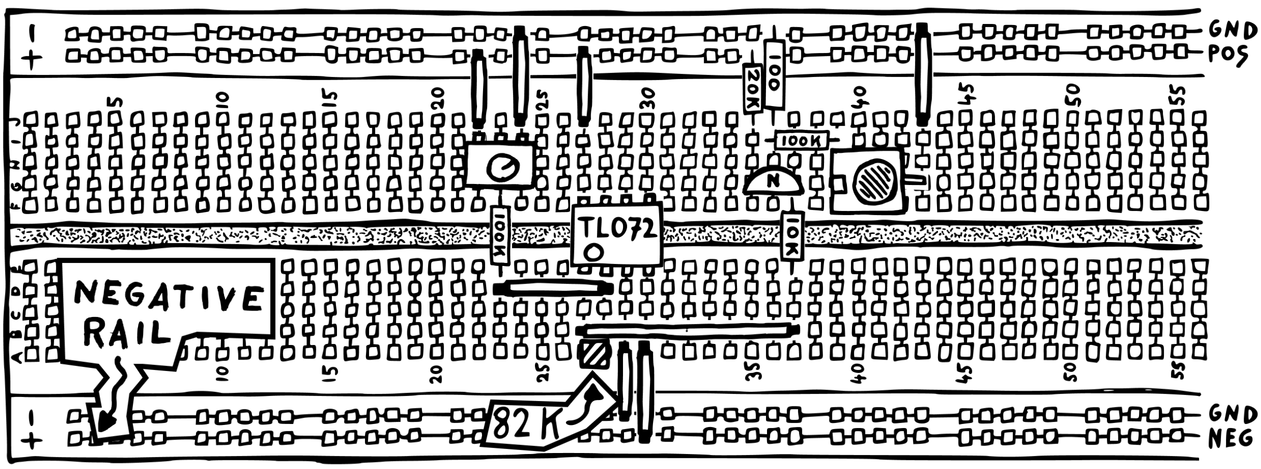

[15] Read more about op amps in the components & concepts appendix (page 39/40).

[16] If you want to use batteries, you can find a quick setup-guide in the components & concepts appendix (page 32). Keep in mind that even though I specify +/-12 V as our supply voltages in these schematics, running the circuits with +/-9 V will still largely work the same way.

After you've built this, connect an oscillator to the input socket, and your oscilloscope to the transistor's collector – which is the amplifier's output.[17] Then try adjusting the potentiometer to change the circuit's gain. **Unfortunately, the results won't be particularly useful**. There are two major problems. For one, the output signal's DC offset seems to be tied to the set CV and is moving around wildly. And also, the gain added by increasing the CV is so minimal it's actually near impossible to see the oscillation without zooming in massively. Bummer – so what's going on here? Well, two things.

First, let's talk about the output signal's DC offset. Why does it move around so much? Actually, we already kind of know the answer. What we're doing by changing the emitter voltage is moving the operating region further up or down the base-voltage-to-collector-current-curve. But by doing this, we're not only changing the amplifier's gain. We are also changing the average amount of collector current that'll be flowing while the amplifier is operating. **And the more current is flowing on average, the more of our 20k resistor's potential is used up on average**. **Meaning that the average voltage at the transistor's collector is changing as well**. So effectively, the amplifier's gain is directly tied to the output signal's DC offset. To try and compensate for that, we could of course turn on the AC coupling on our oscilloscope. This will remove any DC offset from the VCA's output. Since our oscillation is now semi-fixed to the 0 V line, you can turn up the scaling so you can actually see what's going on.

It'll be a bit hard to tell through the noise floor, but changing the CV really does have an effect on the gain. Unfortunately, the range is definitely way too small. This shouldn't be surprising, though. Because if we look at the base-voltage-to-collector-current-curve again, we can see that the steepness (and therefore the gain level) only changes slightly in the area between about 400 mV and 450 mV base voltage. **Beyond that, the emitter resistor is preventing us from increasing the gain at all, since it bends the curve into a straight line**.

On top of that, the AC coupling is not a perfect solution for removing the DC offset. Because as you're changing the gain, you can see that the signal goes off-center for a little while. That's just how AC coupling works. Because of these two problems, our simple circuit just doesn't cut it. Fortunately, we can salvage it by adding in a few extra components.

[17] If you don't have an oscilloscope, you can watch me demonstrate this here: https://youtu.be/yMrCCx6uqcE?t=1907

THE DIFFERENTIAL AMPLIFIER

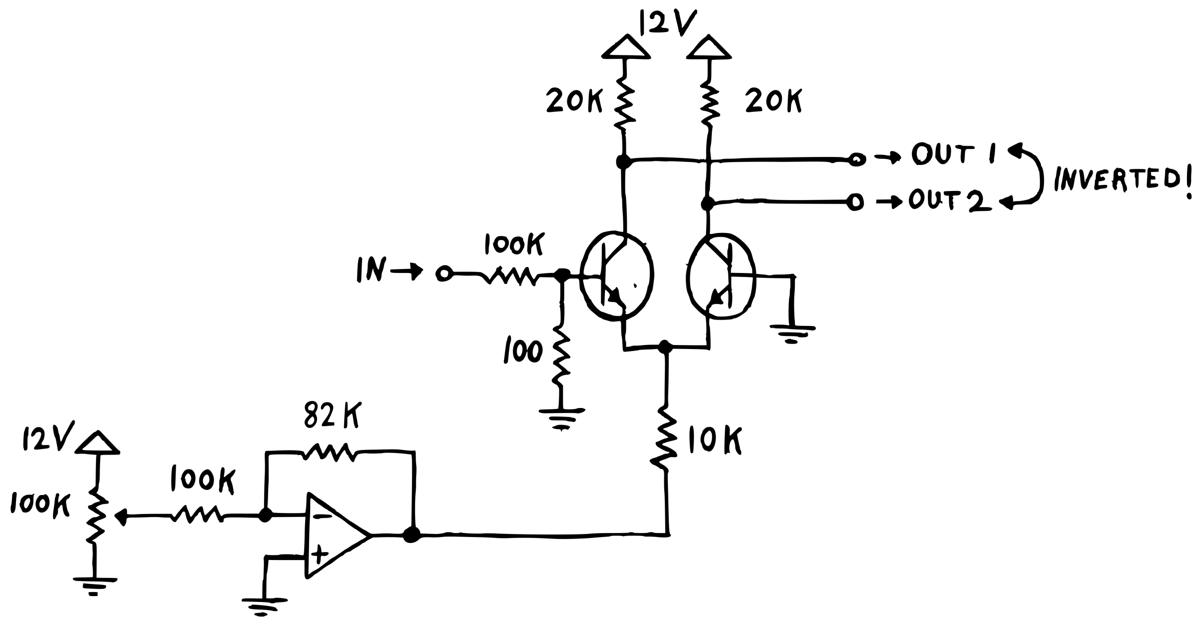

The first change we'll make probably looks very confusing at first glance: we are basically copy-pasting our transistor and its collector resistor, while connecting both emitters to the same, original emitter resistor.[18] Also, we're directly grounding the new transistor's base. What's this supposed to achieve? **By doing this, we are turning our regular amplifier into a differential amplifier**. If that doesn't mean much to you, no worries – I'll walk you through the circuit step by step.

Essentially, we are exploiting two interlinked mechanisms here. First, without any input signal, there are two stable identical currents flowing through our transistors[19] that combine at their emitters and are then sunk by the op amp. **You could call this our total bias current – which directly depends on the voltage at the op amp's output**. The more negative (below 0 V) that output goes, the stronger the total bias current will become. This is because both transistors' bases sit at (or around) 0 V while the VCA is operating. And the bigger the voltage differences between their bases and emitters get, the more the transistors will open up.

Second, and just like before, we apply a heavily scaled-down version of our input signal to the left transistor's base. As the input signal oscillates, the bias current on the left hand path will start to oscillate as well – simply because we're further manipulating the voltage between that transistor's base and emitter. Now, as the bias current oscillates, it forces its twin on the other side to oscillate as well – just in an inverted manner. This is because of

[18] You can try this chapter's circuit in a simulator. I've already set it up for you right here: https://tinyurl.com/yd7shjwz – you can change all values by double clicking on components.

[19] Assuming that the transistors are closely matched – meaning that their base-voltage-to-collector-current-curves should be close to identical.

the negative feedback created by the emitter resistor, which resists the current going through it.

Since it's easier to force the right transistor to close down (or open up) than to in- or decrease the total bias current, the former will happen. **This means that the total bias current will stay roughly constant, while the two individual collector currents fight each other for real estate within it**.

Now, the trick here is that if we in- or decrease the total bias current by changing the op amp's output voltage, we also change the transistors' gain levels. This is simply because the more current is already flowing through a transistor, the more sensitive it becomes to changes in base-emitter voltage. So as the total bias current increases, the two collector currents will oscillate more heavily. Causing the collector voltages to oscillate more heavily as well. This way, we can get as much gain as we need, since we are not relying on the steepness of the base-voltage-to-collector-current-curve anymore.

Also, if the collector resistors are matched, the collector voltage oscillations should be the same – just inverted.[20] Additionally, **their DC offset should be identical, as both collector currents swing across the exact same mid-point**.

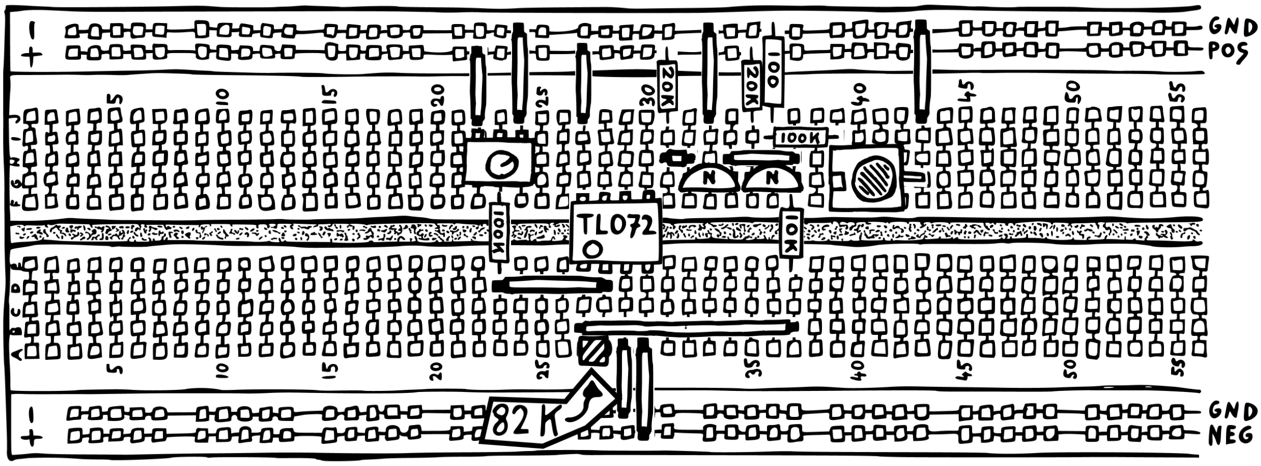

If you connect your oscilloscope to both collectors (on separate channels – or one after the other) and play with the potentiometer, you should be able to verify our assumptions.[21] Both signals are moving downwards as their volume increases significantly. Now if you zoom in a bit, you should see that they indeed are inverted copies of each other. Okay – but you might say: so what? How is having two signals with their DC offsets tied to the gain better than one?

[20] The 20k resistors in this kit have a 0.1% value tolerance – meaning they're already matched out of the box.

[21] Again, you can watch me do this if you don't have an oscilloscope: https://youtu.be/yMrCCx6uqcE?t=2310

SIGNAL SUBTRACTION

Valid point – but here's the twist. **If we take our inverted twin signals and subtract them from each other, the result will be a signal without any DC offset**. Why? Well, since both original signals have the same offset, they'll cross each other exactly at their respective mid-points.

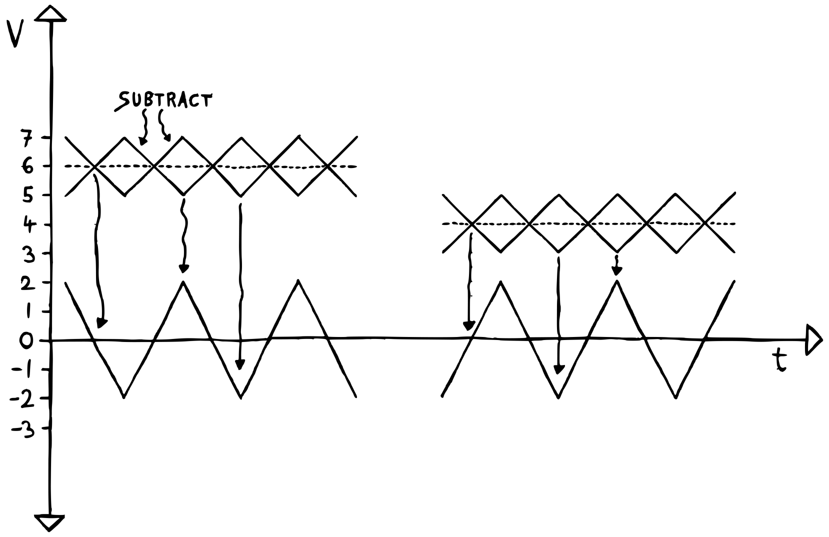

For this example, we'll assume that this would be right on the 6 V line. Now if we subtract 6 V from 6 V, we get 0 V. From here, the two signals will diverge, which means that their difference increases. At the waves' respective peak and valley, that difference is at its most extreme: 7 V versus 5 V. Subtract those numbers, and we get a 2 V peak for our new, centered oscillation. Move half a wavecycle ahead, and the same is true in the other direction, giving us a -2 V valley. **And the best thing about this is that it works regardless of what the exact DC offset is – as long as it's the same for both original signals.** Okay, but how do you subtract two signals from each other? Simple: with another op amp.[22]

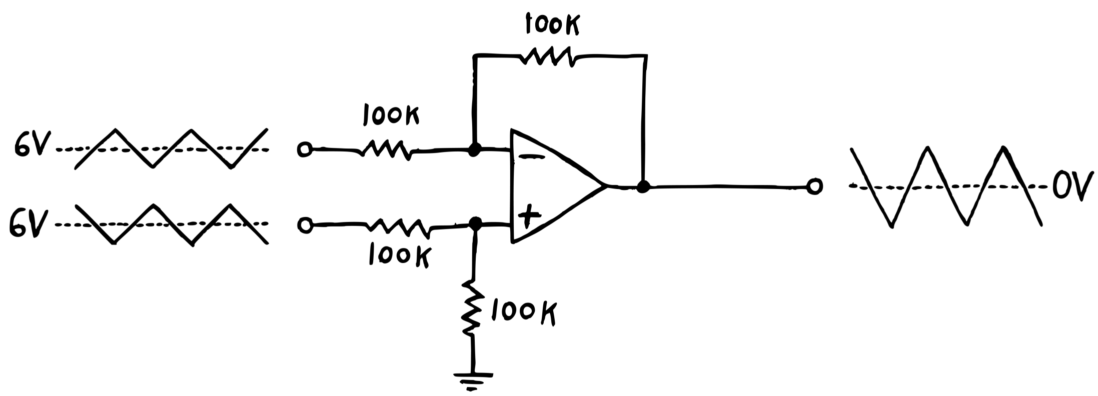

[22] You can try this chapter's circuit in a simulator. I've already set it up for you right here: https://tinyurl.com/y7gxgq23 – you can change all values by double clicking on components.

Because this is what op amps do: they subtract the voltage at their inverting input from the voltage at their non-inverting input. And after that, they multiply the result by a fixed, internal gain. Because an op amp's gain is huge, we have to take some measures to calm it down though. If we really just want to subtract our two signals without changing the resulting signal's amplitude, we can set up our op amp as shown above. Here's how it works. Let's first imagine that both input voltages are the same: 6 V.
Since the two resistors at the non-inverting input form a 50% voltage divider, we'll measure 3 V between them. Now for the op amp to enter a state of balance in this scenario, the voltage at its inverting input would need to be 3 V as well. And the only way that can really happen is if the op amp's output sits at a steady 0 V. Because then, both paths are exactly the same: two 50% voltage dividers with 6 V on one end and 0 V on the other.
Okay, but what happens if the two input voltages start to diverge? To answer that, let's assume that the bottom voltage now sits at 7 V, while the one on top has dropped to 5 V. This means that the op amp's non-inverting input reads 3.5 V. Now, since the op amp wants to get the other input up to the same value, the output really can't stay at 0 V. **Instead, it will have to make up for the difference between the two input signals.** And to do that, it'll push out a steady 2 V. Because the midpoint between 5 V and 2 V is exactly 3.5 V. Cool, so let's add this to our circuit and see how we fare.

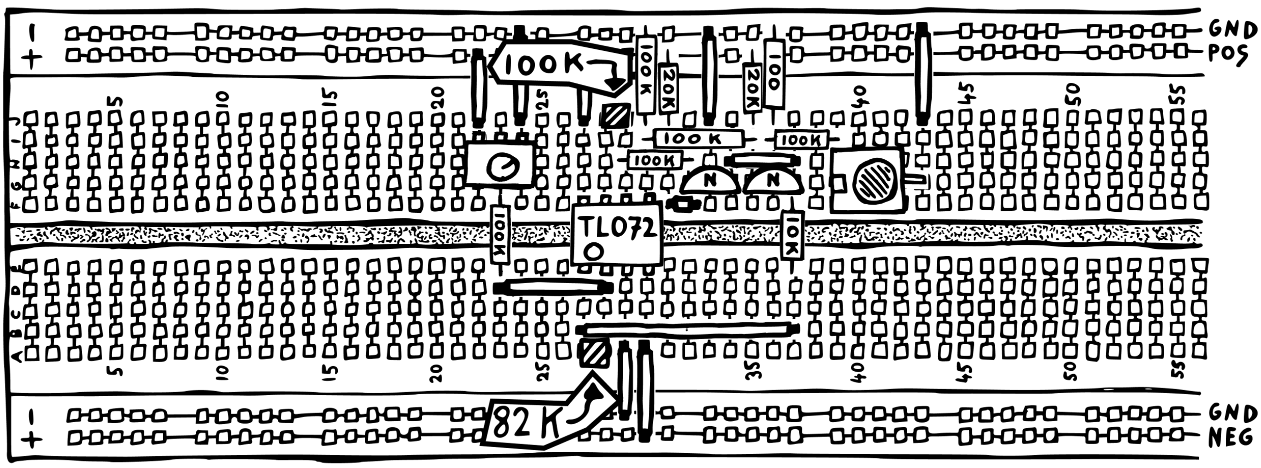

And that's it – we should now be able to pick up a properly centered signal from the subtractor op amp's output (which is the second pin from the left on the TL072's upper pin row). Let's hook up the oscilloscope to check that out.[23] Make sure the AC coupling is turned off, though! If your signal still isn't properly centered, this means that your transistors are not tightly matched. But not to worry – we can compensate for a bit of mismatch with a very simple but effective addition to the circuit.

[23] Again, you can watch me do this if you don't have an oscilloscope: https://youtu.be/yMrCCx6uqcE?t=2567

ELIMINATING TRANSISTOR MISMATCH

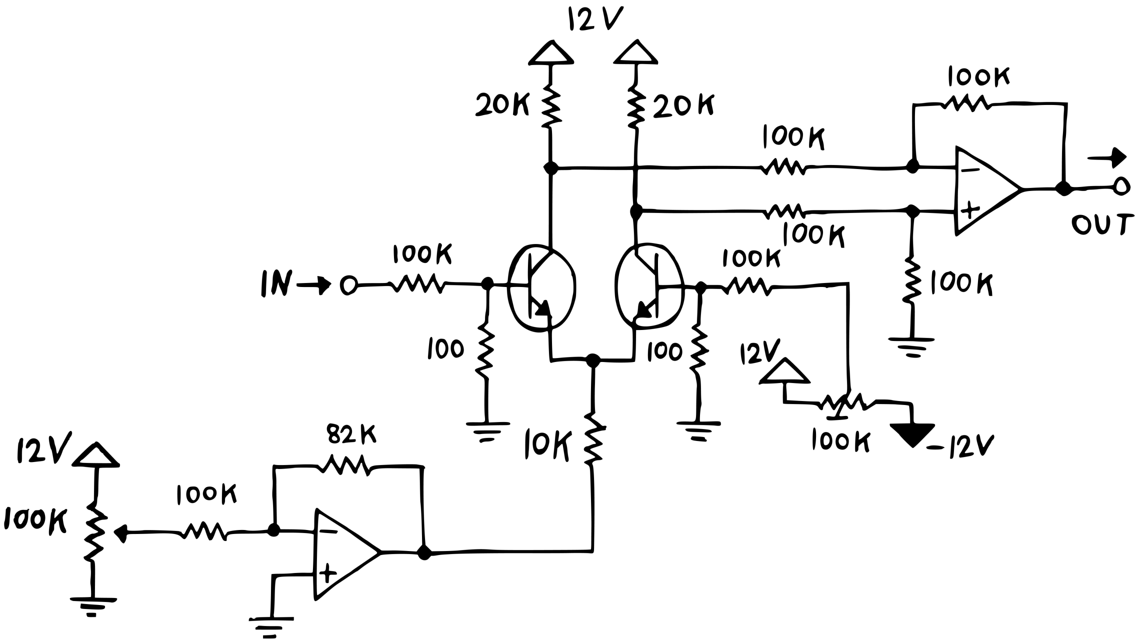

Here's how it works. Instead of connecting our second transistor's base directly to ground, we'll give it a very small, constant, non-zero voltage via a precision trimmer set up as a variable voltage divider.[24] **By doing that, it's as if we're changing the transistor's natural gain level – a positive voltage simulates more gain, while a negative one simulates less.** And it's this natural gain level that would ideally be identical between our two transistors. Because then, their collector voltage swings would have the exact same DC offset – and subtracting them from each other would give us our precious properly centered output signal.

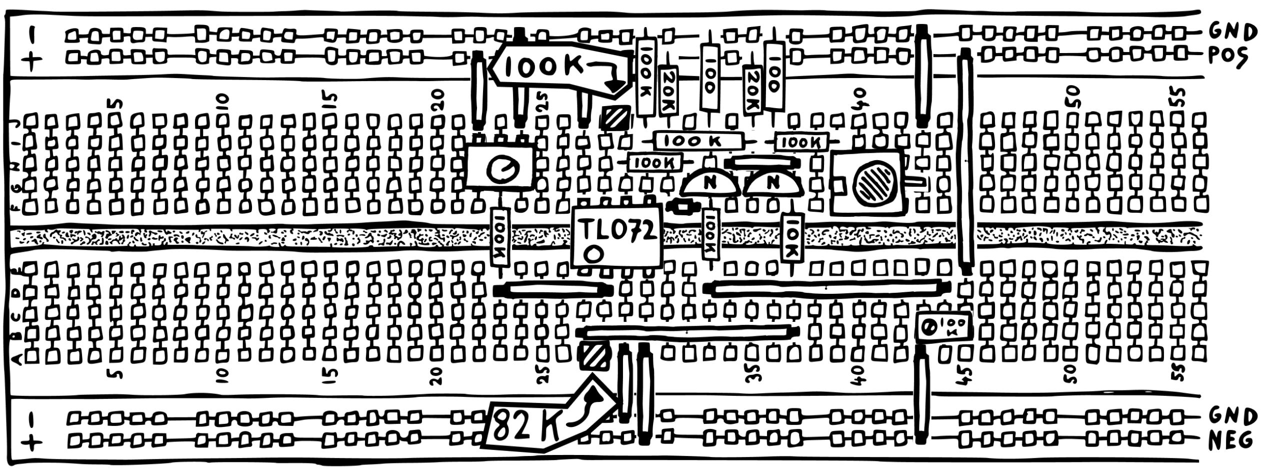

[24] You can try this chapter's circuit in a simulator. I've already set it up for you right here: https://tinyurl.com/y7kh5c72 – you can change all values by double clicking on components.

After setting this up on your breadboard, you should try fiddling with the trimmer's setting. Monitor the output signal on your oscilloscope, and aim to get the waveform's mid axis exactly aligned with the 0 V line.

After this, there's just one minor thing we need to fix. **If you turn the gain way up, you can see that the maximum voltage swing is still pretty far from the 10 V peak-to-peak that is the de-facto standard in the modular synth world.**

BOOSTING THE OUTPUT SIGNAL

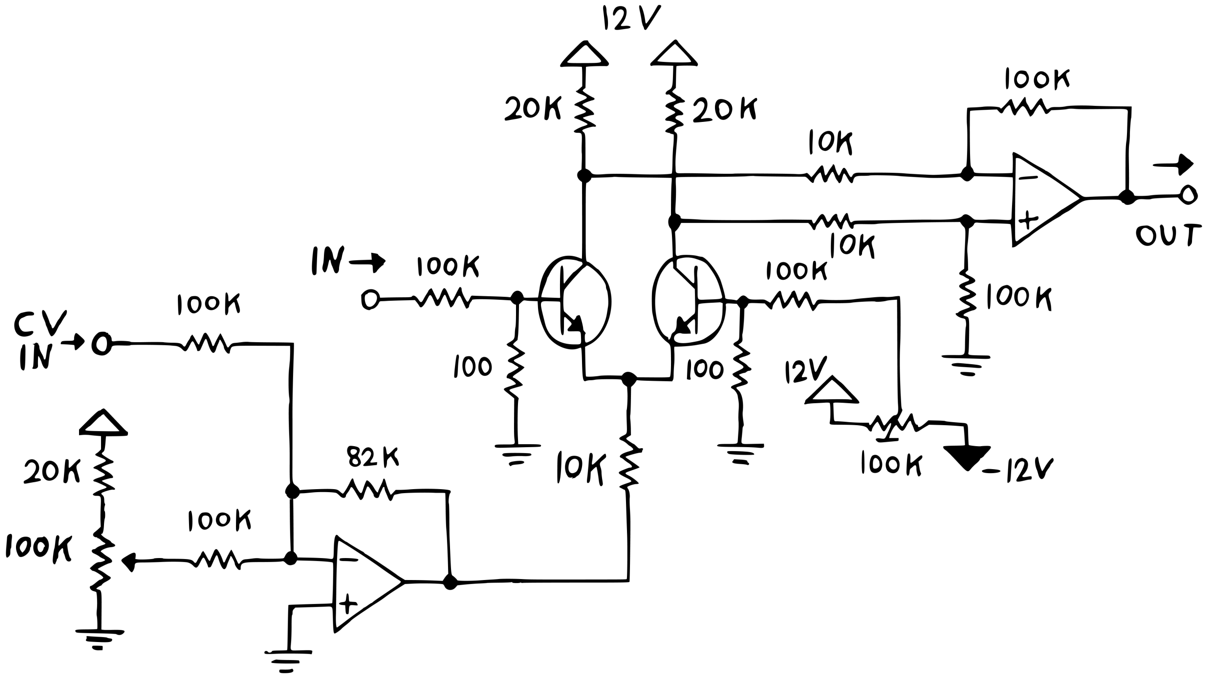

To fix this, we can simply use a bit of our subtractor op amp's gain. Because right now, we are really only using it to subtract our two signals – not to amplify the result. How do we change that? **By making the output work harder to equalize the two input voltages**. And in order to make it work harder, we have to change the relations between the resistors around the op amp.[25]

If we make the two left hand ones smaller than the ones on the right, we're essentially creating two effects that'll synergize. First, the voltage swing at the non-inverting input will get bigger. Because where before, the 50% voltage divider was effectively halving the signal's amplitude, now, we're only cutting it down by a relatively small margin. And second, the op amp's output will have to push and pull a lot harder through the feedback resistor to match the bigger swing at the non-inverting input.

The resistor values I've chosen here will actually give us a slightly louder output for a 12 V CV input than what we were aiming for. I did that because in a eurorack system, modules like envelopes or LFOs won't be pushing out voltages bigger than 10 V, tops. **With these resistor values, we'll get a 10 V peak-to-peak output for a 10 V CV input. Ensuring that we can actually reach the maximum output volume when controlling the circuit with another module.**

Speaking of connecting to other modules: in order to be able to do that, we'll have to add a CV input- and a signal output socket. Also, we'll keep the manual CV control in the

[25] You can try this chapter's circuit in a simulator. I've already set it up for you right here: https://tinyurl.com/yce5ddsh – you can change all values by double clicking on components.

circuit, so that we can dial in a CV offset value that gets summed together with the external CV input – in case you'd like to have your VCA open to varying degrees by default. Because 10 V is the expected maximum CV value, we'll also place a 20k resistor between our potentiometer and the positive rail. Ensuring that our offset will max out at precisely those 10 V.

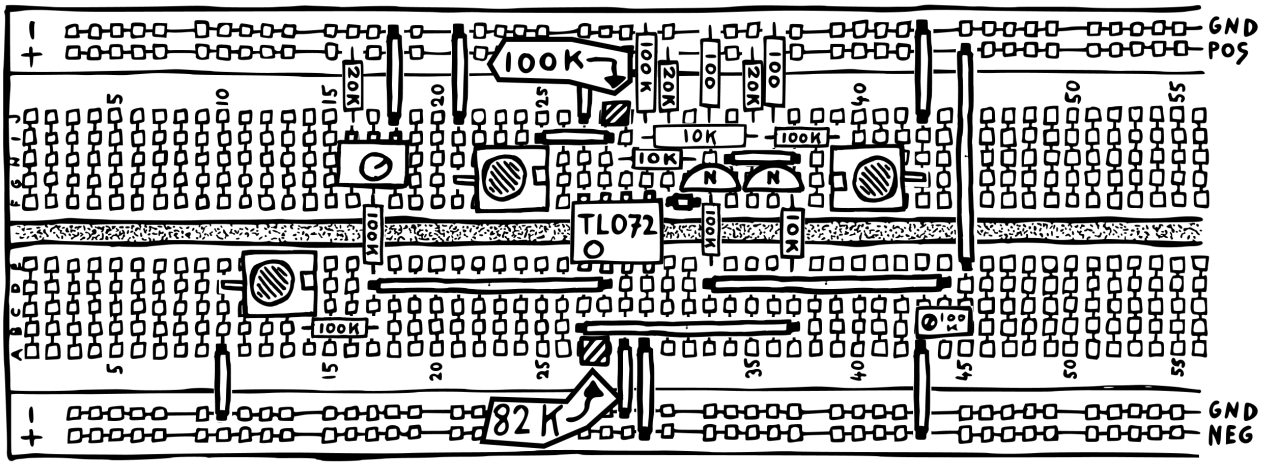

And with that, our VCA is done. Plug in your headphones, an oscillator and an LFO or envelope and see what happens! If you now want to make your creation permanent, dig out the panel and PCB from the kit, heat up your soldering iron and get to building. You can find more information on how to populate the board & how to solder in the enclosed appendices.

COMPONENTS & CONCEPTS APPENDIX

In this section, we'll take a closer look at the components and elemental circuit design concepts we're using to build our module. Check these whenever the main manual moves a bit too fast for you!

THE BASICS: RESISTANCE, VOLTAGE, CURRENT

There's three main properties we're interested in when talking about electronic circuits: resistance, voltage and current. To make these less abstract, we can use a common beginner's metaphor and compare the flow of electrons to the flow of water through a pipe.

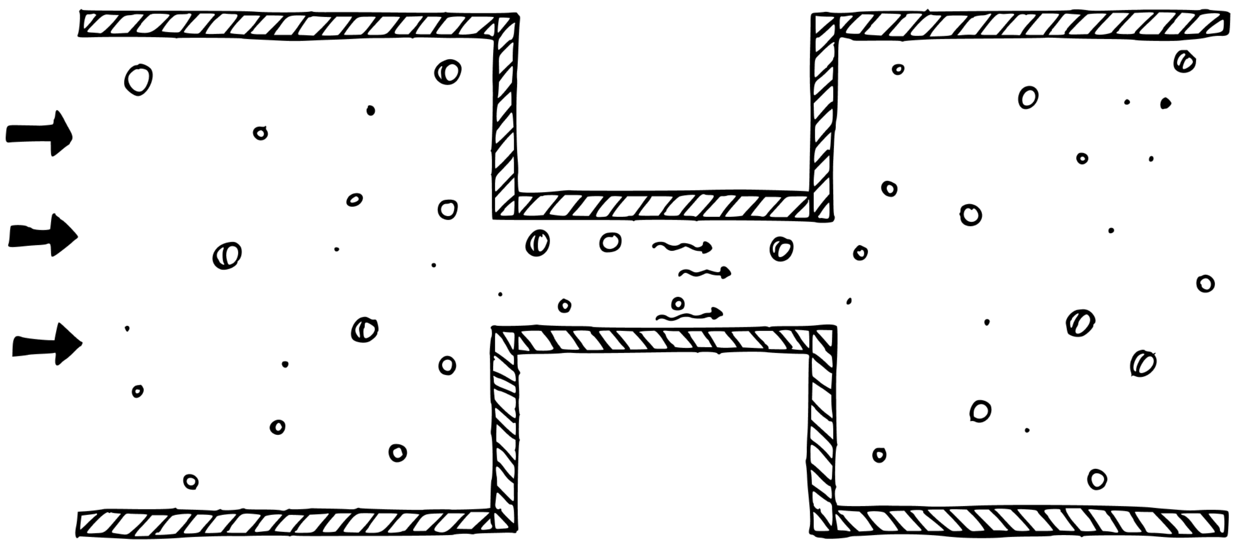

In that metaphor, resistance would be the width of a pipe. The wider it is, the more water can travel through it at once, and the easier it is to push a set amount from one end to the other. Current would then describe the flow, while voltage would describe the pressure pushing the water through the pipe. You can probably see how all three properties are interlinked: **more voltage increases the current, while more resistance to that voltage in turn decreases the current**.

USING TWO 9 V BATTERIES AS A DUAL POWER SUPPLY

Dual power supplies are great – and if you want to get serious about synth design, you should invest in one at some point. But what if you're just starting out, and you'd like to use batteries instead? Thankfully that's totally doable. **You just need to connect two 9 V batteries like shown here.** For this, you should use 9 V battery clips, which are cheap & widely available in every electronics shop.

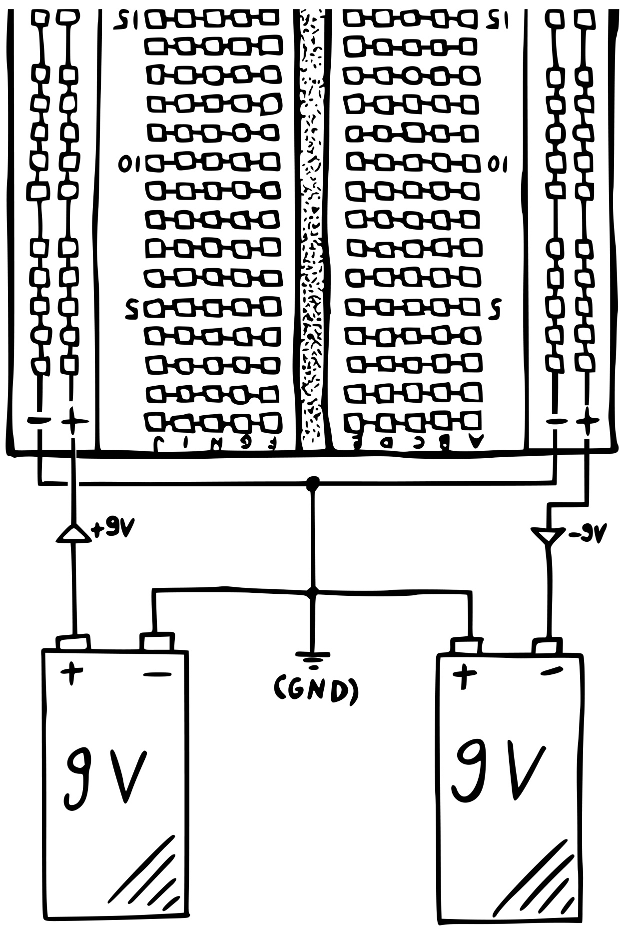

By connecting the batteries like this, the positive terminal of the left battery becomes your +9 V, while the negative terminal of the right one's now your -9 V – and the other two combine to become your new ground.[26]

Please make sure you disconnect the batteries from your breadboard when you make changes to the circuit! Otherwise you run the risk of damaging components.

[26] If you're struggling with setting this up, you can watch me do it here: https://youtu.be/XpMZoR3fgd0?t=742

RESISTORS

While a conductive wire is like a very big pipe where lots of water can pass through, **a resistor is like a narrow pipe that restricts the amount of water that can flow**. The narrowness of that pipe is equivalent to the resistance value, measured in ohms. The higher that value, the tighter the pipe.

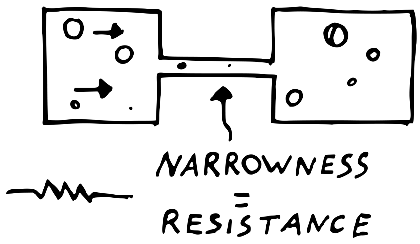

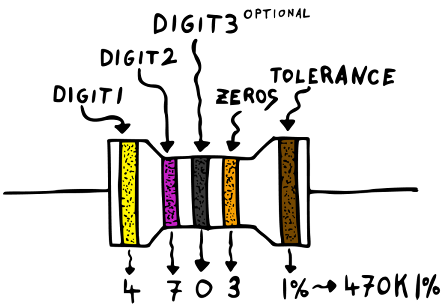

Resistors have two distinctive properties: linearity and symmetry. Linearity, in this context, means that for a doubling in voltage, the current flowing will double as well. Symmetry means that the direction of flow doesn't matter – resistors work the same either way.

On a real-life resistor, you'll notice that their value is not printed on the outside – like it is with other components. Instead, it is indicated by colored stripes[27] – along with the resistor's tolerance rating. In addition to that, the resistor itself is also colored. Sometimes, depending on who made the resistor, this will be an additional tolerance indicator.

For the resistors in this kit, a yellow body tells you that the actual resistance value might be ±5% off. A dark blue body indicates ±1% tolerance. Some kits will also contain light blue ±0.1% resistors to avoid the need for manual resistor matching.

While in the long run, learning all these color codes will be quite helpful, you can also simply use a multimeter to determine a resistor's value.

[27] For a detailed breakdown, look up *resistor color coding*. There are also calculation tools available.

CAPACITORS

A capacitor is a bit like a balloon that you can attach to the open end of a pipe. If there's some pressure in the pipe, the balloon will fill up with water until the pressure equalizes. (Since the balloon needs some space to expand into, both of the capacitor's legs need to be connected to points in your circuit.)

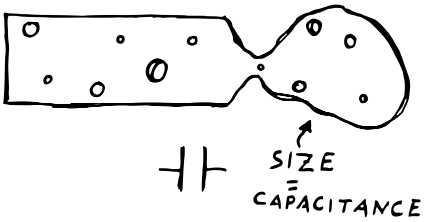

Then, should the pressure in the pipe drop, the balloon releases the water it stored into the pipe. The maximum size of the balloon is determined by the capacitor's capacitance, which we measure in farad. There are quite a few different types of capacitors: electrolytic, foil, ceramic, tantalum etc. They all have their unique properties and ideal usage scenarios – but the most important distinction is if they are polarized or not.

You shouldn't use polarized capacitors against their polarization (applying a negative voltage to their positive terminal and vice versa) – so they're out for most audio-related uses like AC coupling, high- & low pass filters etc.

Unlike resistors, capacitors have their capacitance value printed onto their casing, sometimes together with a maximum operating voltage. **Be extra careful here!** That voltage rating is important. Your capacitors can actually explode if you exceed it! So they should be able to withstand the maximum voltage used in your circuit! If they're rated higher – even better, since it will increase their lifespan. No worries though: the capacitors in this kit are carefully chosen to work properly in this circuit.

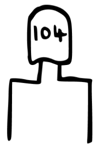

Ceramic capacitors usually come in disk or pillow-like cases, are non-polarized and typically encode their capacitance value.[28] Annoyingly, they rarely indicate their voltage rating – so you'll have to note it down when buying them.

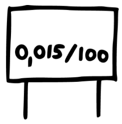

Film capacitors come in rectangular, boxy cases, are non-polarized and sometimes, but not always, directly indicate their capacitance value and their voltage rating without any form of encoding.[29]

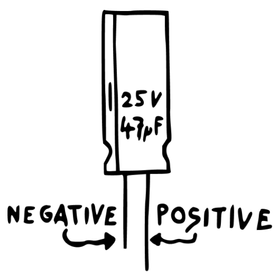

Electrolytic capacitors can be identified by their cylinder shape and silver top, and they usually directly indicate their capacitance value and their voltage rating. They are polarized – so make sure you put them into your circuit in the correct orientation.

[28] For a detailed breakdown, look up *ceramic capacitor value code*. There are also calculation tools available.

[29] If yours do encode their values, same idea applies here – look up *film capacitor value code*.

DIODES

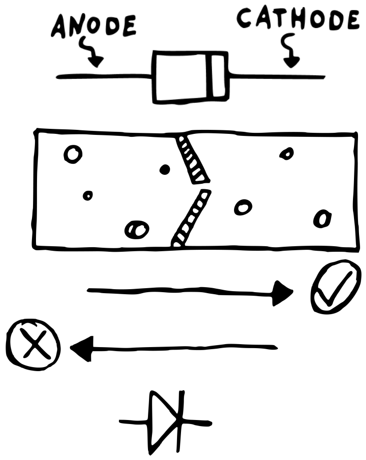

Diodes are basically like one way valves. Current can only pass through in one direction – from anode to cathode. That direction is indicated by the arrow in the diode symbol, and by a black stripe on the diode's casing. So any current trying to move in the opposite direction is blocked from flowing.

There are a few quirks here, though. For one, the diode will only open up if the pushing force is strong enough. Generally, people say that that's 0.7 V, but in reality, it's usually a bit lower. Also, diodes don't open up abruptly – they start conducting even at much lower voltages, although just slightly.

There are a lot of different diode types: zener, schottky, rectifier, small signal etc. They all have their unique properties and ideal usage scenarios – but usually, a generic 1N4148 small signal diode will get the job done.

SCHMITT TRIGGER INVERTERS

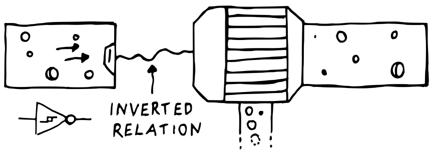

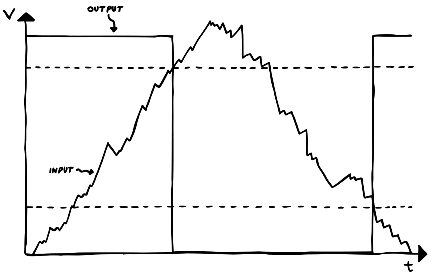

You can think of a schmitt trigger inverter as two separate things. On the left, there's a sensor that measures the pressure inside an attached pipe. On the right, there is a water pump. This pump's operation is controlled by the sensor. Whenever the pressure probed by this sensor is below a certain threshold, the pump will be working. If the pressure is above a second threshold, the pump won't be working.

Here's a quick graph to visualize that. The squiggly line represents the voltage at the input, while the dotted line shows the voltage at the output. So every time we cross the upper threshold on our way up, and the lower one on our way down, the output changes its state. One thing that's very important to keep in mind: no current flows into the sensor! It's really just sensing the voltage without affecting it.

VOLTAGE DIVIDERS

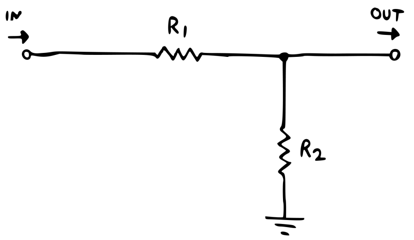

A voltage divider is really just two resistors set up like this: input on the left, output on the right. If R1 and R2 are of the same value, the output voltage will be half of what the input voltage is. How does it work?

Let's use our analogy again: so we have a pipe on the left, where water is being pushed to the right with a specific amount of force. Attached is a narrow pipe, representing R1, followed by another wide pipe. Then at the bottom, there's another narrow pipe, representing R2, where water can exit the pipe system. Finally, imagine we've set up a sensor measuring the voltage in the right hand pipe.

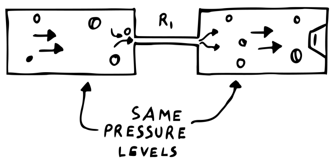

First, think about what would happen if R2 was completely sealed off. Our sensor would tell us that **the pressure on the right side is exactly the same as the pressure on the left**. Because the pushing force has nowhere else to go.

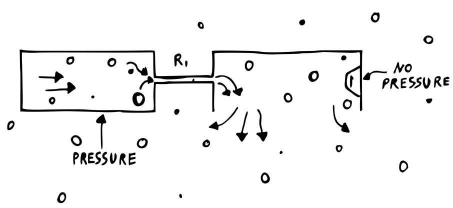

On the other hand, imagine R2 would just be a wide opening. Then **the pressure on the right would be 0**, because it'd all escape through that opening. But what happens if R2 is neither completely closed off nor wide open? Then the pressure would be retained to varying degrees, depending on the narrowness of the two resistor paths.

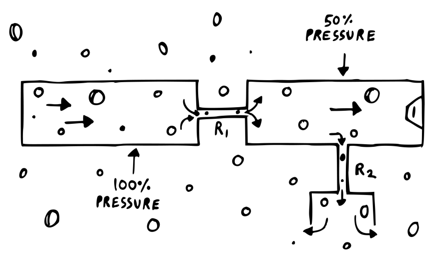

If pipe R1 is wide and pipe R2 is narrow, most of the pressure will be retained. But if it's the reverse, the pressure level will be only a tiny fraction. And if R1 and R2 are identical, **the pressure will be exactly half of what we send in**.

POTENTIOMETERS

Potentiometers can be used as variable resistors that you control by turning a knob. But, and that's the handy part, they can also be set up as variable voltage dividers. To see how that works, let's imagine we open one up.

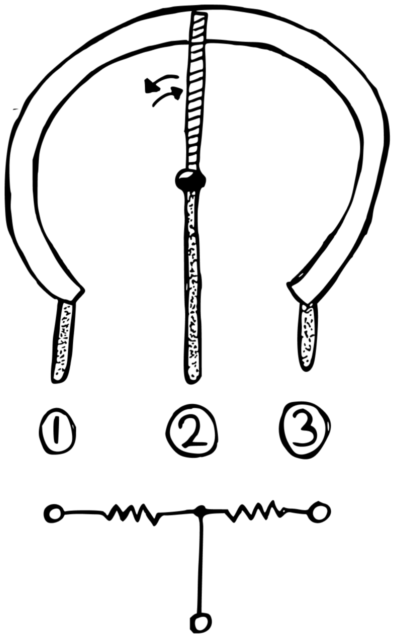

Inside, we would find two things: a round track of resistive material with connectors on both ends plus what's called a wiper. This wiper makes contact with the track and also has a connector. It can be moved to any position on the track. Now, the resistance value between the two track connectors is always going to stay exactly the same. That's why it's used to identify a potentiometer: as a 10k, 20k, 100k etc. But if you look at the resistance between either of those connectors and the wiper connector, you'll find that that's completely dependent on the wiper's position.

The logic is really simple: **the closer the wiper is to a track connector, the lower the resistance is going to be between the two**. So if the wiper is dead in the middle, you'll have 50% of the total resistance between each track connector and the wiper.

From here, you can move it in either direction and thereby shift the ratio between the two resistances to be whatever you want it to be. By now, you might be able to see how that relates to our voltage divider. If we send our input signal to connector 1 while grounding connector 3, we can pick up our output signal from the wiper. Then by turning the potentiometer's knob, we can adjust the voltage level from 0 to the input voltage – and anything in between.

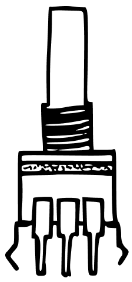

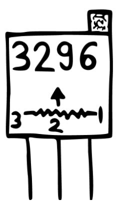

In these kits, you will encounter different types of potentiometers. First, there's the regular, full-size variant with a long shaft on top. These are used to implement user-facing controls on the module's panel and they usually – but not always – indicate their value directly on their casing. Sometimes, they'll use a similar encoding strategy as capacitors, though.[30]

Second, we've got the trimmer potentiometer, which is usually much smaller and doesn't sport a shaft on top. Instead, these have a small screw head that's supposed to be used for one-time set-and-forget calibrations. Trimmers usually encode their value.

[30] Look up *potentiometer value code* for a detailed breakdown.

AC COUPLING

What is AC coupling – and how does it work? Imagine two adjacent pipes with a balloon between them. Now, no water can get from one pipe into the other, since it's blocked by the balloon. But, and that's the kicker, **water from one side can still push into the other by bending and stretching the balloon, causing a flow by displacement**.

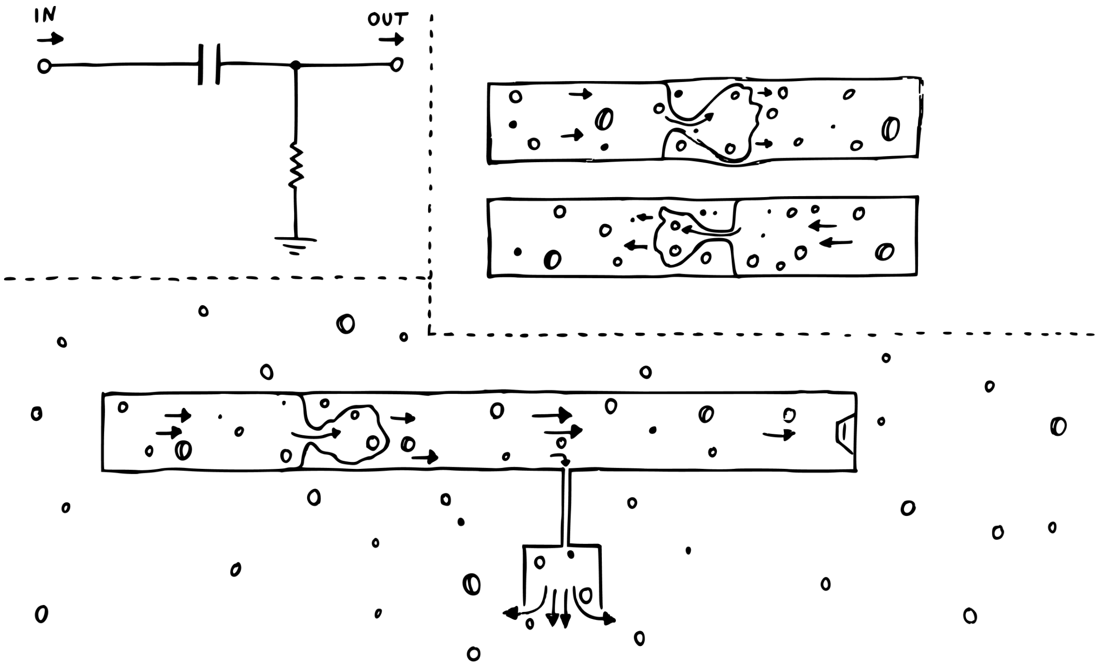

Next, we'll bring in a resistor after the coupling point, going straight to ground. **This acts like a kind of equalizing valve**. Now imagine we apply a steady 5 V from one side. Then on the other side, we'll read 0 V after a short amount of time. Why? Because we're pushing water into the balloon with a constant force, causing it to stretch into the other side, displacing some water. If we didn't have the equalizing valve there, we'd simply raise the pressure. But since we do have it, the excess water can drain out of the system. Until the pressure is neutralized, and no water is actively flowing anymore.

Okay, so now imagine that the voltage on the left hand side starts oscillating, let's say between 4 V and 6 V. When we start to go below 5 V, the balloon will begin contracting, basically pulling the water to the left. This will create a negative voltage level in the right hand pipe – like as if you're sucking on a straw, making the voltage there drop below 0 V. Then, once the pressure on the other side rises above 5 V, the balloon will inflate and stretch out again, pushing water to the right. And the pressure in the right hand pipe will go positive, making the voltage rise above 0 V. **We've re-centered our oscillation around the 0 V line.** Okay, but what about the resistor? If current can escape through it, doesn't that mess with our oscillation? Well, technically yes, but practically, we're choosing a narrow enough pipe to make the effect on quick pressure changes negligible!

OP AMPS

Op amps might seem intimidating at first, but they're actually quite easy to understand and use. The basic concept is this: every op amp has two inputs and one output. Think of those inputs like voltage sensors. You can attach them to any point in your circuit and they will detect the voltage there without interfering. **No current flows into the op amps inputs – that's why we say their input impedance is very high**. Near infinite, actually. Okay, but why are there two of them?

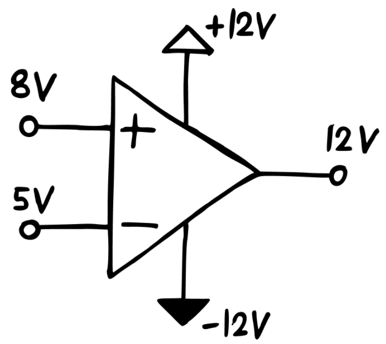

The key here is that op amps are essentially differential amplifiers. This means that they only amplify the difference between their two inputs - not each of them individually. If that sounds confusing, let's check out a quick example. So we'll imagine that one sensor – called the non-inverting input – is reading 8 V from somewhere. The other sensor – called the inverting input – reads 5 V. Then as a first step, the op amp will subtract the inverting input's value from the non-inverting input's value. Leaving us with a result of 3. (Because 8 minus 5 is 3.) This result then gets multiplied by a very large number – called the op amp's **gain**. Finally, the op amp will try to push out a voltage that corresponds to that multiplication's result.

But of course, the op amp is limited here by the voltages that we supply it with. If we give it -12 V as a minimum, and +12 V as a maximum, the highest it can go will be +12 V. So in our example, even though the result of that multiplication would be huge, the op amp will simply push out 12 V here and call it a day.

The handy thing though about op amp outputs is that they draw their power directly from the power source. This means that they can supply lots of current while keeping the voltage stable. **That's why we say an op amp has a very low output impedance.**

OP AMP BUFFERS/AMPLIFIERS

Buffering, in the world of electronics, means that we provide a perfect copy of a voltage without interfering with that voltage in the process. With an op amp-based buffer, the buffering process itself works like this. We use the non-inverting input to probe a voltage, while the inverting input connects straight to the op amp's output. **This creates what we call a negative feedback loop**. Think of it this way. We apply a specific voltage level to the non-inverting input – let's say 5 V.

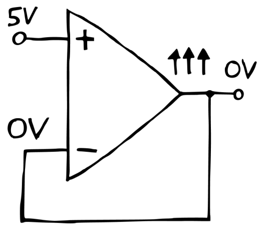

Before the op amp starts processing the voltages at its inputs, the output will be switched off. This means that **output and inverting input sit at 0 V at first**. So then, the op amp will subtract 0 from 5 and multiply the result by its gain. Finally, it will try and increase its output voltage to match the calculation's outcome.

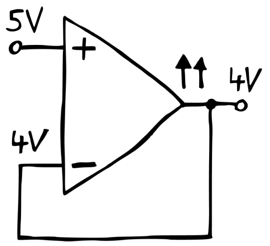

But as it's pushing up that output voltage, the **voltage at the inverting input will be raised simultaneously**. So the difference between the two inputs is shrinking down. Initially, this doesn't matter much because the gain is so large. As the voltage at the inverting input gets closer to 5 V though, the difference will shrink so much that in relation, the gain suddenly isn't so large anymore.

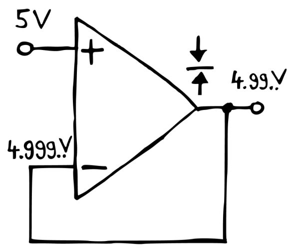

Then, the output will **stabilize at a voltage level that is a tiny bit below 5 V**, so that the difference between the two inputs multiplied by the huge gain gives us exactly that voltage slightly below 5 V. And this process simply loops forever, keeping everything stable through negative feedback. Now if the voltage at the non-inverting input changes, that feedback loop would ensure that the output voltage is always following. So that's why this configuration works as a buffer: the **output is simply following the input**.

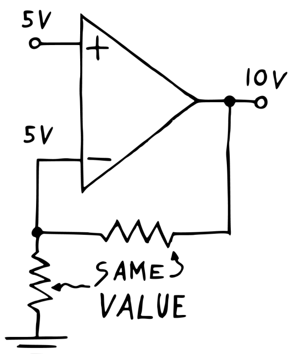

How about amplifying a signal though? To do that, we'll have to turn our buffer into a proper non-inverting amplifier. We can do that by replacing the straight connection between inverting input and output with a voltage divider, forcing the op amp to work harder. Here's how that works. Say we feed our non-inverting input a voltage of 5 V. Now, **the output needs to push out 10 V in order to get the voltage at the inverting input up to 5 V**. We call this setup a non-inverting amplifier because the output signal is in phase with the input.

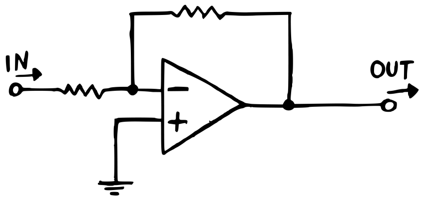

For an inverting buffer/amplifier, the input signal is no longer applied to the non-inverting input. Instead, that input is tied directly to ground. So it'll just sit at 0 V the entire time. The real action, then, is happening at the inverting input. Here, we first send in our waveform through a resistor. Then, the inverting input is connected to the op amp's output through another resistor of the same value.

How does this work? Well, let's assume that we're applying a steady voltage of 5 V on the left. Then as we already know, the op amp will subtract the inverting input's from the non-inverting input's voltage, leaving us with a result of -5 V. Multiply that by the huge internal gain, and the op amp will try to massively decrease the voltage at its output.

But as it's doing that, an increasingly larger current will flow through both resistors and into the output. Now, as long as the pushing voltage on the left is stronger than the pulling voltage on the right, some potential (e.g. a non-zero voltage) will remain at the inverting input. Once the output reaches about -5 V though, we'll enter a state of balance. Since both resistors are of the same value, the pushing force on the left is fighting the exact same resistance as the pulling force on the right. **So all of the current being pushed through one resistor is instantly being pulled through the other**.

And that means that the voltage at the inverting input will be lowered to about 0 V, allowing our op-amp to settle on the current output voltage level. So while we read 5 V on the left, we'll now read a stable -5 V at the op amp's output. Congrats – we've built an inverting buffer! **If we want to turn it into a proper amplifier, we'll simply have to change the relation between the two resistances.** By doing this, we can either increase (if you increase the right-hand resistor's value) or reduce (if you increase the left-hand resistor's value) the gain to our heart's content.

BIPOLAR JUNCTION TRANSISTORS

Bipolar junction transistors (or BJTs for short) come in two flavors: NPN and PNP. This refers to how the device is built internally and how it'll behave in a circuit. Apart from that, they look pretty much identical: a small black half-cylinder with three legs.

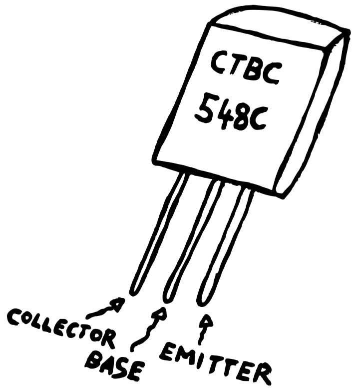

Let's take a look at the more commonly-used NPN variant first. Here's how we distinguish between its three legs. **There's a collector, a base and an emitter.**[31] All three serve a specific purpose, and the basic idea is that you control the current flow between collector and emitter by applying a small voltage[32] to the base. The relation is simple: **more base voltage equals more collector current**. Drop it down to 0 V and the transistor will be completely closed off. Sounds simple – but there are four important quirks to this.

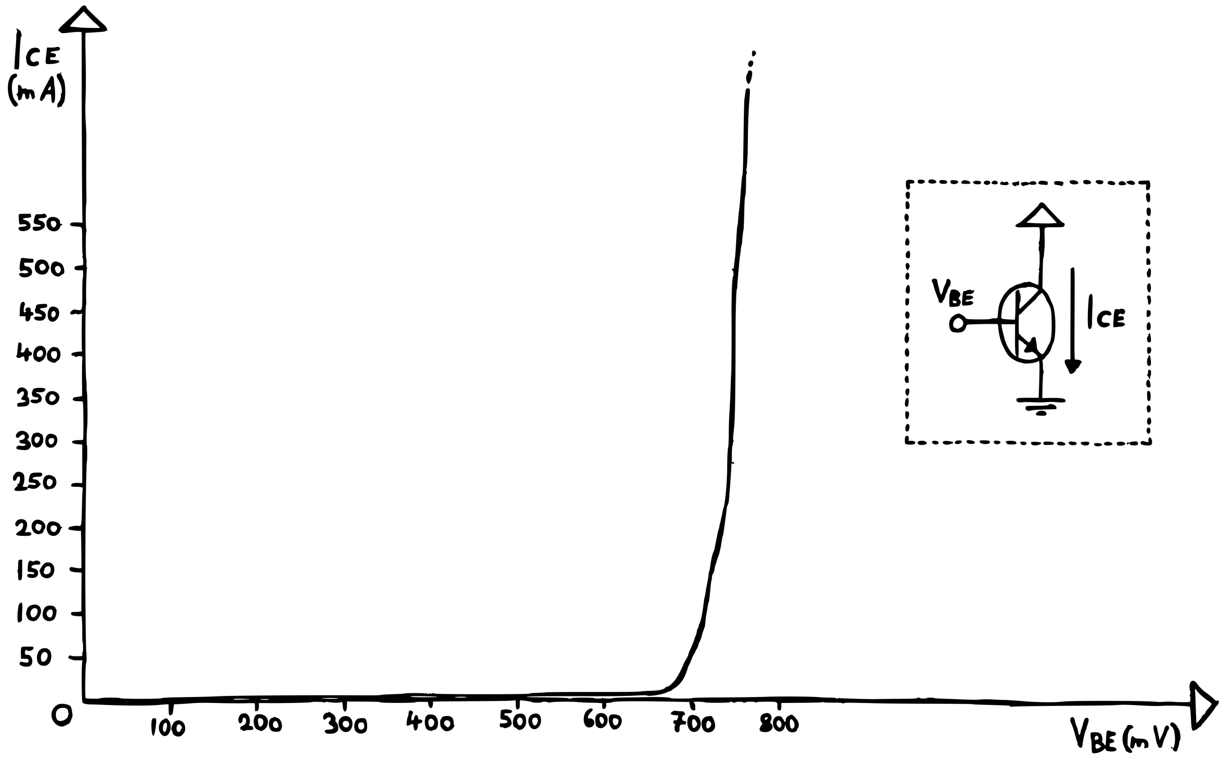

First, the relation between base voltage and collector current is exponential. Second, unlike a resistor, a BJT is not symmetrical – so we can't really reverse the direction of the

[31] Please note that the pinout shown here only applies for the BC series of transistors. Others, like the 2N series, allocate their pins differently.

[32] The voltage is measured between base and emitter. So „a small voltage" effectively means a small voltage **difference** between base and emitter!

collector current. (At least not without some unwanted side effects.) Third, also unlike a resistor, a BJT is not a linear device. Meaning that a change in collector voltage will not affect the collector current. And fourth, the collector current **b** affected by the transistor's temperature! The more it heats up, the more current will flow.

Now, for the PNP transistor, all of the above applies, too – except for two little details. Unlike with the NPN, **the PNP transistor decreases its collector current when the voltage at its base increases**[33]. So you have to bring the base voltage below the emitter to open the transistor up. Also, that collector current flows out of, not into the collector!

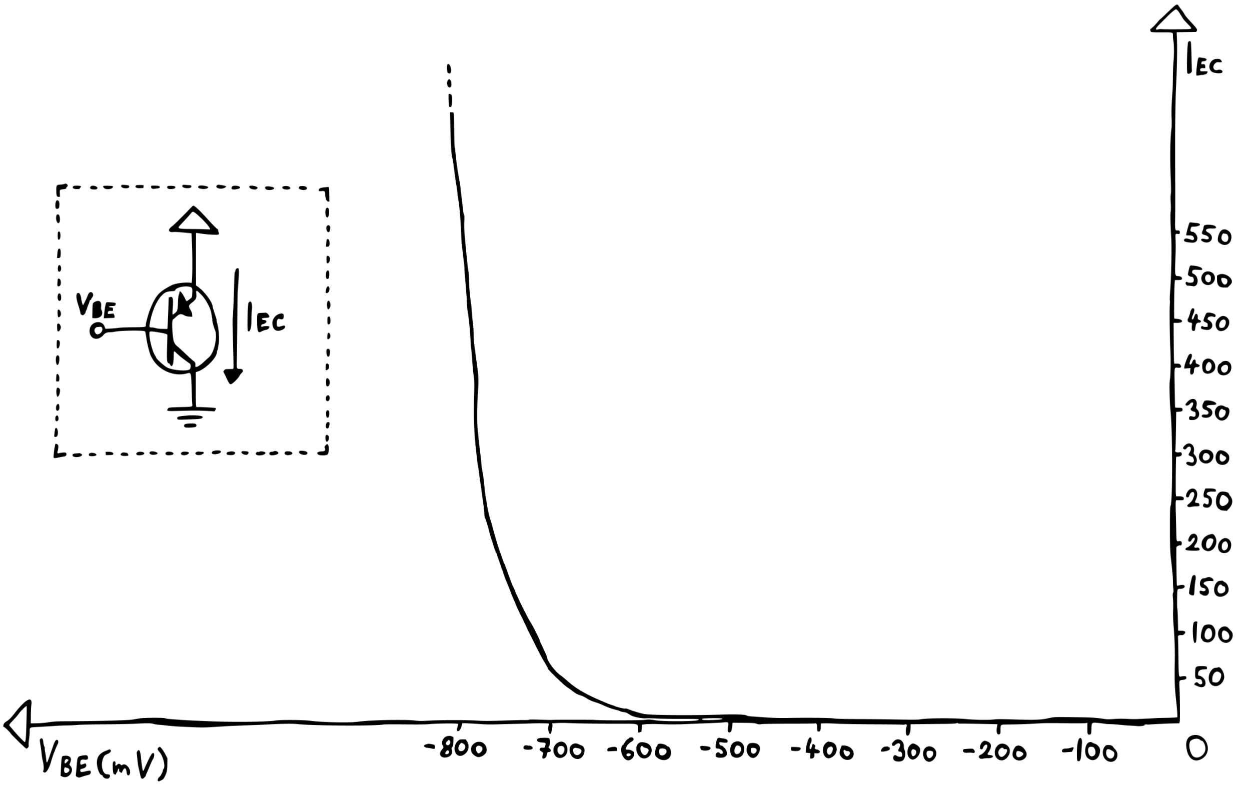

33 Again, the voltage is measured between base and emitter.

TOOLS APPENDIX

There are two types of tools that will help you tremendously while designing a circuit: multimeters and oscilloscopes. In this appendix, we'll take a quick look at each of these and explore how you use them.

MULTIMETERS

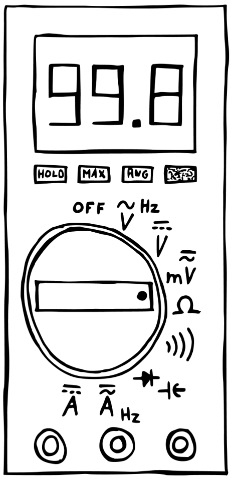

Multimeters come in different shapes and sizes, but the most common type is probably the hand-held, battery powered variant. These can measure a bunch of different things: voltage, current, resistance, continuity. Some have additional capabilities, allowing you to check capacitance, oscillation frequency or the forward voltage drop of a diode.

When shopping for one, you'll probably notice that there are really expensive models boasting about being TRUE RMS multimeters. For our purposes, this is really kind of irrelevant, so don't feel bad about going for a cheap model!

Using a multimeter is actually really straightforward. Simply attach two probes to your device – the one with a black cable traditionally plugs into the middle, while the red one goes into the right connector. Next, find whatever you want to measure and select the corresponding mode setting.

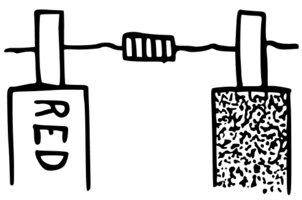

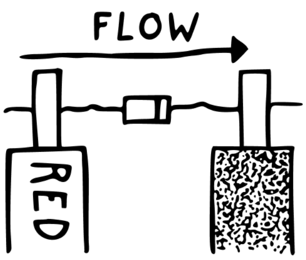

In some cases, it doesn't matter which probe you connect to which component leg or point in your circuit. This is true for testing resistors, non-polarized capacitors (foil/film, ceramic, teflon, glass etc.), continuity[34] or AC voltage.

In others, you'll have to be careful about which probe you connect where. For testing the forward voltage drop of a diode, for example, **the multimeter tries to push a current from the red to the black probe**. Here, you'll have to make sure the diode is oriented correctly, so that it doesn't block that current from flowing. For testing a DC voltage, you want to make sure the black probe is connected to ground, while you use the red one to actually take your measurement.

[34] Just a fancy word for saying that two points are electrically connected.

OSCILLOSCOPES

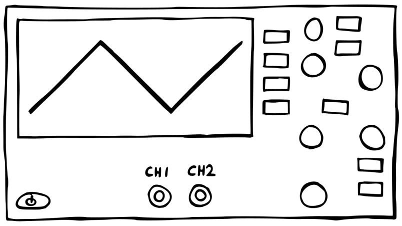

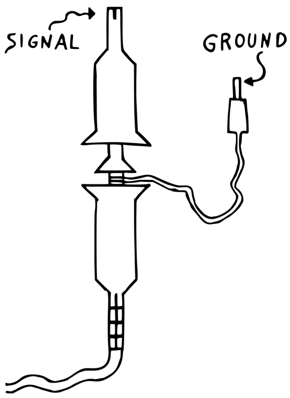

While multimeters are fairly cheap and compact, oscilloscopes are usually somewhat pricey and bulky. **If you're willing to make the investment, they are a huge help with the troubleshooting process, though**. Using one is, again, surprisingly straightforward – if you manage to work your way through the sometimes quite convoluted UI, especially on digital models.

To start using your scope, simply attach a probe to one of the channel inputs. These probes usually have two connectors on the other end: a big one that you operate by pulling the top part back – and a smaller one, which is usually a standard alligator clip. The latter needs to be connected to your circuit's ground rail, while you probe your oscillation with the former. Now what the oscilloscope will do is **monitor the voltage between the two connectors over time and draw it onto the screen as a graph**. Here, the x-axis is showing time, while the y-axis is showing voltage. You can use the device's scaling controls to zoom in on a specific part of your waveform.

Usually, digital oscilloscopes will also tell you a couple useful things about the signal you're currently viewing: minimum/maximum voltage level, oscillation frequency, signal offset. Some even offer a spectrum analyzer, which can be useful to check the frequencies contained in your signal.

MODULE ASSEMBLY APPENDIX

Before we start building, let's take a look at the complete **mki x es.edu dual VCA** schematics (see next page) that were used for the final module's design and PCB fabrication. You'll notice that we've squeezed not just one, but two VCAs into this module – while also adding a few components compared to the breadboard schematics. Most components on the production schematics have denominations (a name – like R1, C1, VT1, VD1, etc.) and values next to them. Denominations help identify each component on the PCB, which is particularly useful during **calibration, modification or troubleshooting.**

XS1 – XS6 are input and output **jack sockets** – these are the very same we've already been using on the breadboard for interfacing with other devices. In our designs, we use eurorack standard 3,5mm jack sockets (part number WQP-PJ301M-12). These sockets come with three lugs – a connector lug to which the tip of the patch cable is connected (and respectively, the audio or CV signal is applied), a ground lug which connects the patch cable to circuit ground and a switching lug. The **switching lug** is normally connected to the connector lug, but as soon you insert a patch cable into the socket, it gets disconnected from the connector lug. This is very handy for grounding inputs when nothing is patched into them. If the input is not grounded, some tiny current may bleed into the module, which can result in audible noise or CV fluctuations.

Another advantage of a switching lug is the possibility of daisy chaining inputs. In the case of the mki x es.edu dual VCA, if nothing is patched into IN2 (XS4), the audio signal from IN1 is routed to IN2, applying the same input to both VCAs. This is called **normalization** – IN2 is normalled to IN1. The same goes for the CV inputs – CV2 is normalled to CV1 and you can control both VCAs with one CV applied to the CV1 input.

XP1 is a standard eurorack **power connector**. It's a 2x5 male pin header with a key (the black plastic shroud around the pins) to prevent accidental reverse polarity power supply connection. This is necessary because connecting the power incorrectly will permanently damage the module.

VD1 and **VD2** are **schottky diodes** that double-secure the reverse polarity power supply protection. Diodes pass current only in one direction. Because the anode of VD1 is connected to +12 V on our power header, it'll only conduct if the connector is plugged in correctly. If a negative voltage is accidentally applied to the anode of VD1, it closes, and no current passes through. The same goes for VD2, which is connected to -12 V. Because schottky diodes have a low forward voltage drop, they are the most efficient choice for applications like this.

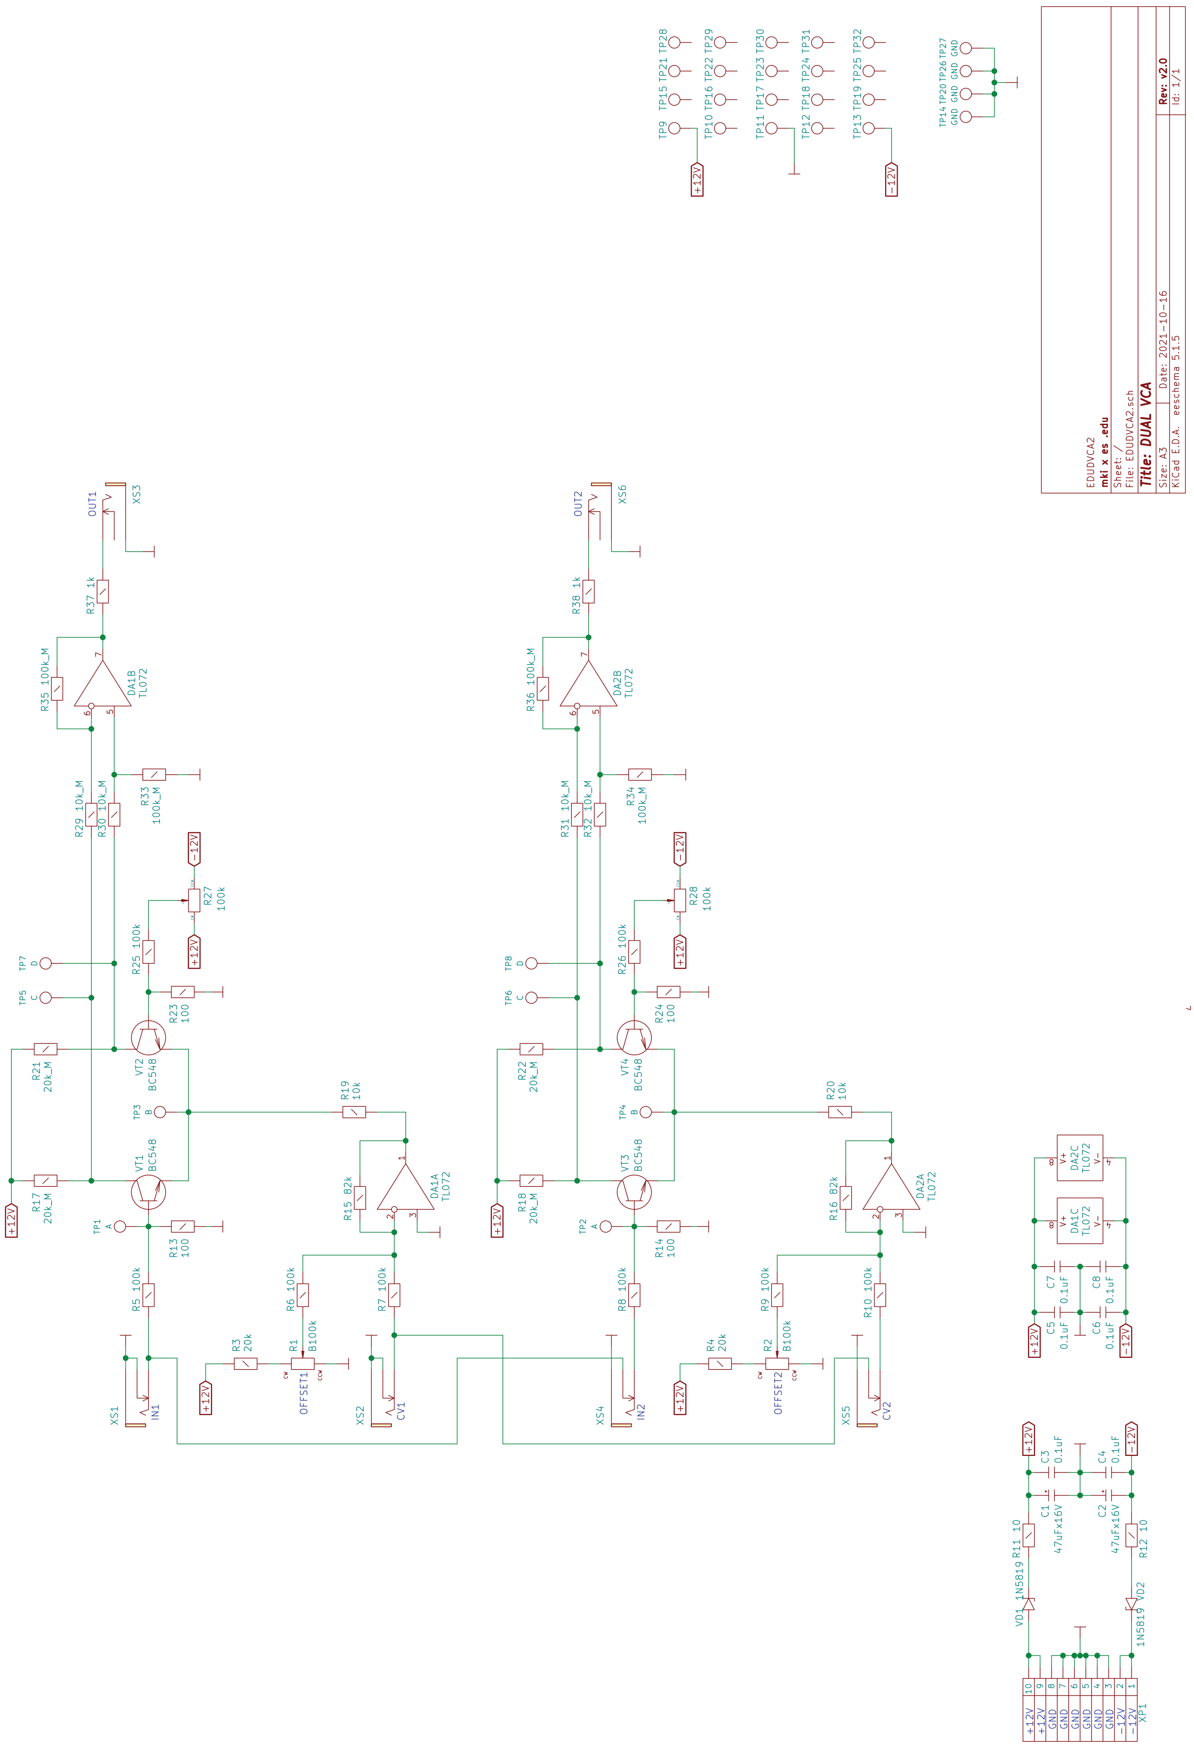
EDUDVCA2
mki x es.edu
Sheet: /
File: EDUDVCA2.sch
Title: DUAL VCA
Size: A3
Date: 2021-10-16
KiCad E.D.A. eeschema 5.1.5
Rev: v2.0
Id: 1/1

Next, we have two **10 Ohm resistors** (**R11** and **R12**) on the + and – 12 V rails, with **decoupling** (or **bypass-**) capacitors **C1 – C4**. These capacitors serve as energy reservoirs that keep the module's internal supply voltages stable in case there are any fluctuations in the power supply of the entire modular system. In combination with R11 and R12, the large 47 microfarads pair (C1 and C2) compensates for low frequency fluctuations, while C3 and C4 filter out radio frequencies, high frequency spikes from switching power supplies and quick spikes created by other modules. Often another component – a **ferrite bead** – is used instead of a 10 Ohm resistor, and there's no clear consensus among electronic designers which works best, but generally, for analogue modules that work mostly in the audio frequency range (as opposed to digital ones that use microcontrollers running at 8 MHz frequencies and above), resistors are considered to be superior.

Another advantage of 10 Ohm resistors is that they will act like **slow "fuses"** in case there's an accidental short circuit somewhere on the PCB, or an integrated circuit (IC) is inserted backwards into a DIP socket. The resistor will get hot, smoke, and finally break the connection. Even though they aren't really fuses, just having them there as fuse substitutes is pretty useful - **you'd rather lose a cent on a destroyed resistor than a few euros on destroyed ICs.**

Capacitors **C5 – C8** are additional decoupling capacitors. If you inspect the PCB, you'll see that these are placed as close to the power supply pins of the ICs as possible. For well-designed larger PCBs you will find decoupling capacitors next to each IC. Like the others, their job is to simply compensate for any unwanted noise in the supply rails. If the input voltage drops, then these capacitors will be able to bridge the gap to keep the voltage at the IC stable. And vice-versa - if the voltage increases, then they'll be able to absorb the excess energy trying to flow through to the IC, which again keeps the voltage stable. Typically, 0.1 uF capacitors are used for this purpose.

Before you start soldering, we highly recommend printing out these part placement diagrams with designators and values. Because some of our PCBs are rather densely populated, this will help you to avoid mistakes in the build process.

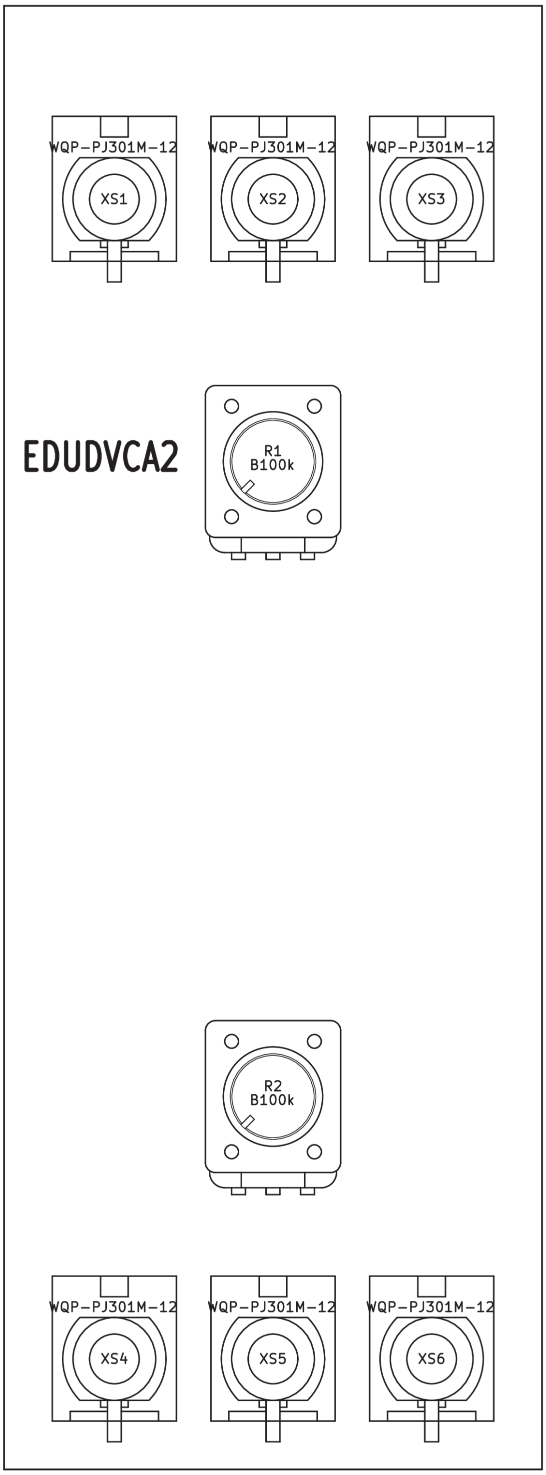
WQP–PJ301M–12
XS1
XS2
XS3
EDUDVCA2
R1
B100k
R2
B100k
XS4
XS5
XS6

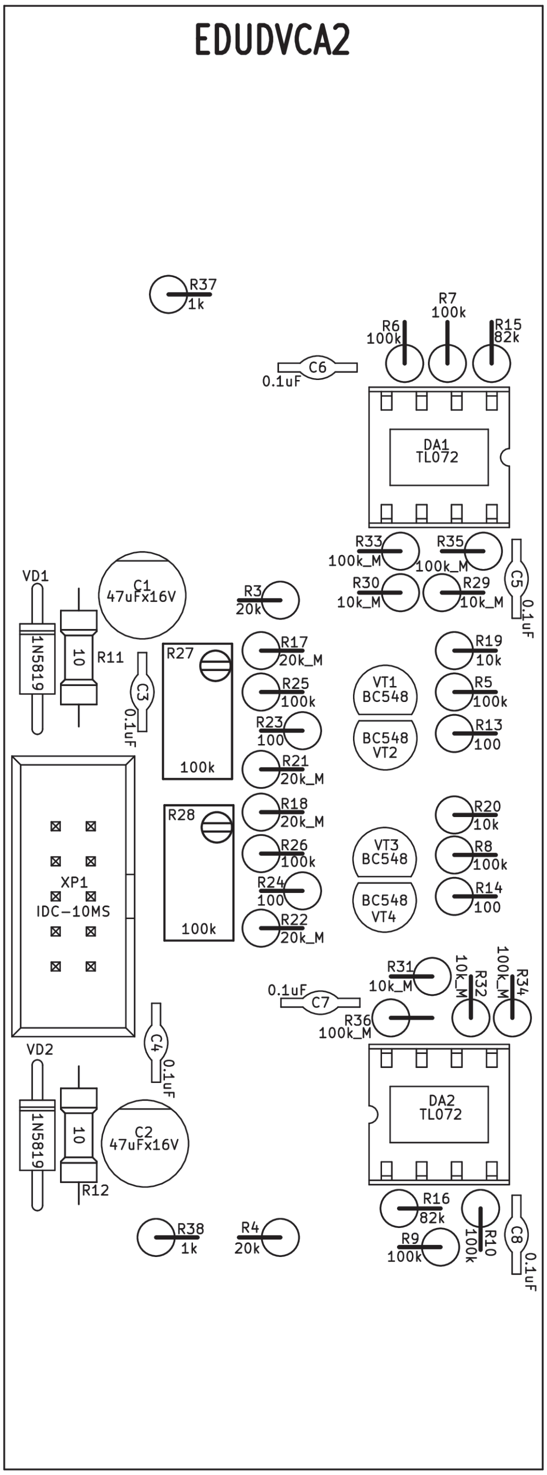
EDUDVCA2
R37
1k
R6
100k
R7
100k
R15
82k
C6
0.1uF
DA1
TL072
R33
100k_M
R35
100k_M
C5
0.1uF
R30
10k_M
R29
10k_M
VD1
1N5819
C1
47uFx16V
R3
20k
10
R11
C3
0.1uF
R27
100k
R17
20k_M
R25
100k
R23
100
R21
20k_M
R19
10k
R5
100k
R13
100
VT1
BC548
VT2
XP1
IDC–10MS
R28
100k
R18
20k_M
R26
100k
R24
100
R22
20k_M
VT3
VT4
R20
10k
R8
100k
R14
100
R31
10k_M
R32
10k_M
R34
100k_M
C7
0.1uF
R36
100k_M
C4
0.1uF
VD2
R12
C2
47uFx16V
DA2
TL072
R16
82k
R9
100k
R10
100k
C8
0.1uF
R38
1k
R4
20k

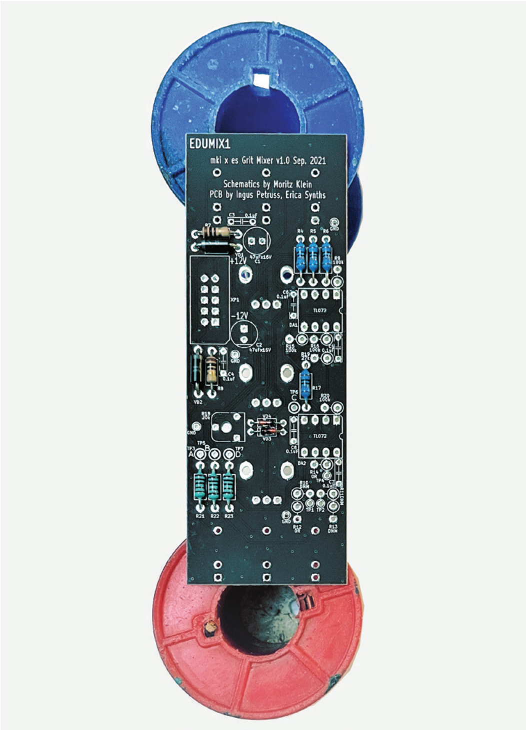

Place the dual VCA PCB in a PCB holder for soldering or simply on top of some spacers (I use two empty solder wire coils here).

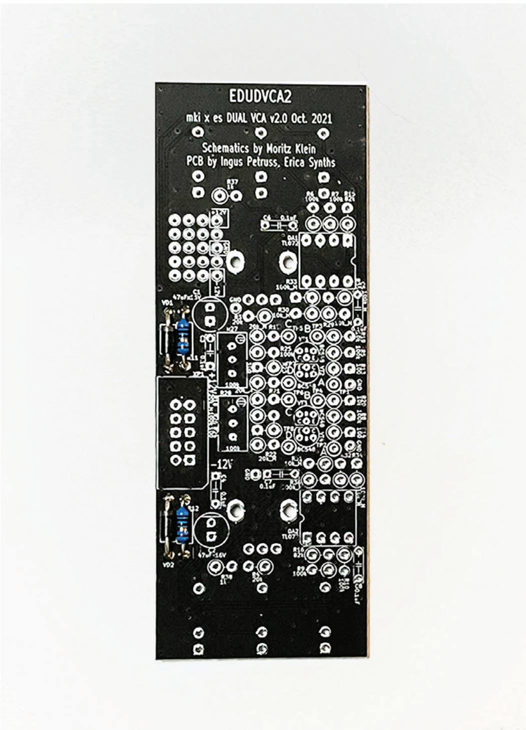

I usually start populating PCBs with lower, horizontally placed components. In this case, these are **R11**, **R12** and the **power protection diodes**. Bend the resistor leads and insert them in the relevant places according to the part placement diagram above. All components on the PCB have both their value and denomination printed onto the silkscreen. If you are not sure about a resistor's value, use a multimeter to double check. Next, insert the diodes. Remember – **when inserting the diodes, orientation is critical!** A thick white stripe on the PCB indicates the cathode of a diode – match it with a stripe on the component. Flip the PCB over and solder all four components. Then, use pliers to cut off the excess leads.

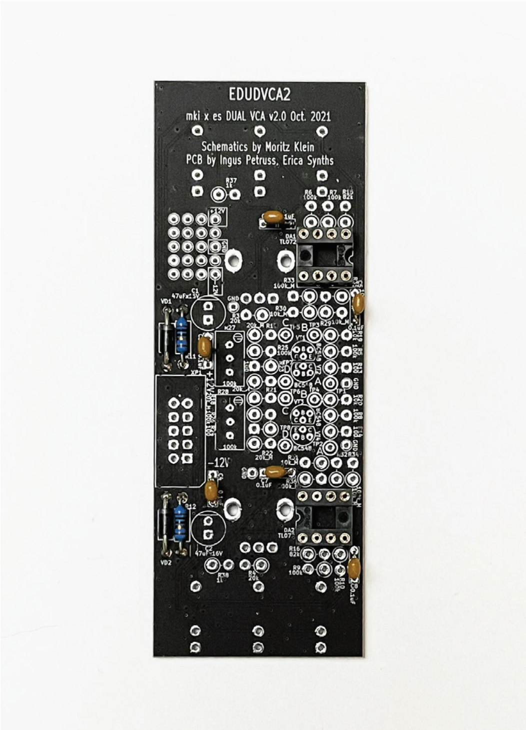

Next, insert the first DIP socket, hold it in place and solder one of the pins. Continue with the next DIP socket. Make sure the DIP sockets are oriented correctly – the notch on the socket should match the notch on the PCB's silkscreen. Now, turn the PCB around and solder all remaining pins of the DIP sockets. Then proceed with the ceramic capacitors. Place the PCB in your PCB holder or on spacers, insert the capacitors and solder them like you did with the resistors & diodes before. Now, your PCB should look like this.

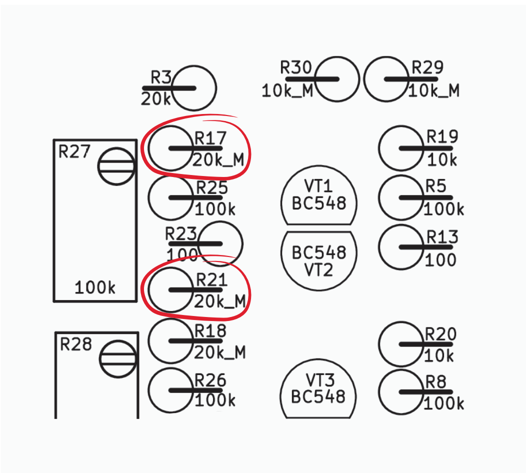

Now take a close look at the part placement diagram. Some resistors are marked with an M next to their value. **These should be matched in pairs with as similar resistances as possible.** Normally, you'd need a bunch of, say, 20 kOhm resistors for this. You would check each one with your multimeter and select pairs with matching resistances. They don't have to all be 20 kOhm exactly – it's just important that the individual values are matching. Two resistors with 19.98 kOhm resistance, for example, would be perfectly okay.

But as this is a pre-assembled kit and we didn't want to just drop a big pile of resistors on your head, we pack it with **±0.1% tolerance resistors** where needed. Which means that there's no real need for matching – they're going to be close enough straight out of the box. You can identify these **±0.1% tolerance resistors by their light blue bodies** – use them wherever you see an M on the part placement diagram. If you're unsure about a resistor's value, use your multimeter to double-check.

In order to save space on the PCB, some of our projects, including the dual VCA, have **vertically placed resistors.** So the next step is to place & solder those. Bend a resistor's legs so that its body is aligned with both legs and insert it in its designated spot. Then solder the longer lead from the top side of the PCB to secure it in place, turn the PCB around and solder the other lead from the bottom. You can insert several resistors at once. Once done with soldering, use pliers to cut off excess leads.

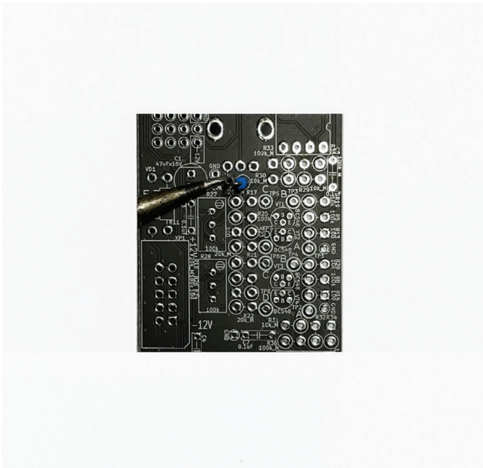

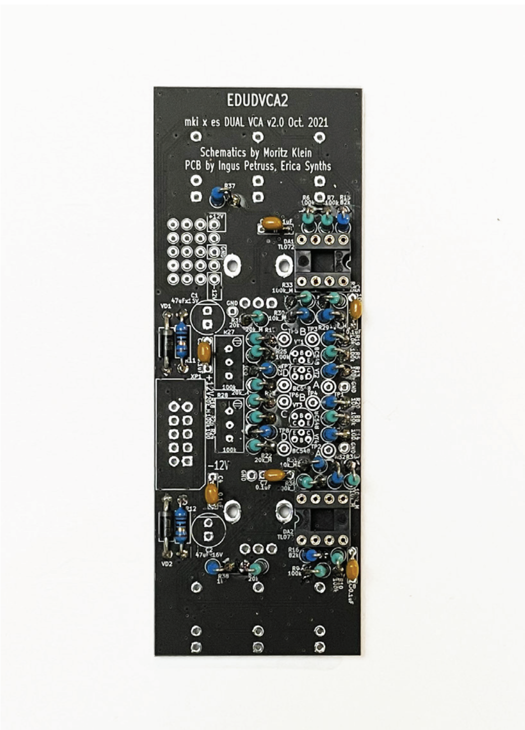

Once you are done with soldering all resistors, your PCB should look like this.

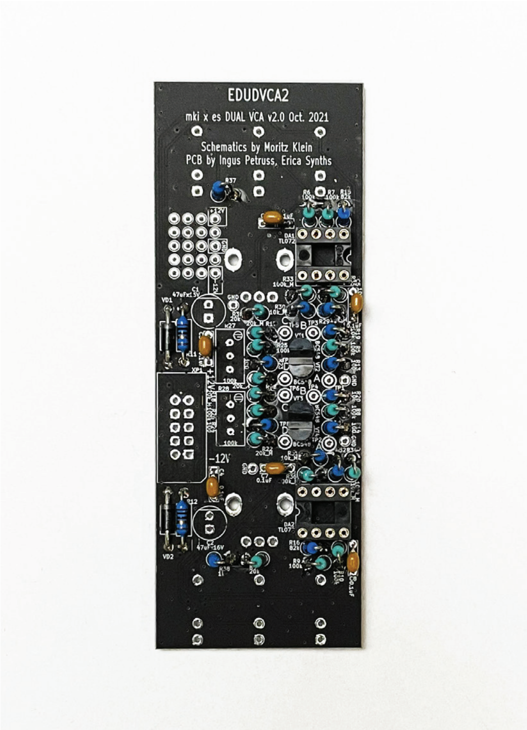

Next up: inserting & soldering the transistors. Make sure you place the transistors in their designated spots. Also, they need to be properly aligned with the marked outline on the silkscreen – orientation is critically important here.

Next, insert & solder the electrolytic capacitors. Electrolytic capacitors are bipolar, and you need to mind their orientation. The positive lead of each electrolytic capacitor is longer, and there's a minus stripe on the side of the capacitor's body to indicate the negative lead. On our PCBs, the positive pad for a capacitor has a square shape, and the negative lead should go into the pad next to the notch on the silkscreen.

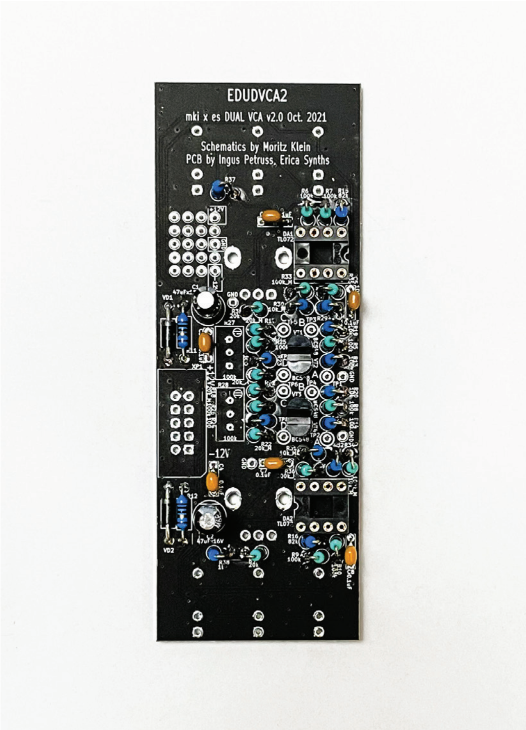

Now your PCB should look like this.

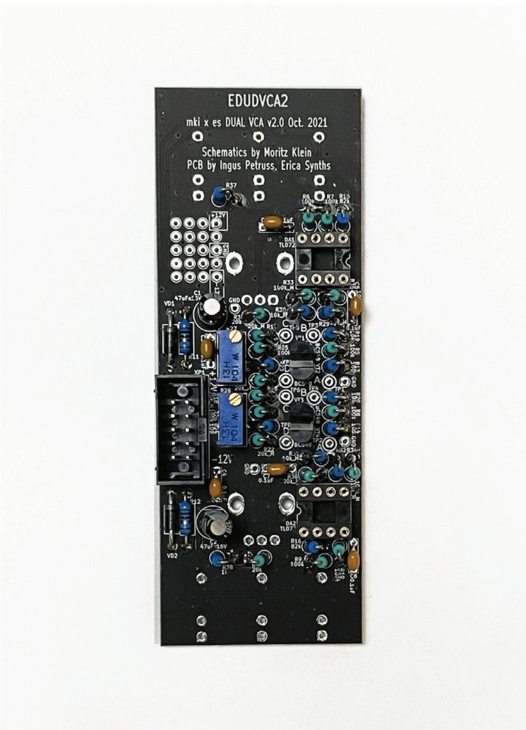

Next, solder the precision trimmer and the power supply socket. Make sure the orientation of the socket is as shown in the picture below – the arrow pointing to the first pin is aligned with a notch on the silkscreen. The key on the socket will be facing inwards to the PCB.

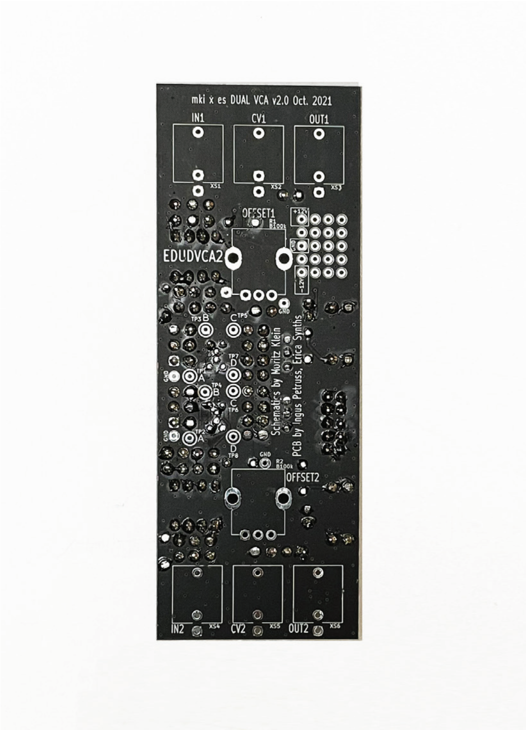

Now, turn the PCB around and inspect your solder joints. **Make sure all components are soldered properly and there are no cold solder joints or accidental shorts**. Clean the PCB to remove extra flux, if necessary.

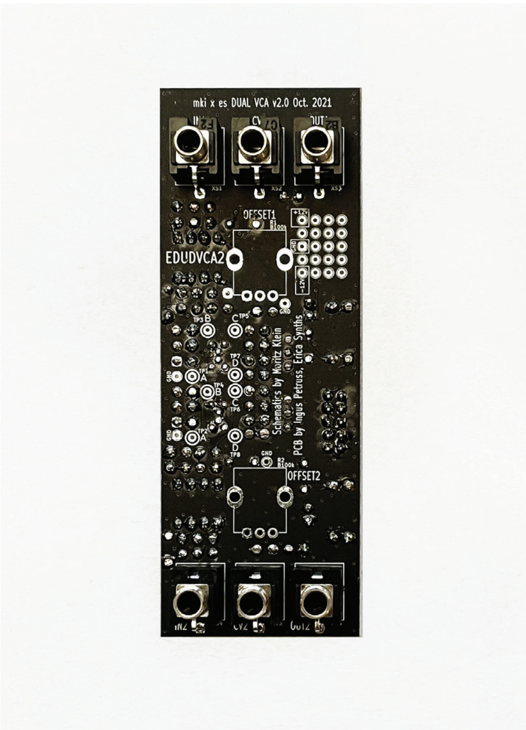

Insert the jack sockets and solder them.

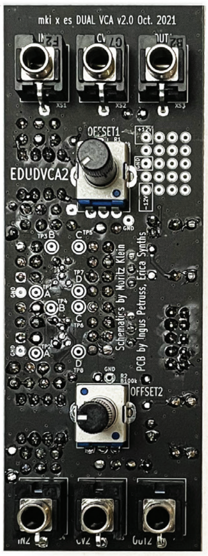

Insert the potentiometers, but don't solder them yet! Fit the front panel and make sure that the potentiometers' shafts are aligned with the holes in the panel – and that they're able to rotate freely. Now, go ahead and solder the potentiometers.

Install the front panel and fix it in place with the 6 hex nuts on the jack sockets. We are almost done!

Now, **insert the ICs into their respective DIP sockets.** Mind the orientation of the ICs – match the notch on each IC with the one on its socket.

Congratulations! **You have completed the assembly of the mki x es.edu dual VCA module!** Now connect it to your eurorack power supply and switch it on. If there's no "magic smoke", it's a good sign that your build was successful. Connect a VCO's SAW output to VCA IN1, and VCA OUT1 to a mixer (connected to monitors or headphones). Rotate the OFFSET1 knob all the way clockwise and check if you hear the oscillator. Then, connect OUT2 to the mixer, turn OFFSET2 all way clockwise and again check if you hear the oscillator. If you do in both cases, you can proceed with calibrating the VCA as described in the **ELIMINATING TRANSISTOR MISMATCH**-section above.

SOLDERING APPENDIX

If you’ve never soldered before – or if your skills have become rusty – it’s probably wise to check out some **THT** (through-hole technology) **soldering tutorials on YouTube**. The main thing you have to remember while soldering is that melted solder will flow towards higher temperature areas. So you need to make sure you apply equal heat to the component you are soldering and the solder pad on the PCB. The pad will typically absorb more heat (especially ground-connected pads which have more thermal mass), so keep your soldering iron closer to the pad on the PCB. It’s critically important to dial in the right temperature on your soldering station. I found that about 320 °C is the optimal temperature for most of parts, while for larger elements like potentiometers and sockets, you may want to increase that temperature to **370 °C.**

Here’s the recommended soldering sequence:

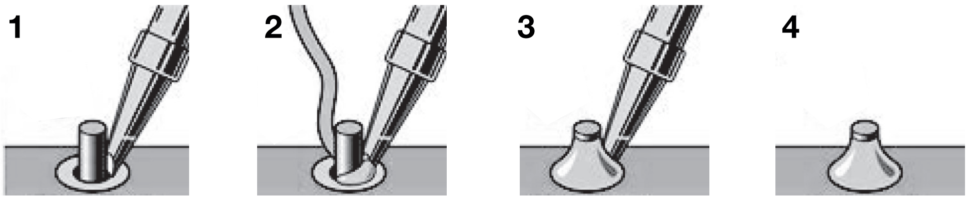

Heat part and pad 2 - 3 sec | Add solder | Continue heating 1 -2 sec. | Let cool

After you have completed soldering, inspect the solder joint:

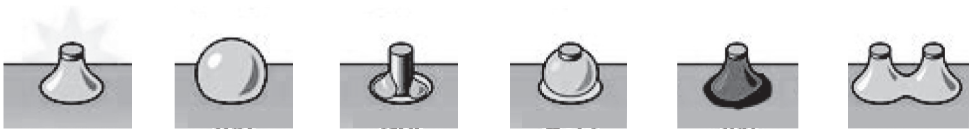

Perfect | Too much solder | Not enough solder | Cold joint | Too much heat | Short

DIY electronics is a great (and quite addictive) hobby, therefore we highly recommend you invest in good tools. In order to really enjoy soldering, you'll need:

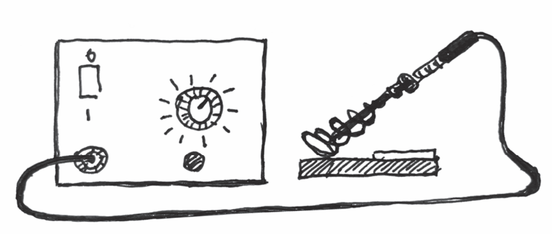

A decent soldering station. Top-of-the-line soldering stations (brands like Weller) will cost 200€ and above, but cheaper alternatives around 50€ are often good enough. Make sure your soldering station of choice comes with multiple differently-sized soldering iron tips. The most useful ones for DIY electronics are flat, 2mm wide tips.

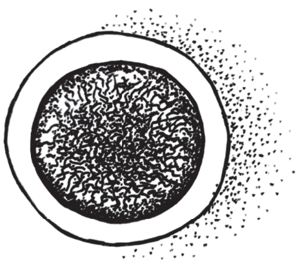

When heated up, the tips of soldering irons tend to oxidize. As a result, solder won't stick to them, so you'll need to clean your tip frequently. Most soldering stations come with a **damp sponge for cleaning the iron tips** – but there are also professional solder tip cleaners with **golden curls** (not really gold, so not as expensive as it sounds). These work much better because they do not cool down the iron.

Solder wire with flux. I find 0,7mm solder wire works best for DIY projects.

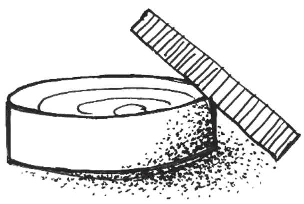

Some **soldering flux** paste or pen will be useful as well.

Cutting pliers. Use them to cut off excess component leads after soldering.

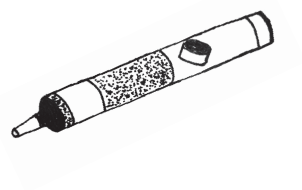

A solder suction pump. No matter how refined your soldering skills are, you will make mistakes. So when you'll inevitably need to de-solder components, you will also need to remove any remaining solder from the solder pads in order to insert new components.

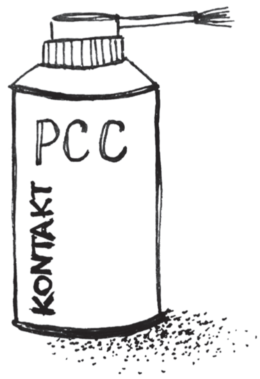

Once you have finished soldering your PCB, it's recommended to remove excess flux from the solder joints. **A PCB cleaner** is the best way to go.

All of these tools can be found on major electronic components retailer websites, like Mouser, Farnell and at your local electronics shops. As you work your way towards more and more advanced projects, you'll need to expand your skillset and your tool belt – but the gratification will be much greater.

"But you have to give your whole life to a cello. When I realized that, I went back to the guitar and just turned the volume up a bit louder."

- Ritchie Blackmore