

Building a noisebox

Experiments in audio synthesis

Nomen Nescio

May 31, 2019

Project: noisebox
Version: 0.3

Contents

1	A bit of theory	3
1.1	Sound	3
1.2	Oscillators	4
1.3	Opamps	6
2	A bit of praxis	10
2.1	The circuit	10
2.1.1	Legenda	11
2.2	Tools for the trade	12
2.3	Building instructions	13
3	Crisis and experiments	14
3.1	Fixing a noisebox	14
3.2	Interventions	14

1 A bit of theory

1.1 Sound

The experience of hearing sound comes from our ears detecting vibrations in the air. To create such vibrations in the air, we use a speaker. As we feed the speaker with a voltage that goes up and down, the speaker moves back and forth accordingly, causing the air to vibrate with it. What this sound sounds like is determined by many factors, but one of the most important is its fundamental frequency. A sound usually consists of many frequencies, with a dominant fundamental frequency that is perceived as the sound's pitch. Hitting the fifth A key on a piano will, if the piano is tuned well, produce a tone with a fundamental frequency of 440 Hz., meaning it vibrates up and down 440 times per second. Higher tones vibrate faster, they have a higher frequency, and lower tones have a lower frequency, vibrating slower. One of the key objectives in audio synthesis is thus to create an electric signal with a voltage that swings up and down at a specific frequency. This is the job of the oscillator and it is the first out of three major building blocks in our noisebox. In fact, our noisebox consists of two oscillators: one whose frequency can be adjusted with a slider and whose frequency can be controlled by allowing more or less light to shine on the sensor.

But fundamental frequency is far from the only factor determining what a sound sounds like. Another factor is the shape along which the voltage goes up and down. In our noisebox, we will mostly be using squarewaves, but we'll also encounter triangle waves, and if you follow the suggestions for modifications at the end of this document, you'll find you can also synthesise sawtooth shapes. I won't go into detail on waveshapes in this document, but if you're interested in how waveshapes relate to frequencies, search for overtones and find out how the different shapes can be seen as combinations of sinewaves at different frequencies...

Another aspect of sound is how loud it is. To make our noisebox scream and squeek loud or more silently, there is a second light sensor that can be used to control the volume of the signal coming from the oscillators. This light sensor — in fact a light dependent resistor — is part of the second major building block of our noisebox: it is connected to a so called opamp, which can reduce or amplify the difference in voltage between the low and the high point of our waves, thereby controlling the volume of the outgoing signal.

Sadly, the signal that comes out of our opamp isn't quite strong enough to drive the speaker, so the third and last building block of our noisebox is the amplifier: a chip that takes an incoming signal and provides the power needed for that signal to make the speaker move. To prevent this document from becoming a book, we'll leave this part as a black box.

For the first two parts — the oscillator and the opamp — I'll now give a rudimentary explanation, trying my very best not to burden you with the theory of electronic engineering, but focussing rather on the concepts that make it work.

1.2 Oscillators

Imagine if you will an irrigation system consisting of a water sensor and a tap. When the sensor is dry, the tap is turned on and water starts to flow. When the sensor is wet, the tap is turned off again.

Now, imagine the tap being connected first to a valve, and then to a balloon with a little hole at the bottom. The sensor is then placed below the balloon. What will happen with the waterflow?




Figure 1: A mechanical water oscillator

At first, the sensor is dry, so the tap will go on and water will start to flow to the valve. Depending on the position of the valve, a certain amount of water will flow through into the balloon. But once it reaches the balloon, it does not immediately fall down onto the sensor — first, the water fills up the balloon and only then does the water start dripping down onto the sensor. Once this happens, the sensor slowly gets wet enough to trigger the tap to go off again. With the tap off, no more water is coming into the balloon, but this does not mean that water immediately stops dripping onto the sensor. Rather, the balloon will slowly contract and keep dripping water down onto the sensor until it is empty. Only then does the sensor dry up and trigger the tap to go back on again, repeating the cycle *ad infinitum*.

So what does this mean for the waterflow? There are two points of interest for this question: 1) the waterflow on the sensor, and 2) the waterflow just after the tap.

If we look at the sensor, the water level keeps slowly rising and falling around the point where the sensor is triggered, hence in line with the filling and emptying of the balloon. Plotted over time, this looks like a triangle-shape figure, just like the one at the top of figure 2.




Figure 2: Water levels: triangle and square wave

If we look at the point just after the tap, we'll see water rushing through as the tap is opened, only to abruptly stop as the sensor closes the tap again. Plotted over time, this looks like a square wave, just like the one at the bottom of figure 2.

The frequency of these waves — how fast the cycle of the tap turning off, on, and back off again happens over time — is determined by both the position of the valve and the size of the balloon. The larger the balloon, the longer it takes to fill up, and thus the longer it takes both before water starts sufficiently dripping on the sensor and before the water stops dripping after the tap is closed. The result being the cycle taking longer and thus the frequency being lower. Vice versa, a smaller balloon fills up faster, but also empties faster, resulting in a shorter cycle and thus a higher frequency.

Changing the balloon, however, is a rather impractical way of adjusting the frequency. That's why we also built in a valve. The more the valve is closed, the longer it takes for the balloon to fill up, hence the longer it takes for the water to drip down on the sensor, and thus the longer the full cycle takes, resulting in a lower frequency. As the valve is opened, the balloon fills up faster, and the frequency is increased.

Hooray, we now have a water oscillator with adjustable frequency! But wasn't this an electronics project? Right, and luckily, the electronic oscillator we are going to build works very similar.

The sensor and the tap together are replaced by an inverter: a chip that reads an incoming voltage and inverts it: if the incoming voltage is low, the outgoing voltage is set to the maximum, if the incoming voltage is high, the outgoing voltage is set to the minimum. The inverter is represented in schematic 3 by a triangle with a small circle at the tip. The left side being the input (the sensor, in our water example) and the right side with the small circle is the output (the tap, in our water example).

Just like in our water example, the output of the inverter is connected to a variable resistor, also known as a potentiometer. This resistor functions very much like the valve. This then connects to a capacitor, functioning like the balloon, and to the input of the inverter (the sensor). You'll notice the analogy of the balloon doesn't work completely to




Figure 3: The oscillator

explain the capacitor: the electricity doesn't go through the capacitor into the inverter like the water goes through the balloon onto the sensor. Rather, the capacitor has one pole connected to both the resistor and the inverter, swallowing up, buffering, and letting go of electricity much like the balloon, while the other pole is connected to ground. Not exactly the same, but close enough to get the general idea.

As mentioned before, our noisebox contains two oscillators: one controlled by a slider and one by light. They can be played separately, each creating their own frequency, but the fun (and the noise!) really starts when you combine them in a process called cross-modulation. When we put a capacitor inbetween the output of the inverter of oscillator 1 and the input of oscillator 2, the result is that the cycle of going up and down in oscillator 2 will be interrupted periodically as oscillator 1 goes up. Imagine two water oscillators next to eachother, where the water coming out of the tap of the first also sprays on the sensor of the second, and try to picture what the waterflows will look like then. What follows is an unstable frequency, causing a rather harsh sound that should appeal to the more dissonantly inclined.

1.3 Opamps

The opamp is a chip that has two inputs and one output, its representation in schematics can be seen in figure 4. You'll notice one input marked with a - and the other with a +, these are called the inverting and the non-inverting input.




Figure 4: An opamp

What an opamp does is compare the voltage between the inverting and the non-inverting

input. This difference is then amplified massively (usually by a factor over 20000) and sent to the output. In practice, this means that when the non-inverting input (the +) has a higher incoming voltage than the inverting input (the -), the output will be the maximum voltage (in our case, the 9 volt coming from the battery of the noisebox). Vice versa, if the inverting input (the -) has a higher input voltage, the output will be the minimum voltage (in our case, the 0 volt or ground from the battery).

Let's try the water analogy explain to explain the opamp. Imagine the contraption depicted in figure 5. If the water pressure on the non-inverting terminal is higher than that on the inverting terminal, the lever will be pushed down, causing the compressor to be connected to the output and water bursting out. Vice versa, if the water pressure on the inverting side is higher, the lever will be pushed up, causing the output to be connected to the vacuum pump, sucking up the water from the output.




Figure 5: A water opamp¹

Now that is all nice and well for comparing two signals, but what we want to do is attenuate or amplify just one signal — we want to control the volume from the output of our oscillators. To achieve this, we start by introducing feedback, as you can see in figure 6. With the output connected to the inverting input, the opamp now has a tendency towards keeping the lever arm in balance.

As soon as the water pressure at the inverting terminal is lower than at the non-inverting, the lever moves down, letting water from the compressor through to the output, which in turn streams back to the inverting input, increasing the water pressure there until it is the same as the pressure on the non-inverting input and the compressor is cut off from the output again. Vice versa, if the water pressure at the inverting side is higher, the lever is pushed up, connecting the vacuum pump to the output, sucking the water pressure away from the inverting input, hence moving the lever back down until it is back in a stable position. Note that a lower water pressure at the inverting input causes water to gush out at the output, whereas a higher water pressure at the inverting input causes the vacuum to suck from the output. The water pressure at the output becomes the exact opposite of

¹image stolen from some forumpost, origin unknown




Figure 6: Inverting water buffer

the water pressure at the input. In electronics, this is called an inverting buffer. Well great, buffers are nice, but we're still not attenuating or amplifying! To accomplish that, we need to introduce two more elements. Imagine two valves, one placed before the inverting input, and one placed in the middle of the feedback loop, as can be seen in figure 7. Imagine the first valve being almost fully open, and the second valve (the one in the middle of the feedback loop) being halfway closed. This time, when the water pressure arriving is higher than the pressure at the non-inverting side, the lever will go up, and the output will start sucking water away, but this time it will have a harder time decreasing the water pressure at the inverting input, as there is a valve in the way.

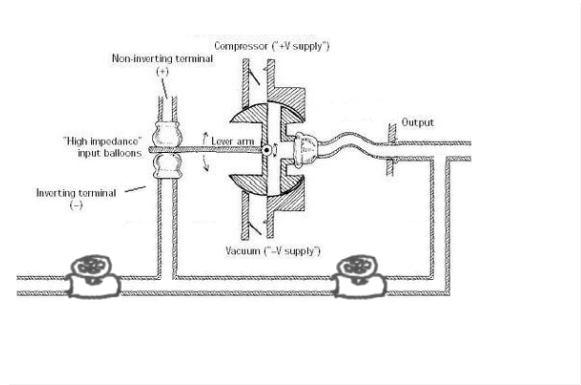


Figure 7: Inverting water amplifier

The pressure at the output will now not just be the inverse of the input, it will be even less. Vice versa, when the pressure at the input is low, the lever will go down, and water will gush out of the output, but it will have to stream out quite a bit harder than usual because of the valve being in the way. The pressure at the output is then not just the

inverse of the input, but even more. How much more (or less) depends on the position of the two valves. The further the valve in the feedback line is closed, or the further the valve at the input is opened, the more the output will be not just inverted, but also amplified. The further the valve in the feedback line is opened, the more the input valve is closed, the more the output will not only be inverted, but also attenuated.

Bringing the analogy back to the realm of electronics, we substitute the valves for resistors and as can be seen in figure 8 we now have our inverting amplifier, where the amount of attenuation c.q. amplification of our signal (that is, the gain of the amplifier) is set by the ratio between the two resistors.




Figure 8: Inverting amplifier

In our noisebox, the first resistor is a light dependent resistor, a sensor that changes value depending on how much light comes in. The second resistor (the one in the feedback path) is adjustable by a trimmer, to make the noisebox operable under different light conditions. The non-inverting input is set to 4,5 volt, which is half of the 9 volt that comes from our battery and should be the middle point of the wave that comes from our oscillator. Like this, the opamp will seek stability and attenuate by contracting the signal nicely towards its midpoint.

2 A bit of praxis

Now that we have covered the most important concepts of the noisebox, the oscillator and the inverting amplifier, it's time to put things into practice and build a box! This section starts of with the schematics of the noisebox, this can be your reference point throughout the whole process of soldering and connecting components. What follows is a list of all the components, the tools you will need, and then a step by step guide to building the noisebox.

2.1 The circuit

Figure 9 shows the complete schematic of the noisebox. Note that the numbers next to the inverters, opamp, and the lm386 designate the pin numbers on the chip. Switches SW1 and SW2 are combined into a 4-position switch.

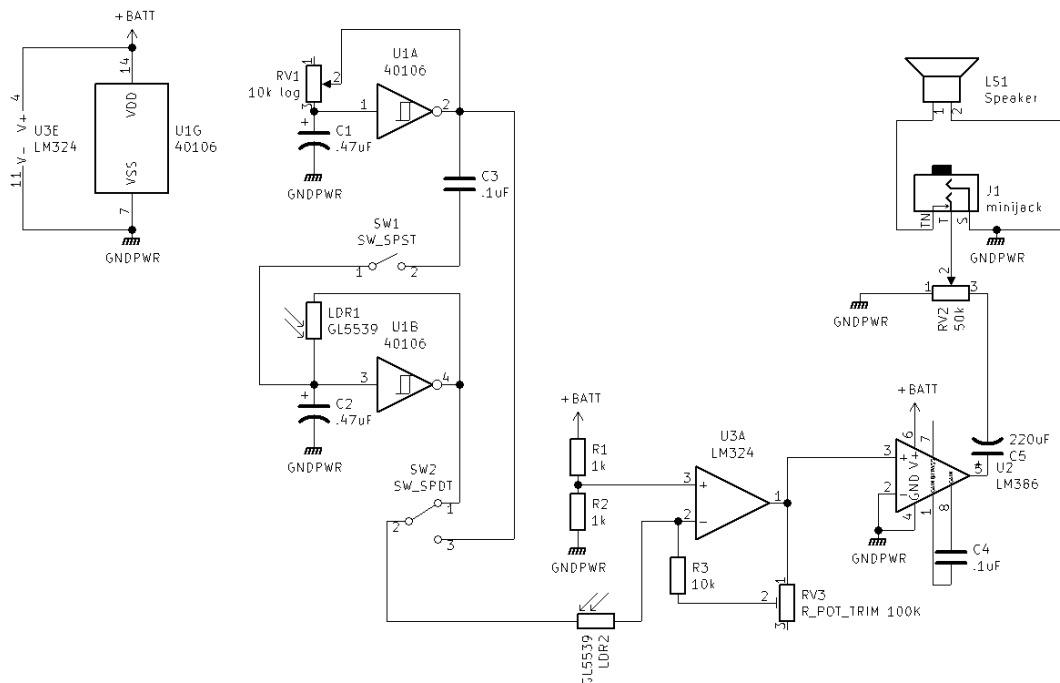













Figure 9: Noisebox v0.3 schematics

2.1.1 Legenda

Making sense of the schematic is easier when you know what the symbols mean. Here's a list to help you:

Symbol	Meaning	Description
	Power	Positive side of the battery
	Ground	Negative side of the battery
	Capacitor	Little yellow knob with two poles
	Electrolytic Capacitor	Cylinder with two poles at the bottom, one side marked as negative, corresponding to the curved side of the symbol
	Light Dependent Resistor (LDR)	Little plate with lines on top and two poles
	Inverter	A chip, in our case the 40106
	Amplifier	A chip, in our case the LM386
	Opamp	A chip, in our case the LM324
	Resistor	Thin cylinder with a pole at each side, marked with coloured rings. Brown/Black/Red = 1kΩ (R1 & R2), Brown/Black/Orange = 10kΩ (R3)
	Potentiometer	In our case, RV1 is the slider, RV2 is the turnknob
	Trimmer	A little turnknob with a slot for a screw-driver

Apart from these items, the noisebox also consists of: a minijack plug, a speaker, 4 pieces of 7x9cm multiplex wood, 1 piece of 7x9cm triplex, a 7x9cm protoboard, 2 9cm strips of wood, 4 insert nuts, 4 bolts, a metal battery clamp, a clip for the battery, and a bunch of wire.

Optional components for modifications are diodes, a 10k trimmer, more resistors and capacitors, and perhaps an NPN transistor if you really want to go wild.

2.2 Tools for the trade

The following tools are indispensable for building the noisebox:

- Soldering iron
- Soldering tin
- Solder sucker
- Wire cutter
- Wire stripper
- Screwdriver
- Drill
- Hot glue gun
- Wood glue

You'll probably also want:

- Multimeter
- Datasheets²³⁴
- Breadboard

And these will make building and experimenting a lot easier:

- Oscilloscope⁵
- Bench top power supply
- External amplifier and speaker/headphones

²40106 datasheet: <http://www.ti.com/lit/ds/symlink/cd40106b.pdf>

³LM324 datasheet: <http://www.ti.com/lit/ds/symlink/lm324-n.pdf>

⁴LM386 datasheet: <http://www.ti.com/lit/ds/symlink/lm386.pdf>

⁵You can use your soundcard as oscilloscope, see: <http://xoscope.sourceforge.net/>

2.3 Building instructions

Sadly, time did not allow me to write up elaborate building instructions. Here is a very rough first version of the instructions. These instructions will be complemented and amended in later versions based on the experience and feedback from the noisebox workshops.

1. Start by drilling a 35mm hole in one of the multiplex pieces, solder two wires onto the speaker and glue the speaker onto the hole.
2. Drill two holes in the triplex board: one for the minijack and one for the potmeter, attach both.
3. Screw the battery clamp onto another multiplex board.
4. Solder a wire from the black wire of the battery clip to the negative pole (-) of the speaker, continue to the ground of the minijack, on to the ground terminal of the potmeter, and end with a fairly long black piece of wire to eventually connect to the protoboard.
5. Connect the positive pole (+) of the speaker to the normalisation pin of the minijack.
6. Connect the input of the minijack to the middle terminal of the potmeter.
7. Connect the red wire of the battery clip to one side of the switch of the potmeter.
8. Connect a fairly long piece of red wire to the other side of the switch of the potmeter.
9. Connect a fairly long piece of green wire to the input terminal of the potmeter.
10. Now, screw and glue the multiplex and triplex pieces together to form a box.
11. Solder all the components onto the protoboard, following the schematic in figure 9
12. Finish by soldering the red, black, and green wires from the box to the power, ground, and signal output positions on the bottom of the protoboard.
13. Fix the protoboard onto the two 9cm wooden strips by drilling and screwing in the insert nuts and attaching the board with bolts.
14. Screw/glue the wooden strips into the box.
15. Unscrew the bolts from the protoboard, insert a battery, screw the protoboard back on and make some noise!

3 Crisis and experiments

3.1 Fixing a noisebox

The noisebox usually don't immediately work after trying to build them. Usually there is some error in the soldering somewhere, maybe a loose contact, perhaps a shortcircuit somewhere. The idea behind this section is to collect common errors, and document ways to find and fix them. This section will be written later based on the experience and feedback from the noisebox workshops.

3.2 Interventions

The noisebox deliberately comes with protoboard and not a ready-made circuitboard where you just have to solder in components like most kits have. This is in part for educational purposes, having to puzzle out yourself how to place and connect the components to match the schematic forces you to think about the schematic and hopefully question why certain parts are connected. It also offers options for modification and experimentation, you are free to add or replace components as you see fit and change the design of the noisebox to fit your interests. Here are some suggestions for alterations you could make:

- Make an inverting amplifier from three of the spare pins on the LM324 and attach the input pin of one of the 40106 oscillators to the inverting input of the opamp. If you set the gain high enough and connect the output to the switch (or directly to the LM386), you should now be able to hear a triangle wave instead of a square.
- Place a diode between the input and output pins of one of the 40106 oscillators, attach the input pin to an inverting amplifier or hook up and oscilloscope to hear/see the waveshape become a sawtooth.
- Notice how pulling the slider to the far end creates sounds at a frequency too high to hear? Place a trimmer inbetween the slider and the inverter to be able to adjust the highest frequency and dial it back down to an audible level.
- Want to play actual melodies with the slider but having a hard time with such a small slider for so many octaves? Place a 10k Ω resistor in parallel with the slider to cut its frequency range in half.