Castle Rocktronics

002 – Output Mixer

Four channel mixer with built in overdrive and output attenuation

Comments, suggestions, questions and corrections are welcomed & encouraged:
contact@castlerocktronics.com

Table of Contents

Version history:

01/November/2014 – First draft

07/Apr/2015 – Minor corrections (credit to Suzy Creamcheeze) and a bigger correction to the decibel section where some of the equations were wrong :/

0. Build documents

0.1 Photos

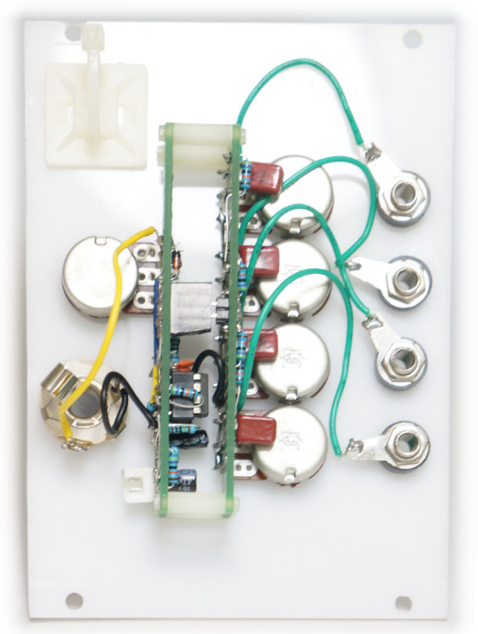

0.2 Layout

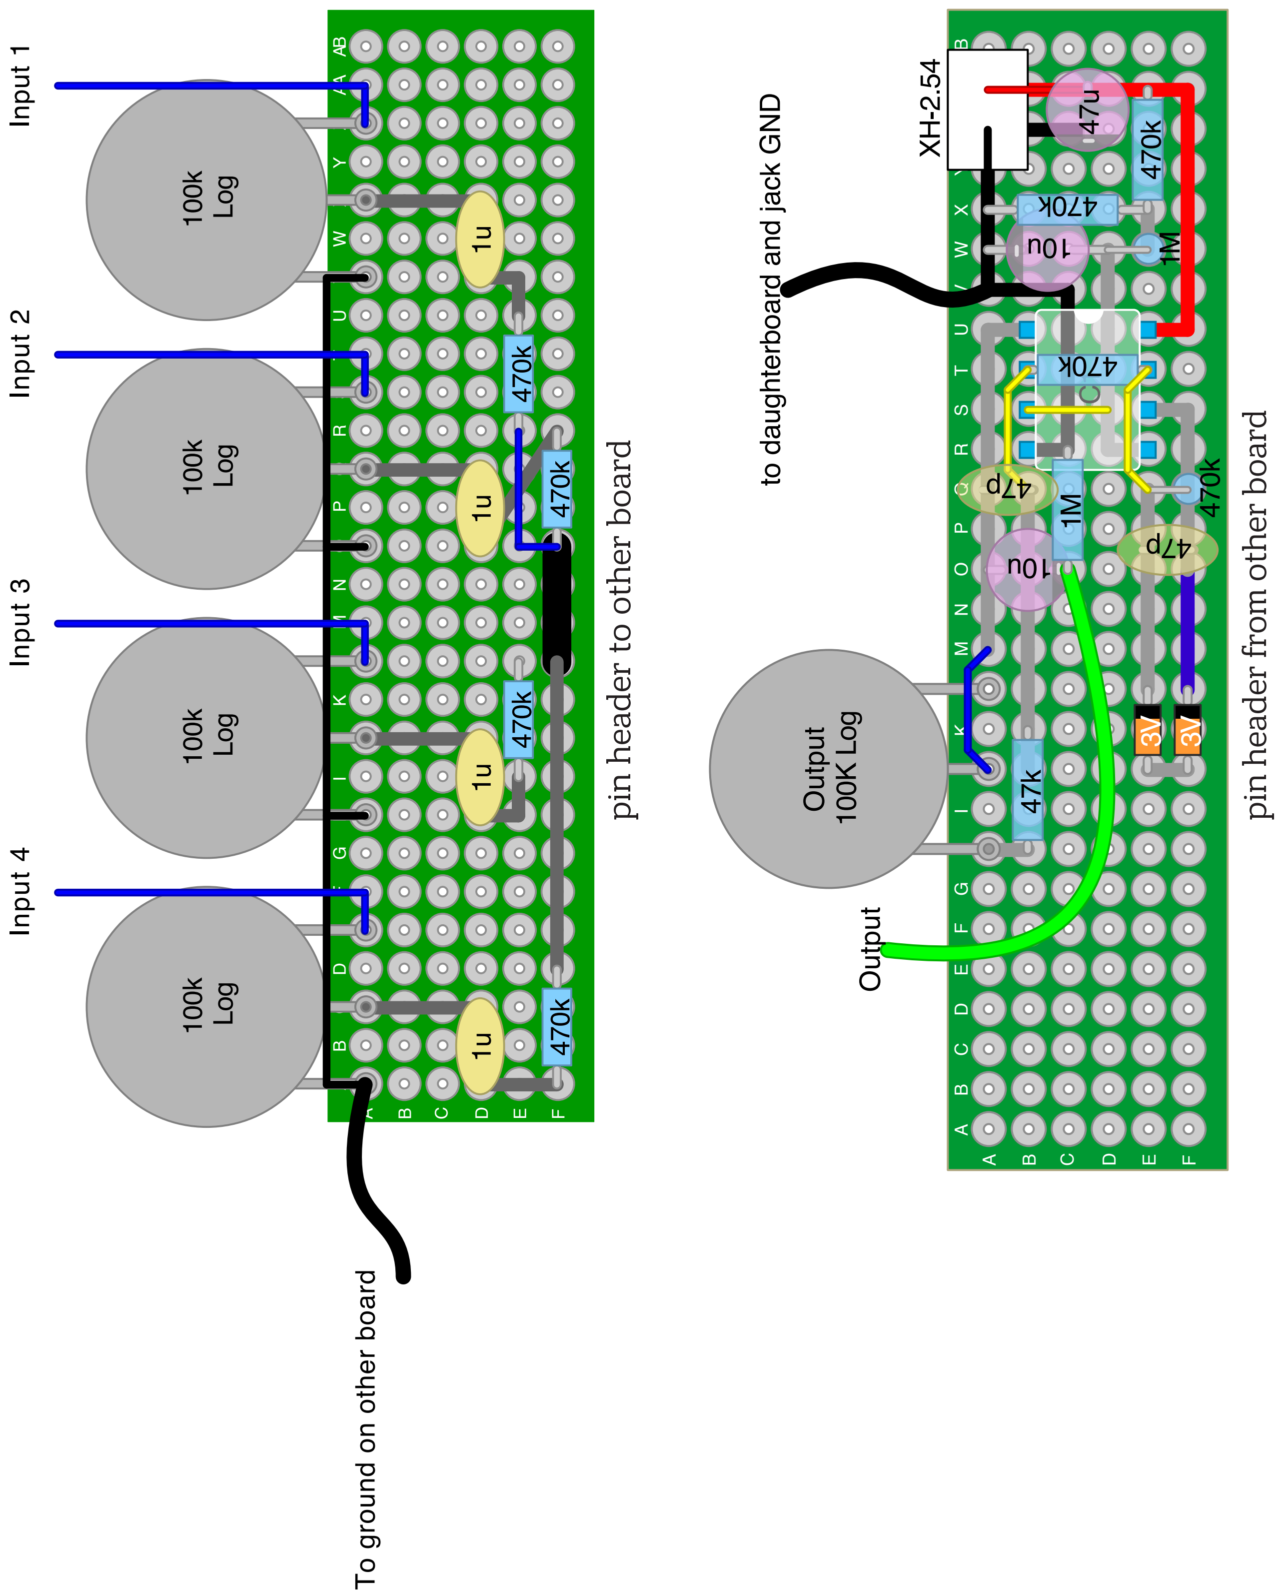

Input 1
Input 2
Input 3
Input 4
100k Log
1u
470k
pin header to other board
To ground on other board
to daughterboard and jack GND
XH-2.54
47u
470k
10u
1M
47p
3V
47k
Output 100K Log
Output
pin header from other board

0.3 Schematic

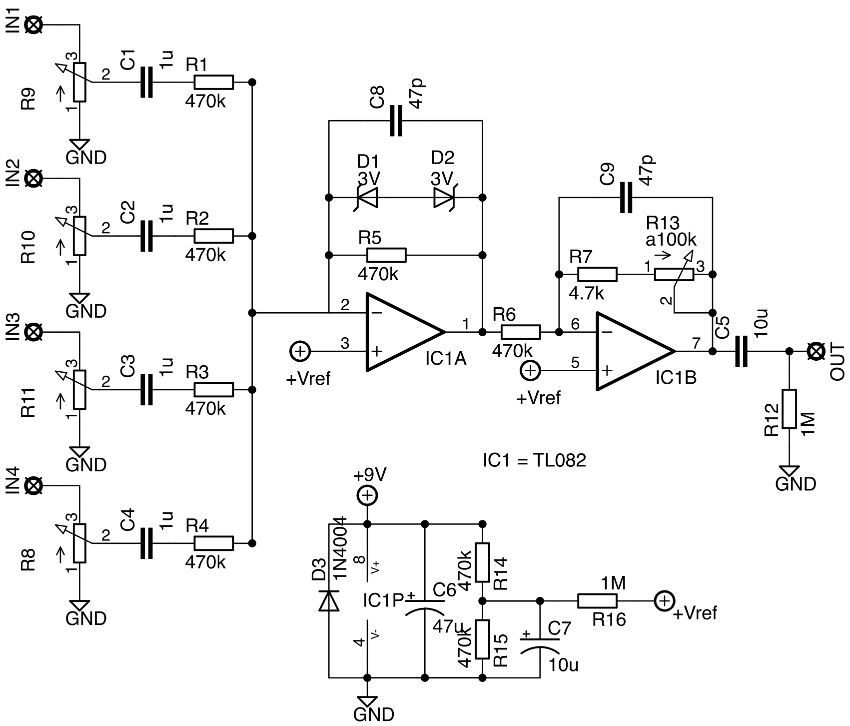

0.4 Bill of Materials

ICs

TL082	*x1*

Pots

A100k	*x5*

Resistors

4.7k	*x1*
470k	*x8*
1M	*x2*

Capacitors

47pF	*x2*
1µF	*x4*
10µF	*x2*
47µF	*x1*

Diodes

3V zener	*x2*
1N4004	*x1*

Connectors

Banana Jacks	
Black	*x4*
XH–2.54	*x1*
Pin header	
Male	*x1*
Female	*x1*

0.5 Drill template

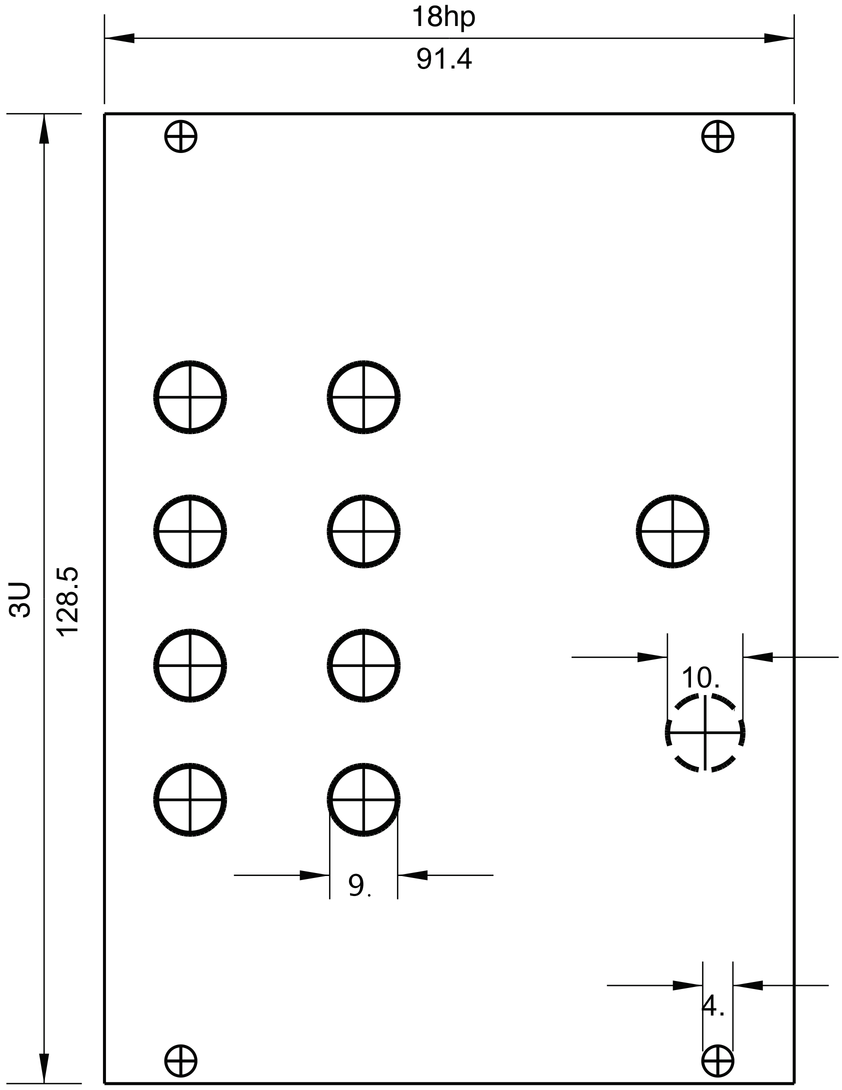

All measurements in millimetres

1:1 scale

Print out to use

1. Explanation

This mixer is designed to be the last thing in a modular set up so it needs to do two things: mix more than one input together; and step them down to line level. Our two op-amps do each of these things separately, with a few extra bonus features – nice overdrive when we slam the inputs hard and a fantastically low impedance output for driving a long cable to a mic or mixer.

1.1 Input stage

The 4-channel input stage, which is where our inputs come together and become one, is a basic unity-gain inverting amplifier. First, each of our inputs comes in through a pot used as a volume control that we can use to tweak the level of each input independently. After this it comes through a capacitor to filter out any DC offset present in the signal. Last, it comes through a 470kΩ resistor to reach the inverting input of the op-amp. This 470kΩ resistor together with the op-amp's feedback resistor, which is also 470kΩ, sets the gain at −1 for each channel.

In the feedback loop there are also two zener diodes, which is where our overdrive comes from. This is classic "not a bug but a feature" design. The diodes let us limit the signal to around +/- 3.6V. The reason for this is that when we sum all our signals together, they can easily exceed what the op-amp is capable of outputting and would result in the op-amp clipping. Op-amp clipping can be OK but usually it sounds pretty fugly. Luckily, using diodes lets us create much a much nicer "overdrive" sound at the expense of a little less headroom.

The last piece of the puzzle is the very tiny 47pF capacitor. Such a small value might seem pointless, but it's to stop the op-amp oscillating or ringing which can be a problem sometimes. Making it this small also has the advantage of not rolling off too much highs.

1.2 Output stage

The output stage is another inverting amplifier that takes the summed outputs from the input stage and then attenuates them so they are brought down to somewhere between line and instrument level. We have a pot to adjust the gain and another tiny cap to add a little more protection from ringing and oscillation. Last, but not least, we have a nice big 10µF cap to preserve all the lows and a 1MΩ pull-down resistor to stop any horrible pops surging into whatever you hook the output to.

1.3 Power filtering and V_{ref} stage

You might have noticed that this circuit lacks the 78L05 we use in most of the other circuits. The reason for this is that we are working with op-amps and op-amps tend not to be able to manage rail-to-rail output like our CMOS chips can.

For this reason, we are running this (and most other op-amp circuits in the modular) off the 9-12V DC input. This allows us more headroom, since we know that our input it usually going to be 0-5V DC, so we don't have to worry about clipping … at least not when we are summing signals together in a mixer like this.

However, op-amps are generally meant to be run on a dual power supply – meaning that there is both a positive and negative power rail. We're trying to keep everything simple, so we're just using a single-sided DC supply. We can get around this though by creating a "virtual ground" voltage that we call V_{ref}. This is made by making a voltage divider that will half the supply voltage, and then we use this half-way point as our reference.

We've also got some nice big capacitors to add some filtering of the power to keep everything behaving nice.

2. Analysis

We're going to mix it up a little here and we'll be looking at the output stage before the input stage. This is because the output stage is much simpler than the input stage as there is only one input to deal with. First, we're going to have to learn a little bit about op-amps and what makes them act the way they do.

2.1 Op-amps as inverting amplifiers

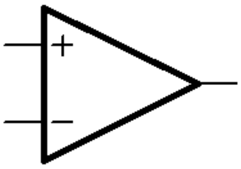

(fig 2.1a)

Op-amps *(fig 2.1a)*, like pretty much all other components, can be boiled down to some massive generalizations about how they behave. These generalizations tend not to hold up to scrutiny, but do tend to be good enough for a basic understanding and some rough analysis. For op-amps the most important golden rule is this:

Op-amps will do everything they can to make the voltage at the inverting input (-) the same as the voltage at the non-inverting input (+).

Using this rule, we can build ourselves some nice, dependable amplifiers.

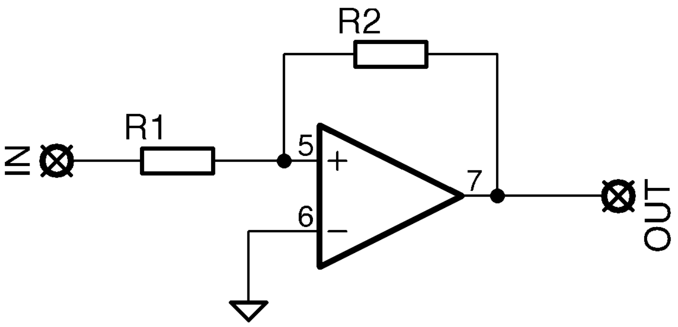

(fig 2.1b)

Here *(fig 2.1b)* we have a schematic for the simplest possible inverting amplifier using an op-amp. You'll notice it looks very similar to the two stages of our mixer but missing the capacitors and diodes, as well as the non-inverting input being tied to ground rather than a V_{ref} voltage like ours. It is possible to create a non-inverting amplifier, but we don't use any in this module.

The most important thing about any kind of amplifier is its gain, often referred to as G, which is a measurement of how much bigger the output is compared to the input. Although most amps boost the voltage, there are other amps that boost current or convert current into a voltage and vice-versa.

Luckily for us, someone else already did the calculation for an inverting amplifier and came up with this:

$$Gain = -\frac{R_2}{R_1}$$

However, algebraic expressions tend to be useless unless we understand the reasoning behind them.

Since we know that an op-amp works as hard as it can to make the voltages at its two inputs the same, it is actually very easy to calculate the amp's output. First, let's remove our op-amp and pretend there's just a voltage divider between the output and the input. This voltage divider has something special about it however. Usually we want to know the value in-between the resistors, but this time we know exactly what it is and want to calculate the voltage at the bottom.

So from here we can make a few statements.

- The current through the two resistors will be the same (Kirchoff's current law)
- The voltage across R_1 will be the difference between the input and V_{ref}
- The voltage across R_2 will be the difference between V_{ref} and V_{out}

Therefore, since $I = {}^{V}/_{R}$ and we know the current through the resistors is the same:

$$\frac{V_{in} - V_{ref}}{R_1} = \frac{V_{ref} - V_{out}}{R_2}$$

Also, for the sake off simplicity, we can consider V_{ref} to actually be DC offset (because it is) which allows us to ignore it when calculating the AC gain of the amp. Thus, for calculating the gain we are going to pretend that V_{ref} is 0V, so we can simplify it to:

$$\frac{V_{in}}{R_1} = -\frac{V_{out}}{R_2}$$

So, to figure out the gain we just need to calculate V_{in}/V_{out}:

$$Gain = \frac{V_{in}}{V_{out}} = -\frac{R_2}{R_1}$$

The negative result just means that the output is inverted, so it will look "upside-down". Thankfully, human ears can't tell the difference so it doesn't matter in an AC amp.

To get V_{out}, we multiply V_{in} by the gain and add V_{ref} back in.

2.2 The output stage

It might seem a little backwards to look at the output first rather than the input, but our output stage is much simpler as it lacks any diode clipping and is not summing multiple inputs together. However, the output stage does something very important. The 0–5V signal running around inside our modular is a fair bit louder than line-level or instrument-level audio, which is what amplifiers are expecting. This means we need to drop our output down to a more manageable level

2.2.1 What the hell is a decibel?

$$xdB = 10^{x/10}$$

Figuring out what line-level actually is turns out to be less simple than you'd think as it involves understanding what a decibel is. What makes this confusing at first is that a decibel is a RELATIVE measurement – which means it tells us about a comparison between two powers or amplitudes. A lot of the time one of these is already defined for us as a very specific measurement of power or voltage, all of these references tend to have names – dBu, dBV, dBm etc.

Line level is defined as being 0dBu, which is in turn defined as the "V_{RMS} equivalent to 1mW through a 600Ω load". RMS stands for "root mean square" and is simply a way to convert a waveform so you can calculate the power dissipated by it. For a sine wave of any frequency, it is defined as:

$$RMS_{a\sin(2\pi ft)} = \frac{a}{\sqrt{2}}$$

$$0dBu = \sqrt{0.6V_{RMS}} \approx 0.7746V_{RMS}$$

Where *a* is its amplitude (the difference between the highest and lowest points). For a square wave the RMS is just *a*.

From here we can calculate what the RMS equivalent of 0dBu is:

$$0dBu \times 10^{-20/10} = \frac{\sqrt{0.6}}{100}V_{RMS}$$

Instrument level is a little harder to pin down as not all instruments output the same. It is in the ballpark of –20dBu. But how much smaller is that? Luckily, converting from dB is fairly simple.

That's a 100th, or 1%, of the original power!

Before we move on, it's worth noting why decibels are useful when we talk about amplification and gain. Since decibels are logarithmic, when we want to figure out what the total gain of a number of stages are, all we have to do is add together their dB gains, which is much easier than multiplying them all together. Decibels also allow us to represent massive changes with very simple numbers. +90dB is an increase of 10^9 or 1,000,000,000! 90 is much easier to write out. This simple adding together also means that memorising a few particularly useful dB ratios can make it very easy to add and subtract cascaded stages of gain to figure out the total gain at the output. Mixing desks and DAWs all use dBs too, so knowing how much louder or quieter you are actually making a track can be really useful. Conveniently, the human ear hears +10dB of change to be twice as loud and –10dB to be half as loud!

The most useful dB gains to remember are:

- –20dB = $1/_{100}$ the power, $1/_{10}$ amplitude
- –12dB = ~$1/_{16}$ power, ~$1/_4$ amplitude
- –10dB = $1/_{10}$ power, ~$1/_3$ amplitude
- –6dB = ~$1/_4$ power, ~$1/_2$ amplitude
- –3dB = ~$1/_2$ power
- 0dB = no change
- +3dB = ~2x power
- +6dB = ~4x power, ~2x amplitude
- +10dB = 10x power, ~3x amplitude
- +12dB = ~16x power, ~4x amplitude
- +20dB = 100x power, 10x amplitude

Notice that when you add dB values together, you MULTIPLY the gains.

2.2.2 Calculating the desired gain

We need to drop the output down to where a square wave would be around ~0.7756V peak-to-peak since square-waves are the highest-power AC waveform possible. Luckily, the clipping diodes at the input stage mean that we already put a known limit on what is coming into the output stage. V_{max} = ±3.6V = 7.2V_{pp}.

$$Gain = \frac{Vout}{Vin} = \frac{0.7746}{7.2} = \ 0.1$$

However, this is about the absolute loudest we want it to be. Line-level can be around −20dBu lower. −20dB is a ratio of 0.1 or $^1/_{10}$th of the original amplitude ($^1/_{100}$th of the original power) which would make our needed gain ~0.001 at the absolute lowest when dealing with square waves. This is a little over the top however, and we'll be better off aiming for a minimum gain of ~0.01.

Using our gain equation, figuring out the resistors needed is painlessly simple. We can choose whatever we want for R_1, and then we just calculate what we need for R_2. R_2 will actually be a resistor and a pot. The resistor will set the minimum gain, and the resistor plus the pot's maximum resistance will set the maximum gain.

Since we're already using a bunch of 470kΩ resistors, lets use another for R_1! Since gain = $-R_2/R_1$ we can figure out R_2 but simple multiplying R_1 (which is 470kΩ) by the desired gain.

R_2 = 0.01 × 470kΩ = 4.7kΩ

Now that we have figured out our minimum gain we just have to slap in a pot to let us add resistance and increase the gain. Knowing that we want the gain to be below 1, we need to pick a pot SMALLER than R_1, which is 470kΩ. A100kΩ pots are common, so we'll throw one of those in. For fun, lets figure out what our max gain is when the pot is up full (so R_2 will be 100kΩ plus the 4.7kΩ of our minimum gain resistor):

Gain = R_2 /R_1 = 104.7kΩ/470kΩ = ~0.2

0.2 is a little higher than what we specified, but that's not necessarily a bad thing. That's a good amount of wiggle room should we need it … or we want to smash the input on a valve amp for lush gain sounds.

2.2.3 So what is our amp's gain in dB?

Calculating gain in dB from a known voltage gain is fairly simple. We can covert our gain using the expression:

$Gain_{dB}$ = 10 x $\log_{10}(Gain_V)$.

It's also possible to do the opposite with:

$Gain_V$ = 10^($Gain_{dB}$/10).

Remember that Vvltage gain is a measurement of AMPLITUDE change and decibels measure a change in POWER.

2.3 The output stage

Now we've learnt a bit about inverting amplifiers and how to calculate their gain, we can take it to the next level and look at a SUMMING amplifier. Mixers for audio use are actually just summing amplifiers. They take a bunch of voltages in their inputs and then use an op-amp to sum them together and output the total. You'll notice that we have changed things a little now. What was R_2 is now R_f (for feedback).

2.3.1 Output of a summing amplifier

The summing amplifier we are using here is another inverting amplifier, so we can once again treat it as a voltage divider (sort of). The issue is that we have more than one voltage that all converge at our op-amps converting input. We already know that this point has to be the same as V_{ref}, but how does it manage that with so many different inputs?

According to Kirchhoff's current law, the current flowing into a point is the same as the current flowing out of it. We can use this observation to our advantage as we know that all of the current flowing through R_f is a combination of all the current through the input resistors R_1 to R_n.

$$I_f = I_1 + I_2 + \cdots + I_n$$

$$\frac{V_{out}}{R_f} = -\frac{V_1}{R_1} + \frac{V_2}{R_2} + \cdots + \frac{V_n}{R_n}$$

$$V_{out} = -R_f \left(\frac{V_1}{R_1} + \frac{V_2}{R_2} + \cdots + \frac{V_n}{R_n} \right)$$

If we choose different values of each of the input resistors,we can tailor the gain for each input. The calculation will be the same as if it were a single inverting-op amp. With our mixer we don't want to do that because pots are much simpler, but it is something to keep in mind.

2.3.2 Clipping diodes in op-amp feedback loops

When put in a feedback loop, diodes have the effect of limiting the output voltage from the op-amp. This works because the op-amp will still deliver the CURRENT necessary to keep the two inputs at the same voltage – by "pushing back" against the input through the resistor connecting it to the input – but this doesn't require the op-amp to raise the voltage at the output like when there is only a resistor in the feedback loop. Remember that diodes experience very tiny changes in the voltage difference across them for much larger changes in the current through them. This means that as soon as the voltage drop between the output and the input gets large enough to switch the diodes on, it will pretty much stay pinned at that point.

Luckily, diode clipping sounds much nicer than the clipping that happens when the op-amp reaches its maximum output. This can happen very easily in our circuit too. If we have all four the inputs going $5v_{pp}$, then this would add up to $20v_{pp}$ and we only less than 9v to work with! So the diodes do us two favours: they give us a nice built in overdrive effect for cheap; and they prevent the op-amp hitting the rails and making a fugly mess out of our inputs!

For this set up we are using zener diodes.

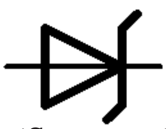

(fig 2.3.2a)

Zener diodes *(fig 2.3.2a)* are special diodes that work both ways. This might seem to defeat the point of a diode, but the voltage needed get them to conduct backwards is often much bigger than the 0.6v needed to get them to conduct forwards. It's also a voltage that we can choose by picking a specific diode. Zener diodes come in a wide range of reverse voltages and they have a number of great uses. In this circuit we are using two facing opposite directions.

This means that the total voltage drop is the reverse-breakdown voltage of one zener plus the 0.6v forward voltage of the other (making 3.6v in our circuit). This allows us to use only two diodes to get much more headroom than would be possible even with LEDs. With LEDs we'd need around four (two in each direction), and with normal silicon diodes we'd need twelve(!). However, with zeners we only need two. How delightfully frugal!

3. Modifications

3.1 More inputs

Adding more inputs is as simple as copying the other input stages. You can go wild and add as many as you want. However, bear in mind that the more inputs you have the more likely you are to have things start overdriving. Twenty $1V_{p.p.}$ inputs need 20V of headroom, four only need 4V!

3.2 Headphone out

The LM386 is an awesome, cheap, single chip, low-voltage power amp. You don't get much better than that combination of adjectives. Using one to build a little headphone amplifier is not tough and it will probably help keep your neighbours and flatmates happy. Bear in mind that you will need two chips for stereo headphones (do mono headphones even exist anymore?). This add on could triple your part count, but it's not much more difficult than what we've already built. Check out HeadWize for a great LM386 head-phone amp schematic, just remember to build the right channel too!

http://headwize.com/?page_id=779

4. Lessons learned

4.1 What not to do when making a layout

My unnecessary obsession with building everything on as small a board as possible ended up really backfiring on me on this one. The main board with the op-amp on it was really tough to build and debug.

Similarly, using PCB spacers ended up having the opposite effect from what I wanted. By having to put the pot for the master gain on the underside of the board, it ended up making the layout bigger (from an HP perspective) and more fiddly.