MSK 006 Exponential VCA

Matthew Skala

May 18, 2016

Contents

General notes

This is documentation for the MSK 006 Exponential VCA, which is a voltage-controlled amplifier built entirely out of discrete components and intended for use in a Eurorack modular synthesizer. The design is more or less the circuit of an IC operational transconductance amplifier (like an LM13700) with a simple exponential converter on the control voltage input, plus a pure class B output stage with a special circuit for compensating the crossover distortion usually typical of pure class B amplifiers.

The circuit and PCB design included in this package is for a single channel. This documentation describes combining three of the boards into a single 12HP module, largely because I had OSHPark make my PCBs and they do them three at a time. However, many other configurations are possible.

There is a mystique surrounding "discrete" audio circuits, and a lot of outlandish claims are made about them. These claims are mostly bullshit. Especially in a VCA circuit, you get better performance if your transistors are well-matched and kept at the same temperature, and that is much easier if they are parts of the same chip. *Integrated circuit amplifiers provide better performance than "discrete" circuits in almost all cases.* The circuit described in this document will *not* outperform standard IC designs by any objective measure. It has several disadvantages that they avoid. Similar comments can be made about through-hole construction, as used by this module, when compared to surface-mount.

Then why design and build it?

The main reason is simply that it's fun. It may also be educational. One of my main reasons for building it was to gain a better understanding of how IC OTAs actually work, and to test and perfect some techniques I will use in other designs. This is intended to be a VCA you can hear—one that adds something of its own to the sound instead of just being a pure colourless volume control.

Specifications

This module, as I built it and in normal use, draws up to about 10mA *per channel* (30mA overall for a triple module) from both the +12V and -12V Eurorack supplies. In output-clipping conditions with a low impedance connected to the output, it may draw a little more. It does not use +5V power.

The input impedance is 100kΩ for the audio input, and unspecified, but high, for the control voltage input. It will not be damaged by shorting any input or output to any fixed DC voltage at or between the power rails. Patching outputs to outputs is inadvisable, but unlikely to damage the MSK 006; it may possibly (but probably not) pose some threat to the *other* module if an MSK 006 output is patched into some other module's output.

By default, the output is capable of driving about ±6V into a 1kΩ impedance. It could be modified to provide a lot more power by reducing the values of R12 and R13, and possibly substituting higher-power types for Q17 through Q20.

This module (assuming a correct build using the recommended components) is protected against reverse power connection. It will not function with the power reversed, but will not suffer or cause any damage. Some other kinds of misconnection may possibly be dangerous to the module or the power supply.

The sensitivity of the control voltage input is approximately 10dB per volt. It is a little more sensitive above zero and a little less below zero, and in principle any control voltage between the power supplies should be usable. The *recommended* range of input control voltages is about 0 to 8V.

This module exhibits a significant amount of control voltage feedthrough, which will be apparent for instance as a "thump" or "click" noise when used with a fast envelope. This is a normal part of the sound of the module; if you don't like it, you should be using an IC instead of a discrete OTA.

The output of this module is intended to be AC coupled. It can be modded for DC coupling by installing a wire jumper instead of C6, but be aware that then there is likely to be a DC offset on the output, and the sign and magitude of this offset will be difficult to predict or control. Such a modded module may also be more sensitive to output short circuits than the AC version.

It will probably work fine on ±15V power without any changes to the circuit, but I have not tested that.

Source package

You most likely received this document either printed on paper or as a PDF file, but if you download the source code package from `http://ansuz.sooke.bc.ca/electronics.php`, it should include:

COPYING a copy of the GNU General Public License, version 3.

Makefile a configuration file for GNU Make, to assist in building the documentation; be warned that this is actually a semi-manual process, and the enclosed Makefile depends on access to my internal SVN repository. It will probably not work in a completely standalone way, but is provided for whatever value it may have.

docs/ subdirectory of source files for this PDF document. Actually building it depends on LaTeX, LaTeX.mk (Danjean and Legrand), and Circuit_macros (Aplevich). There are a few PDF files that must be manually generated using Kicad from the Kicad sources elsewhere in the package, if the relevant designs change.

gerbers/ subdirectory containing the PCB designs in the format that fabrication services usually require; manually generated using Kicad.

kicad-symbols/ subdirectory containing my customised schematic symbol and PCB footprint libraries for Kicad. Kicad doesn't normally keep dependencies like symbols inside a project directory, so on my system, these files actually live in a central directory shared by many projects. As a result, upon unpacking the ZIP file you may need to do some reconfiguration of the library paths stored inside the project files, in order to allow the symbols and footprints to be found. Also, this directory will probably contain some extra bonus symbols and footprints not actually used by this project, because it's a copy of the directory shared with other projects.

layout.kicad_wks customised title block and drawing border for use with Kicad.

panel10.jpg, panel12.jpg photos of the front panel (10HP and 12HP versions) for use in places like ModularGrid; the 12HP version is actually edited from the 10HP photo of one of my prototypes, because as of this writing, I don't have a physical 12HP panel to photograph.

stripboard.tex LaTeX source for an experimental stripboard layout, *warning*, this is not tested and may not be up to date; it was mostly just an exercise in using the graphics software.

vca.fpd Schaeffer/FPE "Front Panel Designer" file for a 12HP front panel.

vca.kicad_pcb Kicad file for the PCB design.

vca.pro Kicad "project" file.

vca.sch Kicad file for the schematic diagram.

PCBs and mounting

The enclosed PCB design is for a basic two-layer board with top and bottom ground planes. There shouldn't be anything terribly difficult about getting it fabricated by OSH Park or any other contractor. The circuit could, of course, also be built on stripboard, with point-to-point wiring, or by any other method. The PCB design includes a bunch of wire pads for connecting boards to each other (for CV normalling, and to share a single power connection for a three-board module) or to allow

Each PCB is $3.90'' \times 1.45''$, or 99.06mm× 36.83mm. They are designed to mount perpendicular to the panel (using the jacks and pots for support) in a Eurorack case at least 39mm deep (including a couple of extra millimetres for board to board wiring), but the boards include mounting holes that could be used to support them in some other configuration, for instance in another format.

Panel

The enclosed panel design is compatible with Eurorack (12HP for a three-channel module) and Schaeffer AG's manufacturing service. Front Panel Express is the US arm of the same business and should also be able to fabricate this design. There are plenty of other options for making a panel and this one is not necessarily the cheapest or best, but the result is very pretty, anyway. I used 1.5mm anodized aluminum, which is optimal for the jacks I used; they can just barely work with 2.0mm, but it's pushing it.

My panel layout is based on mounting both the jacks and the pots directly on the PCB. That puts their centres at different distances from the PCB

(5.5mm for the jacks and 6.95mm for the pots); so the holes in the panel are placed accordingly and these holes do not line up horizontally when viewed from the front in the finished module. If you want the alignment on the front panel to be perfect, you may have to do something more complicated.

Jacks

My PCBs were designed for use with Lumberg 1503 12 switched mono 1/8″ barrel jacks; CUI Inc. MJ-3536 jacks should also work fine. The CUI jacks are slightly deeper, but still fit reasonably well in the PCB footprint designed for the Lumberg jacks; there is enough extra space in the circular holes to accommodate the change in pin location. One could instead mount jacks of any kind to the panel and not the board, with wires connecting them to the appropriate locations on the board.

Transistors

I designed the module around 2N3904 NPN and BC557C PNP silicon transistors, but it should work with almost any common transistor types. You can even use a mixture of different transistor types. However, NPNs should only be replaced with NPNs and PNPs with PNPs, and within each of the following sets, all the transistors should be of the same type:

- {Q3,Q6,Q7}
- {Q4,Q10}
- {Q5,Q8,Q9}
- {Q11,Q15,Q16}
- {Q12,Q13,Q14}
- {Q17,Q19}
- {Q18,Q20}

These sets also describe how the transistors should be "matched" if you want to match transistors on V_{BE}, but I don't recommend trying too hard to do that.

I built a prototype with a 2N5087/2N5088 combination, and that worked fine. The transistors used for Q4 and Q10 probably affect the sound most, and those are the ones to substitute first if you want to use special transistors to produce a special customized sound. I also prototyped a more ambitious mod (described in its own chapter of this document) to use N-channel JFETs such as type MPF102 for Q4 and Q10.

Watch out for the pinouts when substituting transistors; the 2N5087, in particular, has the opposite pinout to the BC557 (EBC where the BC557 is CBE) and so must be mounted backward relative to the PCB markings.

Capacitors

This module uses two 10μF electrolytic capacitors (per module) and two 0.1μF ceramic capacitors (per channel) for power supply decoupling; this is an undemanding application and any common types rated for the voltage may be used. I have also specified a 220pF capacitor (C1) for frequency compensation and to help prevent oscillation, and any common ceramic type should be fine. The value 220pF performed best in my tests, but does not make much difference, and the module would probably still perform acceptably even if this capacitor were omitted entirely. I did not actually observe any oscillation problems, but want to be on the safe side, especially if people will be randomly substituting transistors.

For the output coupling capacitor C6, I suggest using a film capacitor; not an electrolytic because of distortion, and the unpredictable and varying polarity of the DC offset. I have written 2.2μF on the schematic and board, but in fact I used 1.0μF in my prototype and it worked fine. I would not go much below 0.5μF for fear of harming the low-frequency response.

Resistors

I used the resistors I had in the specified values, which were an assortment of many different types ranging from 1/8W to 1/2W. None of the resistance values is critical; 10% tolerance would probably be fine, though in practice you are unlikely to easily find any worse than 5% unless they are very old indeed.

If you anticipate mis-patching the output a lot, then it might be worth calculating the power expected to be dissipated in the output current limiting resistor R12, and making sure it is within the rating of the resistors. With the specified 12V supply and 1.2kΩ resistance, assuming it goes into a dead short, the power dissipation will be 0.12W and just within the limits of a 1/8W resistor; it is possible to imagine even worse cases than a dead short, but you also will probably use a bigger than 1/8W resistor anyway.

Output coupling modification

This module is designed for audio signals, with an exponential response curve and an AC-coupled output. It should be good down to moderate LFO speeds, but it is not intended for processing control voltages. If you wish to use DC coupling instead, then install a wire jumper across across the circuit board pads for C6 instead of the specified capacitor, and use a 1/2W resistor for R12.

I do not particularly recommend this modification. The resulting DC-coupled module will likely have a significant amount of offset (zero input giving nonzero output voltage), so it still won't be very good for DC control voltages. It will also have an exponential response, whereas people usually want linear response for control voltages. There are many other VCA designs available which are better suited to control voltage applications, and one of them might be a better choice than a DC-modified MSK 006. Nonetheless, it's an easy modification to do, and some people really think they want DC coupling, so there it is for what it's worth.

Use and contact information

This module design is released under the GNU GPL, version 3, a copy of which is in this ZIP package in the file named `COPYING`. One important consequence of the license is that if you distribute the design to others—for instance, as a built hardware device—then you are obligated to make the source code available to them at no additional charge, including any modifications you may have made to the original design. Source code for a hardware device includes without limitation such things as the machine-readable, human-editable CAD files for the circuit boards and panels. You also are not permitted to limit others' freedoms to redistribute the design and make further modifications of their own.

I plan in the near future (as of this writing, mid-2016) to start a business and sell kits or assembled versions of this design, but at the moment, I have no such business. If you want one of these modules, then you'll have to either wait, build it yourself, or hire some third party to do it.

I'm posting this and other electronics projects at `http://ansuz.sooke.bc.ca/electronics.php`. That would be a good place to look for updated versions or other related material. My Soundcloud account, which usually includes tracks recorded with this and other homemade electronics, is at `https://soundcloud.com/matt-skala/`. I can be found on the Muff Wiggler Forum as "mskala," but although others are welcome to do so, I don't plan to host "build threads" for this or my other projects there. The instructions for building it are in this document instead.

Email should be sent to `mskala@ansuz.sooke.bc.ca`.

Safety and other warnings

I offer no warranties whatsoever regarding these instructions and you follow them at your own risk.

Ask an adult to help you.

Soldering irons are very hot.

Solder splashes and cut-off bits of component leads can fly a greater distance and are harder to clean up than you might expect. Spread out some newspapers or similar to catch them, and wear eye protection.

Lead solder is toxic, as are some fluxes used with lead-free solder. Do not eat, drink, smoke, pick your nose, or engage in sexual activity while using solder, and wash your hands when you are done using it.

Solder flux fumes are toxic, *especially* from lead-free solder because of its higher working temperature. Use appropriate ventilation.

Some lead-free solder alloys produce joints that look "cold" (i.e. defective) even when they are correctly made. This effect can be especially distressing to those of us who learned soldering with lead solder and then switched to lead-free. Learn the behaviour of whatever alloy you are using, and then trust your skills.

Water-soluble solder flux must be washed off promptly (within less than an hour of application) because if left in place it will corrode the metal. Solder with water-soluble flux should not be used with stranded wire because it is nearly impossible to remove from between the strands. Stranded wire is recommended for the connections between boards in this module, and those connections, at least, should *not* be soldered with water-soluble flux.

This module does not use any especially high impedances, but in general, residue from traditional rosin-based solder flux can result in undesired leakage currents that may affect high-impedance circuits. If your soldering leaves a lot of such residue then it might be advisable to clean that off. I used a "no-clean" rosin flux throughout my module, it produced very little residue, and this was not an issue for me.

Voltage and current levels in some synthesizer circuits may be dangerous.

Building your own electronic equipment is seldom cheaper than buying equivalent commercial products, due to commercial economies of scale from which you as small-scale home builder cannot benefit. If you think getting into DIY construction is a way to save money, you will probably be disappointed.

Bill of materials

These components are needed for each board:

Qty	Ref	Value/Part No.	
1	C1	220pF	ceramic, 0.2″ lead spacing
2	C4, C5	0.1μF	axial ceramic
1	C6	2.2μF	film capacitor, 0.2″ lead spacing
3	J1, J2, J3	1503 12	audio jack (Lumberg; or use CUI MJ-3536)
10	Q1, Q2, Q5, Q8, Q9, Q11, Q15, Q16, Q18, Q20	2N3904	NPN general-purpose transistor, TO-92 EBC
10	Q3, Q4, Q6, Q7, Q10, Q12, Q13, Q14, Q17, Q19	BC557C	PNP general-purpose transistor, TO-92 CBE
1	R1	27kΩ	
3	R2, R4, R11	10kΩ	
1	R3	100kΩ	linear PCB-mount panel pot, conductive plastic (TT Electronics P260 series)
1	R5	39kΩ	
3	R6, R9, R10	560Ω	
1	R7	100kΩ	
1	R8	4.7kΩ	
1	R12	1.2kΩ	
1	R13	3.0kΩ	

For each module of one or more boards sharing a power connection, add:

Qty	Ref	Value/Part No.	
2	C2, C3	10μF	radial aluminum electrolytic, 0.1″ lead spacing
2	D1, D2	SBA130	Schottky diode
1	P4	2×5	0.1″ pin header (Eurorack power connection)

Also necessary: one PCB per channel, hookup wire, solder, front panel, knobs, Eurorack power cable, etc. If building exactly to my physical design, you will need four hex spacers with 17mm body length, male M3 threads on one end and female on the other, plus two machine screws and two nuts matching these. If I counted correctly, the total number of pieces for one three-channel module, excluding bulk supplies like wire and solder, is 141; the same as the number of pieces in the Lego "Horse Trailer" set (model 6359, released in 1986).

Tolerances, temperature coefficients, and so on are not critical on any components. See notes elsewhere in this document regarding optional substitutions, modifications, and transistor matching. Resistors should be at least 1/8W; the board will reasonably accommodate up to 1/2W.

Build step-by-step

Note that although I'm describing a separate step for each component value, and that's how I built mine so as to have plenty of photo opportunities, if you are reasonably confident about your skills you may find it easier to populate all or most of the board (i.e. put the components in place) and then solder them in a single step. Except where noted, the order in which you add components does not matter much.

Preliminaries

Count out the right number of everything according to the bill of materials.

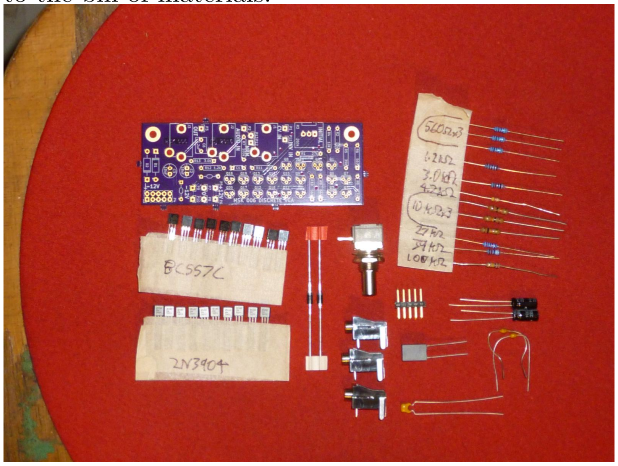

I recommend assembling the structural components (panel pots, jacks, spacers, and so on) first to make sure they all fit together properly. Not shown in this photo: the M3 machine screws and nuts (two of each) to mate with the spacers, shown. Another reason to do it this way is to prevent the relatively large amounts of soldering heat needed for the jacks, from damaging more delicate components that will be installed near the jacks. The disadvantage is that it means you will not be able to install the lowest-profile components first, which some people like to do on through-hole boards so that the board will lie flat component-side down on a table during soldering.

Also note: I built the prototype using a 10HP panel, and that is what's shown in the photos, but I don't recommend it because that size placed the potentiometer knobs too close together to use the knobs I wanted to use, and the boards close enough to the edges that component legs hang out into space that might be taken up be the adjacent module. The front panel design included in this package of plans is for a 12HP module (panel slightly wider, but otherwise as shown), which is what I recommend: it allows for a few extra millimeters between the boards and on each side.

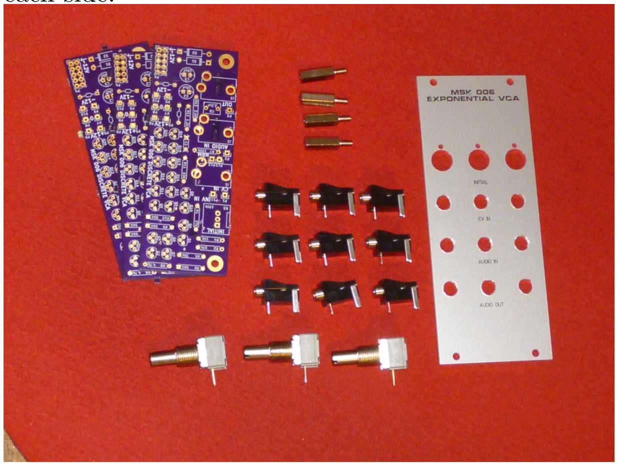

Your PCBs may have little chunks of extra material attached to the edges as a result of the manufacturing process; they are left over from when the boards were separated from larger panels. These will get in the way of other components, so break them off with pliers and file down any remaining protrusions around the edges of the boards.

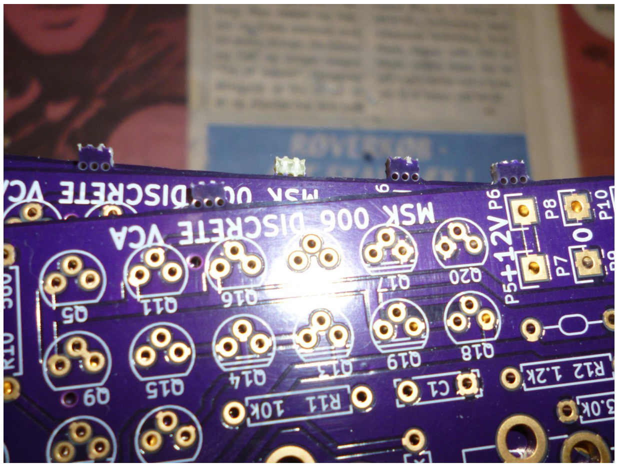

Place (do not solder yet) the jacks and potentiometers in their spots on the PCBs, and assemble the PCBs into a stack using the spacers and threaded hardware. It may be necessary to trim the legs of the potentiometers (just clip 1mm or so off with diagonal cutters—be careful not to remove too much) to get them to fit nicely together.

Attach the front panel to the panel components. Screw on all the nuts, but do not use any lockwashers at this time, and only screw the nuts on as tight as you can make them with your fingers. This is only a temporary assembly. The idea is just to hold all the panel components in place while you solder them, so that they will be soldered in perfect alignment with the holes in the panel.

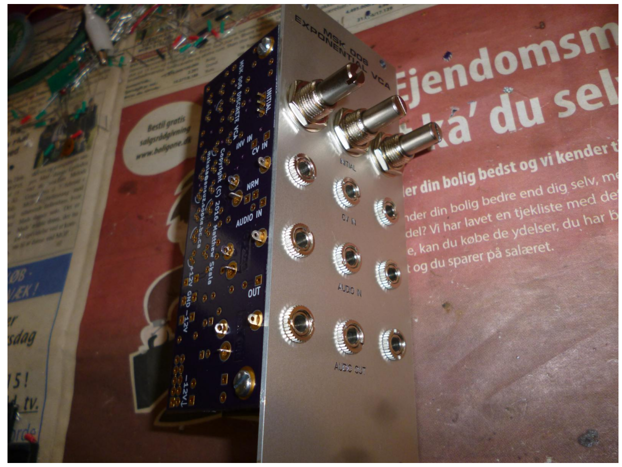

I don't have photos for every stage of this process, but from here you can solder the three boards one at a time. First, the one whose solder side is exposed when the module is fully assembled (the left-hand board when the module is installed in a case). Then, if you stacked up the boards exactly as shown, it should be possible to undo the two machine screws and the front-panel nuts holding that board in place, and carefully pull away the just-soldered board while leaving the other two attached to the panel. Then you can solder the second board, and similarly remove it, before doing the third. This way, each board is attached to the panel when you solder it, ensuring alignment will be good in the final assembly.

Disassemble the module in preparation for soldering the other parts to the boards.

Resistors

Resistors are never polarized. I like to install mine in a consistent direction for cosmetic reasons, but this is electrically unnecessary. They will bear colour codes of three or four bands representing the value, plus an additional band for the tolerance. Most commonly, the patterns are three value bands plus metallic gold for a 5% resistor, or four value bands plus brown for a 1% resistor, which may often lead to ambiguity in which direction you should read the code; other patterns are also possible, for instance if you happen to have higher-precision resistors. I used a mixture of resistor types in the prototype build shown in the photos, and describe in the text only the value band colour codes. If in doubt, always check resistors with an ohmmeter.

Install the three 560Ω resistors (colour code blue-green-brown or blue-green-black-black) R6, R9, and R10. These provide bias voltages for transistor bases in the exponential converter and the differential pair.

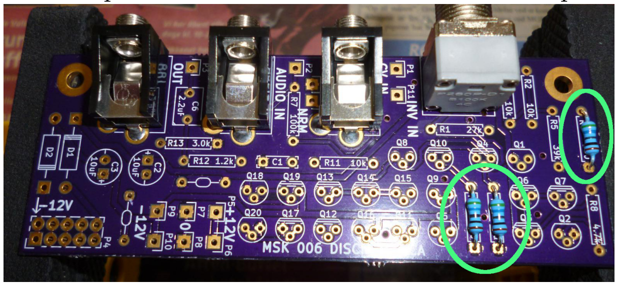

Install the 1.2kΩ resistor (colour code brown-red-red or brown-red-black-brown) R12. This limits the output current.

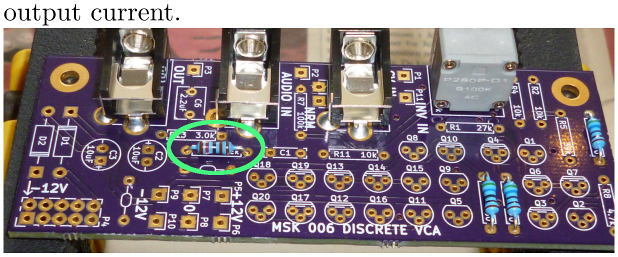

Install the 3.0kΩ resistor (colour code orange-black-red or orange-black-black-brown) R13. This provides a DC load for the output transistors, which limits the effect of imbalances between them.

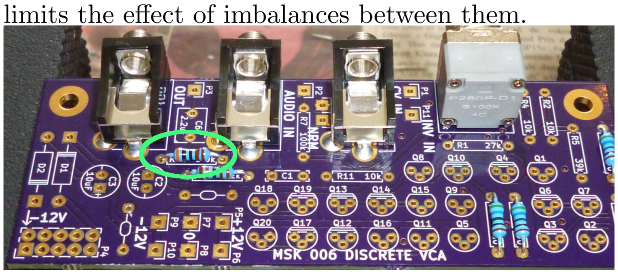

Install the 4.7kΩ resistor (colour code yellow-violet-red or yellow-violet-black-brown) R8. This sets the upper limit on the current that the exponential converter can generate.

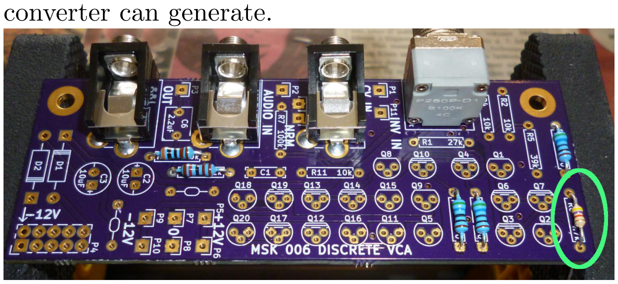

Install the three 10kΩ resistors (colour code brown-black-orange or brown-black-black-red) R2, R4, and R11. R2 and R4 control the response of the exponential converter to control voltage and knob adjustment; and R11 converts current to voltage in the interface between the VCA's input and output sections.

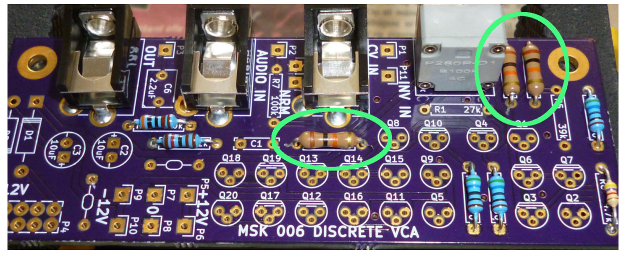

Install the 27kΩ resistor (colour code red-violet-orange or red-violet-black-red) R1. This sets the bias current for the control voltage input buffer.

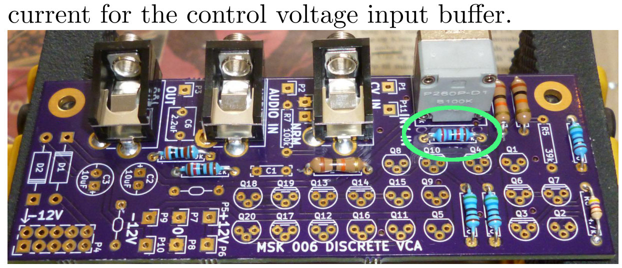

Install the 39kΩ resistor (colour code orange-white-orange or orange-white-black-red) R5. This sets the response of the exponential converter to the control knob.

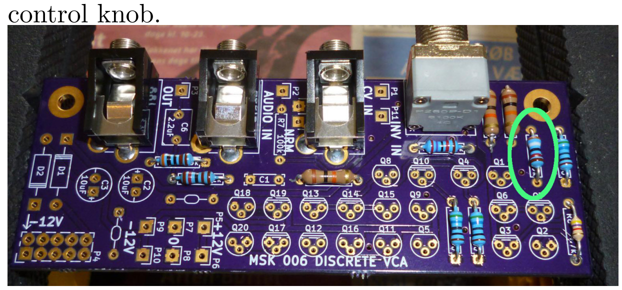

Install the 100kΩ resistor (colour code brown-black-yellow or brown-black-black-orange) R7. This sets the impedance and limits current on the audio input.

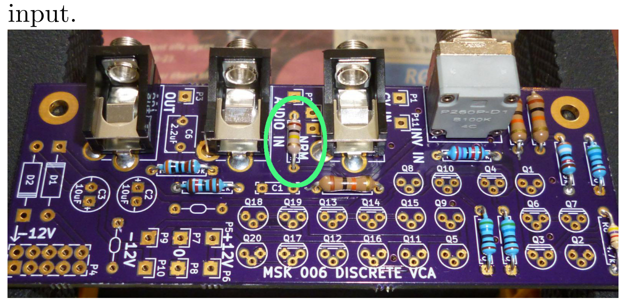

Capacitors

None of the capacitors installed *at this point* are polarized; there will be some polarized ones discussed later.

Install the 220pF capacitor C1. This capacitor helps compensate for the glitch when the class B output stage switches between "push" and "pull" driving modes; it also is intended to help prevent parasitic oscillations, though it's not clear that they would ever be an issue in this circuit anyway.

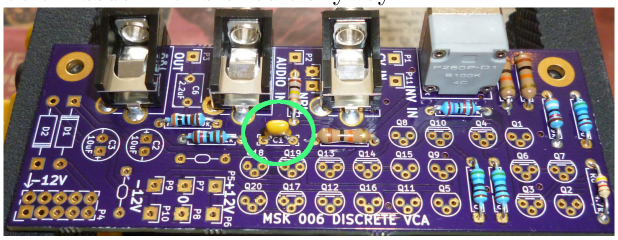

Install the 0.1μF capacitors C4 and C5. These are general power supply bypass capacitors, and reduce crosstalk between this channel, other channels, and other modules in a system. Their reference designators ("C4" and "C5") are not shown on the PCB silkscreen, only an outline symbol reserved for decoupling capacitors. This is for consistency with some of my other projects, where decoupling capacitors are so numerous that marking a designator for each one would be inconvenient.

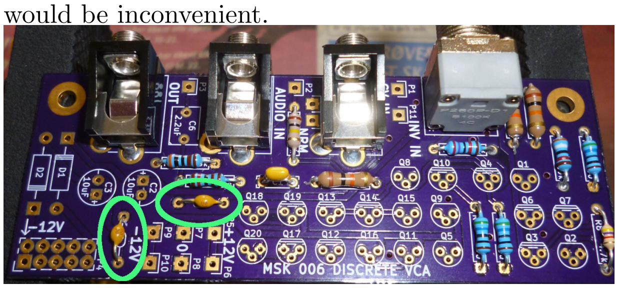

Install the 2.2μF film capacitor C6. This blocks DC on the audio output.

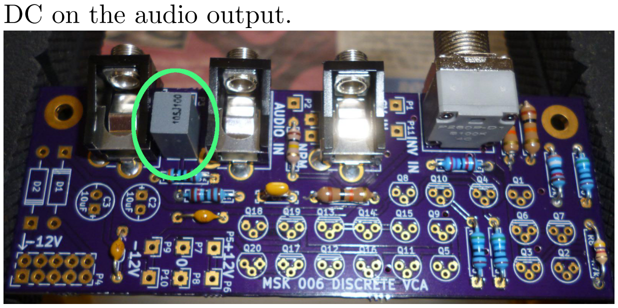

Semiconductors

This module's core consists of 20 silicon bipolar junction transistors, 10 each of NPN and PNP. The recommended types come in TO-92 packages, which are little epoxy beads flattened on one side. My 2N3904 NPN transistors shown in the photos happened to be marked with silver paint[*] and thus they are easy to distinguish from the all-black BC557C PNP transistors, but the silver paint marking is relatively uncommon. Both types of transistors are more often all black with barely-visible etched lettering identifying the type and manufacturing date. Be careful not to mix them up. Installing the wrong type of transistor in a location on the board, or the right type in the wrong orientation, is likely to destroy at least that transistor and possibly others too when power is applied.

As shown in the photo, there are two different silkscreen symbols for the transistors to help keep the two types separated. NPN transistors like Q8 and Q9 are denoted by a circle with a flat side indicating the orientation of the flat side of the plastic package. PNP transistors like Q13 and Q14 have two additional lines.

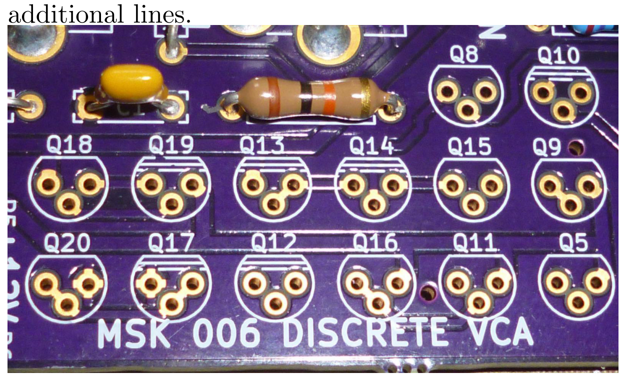

The pins of a TO-92 transistor and therefore their solder pads are close enough together that it can be hard to avoid creating solder bridges between them. In this project the situation is further complicated by the fact that several of the transistors are "diode-connected," that is with base and collector intentionally joined by copper on the board. The places where this occurs are marked by arcs on the solder-side silkscreen, as shown. These arcs should remind you, when you inspect the completed board, *not* to attempt to remove any solder bridges that may form between those pins—with PCB copper making the connection, removing the solder bridge would be both difficult and unnecessary. Of course, solder bridges among any *other* transistor pins remain a problem and should be removed.

[*] It is not clear what the significance of the silver paint marking may be; we had an interesting discussion about it on the synth-DIY mailing list and did not come to a convincing conclusion.

Install the 10 NPN type 2N3904 transistors (I won't list all their designators here) according to the PCB silkscreen markings.

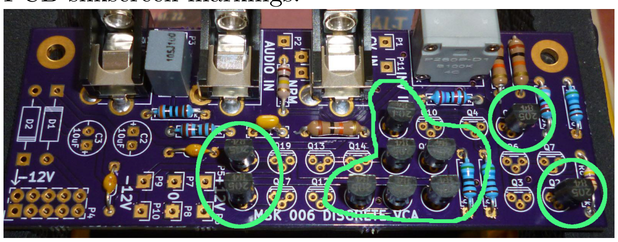

Install the 10 PNP type BC557C transistors according to the silkscreen markings.

Power inlet components

I built my prototype modules each with three VCA boards connected together and sharing a single Eurorack power connection. These instructions assume you are building the module the same way; modify as needed if you are doing something else. The power inlet components (C2, C3, D1, D2, and P4) only need to be installed on one board, the one that will be rightmost in the final assembled module.

Install the two Schottky diodes D1 and D2. These protect the module against reverse connection of the power supply. They are polarized and must be installed in the correct direction; otherwise they will prevent the module from operating. One end of each diode will be marked, usually with a stripe of grey paint around the black plastic body of the diode. That end is the cathode. The diode outline on the PCB silkscreen is marked with a similar stripe showing the direction of the cathode, and the solder pad for the cathode is square instead of round.

Install the two 10μF electrolytic capacitors C2 and C3. These filter incoming power to prevent noise in the case power system from affecting the VCAs. They are polarized components, and may explode if connected backwards. As such, there are multiple clues to make sure they end up in the right direction. The negative leg of each capacitor will be marked in some way, usually with a printed stripe and minus signs on the plastic wrapping of the capacitor body. The negative leg of the capacitor will usually also be shorter, though that is less reliable than the body markings. Some capacitors are also marked with red text (traditionally a "positive" colour) on the positive side, but I would trust that sign least of all. On the PCB, the positive and negative pads are marked with positive and negative signs in the silkscreen, and the solder pads themselves are round for negative and square for positive.

Install the 10-pin Eurorack power header P4. It is not polarized in the horizontal plane. However, if it has shorter legs on one side, then those are the ones that should go through the PCB (leaving the longer legs sticking up to mate with the connector on the power cable), and if it has tin plating on one end of the pins and gold on the other, then the tin side should be the one soldered through the board. Secure the header carefully to the board, possibly with tape, before soldering it. It is easy to accidentally solder it at an angle, which is a difficult error to fix and may cause trouble when you later attach the power cable.

Note that Eurorack power connections are polarized even if the connectors are not. The cables are usually grey ribbon type with a red stripe along one side indicating pin 1, which carries -12V power. For most modules including the MSK 006, the red stripe should be at the *bottom* when the module is mounted vertically in a case. On the MSK 006, the correct location of the -12V supply is also marked with the text "-12V" and an arrow on both sides of the PCB silkscreen. This module is also protected (by the Schottky diodes you just installed) from damage in case of a reversed power connection; if you connect the power backwards and nothing else is wrong, then the module will not power up but will be fine once you connect the power correctly. However, many other modules are not so protected, and it is dangerous to get into the habit of depending on protection diodes. Destroying a module by connecting power backwards is almost a rite of passage for Eurorack users.

Final assembly

The three boards will be referred to as Board 1, Board 2, and Board 3; they are the ones at the left, in the middle, and at the right in the final assembled module. They will be connected with wires that curve around the back of the module, as shown. This photo is from another prototype, and some of the components don't match the ones in the build instructions above, but the board to board wiring is the same.

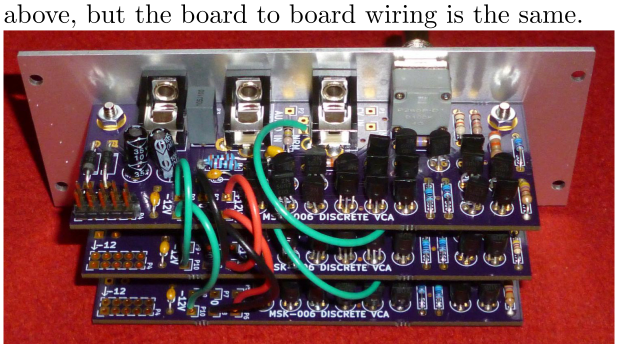

Cut lengths of hookup wire about 8cm to 10cm (3″ to 4″) and use them to connect P1, labelled "CV IN," on Board 1 to P12, labelled "NRM," on Board 2; and similarly, P1 on Board 2 to P12 on Board 3. This is for CV normalling: each board's CV automatically controls the next board if no inserted cable overrides it. Use a very short jumper (perhaps a chunk of cut-off resistor lead) to join P12 and P13 on Board 1; this sets the default input control voltage to zero.

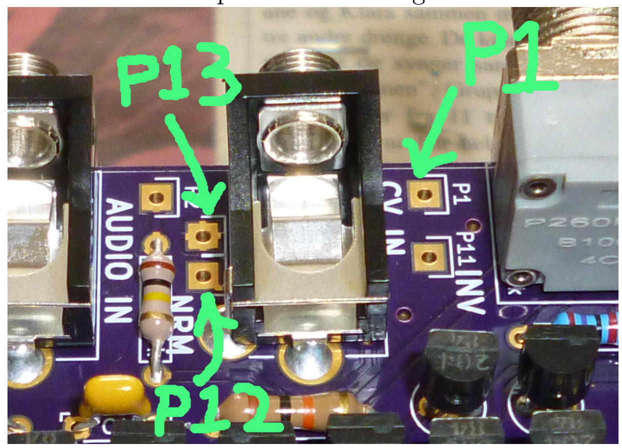

Connect the +12V wire pads (P5 and P6) on Board 3 to these pads on Boards 1 and 2 (one on each board). Similarly, connect the ground pads (P7 and P8, labelled "0") and the -12V pads (P9 and P10) on Board 3 to the corresponding pads on the other boards. These connections share power among the boards from the inlet on Board 3.

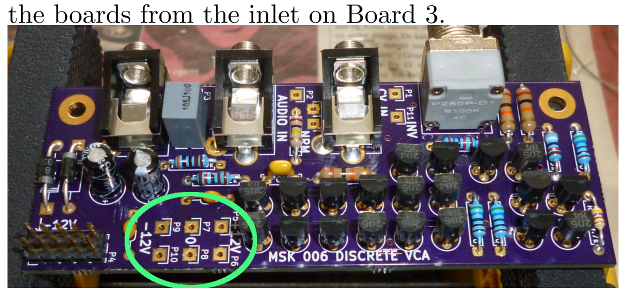

Assemble the stack of boards with the hex spacers, machine screws, and nuts. Attach the front panel, and tighten the nuts on the jacks and potentiometers. This time, use the spring washers and lock washers that came with the pots to make a good, tight installation. Add knobs. You have a finished module.

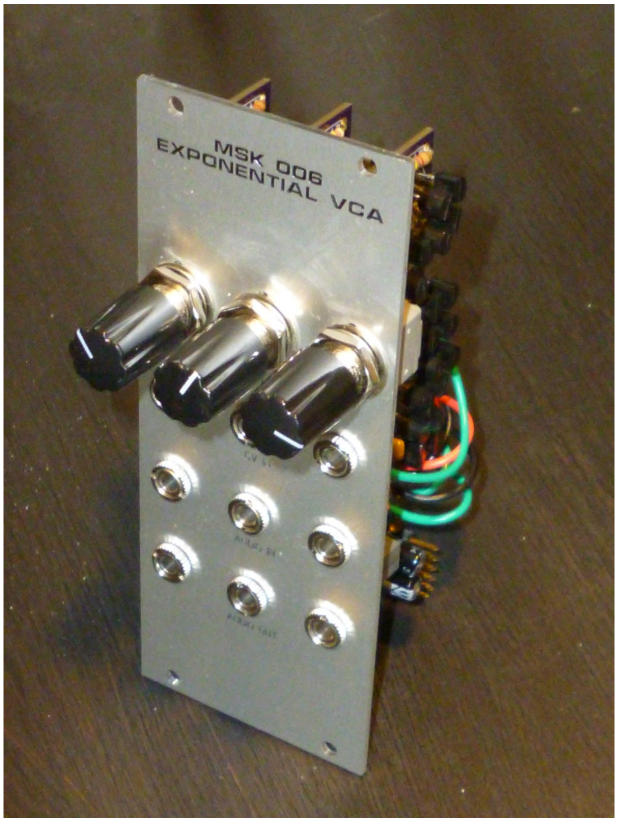
MSK 006
EXPONENTIAL VCA

JFET modification

With a few changes to the bill of materials and build procedure, the MSK 006 can be built to use N-channel JFETs for the differential pair instead of the standard PNP BJT. This makes a small difference to the sound. It's not clear that the sound with JFETs is *better* or *worse*; it's just *different*. If building a module with multiple channels, it might be fun to have the channels not all the same, so as to provide a wider sonic palette.

Doing the modification requires running some extra wires on the board, and it will be described in a little less detail than the standard build. Thus, it is appropriate for builders with a little more skill and experience, and if you are building more than one channel including both standard and modified builds, then you might want to do the standard builds first so that you'll have some practice before attempting the modified version.

A channel made with this modification will consume about four times as much power as one made with the original all-BJT circuit.

Other modifications described in this document (swapping the BJTs for other types, using DC-coupled output) are compatible with this one; you can do any combination of them.

Parts substitutions

For the JFET version, make the following changes to the bill of materials.

- Replace two of the PNP transistors (Q4 and Q10) with N-channel JFETs; see discussion below.
- Replace the 100kΩ resistor R7 with a series combination of a 91kΩ resistor and a 0.47μF capacitor.
- Replace the 4.7kΩ resistor R8 with a 1.2kΩ resistor.
- Replace two of the 560Ω resistors (R9 and R10) with 12kΩ resistors.

There are also some necessary wiring changes, discussed below.

Selecting and matching the JFETs

It is quite hard for manufacturers to do accurate quality control when manufacturing JFETs, and so there will be a lot of variation of parameters among individual JFETs of the same type. That can be a problem in the modified MSK 006 circuit, where the matching affects performance. Thus, I strongly recommend testing and matching individual JFETs when doing this mod, instead of just using two of the same type and hoping for them to be close. Buy more of the transistors than you need (they are cheap) and plan to only use about a third of them. Unmatched JFETs can be saved for use in other projects that don't require matching. If you have bought an MSK 006 JFET mod kit, then it should come with pre-matched transistors and you can skip the rest of this section—just be sure to use the JFETs in the pairs that are labelled in the kit packaging, instead of mixing up those from two or more channels.

Most low-power N-channel JFETs should work well in this circuit, assuming suitable matching. I used type MPF102 for my prototype because I had some old ones left over from another project, but I think that type is no longer manufactured. You can still find them for sale as new old stock because they were extremely popular and only recently discontinued; you can probably also find Chinese recent-manufacture equivalents which are in some sense "counterfeit" but which will perform well in this circuit when properly matched. We are not pushing the MPF102 anywhere near the limits of its specification, so even a poor imitation will probably be just fine. *Don't* pay a premium for "vintage" parts. You'd be better off buying a more recent type that is still in production.

Reasonable substitutions include: 2N3819 (near-identical specs to MPF102, but it may also be hard to find); J112 (readily available, specs seem okay but are not identical); J113 (also easy to find, but cutoff voltage and optimal current level may be too low); J111 (designed for higher power levels, should work but the gain may be less than optimal). This is not an exhaustive list. Check the data sheets for the exact pinouts, which vary among types. Note that matching and selection of the individual transistors is more important for this circuit than the type number. The two transistors in the pair should be the same type, but the natural variation among individual transistors is a fair bit more than the variation expected between reasonably equivalent types. For some types, it may be that *some* but not *all* individual transistors will be suitable.

Connect each JFET in the following circuit and measure the voltage as indicated.

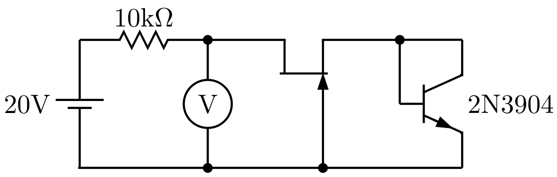

If you don't have a power supply adjustable to 20V, then 24V (such as the ±12V rails of a synthesizer supply, ignoring ground in the middle) will do. The specific bipolar transistor used is not critical, but you should use the same one between tests of different JFETs, so that you will only measure variations in the JFETs and not in the BJTs too. The measured voltage in this circuit should be between 1.2V and 10V for the JFET to be usable in the MSK 006, and it is likely to be near the low end of that range. I tested nine MPF102 JFETs and saw voltages ranging from 1.23V to 1.82V, including two very good pairs around 1.25V and 1.45V.

Choose two JFETs with measured voltages in the matching circuit as near as possible to each other.

Wiring changes

The JFETs should be wired in place of Q4 and Q10 so that equivalent pins connect to equivalent pins, bearing in mind the opposite polarity of the N-channel JFET compared to the PNP transist. That is, wire the source of the JFET where the collector of the BJT would go; the gate of the JFET where the base of the BJT would go; and the drain of the JFET where the emitter of the BJT would go. *Check the data sheets for your transistors* to be sure of which pin is which; this varies a lot among different JFET types. On many JFETs, the source and drain are interchangeable; this will be noted on the data sheet if it applies, and it may give you more options for wiring. Nonetheless, the necessary orientation of the JFET package will probably *not* match the silkscreen markings on the PCBs, which were designed for BC557C transistors.

The ends of R9 and R10 which would go to ground in a standard build must instead go to -12V power in the modified build. I achieved this in the prototype by soldering only the base/gate ends of each resistor to the board and creating a floating connection among the other ends and a piece of wire which ran to a spare -12V pad (P9) elsewhere on the board, as shown. Without this wiring change the transistor gates will be forward-biased with about 1mA of current, which may damage them and will certainly prevent the module from operating properly.

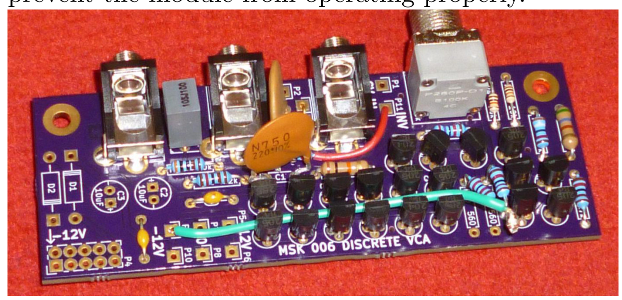

The series combination of a 91kΩ resistor and 0.47μF capacitor, which replaces R7, should be connected from P2 (audio in) to P11 (which would be "inverting in" in a standard build). This is because the gain characteristic of the JFETs is inverted relative to what BJTs would do; so connecting input to Q4 instead of Q10 inverts it again and keeps the overall behaviour of the module non-inverting. Using these pads on the board also allows a little more space for the necessary wiring, instead of trying to fit two components into a footprint laid out for just one. I used a big old-fashioned disc capacitor in the prototype shown in the photo, but I don't recommend that; it was hard to fit into the module without interfering with the next board above, and I ended up damaging the coating while bending the leads into shape. A modern ceramic leaded-chip style of capacitor might be a better choice. A film capacitor in this value would likely be too physically large to fit well, and is probably overkill for this undemanding application, but if you want to use one and it fits, it will work fine from an electrical point of view.

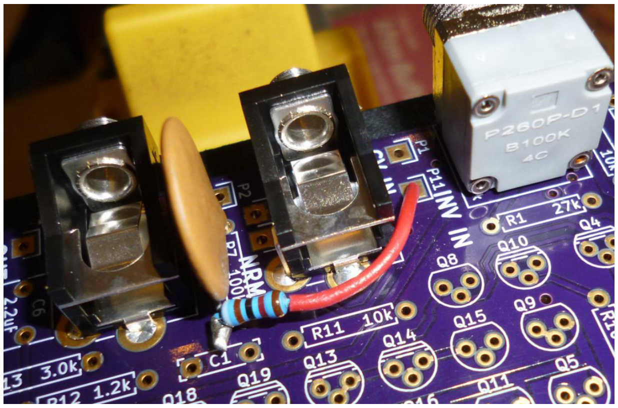

In this photo, I slipped a piece of insulation from a chunk of hookup wire over the resistor lead that goes to P11, so as to reduce the chance of it shorting on other things, and I plugged the capacitor into the pad for R7 nearest P2 instead of actually into the P2 pad, because that made the angle work better. These two pads are directly connected by a trace on the board, so either may be used.

Patch ideas

Here's a classic subtractive-synthesis patch. A MIDI interface controls an oscillator, filter, and envelope generator, and the MSK 006 applies the envelope to the filtered audio.

The same modules with a couple more connections: MIDI velocity controls a second channel of the MSK 006, so that harder-struck notes will be louder.

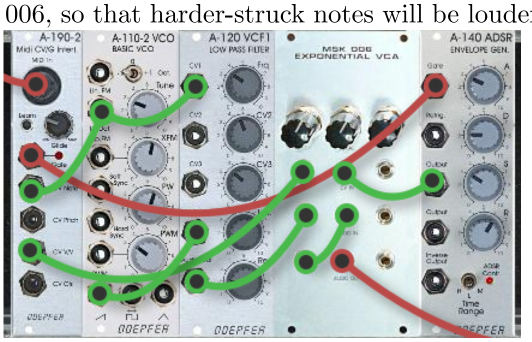

Synth hobbyists call a module "passive" if it is powered by the input signals, regardless of whether it's really passive in the electronic sense of containing no active devices. When you use such modules, they inevitably reduce the power level of the signal; and an MSK 006 channel without a control voltage can be used as a simple amplifier to raise the signal back to standard modular level.

The MSK 006 can be used to raise and lower the levels of signals, for interfacing modular effects to external equipment that runs at less than modular level. The output can also drive headphones and small speakers, at least to moderate volume.

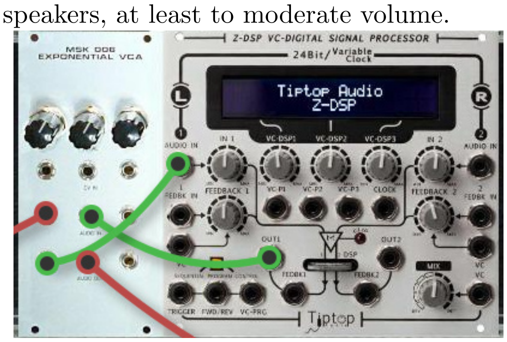

Circuit explanation

The core of this module is an OTA very similar to that in an LM13700. The control voltage feeds an exponential current sink; output from that goes through a PNP Wilson current mirror; and then a differential pair of PNP transistors splits the control current proportionally to the input voltage. The two halves of the control current each go through NPN current mirrors, and one of them goes through a second PNP current mirror, so that we end up with a complementary current source and current sink, both in linear proportion to the bipolar audio input voltage and in exponential proportion to the control voltage. Their difference is a voltage-controlled current representing the input multiplied by the CV-chosen gain.

The current from the OTA core drives a resistor which converts it back to a voltage, and then an NPN and PNP transistor in push-pull configuration amplifies the output. The current-to-voltage resistor is actually in series with two diode-connected transistors matching the output transistors, and that causes a discontinuity in the current-to-voltage response which in principle ought to be equal to and opposite from that introduced by the class B output stage, thus cancelling out crossover distortion.

Now, each of the sections in more detail.

CV processing

Control voltage from the outside world is applied directly to the base of Q1, which is connected as an emitter follower. It will conduct just enough current to lift its collector up to one diode drop below the input control voltage. That current is split between R2, which is the link to the rest of the circuit, and R1, which pulls the emitter below ground when the input control voltage is near or below zero. The relatively high value of R1 relative to R2 means that the emitter follower's current output is somewhat asymmetrical: it can source more current than it can sink, which means positive control voltages have a greater effect than negative ones. These resistor values were chosen by experiment for the best-sounding curve, even though it deviates a bit from a theoretically correct exponential.

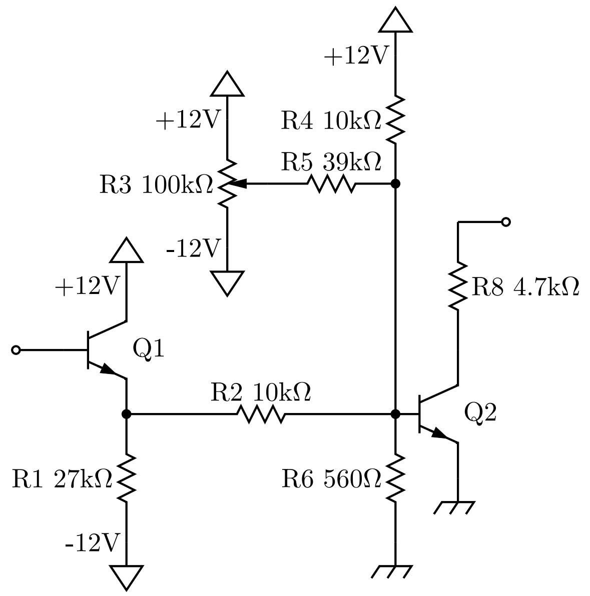

Multiple sources of control are mixed at the base of Q2: the buffered input control voltage from Q1 through R2, the "initial" knob setting from R3, through R5, and constant bias from R4 and R6, which set the range and scaling for the others. The resistor values set the relative weights of these influences, so that the overall total voltage at this node will be around 0.65V (the base-emitter diode drop) plus or minus about 60mV per 10dB of control current variation (note that is 20dB of overall power gain through the amplifier).

The transistor Q2 then does its usual exponential-response thing, conducting a current through its emitter that increases by a factor of 10 for every 60mV increase in base voltage. The resistor R8, connected between this transistor and the input of a current mirror at a fixed voltage slightly below the power line at about 11.5V, serves to limit the maximum possible control current to about 2.4mA.

Wilson current mirror

There are four Wilson current mirrors (two PNP and two NPN) in this circuit, encompassing most of the transistors in the design. I describe the mirror using the reference designators of just one of the NPN ones, but the other three work the same way, flipped over in the case of the PNP mirrors. This section's function is to provide an input held at a constant voltage, into which other sections can feed a current, and then source or sink an equal current in the opposite direction at a high-impedance output, where the voltage is allowed to vary but the current is fixed equal to the input current. Thus, one way it can be thought of is as very low to very high impedance transformer.

Correct operation of this kind of circuit requires that the transistors match. That is one reason for using the three-transistor Wilson configuration here instead of a simpler two-transistor mirror; the third transistor improves the accuracy when the transistors do not match well, as is typical of discrete transistors. IC amplifiers are often designed with two-transistor mirrors, making use of the better matched transistors available in an IC design.

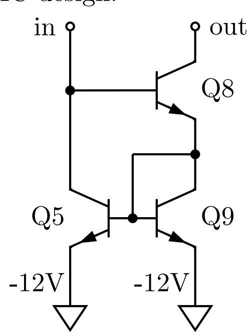

First, note that the input is connected to the base of Q8, which has its emitter connected to the bases of Q5 and Q9, which in turn have their emitters connected to negative power. That right there means the input voltage is fixed at two diode drops, roughly 1.2V (with a small variation logarithmic in the input current), above the negative supply. It looks like a constant voltage with low impedance into which an external source can dump arbitrary amounts of current.

Suppose the external source does feed some current into the input. Some of the input current (actually very little, but we'll get to that) goes into the base of Q8, causing it to conduct a much larger current from collector to emitter. Now, where will Q8's collector current go?

The collector current from Q8 must go into the bases of Q5 and Q9, and through the collector of Q9. Q9, being diode-connected, will conduct as much current as it can through its collector, and bring its collector and base (which are the same node in the circuit) to whatever voltage it takes to make a transistor of this type conduct that much current.

But then that voltage is also applied to the base of Q5, causing it to conduct as much current from collector to emitter as Q9 is conducting. And that current has to come out of the share that would otherwise be going into Q8. So, if Q5 (and therefore Q9 and Q8) tries to conduct more current than the input, not only will that current have nowhere to come from, but it will shut off Q8 and therefore the entire section. Q5 ends up stealing about 99% of the input current, leaving just enough remaining to feed the base of one transistor of this type (regardless of whether it's Q5 itself, or Q8 or Q9) when it's conducting this much collector current. The feedback among the transistors means that the output current ends up being almost exactly equal to the input current, while the input current and output voltage are free to vary over a wide range.

To reiterate: input current causes Q8 to conduct. Q8 tries to conduct h_{FE} times the input current, traditionally thought of as about 100 (actually more, in the MSK 006's BC557C PNP transistors). Almost all of that current goes into Q9's collector, causing it to develop a voltage on its base that drives Q5 to pull the same current out of the input and away from Q8's base, which reduces the current through Q8 until the whole thing stabilizes. In the stable state, Q5 and Q9 are conducting identical collector and base currents, with the collector current 99% or more of the input and the base 1% or less. Q5's collector is driven by the input, Q9's collector is driven by the output, the base currents are added and drawn through Q8, and then Q8's base draws one transistor's worth of base current from the input while the other unit of base current comes from the output through Q8's collector. The end result is that input and output *both* drive one emitter current and one base current, which should make them identical. The input is held at constant voltage two diode drops above negative power (low impedance), and the output voltage is free to vary over whatever range Q8's voltage rating allows (high impedance).

For it to work perfectly, Q5 and Q9 should be as similar to each other as possible; they should be at the same temperature, because their current-voltage function is (in fact, is highly) temperature-dependent; and to a lesser extent, Q8 should be similar to them and its gain should be high. One remaining source of inaccuracy is the fact that Q8's collector includes two base currents while the other transistors' collectors only carry one, which means that its own base current will be off from what it should be by a factor of 1% or less—not enough to worry about and dwarfed by other sources of error.

There are other things going on in the circuit, too. By carrying most of the output voltage, Q8 keeps Q9 and Q5 at almost the same collector voltage; differing by one diode drop, but not by most of the output voltage, as would be seen in a two-transistor mirror. That reduces the Early effect (deviation from perfect current-sink behaviour) and self-heating (which makes temperature matching harder) in the core mirror transistors. These effects still occur relative to Q8, but are reduced basically by a factor of h_{FE} because Q8 is not involved in the log/exponential mirroring process.

The transistors Q9 and Q5 are themselves a classic two-transistor current mirror in the backward direction (from the overall section's output back to its input). Adding Q8 is the innovation of Wilson, which makes the current mirror much more accurate, and is important given we're building it out of discrete transistors that are likely to be poorly matched. In principle, one could substitute the two-transistor Q9/Q5 current mirror recursively with a Wilson three-transistor mirror, to make a four-transistor mirror, and carry this out to as many transistors as desired. There are a number of serious disadvantages to that, though—among others, that each added transistor adds a diode drop to the input voltage. In practice, it is rarely beneficial to go beyond three when building a current mirror out of bipolar transistors. Some MOS mirror designs used in ICs are much more complicated.

Differential pair

The exponential converter *sinks* a current that increases exponentially with input control voltage, then it goes through the first PNP current mirror (Q3, Q6, and Q7) which acts as a current *source*. That current feeds the differential pair Q4 and Q10, a very simple but important section of the VCA. It is here that the variable gain actually occurs.

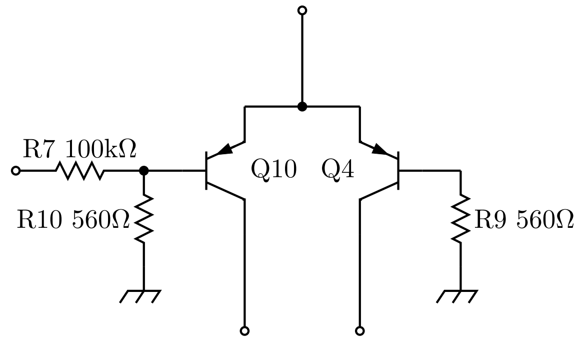

The bases of the transistors are always held near ground by the low resistances of R9 and R10. Their emitters are tied together, and that node is fed by the current mirror output, free to vary in voltage. Then in the absence of any audio input through R7, these transistors will equally split the control current, sending half of it to each of their collectors, with the emitter voltage floating one diode drop above ground. The base currents will be equal, at a small fraction of the control current, and basically negligible.

The input side of R7 is connected directly to the audio input jack. Suppose we apply some voltage there. The voltage is divided by about 180 (the ratio of the R7/R10 voltage divider) and applied to the base of Q10. Now there is an *additive* voltage difference in the base voltages of the transistors, and since their emitters are tied together, that becomes an additive difference in their base-collector voltage differences. By the exponential characteristic of the transistors, the additive voltage on the audio input becomes a *multiplicative* difference in the currents leaving the section through the transistor collectors. Adding or subtracting about 3.2V of voltage (that is, the 18mV per octave characteristic of the transistors, times 180) on the audio input will multiply or divide the current ratio by 2. But the total current exiting through the two collectors is still set by the control current from the exponential converter.

So the differential pair generates two positive currents whose sum is determined exponentially by the control voltage input but whose ratio is determined linearly by the audio input. The difference between these two currents is a bipolar (could be positive or negative) current that will be exactly proportional to the voltage we want to get out of the module.

There follow three current mirrors (as described in the previous section) which together change that difference current into a more usable form. The differential pair can only provide positive currents, its outputs need to be at least a little below ground voltage, and it performs more accurately if its output voltages can be kept constant. The Q11/Q15/Q16 mirror reflects one of the positive output currents into a negative current, supplied from near the negative power rail. The Q5/Q6/Q9 and Q12/Q14/Q15 mirrors reflect the other current twice, so it will be a positive current again but supplied from near the positive rail instead of near ground. The collectors of Q14 and Q15 are connected directly together to create a bidirectional current source/sink whose voltage can freely range about ±10V. This is substantially the same topology used in the LM13700 and similar chips.

Any imbalance between the transistors of the differential pair will cause the output current, at zero input voltage, to be non-zero. The nonzero current will create a DC offset when fed into the output stage, but there are also sources of offset within the output stage itself. It was a conscious decision in designing this circuit not to attempt to remove the offset here. One advantage is that this way the module won't require trimming; it was also desired to embrace rather than eliminating some of the quirks associated with the discrete design. If I wanted to remove the offset in the input, though, I would do it by feeding some current into the inverting input (base of Q7) via a trimmer potentiometer and series resistor. Another design often seen is to replace the resistors from both transistor bases to ground by the two halves of a trim pot, with ground connected to the wiper; that is in some sense more elegant because it applies the compensation to both bases, but it has the disadvantage of making the gain change with the offset trim.

Output stage

The differential pair and current mirrors produce a signal representing what we want to be the output voltage of the module, but in the form of a current at high impedance. We want to convert that linearly to a voltage; and linear current to voltage conversion is exactly what a resistor does. The actual circuit is a little more complex than just a resistor, but it has R11 at its heart to perform the conversion.

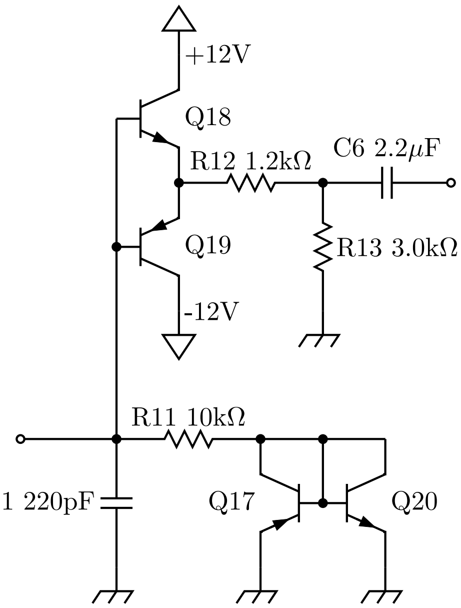

Temporarily ignoring C1, Q17, and Q20, we have current from the OTA mirrors going through R11 to ground. The voltage across R11 will be exactly proportional to the current, whether positive or negative. Then Q18 and Q19 are both emitter followers, copying the voltage across R11 to their shared emitter. When their shared base goes above the output voltage on the shared emitter, then Q18 will start to conduct and pull that voltage up; Q19 is at that point reverse-biased on its emitter-base junction and presents a high impedance. On the other hand, when the shared base goes below the output voltage on the shared emitter, then Q18 is turned off and Q19 starts to conduct. The two transistors work together to keep the output voltage the same as the input, while drawing very little current from the input (so as not to disrupt R11's current to voltage conversion) and supplying as much current to the output as is needed. This is a classic push-pull pure class B amplifier design.

There are some wrinkles, however. In order to turn on Q18, the voltage on R11 must actually go a diode drop *above* the output voltage; and to turn on Q19, it must go a diode drop *below*. So unless we do something more, there will be a gap in the middle through which the input voltage can vary without any power being supplied to the output. On an oscilloscope, it can look like a chunk was cut out of the middle of the waveform and then the halves pasted together. This effect is called "crossover distortion" (named for voltages crossing over zero, nothing to do with speaker crossovers) and it sounds terrible.

Crossover distortion shows up as undesired odd harmonics at a more or less constant voltage level regardless of the input, thus greater as a proportion when the input level is low. That is the opposite of most distortion processes in nature (which show up more at high levels) and so it is hard for our ears to hear it as natural "warmth" integral to the desired signal; instead, it sounds like an intrusive completely separate signal added to the desired signal. Crossover distortion is the bugaboo of pure class B amplifiers and one of the main reasons they are seldom used.

One way to prevent crossover distortion is to make the amplifier class AB instead: use a biasing arrangement that keeps both transistors turned on a little bit at all times, so that there is never an issue of the input voltage having to cross a gap between turning one off and turning the other on. Typically, that would involve using a network of resistors and diodes to keep the NPN base one diode drop above the input voltage and the PNP base one diode drop below. The MSK 006 does something a little different, making use of the fact that the input to this circuit is actually a current from a pair of high-impedance mirrors, not a voltage. Here the compensating diodes are Q17 and Q20, in series with R11, and the bases of the power transistors remain directly connected to each other.

Note that current from the OTA must go not only through R11 but also the parallel combination of Q17 and Q20, which are connected as diodes in parallel, in opposite directions. The voltage drop from the top of R11 to ground is then 10kΩ times the input current, for R11, *plus* one diode drop either side of ground. When the input current changes sign, either Q17 or Q20 will turn off, the other will turn on, and the voltage will jump two diode drops as fast as the transistors can switch (which is much faster than audio). Thus, the waveform at the top of R11 is predistorted, with a gap inserted around zero, in a way that should (if Q17 matches Q19 and Q20 Q18) compensate exactly for the crossover distortion of the class B stage.

In practice, it is not perfect. The output transistors may not match exactly with their corresponding compensation transistors at the best of times, and they are being operated at different voltage and current levels, so that the necessary base voltage to compensate the distortion may not perfectly match the compensation provided by the diode-connected transistors. There is also a nonzero amount of time required for each transistor to switch, which can result in a notch or step in the waveform while one part of the circuit waits for another. Nonetheless, it works well in practice, reducing the crossover distortion to a level that is not annoying. What distortion remains can be called a "quirk" rather than a failure of the design, bearing in mind that this is meant to be a VCA *that you can hear*.

The capacitor C1 is intended to help reduce oscillation and distortion. I have not observed, and don't expect, any particular issues with parasitic oscillations in this circuit, but (especially bearing in mind that I'm putting it out as a DIY project for people to build with more or less random components) it seems worth using some caution. The time constant of a 220pF capacitor against the 10kΩ impedance at that point works out to an oscillation frequency of about 72kHz. If by whatever means a signal higher than that should appear at that point in the circuit, then it will be dumped into ground instead of being available for amplification by the output stage or (via unintended feedback paths) the input stage. It also has an effect, which I have not analysed carefully, on the overall distortion, by changing the shape of the crossover notches. The specific value of 220pF was chosen by trial and error with the actual circuit to get the least amount of total harmonic distortion on sine waves in the mid-audio range. However, it doesn't actually have much effect at all; anything from 0pF (omitting the capacitor entirely) to 2000pF would probably be acceptable.

The class B output transistors can drive their shared emitter close to either power rail, something like 23V peak-to-peak, with (in principle, using the specified transistor types) as much as 200mA of current in both directions before they burn up, and no current limiting in place to prevent it from delivering even more power. That could certainly cause some problems if connected to a power rail, an overly sensitive input, or similar. So to prevent the MSK 006 from frying other equipment, the resistor R12 is connected between the output transistors and the actual output. It limits current to about 10mA, which is still a lot.

Remember that there is a significant voltage offset as a result of imbalances among the transistors. To prevent that from appearing on the output, C6 serves as an AC coupling capacitor. Note the original design value was 2.2μF, but I actually built my prototypes with 1.0μF because the larger film capacitors are hard to source; I may end up changing the silkscreen to read 1.0μF in some future revision of the board. Home builders, use whichever is convenient—the capacitance value is not critical—but I would recommend no less than 0.5μF, and I would *not* recommend using an electrolytic because of the unpredictable direction of the bias.

Adding C6 creates a new problem with *current* offset. Leakage through Q18 and Q19 (driven by the OTA's offset current) when they are both nominally turned off will probably not be identical, and so the output tries to drive what may be a relatively significant constant current when the input is at zero. My estimate based on the observed time constants is that it's about 1μA; I haven't attempted exact measurements.

Unless we do something, if the signal is silent or quiet for a few seconds then the output leakage current will charge up C6, bringing the emitters of Q18 and Q19 some significant distance from ground. Then the crossover distortion compensator no longer works, and we get severe distortion *until* the next loud bit in the signal, which causes the transistors to turn fully on long enough to drain away the charge, allowing the compensation to work again, at which point the distortion goes away and the amp sounds good, even on quiet signals, until there is another stretch of a few seconds of quiet, the capacitor recharges, and the distortion returns.

This was an interesting problem to debug. I had the circuit working well with a DC-coupled output, switched to AC coupling thinking that would make no difference, and then had to figure out why it *did* make a difference. Anyway, R13 solves the problem by providing a DC path for the leakage current to ground, so that it will result in a constant-voltage DC offset (to be blocked by C6) instead of a constant current that can charge C6 to arbitrary voltages. R13 also forms a voltage divider with R12, limiting the maximum signal output to something a little more reasonable than 23V peak-to-peak.

MSK 006 EXPONENTIAL VCA

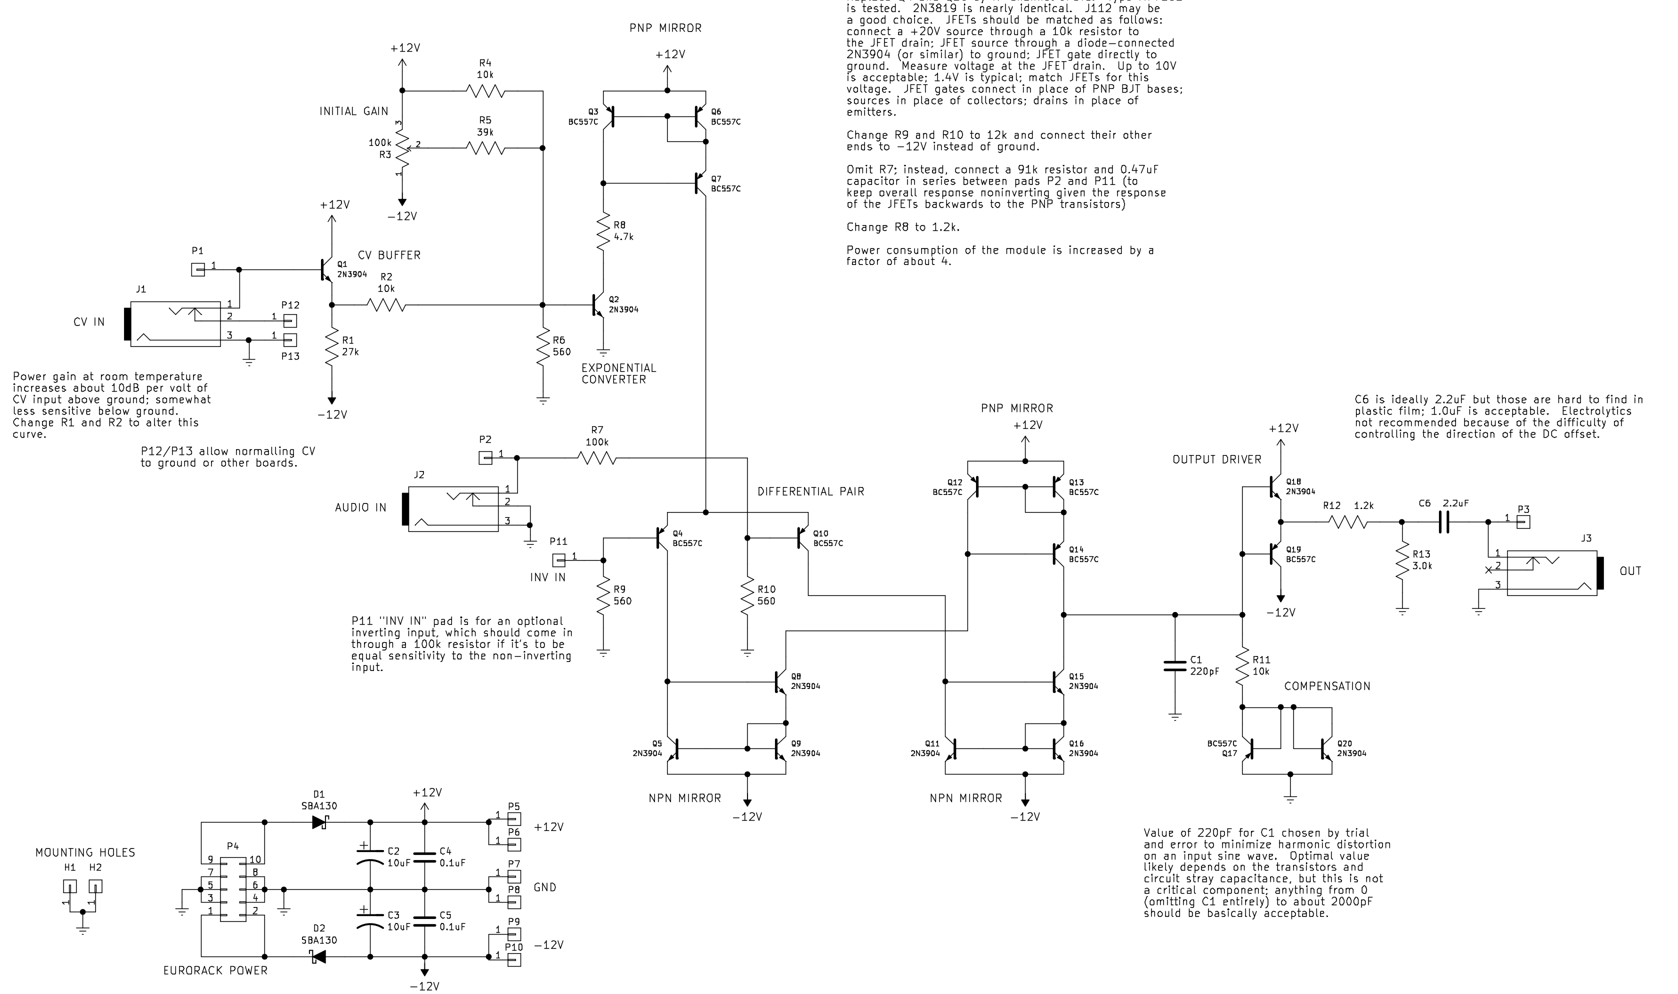

Transistor model is not critical and different models may be mixed as long as they are the right NPN/PNP type, but parameters should match:
* within any three-transistor "mirror," as marked
* within the two-transistor "differential pair"
* between each output driver transistor and the corresponding NPN/PNP "compensation" transistor

Also, beware of different pinouts!

Optional JFET modification:

Replace Q4 and Q10 by N-channel JFETs. Type MPF102 is tested. 2N3819 is nearly identical. J112 may be a good choice. JFETs should be matched as follows: connect a +20V source through a 10k resistor to the JFET drain; JFET source through a diode-connected 2N3904 (or similar) to ground; JFET gate directly to ground. Measure voltage at the JFET drain. Up to 10V is acceptable; 1.4V is typical; match JFETs for this voltage. JFET gates connect in place of PNP BJT bases; sources in place of collectors; drains in place of emitters.

Change R9 and R10 to 12k and connect their other ends to -12V instead of ground.

Omit R7; instead, connect a 91k resistor and 0.47uF capacitor in series between pads P2 and P11 (to keep overall response noninverting given the response of the JFETs backwards to the PNP transistors)

Change R8 to 1.2k.

Power consumption of the module is increased by a factor of about 4.

Power gain at room temperature increases about 10dB per volt of CV input above ground; somewhat less sensitive below ground. Change R1 and R2 to alter this curve.

P12/P13 allow normalling CV to ground or other boards.

P11 "INV IN" pad is for an optional inverting input, which should come in through a 100k resistor if it's to be equal sensitivity to the non-inverting input.

C6 is ideally 2.2uF but those are hard to find in plastic film; 1.0uF is acceptable. Electrolytics not recommended because of the difficulty of controlling the direction of the DC offset.

Value of 220pF for C1 chosen by trial and error to minimize harmonic distortion on an input sine wave. Optimal value likely depends on the transistors and circuit stray capacitance, but this is not a critical component; anything from 0 (omitting C1 entirely) to about 2000pF should be basically acceptable.

$Id: vca.sch 4814 2016-05-02 14:43:42Z mskala $

1/1

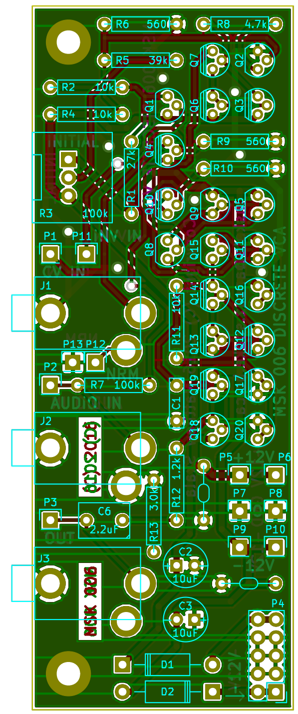

MSK 006 EXPONENTIAL VCA

$Id: vca.kicad_pcb 4741 2016-03-29 15:08:53Z mskala $

1/1

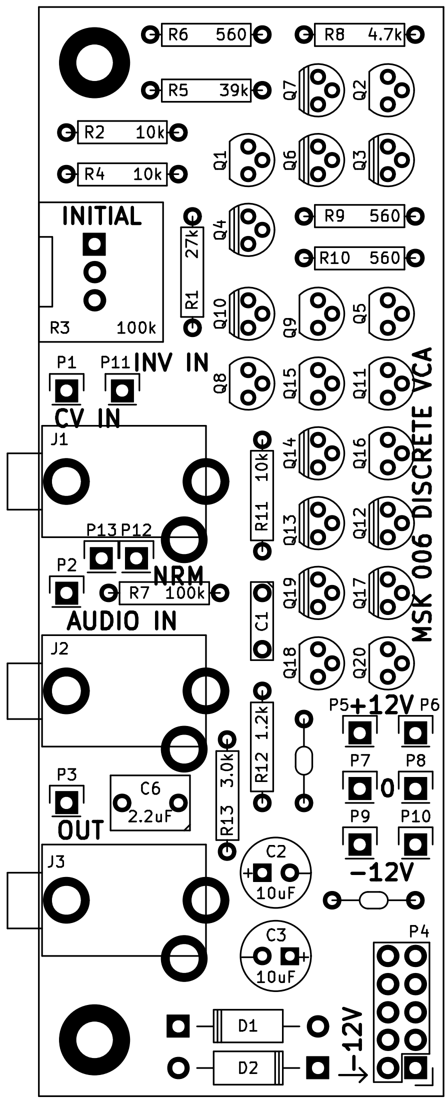
MSK 006 DISCRETE VCA
R6 560
R8 4.7k
R5 39k
R2 10k
R4 10k
INITIAL
R3 100k
R1 27k
R9 560
R10 560
Q1
Q2
Q3
Q4
Q5
Q6
Q7
Q8
Q9
Q10
Q11
Q12
Q13
Q14
Q15
Q16
Q17
Q18
Q19
Q20
P1
P11
INV IN
CV IN
J1
R11 10k
P13
P12
NRM
P2
R7 100k
C1
AUDIO IN
J2
R12 1.2k
R13 3.0k
P3
C6
2.2uF
OUT
J3
C2
10uF
C3
10uF
P5
+12V
P6
P7
P8
0
P9
P10
-12V
P4
D1
D2
-12V

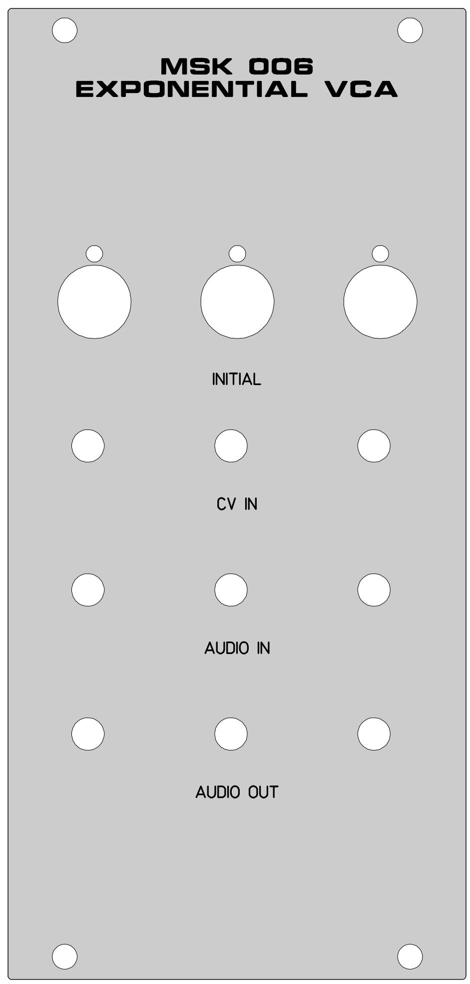

MSK 006
EXPONENTIAL VCA
INITIAL
CV IN
AUDIO IN
AUDIO OUT