MSK 008 Dual VC Octave Switch

North Coast Synthesis Ltd.
Matthew Skala

September 22, 2023

Contents

General notes

This manual documents the MSK 008 Dual VC Octave Switch, which is a control voltage processing module for use in a Eurorack modular synthesizer. Each of the two channels in the module contains a specialized quantizer, a toggle switch that adds or subtracts a one-volt offset, and a precision voltage adder with optional subtraction function. The module's main function is to transpose a melody represented by a control voltage, but it can also be used for several other purposes.

Specifications

The impedance of the quantizer inputs is high, several hundred kiloohms at least, depending on the setting of the offset switch. For the other inputs, the impedance is 100kΩ per input; when a signal drives two inputs through normalization, it will see a 50kΩ impedance. The impedance of the outputs is very low, with current limited by 1kΩ resistors inside the op amp feedback loops.

Any input voltage between the power supply rails (nominally ±12V) is safe for the module; output voltages are limited by the capabilities of the op amps to about ±10V and will clip if the input voltages sum to a result outside that range.

Shorting any input or output to any fixed voltage at or between the power rails, or shorting two to each other, should be harmless to the module. Patching the MSK 008's output into some other module's output should be harmless to the MSK 008, but doing that is not recommended because it is possible the non-MSK 008 module may be harmed.

Although primarily intended to operate on DC control voltages, this module should be usable at all audio frequencies. Ultrasonic frequencies are rolled off to ensure op amp stability.

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 cause or suffer any damage. Some other kinds of power misconnection may possibly be dangerous to the module or the power supply.

In normal operation the peak current demand of this module is 45mA from the +12V supply and 50mA from the -12V supply. Placing a heavy load on the outputs (for instance, with so-called passive modules) will increase the power supply current.

Voltage modification

This circuit was designed for ±12V power and will not work properly on ±15V power unless modified. To modify it for ±15V, as well as using the proper power connector and making sure all components are rated for the increased voltage, make the following resistor changes:

- change R6 and R7 (voltage regulator ballast) from 1.8kΩ to 2.4kΩ;
- change R13 and R14 (quantizer input weights) from 75kΩ to 62kΩ;
- change R49 and R51 (LED current control) from 1.2kΩ to 1.8kΩ; and
- change R50, R52, R53, and R54 (LED current control) from 910Ω to 1.2kΩ.

I have calculated but not tested these resistor changes.

Source package

A ZIP archive containing source code for this document and for the module itself, including things like machine-readable CAD files, is available from the Web site at https://northcoastsynthesis.com/. Be aware that actually building from source requires some manual steps; Makefiles for GNU Make are provided, but you may need to manually generate PDFs from the CAD files for inclusion in the document, make Gerbers from the PCB design, manually edit the .csv bill of materials files if you change the bill of materials, and so on.

Recommended software for use with the source code includes:

- GNU Make;
- LaTeX for document compilation;
- LaTeX.mk (Danjean and Legrand, not to be confused with other similarly-named LaTeX-automation tools);

- Circuit_macros (for in-document schematic diagrams);
- Kicad (electronic design automation);
- Qcad (2D drafting); and
- Perl (for the BOM-generating script).

The kicad-symbols/ subdirectory contains 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.

The package is covered by the GNU GPL, version 3, a copy of which is included in the file COPYING.

PCBs and physical design

The enclosed PCB design is for two boards, each $3.90'' \times 1.50''$, or 99.06mm×38.10mm. They are intended to mount in a stack parallel to the Eurorack panel, held together with M3 machine screws and male-female hex standoff hardware. See Figure 1. Including 18mm of clearance for the mated power connector, the module should fit in 43mm of depth measured from the back of the front panel.

Functional description

Users will probably want to think about this module in terms of the applications in which it is *used*: octave switching, wavefolding, mid-side encoding, and so on. Even the name "octave switch" refers to one application, and some of the others are described in the section on suggested patches. But all these applications, and those yet to be invented, can be understood in terms of what the module *actually does*, as described here without direct reference to applications. See the block diagram in Figure 2.

There are two channels. Each channel has a QUA input, which goes into a quantizer that rounds it to the nearest whole volt in the range -2V…+2V. So any input voltage from -12V to -1.5V will round to -2V; any input voltage from -1.5V to -0.5V will round to -1V; any input voltage from -0.5V to +0.5V will round to 0V; and so on. There is also a three-position toggle switch, which can add or subtract one volt at the quantizer *input*; the quantizer output remains in

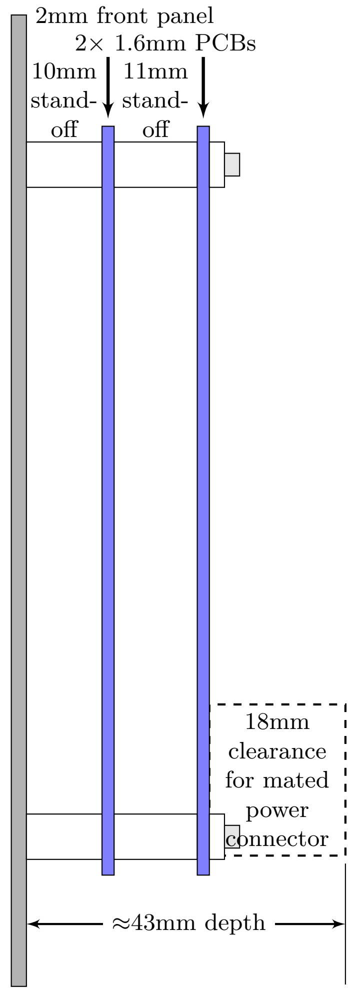

Figure 1: Assembled module, side view.

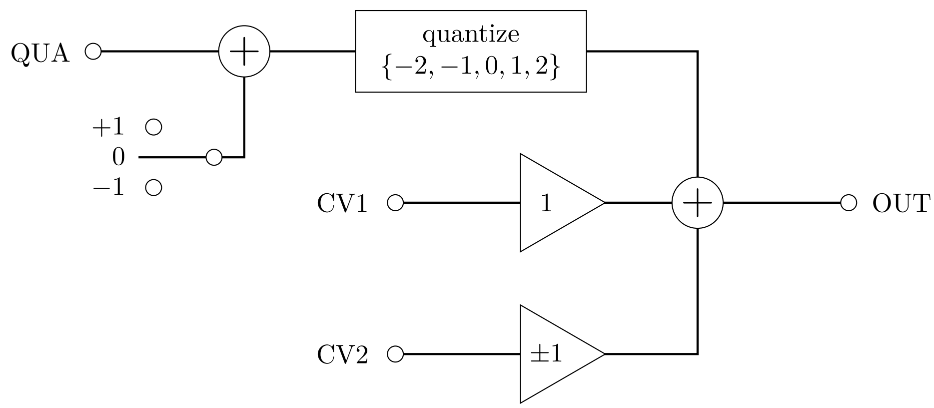

Figure 2: Block diagram (one channel).

the range -2V. . . +2V.

Although the quantizer output voltages are intended to be accurate (typically 1mV error with decent adjustment; this specification is not guaranteed), the input thresholds are not meant to be very accurate. Whole volts will round correctly, but the exact positions of the boundaries may not be exact half-volt values, and may shift with the position of the toggle switch. Allowing some tolerance on these was a necessary trade-off to fit the module in 8HP with the desired technology.

When the quantizer output (not the final output of the channel) is positive, whether +1V or +2V, the channel's LED glows red; when it is negative, green; when zero, the LED remains dark.

The precision adder sums the quantizer output, the CV1 input, and either the CV2 input or the negative of the CV2 input. The result of the summation appears at the channel output.

Whether a channel adds or subtracts the CV2 input is determined by a jumper built into the circuit board. The default, standard for our assembled modules and any built from kits without modifying the boards, is for the left channel (as viewed from the front when the module is installed) to add and the right channel to subtract. This selection for a channel could be changed during build by cutting the jumper trace and replacing it in the opposite direction with a blob of solder. A more ambitious project might cut the trace and wire an SPDT switch to the three contact points provided, allowing front-panel selection between the two modes. But North Coast Synthesis Ltd. has no current plans to market an "expander" for doing that; it is left as an idea for more advanced hobbyists to pursue on their own.

The CV1 and CV2 inputs of each channel, but not the QUA input, are normalized to and from the corresponding inputs on the other channel. Plugging a cable into CV1 on either channel will drive both CV1 inputs if there is no cable plugged into the other CV1, and the same for CV2. Thus, in a default jumper configuration, plugging cables into CV1 and CV2 on one channel will give the sum and difference of the two input voltages on the two channel outputs.

Each QUA input is normalized to 0V. Plugging a cable into one of the QUA inputs with nothing on the other end breaks this normalization and leaves the comparator inputs to float; that may have unexpected effects. Typically, the unconnected input will drive the quantizer to +2V with the toggle in the 0 or +1 positions and −2V with the toggle in the −1 position, but that behaviour is not guaranteed and half-patching the input this way is not recommended. Some module "trigger" outputs, notably those on the Befaco Rampage, do not produce any specific voltage but expose a high impedance like unpatched cables when they are not triggering; those may produce unexpected results if patched into the MSK 008's QUA input.

Use and contact information

This module design is released under the GNU GPL, version 3, a copy of which is in the source code 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 sell this and other modules, both as fully assembled products and do-it-yourself kits, from my Web storefront at http://northcoastsynthesis.com/. Your support of my business is what makes it possible for me to continue releasing module designs for free. The latest version of this document and the associated source files can be found at that Web site.

Email should be sent to
mskala@northcoastsynthesis.com.

Safety and other warnings

Ask an adult to help you.

North Coast Synthesis Ltd. does not offer warranties or technical support on anything we did not build and sell. That applies both to modules built by you or others from the kits we sell, and to fully-assembled modules that might be built by others using our plans. Especially note that because we publish detailed plans and we permit third parties to build and sell modules using our plans subject to the relevant license terms, it is reasonable to expect that there will be modules on the new and used markets closely resembling ours but not built and sold by us. We may be able to help in authenticating a module of unknown provenance; contact us if you have questions of this nature.

For new modules purchased through a reseller, warranty and technical support issues should be taken to the reseller *first.* Resellers buy modules from North Coast at a significant discount, allowing them to resell the modules at a profit, and part of the way they earn that is by taking responsibility for supporting their own customers.

We also sell our products to hobbyists who enjoy tinkering with and customizing electronic equipment. Modules like ours, even if originally built by us, may be quite likely to contain third-party "mods," added or deleted features, or otherwise differ from the standard specifications of our assembled modules when new. Be aware of this possibility when you buy a used module.

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.

Residue from traditional rosin-based solder flux can result in undesired leakage currents that may affect high-impedance circuits. This module does not use any extremely high impedances, but small leakage currents could still reduce its accuracy. If your soldering leaves a lot of such residue then it might be advisable to clean that off.

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

This table is not a substitute for the text instructions.

Qty	Ref	Value/Part No.	
2	C5, C6	33pF	radial ceramic, 0.2″ lead spacing
6	C1–C4, C7, C8	0.1μF	axial ceramic
4	C9–C12	10μF	radial aluminum electrolytic, 0.1″ lead spacing
16	D1–D8, D13–D20	1N4148	or 1N914; switching diode
2	D9, D10	1N5818	or SB130; Schottky rectifier
2	D11, D12	MCL056PURGW	bi-colour LED, Multicomp
2	H5, H6		nut for M3 machine screw
2	H7, H8	M3x10	M3 male-female standoff, 10mm body length
2	H9, H10	M3x11	M3 male-female standoff, 11mm body length
6	H11–H16	M3x6	M3 machine screw, 6mm body length
4	H17–H20		nylon washer for M3 machine screw
1	J9		female single-row socket, 12 pins at 0.1″
8	J1–J8	1502 03	switched mono 3.5mm panel jack, Lumberg
1	P2		male single-row header, 12 pins at 0.1″
1	P1		male Eurorack power header, 2×5 pins at 0.1″
4	R50, R52–R54	910Ω	
1	R3	1kΩ	vertical multiturn, Bourns 3296Y/Vishay T93YB
2	R47, R48	1kΩ	
2	R49, R51	1.2kΩ	
2	R6, R7	1.8kΩ	
3	R9–R11	2.7kΩ	
2	R2, R4	6.8kΩ	
4	R1, R5, R8, R12	8.2kΩ	
8	R18–R21, R23–R26	47kΩ	
2	R27, R32	50kΩ	vertical multiturn, Bourns 3296Y/Vishay T93YB
2	R13, R14	75kΩ	
2	R17, R22	200kΩ	
2	R15, R16	1MΩ	
10	R37–R46	100kΩ 0.1%	
8	R28–R31, R33–R36	499kΩ 0.1%	
2	SW1, SW2	100SP3T1B1M1QEH	E-Switch 100-series SPDT on-off-on toggle
2	U3, U4	LM339	quad comparator, open collector output
1	U5	TL074B	quad JFET-input op amp, -B low offset version
2	U1, U2	TL431	2.495V reference in TO-92 package

Not listed above, but also needed: front panel, two circuit boards, three 14-pin DIP sockets, solder, etc. Fixed resistors should be 1% metal film, except those specified as 0.1%. North Coast Synthesis Ltd. kits or assembled modules sometimes contain other parts with equivalent or better specifications rather than exact manufacturers and part numbers shown here.

Building Board 2

The recommended order for building this module is to assemble Board 2, the one further from the front panel, first. That will make it easier to get all the physical positioning right for the components that bridge between the boards or pass through the panel.

Note that although I'm describing a separate step for each component value, and that's how I built my prototype 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. There is an abbreviated BOM for Board 2, and a few items from Board 1 used during the assembly of Board 2, in Table 2.

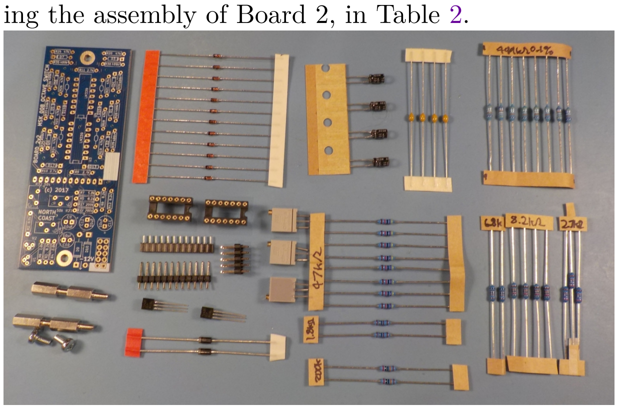

There are three multiturn trimmers to be installed on this board. Before installing them, use an ohmmeter to adjust each one to 50% of its range. Measure the resistance along the track, then measure the resistance from the wiper to one end and adjust to make the wiper half the total track resistance. This need not be exact, but having them start near their midpoints will help with adjustment later, by reducing issues with interaction among the different settings. With all trimmers pre-set to 50%, the module should basically work even if it is not at its best, whereas if they are installed at extreme values instead, then you may have trouble getting it up and running enough to adjust it more accurately.

Be aware that I sometimes discover trimmers with incorrect printed markings in incoming parts shipments, in particular trimmers sold to me as 50kΩ that on testing seem to really be 50kΩ, but the factory printing says "10K." I think many people who buy trimmers simply feed them directly into assembly robots without ever reading the printed markings, so it's easy for this kind of mistake to get into the supply chain. If you find a trimmer with unexpected markings in your MSK 008 kit, or if you have one whose value is in doubt for whatever reason, *test it* with an ohmmeter and make sure of what it actually is (bearing in mind the tolerance on track resistance, which may be up to 20%) before you solder it into a board. This centering pre-adjustment step is a good opportunity to do that. If you buy a kit from North Coast and verify that it does not contain one 1kΩ (that is, measured value between 800Ω and 1200Ω) and two 50kΩ trimmers (measured value between 40kΩ and 60kΩ, regardless of the markings), please contact us for help sorting the situation out.

Decoupling capacitors

The four axial ceramic 0.1μF decoupling capacitors, C1 to C4, are shown on the board by a special symbol without their reference designators.

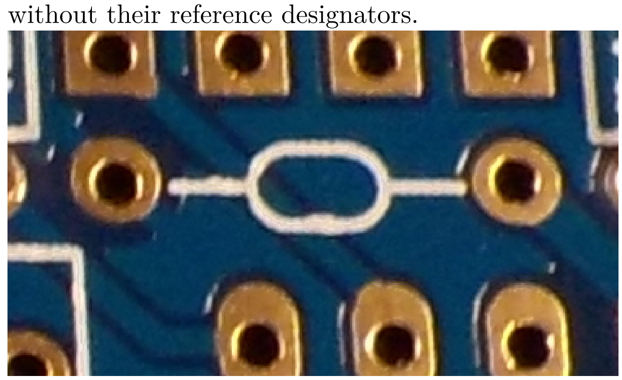

Install these four capacitors where the symbol appears. They are not polarized and may be installed in either orientation. These capacitors act as filters for the power supplies to the comparator chips, pre-

This table is not a substitute for the text instructions.

Qty	Ref	Value/Part No.	
4	C1–C4	0.1μF	axial ceramic
4	C9–C12	10μF	radial aluminum electrolytic, 0.1″ lead spacing
12	D1–D8, D13, D15, D17, D19	1N4148	or 1N914; switching diode
2	D9, D10	1N5818	or SB130; Schottky rectifier
2	H7, H8	M3x10	M3 male-female standoff, 10mm body length
2	H9, H10	M3x11	M3 male-female standoff, 11mm body length
2	H11, H12	M3x6	M3 machine screw, 6mm body length
1	J9		female single-row socket, 12 pins at 0.1″
1	P2		male single-row header, 12 pins at 0.1″
1	P1		male Eurorack power header, 2×5 pins at 0.1″
1	R3	1kΩ	vertical multiturn, Bourns 3296Y/Vishay T93YB
2	R6, R7	1.8kΩ	
3	R9–R11	2.7kΩ	
2	R2, R4	6.8kΩ	
4	R1, R5, R8, R12	8.2kΩ	
8	R18–R21, R23–R26	47kΩ	
2	R27, R32	50kΩ	vertical multiturn, Bourns 3296Y/Vishay T93YB
2	R17, R22	200kΩ	
8	R28–R31, R33–R36	499kΩ 0.1%	
2	U1, U2	TL431	2.495V reference in TO-92 package
2	U3, U4		14-pin DIP socket

Table 2: Bill of Materials for assembling Board 2.

venting any current spikes associated with their rapid switching from affecting other things on the same power supply. An MSK 008 kit should include six of these capacitors, and only four are used on this board; save the remaining two for use on Board 1.

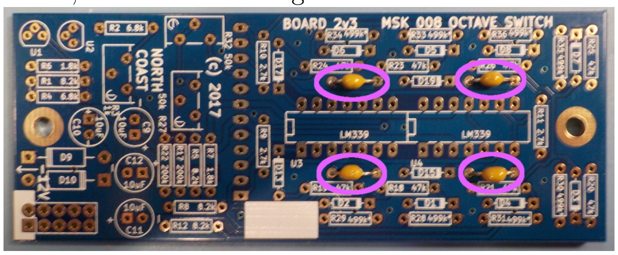

Fixed resistors

Resistors are never polarized. I like to install mine in a consistent direction for cosmetic reasons, but this is electrically unnecessary. In this module, most resistors are metal film 1% type; a few are 0.1% precision metal film resistors. Both kinds will usually have blue bodies and four colour bands designating the value, plus a fifth band for the tolerance, and these are the resistors shipped in the North Coast kits. The tolerance band is brown for 1% and violet for 0.1%, but note that we may occasionally ship better-tolerance resistors in the kits than the specifications require, if we are able to source them at a good price. Accordingly, I mention only the four value band colours for this type of resistor; if you are using resistors with other codes, you are responsible for knowing them. Note that colour codes on metal film 1% resistors are often ambiguous (reading from one end or the other end may give two different values, both plausible) and some of the colours are hard to distinguish anyway. If in doubt, always measure with an ohmmeter before soldering the resistor in place.

There are no cases in this module of the same nominal resistance value being used at more than one tolerance, but the 0.1% resistance values are marked with asterisks (*) on the board silkscreen as a reminder that these positions require special resistors.

The physical size of the resistors may vary, and details like the exact colour of the bluish background. You can see some of that variation in the photos in these instructions. Some of the resistance values used in this module are hard to find, and we source different values from different suppliers, so not all the resistors in a kit will necessarily be from the same manufacturer, nor match on non-critical specifications like power rating and physical size.

Install the two 1.8kΩ (brown-grey-black-brown) resistors R6 and R7. These control the overall current level through the reference voltage regulators.

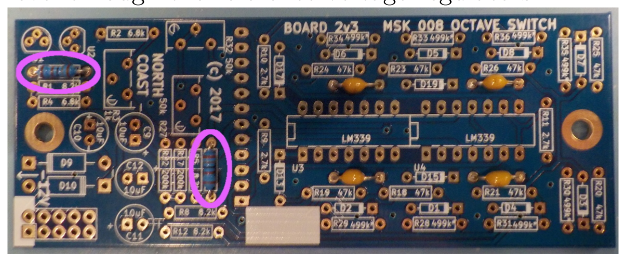

Install the three 2.7kΩ (red-violet-black-brown) resistors R9, R10, and R11. These are part of the divider that generates reference voltages for the quantizer.

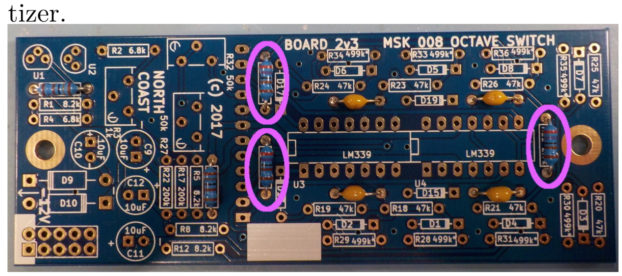

Install the two 6.8kΩ (blue-grey-black-brown) resistors R2 and R4. These are used in programming the voltage regulators.

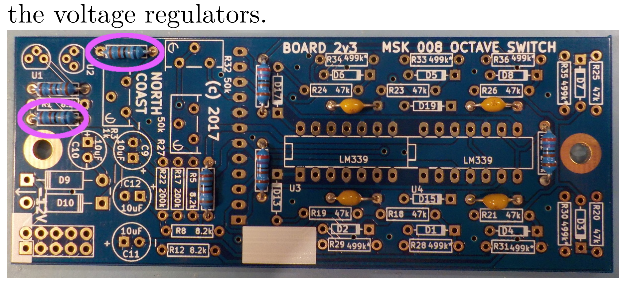

Install the four 8.2kΩ (grey-red-black-brown) resistors R1, R5, R8, and R12. These are used in multiple places in the reference voltage generator.

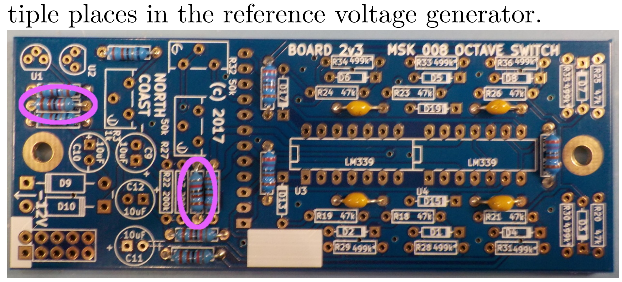

Install the eight 47kΩ (yellow-violet-black-red) resistors R18 to R21 and R23 to R26. These limit the current drawn from the -4.5V reference bus, inside the quantizers.

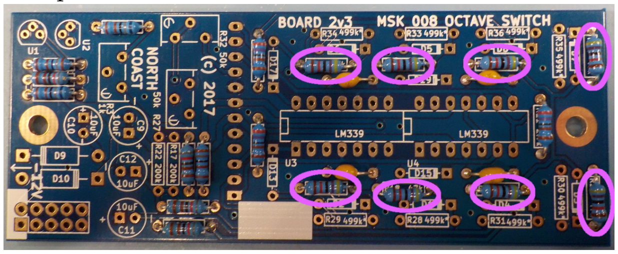

Install the two 200kΩ (red-black-black-orange) resistors R17 and R22. These (along with the 50kΩ trimmers) control the offset of the quantizer outputs.

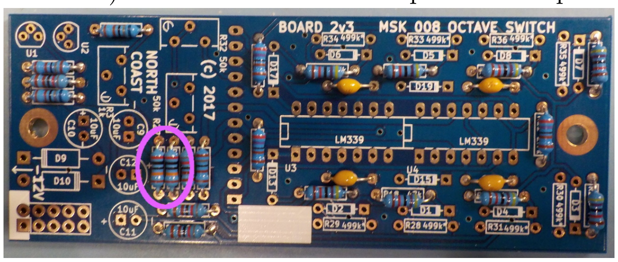

Install the eight 499kΩ 0.1% (yellow-white-white-orange) precision resistors R28 to R31 and R33 to R36. These control the voltages of individual steps in the quantizer outputs. When building the prototypes I noticed that either these resistors are just a little larger than others of the same approximate body size, or maybe their leads are a little stiffer; I found that I needed to bend the leads tighter (closer to the bodies) than usual in order to have them fit nicely on the board. Take it slowly and carefully.

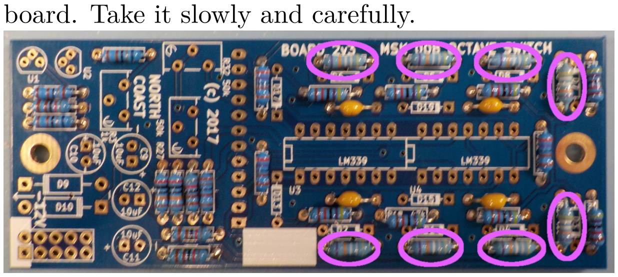

Semiconductors

Install the twelve 1N4148 or 1N914 switching diodes D1 to D8, D13, D15, D17, and D19. These act as switches to generate the proper output currents from the quantizers. They are polarized components and it is important to install them right way round. Each diode is packaged inside a pink glass bead with a black stripe at one end; that end is the *cathode*. The silkscreen markings on the board have a corresponding stripe and the diodes should be installed with their stripes matching the markings on the board. The solder pads for the cathodes are also square instead of round. Installing one or more of these diodes backwards will result in incorrect output voltages or input thresholds for the quantizers.

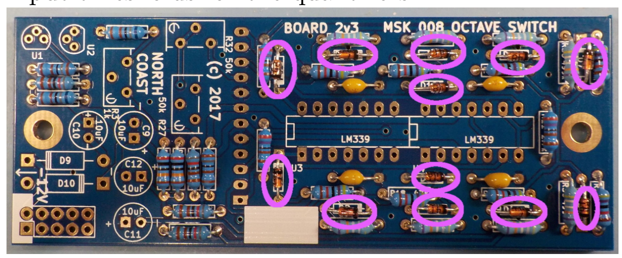

The switching diodes are the only small glass diodes in an MSK 008 kit, but do not confuse them with other kinds of small glass diodes (such as Zeners) that you might have on hand from other projects. All diodes in this kind of package look pretty much identical, distinguished only by their electrical properties and near-microscopic code numbers etched onto the glass.

Note that a kit should include sixteen of these diodes, and only twelve are used on this board; save the remaining four to install on Board 1.

Install the two 1N5818 or SBA130 Schottky rectifier diodes D9 and D10. These are for reverse-voltage protection; they cut off power to the module when the power plug is backwards. They are polarized and it is important to install them in the right direction. As with the switching diodes, these diodes will be marked with stripes indicating their cathodes (here, probably white or light grey paint on a black or dark grey plastic package) and those stripes should match the stripes on the PCB silkscreen. The cathode solder pads are also square. Installing these backwards means they will have the opposite of the intended protective effect.

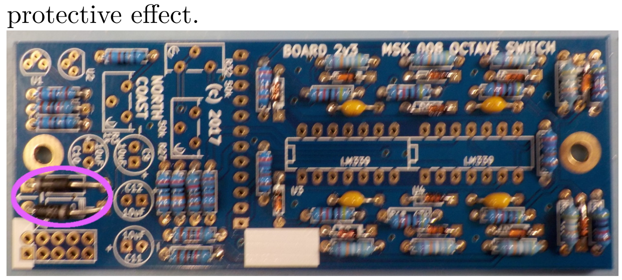

Install the two 14-pin DIP sockets for the LM339 quad comparator chips, U3 and U4. These chips test input voltages against reference values in the quantizers. The sockets themselves do not care which direction you install them, but it is critically important that the chips installed in the sockets should be installed in the right direction. To help with that, the sockets will probably be marked with notches at one end (indicating the end where Pin 1 and Pin 14 are located) and you should install the sockets so that the notched ends match the notches shown on the PCB silkscreen. The solder pad for Pin 1 is also distinguished by being rectangular instead of rounded.

Installing DIP sockets without having them tilted at a funny angle can be tricky. I recommend inserting the socket in the board, taping it in place on the component side with vinyl electrical tape, then soldering one pin on one corner and checking that the socket is snug against the board before soldering the other pins. That way, if you accidentally solder the first pin with the socket tilted, it will be easier to correct (only one pin to desolder instead of all of them).

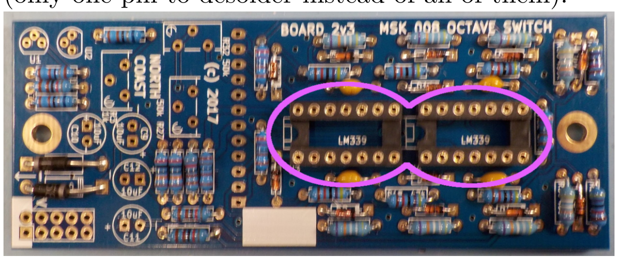

If you somehow manage to solder an entire socket in backwards, don't try to desolder it to turn it around. Just leave it as it is and remember that when you insert the chip, you must insert it so the chip matches the markings on the *board*, not the turned-around socket.

Install the two TL431 voltage regulator chips U1 and U2, which provide the ±4.5V reference voltages for the quantizer. These are packaged in epoxy plastic TO-92 pills like transistors, but they are actually complete integrated circuits each comprising about ten transistors and some other components. They are polarized and will be destroyed if installed in the wrong direction. Each chip's place on the board is shown by a silkscreened circle with one flattened side, and the chip packages are similarly rounded with one flat side; install the chips so that their flat sides match the ones shown on the board. The middle one of the three legs should be carefully bent backward to match the triangular arrangement of holes on the board, and the body of the package should sit at least a few millimetres above the board. Push it down far enough for the legs to be snug in the holes, but do not attempt to seat the package flush to the board.

Be aware that the solder holes for these chips are small and close together, and be careful to avoid making solder bridges between them.

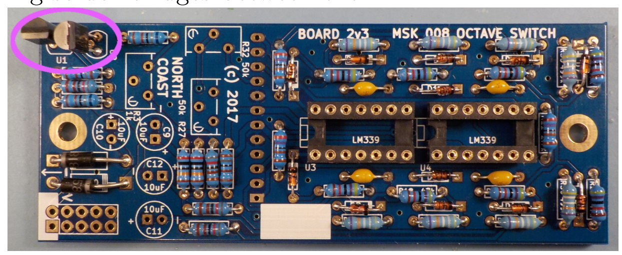

Board to board connectors

Although the sequence is not critical, I recommend installing the male header connector (P2) that links Board 2 to Board 1 at this time, before the taller components like electrolytic capacitors and trimmers, because once those are in place they may restrict access to the solder pads for the header. For best alignment, you should solder the male connector while it is mated with the female connector (J9) on Board 1, and it's convenient to solder the Board 1 connector at this time too.

Mate the two connectors firmly, then assemble the two boards and the connectors using the M3 machine screws and 10mm and 11mm standoffs, as shown. The 11mm standoffs should separate the two boards; I suggest using the 10mm standoffs instead of hex nuts for this temporary assembly because they're easier to tighten by hand. Do not confuse the two lengths. Solder the connectors on both boards. Then disassemble them, and set aside the hardware and Board 1 for later.

Note that the first batch of female header connectors I bought for this project were an expensive low-profile type that I decided not to continue using in the long term, because their main advantage would be allowing the boards to fit closer together, and other components make it hard to take full advantage of that. If you have an early kit, you may have one of these connectors, and they are the type shown in the photos. The low-profile female connectors, mated with the male headers, add up to slightly less than 11mm; so there will be a little bit of play in the temporary assembly, and you should try to adjust the connector positions before soldering so that both the male and female sides fit nicely in their corresponding holes in the PCBs. Later kits will have a taller female connector that fits the design spacing more precisely.

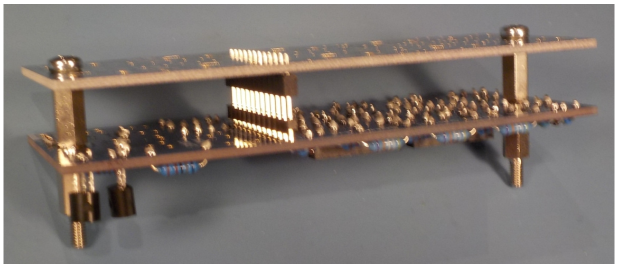

Electrolytic capacitors

Install the four 10μF electrolytic capacitors C1 through C4, which filter the power supply and reference bussesfor the module as a whole. These are polarized components and they may explode if installed backwards. Each one will be marked on its casing with a stripe and minus signs to indicate the negative lead; the positive lead will probably also be longer. These clues should be matched with the markings on the PCB: plus and minus symbols in the silkscreen and a square solder pad for the positive (long) lead.

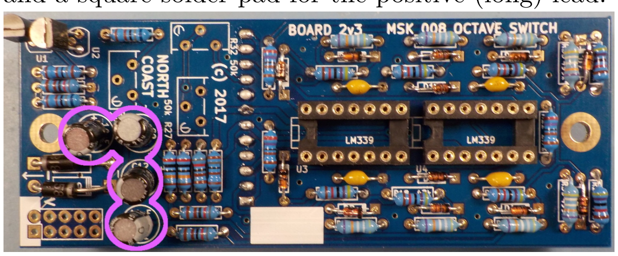

Trimmer potentiometers

If you have not already set the trimmers to 50% of their full scale value as described under "Preliminaries" above, then do it now. Also, take note of the comments about incorrectly-labelled trimmers in that section.

Trimmers usually are not washable, so if you plan to clean your boards by full immersion in water or other solvent,* your last chance is now; future cleaning will have to be done with a brush and some care to avoid letting liquid seep into the trimmers. Even now you should take some care with the DIP sockets, because solvent can carry flux residue into them and form a varnish-like layer if not carefully rinsed away.

Trimmers are not exactly polarized, but the three legs of each trimmer serve different functions and need to be connected to the right holes. The physical arrangement of the legs and corresponding holes should make it impossible to install the trimmers wrong way round.

*Did you hear the one about the hipster who was told to clean circuit boards with "IPA," and used India Pale Ale?

Install the 1kΩ trimmer R3. This trimmer's main purpose is to adjust the higher-voltage output levels of both quantizer channels. What it directly controls is the voltage on the -4.5V reference bus.

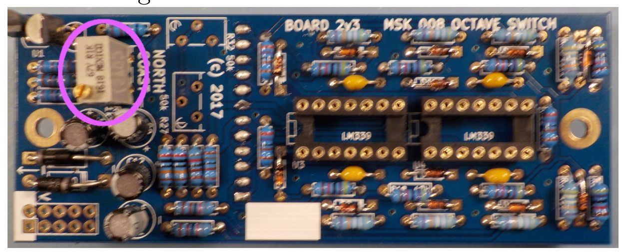

Install the two 50kΩ trimmers R27 and R32. These trimmers adjust the lower-voltage output levels of the quantizers, one trimmer for each channel. Their direct function is to control the offsetting current fed into the final summing amplifiers.

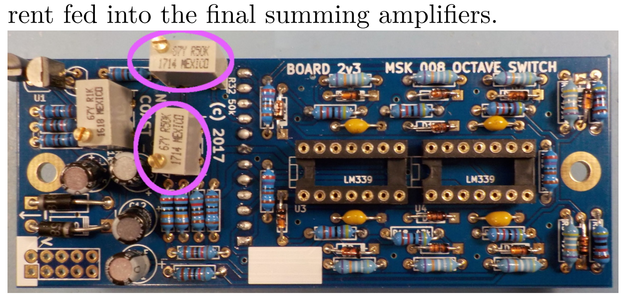

Eurorack power connector

Install the 2×5-pin Eurorack power connector. This connector is not polarized in itself, although the connection it makes is polarized. As with the DIP sockets, you should be careful to get it installed snugly against the board, not tilted at an angle. Use vinyl tape or similar to hold it in place, solder one pin, then check that it is straight before you solder the other pins.

Be aware that both the connector and the copper connections to it on the PCB have relatively large thermal mass. These solder joints will need more heat than usual; and after you have soldered it, the connector will remain hot longer than recently-soldered components usually do. Don't burn yourself.

The six pins in the centre of the connector, that is all except the four corner pins, are for grounding and they are all connected together on the board. Thus, if you accidentally form solder bridges among these six pins while installing the connector, don't waste effort trying to remove them; they will have no electrical effect.

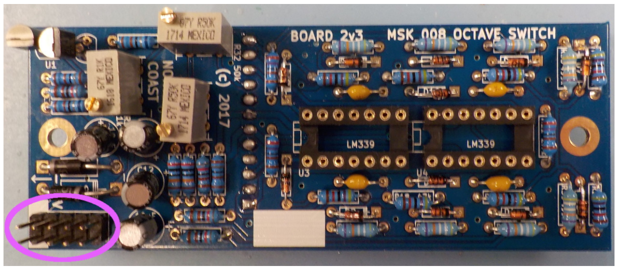

In between completed boards is a good time to take a break.

Building Board 1

Board 1 has components on both sides, and for best results, it is important to install them in the right order. Build Board 2 first, and see the general comments in the Board 2 chapters about how to approach the task.

Preliminaries

Count out the right number of everything according to the bill of materials. There is an abbreviated BOM for the items needed in this chapter (including the final assembly of the module) in Table 3. It is also assumed you have a finished Board 2 from the previous chapter.

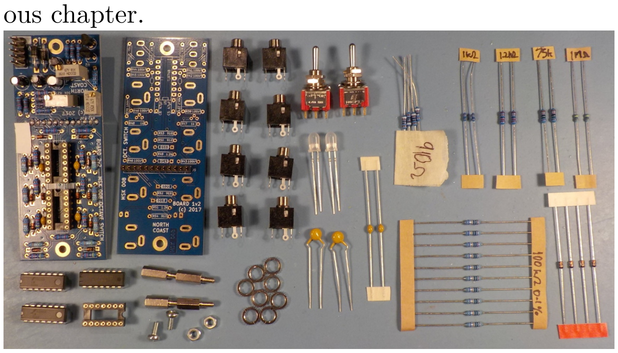

If you wish to reconfigure the CV2 inputs, you can change the solder jumpers on Board 1 at this point. In a standard build, the CV2 input on the left channel *adds* its voltage to CV1 and the quantizer output, while the CV2 input on the right channel *subtracts*. This configuration is the most useful one for most users and it is not particularly recommended to change it. But if you want to, first locate the jumpers on the board. JP1 on the left is for the left channel and JP2 on the right is for the right channel.

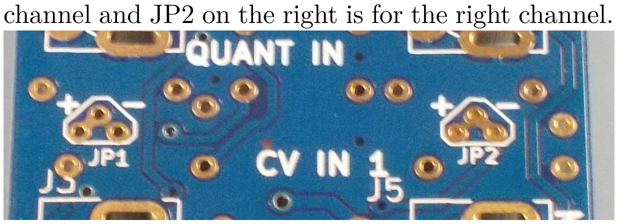

Each jumper has three terminals. The middle one is connected to the one on the left or the right for positive (adding) or negative (subtracting) respectively, with a default set by a trace built into the circuit board. To change the jumper, use a sharp knife to carefully cut the connecting trace, and then apply a blob of solder connecting the middle terminal to the one for the opposite selection. Use an ohmmeter to check that you have really broken the connection on one side and recreated it on the other.

For a more advanced modification, you can solder fine wires into the holes provided and run them into other circuitry of your own construction. If you add an SPDT switch, you could use it to change the add/subtract option by flipping the switch. You could also connect any number of additional inverting and noninverting input jacks through 100kΩ precision resistors (0.1% or better) to the + and − terminals, which as shown on the schematic diagram are op amp virtual ground summing nodes.

Decoupling capacitors

The two axial ceramic 0.1μF decoupling capacitors C7 and C8 are shown on the board by a special symbol without their reference designators.

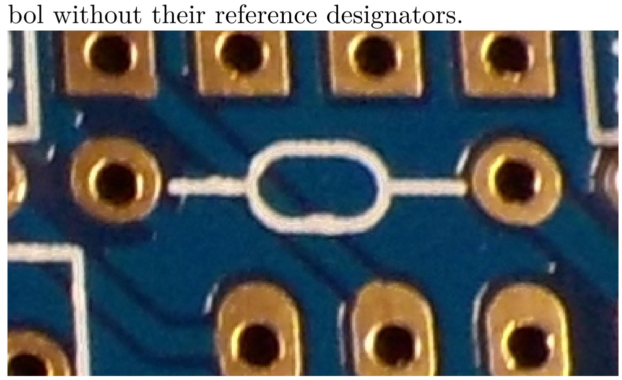

This table is not a substitute for the text instructions.

Qty	Ref	Value/Part No.	
2	C5, C6	33pF	radial ceramic, 0.2″ lead spacing
2	C7, C8	0.1μF	axial ceramic
4	D14, D16, D18, D20	1N4148	or 1N914; switching diode
2	D11, D12	MCL056PURGW	bi-colour LED, Multicomp
2	H5, H6		nut for M3 machine screw
2	H7, H8	M3x10	M3 male-female standoff, 10mm body length
2	H9, H10	M3x11	M3 male-female standoff, 11mm body length
2	H11, H12	M3x6	M3 machine screw, 6mm body length
8	J1–J8	1502 03	switched mono 3.5mm panel jack, Lumberg
4	R50, R52–R54	910Ω	
2	R47, R48	1kΩ	
2	R49, R51	1.2kΩ	
2	R13, R14	75kΩ	
2	R15, R16	1MΩ	
10	R37–R46	100kΩ 0.1%	
2	SW1, SW2	100SP3T1B1M1QEH	E-Switch 100-series SPDT on-off-on toggle
2	U3, U4	LM339	quad comparator, open collector output
1	U5	TL074B	quad JFET-input op amp, -B low offset version
1	U5		14-pin DIP socket

Table 3: Bill of Materials for Board 1.

Install these two capacitors where the symbol appears. They are not polarized and may be installed in either orientation. These capacitors act as filters for the power supplies to the op amp chip, protecting them from high-frequency crosstalk.

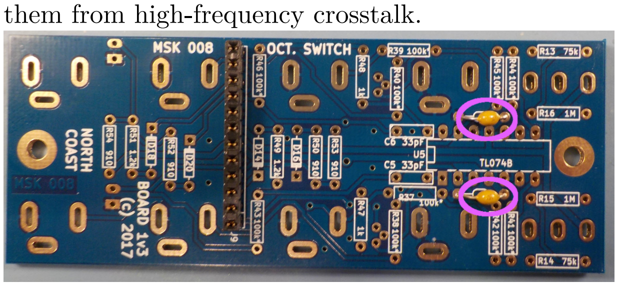

Fixed resistors

Resistors are never polarized. I like to install mine in a consistent direction for cosmetic reasons, but this is electrically unnecessary. In this module, most resistors are metal film 1% type; a few are 0.1% precision metal film resistors. Both kinds will usually have blue bodies and four colour bands designating the value, plus a fifth band for the tolerance, and these are the resistors shipped in the North Coast kits. The tolerance band is brown for 1% and violet for 0.1%, but note that we may occasionally ship better-tolerance resistors in the kits than the specifications require, if we are able to source them at a good price. Accordingly, I mention only the four value band colours for this type of resistor; if you are using resistors with other codes, you are responsible for knowing them. Note that colour codes on metal film 1% resistors are often ambiguous (reading from one end or the other end may give two different values, both plausible) and some of the colours are hard to distinguish anyway. If in doubt, always measure with an ohmmeter before soldering the resistor in place.

There are no cases in this module of the same nominal resistance value being used at more than one tolerance, but the 0.1% resistance values are marked with asterisks (*) on the board silkscreen as a reminder that these positions require special resistors.

The physical size of the resistors may vary, and details like the exact colour of the bluish background. You can see some of that variation in the photos in these instructions. Some of the resistance values used in this module are hard to find, and we source different values from different suppliers, so not all the resistors in a kit will necessarily be from the same manufacturer, nor match on non-critical specifications like power rating and physical size.

Install the four 910Ω (white-brown-black-black) resistors R50, R52, R53, and R54. These form part of the current-controlling network for the LED drivers.

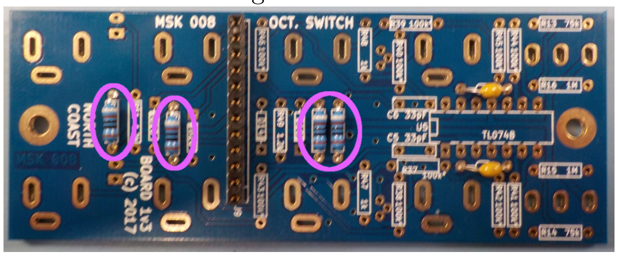

Install the two 1kΩ (brown-black-black-brown) resistors R47 and R48. These are current-limiting resistors to protect external circuits from excessive output power on the module outputs; they also serve to isolate the op amp chip from reactive (capacitive or inductive) loads that could make it unstable. Do not confuse these resistors with other power-of-ten values such as 100kΩ, which also have codes starting brown-black-black but with different colours for the fourth, exponent-indicating, colour band.

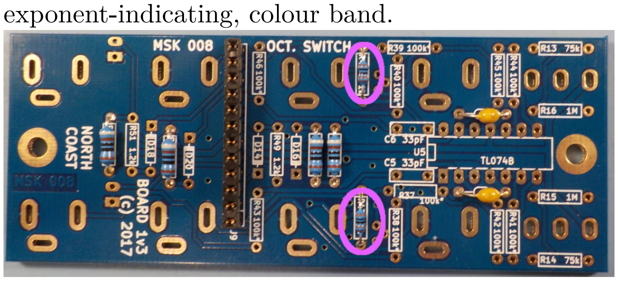

Install the two 1.2kΩ (brown-red-black-brown) resistors R49 and R51. These form part of the current-controlling network for the LED drivers.

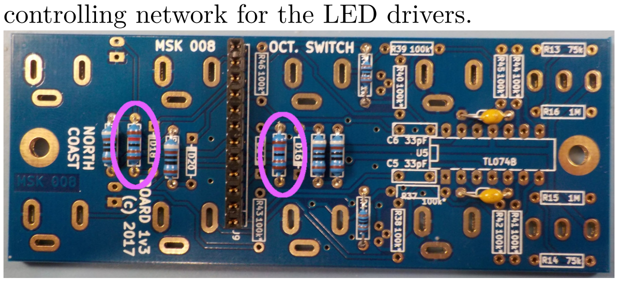

Install the two 75kΩ (violet-green-black-red) resistors R13 and R16. These set the weight of the quantizer voltage inputs in comparison to the manual octave switches.

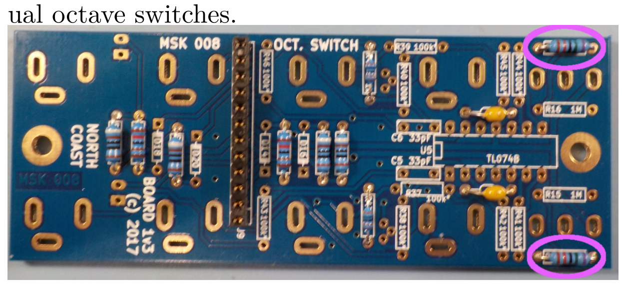

Install the two 1MΩ (brown-black-black-yellow) resistors R15 and R16. These set the weight of the manual octave switches in the quantizers. Do not confuse them with other power-of-ten resistance values.

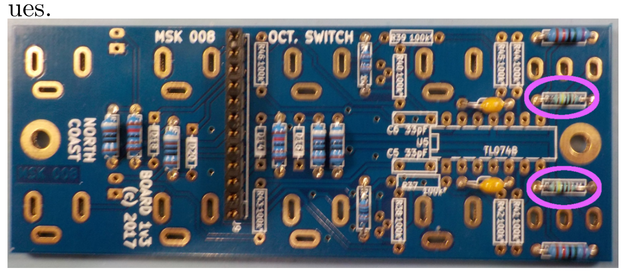

Install the ten 100kΩ 0.1% precision resistors. These control the gain of the summing and inversion op amp circuits. When building the prototypes I noticed that either these resistors are just a little larger than others of the same approximate body size, or maybe their leads are a little stiffer; I found that I needed to bend the leads tighter (closer to the bodies) than usual in order to have them fit nicely on the board. Take it slowly and carefully.

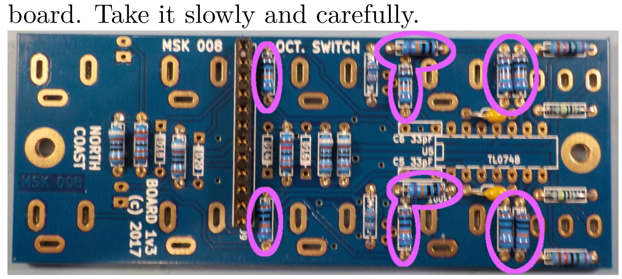

Semiconductors

Install the four 1N4148 or 1N914 switching diodes D14, D16, D18, and D20, and D19. These switch the LED drive currents. They are polarized components and it is important to install them right way round. Each diode is packaged inside a pink glass bead with a black stripe at one end; that end is the *cathode*. The silkscreen markings on the board have a corresponding stripe and the diodes should be installed with their stripes matching the markings on the board. The solder pads for the cathodes are also square instead of round. Installing one or more of these diodes backwards will result in incorrect operation of the LEDs.

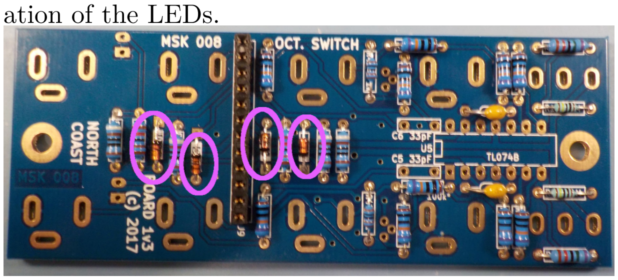

Install the 14-pin DIP socket for the TL074B quad low-offset operational amplifier (op amp) U5. This chip processes the CV1 and CV2 inputs, inverting them as necessary, and does the final summation to drive the module output. The socket itself does not care which direction you install it, but it is critically important that the chip installed in the socket should be installed in the right direction. To help with that, the socket will probably be marked with a notch at one end (indicating the end where Pin 1 and Pin 14 are located) and you should install the socket so that the notched end matches the notch shown on the PCB silkscreen. The solder pad for Pin 1 is also distinguished by being rectangular instead of rounded.

Two of the holes for the capacitors C5 and C6 are located just off one end of the DIP socket's footprint at the same spacing, making it possible to insert the DIP socket shifted over one space with two of its pins in the capacitor holes. Be careful not to do that.

Installing DIP sockets without having them tilted at a funny angle can be tricky. I recommend inserting the socket in the board, taping it in place on the component side with vinyl electrical tape, then soldering one pin on one corner and checking that the socket is snug against the board before soldering the other pins. That way, if you accidentally solder the first pin with the socket tilted, it will be easier to correct (only one pin to desolder instead of all of them).

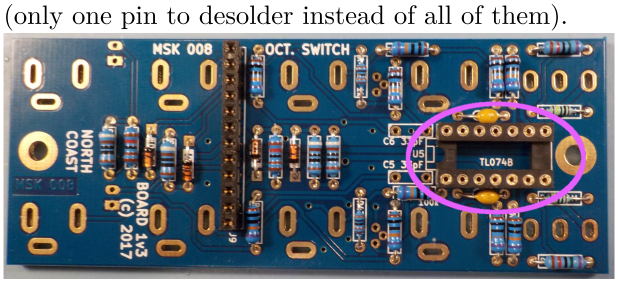

Compensation capacitors

Install the two 33pF radial ceramic capacitors C5 and C6. They will probably be marked "330," which means 33pF in a scheme similar to the resistor colour code: significant digits 3 3 to be followed by 0 zeroes. These capacitors help ensure stability of the op amp drivers by killing the frequency response at ultrasonic frequencies, making parasitic oscillation harder to sustain.

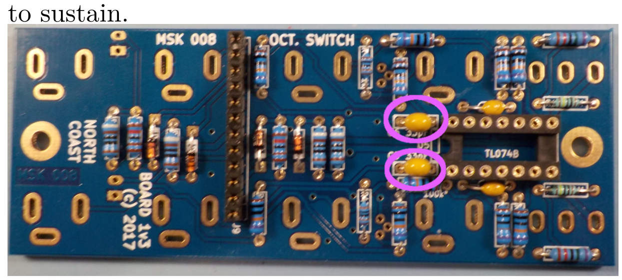

Panel components

Fasten the 10mm and 11mm standoffs to Board 1. The 11mm standoffs should be on the same side of the board as the resistors and other components, with their male ends sticking up as shown; they will separate the two boards when the module is fully assembled. The 10mm standoffs should be on the other side, where they will separate the circuit boards from the panel, with their male ends through the holes in the board to mate with the 11mm standoffs. *Do not confuse the two similar lengths of standoffs.* Arranging them wrong at this stage may result in soldering panel components at the wrong spacing in a way that will eventually make it impossible to assemble the module correctly. There is an exploded diagram for the final module on page 45. You may wish to consult it if the way the hardware fits together is at all unclear.

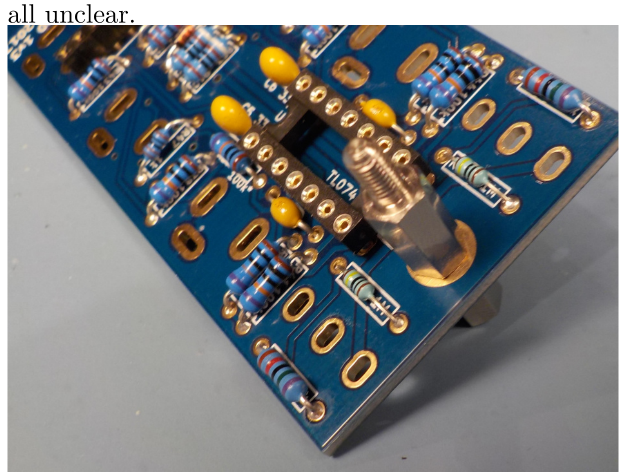

Place (do not solder yet) the LEDs D11 and D12 in their corresponding holes on the board. Single LEDs are polarized and can be destroyed by reverse voltage. These ones here are special bi-colour devices with two separate LEDs in each package. The internal connection is such that each one protects the other from reverse voltage; so if connected backwards, they will not be destroyed, but the intended green and red colours will be swapped. Each LED lens has one flat side, and one leg shorter than the other on that side. The short leg is Pin 1. Its proper place on the board is marked by a circle with a flattened side matching the direction of the flattened side on the LED lens, and an oval solder pad. The other leg (Pin 2, long, farther from the flat side) goes into the rectangular solder pad. Be sure both LEDs are placed right way around according to these clues.

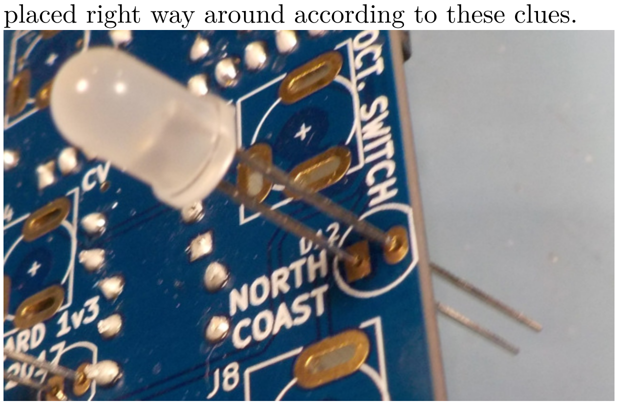

Place (do not solder yet) the eight jack sockets J1 to J8 in their corresponding holes on the board. Their arrangement of legs and corresponding slotted holes is such that they will only fit one way.

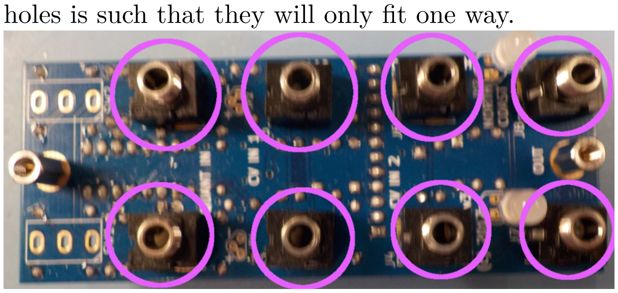

Place (do not solder yet) the two toggle switches SW1 and SW2 in their corresponding holes on the board. The electrical connections on these switches are symmetrical, but there is a keyway or groove on the threaded bushing of each switch, and the keyway must be oriented downward for the mounting hardware to fit properly later.

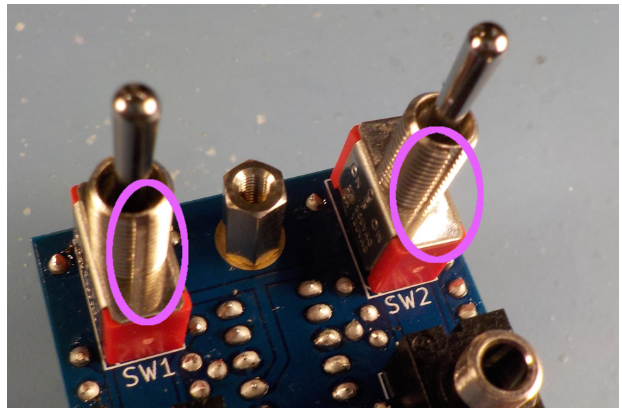

Put the panel in place over the board so that the jack socket and toggle switch bushings go through the corresponding holes in the panel and the standoff line up behind their corresponding holes. It may require some careful adjustment of the jack socket locations to get them all through the panel properly. Check that the keyways on the toggle switch bushings are facing downward, towards the small panel holes that will accept the tabs from the locking rings. Attach the panel firmly with the machine screws.

Use the knurled nuts that came with the jacks, to attach those to the panel. Beware of damaging the panel with wrenches, pliers, and similar. If you must use pliers, wrap them with tape to reduce the risk of scratches; but just screwing the nuts on with finger pressure should be sufficient.

Use the hardware that came with the switches to attach them to the panel. In order starting closest to the panel, there should be a locking ring with a tab that fits into the keyway on the switch bushing and another tab that fits into the small hole on the panel; a toothed lockwasher; and at least one hex nut.* If the locking ring will not fit because you mounted the switch backwards with the keyway on the opposite side, this is your last chance to fix it.

Turn the assembly over so that the LEDs fall into place. Adjust them by gently pulling or pushing their legs until they all pass through the panel holes by whatever amount you prefer. If they are very loose, it may be necessary to use sticky tape to temporarily hold them at the right depth, but such a measure is unlikely to be needed; the holes are designed to be a moderately snug friction fit on the LED lenses.

Solder all the jack sockets, LEDs, and switches. The jack sockets and switches may require a relatively large amount of solder to fill the correspond-

*These switches usually come with two nuts, but one is probably sufficient given the use of the toothed lockwasher.

ing holes, but these joints are structural and should not be neglected. The LEDs, on the other hand, are sensitive to soldering heat and should not be given excessive amounts of solder and heating, though the joints must at least be strong enough not to break if users should accidentally press on the LED lenses.

Final assembly

Insert the TL074B chip in its socket on Board 1. Be careful to insert it right way round: the end with Pin 1 will be marked by an indentation at one corner or a notch in the end and this end of the chip should be inserted to match the notch in the socket and on the board silkscreen and the rectangular Pin 1 solder pad. The Pin 1 end of the chip is at the bottom when the module is inserted in a rack.

Also be careful that all the legs of the chip go into the corresponding holes in the socket. These chips, when brand new, usually have their legs splayed outward a little bit (a measure intended to help them fit snugly into circuit boards when used without a socket) and you must gently bend the legs inward in order to fit them in the sockets. If you apply pressure to a chip prematurely, without all the legs properly fitting into the holes, it is easy to have the legs fold up or even break off.

It should not be necessary to remove the panel from Board 1 again. Just attach Board 2, carefully fitting its header plug into the header socket on Board 1 and the male ends of the spacers through the corresponding holes in Board 2. Then use the hex nuts to fasten Board 2 in place.

Insert the two LM339 chips in their sockets on Board 2. Be careful to insert them right way round, with the Pin 1 marking on the chip matching those on the board, pointing downward when the module is inserted in a rack. As with the TL074B, be careful all the legs are in the holes of the socket before you press the chip down, lest you fold up the delicate legs.

There is a rectangular white area at the lower left corner of Board 2 reserved for adding a serial number, signature, quality control marking, or similar. Use a fine-tipped permanent marker to write whatever you want there. Isopropyl alcohol will probably dissolve marker ink, so do this step after any board-cleaning.

Your module is complete.

Adjustment and testing

The MSK 008 is designed to produce accurate output voltages, a goal partly achieved through the use of precision components and partly by adjusting trimmers to compensate for any remaining inaccuracies. To perform the adjustment steps you will need an accurate voltage standard (usually, a good-quality multimeter) to compare against, and at least one patch cable.

For general troubleshooting and testing, you will also need a suitable power supply and a multimeter. An oscilloscope is not required, though optional testing steps using one will be described.

Short-circuit test

With no power applied to the module, check for short circuits between the three power connections on the Board 2 Eurorack power connector. The two pins at the bottom, marked with white on the circuit board, are for -12V. The two at the other end are for +12V; and the remaining six pins in the middle are all ground pins. Check between each pairing of these three voltages, in both directions (six tests in all). Ideally, you should use a multimeter's "diode test" range for this; if yours has no such range, use a low resistance-measuring setting. It should read infinite in the reverse direction (positive lead to −12V and negative lead to each of the other two, as well as positive lead to ground and negative to +12V) and greater than 1V or 1kΩ in the forward direction (reverse those three tests). If any of these six measurements is less than 1kΩ or 1V, then something is wrong with the build, most likely a blob of solder shorting between two connections, and you should troubleshoot that before applying power.

Optional: Although we test all cables before we sell them, bad cables have been known to exist, so it might be worth plugging the Eurorack power cable into the module and repeating these continuity tests across the cable's corresponding contacts (using bits of narrow-guage wire to get into the contacts on the cable if necessary) to make sure there are no shorts in the cable crimping. Doing this *with the cable connected to the module* makes it easier to avoid mistakes, because the module itself will short together all wires that carry equal potential, making it easier to be sure of testing the relevant adjacent-wire pairs in the cable.

Plug the module into a Eurorack power supply and make sure neither it nor the power supply emits smoke, overheats, makes any unusual noises, or smells bad. If any of those things happen, turn off the power immediately, and troubleshoot the problem before proceeding.

Optional: Plug the module into a Eurorack power supply *backwards*, see that nothing bad happens, and congratulate yourself on having assembled the reverse-connection protective circuit properly. Reconnect it right way round before proceeding to the next step.

Blinking lights

Plug the module into a Eurorack power supply, look at the lights on the front, and flip the switches. With a channel's switch down (−1 position) the light for that channel should be green; with it up (+1 position), the light should be red; and with it in the middle (0), the light should be unilluminated.

Output voltage adjustment

Identify the three trimmers on the back of the module (Board 2). The basic concept here is that each channel has five output voltages from −2V to +2V. There is a trimmer (R27 for the left channel and R32 for the right) which controls the level of the −2V output voltage separately per channel. Then there is also a third trimmer (R3) which adjusts the spacing between the steps, for both channels; as a result its effect is greatest on the +2V levels. Nonetheless we will do the finest adjustments on these trimmers by examining the −1V and +1V levels, because those are easier to generate and test (given that ±2V is full scale on many multimeters) and choosing points nearer the middle helps distribute residual errors.

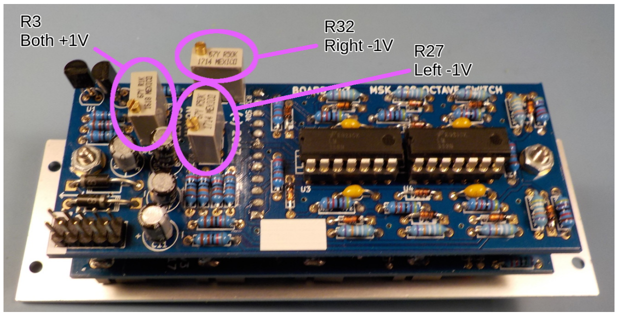

For this adjustment procedure you will need to apply power to the module, and check the output voltages on the two channels under different switch settings. If you have a spare 3.5mm jack socket it may be convenient to make a little adapter for testing the voltage on an output jack using the socket, a patch cable, and clip-on multimeter probes. Otherwise, you can just press the probes against the bare end of a patch cable, or use alligator clips or similar. Attempting to probe into the back of Board 1 with pointed probes will probably work, but may be annoying and fiddly.

Please do not pass the output voltages through any other piece of electronic equipment before measuring them, especially not the so-called "buffered outputs" of a Mordax Data. Such devices often introduce inaccuracies in the voltages they pass through, so if you adjust the MSK 008's outputs to produce correct results on the far side of the buffer, then you will be introducing an error in your MSK 008 equal and opposite to the error introduced by your buffering device, resulting in problems that are difficult to debug.[*] Measure the voltages, using a real voltmeter and not a Data or similar, directly at the outputs of the octave switch without other modules in between.

Step 1: First you will adjust R27 on the −2V output of the left channel. Patch the output of the right channel into the quantizer input of the left channel, and switch both toggle switches to the −1 position (both LEDs should glow green). Measure the output voltage of the left channel and adjust R27 to bring it close to −2V. Perfection is not necessary here. You will be adjusting it again later; but getting it nearly right at this stage will make the later adjustments, which interact, much easier.

[*]I certainly found it difficult to debug when a customer attempted this and then contacted me for technical support.

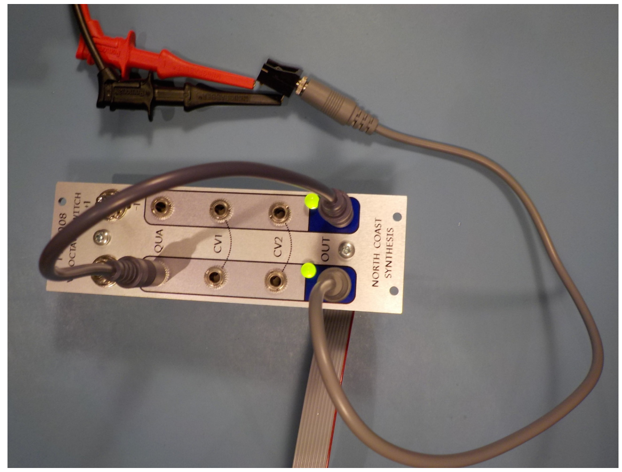

Step 2: Switch both toggle switches to the +1 position (both LEDs should glow red). Measure the output voltage of the left channel and adjust R3 to bring it close to +2V.

Step 3: Reverse the patch: output of left channel into QUA input of right channel, voltmeter to measure output of right channel. Switch both toggle switches to −1, and adjust R32 to bring it close to −2V. *Then remove the patch cable between the two channels.*

Step 4: Measure the output of the left channel. Switch its toggle switch to −1, and adjust R27 to bring its output as close to −1.000V as possible.

Step 5: Switch the left channel toggle switch to +1, and adjust R3 to bring its output voltage as close to +1.000V as possible.

Step 6: Connect the voltmeter to the right channel. Set the right channel's toggle switch to −1, and adjust R32 to bring its output voltage as close to −1.000V as possible.

Step 7: Repeat steps 4, 5, and 6 a second time.

Completing these steps as described should be enough to get all the quantizer output voltages as accurate as the components and your test equipment allow. Without guarantee, this will normally be better than ±1mV on all the output voltages. If you have a lot of patience, you can attempt to split any remaining error on R3 between the two channels.

If it seems necessary to turn a trimmer all the way to the end of its range, and especially if even doing so does not allow you to hit the desired output voltages, then something is wrong; see the troubleshooting suggestions next section.

Troubleshooting

It would require several books to convey all the skills and knowledge useful in troubleshooting even a simple electronic circuit like this one, but here are some possible symptoms and some suggestions on diagnosis and treatment.

No response from the module at all; *none* of the lights light up, no signal on the output. Most likely a power problem, such as a power cable plugged in wrong or a short circuit. It might even be a problem in the power supply and not the module itself.

Module responds, but not as expected: first attempt to localize the problem. Do signals on the CV1 and CV2 inputs appear correctly at the output (though possibly with DC shifts, and with inversion on inverting CV2 signals)? If not, check the summing circuits (U5 and associated components) on Board 1 for problems. Do voltages on the QUA inputs result in the expected quantized voltages on the outputs? If not, look at the quantizers on Board 2, involving U3 and U4. Note U3 detects the positive switching thresholds for both channels and U4 detects the negative thresholds for both channels; it is not one channel per chip.

If the quantizer appears to operate at all, but input thresholds or output voltages seem to be wrong: check the voltages along the R8–R12 string for the nominal values shown on the schematic, though these may correctly *not* have exactly the values shown if the adjustment procedure has modified them to compensate for inaccuracies elsewhere. They should at least be close. Wildly incorrect voltages along this string may indicate problems with the voltage generator section (see circuit explanation), including U1, U2, and the components near them on Board 2. Check for solder bridges on U1 and U2, with an ohmmeter as well as visually.

If the voltages along the resistor string are correct but the quantizer output or LED responses are wrong, especially if the chips or LED-associated components get noticeably hot: check for reversed switching diodes.

General tips for debugging DIP ICs: make sure for, for each IC, that

- it really is the type of IC it's supposed to be, not something else (beware of cheap ICs you buy from Chinese sellers on eBay, though the ones in this project are common enough in more reputable channels that you probably wouldn't have attempted that anyway);
- it is plugged in snugly;
- all the legs of the chip go nicely into the corresponding holes in the socket, with none bent outside or folded up under the chip;
- it is plugged in *at all* (forgetting to do so is a surprisingly common mistake!);
- it is plugged in the right way around, with the Pin 1 indentation or notch at the top and matching the other clues on the board (if this is wrong, the chip is probably destroyed and will need to be replaced);
- there are no solder bridges on the chip socket, unsoldered pins, debris clogging the socket holes, or similar;
- its decoupling capacitors (the small ceramic ones) are installed and there is nothing wrong with their solder joints; and
- in the case of the two LM339 chips, try swapping them and see if that makes any difference.

Patch Suggestions

The MSK 008 provides simple building blocks that can be combined and applied in a number of ways, some of which are hard to describe or very much specific to individual larger patches. It's what is sometimes called a "patch-programmable" module. The suggestions in this chapter are intended to inspire readers to devise their own and *not* to be a complete list of every possible way to use the module.

Basic octave switch and transposition

Basic use as a manually-controlled octave switch: pitch CV into the CV1 input, output into oscillator, switch the toggle switch up or down to go up or down an octave. Either channel can be patched like this and the other remains available for some other use; with no other inputs patched, the second channel can (through normalization) drive a second oscillator with independent manual octave switching.

In this patch it's also possible to do a CV-controlled transposition by feeding the transposition voltage into the CV2 input. On the left, it transposes in the same direction as the input; on the right, it is inverted; and the same transposition voltage can be used for both at once through the normalization.

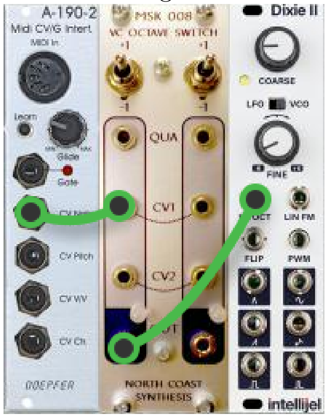

CV-controlled octave switch

The QUA input (here driven by an attenuated sine wave LFO) has a similar effect to switching the manual switch, but can go from −2 to +2 octaves.

Switching by other musical intervals

An oscillator or other module with an "exponential FM" input can switch up or down by some other interval, such as a perfect fifth, by adjusting the attenuation on the exponential FM. Here the MSK 008 is functioning as a manually controlled ±1V offset generator; as before, patching an LFO into the QUA input would allow the same thing under CV control.

Wavefolding

Basic wavefolder patch. The input signal is mult-ed to both the QUA input and CV2 on the inverting channel. To understand its operation, imagine the rising slope of a sawtooth wave on the input. The quantizer changes that into a stair-step, and then the subtraction changes each step of the stair into a little falling sawtooth slope. The picture shows the input going through a VCA, which could be used as a manual attenuator or under CV control to modulate the amount of folding effect; a little slow modulation makes the sound much more animated and interesting. Output voltage for this patch is low (1V peak to peak, unless the input is hot enough to exceed the quantization range) and it may need amplification in some applications. For a stronger output (up to 4V peak to peak), try feeding the input only into the QUA input (not also CV2), which gives a slightly less aggressive bit-crushing sound and works in either channel.

Scale degree detector

This is an extended application of the wavefolder patch, originating in a discussion on a Web forum about how to detect a given scale degree for switching sequencer patterns in a self-generating melody. It's to be understood that the details would vary a lot depending on the exact application, but this example illustrates some useful tricks for getting the most out of the module.

The sequencer generates a pitch control voltage, here assumed to be from 0V to 5V (an actual ER-101 is capable of a wider range, but I'm only using it as a recognizable example of a sequencer, not meaning to imply that the input would necessarily really be an ER-101 in practice). That voltage goes into the CV2 input of *both* channels through the normalization. So in the right channel, we are subtracting the pitch voltage from an offset generated by the Triatt, which will be nominally +2.5V. Ignore the QUA input of the right channel for the moment. The output on the right ranges from +2.5V down to −2.5V as the pitch CV *increases*, and that goes into the QUA input of the left channel.

The left channel's quantizer recognizes which octave we are in, with a quantizer output of +2V for the first octave (pitch CV from 0V to 1V, right channel output from +2.5V down to +1.5V), +1V for the second octave, and so on. To the quantizer output we add the original pitch CV. So as long as the pitch CV remains in the range 0V to 5V, the output of the left channel is in the range 2V to 3V, expressing what can be called the "pitch class," pitch modulo octave, or scale degree within the octave. For example, if C in one octave is 0V on the input, then C in any octave will be 2V on the output, and F♯ in any octave will be 2.5V on the output. This output then goes into other modules, here depicted by an A-167 comparator, to detect when the sequencer hits a given scale degree. That might be used to trigger a change in the sequencer pattern or something.

In order to have the output range be 2V to 3V and not the less convenient 4.5V to 5.5V, we need to defeat the normalization on the left channel CV1 input. It would work to just plug in a patch cable there with nothing on the other end, but people don't like having loose cable ends floating around, so the diagram shows this input patched into the right channel's QUA input, which (with the offset switch in the centre 0 position) has a very high impedance and is effectively the same as leaving the cable unpatched, but looks tidier. The CV1 input itself is 100kΩ into a virtual ground and will not have any unexpected effects on the right channel. For other output ranges instead of 2V to 3V, one could apply an appropriate positive or negative offset to the left channel CV1 input instead.

The exact cut-off point for the octaves can be adjusted by adjusting the offset into the right-channel CV1. In the simplest case one could just mult the pitch CV into QUA and CV2 of the right channel in the basic wavefolder patch, leaving the left channel free for other purposes; but that would give a ±2.5V input range, ±0.5V output range, inversion of the voltages, and no control over the cut-off point; so it might not be useful in as many patches.

Mid-side and fake stereo

With signals fed into the CV1 and CV2 inputs, normalled across to both channels, the two outputs of the MSK 008 are the sum and difference of the two signals. As a result it functions as a "mid-side" encoder or decoder, with 6dB of gain across the pair if you use two of them with nothing in between. In this example we've got a mono signal going into a Doepfer A-188-1 bucket brigade delay. The dry signal (through the mult built into the Doepfer module) goes into CV1 on the MSK 008 to be the "mid" signal applied to both stereo channels. The BBD's mix knob is set to fully wet output, and that goes into CV2 to be the "side" signal for mid-side decoding, applied with opposite phase to the right and left.

The result is a fast echo and phase cancellation on the two stereo channels; a "fake stereo" effect that spreads mono material across the sound stage as if it had been recorded in stereo. It's not very realistic, but a similar effect was historically important. When stereo recording was new, Capital Records sold a lot of vinyl that they called "Duophonic," with essentially this effect applied to mono masters in order to appeal to people with new stereo hi-fi equipment.

Other interesting effects can be had by using filters, distortion, and other things instead of or in addition to the BBD.

Basic Schmitt trigger patch

With the output fed back into the QUA input, either channel can function as a Schmitt trigger: that is, a comparator with hysteresis. When the sum of the other inputs (taking into account any inversion) is greater than +2.5, the module switches into the positive state (quantizer output +2V). When the sum is less than −2.5V, the module switches into the negative state (quantizer output −2.5V). In between, it retains its current state. This behaviour can be used for a number of switching and distortion effects, detecting peaks, and so on. However, in the simplest single-channel version, it is not particularly well-behaved; in particular, the input voltages appear superimposed on the output, and the basic states are ±2V, which may not be what other logic modules expect.

Gate stretcher/delay

Here's an application of the Schmitt trigger patch. An ADSR envelope generates a voltage input for the Schmitt trigger. The toggle switch for that channel is set to −1 to ensure that when the ADSR is at 0V, the module output will be biased into the negative state. When the ADSR receives a gate signal, its output voltage starts to increase; eventually it rises enough to trip the MSK 008 into the positive state, and the delay before that happens is controlled by adjusting the attack time. Similarly, when the gate disappears, the ADSR generator goes into its release phase, and there is a delay controlled by release time before it will switch the MSK 008 into the negative state.

The Doepfer A-140 is shown as an example of a typical envelope generator, but one could also use any common "AD"-style envelope generator, and if it is desired to use this patch for stretching a trigger signal (which is nothing but a very short gate) into something longer, with a delay longer than the trigger duration, then it might be important to choose an envelope generator that will go through its full attack cycle despite having a very short trigger input.

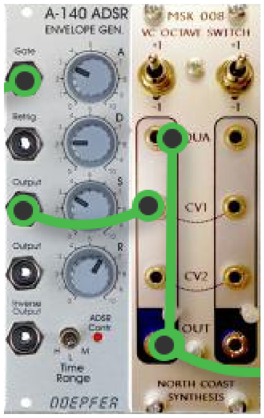

The deluxe version of the patch uses the second channel of the MSK 008, and an offset input (typically 2V), to clean up the output voltage a bit. Minor stair-stepping is still possible.

Circuit explanation

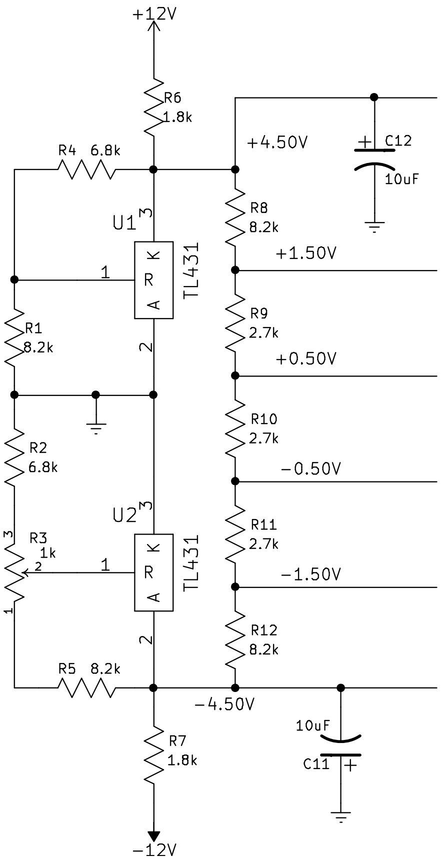

Figure 3: The voltage generator section.

Reference voltage generator

At the heart of the MSK 008's quantizer circuit, the section shown in Figure 3 generates six reference voltages: ±4.50V, ±1.50V, and ±0.50V. This section is shared between the two channels in the module.

The voltages shown in Figure 3 are nominal values only. The actual voltages of these references will vary because of component tolerances, and even the −4.50V reference, which can be adjusted more or less directly by R3, may actually end up adjusted to a voltage not exactly −4.50V in order to help compensate for inaccuracies elsewhere in the module. The function of these references is primarily to be stable, and close enough to the nominal values to allow the rest of the circuit to work correctly, rather than for them to be specific accurate voltages. For this reason, it's not necessary to use a close-tolerance version of the TL431 regulator; the default standard version (which might be 2% tolerance) is fine.

The first thing to understand in analysing this circuit is the function of the TL431. This IC has a fairly complicated circuit inside it, but to the outside world it appears relatively simple. It's a shunt voltage regulator, used very much like a Zener diode. In fact, some versions of the circuit symbol depict it as a Zener diode with an extra pin, and the letters "K" and "A" in the diagram, for "cathode" and "anode," reflect the Zener-diode analogy. The critical difference is that it also has that third pin, labelled "R" for "reference."

When operating normally in a circuit that allows it to do so, the TL431 will pass whatever current from cathode to anode is necessary to keep the reference pin 2.495V above the anode. If you connected the reference pin and cathode directly together, it would be just like a 2.495V Zener diode with unusually good temperature stability and turn-on characteristics. But here we are driving the reference pin from the cathode through a voltage divider. In the case of U1, the divider is R1 and R4, and the voltage on the reference pin is $8.2\text{k}\Omega/(6.8\text{k}\Omega + 8.2\text{k}\Omega) = 0.5467$ times the voltage on the cathode. Then in order for the reference voltage to be 2.495V, the voltage on the cathode needs to be 4.564V. The TL431, if it can, will draw enough current through its cathode to make that true. By adding the voltage divider, we have in effect programmed U1 to behave like a 4.564V Zener diode. Tolerance variation in any of the three components may push this voltage off by a few percent in either direction.

In the case of U2, there is the 1kΩ potentiometer R3 inserted into the middle of the voltage divider to make the ratio adjustable, but it's otherwise the same circuit, and we'd normally expect that pot to be adjusted to roughly the same ratio as in the U1 circuit; the nominal −4.50V reference will be close to the negative of the nominal 4.50V reference.

Like a Zener diode, the TL431 only works well over a specified range of shunt currents. That's partly because it needs a source of energy for its complicated internal circuit (error amplifier, precision reference, and so on). Since it functions as a sort of amplifier with feedback, there may also be a concern about stability: when the load is capacitive and within a certain range, a TL431 may go into parasitic oscillation. So we need to ensure that the chips stay within the operating range where they can work normally, and that is accomplished by R6, R7, C11, and C12.

If the cathode of U1 is at about 4.50V, and the positive power supply is about 12V, then we have about 7.5V across R6, resulting in a current of 4.167mA. Out of that, we must subtract about 367μA for the resistor chain R8…R12, another 300μA for the voltage divider driving the TL431 reference input, and whatever current is drawn by the other parts of the circuit that use this reference. That is up to 45μA for the 4.50V reference (because it drives two adjustable resistances each at least 200kΩ into virtual ground). For the −4.50V reference the other parts of the circuit actually draw current (worst-case about 1mA of it) into negative supply instead of into ground, with the effect of increasing the current through U2 and making the regulator's job even easier. The net current through each regulator stays between about 3.5mA and 5.0mA, which is comfortably within their advertised operating range of 1mA to 100mA.

As for stability, the 10μF capacitors, in parallel with whatever other capacitances may exist in the circuit, ensure that each regulator sees an effective capacitance of at least 10μF, placing it safely out of the danger zone which for these circuit parameters would be from around 0.01μF to 2μF. This technique could still fail in case of a really badly behaved load that acts like an inductor and cancels out part of the capacitance, but there is nothing like that in the actual circuit that uses these reference voltages. In fact, the stray capacitance that does exist is probably small enough that we would probably be safe on the other side of the danger zone without C11 and C12; they are included out of an abundance of caution and to help filter out any high-frequency garbage that might get into the reference busses from other sources.

Given the ±4.50V references, the remaining reference voltages are generated by a straightforward resistor chain, R8…R12. For perfectly correct ratios the 8.2kΩ resistors ought to be exactly three times the 2.7kΩ resistors (thus 8.1kΩ), and that is not quite true, but it's close enough. If you do a more precise calculation, you may note that the error from using 8.2kΩ resistors here is in the opposite direction from, and will to some extent cancel out, the earlier error of the ±4.50V references being really more like ±4.564V. However, that fortunate cancellation is something of an illusion, because the component tolerances are not precise enough to really say anything more than "close enough." The ±4.50V references can each source or sink more than 1mA without challenging the regulators; the ±1.50V and ±0.50V references cannot handle so much without being pulled off voltage, but they drive only the very-high-impedance inputs of LM339 chips with bias currents measured in nA.

Manual octave switch

The QUA input jacks go into resistor voltage dividers like this one:

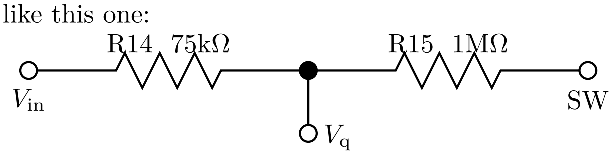

This is a passive mixer intended to add or subtract about 1V from the input under control of the three-position toggle switch; it's not very accurate but produces results that are close enough given we only really care about the positions of the four quantizer rounding boundaries, which are ±0.5V and ±1.5V. The input jack connects to V_{in} and the toggle switch either is an open circuit (centre position) or connects the node labelled SW to one of the power rails, which (taking into account a nominal drop of 0.2V for the Schottky protection diodes) are ±11.8V. The quantizer will switch states when V_{q} crosses the ±0.5V and ±1.5V boundaries. From these known voltages it's straightforward to apply Ohm's and Kirchhoff's Laws to find the values for V_{in} to bring V_{q} to each of these boundaries in each state of the switch.

switch state			bin
−1	0	+1	
			}+2V
+2.50V	+1.50V	+0.73V	
			}+1V
+1.42V	+0.50V	−0.35V	
			} 0V
+0.35V	−0.50V	−1.42V	
			}−1V
−0.73V	−1.50V	−2.50V	
			}−2V

Note that the boundaries shift in the opposite direction from the switch: to have the effect of adding 1V to the input, we *subtract* approximately 1V from the quantizer boundaries, so the same input will fall into a bin 1V higher.

As an intuition for where the design came from, consider that if we just connected V_{in} to V_{q} with a 75KΩ resistor, we could shift the voltage at V_{q} up or down perfectly by 1V by adding or subtracting a $1\text{V}/75\text{k}\Omega = 13.33\mu\text{A}$ current at the V_{q} node. A true constant-current source would be more complicated to build, but a high voltage applied through a high resistance is a way of faking it. Then the ±11.8V power supply rails, applied through 1MΩ, look approximately like a ±11.8μA current source (exact only when V_{q} is at 0V), which is close enough. The resistor values were chosen by choosing R15 to be 1MΩ as a convenient high resistance, and then running the calculation backwards and choosing a nearby standard value to get 75kΩ for R14.

Quantizer building blocks

The first thing to understand in the quantizer circuit is the basic function of the LM339 comparator. Its symbol looks like an op amp symbol, and much of its internal circuitry resembles that of an op amp.

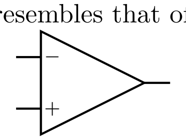

Comparators are optimized for different parameters compared to op amps, but the most important difference between the LM339 in particular and a standard op amp is in its output. Instead of giving its output as a voltage or as a current, the LM339 presents an *output impedance* that depends on the difference between its inputs. If the positive input is at a higher voltage, the output is high-impedance, looking like an open circuit. If the negative input is at a higher voltage, then the output is low-impedance, looking like a short circuit into the negative power supply, both conditions only being true in the ideal and approximated up to the limits of the chip's capabilities.

This kind of output is implemented using an NPN transistor with its collector pointing at the output pin and no pull-up circuitry, so it's called an *open collector* output. (*Open drain* for a similar kind of output implemented with a MOSFET.) They're commonly used for devices like comparators which might be applied at the interfaces between systems with different power supplies, because one can attach an external pull-up circuit referenced to a different power supply from the comparator's own power and get conversion between the different voltage standards at very little cost.

Here, we're going to use an external pull-up circuit to massage the output into a format that will be useful by the output summing amplifier and LED driver circuits. A section of the following form occurs in four places in the MSK 008.

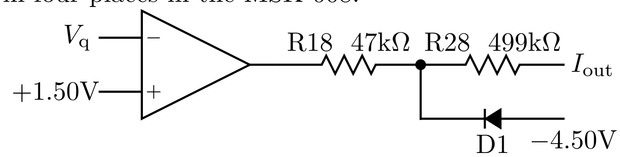

The node labelled I_{out} is the virtual-ground summing node of the output amplifier, which will be discussed later. It is held at a constant potential of 0V while other sections apply different currents to it. The +1.50V and −4.50V voltages come from the reference-voltage generator. And V_{q} is the quantizer input voltage that comes from the QUA jack with a possible adjustment for the manual octave switch.

When V_{q} is less than the +1.50V reference voltage, the output of the comparator looks like an open circuit. Then R18 is out of the circuit, and the diode is reverse-biased, so it also looks like an open circuit, and R28 is also taken out. No current, or only a very small leakage current, flows.

But when V_{q} exceeds +1.50V, the comparator turns on. Its output voltage drops to about −11.6V; that is the −11.8V power supply voltage plus about 0.2V consumed by the saturated output transistor. Remembering that the I_{out} node is held at 0V, the voltage divider of R18 and R28 will try to bring the diode's cathode to a voltage of about −10.6V. But the diode is now forward biased, so it will conduct as much current from the −4.50V bus as necessary to prevent its forward voltage from growing too large. Suppose the diode's forward voltage is 0.49V. Then the diode's cathode will be at −4.99V; and then the current sunk through the circuit output is exactly 10μA by Ohm's Law on R28.

So the overall action of this section is that it sinks a tightly controlled amount of current through I_{out}: either 10μA when V_{q} exceeds +1.50V, or zero when V_{q} does not exceed +1.50V. The forward voltage of D1 under these conditions is probably not really 0.49V, but we can correct for that, assuming all the diodes in the different instances of this section are reasonably similar to each other, by fudging the value of the −4.50V reference. Recall that the circuit which supplies that bus includes the trimmer R3 for this purpose.

Another very similar section also occurs in four places in the MSK 008; it differs by the addition of two more diodes.

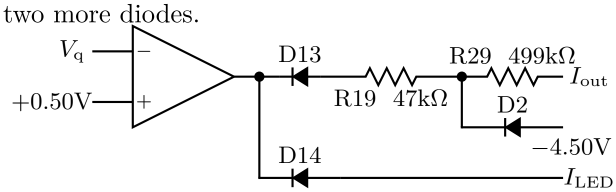

When V_q is less than the reference voltage (here +0.50V), the output of the comparator is high-impedance. Then D13 and D14, pointing nose to nose, isolate I_{LED} from the rest of the circuit, and the rest functions as before. No current flows.

When V_q exceeds +0.50V, then (as before) the comparator output drops almost to the negative supply voltage. D13 conducts, the resistors lower the cathode of D2 far enough for it to become forward-biased, and as before, the forward voltage of D2 limits how negative the cathode voltage can become; it goes to nominally −4.99V and the circuit draws 10μA through I_{out}. However, we now have the additional output I_{LED}. This node is connected to a voltage near the negative supply though D14 alone. So this is a two-output version of the comparator current sink. Independently of the precise current sinking on I_{out}, the circuit will also attempt to sink an unspecified large amount of current through I_{LED}. The maximum depends on the rating of the LM339, which varies a little by manufacturer but is typically about 15mA or 18mA.

Putting it together

The skeleton of the circuitry downstream of the comparators looks something like this. Note that these amplifier symbols are regular op amps with voltage output (TL074B type), not open-collector LM339 comparators.

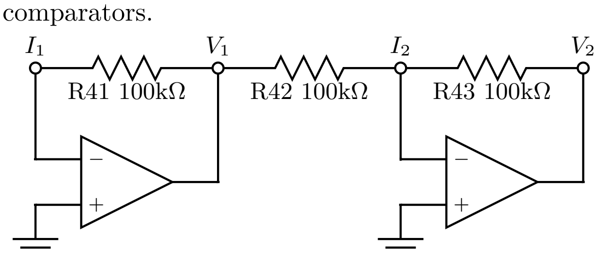

Recall the basic rules of op amps in negative feedback: the positive and negative inputs are at equal voltage, and they pass no current, all subject to the amplifier's ability to drive its output to make that true. So because the positive inputs of the two op amps are fixed at 0V, the nodes marked I_1 and I_2 are virtual grounds, also held at 0V. If no external currents are applied, then the voltages at V_1 and V_2 are also quickly seen to be zero.

Suppose we inject some current into the node labelled I_1 from some other circuitry connected to that node. The current must flow through R41 (because it can't flow into the op amp), and so we can derive the voltage at V_1 by Ohm's law. For every 10μA flowing into I_1 from the external circuit, the voltage at V_1 must decrease by 1V, and similarly it must increase by 1V for every 10μA flowing out of I_1 into the external circuit.

Now look at R42, which connects the V_1 node with the virtual ground at I_2. For every 1V V_1 goes below ground potential, R42 draws 10μA out of I_2, which must also flow through R43, forcing V_2 to a potential 1V *above* ground. The voltage on V_2 in the absence of external tampering at I_2 is always the negative of that at V_1. But we can also force current into or out of I_2 by means of external circuitry, and in that case, every 10μA injected pushes V_2 down by 1V.

The overall story here is that V_2 carries a voltage that reflects the sum of all current sources injected into I_1, minus the sum of all current sources injected into I_2, at a conversion factor of 10μA per volt. That's the kind of calculation we need in order to add up the noninverting, inverting, and quantizer inputs of the MSK 008.

The CV1 input is easy: it just connects through a 100kΩ resistor (R37, for the first channel of the module) to the node labelled I_1. Each volt on this input drives 10μA through the resistor into I_1, which translates to −1V on V_1, 10μA *removed from* I_2, and 1V on V_2, which is the module's output.

The CV2 input offers a configurable option: depending on the channel and modifications made to the circuit board, it may connect through a 100kΩ resistor (R38 for the first channel) either to I_1 or I_2. If I_1, it functions just like, and sums with, the CV1 input. If CV2 is connected to I_2, then it operates as an inverting (subtracting) input. Each +1V on CV2 in this configuration drives 10μA into I_2, forcing a change of −1V in the final output V_2.

As for the quantizer output, each channel contains four of the current-output comparator sections described in the "Quantizer building blocks" section, with threshold voltages set to −1.5V, −0.5V, +0.5V, and +1.5V. (See the full schematic at the back of this manual for the exact configuration.) All of their I_{out} nodes are connected to I_2 in the summing skeleton. So at any given time, a whole number from zero to four of these sections will be removing 10μA of current each from I_2. That contributes a whole number of volts from 0V to 4V to the final output, which is the source of the quantizer stair-steps. As the QUA input becomes more and more positive, it will turn on the comparator sections one by one as it crosses the quantization boundaries, and drive up the module output in reasonably accurate 1V steps.

But the quantizer output ought to go from −2V to +2V, not 0V to +4V, in order to match the input ranges as nearly as possible. So we need to apply an offset of −2V. That is accomplished by injecting 20μA into I_2, which comes from the +4.50V reference bus through a 225kΩ resistance composed of a 200kΩ fixed resistor (R17 for the first channel) in series with a 50kΩ trimmer set near its midpoint (R27 for the first channel). The trimmer allows adjusting the precise amount of the offset to take into account variation in the actual voltage of the reference bus, offsets resulting from imperfect op amps, and component tolerances in general.

On the full schematic you can note that the feedback circuits for the final output amplifiers are a little more complicated than shown here: in each channel there is an additional 1kΩ resistor between the op amp output and the output jack (the node called V_2 here) and a 33pF capacitor directly from the op amp output to its negative input. These components are intended to help guarantee stability and protect external circuits. The TL074B chips can source or sink a fair bit of current, possibly enough to make trouble for a very low-impedance input that might accidentally be connected; the 1kΩ resistors limit the maximum current possible. There is also a potential issue of parasitic oscillation if something capacitive (like a very long unpatched patch cable) should be connected directly to the op amp output. The resistors provide some isolation between the op amp and reactive loads in the outside world, reducing the possibility for oscillation to occur, and the 33pF capacitors further support that effort by killing the op amp gain above a few tens of kHz (ultrasonic frequencies).

LED driver

Four of the eight comparator sections, namely those associated with the ±0.5V thresholds on both channels, are of the more elaborate type with high-current I_{LED} outputs as well as the precise 10μA outputs applied to the final summing amplifier. These outputs act as strong current sinks near the negative supply when activated by the QUA input going above the corresponding threshold, and otherwise act as open circuits. They feed into a resistor network used for driving the channel's bicolour LED. The bicolour LED is really two, of different colours, built into a single package and wired back to back so that on "forward" bias according to the schematic symbol, the green one will light, and on "reverse" bias the red one will light.

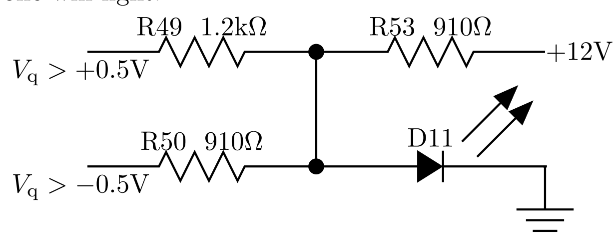

When the quantizer is in the −1V or −2V state, neither of the comparator sections is active. Current flows only from the positive supply through the LED, making it glow green. When the quantizer is in the 0V state, the −0.5V comparator section activates. Then R53 and R50 form a voltage divider between the positive and negative supplies (more or less); the centre voltage is near enough to 0V that neither side of the LED glows. Finally, when the quantizer is in the +1V or +2V state, both comparator sections are active. The voltage divider has 910Ω on top and 910Ω in parallel with 1200Ω (combined resistance about 517Ω), which makes its centre voltage low enough to turn on the red side of the LED.

The basic operation of the circuit is simple enough, but choosing the resistance values to make it work is tricky because of the competing requirements: we need the right amount of current through the LED in both colour states, a low enough voltage across the LED to keep it dark in the "off" state, current levels that the LM339 can handle, and so on. Rather than attempting to do the design entirely by hand, I used the $\mathrm{ECL}^i\mathrm{PS}^e$ constraint logic programming system (http://eclipseclp.org/) to find ranges of component values that would satisfy all the constraints.

I won't go into a detailed tutorial on how to use $\mathrm{ECL}^i\mathrm{PS}^e$ here, but will go through the constraints in neutral mathematical terms and then refer readers to the file `lednetwork.ecl` in the source code distribution, which expresses those constraints in a machine-readable form. After writing down all the constraints on paper and coding them up for the solver, I loaded that file in $\mathrm{ECL}^i\mathrm{PS}^e$ and used its interactive command line to try different values for the components until I found a combination that would work. Then I tried building it on a breadboard to verify that it would really work with the chosen components. The constraint programming approach saved a lot of time building and testing circuits in real life or even in a more conventional simulator, especially when during development I changed some of my decisions on which LEDs to use, which colours should have which polarity, and so on.

See Figure 4, which gives names to all the resistances, voltages, and currents relevant in designing the LED network. Note the variable names in this figure are *not* aligned with the reference designators in the module schematic; they are only for use in this design procedure.

The figure includes an additional resistance, R_2, in series with the LED because at one point I thought that might be useful, even though the final design ended up setting it to 0Ω. At left is the "green" state: current flows only from the positive supply through R_1, R_2, and the LED. At centre, one of the comparator sections has turned on, diverting the current into itself. The voltage labelled V_1 in this state ought to be small enough that the LED will not turn on, in which case no current (to within the precision of our model) flows through it. At right, in the "red" state, both comparators are active and current flows both from the positive supply and in the reverse direction through the LED.

I'm using a voltage of +11.8V for the positive supply, figuring it's +12V minus 0.2V for the Schottky reverse-protection diode at the module's power input. For the outputs of the comparator sections when activated, I use −10.6V: starting from −12V, adding 0.2V for the protection diode, then two diode drops of 0.6V each for the LM339's output transistor and our own switching diode. That may be an overestimate of how much voltage the LM339 output will consume, because it's supposed to be a saturated transistor (which will be less than one diode drop), but we're going to push it near its maximum current and it may not be able to keep the voltage as low as might be ideal. This number is consistent with my experimental measurements, and it's not really necessary that it be very precise anyway. The LED forward voltages of 2.15V for green and 1.75V for red are also from experimental measurements; they are very close to the MCL056PURGW data sheet's "typical" values.

With those voltages specified, we can start placing constraints on the unknown variables. First, all the resistances should be within the range of fixed resistors we can easily (and cheaply) buy; R_2 is exceptional because we might (and, it turns out, will) set it to 0Ω, whereas it seems safe to assume all the others are at least 10Ω.

$$\begin{aligned} 10\Omega \le R_1 &\le 1\mathrm{M}\Omega \\ 0\Omega \le R_2 &\le 1\mathrm{M}\Omega \\ 10\Omega \le R_3 &\le 1\mathrm{M}\Omega \\ 10\Omega \le R_4 &\le 1\mathrm{M}\Omega \end{aligned}$$

All the currents ought to be in the directions shown in the figure, and should be of reasonable magnitude. Note that the high ends of these, and several other constraints we will specify, may seem like they would follow obviously from other constraints, but it's worth having them anyway. The constraint programming system is not specific to electrical engineering, it *only* solves systems of numerical constraints, so any "obvious" constraint we fail to include may result in bogus solutions or no solution at all.

$$0.1\mu\mathrm{A} \le I_i \le 100\mathrm{mA} \text{ for } i \in \{1,2,3,4,5,6\}$$

Both unspecified voltages have to be within the power supply range, and indeed, not too close to the rails.

$$\begin{aligned} -10\mathrm{V} \le V_1 &\le +10\mathrm{V} \\ -10\mathrm{V} \le V_2 &\le +10\mathrm{V} \end{aligned}$$

For the green state, we have current I_1 through the series combination of R_1 and R_2, as a result of the voltage drop from +11.8V to +2.15V; Ohm's Law applies.

$$11.8\mathrm{V} - 2.15\mathrm{V} = I_1(R_1 + R_2)$$

For the other states, we apply Ohm's Law individually to each resistor and its associated current.

$$\begin{aligned} 11.8\mathrm{V} - V_1 &= I_2 R_1 \\ 10.6\mathrm{V} + V_1 &= I_2 R_3 \\ 11.8\mathrm{V} - V_2 &= I_3 R_1 \\ -1.75\mathrm{V} - V_2 &= I_4 R_2 \\ 10.6\mathrm{V} + V_2 &= I_5 R_3 \\ 10.6\mathrm{V} + V_2 &= I_6 R_4 \end{aligned}$$

In the red state, we have two nonzero currents combining and then splitting, which must add up by one of Kirchhoff's Laws.

$$I_3 + I_4 = I_5 + I_6$$

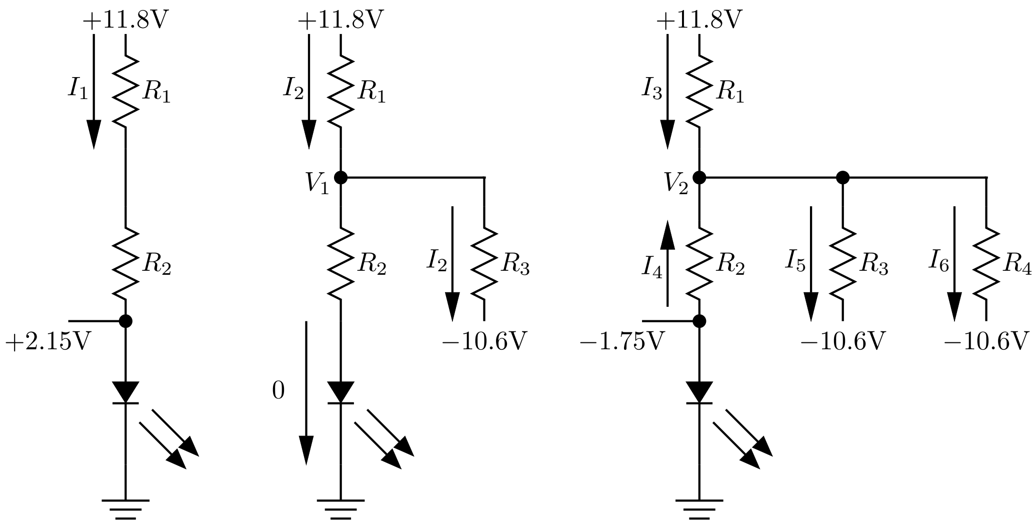

Figure 4: Variable names for designing the LED network.

The LED data sheet recommends 20mA of current in either direction, and says that that much current will give a typical luminous intensity of 50mcd (millicandelas) red and 15mcd green. In fact we want the LED (especially in the red state) to be somewhat dimmer than that so as not to dazzle the users, and we want it to be roughly equally bright in both states. I did some tests on the breadboard and found that these LEDs give a brightness ratio that looks good to me when the red current is between about 20% and 26% of the green current. The best green brightness seems to be with a current between 10mA and 20mA, and in order to be sure that the LED will not perceptibly glow in the off state, the voltage applied when it's meant not to glow should be no more than 1V in either direction. Those translate into constraints on the variables as follows.

$$10\text{mA} \le I_1 \le 20\text{mA}$$
$$0.20 I_1 \le I_4 \le 0.26 I_1$$
$$-1\text{V} \le V_1 \le +1\text{V}$$

The LM339's current-sinking capability is quoted as 15mA on some data sheets and 18mA on others. In the interest of keeping everything safe, we will not ask it to sink more than 13mA.

$$I_2 \le 13\text{mA}$$
$$I_5 \le 13\text{mA}$$
$$I_6 \le 13\text{mA}$$

Finally, we would like to use no bigger than 1/4W resistors, so to allow some safety margin, none of them should dissipate more than 230mW of power in any of the three states.

$$I_1^2 R_1 \le 230\text{mW}$$
$$I_1^2 R_2 \le 230\text{mW}$$
$$I_2^2 R_1 \le 230\text{mW}$$
$$I_2^2 R_3 \le 230\text{mW}$$
$$I_3^2 R_1 \le 230\text{mW}$$
$$I_4^2 R_2 \le 230\text{mW}$$
$$I_5^2 R_3 \le 230\text{mW}$$
$$I_6^2 R_4 \le 230\text{mW}$$

These constraints are all coded up in the file `lednetwork.ecl`, along with some sets of standard resistor values (which may be of use in automating searches). Figure 5 shows the start of an interactive session with ECL^{i}PS^e, first loading the file and then checking the immediate consequences of the constraints. The system says that in order for all the constraints to be satisfied, R_1 needs to be between about 831Ω and 965Ω, R_2 can be at most about 134Ω, R_3 needs to be between about 758Ω and 1036Ω, and R_4 between about 734Ω and 2026Ω. It has also concluded that we won't be able to have more than 11.6mA current through the LED in the green state. Further queries at the interactive prompt allow narrowing the ranges further and testing different standard values within the ranges.

```
ECLiPSe Constraint Logic Programming System [kernel]
Kernel and basic libraries copyright Cisco Systems, Inc.
and subject to the Cisco-style Mozilla Public Licence 1.1
(see legal/cmpl.txt or http://eclipseclp.org/licence)
Source available at www.sourceforge.org/projects/eclipse-clp
GMP library copyright Free Software Foundation, see legal/lgpl.txt
For other libraries see their individual copyright notices
Version 6.1 #168 (x86_64_linux), Fri Sep 27 17:21 2013
[eclipse 1]: ['lednetwork.ecl'].
source_processor.eco loaded in 0.01 seconds
hash.eco   loaded in 0.00 seconds
compiler_common.eco loaded in 0.01 seconds
compiler_normalise.eco loaded in 0.01 seconds
compiler_map.eco loaded in 0.00 seconds
compiler_analysis.eco loaded in 0.00 seconds
compiler_peephole.eco loaded in 0.01 seconds
compiler_codegen.eco loaded in 0.01 seconds
compiler_varclass.eco loaded in 0.00 seconds
compiler_indexing.eco loaded in 0.01 seconds
compiler_regassign.eco loaded in 0.00 seconds
asm.eco    loaded in 0.02 seconds
module_options.eco loaded in 0.00 seconds
ecl_compiler.eco loaded in 0.08 seconds
ic_kernel.eco loaded in 0.01 seconds
linearize.eco loaded in 0.00 seconds
ic_constraints.eco loaded in 0.01 seconds
ic.eco     loaded in 0.00 seconds
ic_generic_interface.eco loaded in 0.00 seconds
ic_search.eco loaded in 0.02 seconds
ic.eco     loaded in 0.05 seconds
lednetwork.ecl compiled 21112 bytes in 0.06 seconds

Yes (0.14s cpu)
[eclipse 2]: resistor_problem([R1,R2,R3,R4,V1,V2,I1,I2,I3,I4,I5,I6]).
lists.eco  loaded in 0.01 seconds

R1 = R1{830.76923076923083 .. 965.00000000000057}
R2 = R2{0.0 .. 134.23076923076962}
R3 = R3{758.07692307692253 .. 1036.4814814814822}
R4 = R4{733.57553218073247 .. 2026.4822252159038}
V1 = V1{-0.74500000000000632 .. 1.0}
V2 = V2{-2.1553893518518534 .. -1.75}
I1 = I1{0.01 .. 0.011615740740740746}
I2 = I2{0.01119170984455958 .. 0.013}
I3 = I3{0.014041450777202064 .. 0.016638865702079935}
I4 = I4{0.0019999999999999879 .. 0.0030200925925925942}
I5 = I5{0.0081473820796855376 .. 0.011674277016742781}
I6 = I6{0.004367173760459267 .. 0.011511576214986993}

There are 63 delayed goals. Do you want to see them? (y/n)
Yes (0.01s cpu)
```

Figure 5: An interactive session with ECL^{i}PS^e.

Mechanical drawings

On the following pages you will find:

- the schematic diagram for the module;
- a mock-up of what the completed module looks like from the front panel;
- the top-side silk screen art showing component placement;
- the bottom-side silk screen art showing component placement (*note this drawing is mirrored, and shows what you actually see looking at the board, not the X-ray view used in other Kicad output*);
- a mechanical drawing of the front panel showing the locations and sizes of the holes in it; and
- an exploded isometric drawing showing how the boards and hardware fit together.

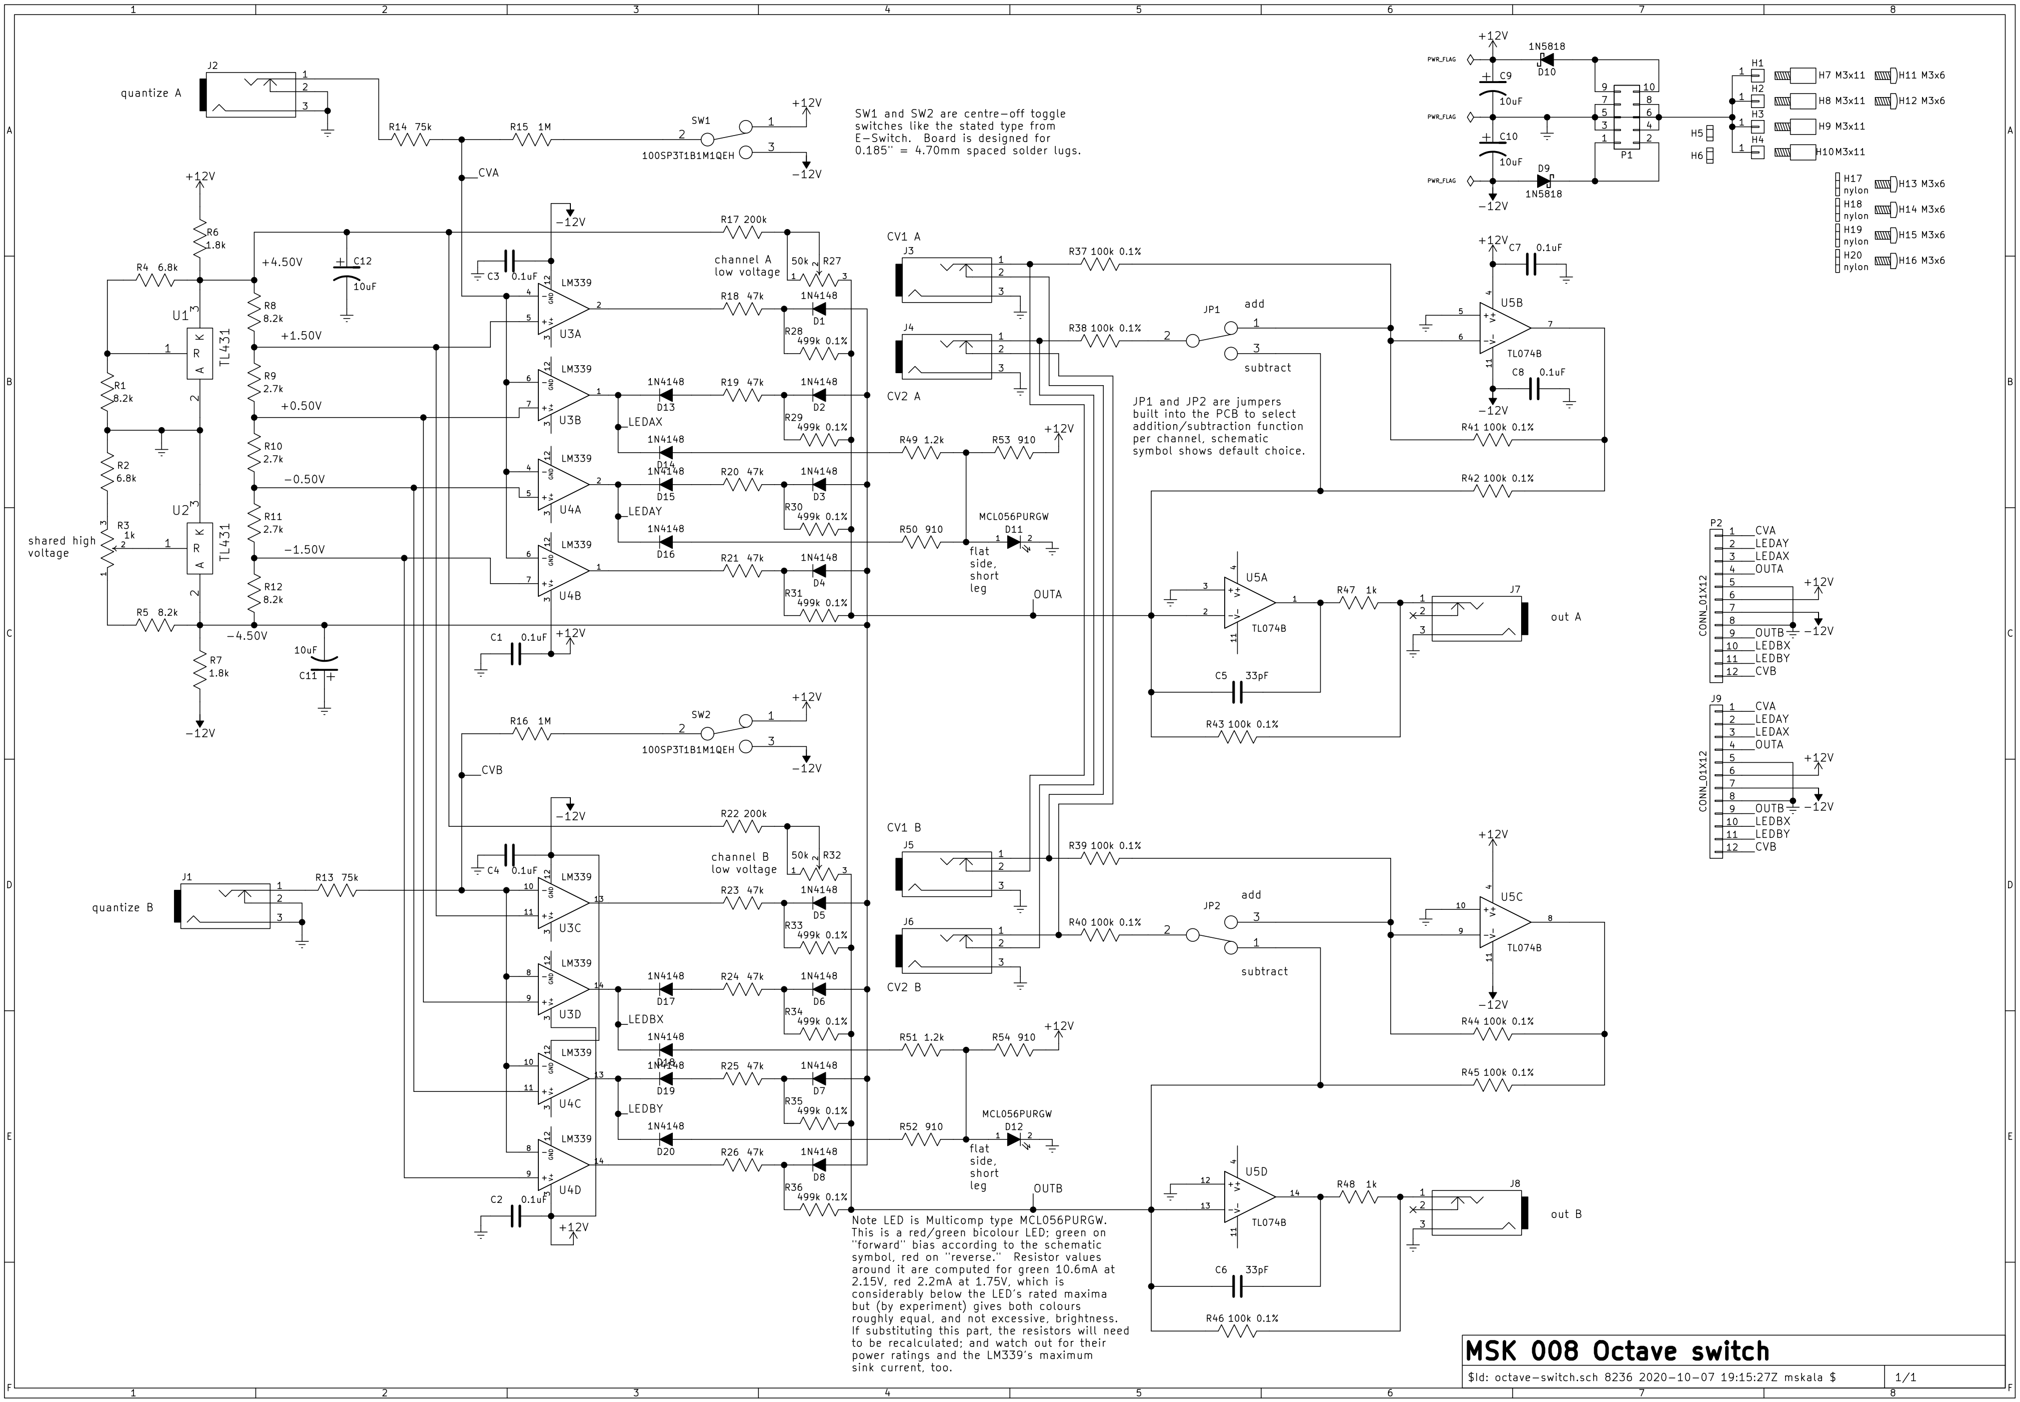
MSK 008 Octave switch
$Id: octave-switch.sch 8236 2020-10-07 19:15:27Z mskala $
1/1
SW1 and SW2 are centre-off toggle switches like the stated type from E-Switch. Board is designed for 0.185" = 4.70mm spaced solder lugs.
JP1 and JP2 are jumpers built into the PCB to select addition/subtraction function per channel, schematic symbol shows default choice.
Note LED is Multicomp type MCL056PURGW. This is a red/green bicolour LED; green on "forward" bias according to the schematic symbol, red on "reverse." Resistor values around it are computed for green 10.6mA at 2.15V, red 2.2mA at 1.75V, which is considerably below the LED's rated maxima but (by experiment) gives both colours roughly equal, and not excessive, brightness. If substituting this part, the resistors will need to be recalculated; and watch out for their power ratings and the LM339's maximum sink current, too.
quantize A
quantize B
CV1 A
CV2 A
CV1 B
CV2 B
out A
out B
channel A low voltage
channel B low voltage
shared high voltage
add
subtract
flat side, short leg
100SP3T1B1M1QEH
+4.50V
+1.50V
+0.50V
−0.50V
−1.50V
−4.50V
CVA
CVB
OUTA
OUTB
LEDAX
LEDAY
LEDBX
LEDBY
CONN_01X12

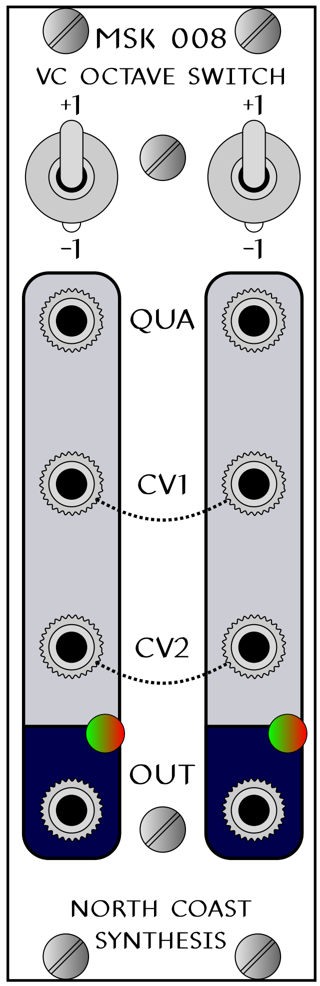
MSK 008
VC OCTAVE SWITCH
+1
-1
+1
-1
QUA
CV1
CV2
OUT
NORTH COAST
SYNTHESIS

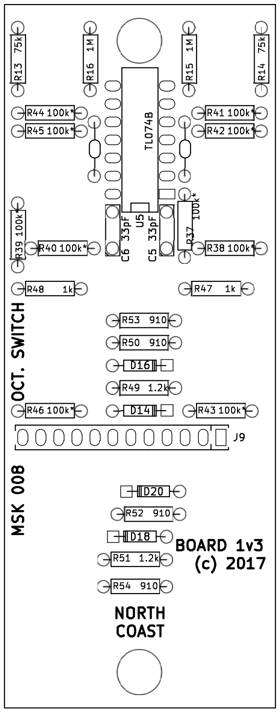

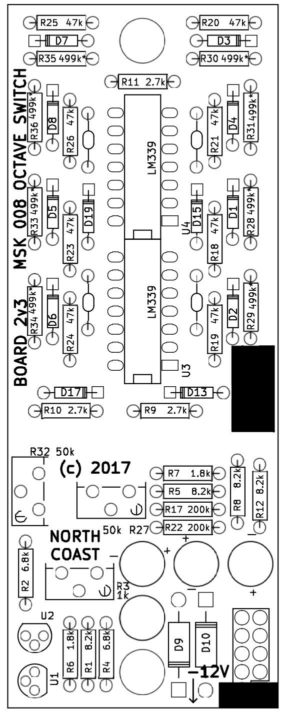

MSK 008 Octave switch

$Id: octave-switch.kicad_pcb 5603 2017-08-17 19:02:50Z mskala $

1/2

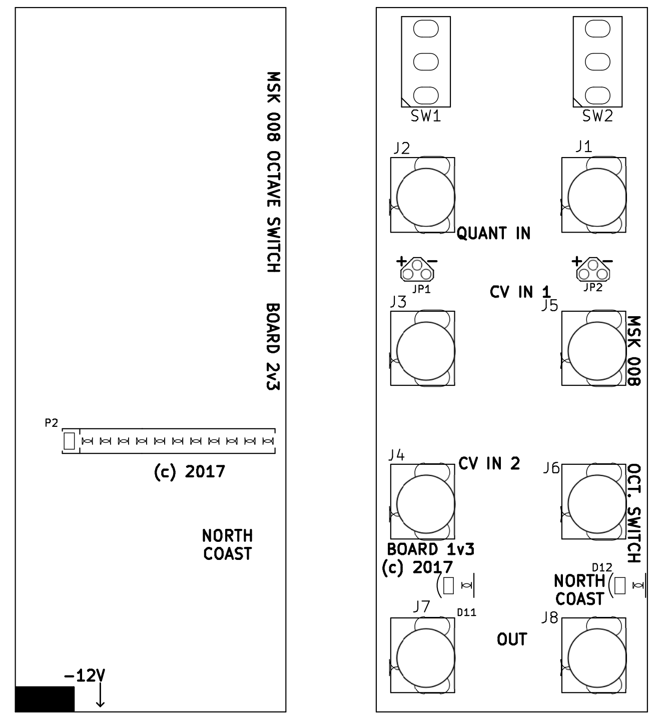

MSK 008 Octave switch

$Id: octave-switch.kicad_pcb 5603 2017-08-17 19:02:50Z mskala $

1/2

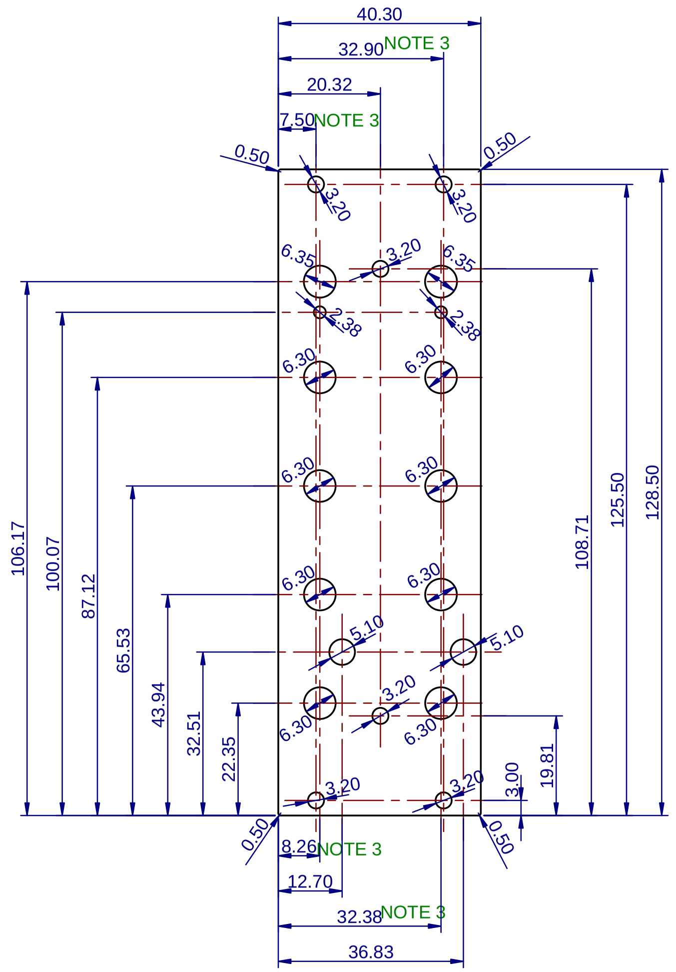

NOTES

1: All dimensions are millimetres with tolerance +/- 0.10mm.
Do not scale off drawing.

2: Material is aluminum, 2.00mm thick. Overall dimensions are 128.50mm by 40.30mm, per Doepfer spec for an 8HP Eurorack synthesizer module.

3: Do not confuse similar horizontal positions. The 3.20mm holes are centred 7.50mm and 32.90mm from the left edge. The 2.38mm, 6.30mm, and 6.35mm holes are centred 8.76mm and 32.38mm from the left edge.

MSK 008 VC Octave Switch panel mechanical	
$Id: panel-mechanical.dxf 5279 2016-12-20 03:09:11Z mskala $	1/1

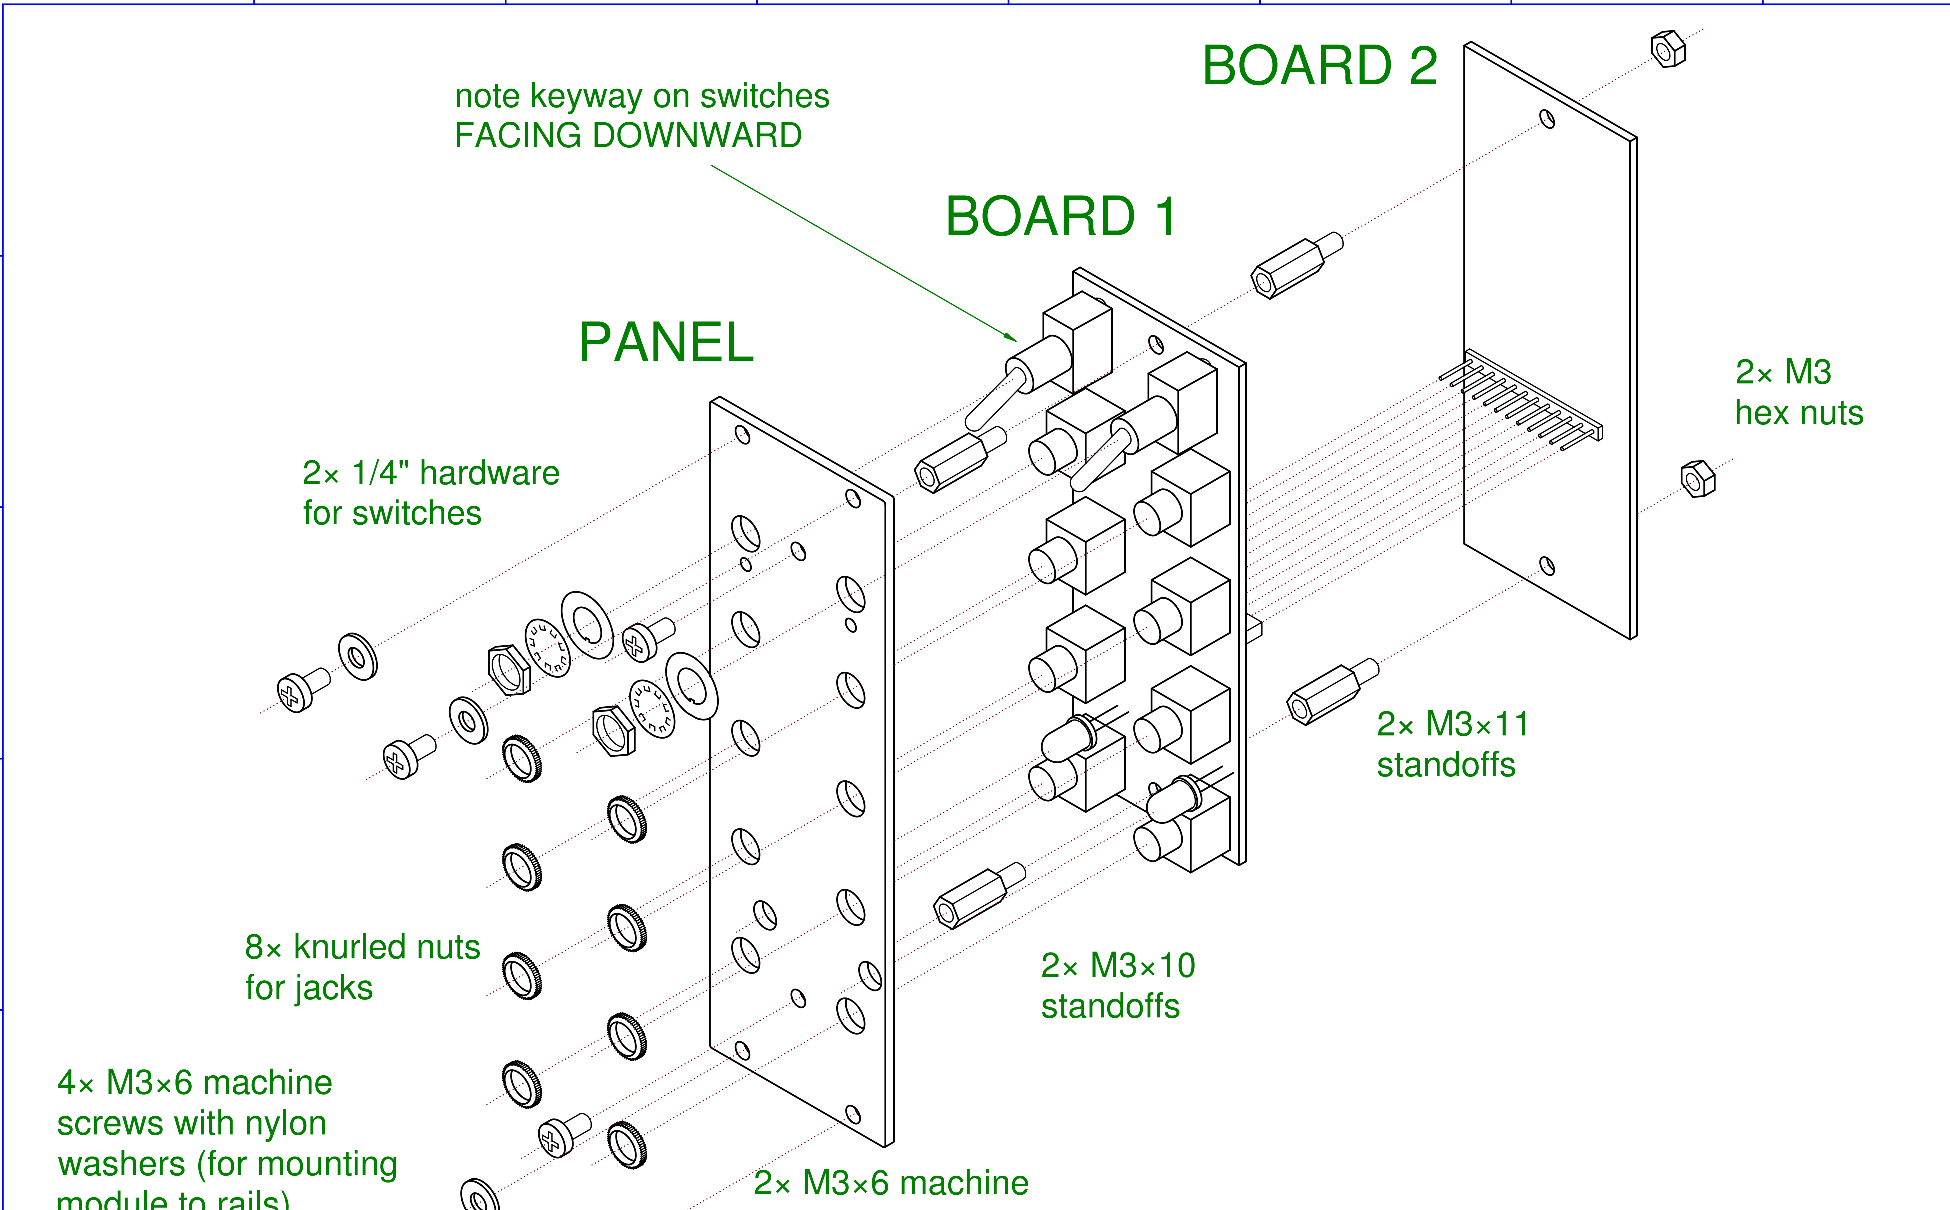

MSK 008 VC Octave Switch assembly diagram

$Id: exploded.dxf 5604 2017-08-17 19:13:53Z mskala $

1/1