

Analysis of the SMR-4 VCF/VCA board

Second revision

Émilie Gillet

emilie.o.gillet@gmail.com

The SMR-4 mkII board is the second revision of the “default” analog signal processing board for the Shruthi-1. It includes a 4-pole VCF (presented in section 2) with VC-resonance and a linear VCA (presented in section 3); and only makes use of inexpensive and widely available ICs.

1 Generalities

1.1 Power

The board is powered from a non-regulated single supply rail in the 7.5V-15V range. A negative rail is obtained from this by means of a LT1054 DC-DC converter; charged into a $100\mu F$ capacitor – or $110\mu F$ equivalent when a $10\mu F$ tantalum cap is added in parallel – and set to a switching frequency of 50kHz. The input positive rail and the negative rail from the LT1054 are regulated by 7805/7905 linear regulators and generously filtered.

The main advantage of working with +/- 5V rails is that no extra regulator is required for powering the digital section. Furthermore, the device can be powered by a 9V battery, or even a pack of five 1.5V batteries. The main drawback is the smaller headroom for all the op-amps – the LT074 clips at around $\pm 3.7V$.

1.2 Input signals

All signals generated by the Shruthi-1 digital control board, be it the raw oscillators signal or the control voltages, are 0 / 5V PWM signals with a carrier frequency of $\frac{20MHz}{510} = 39215Hz$.

This implies that the control signals have to be smoothed. You should not feel bad about it: the MCU generates the control signals at a “slow” rate of 980Hz anyway, so a simple 1-pole low-pass with a cutoff frequency of 723Hz kills the PWM carrier by 35dB while still tracking fast enough the fastest transitions the MCU can create.

The oscillators signal also contains this 39kHz carrier. The carrier is attenuated by the main 4-pole LPF itself, and to some extent by the excessively large “stability” caps in the mixer and VCA I to V converters. In the worst case (cutoff set to its maximum value), the carrier is thus attenuated by 30dB. In most cases, however, the cutoff frequency is set to a lower value, and the carrier is attenuated more strongly. Anti-aliasing filters at the input of your soundcard, speakers with a limited bandwidth, or your ears will be doing the rest of the filtering. This residual of the PWM carrier might however cause a bit of parasitic hiss if the output signal is fed into a “vintage” sampler of FX processor with no anti-aliasing filter (the frequency of the hiss being dependent on the difference between the sampling rate and 39kHz).

$$V_{out}(p) = -I_c(p)Z_{feedback}(p) \quad (5)$$

$$= -\frac{I_c(p)}{Cp} \quad (6)$$

$$= -19.2I_{freq}\frac{R_s}{(R_b + 2R_s)Cp}(V_{in}(p) + V_{out}(p)) \quad (7)$$

Solving for $V_{out}(p)$ and dividing by $V_{in}(p)$, we find the transfer function:

$$H(p) = \frac{-1}{1 + \frac{(R_b + 2R_s)}{R_s}Cp\frac{1}{19.2I_{freq}}} \quad (8)$$

The pole is at the frequency f such that:

$$0 = 1 + \frac{(R_b + 2R_s)}{R_s}C2\pi f\frac{1}{19.2I_{freq}} \quad (9)$$

$$f = -\frac{19.2I_{freq}}{2\pi\frac{(R_b + 2R_s)}{R_s}C} \quad (10)$$

At this stage, we can already pick values for R_b , R_s , C and design the exponential current source so that its current range translate into a target frequency range using the formula above.

The values which have been chosen for R_b , R_s and C follow closely those given in the application schematics of the SSM2040 datasheet (Figure 9 of the datasheet):

$$R_b = 10k\Omega$$

$$R_s = 220\Omega$$

$$C = 1nF$$

$R_s = 200\Omega$ in the SSM2044 datasheet, but we have privileged here values in the E12 series.

2.2 Exponential current source

As explained in the previous section, the cutoff frequency is proportional to the biasing current I_{freq} of the OTAs. Since we want to achieve an exponential frequency control, we need to produce a control current which is an exponential function of the input CV. This is achieved by the exponential current source in figure 2.

Let us start with the core of the circuit – the PNP transistor pair Q_2 and Q_3 . The Ebers-Moll equations for Q_2 and Q_3 give the collector currents for both transistors:

$$I_{c2} = I_{s2} \left(\exp\left(\frac{V_{eb2}}{V_T}\right) - 1 \right) \quad (11)$$

$$I_{c3} = I_{s3} \left(\exp\left(\frac{V_{eb3}}{V_T}\right) - 1 \right) \quad (12)$$

V_T is the thermal voltage, equal to $26mV$ at room temperature. We have $|V_{e2}| < 200mV$ (we will see later that this is the range giving cutoff frequencies in the audio range), $V_{b2} \simeq 0.7V$ (1

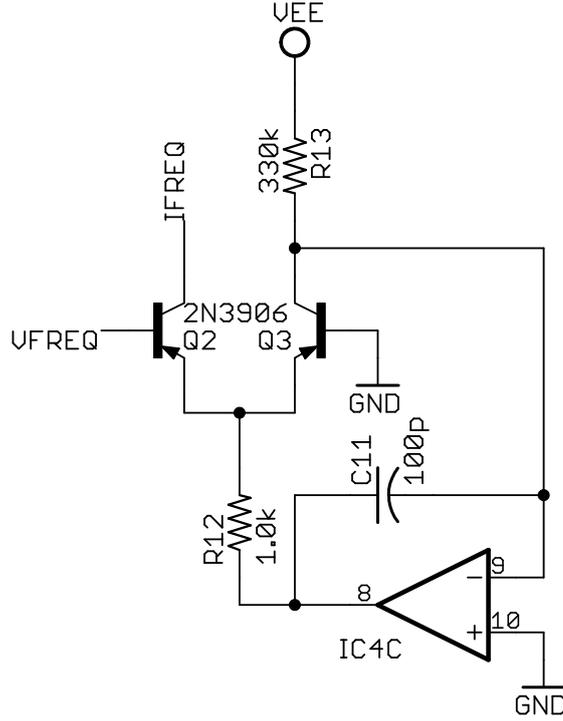


Figure 2: Exponential current source.

diode drop) ; so $V_{eb2} \gg V_T$. It is thus a sane approximation to ignore the -1 term. Assuming that the two transistors have the same characteristics, $I_{s2} = I_{s3}$. Thus we have, by substituting in the first equation the value of I_{s2} derived from the second one:

$$I_{c2} = I_{c3} \exp\left(\frac{V_{eb2} - V_{eb3}}{V_T}\right) \quad (13)$$

Since the base of $Q3$ is grounded, and since the emitters of $Q2$ and $Q3$ are at the same potential:

$$V_{eb2} - V_{eb3} = V_{e2} - V_{b2} - V_{e3} + V_{b3} = -V_{b2} = -V_{freq} \quad (14)$$

The value of I_{c3} is simply:

$$I_{c3} = \frac{V_{ee} - V_{c3}}{R_{13}} = \frac{V_{ee}}{R_{13}} \quad (15)$$

V_{c3} is null because the collector of $Q3$ is virtually grounded. Indeed, the whole point of the operational amplifier $IC4C$ is to keep the collector of $Q3$ at a constant potential. C_{11} is only here to stabilize the op-amp. R_{12} works as a current limiter; we can assume that it has no influence in normal operation.

To summarize the previous steps:

$$I_{freq} = -\frac{V_{ee}}{R_{13}} \exp\left(\frac{-V_{freq}}{V_T}\right) \quad (16)$$

Observe that the biasing current I_{freq} varies in the inverse direction of V_{freq} – when V_{freq} increases, I_{freq} decreases. This is not a problem because V_{freq} is generated by an op-amp in an inverting configuration.

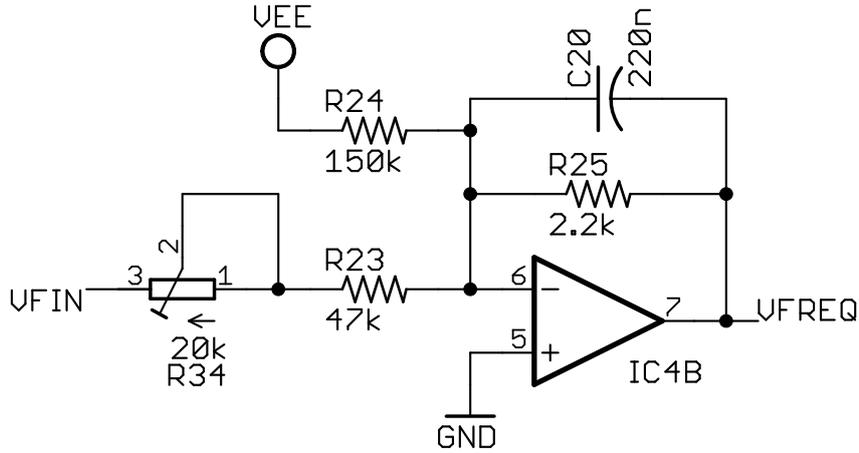


Figure 3: CV-scaling.

One thing not discussed here is temperature compensation. While the main temperature dependent factors (saturation currents) have been neutralized by the “transistor pair” design (as long as the transistors are of a similar kind and thermally coupled), there remains a temperature dependent term, V_T . This term is not compensated. A rough estimate of the variation in frequency is less than $\pm 10\%$ for a change in temperature of $\pm 5^\circ C$ – not good for a VCO, but acceptable for a VCF.

2.3 CV scaling

The CV scaling stage is fed with a 0/5V PWM signal generated by the microcontroller, and generates a scaled voltage which will translate into an appropriate cutoff frequency range. In addition to the scaling itself, this stage also actively smoothes the high-frequency PWM control signal into a near-DC voltage. The schematics of this stage are given in figure 3.

The voltage at the output of the op-amp is given by:

$$V_{freq}(p) = -\frac{R_{25}}{R_{24}}V_{ee} - \frac{Z_f(p)}{R_{20} + R_{23}}V_{fin}(p) \quad (17)$$

Where $Z_f(p)$ is the equivalent impedance of the feedback loop, consisting of R_{25} in parallel with C_{20} . Note that in the first term, this impedance is simply R_{25} since the capacitor in parallel with the resistor has no effect on DC currents. In the second term, it is equal to:

$$Z_f(p) = \frac{R_{25}}{1 + R_{25}C_{20}p} \quad (18)$$

Thus, the cutoff frequency of the filter smoothing the PWM control signals is $\frac{1}{2\pi R_{25}C_{20}} = 329Hz$ with the values given on the schematics. This means that the PWM carrier is attenuated by $6 \log_2 \frac{39215}{329} = 40dB$. Given this (moderately) high attenuation factor, it is safe to consider that when the cutoff value is constant, the circuit behaves as if V_{fin} is a pure DC component in the $[0, 5]V$ range instead of a PWM signal.

Note that the use of a supply rail (V_{ee}) to generate a negative reference voltage is not a very good practice. This makes the circuit vulnerable to ripples/noise in the negative rail ; and requires accurate voltage regulators. In practice, we have not observed significant noise/ripples on the negative rail ; and small random fluctuations of the cutoff are much less perceptible than noise.

2.4 An important consideration

All the computations in the previous steps are done under the assumption that the exponential current source is connected to only one OTA. In the SMR4 mkII filter, the exponential converter is connected to 4 OTAs. Assuming the OTAs have similar characteristics, the biasing current I_{freq} received by each OTA is only one fourth of the value of I_{c2} . A more common solution seen in other designs is to duplicate 4 times the left side of figure 2. Not doing so saves 3 transistors but has a drawback: in case one of the OTA cells fails, the current received by the other cells will be higher. This put a constraint on the current limiting resistor R_{12} : if it is too low, the maximum current emitted by the exponential converter will be acceptable for 4 cells but not for 2 or 3 in case of failure ; and if it is too high, the maximum current will not give a sufficiently high cutoff frequency.

2.5 Component values

2.5.1 Constraints

How are the different values of the component selected? From the past sections, some characteristics of the ICs used in the circuit, and some common sense, we can gather a set of constraints the component values must satisfy:

Input impedance The input impedance of the CV scaling circuit $R_{23} + R_{34}$ should not be too low, at least $10k\Omega$.

OTA small biasing current constraint The biasing current of the OTA should not exceed $2mA$.

Cutoff frequency constraint The frequency range of the filter must span the audible spectrum ($20Hz$ to $20kHz$), and it should be easy to "tune" it to the fundamental frequency of the oscillators. To this effect, we have decided to eschew the $1V/octave$ standard (which would have given a very narrow frequency range given that the input signal is in the $[0, 5]V$ range, and to use another approach.

In the Shruthi-1 firmware, the filter tracking code works the following way: the MIDI note value (an integer between 0 and 127) minus 64 is added to the filter cutoff value (which is itself represented as a 7-bit integer). In other words, if the base cutoff setting of the Shruthi-1 is set to n , the PWM control signal corresponds to a smoothed DC voltage of:

$$V_{fin} = V_{cc} \frac{n + \text{Midi note number} - 64}{128} \quad (19)$$

From this, it appears that if we want a 1:1 tracking of the filter frequency to the fundamental frequency of the oscillators, the cutoff frequency should be multiplied by a semitone ($2^{\frac{1}{12}}$) every time the Midi note number is increased by 1, that is to say every time V_{fin} steps by a $\frac{5}{128}$ increment. This yields the rather odd scale of $0.46875V.oct^{-1}$, which is the target $V.oct^{-1}$ scale of the filter.

The filter is thus expected to cover a range of $\frac{128}{12} = 10.66$ octaves – the highest cutoff value will be 1625 times the lowest cutoff value. The $14.9Hz - 22.8kHz$ range has this exact ratio and spans the audible range.

Pragmatic constraints Resistor values should be preferably in the E12 series ; and the number of distinct component values should be minimized.

2.5.2 Validation

We provide in the rest of this section a validation of the selected values.

Input impedance Greater than $47k\Omega$, all good!

OTA small current constraint At worst, the output of $IC4C$ is $3.7V$ (clipping output voltage for a $TL074$) ; and we have $V_{e2} = V_{e3} = 0.7V$. The maximum current produced by the exponential converter is thus $3.0mA$ with $R_{12} = 1k\Omega$. This is enough to protect one of the LM13700 (2 OTA cells) in case the other one is not inserted in its socket.

Cutoff frequency range Let us check first that the lowest cutoff frequency attainable with the filter has the expected value. When $V_{fin} = 0V$, $V_{freq} = -73mV$ and $I_{freq} = 0.23\mu A$. The lowest cutoff frequency is thus: $14.5Hz$.

Let us now check what is the maximum cutoff frequency that can be reached. V_{fin} is now set to $5V$. Depending on how R_{34} is trimmed, we have:

$$0\Omega \leq R_{34} \leq 20k\Omega \quad (20)$$

$$91mV \leq V_{freq} \leq 161mV \quad (21)$$

$$0.125mA \leq I_{freq} \leq \frac{3.0}{4}mA \quad (22)$$

$$8kHz \leq f \leq 48kHz \quad (23)$$

The target frequency range is thus feasible. Note that the exact 1:1 tracking is achieved with $R_{34} = 10.2kHz$, so the default factory position of the trimmer (middle) is roughly the right one.

3 VCA

3.1 Linear current source

The linear current source in figure 4 generates the biasing current for the VCA OTA cell. Let us make a few assumptions:

- The input signal is DC, and thus, $C9$ has no effect (actually its purpose is to smooth the output).
- Q_1 has an α equal to 1, which implies that $I_{gain} = I_{c1} = I_{e1}$.
- The op-amp is not operating in saturation mode, which implies that $V_- = V_+ = 0$.
- The op-amp has infinite input impedance, which implies that $I_- = 0$.

Thus, we have:

$$I_{gain} = I_{e1} = I_{R8} = \frac{V_{gain}}{R_9} + \frac{V_{ee}}{R_{10}} \quad (24)$$

Intuitively, this circuit works the following way: through feedback, the op-amp ensures that the current flowing through R_8 is the same as the input current at the summing node – it does so by

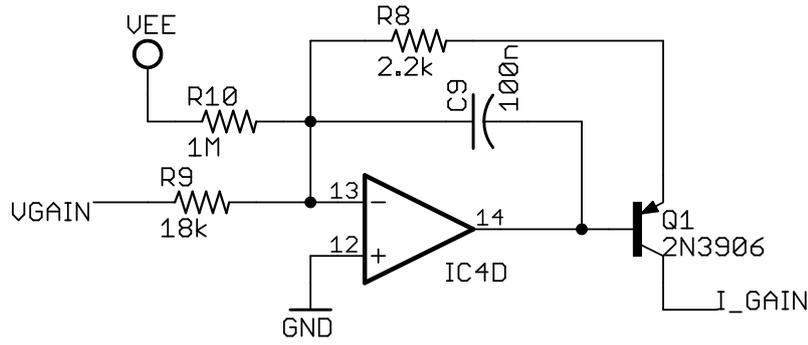


Figure 4: Linear current source.

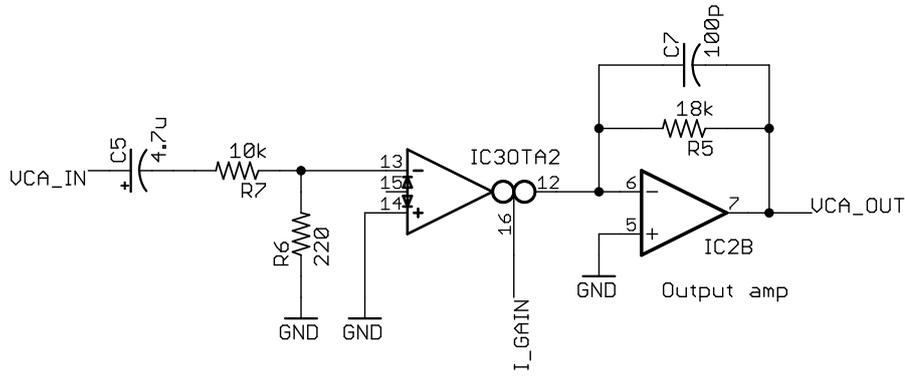


Figure 5: Current controlled amplifier.

controlling the base voltage of Q_1 . It is as if the op-amp is solving the Ebers-Moll equation for V_{b1} to give the target I_{e1} !

I_{gain} is thus directly proportional to V_{gain} . The additional offset term $\frac{V_{ee}}{R_{10}}$ helps “silencing” the VCA – the reason behind this is that the Shruthi-1 control board cannot actually output null voltages – when a PWM output on the MCU is set to 0 it outputs a few mV . R_{10} can be left unpopulated.

Observe that R_8 ; in conjunction with the op-amp’s clipping; can play the role of a current limiter on the output current, yielding a kind of reverse-exponential response with a “knee” when the op-amp starts failing to set its output high or low enough to guarantee that a large enough current flows through the transistor. The value chosen here is low enough to not cause current limiting.

3.2 Current controlled amplifier

The schematics in figure 5 represents a current controlled amplifier.

C_5 AC-couples the input. Assuming the input signal VCA_IN has no DC-offset, this capacitor can be ignored. The voltage at the inverting input of the OTA is:

$$V_- = \frac{R_6}{R_6 + R_7} VCA_IN \quad (25)$$

The current at the output of the OTA is:

$$V_o = g_m(V_+(p) - V_-(p)) \quad (26)$$

$$= -19.2I_{gain} \frac{R_6}{R_6 + R_7} VCA_IN \quad (27)$$

Finally, this current is converted into a voltage through *IC2B* which works as a transimpedance amplifier:

$$VCA_OUT = -R_5V_o \quad (28)$$

Using all these, the gain of the VCA is:

$$G = 19.2R_5 \frac{R_6}{R_6 + R_7} \left(\frac{V_{gain}}{R_9} + \frac{V_{ee}}{R_{10}} \right) \quad (29)$$

This yields a gain of 2 when the V_{gain} control signal is equal to its maximum value of $5V$.

4 Resonance control

Resonance is achieved by feeding back a fraction of the filter output signal into the filter input. Traditionally, this is done through a potentiometer (voltage divider). Here, in order to provide voltage control over resonance, we use a VCA similar in design to that of the main signal VCA. There are two important differences in the design:

First, the current output by the OTA is not converted back into a voltage by an op-amp. Instead, it is directly injected at the summing node of the first OTA-integrator cell. A way to look at it: the 220Ω resistor to ground at the OTA-integrator cell inverting input does the current to voltage conversion.

Then, the resonance VCA is fed not only with the filter output, but also with a fraction of the unfiltered input signal itself. This causes the level of the signal entering the filter to increase as resonance is increased, and helps in compensating the dreaded “resonance loudness drop” of 4-pole filters. An extra refinement here is that the the input signal is coupled into the resonance VCA through an undersized capacitor. This will cause the resonance loudness drop to be compensated only in the highest regions of the spectrum ; creating brighter and shinier sounds when the resonance is increased. This trick was inspired by the Jupiter-8 filter circuit.

Foreword: learning from mistakes

18 months and roughly 1000 units have happened between the design of the original filter (which was my first shot at doing analog synth circuits) and this revision. This section briefly discusses some of the changes:

- The main “awkwardness” in the original design was the exponential converter. All the calculations had been done with a pair of NPN; until I realized in shock that the current biasing pin on the LM13700 is below ground. My attempt to save the design with PNPs “mirroring” the current from the exponential converter was rather inelegant.
- During the design of the first version, many efforts had been done to get all the part values exactly right, which sometimes requested additional voltage dividers to get required voltage ratios using E12 or E24 values – this can be seen in the filter cutoff CV scaler. This ultimately proved futile as some parts have rather wide tolerances (at best 2.5% for capacitors); and thus came the trimmers.

- **Trimmers are evil, do not use them.** Some builders are obsessed by them and are chasing a magical combination of settings giving "the perfect sound". They distract the attention of some builders trying to identify a build problem; while some builders having committed grave assembly mistakes do not bother fixing them because they wrongly assume the bad sound comes from inadequate trimming. Finally, giving choices to users is not very compatible with the diffusion and sharing of patch libraries ("how my patch will sound on someone else's device which has a different cutoff range?"). I am now keeping trimmers for the strict minimum – V/Oct – tuning and would rather rely on sane defaults even if they are at $\pm 5\%$.
- I made a point of having a DC-coupled VCA in the first version to have the flattest response in the low-end. This requested the adjustment of a trimmer to compensate for a slight DC offset at the output of the filter – and I suspect that a significant proportion of builders ignored it or got it wrong. I am no longer scared of AC-coupling and losing half a dB at 10 Hz.
- On the original filter, the compensation for the resonance loudness drop was done at the final VCA. Here it is done directly within the resonance VCA - a simpler design I learnt from looking at the schematics of IR3109-based filters.
- The original design had a few "current to voltage to current" blocks that came from piecing together basic bits of circuits but could ultimately be eliminated.
- Some better choices of parts: DC jacks with thick legs (easier to source), trimmers with side adjustment (trimming can be done without disassembling the unit), 0.1in space solder pads for pots.

To conclude, a list of things that did not work well during the design of this filter and that I had to throw away:

- Unipolar power supply and virtual ground. In my quest of getting rid of the expensive LT1054, I worked on a first design with a 9V unipolar supply and a virtual ground. In spite of trying many different techniques for getting a clean middle rail (buffered rail splitter IC ; precision voltage reference ; voltage regulator), noise was always a problem and the virtual ground circuit always performed 5 or 6 dB worse than its split supply equivalent. Even with copious (to the point of being more expensive than a LT1054) decoupling. Worst part: those super useful 9V PP3 battery to 2.1mm jack clips could not be used, since an input voltage above 11V had to be provided to the filter board.
- AC power supply. What's more expensive than a LT1054? A pair of heatsinks and big caps.
- LM13700 Darlington buffers. The "OTA-C-buffer" configuration is actually very common, but I could not get it to work with the LM13700 Darlington in a satisfying way. The main problem was self-oscillation, which, for different values of the cutoff frequency, did not appear at the same values of the resonance VCA gain. Either a loud, Curtis-like, distorted sine wave for high cutoff values and a normal sine wave for the lowest values ; or a normal sine wave for high cutoff values and no self-oscillation for the lowest values. It was as if the Darlington buffers had a slight high-pass effect.
- LM13700 linearization diodes. Where's the fun gone?