AN701 APPLICATION NOTE

Designing Voltage Controlled Filters for Synthesizers with the SSI2164

By Jeroen Allaert

INTRODUCTION

The SSI2164 was originally intended as a voltage controlled amplifier, but is in fact a voltage controlled current amplifier that's useful for countless other applications. Four matched VCA's on a single piece of silicon with exponential control makes this IC especially suitable for synthesizer VCF design. Its low noise and exceptional linearity are a huge benefit while its main weakness – the limited compliance of its outputs – is actually its greatest strength from a sound designer's point of view.

In this application note you'll read how to follow the path by sticking to what's advised in the datasheet, avoid some pitfalls, and design VCFs with terrific audio performance. We'll also stray from the path to design VCFs offering a bit more character, including re-creation of several classic synthesizer filter topologies with excellent results.

ESSENTIAL FILTER BUILDING BLOCKS

Low-Pass

Figure 1 shows ***the*** essential building block[1] for designing VCFs with the SSI2164. It's a single pole low-pass stage, using a structure recognizable to those familiar with OTA-C LPFs found in many synthesizers using ICs such as the CA3080, LM13700, BA662, IR3109 and SSM2040. Differences are the lack of a voltage divider as the input is a ground-referenced current instead of a mV-range voltage, and inverting op-amp integrator instead of a ground-connected capacitor followed by a high-impedance unity-gain buffer.

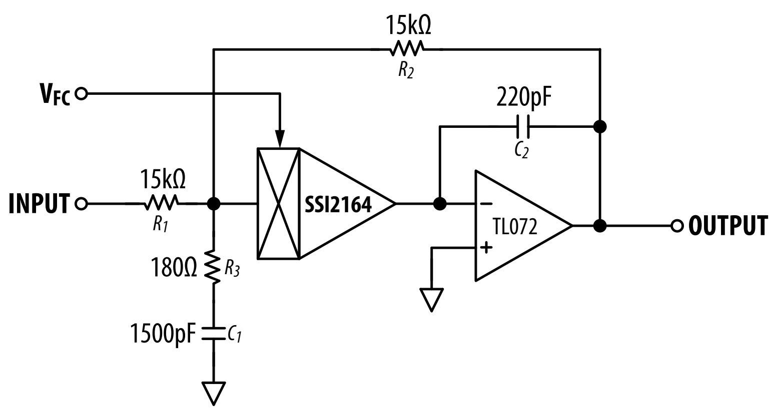

Figure 1: A Single-Pole Low-Pass Filter Stage

This inverting op-amp integrator serves two purposes. It facilitates necessary negative feedback because the SSI2164's inputs are non-inverting, and maintains the output current pin at virtual ground to stay within output compliance as specified in the data sheet.

[1] This is basically US Patent #3.805.091: "Frequency sensitive circuit employing variable transconductance circuit" by Dennis P. Colin, ARP Instruments Inc. from April 16th 1974, but done here with a current amplifier instead of a transconductance amplifier.

The corner frequency of this LPF stage is set by the RC time-constant of R_2C_2 and adjustable by the frequency control voltage at the SSI2164's control pins using an exponential ratio of roughly -200mV per octave[2]:

$$f_c = \frac{2^{-5V_f}}{2\pi R_2 C_2}$$

The gain is negative and set by R_1 and R_2.

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

Its Laplace domain transfer function, with

$$S = \frac{j\omega}{2\pi f_c}$$

is:

$$-\frac{1}{S+1}$$

When f_c is set to 1 kHz, the following Bode plots result[3]:

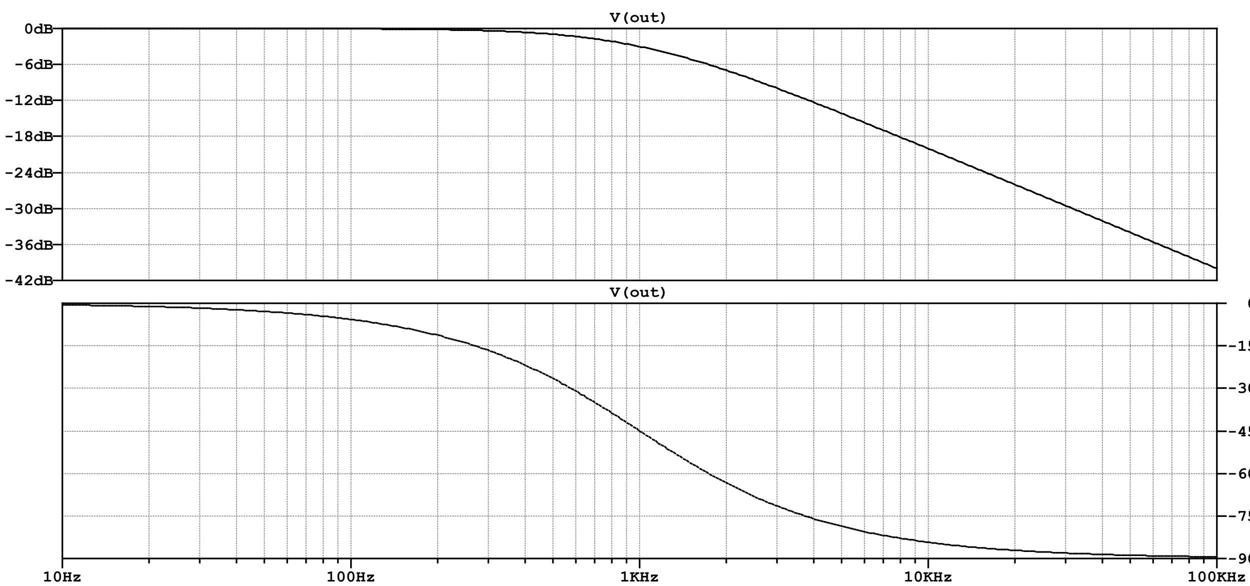

Figure 2: Bode Plots of the Single-Pole Low-Pass Filter Stage (shown non-inverting)

Component values shown in Figure 1 are dimensioned for unity gain, signal swing of ±15V, and a corner frequency of 50kHz with 0V at the control pin. With these values, V_{FC} should be positive-only voltages to maintain stability and prevent excessive power consumption. See "Dimensioning and Component Selection" later in this note.

High-Pass

"***The***" is emphasized at the beginning of this section because one might be tempted to translate familiar OTA-C high-pass, band-pass, etc. filter circuits to the SSI2164. Figure 3 shows an OTA-C high-pass filter stage built around a LM13700. Due to the virtually infinite output impedance of the LM13700 and input impedance of the op-amp, the input voltage is copied to the output of the op-amp. By negative feedback, the low-pass filtered version is subtracted from it, leaving only the high-frequency content at the output.

---

[2] This value of 200mV is calculated from the 33mV/dB Gain Constant in the datasheet, which takes into account self-heating of the device. The value at room temperature is 180mV/octave, which is exactly 10x (there is a 1/10 resistive voltage divider at the control ports inside the SSI2164) the scaling of a standard exponential converter. Use this value for spice simulations and thermistor calculations.

[3] In this document all Bode plots will be shown non-inverting (even when the circuit itself is inverting) and idealized to avoid confusion when combining multiple single filter stages as higher-order filters.

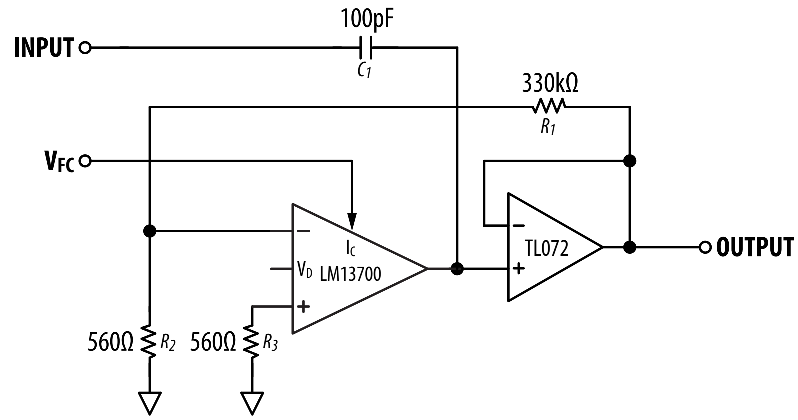

Figure 3: OTA-C High-Pass Filter Stage

(Dimensioned for ±15V signals and 50kHz @ I_F = 1mA

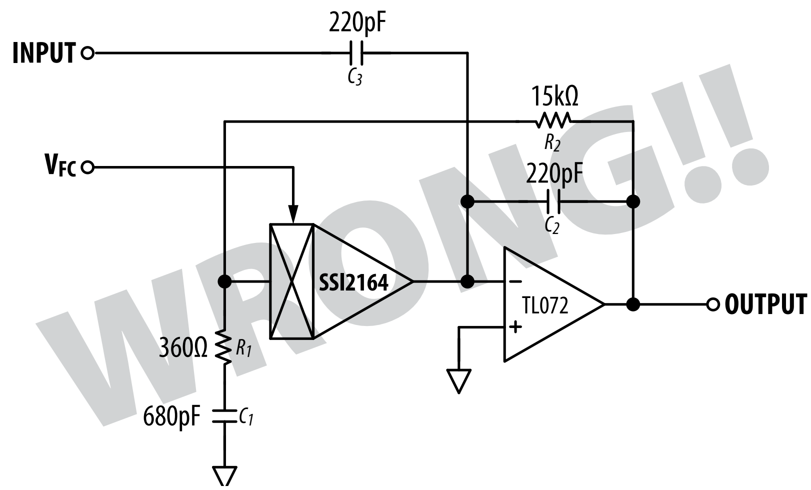

Figure 4: Non-Working High-Pass Filter Stage Using SSI2164

Another way to look at it: the node between the output of the OTA and input of the buffer behave like a current-controlled ground-connected resistor. The capacitor turns this in an RC high-pass filter.

Figure 4 shows the unwary translated version of the circuit in Figure 3, to the required ground-referenced input current of the SSI2164 and its output maintained at virtual ground. The output impedance of the SSI2164 is also virtually infinite, so the input voltage is inverted at the output of the op-amp, which forms a unity inverter with C_2 and C_3. Again, by negative feedback the low-pass filtered version is subtracted from this and should leave only the high-frequency content at the output. Unfortunately this does not work.

The output will saturate to either the negative or positive supply when increasing the frequency control voltage. The reason is as V_{FC} is increased the SSI2164's gain decreases, which in turn lowers the DC-feedback until it's not strong enough to compensate for the DC error that builds up on the integrating capacitor.

The DC error comes mainly[4] from the intrinsic asymmetric output impedance of op-amps. This causes overshoot and undershoot to be unequal, which results in a net DC-charge to be injected in the integrating capacitor with each cycle of the input signal. Higher input amplitude and frequency content increases this effect and causes saturation to happen sooner.

Note for Spice users: the output of op-amps is almost never properly modeled. These circuits might seem to work when simulated. In real-life they do not. This issue will affect any circuit where the input signal is directly fed into the op-amp integrator. As a result, the translation of an OTA-C all-pass filter circuit, shown in Figure 5, has the same issues.

[4]Some offset current is present at the SSI2164's output, but that only comes into play at a very low corner frequency and some current can also be sourced from, or sinked into it, if the op-amp is too slow to follow the input signal and significant voltage becomes present at the virtual ground summing node.

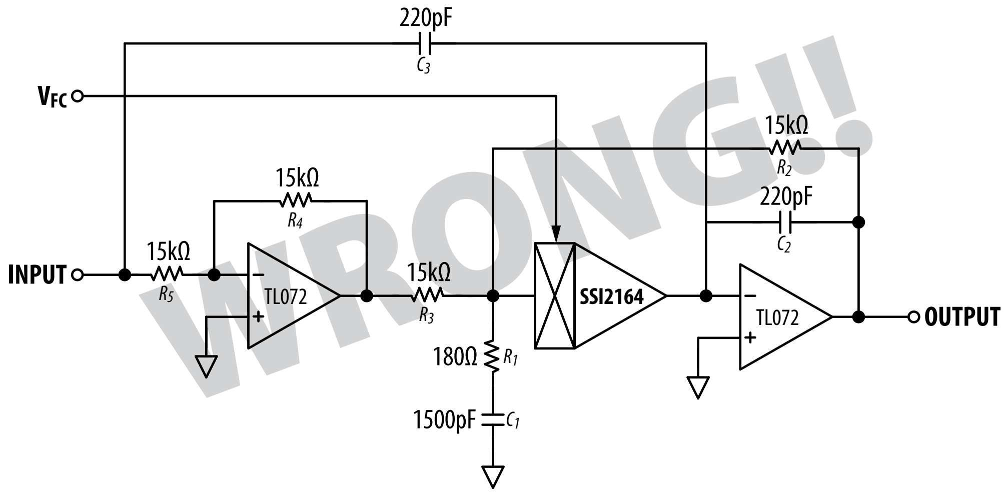

Figure 5: Non-Working All-Pass Filter Stage Using SSI2164

The low-pass circuit of Figure 1 is not affected by this problem because the input signal is already filtered enough to prevent excessive over- and undershoot.

A secondary effect may also come into play as the input signal will not be provided directly, but instead buffered by an op-amp to isolate unwanted extra input impedances that would otherwise affect the filter's response. With these circuits the input is purely capacitive and most op-amps are unstable when loaded with a significant capacitive load.

Sallen-Key Low-Pass

The same happens in the translation of the OTA-C Sallen-Key low-pass filter shown in Figure 6. The feedback op-amp that sets this filter's Q is capacitively loaded by C_5. Next to that, the same issue of saturating to one of the supplies will happen at high Q with low control voltage.

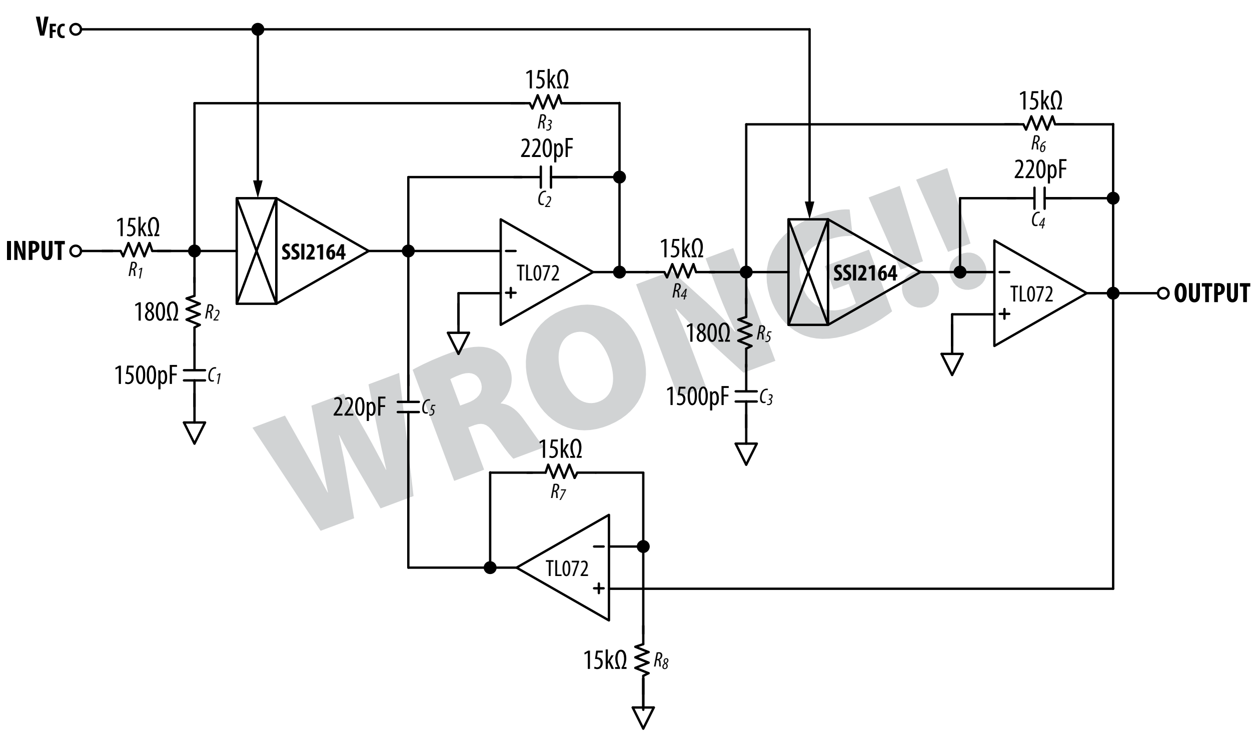

Figure 6: Non-Working Sallen-Key Filter Using SSI2164

SSI2164-OPTIMIZED BUILDING BLOCKS

"***The***" is also emphasized in the first sentence of the previous section because one might be tempted to 'cure' the symptoms of those problematic circuits by, for example, adding discharge resistors to drain unwanted DC charges or adding resistors to isolate these capacitive loads a bit. While this can be done, the filter's response is changed with some of its range sacrificed and extra poles introduced. These will likely wreak havoc when building higher order-filters with feedback. It is better to start anew and design these circuits from scratch instead of translating the OTA-C circuits.

High-Pass

Simply put, a high-pass filter is the output of a low-pass filter subtracted from its input. Figure 7 shows a single-pole high-pass filter stage.

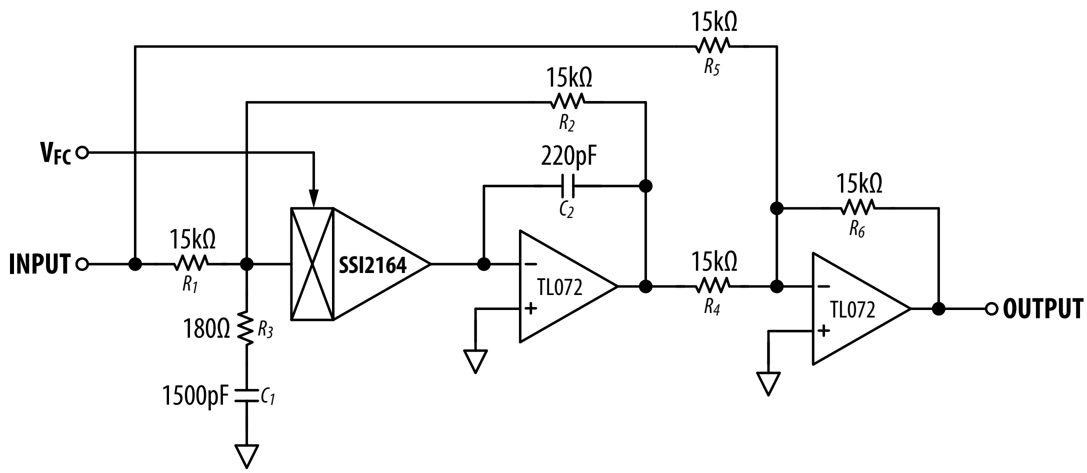

Figure 7: Single-Pole High-Pass Filter Stage

This is basically the LPF circuit of Figure 1, which is inverting, with an extra summing amplifier. It is also inverting and the equation for corner frequency remains valid. Its transfer function is:

$$-\frac{S}{S+1}$$

With bode plots:

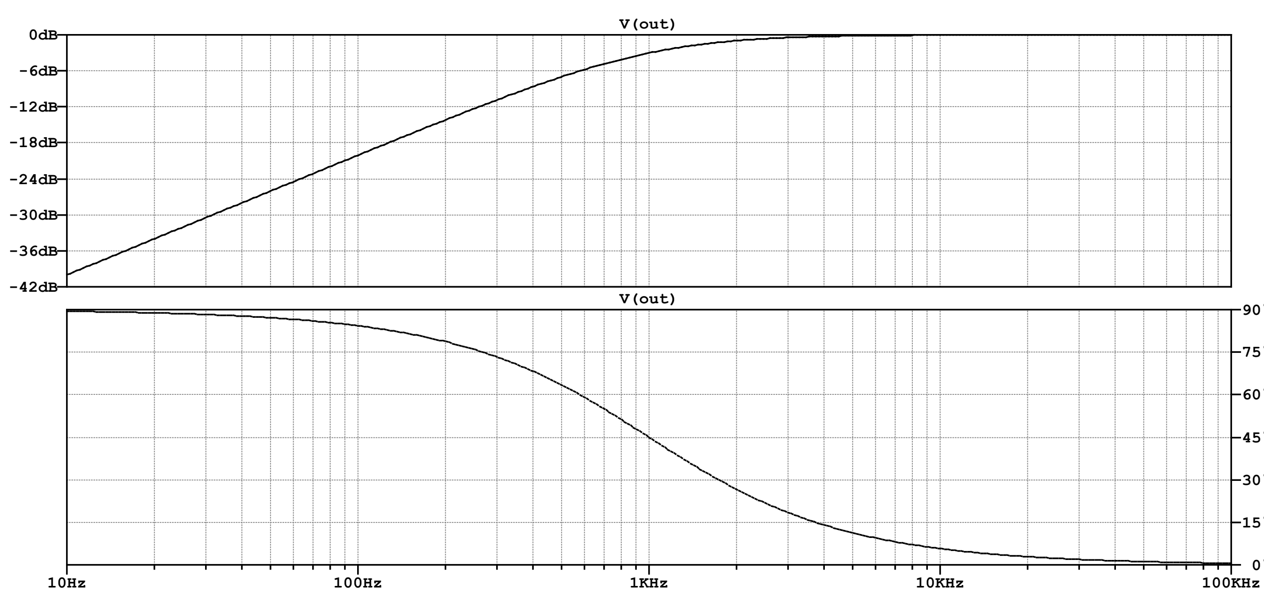

Figure 8: Bode Plots of the Single-Pole High-Pass Filter Stage (shown non-inverting)

All-Pass

An all-pass filter, better known as phase-shifter, is the output of a high-pass filter subtracted from a low-pass one. Figure 9 shows a single-pole all-pass stage. It's exactly the same as the HPF circuit of figure 6, except R_4 is halved, so that the low-pass output is added to the inverted high-pass output. This is the same as subtracting the high-pass from the low-pass. This one is non-inverting.

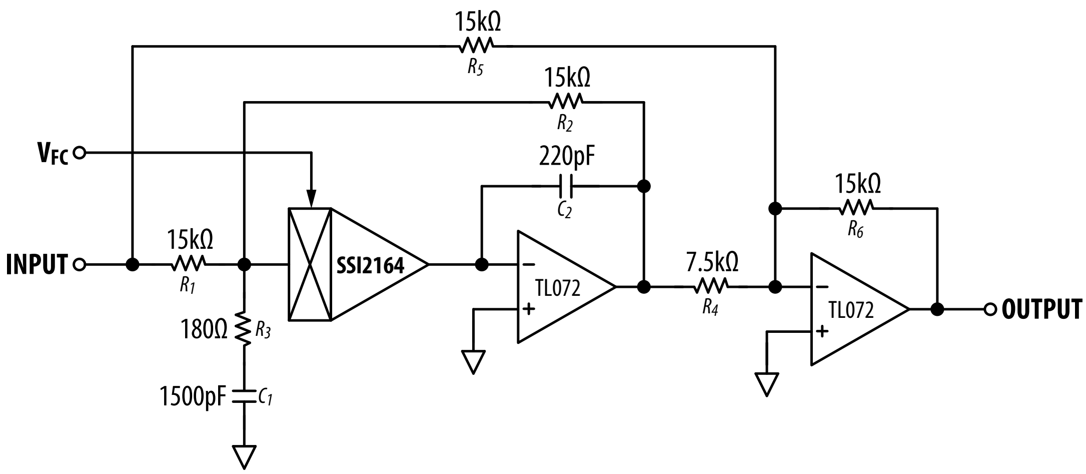

Figure 9: Single-Pole All-Pass Filter Stage

Its transfer function is:

$$\frac{1-S}{S+1}$$

With bode plots:

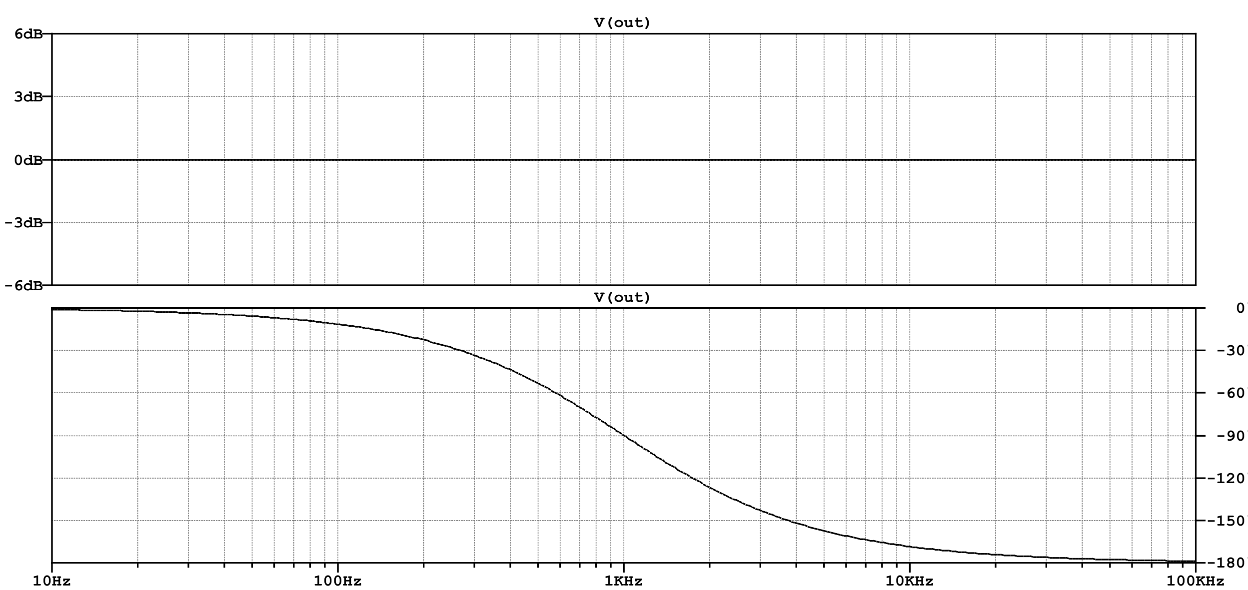

Figure 10: Bode Plots of the Single-Pole All-Pass Filter Stage

Multi-Mode Filter

The circuit of Figure 7 (and Figure 8) can easily be turned in a multi-mode filter by adding 2 switches. With both switches open, the circuit of Figure 11 is a non-inverting low-pass filter. Closing S_1 makes it an inverting high-pass filter and also closing S_2 turns it into a non-inverting all-pass filter. FET-switches or digital potentiometers can also be used to make this programmable.

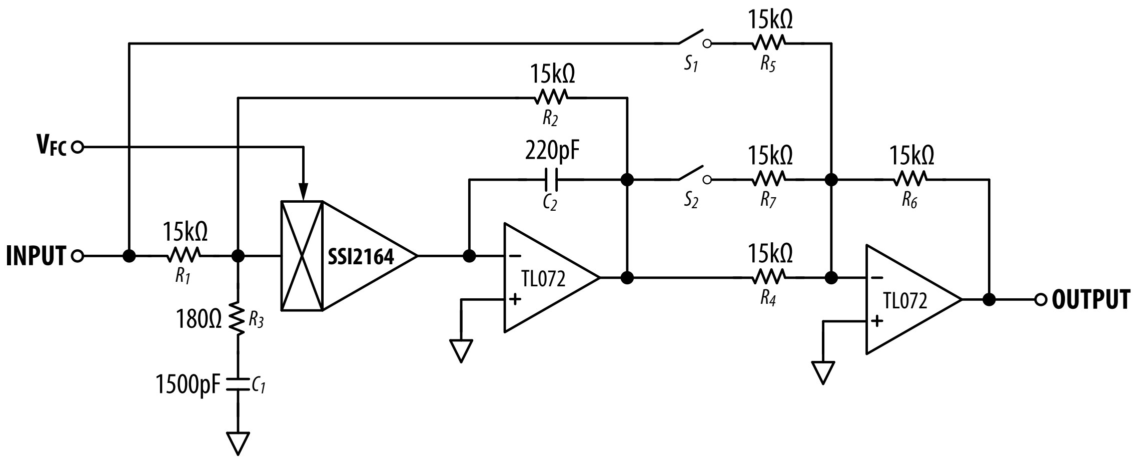

Figure 11: Multi-Mode Filter Stage

All the SSI2164-optimized circuits in this section are based on the same principle as the multi-mode filters described later and require precision resistors. Consider it wise to use a resistor array with tight ratio tolerances.

Unlike OTA-C-based high-pass and all-pass filters, the described SSI2164 filters don't suffer from unwanted attention resulting from parasitic capacitances of the OTA's output, op-amp input, and connecting PCB traces that form a capacitive divider with the filter's timing capacitor. The lack of exact specification (OTA output capacitance), voltage dependency (for the input capacitance of JFET op-amps) and complex calculation (PCB trace capacitance) in combination with the limited availability of precision capacitors makes such attenuation hard to compensate for. SSI2164 filter gain errors depends almost exclusively on resistor tolerance.

STATE-VARIABLE FILTERS

State-variable filters were used in early ARP synthesizers[5] and later popularized as the Oberheim SEM. The main selling points are the simultaneous availability of low-pass, high-pass and band-pass outputs, and the independence of Q and frequency adjustments.

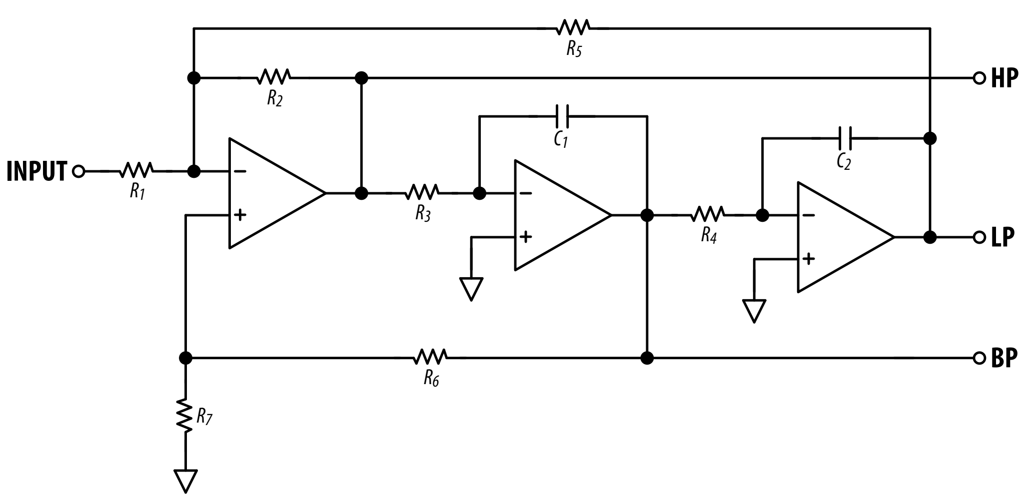

Figure 12: Kerwin-Huelsman-Newcomb Filter

5 'Electrical Design and Musical Applications of an Unconditionally Stable Combination Voltage Controlled Filter/Resonator' by Dennis P. Colin, Tonus, Inc. (ARP), Newton Highlands, Mass. Presented April 29, 1971 at the 40th Convention of the Audio Engineering Society, Los Angeles.

Implementing these with the SSI2164 is very straightforward. Take the textbook Kerwin-Huelsman-Newcomb filter[6] from Figure 12 and insert SSI2164 VCAs in-between the resistors and the integrators. Figure 13 shows the schematic. A potentiometer was added to provide adjustable resonance and also the mixer circuit from ARP's 1047 Mutimode Filter/Resonator and Oberheim SEM that allow continuous fading from low-pass to high-pass response with a notch filter in the center position.

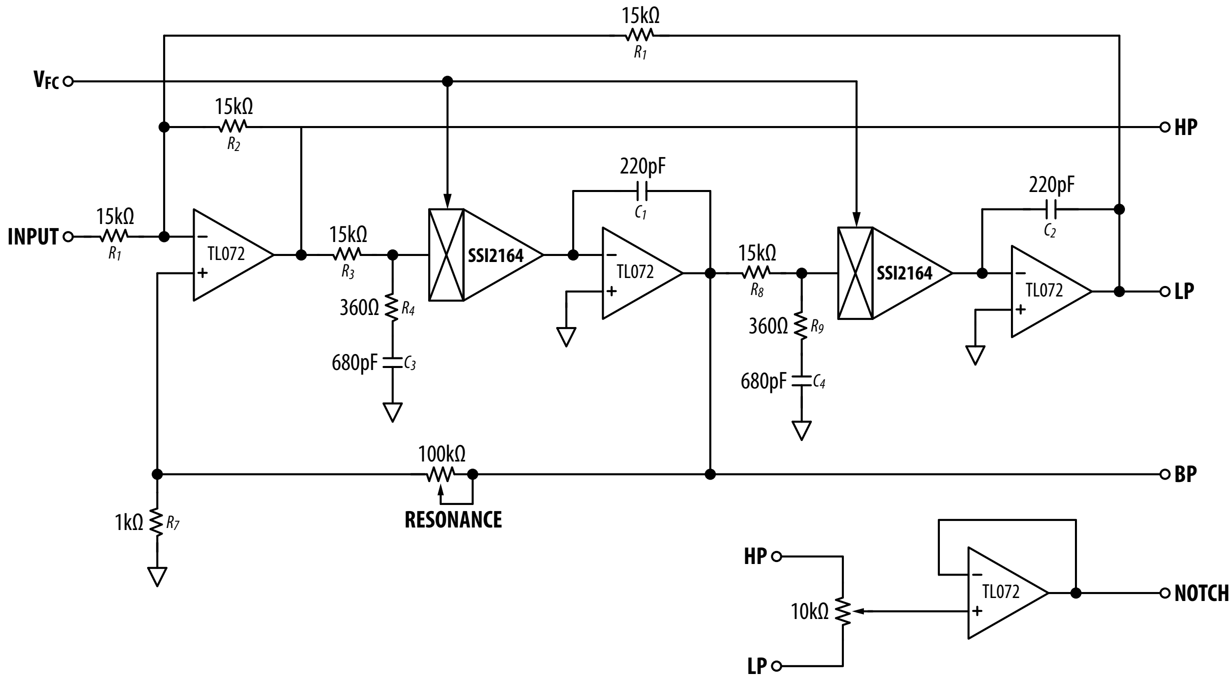

Figure 13: SSI2164-Based State-Variable Filter

The Bode plots of each output are shown in Figure 14 without resonance and with the notch potentiometer in the center position.

By increasing the amount of resonance, these filters are able to peak at the corner frequency. However, by the lack of sufficient phase shift in the resonance feedback loop, they are incapable of self-oscillating – see later section "Resonance." The Bode plot of the high-pass output with resonance is shown in Figure 15.

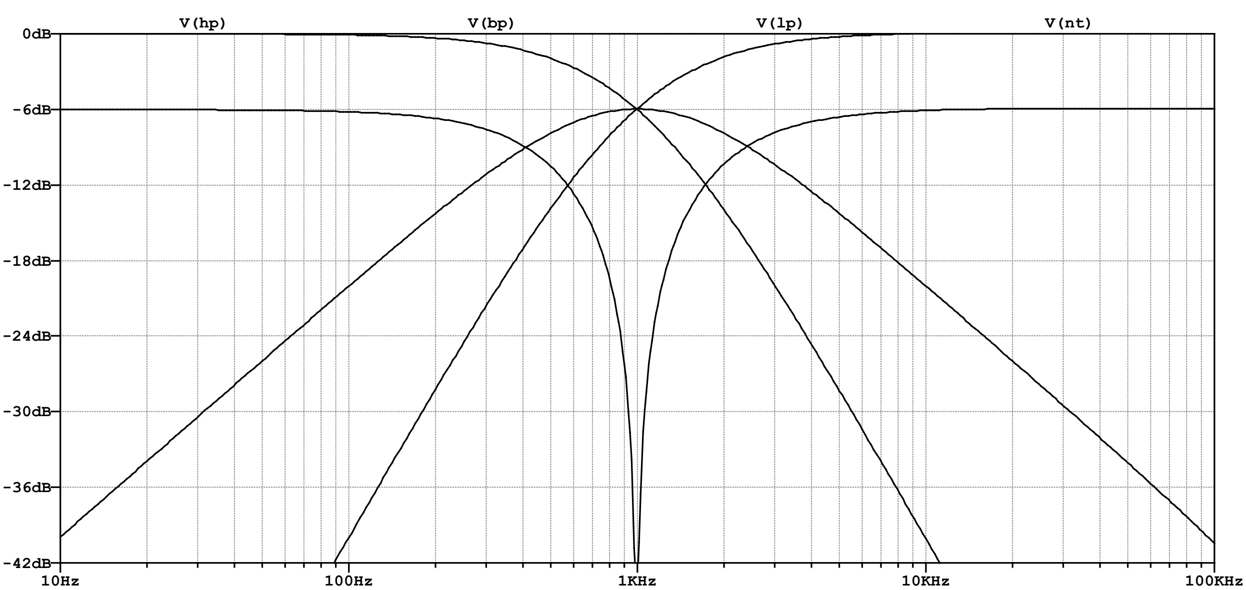

Figure 14: Magnitude Bode Plots of State-Variable Filter Outputs

[6] W. Kerwin, L. Huelsman and R. Newcomb, State variable synthesis for insensitive integrated circuit transfer functions, IEEE J. Solid State Circuit 2 (1967) p87–92

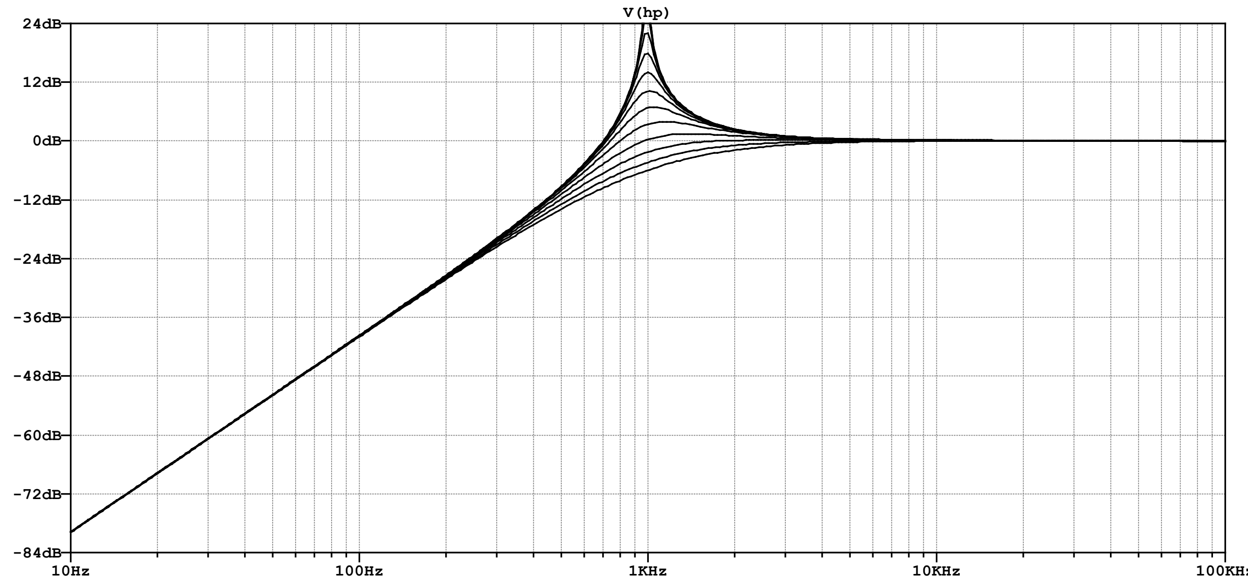

Figure 15: Magnitude Bode Plots of State-Variable Filter High-Pass Output with Increasing Resonance

HIGHER ORDER FILTERS

First-order, single-pole filter stages provide a roll-off of 6dB per octave (or 20dB per decade) with 45° (90° for the APF stage) phase shift at the corner frequency. Higher order and steeper filters can be made by cascading several stages in series. Band-pass filters can be made by cascading high- and low-pass filters. For example, the cascade of two each low-pass and high-pass stages[7] creates a 4th order band-pass filter with a 12dB per octave roll-off on each side. See Figures 16 and 17.

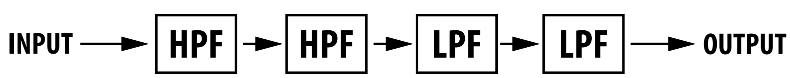

Figure 16: 4th Order 12dB per Octave Band-Pass Filter

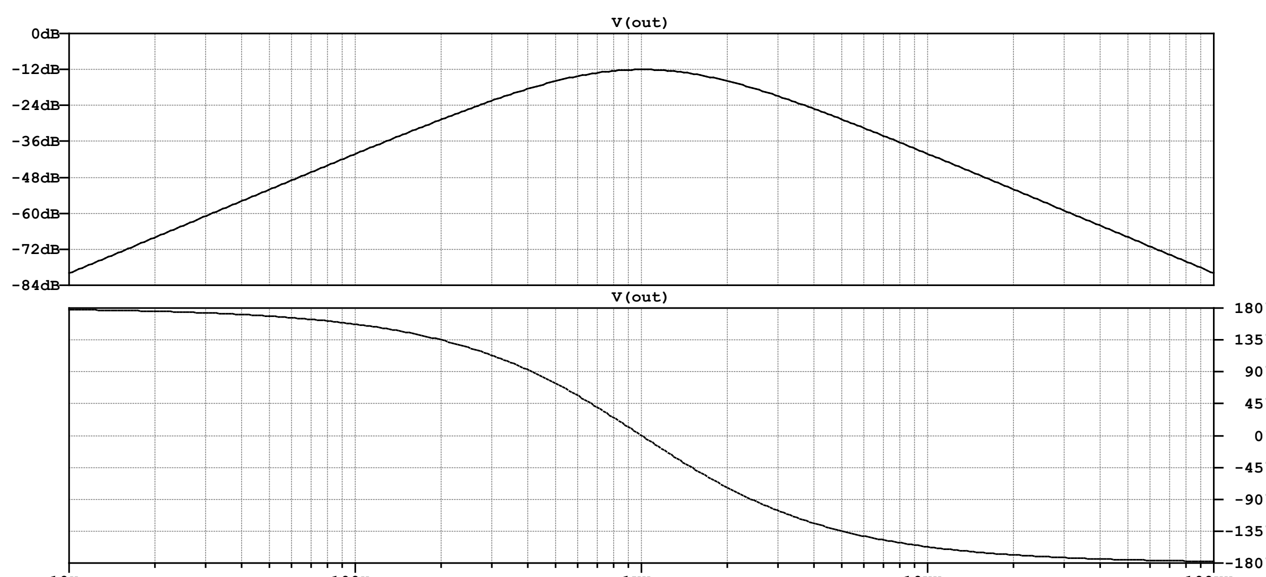

Figure 17: Bode Plots of 4th Order Octave Band-Pass Filter

[7] When only considering ideal filters, the order in which the low-pass and high-pass stages appear makes no difference. Here, in real-life, it makes sense to put the low-pass stages at the end to filter out the high-frequency content of noise generated by each filter stage, as high-frequency noise is perceived to be louder than low-frequency noise.

There's no need to keep things symmetrical – differing quantities of LPF and HPF stage cascading will make one side of the spectrum roll-off steeper than the the other.

RESONANCE

The behavior of a filter can be significantly changed by feeding a certain amount of a filter stage's output into the input of a preceeding stage.

If the total[8] phase-shift in this feedback loop is an integer multiple of 360°, at one or more non-zero[9] frequencies, and the total[8] loop-gain at these frequencies[10] is greater than one, then the filter will self-oscillate at these frequencies.

For example, Figure 17's Bode plot of the filter from Figure 16 shows the phase-shift is 0° (an integer multiple of 360°) at 1 kHz and the attenuation at that frequency is -12dB. Feeding back its output to its input, as shown in Figure 18, will make it self-oscillating if the feedback gain is more than 12dB.

This can be seen in its Bode plot (Figure 19) where the feedback gain is varied from 0 to a little over 12dB.

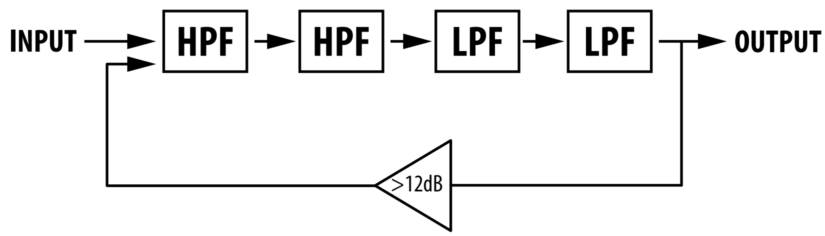

Figure 18: Resonant 4th Order Band-Pass Filter

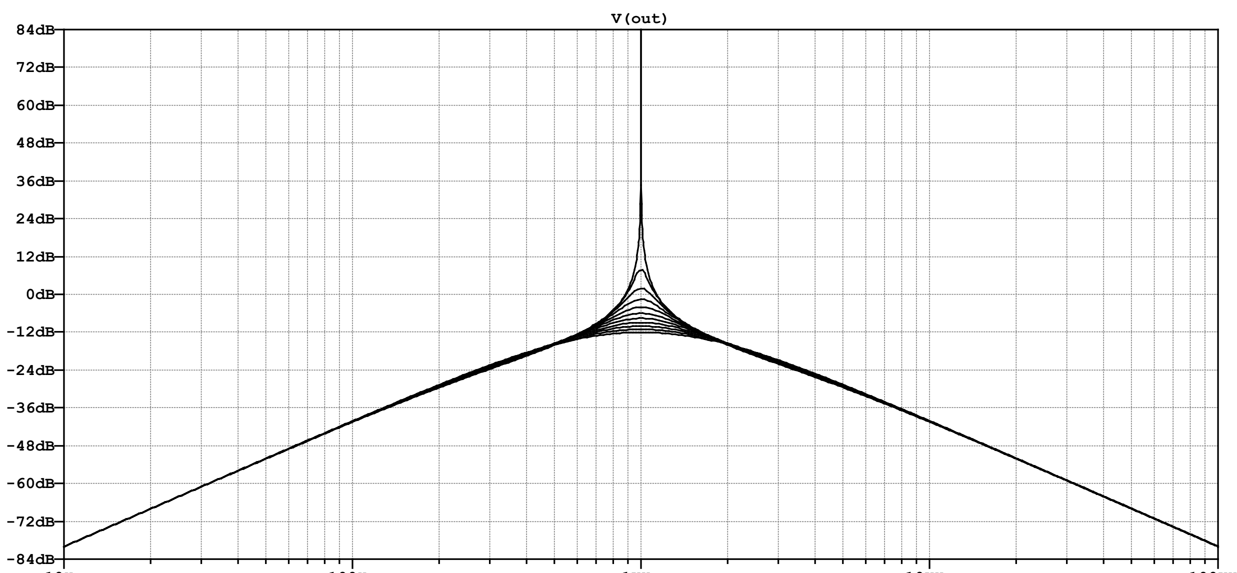

Figure 19: Magnitude Bode Plots of a Resonant 4th Order Band-Pass Filter with Increasing Feedback Gain

THE CLASSIC 24dB PER OCTAVE LOW-PASS VCF

The most used VCF in synthesizers is the cascade of four low-pass filter sections. First seen in Bob Moog's ladder filter[11] in the late 60's and popularized as a cascade of OTA-C filters in the iconic polyphonic analog synths of the early 80's.

---

8 This also includes the feedback-buffer/inverter/amplifier/attenuator/filter.

9 Beware of excessive DC-feedback!

10 Limit the feedback gain beyond a certain frequency with high-pass filters or they will self-oscillate in the RF range due to phase shifts that occur because of the non-infinite speed and bandwidth of the amplifiers used.

11 US Patent # 3.475.623: 'Electronic High-pass and Low-pass Filters Employing the Base to Emitter Diode Resistance of Bipolar Transistors' by Robert A. Moog from October 10th 1966

Its popularity comes from its steep 24dB per octave slope, the amount of resonance feedback not affecting the adjusted corner frequency and its capability of self-oscillation, at that corner frequency, when resonance is increased. This can be seen in its bode plot (Figure 21).

Figure 20 shows its basic structure. In the Bode plots of Figure 2 can be seen that each low-pass stage contributes 45° phase shift at the corner frequency and that the attenuation is there -3dB. The four stages add up to 180° and -12dB. To achieve self-oscillation greater than12dB feedback gain is needed and it has to be inverting. This inversion adds the needed extra 180° phase shift to come to the required integer multiple of 360° for self-oscillation.

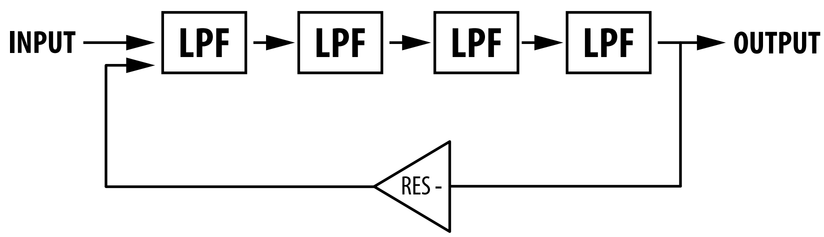

Figure 20: Resonant 24dB per Octave Four-Pole Low-Pass Filter

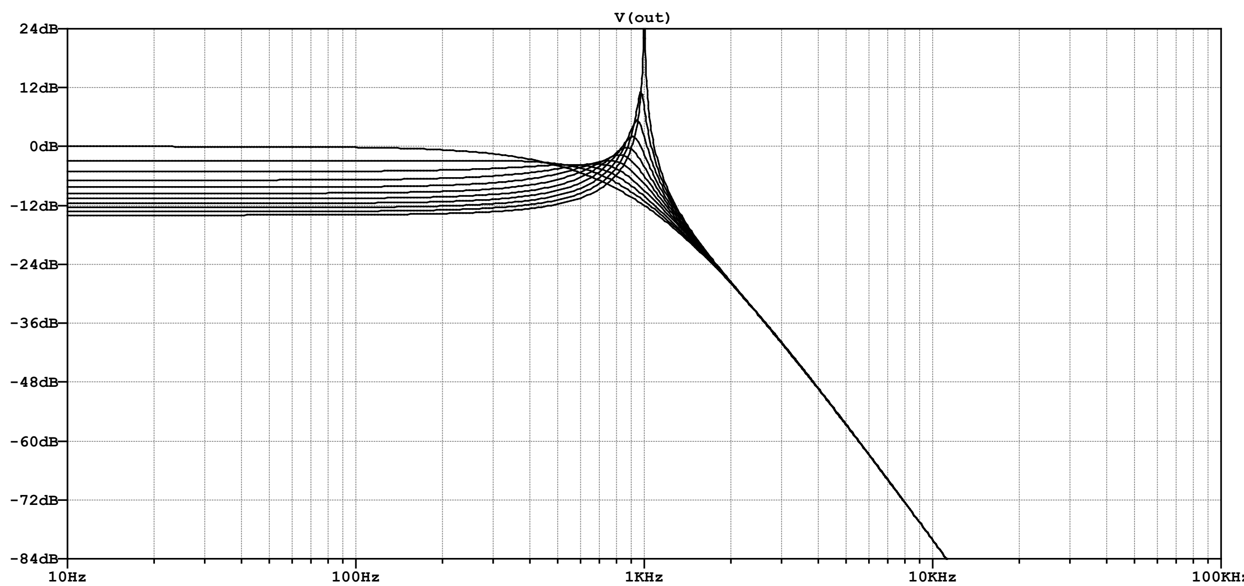

Figure 21: Magnitude Bode Plots of 24dB per Octave Four-Pole Low-Pass Filter with Increasing Resonance

A schematic of such a VCF, built around a single SSI2164 IC, is shown in Figure 22. The cascade on top is four times the circuit of Figure 1 in series, each using one of the four VCAs in a SSI2164 IC. The lower op-amp performs the inversion[12], attenuated by the RES potentiometer and R15 in series with R17 provide the feedback with a gain (x5) of 14dB. 2dB gain headroom is provided as compensation for potential losses and tolerances.

In the bottom-left corner is a variable clipping circuit, courtesy of the RES_LIM potentiometer. Its sole purpose[13] is to set the amplitude at which this VCF will self-oscillate, by limiting the feedback gain beyond a certain voltage level. Without it, this would be set by the maximum output swing of the integrating op-amp of the first low-pass stage, leaving no room for the input signal itself. This function is implemented in many synths by (strings[14] of) diodes in anti-parallel, either to ground or in the feedback loop of an op-amp overdrive circuit. Or, it is inherent to the differential input pair of ladder-filters, discrete amplifiers and OTAs.

[12] This inverter deliberately has a maximum gain of one to prevent it from causing extra phase-shift, due to gain coming at the cost of bandwidth.
[13] Contrary to popular belief it does not 'shape' the self-oscillation wave-form. That function is performed by the very strong frequency-selectivity of the filter.
[14] To increase the forward voltage while simultaneously decreasing their parasitic capacitance. Zeners are not a good idea because of high capacitance and noise

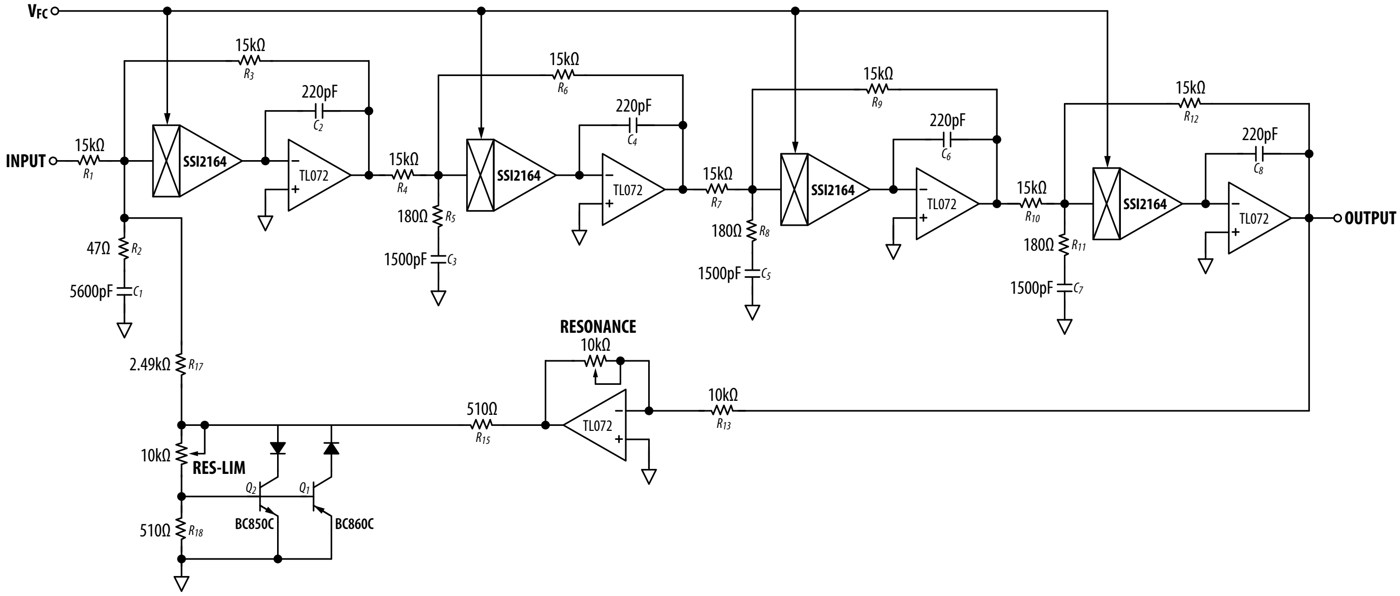

Figure 22: Schematic of 24dB per Octave Four-Pole Lowpass Filter

Due to excellent linearity of the SSI2164, it is necessary to accommodate a self-oscillation amplitude tailored to the range of signals that may be encountered in modular synthesizers. These can vary from sub $1V_{P-P}$ coming from consumer equipment to $20V_{P-P}$ when connected to vintage E-mu modules.

DIMENSIONING AND PARTS SELECTION

According to the datasheet, the input resistance should be chosen to limit the maximum input current to 1mA.

However, some extra sensibility is needed here because that 1mA is multiplied by 10 (+20dB gain) at the output when -660mV is present at the control pins. This means that, when the above VCF is presented with this control voltage and a square wave at maximum input voltage, each filter section will consume 150mW (if run from ±15V rails) extra. Multiply that by 4 (for each stage) and the 118°C/W thermal resistance of the package and you come to a 70°C temperature increase of the SSI2164's die. This not only brings it close to the maximum junction temperature. But, it also shifts the frequency of this VCF audibly. External temperature compensation schemes fail to catch this internal temperature change.

It is OK to dimension the input resistor like this, if the control voltage is limited. All above circuits are dimensioned for positive-only control voltages and a maximum input amplitude of 15V, which should cover everything encountered between ±15V rails

If you do want to benefit from this extra gain at negative control voltages, multiply the input resistors accordingly. Note that this comes at the cost of increased Johnson noise.

Dimensioning the RC compensation network is done according to data sheet instructions. Note that the input sees all connected resistances in parallel and the equivalent total R_{IN} needs to be used for calculations.

The maximum corner frequency needs to be set high enough to minimize the filter's effect when completely open and the SSI2164's output offset current needs to be taken into account at the low end. When dimensioned to 50 kHz at zero control voltage, the VCF covers the whole audible range.

Use low tolerance; low drift capacitors for the integrators. C0G ceramics do the trick here.

Use FET-input op-amps for the integrators. Bipolar ones have too low input impedance and a too high input bias current. The TL07x series will do just fine.

PERFORMANCE IMPROVEMENTS

DC

A DC-offset, in the range a few 100µv to a couple of mV, is to be expected at the output. Compared to signal amplitudes of several volts this is negligible for most purposes. If for any reason it needs to be attenuated, note that it has both static and dynamic components.

The static component can be cancelled out completely by adding a small, trimpot adjustable, DC offset voltage to the input signal or by inserting an equivalent offset current directly in the first section. It is advisable to generate this offset from a clean and filtered reference voltage to avoid pollution from junk on the main supply lines. Alternatively, it can be added after the filter but then you lose the benefit of the extra noise filtering that comes for free in the stop-band of the filter.

The dynamic component cannot be cancelled out by adding a trimpot due to its non-linear relationship with the control voltage. Fortunately, most comes as an output offset current of the SSI2164 that gets cancelled by negative feedback associated with the closed loop nature of each low-pass stage. However, this self-compensation mechanism vanishes at very low cutoff frequencies (high control voltage) because the feedback gain then diminishes. This increasing offset at high control voltages can be prevented by limiting the control voltage to a sensible range and or attenuated by adding high-ohmic discharge resistors in parallel with the timing capacitors. Note this RC combo then sets the low frequency limit of the filter.

Control Feedthrough

Some control feedthrough is present and cannot be cancelled by straightforward means due to its non-linear behavior.

Noise

In the pass-band noise is entirely dominated by the SSI2164, whose noise contribution depends on the resistors at its inputs. The only way to lower it is by decreasing their value and that means either lowering (due to 1mA input current limit of the SSI264) the maximum signal amplitude, which will have no effect on the overall SNR, or running multiple SSI2164 IC's in parallel to distribute the input current. Take into account the latter is less straight forward than it sounds because separate SSI2164's will have to be used which don't have the benefit of being matched.

In the stop-band only unity-gain noise of the last op-amp remains. With a TL07x this noise floor sits at about -103dBu (20kHz unweighted) and taking into the account the over 24Vp-p peak voltage swing of the circuit, when sourced with ±15V, allows for a dynamic range of over 120dB.

Distortion

This circuit will significantly outperform your distortion expectations if they are only based on the SSI2164's THD specs. Two mechanisms are at work that improves things. The first comes inherent with filtering: harmonics in the stop-band are progressively filtered out. The second works in the pass-band: enormous amounts of negative feedback gain of the op-amps irons out all remaining non-linearities.

Class A?

You might be tempted to run the SSI2164 in class A-mode by injecting current in its mode pin. However, contrary to VCAs, with VCFs there is no benefit in doing so. The above mentioned mechanisms do a better job at reducing distortion and increased bias current in the SSI2164s cores will only increase noise, DC-offset, and control feedthrough. In fact, you can sink a few 10µAs from the mode pin to bring the operation closer class-B and improve these things a bit. However, gains are minimal here and this current will have to be trimmed manually because sinking too much will cause unwanted oscillation.

Resonance Stability at Higher Frequencies

If only ideal components were considered, the filter's resonance or Q would remain the same for all frequencies and only depend on the feedback gain. See Figure 23.

However, in the real world, the limited speed of the op-amps and VCA sections introduce delays and extra phase-shifts to the filter, and have as most noticeable effect, which is known as Q-enhancement[15], that the filter's resonance or Q increases at higher frequencies.

This effect, shown in Figure 24, will most likely be of little concern when the self-oscillation amplitude is hard-limited, like in the circuit of Figure 22 or in monophonic and modular synths with sufficient headroom. In polyphonic synths, where all the voices' waveforms and filter self-oscillations have eventually to be summed together, it can become an issue when the self-oscillation amplitude increases beyond its assigned slice of the available headroom and causes the summing amp to clip.

It can be adequately compensated for by adding a few pF capacitors in parallel with R4, R7, R10 and R13 in the circuit of Figure 22. These parallel RC circuits then serve as lead compensators. The exact value (in the range of 1-10pF) of these compensation capacitors has to be determined experimentally as PCB and passive component parasitics also come into play here.

This compensation method is known as Q-compensation, but the term is also used for something entirely different now.

[15]Explained in detail in Electronotes Volume 8, Number 71, November 1976, p14-21 and Electronotes Volume 14, Number 141, September 1982, p3-18

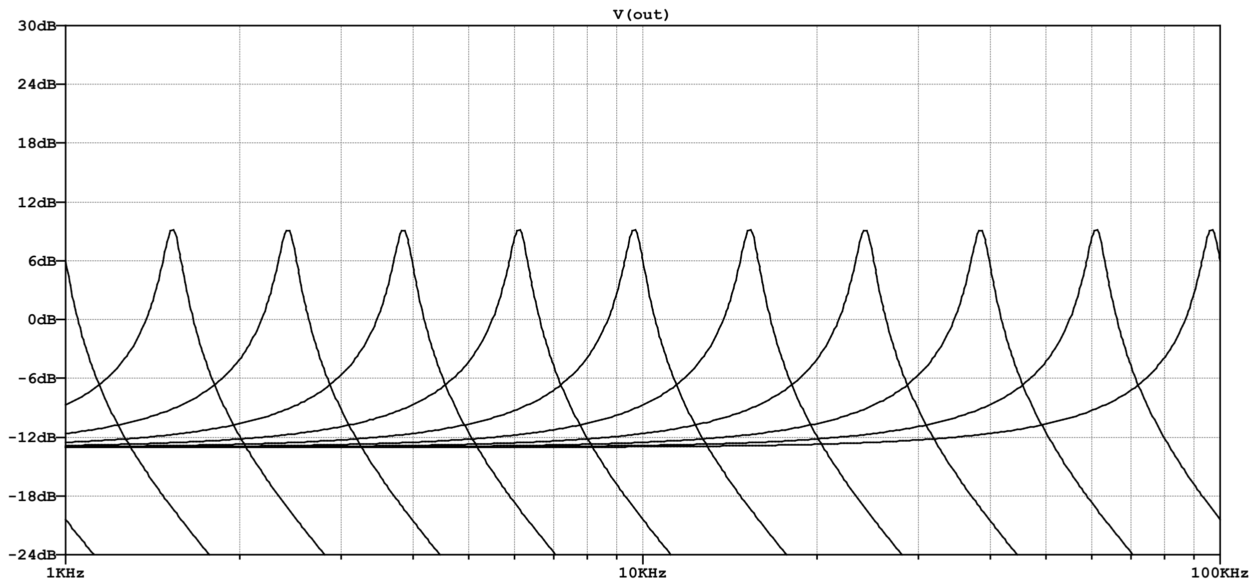

Figure 23: Magnitude Bode Plots Showing Constant Q at all Frequencies with Ideal Components

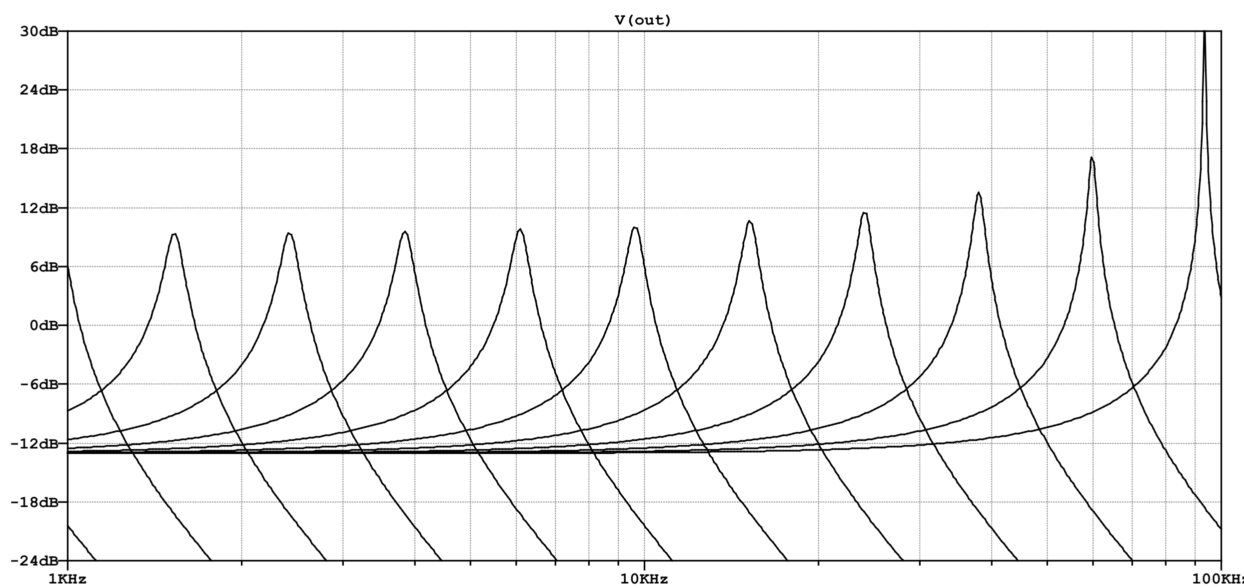

Figure 24: Magnitude Bode Plots of Q Enhancements

Q COMPENSATION

In the Bode plot of Figure 25, it can be seen that the frequency content in the passband gets attenuated when resonance is increased. This is the kind of resonance behavior as in the Minimoog, Prophet 5, Polysix, RSF Kobol and so on.

This attenuation of low-frequency content with increased resonance was perceived by some as a problem. A sort of Band-Aid solution to solve this problem is now known as Q-compensation[16]. It has been implemented in a multitude of ways which all come down to the exact same thing: Amplification that tracks the amount of resonance feedback so that pass-band level is restored. This is typical Roland poly sound.

The most popular and elegant implementation[17] is adding the filter's input signal to the input of the resonance amplifier as shown in Figure 26.

[16] This term was coined by Tony Allgood of Oakley Sound Systems around the turn of the millennium and seems to have stuck with contemporary designers. At the time of writing no prior names or documents that describe this method were found. The first implementation found so far is in the 1975 Roland System-100's VCF. But, no description can be found in its literature. As a matter of fact its instruction manual even contradicts its existence by stating that frequencies below cutoff are attenuated and even accompanies this by 5 Bode plots that gravely exaggerate this attenuation. Please contact the author if you know of earlier implementations or documents that describe it.

[17] Used in Jupiter-8, Juno-6, 60 and 106,. . .

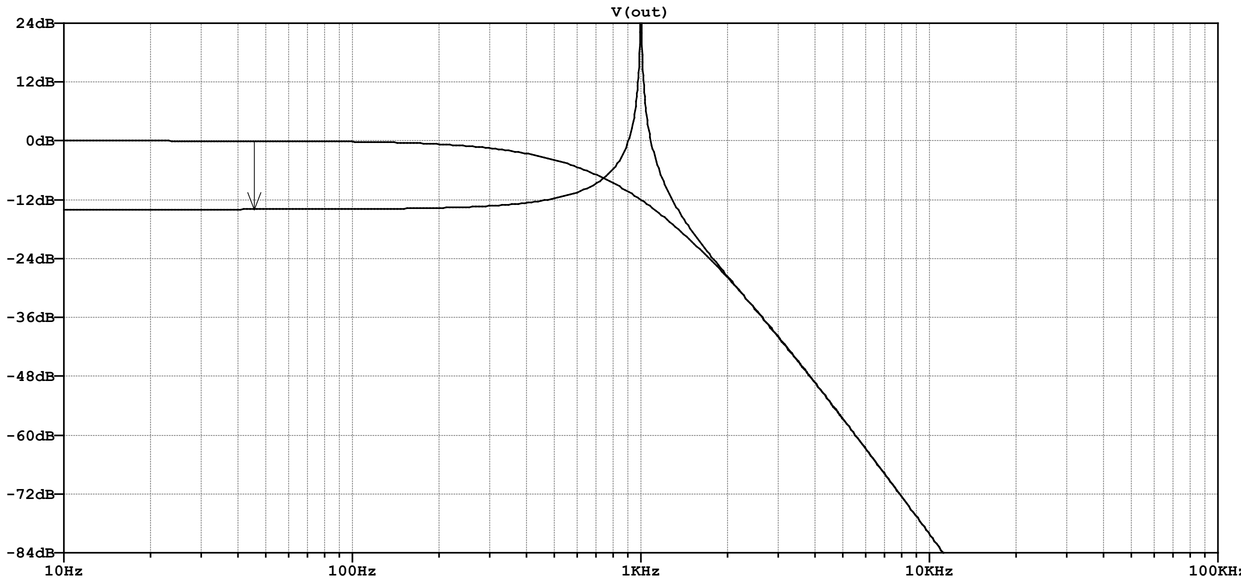

Figure 25: Passband Attenuation When Resonance Gain Increases in Four-Pole Filters

Other methods are, for example, increasing the input signal amplitude[18] with a dual-gang resonance potentiometer or summing the output of the resonance amplifier[19] with the filter's output in the VCA that follows it.

A complete schematic is shown in Figure 27. It is exactly the same as the one in Figure 22. Only the inverter is extra. Note that this can be done somewhat more elegant if an OTA is used as (voltage controlled) resonance amplifier, since both inverting and non-inverting inputs are available. This extra inverter is not needed then.

However, this Band-Aid solution is not without its issues, as all this amplification is not beneficial for signal quality and it quite literally shifts the problem, as the pass-band attenuation is traded in for a shift in the filter's cutoff slope, shown in Figure 28.

This shift is, at the onset of self-oscillation, a frequency multiplication of 3/2 or an interval of a perfect fifth or 7 semitones. As both standard and Q-compensated 4-pole low-pass filters are problematic, a novel approach is needed.

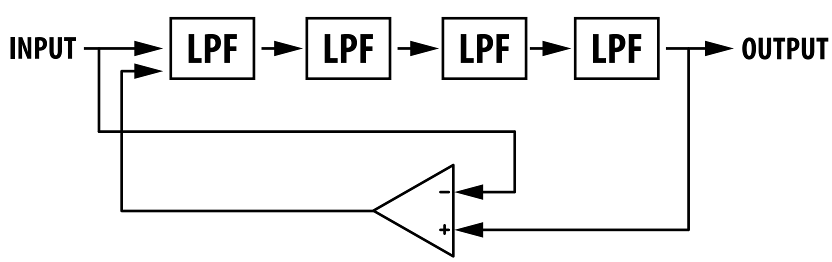

Figure 26: Resonant 24dB per Octave Four-Pole Filter with Q Compensation

[18]Used in the System-100 model 101+102 and SH-5.
[19]Used in the SH-101.

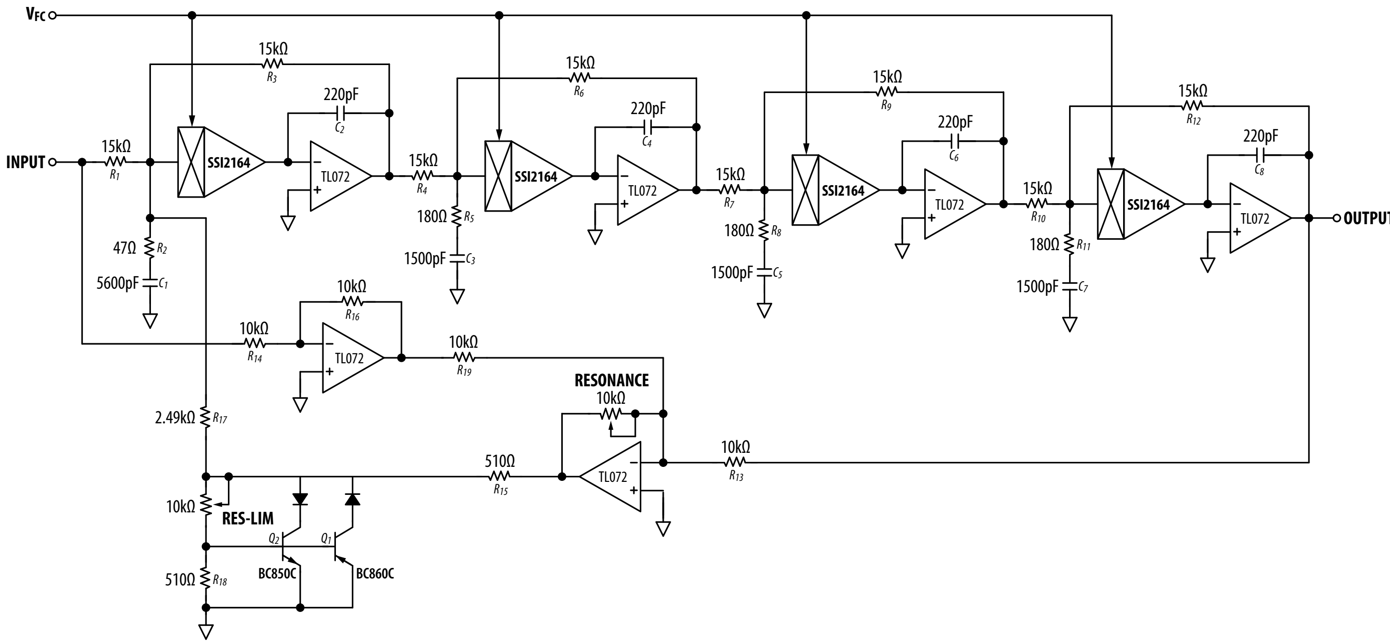

Figure 27: Schematic of a Resonant 24dB per Octave 4-Pole Low-Pass Filter with Q Compensation

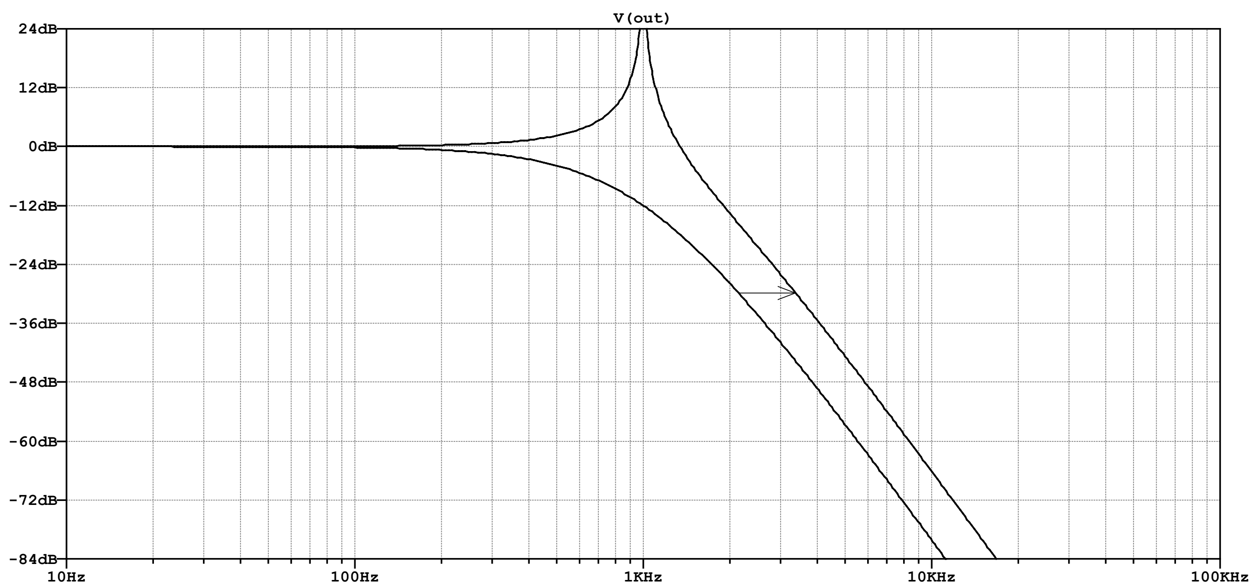

Figure 28: Cutoff Frequency Shift when Resonance Increases in Q-Compensated 4-Pole Low-Pass Filters

SOLVING THE PASSBAND ATTENUATION PROBLEM

To solve this problem, it's root cause has to be analyzed. Observing the Bode plot of the resonance feedback in Figure 29 shows this immediately: The root cause is: inverted and un-attenuated pass-band signal, which cancels out the input signal, when it is summed with it at the filter's input.

The problem can thus be solved by only feeding back signal around the corner frequency, where it is needed to provide resonance peaking. Or in other words: band-pass feedback[20].

In the previous section on resonance, it was already shown that a 4th order band-pass filter has the necessary requirements to peak at the corner frequency. The complete filter structure of a 4-pole low-pass filter with this resonance feedback is shown in Figure 30.

[20]The author had a chat with Joran van Gremberghe and it became apparent he independently, came up with the same solution at about the same time.

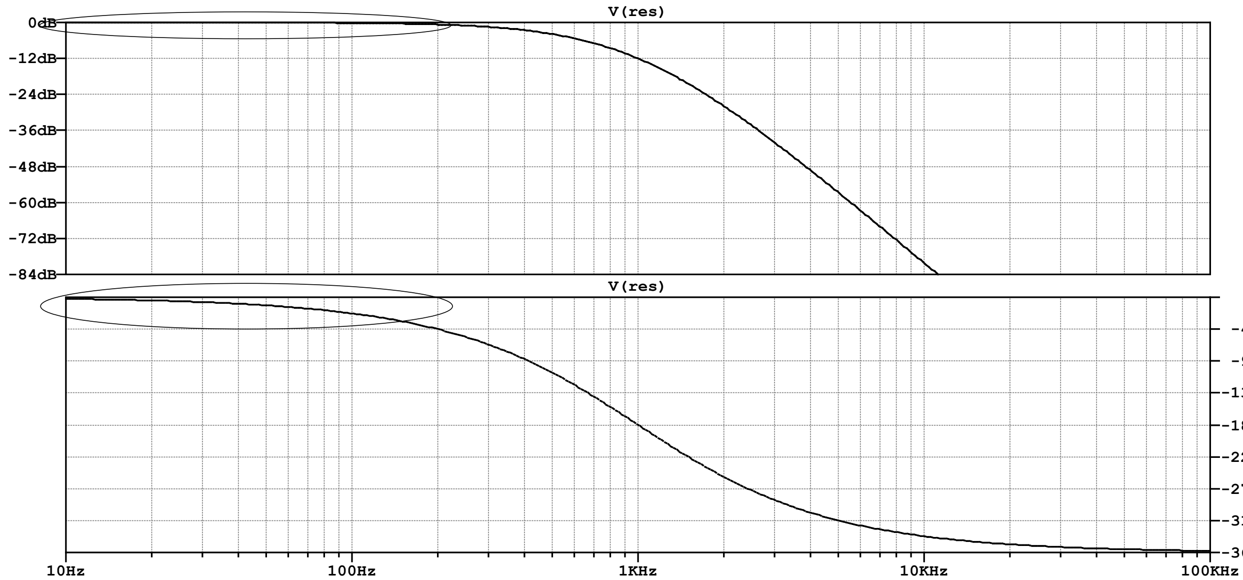

Figure 29: Bode Plots of a 4-Pole Low-Pass Filter Resonance Feedback

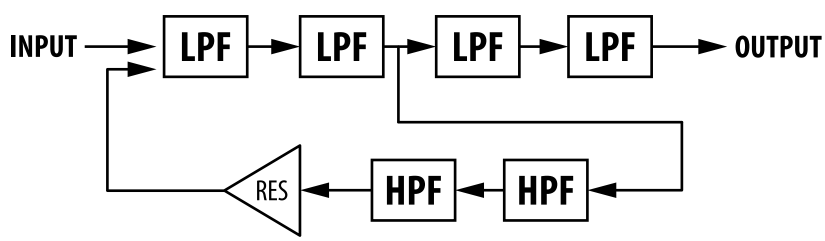

Figure 30: Resonant 24dB per Octave 4-Pole Low-Pass Filter with Bandpass Feedback

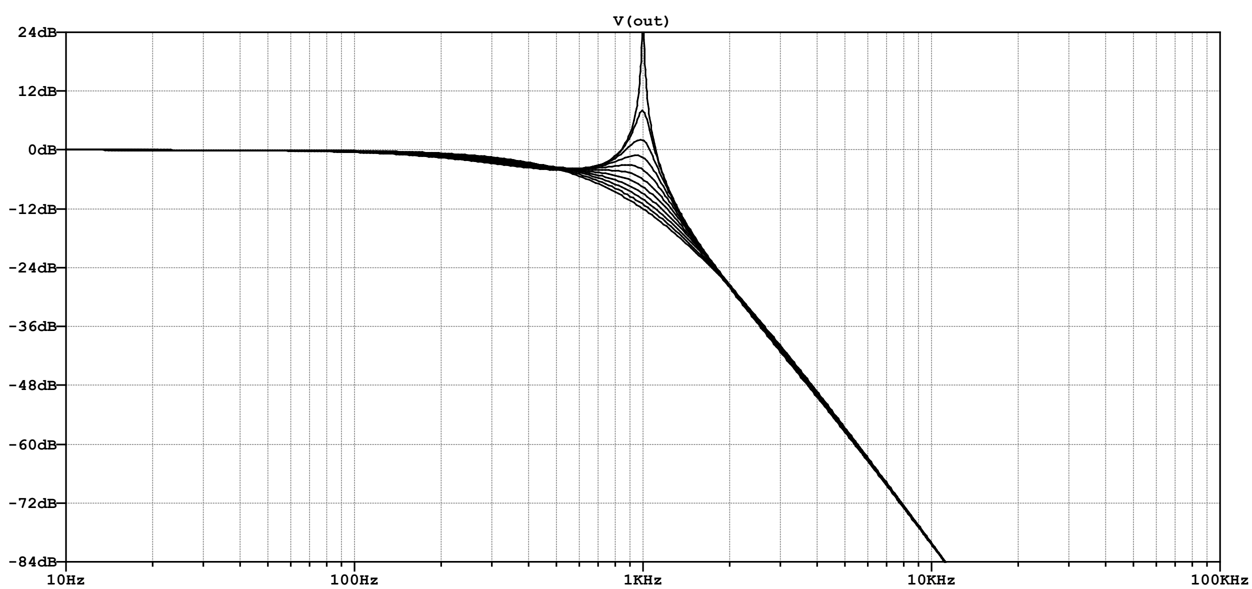

Figure 31: Bode Plots of 4-Pole Low-Pass Filter with Increasing Amounts of 4th Order Bandpass Resonance

This solves the problem completely by preventing it from happening and it is without side-effects. There is no pass-band attenuation and no shift of the cutoff slope, as shown in the Bode plot of Figure 31.

The structure of this filter might seem more complex and to require two extra gain cells (or half of another SSI2164). However, the required 4th order band-pass response can be generated as weighted sum of the outputs of the 2nd, 3rd and 4th low-pass stages. More details about this in the section on Multi-Mode Filters that follows.

The complete schematic is shown in Figure 32 and comes down to adding two extra resistors and an inverter to the schematic of the original uncompensated VCF of Figure 22. The same remark as for the Q-compensated filter of Figure 27 applies here: The inverter is not needed here if an OTA is used as resonance amplifier.

Alternatively, 2nd order band-pass can be used for resonance. In that case it comes down to adding only one extra resistor to the schematic of Figure 22. Its schematic is shown in Figure 33. Note that the resonance gain has been decreased because there is only 6dB attenuation instead of 12dB at the cutoff frequency here. This was done by increasing R17. The compensation network at the input of the first SSI2164 section was also adjusted for this higher input resistance.

This filter has mixed characteristics. The wide resonant peak associated with a 2-pole Sallen-Key[21] filter, combined with the steep cutoff slope of a 4-pole filter, is shown in its Bode plot of Figure 34.

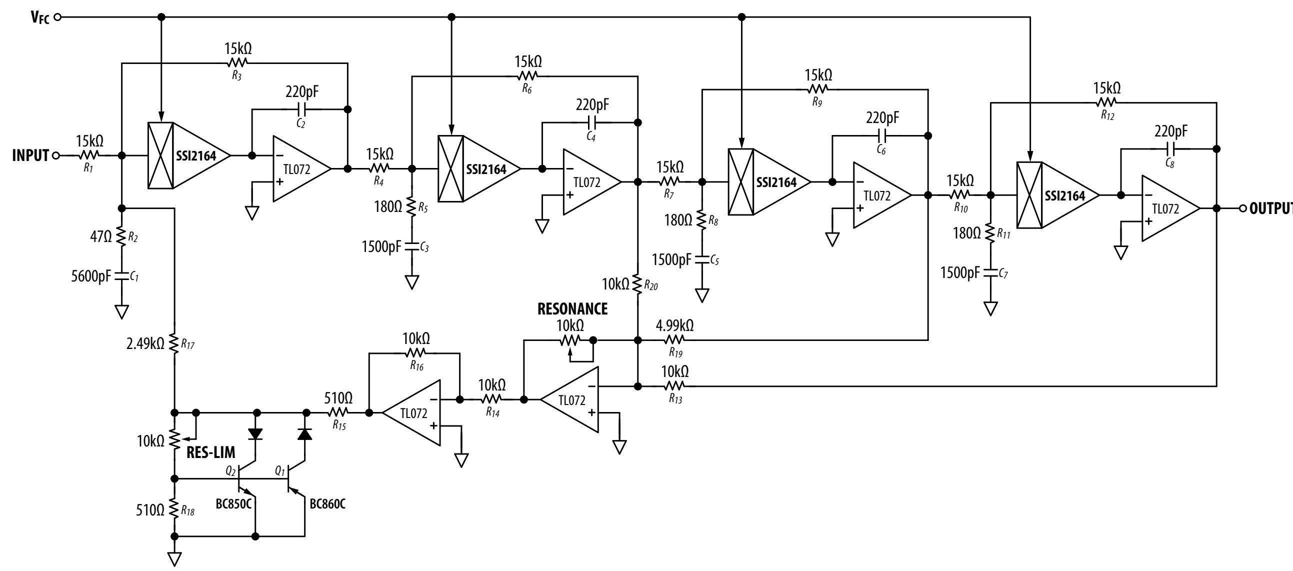

Figure 32: Schematic of a Resonant 24dB per Octave 4-Pole Low-Pass Filter with 4th Order Bandpass Resonance

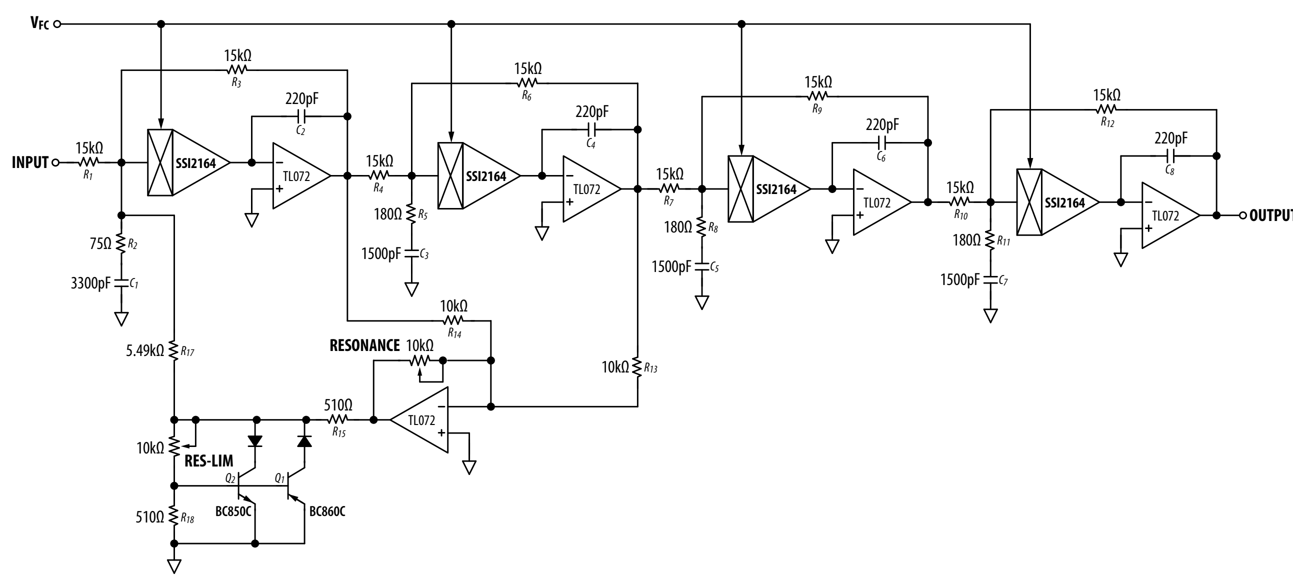

Figure 33: Schematic of a Resonant 24dB per Octave 4-Pole Low-Pass Filter with 2nd Order Bandpass Resonance

[21]The output of the 2nd low-pass stage is in fact a Sallen-Key filter.

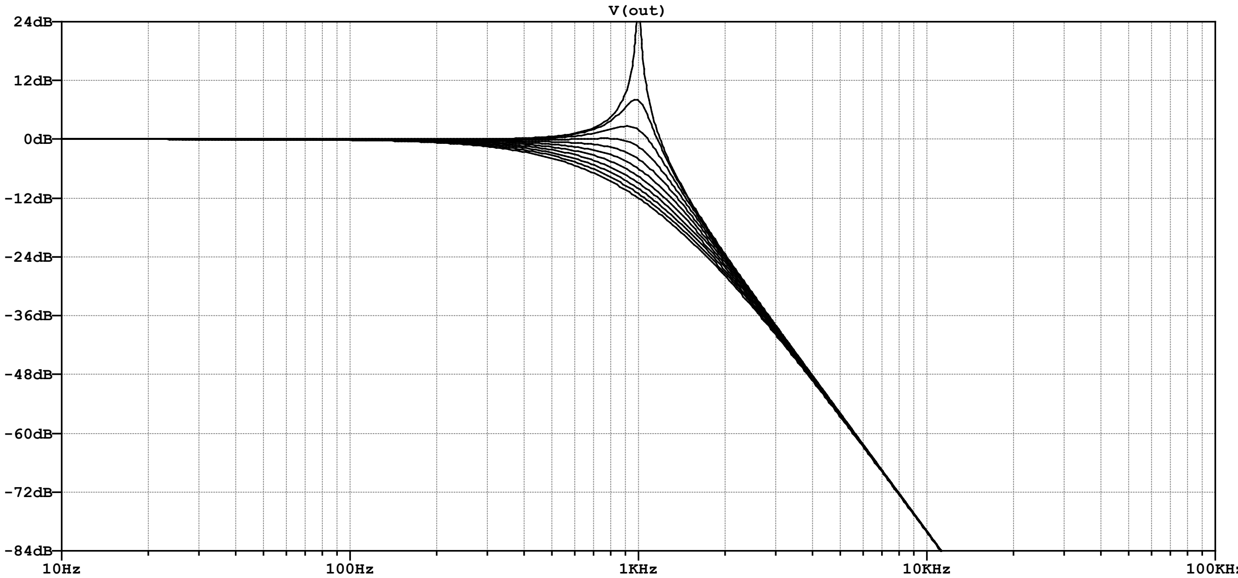

Figure 34: Bode Plots of 4-Pole Low-Pass Filter with Increasing Amounts of 2nd Order Bandpass Resonance Gain

MULTI-MODE FILTERS

Such filters were first implemented in the Siel Mono[22], popularized by the multi-mode VCF in the Oberheim Xpander and Matrix-12, and originating in Electronotes[23].

These are based on taking the weighted sum of the input signal and the outputs of each of the four 1st order low-pass stages as shown in Figure 35.

The resulting filter has as transfer function:

$$\frac{AS^4+(4A+B)S^3+(6A+3B+C)S^2+(4A+3B+2C+D)S+A+B+C+D+E}{(S+1)^4}$$

Any filter with transfer functions up to 4th order that can be represented with a polynomial in S as nominator and $(S+1)^n$ as denominator, with *n* equal to the filter's order, can be created with this method. More info can be found in relevant Electronotes articles and calculating the gains is just a matter of solving a system of five equations with the substitution method:

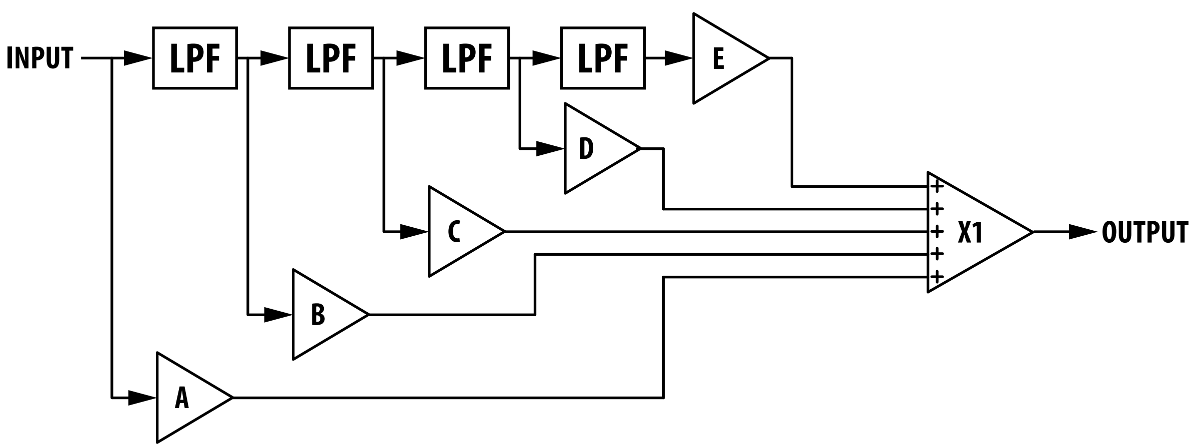

Figure 35: Multi-Mode Filter Based on Weighted Sum of Input and Each Lowpass Stage

[22] This is the earliest implementation found so far. Please contact the author if you know of an earlier one.

[23] Electronotes Application Note 71, January 18 1978 "Multi-Mode Filter based on first-order low-pass" and Electronotes Volume 10, Number 5, January 1978, p5-17 "Additional Design Ideas For Voltage-Controlled Filters"

Re-write the wanted filter's transfer function in the following form as polynomial over $(S+1)^4$

$$\frac{aS^4 + bS^3 + cS^2 + dS + e}{(S+1)^4}$$

And solve

$$AS^4 + (4A+B)S^3 + (6A+3B+C)S^2 + (4A+3B+2C+D)S + A+B+C+D+E = aS^4 + bS^3 + cS^2 + dS + e$$

With

$$\begin{aligned} A &= a \\ B &= b - 4A \\ C &= c - 6A - 3B \\ D &= d - 4A - 3B - 2C \\ E &= e - A - B - C - D \end{aligned}$$

The gains for all basic filters that are a cascade of single low-pass, high-pass, all-pass and notch filter stages with the same frequency can be found in Table 1. Note that the gains B and D, for the 1st and 3rd low-pass sections, have been inverted in the table because the used low-pass stages are inverting. The numbers[24] represent the actual gains to use here.

The SSI2164 is especially suitable for this task, because these kinds of filters stand or fall with the accuracy of each low-pass section and their mutual frequency tracking; the SSI2164 excels in this manner due its low distortion and excellent gain matching of its VCA sections.

The weighted summing can be done in many ways: static op-amp summing amps as in the Siel Mono and the circuits in Figure 32/33, CPU controlled CMOS switched as in the Xpander and Matrix 12, but VCA's or digital potentiometers are also options.

Schematic examples of op-amp summing with gains from Table 1 can be found for "2LP 2HP" in Figure 32, and "1LP 1HP" Figures 33 and 46. Note that the single high-pass and all-pass filter stages of Figure 7 and Figure 9, and their multi-mode combination in Figure 11 also fall in this category, albeit as 1st order examples.

Far more interesting things can be done with this method such as a 6dB per octave filter that's capable of self-oscillation based on a frequency shifted all-pass filter. See the last section "Impossible Filters" for further commentary.

A Few Warnings

Accuracy is the key here. Use precision resistors and low tolerance C0G capacitors. The reason for this is that the physical phenomenon on which this type of filter operates is interference of coherent waves. The resulting filter's curve is drawn by the complex interplay of constructive and destructive interference. Maximum achievable attenuation depends entirely on how well things are matched. For example, the worst case scenario (for which you should design) maximum attenuation of the 1st order high-pass filter, built with an op-amp summer (this comes down to the circuit in Figure 7) and 1%, 5% or 10% tolerance resistors is only -27.8dB, -13.4dB and -6.67dB respectively, when all other parts are considered to be ideal. This effect on the filter's slope is shown in Figure 36. Likewise, the depth of a single notch filter is reduced, from total attenuation, to only -46.1dB, -32.3dB and -26.4dB with timing capacitor tolerances of 1%, 5% and 10% respectively. Due to propagation of uncertainty, this sensitivity to component tolerances gets worse for higher order filters. Of interest is to note that the digital counterpart[25] of this type of filter, the FIR filter, can be realized in arbitrary high orders only because exact summation is possible in the digital domain.

The cost-saving trick used in the Xpander and Matrix 12 does not work properly here. This trick was disconnecting the timing capacitors of the first section with a CMOS switch, so that 8 additional modes were available with the same resistor network. This could be done because of the unique structure of the CEM3372's gain cells that require fairly large (33nF) timing capacitors and the few 10s to 100s of Pico farads capacitance of the CMOS switch are insignificant compared to them. The timing capacitors here (220pF) are much smaller and the extra capacitance from the switch will shift the frequency too much. Over 40dB more junk, from the digital signal that controls the switches, will also be injected through capacitive coupling from the gate. And, there is no need for it anymore. Those 1% resistors now cost less than 1/1000th of their price back then.

[24] These were calculated for the author's master thesis about 15 years ago and many, but not all, have been verified by building actual filters since. As a good practice, it's recommended that one simulate or generate Bode plots from the transfer function before circuit construction to avoid unexpected things.

[25] The author first considered these multi-mode filters as possible analog equivalents of FIR filters, while paying attention in DSP class, realizing that the Laplace transform is somewhat the continuous-time equivalent of the Z-transform and that transfer functions as polynomials in S over (S+1)^n look remarkably similar to those in Z over Z^n, only to discover that they already existed for about 3 decades.

Type	A	-B	C	-D	E
4LP	0	0	0	0	1
3LP 1HP	0	0	0	1	1
3LP 1AP	0	0	0	1	2
3LP	0	0	0	1	0
2LP 2HP	0	0	1	2	1
2LP 1HP 1AP	0	0	1	3	2
2LP 1HP	0	0	1	1	0
2LP 2AP	0	0	1	4	4
2LP 1AP	0	0	1	2	0
2LP	0	0	1	0	0
2LP 1NT	0	0	1	2	2
1LP 3HP	0	1	3	3	1
1LP 2HP 1AP	0	1	4	5	2
1LP 2HP	0	1	5	8	4
1LP 1HP 2AP	0	1	5	8	4
1LP 1HP 1AP	0	1	3	2	0
1LP 1HP	0	1	1	0	0
1LP 1HP 1NT	0	1	3	4	2
1LP 3AP	0	1	6	12	8
1LP 2AP	0	1	4	4	0
1LP 1AP	0	1	2	0	0
1LP 1AP 1NT	0	1	4	6	4
1LP	0	1	0	0	0
1LP 1NT	0	1	2	2	0
4HP	1	4	6	4	1
3HP 1AP	1	5	9	7	2
3HP	1	3	3	1	0
2HP 2AP	1	6	13	12	4
2NP 1AP	1	4	5	2	0
2HP	1	2	1	0	0
2HP 1NT	1	4	7	6	2
1HP 3AP	1	7	18	20	8
1HP 2AP	1	5	8	4	0
1HP 1AP	1	3	2	0	0
1HP 1AP 1NT	1	5	10	10	4
1HP	1	1	0	0	0
1HP 1NT	1	3	4	2	0
4AP	1	8	24	32	16
3AP	1	6	12	8	0
2AP	1	4	4	0	0
2AP 1NT	1	6	14	16	8
1AP	1	2	0	0	0
1AP 1NT	1	4	6	4	0
1NT	1	2	2	0	0
2NT	1	4	8	8	4

Table 1: Gains for Weighted Sum Filters

For the same reason, it is to be avoided to connect anything directly to the integrators inputs (= SSI2164 outputs). See Figure 37. These are very high impedance nodes and thus very sensitive. Any resistive load will ruin things at low frequencies and the extra 10pF load, if you were, for example, to connect another TL07x buffer, is enough for a 5% frequency shift. Only connect to the buffer outputs.

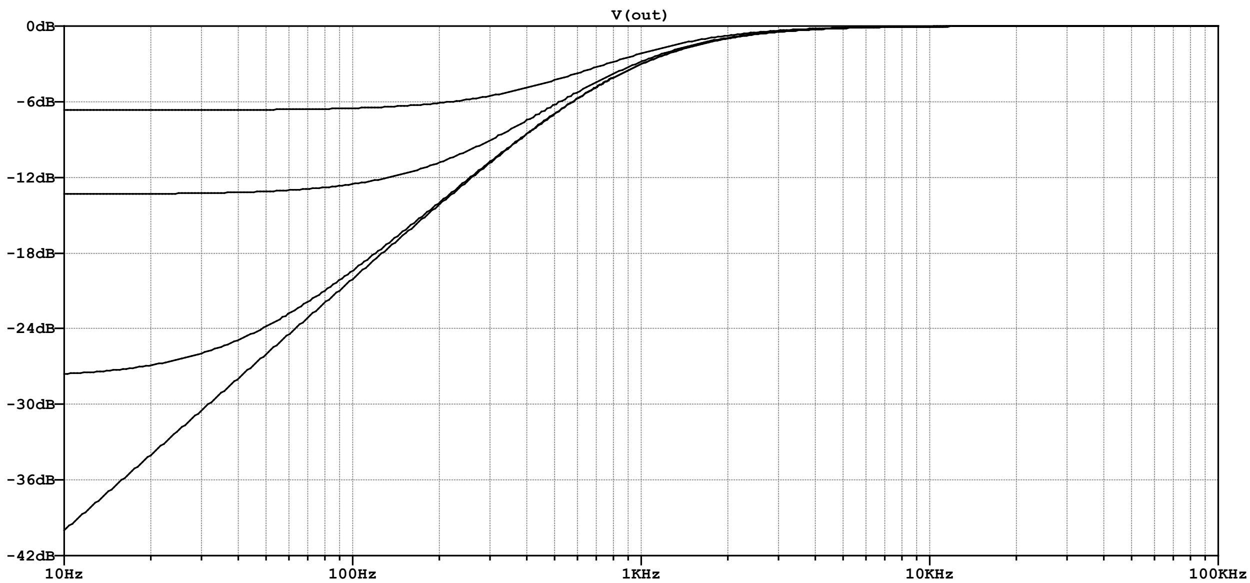

Figure 36: Resistor Tolerance Dependency of the 1st Order High-Pass Filter

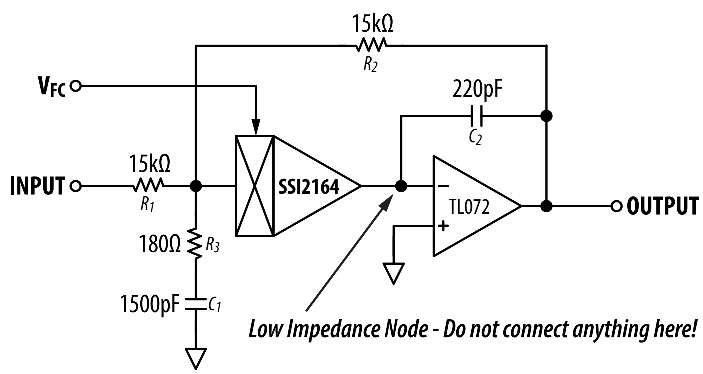

Figure 37: Low Impedance Node to Avoid

A FILTER WITH A BIT MORE CHARACTER - EMULATING THE SSM2040

The previous filter circuits all sound ultra-clean due to the already low distortion of the SSI2164, that gets improved even further by negative feedback and filtering (see 'Distortion' earlier).

What if you want one with a bit more character? A VCF's character is defined by its non-linearities and where they happen. A particularly interesting one in that regard is the SSM2040 filter, invented by Dave Rossum[26] and found in a few very rare synthesizers like the RSF Kobol and Prophet-5 revs 1 & 2. Some consider it the biggest sounding filter of all.

The main mechanism behind its sonic character is the incapability of its gain cells to sink current much below ground. This can be seen in Figure 38. Q1 and Q2 form a differential pair and with Q1's base tied to ground, the output pin, tied to Q2's collector can only go about 500mV[27] below ground before Q2 which sinks the output current, stops working. The Wilson mirror[28] on top that sources the output current allows

[26] US Patent 3969682, Circuit for dynamic control of phase shift, Dave Rossum, inventor.

[27] The buffers in the SSM2040 contribute a significant offset which allows the signal to swing about ±1V.

[28] This differs from the 2040 circuits circulating on the net. The author's thought behind it was, that it had to be there because a simple current mirror would cause significant offset current due to Base-Width Modulation. The Wilson mirror cancels out the BWM of Q4 and its inherent error cancels out the BWM of Q1. An improved 4-transistor Wilson mirror would be too perfect and leave the current error (and thus significant CV-bleedthrough) caused by BWM of Q1. This was later confirmed by Dave Rossum.

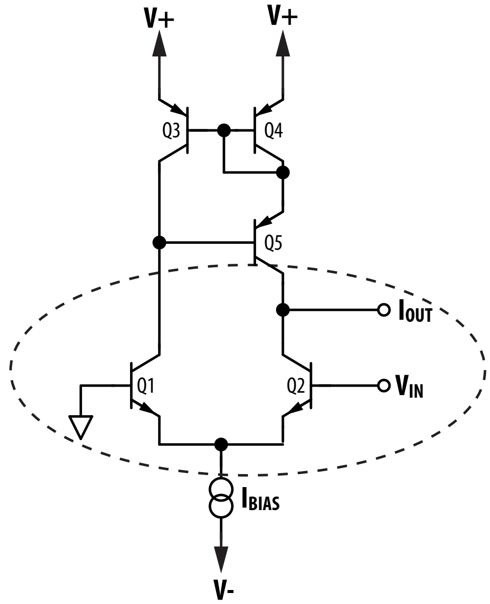

Figure 38: SSM2040 Gain Cell Structure

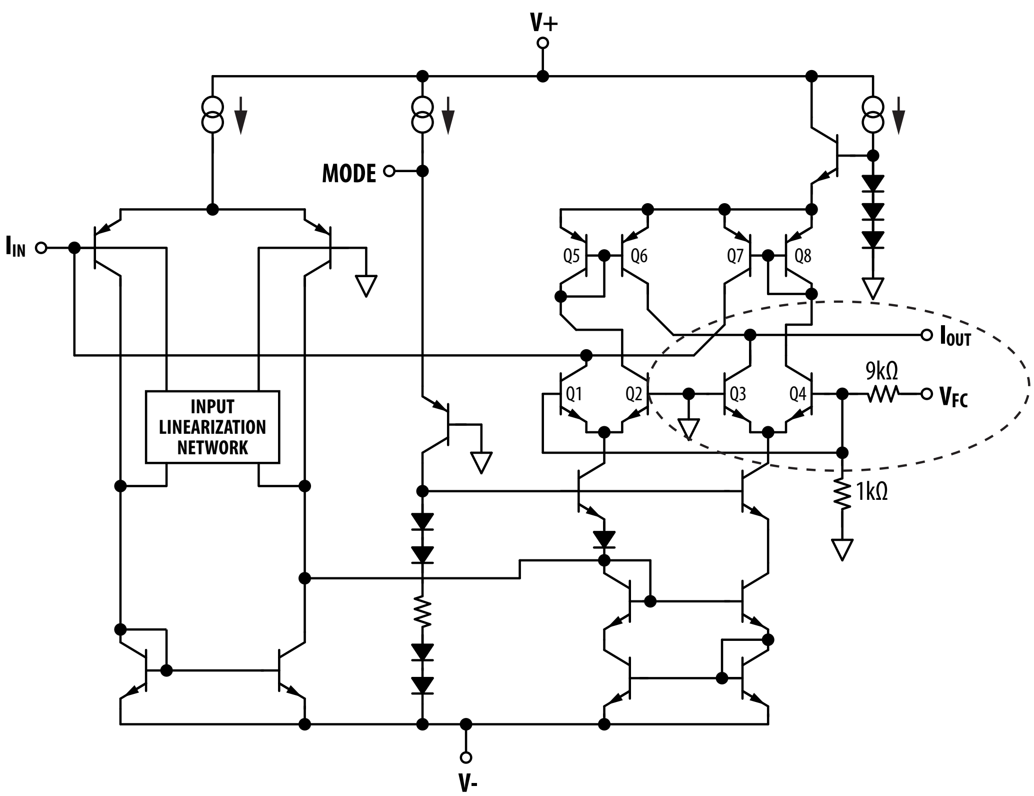

Figure 39: SSI2164 Simplified Schematic

almost full swing to the positive rail, so the overall distortion is asymmetric. Due to the inverting nature of each gain cell, this asymmetric distortion switches direction after each stage, while the input signal propagates through the SSM2040 and is meanwhile also progressively filtered.

The SSI2164 is also capable of doing this. In Figure 39 can be seen that the output current is also sinked by a transistor of a differential pair with one base tied to ground and sourced by a current mirror (not shown in this simplified schematic, but it is also a Wilson mirror) on top. An extra distortion mode is accessible here because the Wilson mirror is not connected to the positive rail; it is connected to a 1.8V bias. So, this one will stop working at about +1.2V. Here the clipping is much harder than at the bottom.

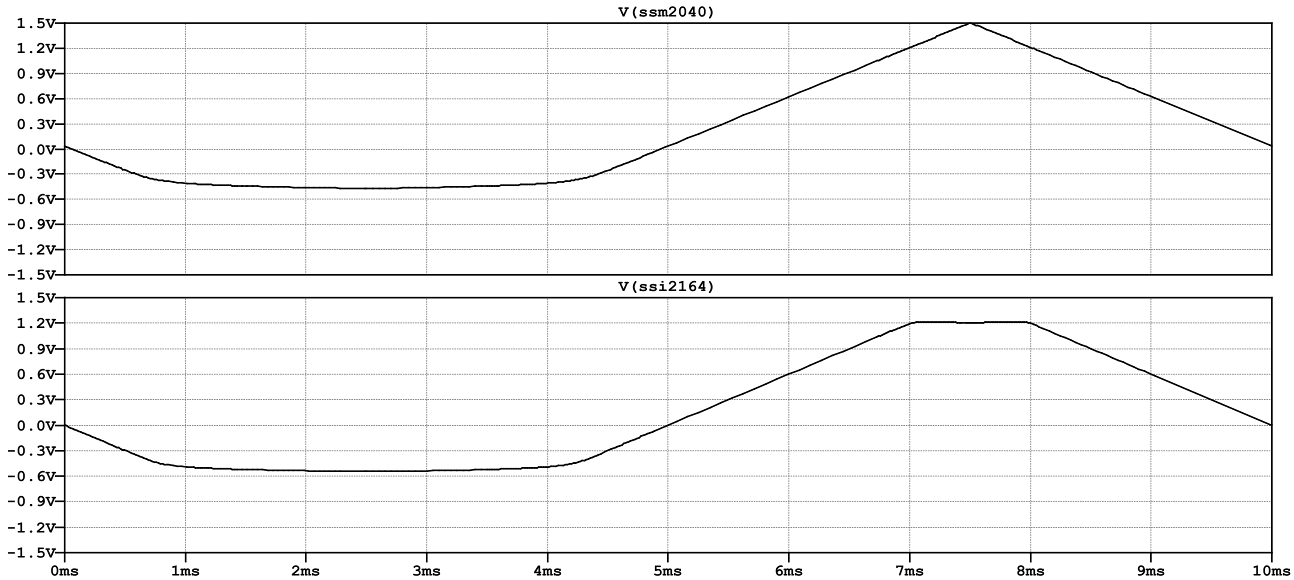

Figure 40: Asymmetric Distortion of SSM2040 and SSI2164

A plot that shows this asymmetric distortion of both the SSM2040 and SSI2164 is shown in Figure 40. The source waveform is a $3V_{P-P}$ triangle.

To actually do this we must ignore "*The output current pin should be maintained at virtual ground*" instruction from the datasheet, taking the ±100mV output compliance spec with a grain of salt and just let the output swing.

Figure 41 shows a single low-pass filter stage that provides this 2040-style asymmetric distortion. Function is the same as the circuit of Figure 1, except integrating is done by a buffered, ground-connected capacitor instead of the op-amp integrator and an extra inverter is needed because the output at the buffer is non-inverting.

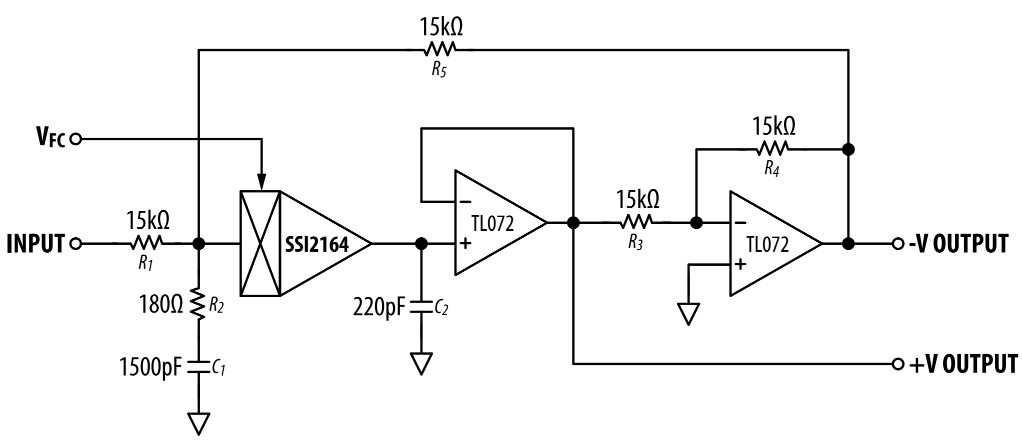

Figure 41: Single SSM2040-Style Low-Pass Filter Stage

If you provide input signals below $1V_{P-P}$, almost no distortion is present. Increase this a bit and you'll get 2040-style asymmetric soft clipping. Increase it more and asymmetric hard clipping is added. So, by changing the input signal level the filter's sonic character changes dramatically.

This circuit has the benefit of both non-inverted and inverted versions of the output being present. This can be of great use when designing multi-mode filters, but comes at the small cost of added noise of the inverter being always present at the output, even if the filter is closed.

By moving the inverter to the input in the circuit of Figure 42, you remove this always present noise at the cost of only having the inverting output being present. This might make sense because of the much smaller signal amplitudes compared to the previous filter circuits and re-amplification could be needed to restore the signal to modular levels for example.

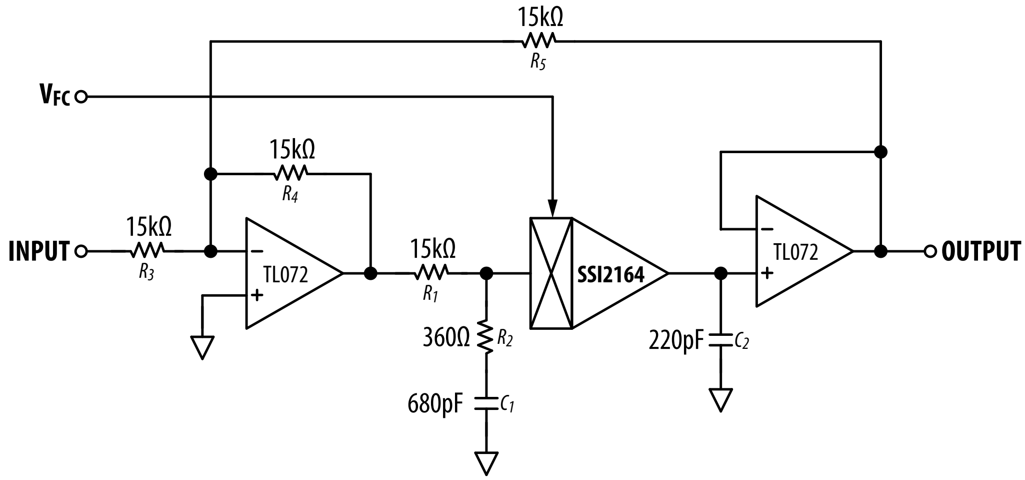

Figure 42: Less Noisy SSM2040-Style Low-Pass Filter Stage

The datasheet says to keep input resistances above 7.5kΩ, but much smaller values can be used at these lower amplitudes as long as the input current is kept below 1mA, with benefit of much lower Johnson noise. Op-amps with suitable drive capacity are then required with matching noise specs.

However, you might ask yourself if dropping the noise floor of the op-amp buffers from -110dBu to -120dBu is worth parts that cost over five times as much.

SALLEN-KEY FILTERS

These are the kind of filters for which the Korg MS-10 and MS-20 are known.

Their basic fixed frequency version is shown in Figure 43, and in-depth info on their MS-10 and MS-20 implementations can be found in the service manuals and their analysis in a document written by Tim Stinchcombe.[29] Both these documents are considered mandatory literature before reading the rest of this section.

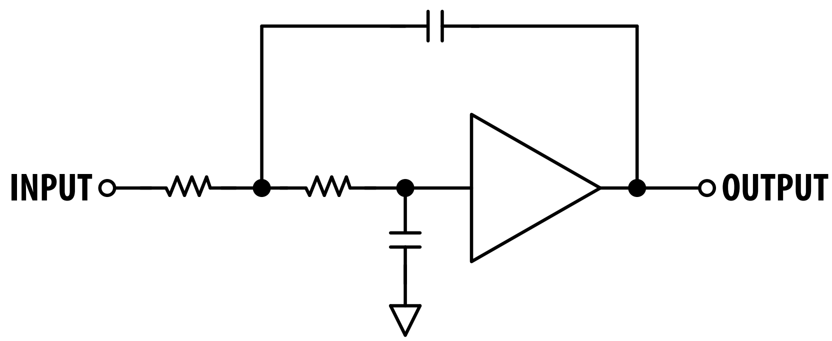

Figure 43: Basic Sallen-Key Low-Pass Filter

The OTA-based version from the late version MS-20 has the structure shown in Figure 44. The first LPF and HPF stage are combined around one OTA VCF stage which has both low- and high-pass inputs.

Unfortunately this HPF-LPF combination does not work with the SSI2164 (see the non-working version in Figure 6). It can be successfully constructed using LPF and HPF stages separately, around 3 sections of a SSI2164. However, by re-arranging things to the equivalent[30] structure in Figure 45, it can be done with only half a SSI2164.

Extra benefits of this structure are that noise contribution of the HPF stage is now filtered through two LPF stages instead of one (see footnote 7) and it is completely removed when the resonance amplifier is closed.

[29] A Study of the Korg MS10 & MS20 Filters, by Timothy E. Stinchcombe, 30 August 2006

[30] The transfer function of the filter in Figure 43 and the one in Figure 44 is exactly the same.

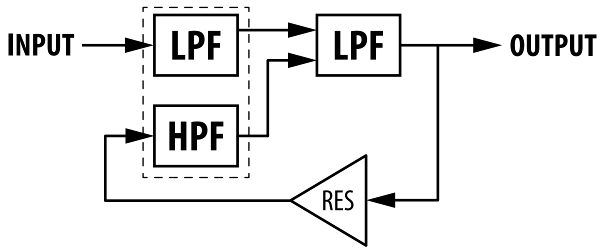

Figure 44: Sallen-Key Low-Pass Filter as Implemented in the Korg MS-20

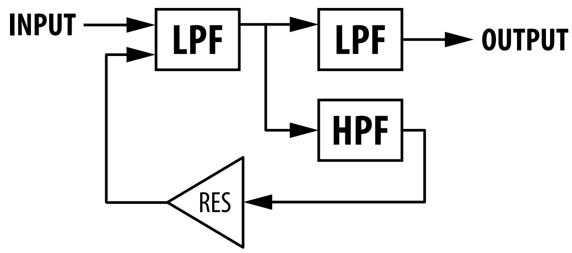

Figure 45: Basic Sallen-Key Low-Pass Filter

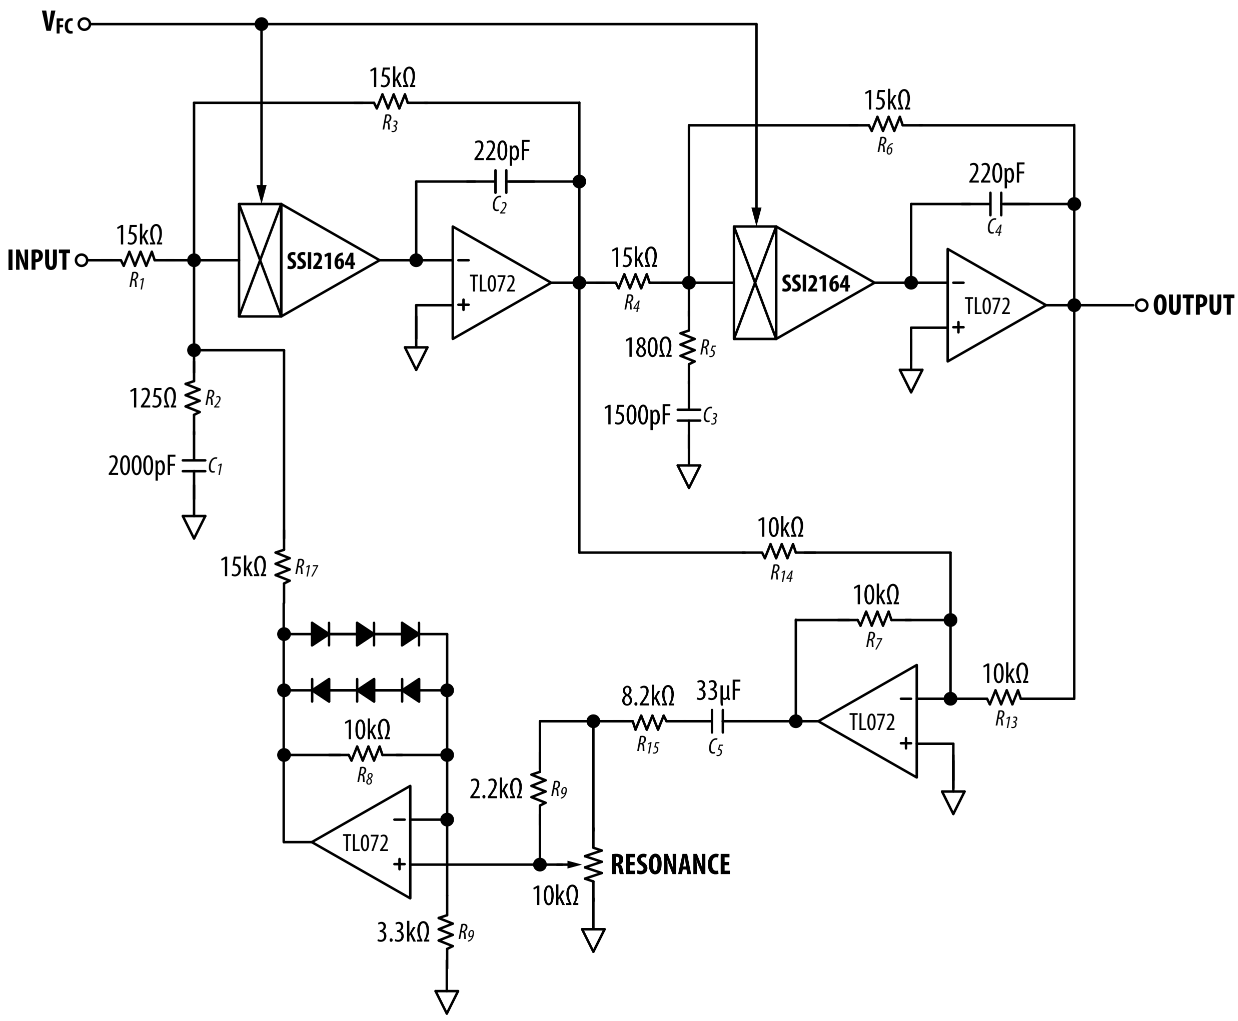

Figure 46: Schematic of a Sallen-Key Low-Pass Filter with the MS-20's Overdrive Circuit

The trick of designing this with only two sections is generating the feedback response as weighted sum of the outputs of the two low-pass stages, as explained in 'Multi-Mode filters' and already used in 'Solving the pass-band attenuation problem. See Figure 33 and footnote 21.

A schematic is shown in Figure 46. This is the same as the one in Figure 33, but without the last two low-pass stages and the clipping circuit has been replaced with the overdrive[31] circuit of the late MS-20 filter to somewhat mimic[32] its main source of extra harmonics.

The old Korg-35 based filter is something of a different caliber. Emulating its un-buffered and open-loop nature is beyond the scope of this document, but some of it quirks are easy to implement, like its distortion, by moving the overdrive circuit to the output and feeding some signal into the control ports (which self-modulates the filter) for its asymmetric resonance.

Contrary to the late MS-20 filter, it is necessary here to use a separate high-pass stage, because the clipping circuit sits in the forward path and has to feed both the output and resonance loop. This generates far more harmonics than the late version, because they are provided unfiltered at the output.

A schematic is shown in Figure 47. Here the high-pass stage of Figure 7 is added and an extra op-amp summer, to add some audio-signal from its output[33] to the frequency control voltage. The amount needed will largely depend on signal strength. As a result, values are not provided here. Of note is also the attenuator just before the overdrive circuit. In the MS-20, this was needed to bring down the signal levels to what the filter could handle and was placed at the input. Here the filter can handle almost rail signals and its only function is to match the original distortion behavior. By moving it to the output of the filter, it also attenuates the noise generated by the filter, before the re-amplification of the overdrive circuit. This move results in 35dB less noise.

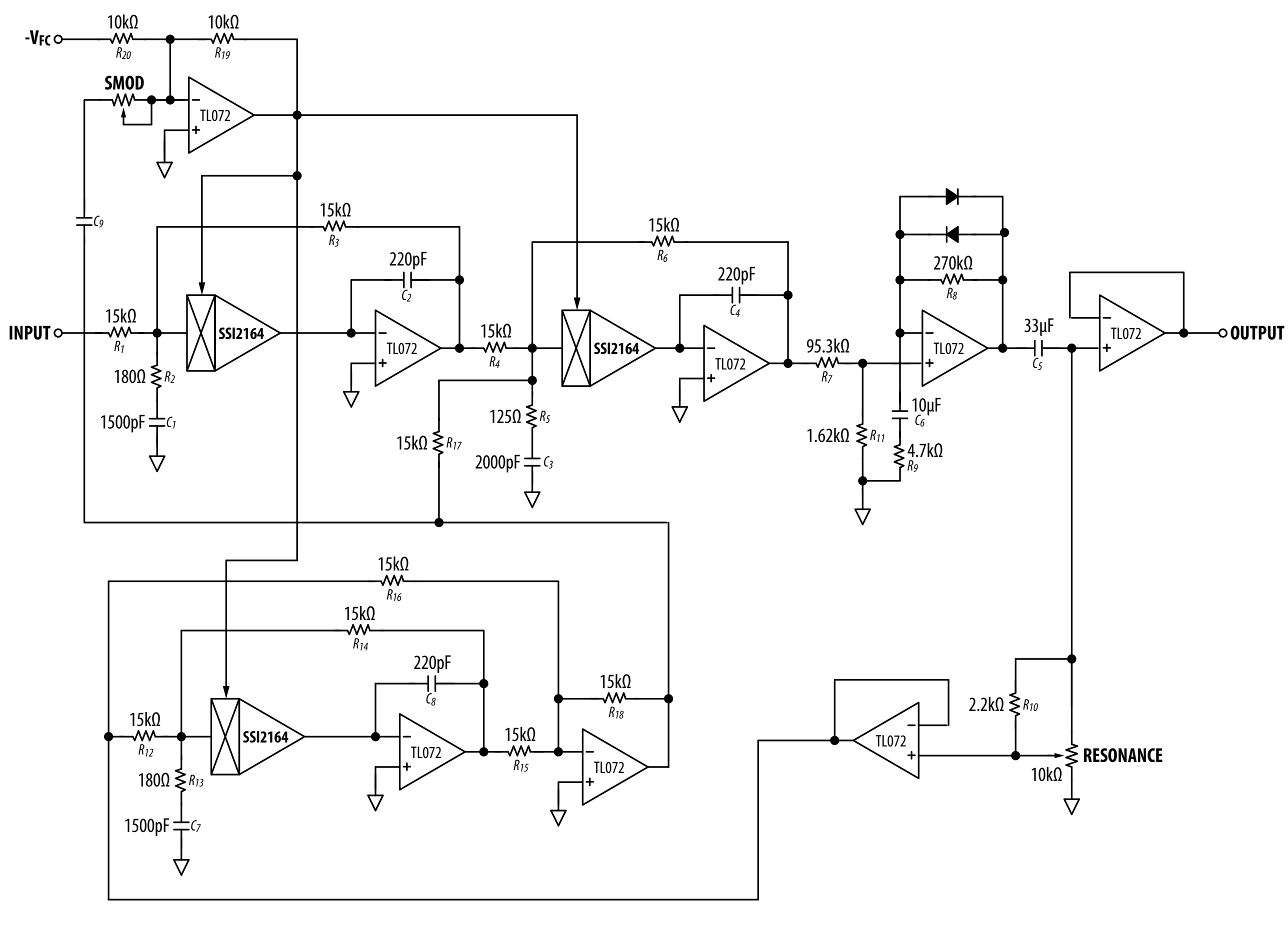

Figure 47: Schematic of a Sallen-Key Low-Pass Filter with the MS-20's Overdrive Circuit and Audio Feedback in Control Ports

[31]This is basically the same as the core of the Ibanez TS-808 Tube Screamer, which is the symmetric copy of very first Boss pedal, the OD-1.

[32]It has the same kind of distortion and sits in the feedback path. However, not in the entirely correct position because that is only accessible if the structure of Figure 42 was implemented, with 3 sections of the SSI2164.

[33]Alternatively the input signal or any filter stage output can be taken, or even the one from the overdrive for different characteristics.

IMPOSSIBLE FILTERS?

Filters whose transfer functions that cannot be represented as a polynomial in S over $(S+1)^n$ are, from a mathematical perspective, impossible to create with the weighted sum of low-pass stages (See "Multi-Mode filters"). However, often they can be approximated fairly accurate.

Useful ones are frequency shifted filters. They can be used for variable slope filters and even for seemingly impossible things, like for example a resonant 6dB per octave (single-pole) low-pass filter that peaks and self-oscillates at its cutoff frequency:

The trick is to combine a normal low-pass stage, which provides 45° phase shift at the cutoff frequency, with a frequency shifted all-pass one. The all-pass provides the required extra 135° phase shift at the cutoff frequency.

The transfer function of such a filter will be:

$$\frac{1}{S+1}\frac{r_f S-1}{r_f S+1}$$

With r_f the ratio of the shifted frequency, f_s, at which the all-pass operates over the cutoff frequency, f_c, of the low-pass stage:

$$r_f=\frac{f_s}{f_c}$$

The phase shift ø of the all-pass stage is:

$$ø=atn\left(\frac{\Im\left(\frac{r_f S-1}{r_f S+1}\right)}{\Re\left(\frac{r_f S-1}{r_f S+1}\right)}\right)$$

This expression can be simplified to:

$$ø=atn\left(\frac{2r_f}{r_f^{\,2}-1}\right)$$

From this, the value of r_f is calculated, for which the all-pass phase shift is -135° at f_c, by solving the above equation with $ø=\frac{\pi}{4}$:

$$r_f=1+\sqrt{2}\cong 2.414$$

Note that 45° is used here instead of -135°, because that lies in the 3rd quadrant[34], and that the second, negative frequency, solution is ignored. The transfer function of the all-pass stage is then:

$$\frac{\left(1+\sqrt{2}\right)S-1}{\left(1+\sqrt{2}\right)S+1}$$

And the transfer function of the complete filter becomes:

$$\frac{\left(1+\sqrt{2}\right)S-1}{(S+1)\left(\left(1+\sqrt{2}\right)S+1\right)}$$

To bring this fraction to denominator $(S+1)^4$, the following division is attempted:

$$\frac{(S+1)^4}{(1+S)\left(\left(1+\sqrt{2}\right)S+1\right)}$$

[34] Arctangent is undefined outside -90° to 90° range.

It has as quotient:

$$\left(\sqrt{2}-1\right)S^2+5\sqrt{2}S-6S+14\sqrt{2}-19$$

And as remainder:

$$20-14\sqrt{2}\cong 0.201$$

The fact that there is a remainder means that this filter's transfer function cannot represented as polynomial over $(S+1)^4$. However, this remainder is a low-valued constant and it will thus only have a minor contribution to the filter's response.

By ignoring this remainder, the transfer function can be approximated by:

$$\frac{S^3+\left(5-2\sqrt{2}\right)S^2-5\left(2\sqrt{2}-3\right)S-14\sqrt{2}+19}{\left(S+1\right)^4}$$

Solving :

$$\begin{aligned}&AS^4+\left(4A+B\right)S^3+\left(6A+3B+C\right)S^2+\left(4A+3B+2C+D\right)S+A+B+C+D+E\\&=S^3+\left(5-2\sqrt{2}\right)S^2-5\left(2\sqrt{2}-3\right)S-14\sqrt{2}+19\end{aligned}$$

Gives the gains[35] for the Multi-Mode filter of Figure 35:

$$\begin{aligned}&A=0\\&B=1\\&C=\left(5-2\sqrt{2}\right)-3=\left(2-2\sqrt{2}\right)\cong -0.83\\&D=-5\left(2\sqrt{2}-3\right)-3-2\left(2-2\sqrt{2}\right)=8-6\sqrt{2}\cong -0.49\\&E=-14\sqrt{2}+19-1-\left(2-2\sqrt{2}\right)-\left(8-6\sqrt{2}\right)=8-6\sqrt{2}\cong -0.49\end{aligned}$$

Its Bode plots are shown (non-inverting) in Figure 48 and demonstrate this approximation is indeed quite accurate. The cutoff slope is an exact 6dB per octave, phase shift at the cutoff frequency is only 3° off from the expected 180° and attenuation is half a dB from the expected -3dB there. Next to that there is less than 1dB ripple below cutoff and the pass-band has only 1.8dB attenuation.

Figure 49 shows its resonant behavior if feedback is added and it is indeed capable of self-oscillation at the cutoff frequency. The schematic is shown in Figure 50.

CONCLUSION

The SSI2164 is an extremely versatile IC useful for designing anything from ultra-clean filters to those with the type of imperfections and character that make analog circuits so desirable.

Due to its improvements over the late 1980's SSM2164 (as well as inferior clones) and advanced IC manufacturing technologies the author believes the SSI2164 is the most suitable component for designing multi-mode filters. Until now, methods have not gone much beyond what was done in the Xpander. But, as shown in this document, the SSI2164 can also be used inside the feedback loop to construct filters beyond the obvious low-pass, high-pass, band-pass, and notch combinations.

One final note: interesting things happen if mathematical accuracy is abandoned in favor of analog character, by using the SSM2040-like circuits of Figures 41 and 42 to construct Multi-Mode filters.

[35] When generating the Bode plots, it was revealed that a small mistake crept in. The transfer function of an inverting all-pass filter was used instead of the non-inverting one to design the filter. This does not affect the math at all and only means that the calculated gains have to be inverted.

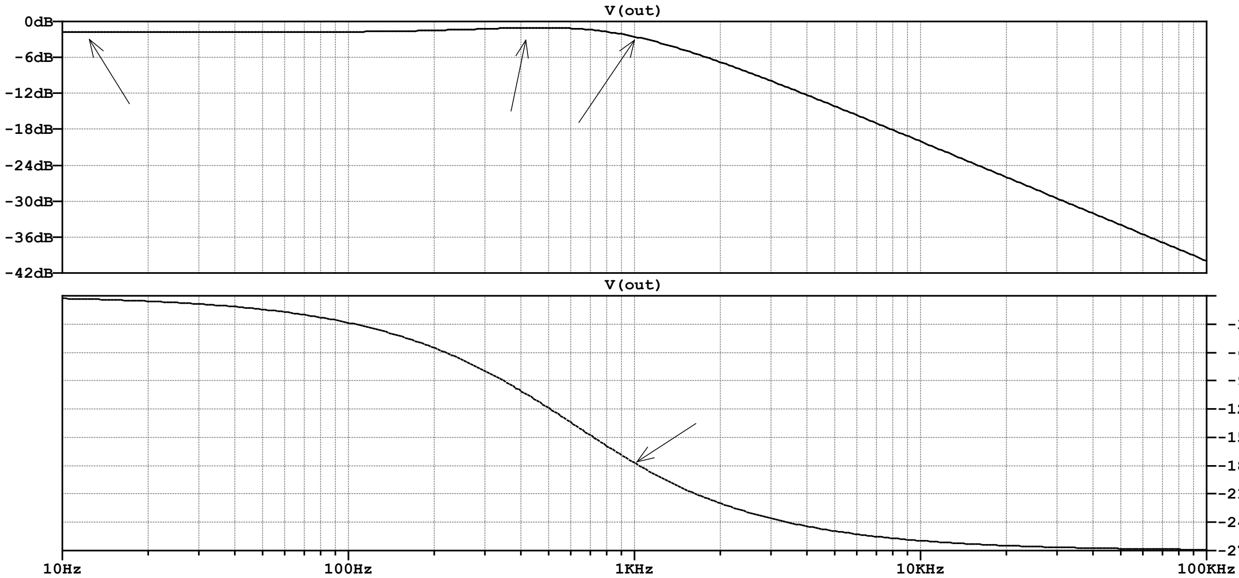

Figure 48: Bode Plots of the Approximated Resonant 6dB per Octave Low-Pass Filter

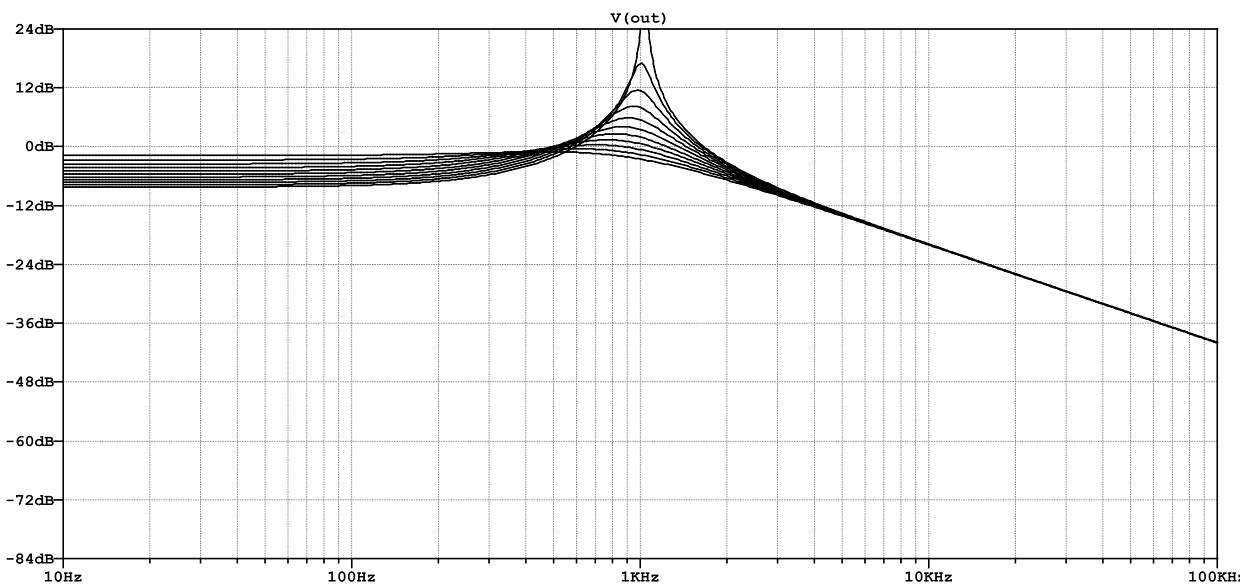

Figure 46: Magnitude Bode Plots of the Approximated Resonant 6dB per Octave Low-Pass Filter, with Increasing Amounts of Resonance Gain

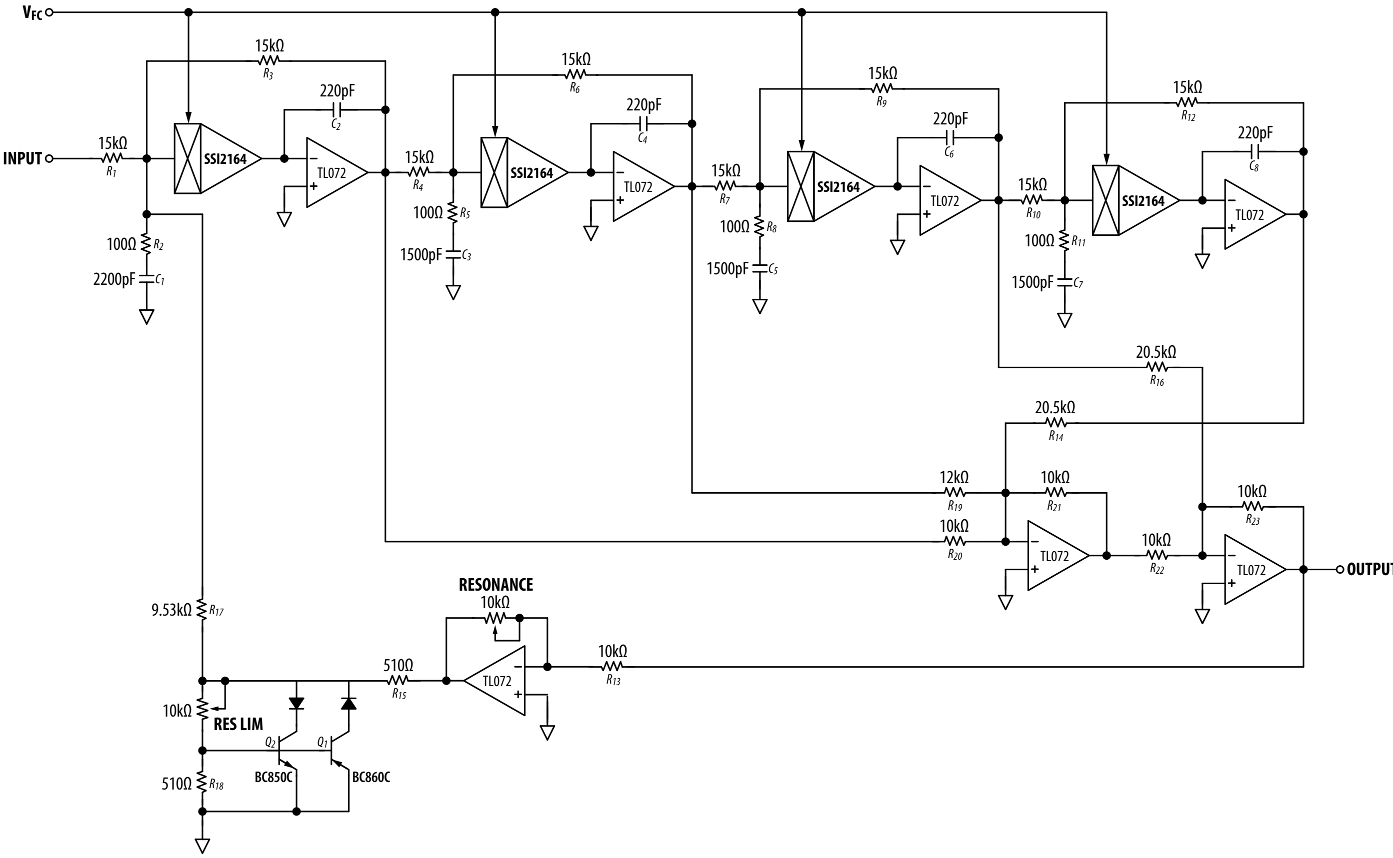

Figure 50: Schematic of a Resonant 6dB per Octave Low-Pass Filter

Jeroen Allaert, despite his relatively young age, has extensive experience in the design of analog synthesizers. His interest in electronics started at age ten, and shortly afterward started collecting electronic musical instruments. He attended Ghent University College (Belgium), writing his thesis on analog synthesis, and received a M.Sc. degree in Electronics with a specialization in Design Techniques. After graduation he worked for Jabil Circuit and then moved to the Photonics Research Group, an associate lab of IMEC at Ghent University. In 2007, Jeroen followed his dream by founding Analogue Renaissance (now Cask Strength Electronics), a specialist in electronic music and custom design. He is also a contributor to Sound Semiconductor. Jeroen is an active DJ and organizes raves in his spare time.