

The valve DAC: the printed circuit boards

Marcel van de Gevel, 15 January 2017

Main board

In a sigma-delta DAC, any crosstalk from the sigma-delta modulate to the voltage reference or to the clock increases the noise floor. To minimize crosstalk, I designed a multilayer PCB for the valve DAC. A four-layer board provides two layers of shielding when you use the inner layers as supply and ground planes and put all the digital stuff on the back side and all the analogue stuff on the front. Inner layer 2 is used as a ground plane. Inner layer 1 is a -300 V supply plane below the valve circuitry, on the rest of the board it is a second ground plane. To give more control over the low-frequency (50 and 100 Hz) return currents, there are no planes below the supply section. The complete schematics and the PCB design can be found in the KiCAD archive

ThevalveDAC_PCB_FPGAfiles\mainPCBfiles\mainPCB_KiCAD\KiCadtop.zip. A PDF version of the schematics and of a part of the layout (not to scale) can be found in ThevalveDAC_PCB_FPGAfiles\mainPCBfiles\mainPCB_pdf_csv. Gerber files are in ThevalveDAC_PCB_FPGAfiles\mainPCBfiles\mainPCB_Gerbers.

Ideally, everything should be made in surface mount technology to minimize the number of wires sticking out of the wrong side of the board. This seemed inappropriate for the valve circuits, though¹. As a compromise, I used SMD technology for the active digital circuitry and through-hole components for the analogue and mixed-signal circuitry, and kept some distance between the FPGA module and the actual DACs. To keep the board costs down (it is still the most expensive part of the DAC), I used a fairly straightforward PCB technology without buried vias. This made it necessary to keep some distance between analogue and digital anyway, as all vias extend to the other side of the board. Where possible I used relatively large SMD components (like 1206 or 0805 size) to make soldering them a bit easier for people over 40, such as myself.

Generally, the through-hole components are on the top side and the SMD components on the bottom side. There are three exceptions; the 24C512 EEPROM, the input connector and an extension connector (which is only needed if you change the FPGA code and need extra I/Os) are through-hole components mounted on the bottom side. They have to be on the bottom because they handle switching digital signals.

The point-to-point wired prototype had several problems related to transmission line reflections. Terminating the lines at the source side with roughly the right impedance was enough to solve these issues. The PCB also contains several resistors meant as transmission line terminators, namely R_{124} , R_{95} , R_{96} , R_{35} and R_{56} . R_{35} and R_{56} are 0 ohm because the estimated output resistance of the DIX4192 is close to the line impedance (including the capacitive load from the various inputs that effectively lower the impedance). Even though they are 0 Ω , I placed these resistors anyway just in case my estimate for the output resistance of the DIX4192 would turn out to be wrong. The clock and clockn lines are necessarily driven from an impedance that is higher than the characteristic impedance. These lines are treated as

¹ People who have seen the rebuilt Colossus in the National Museum of Computing, block H, Bletchley Park, England, know that Colossus was actually made with surface mounted valve holders so valves could be mounted on both sides of the huge mounting panel. The resistors and capacitors were normal wired components, though.

lumped capacitors. Inductors L_2 and L_4 tune out the capacitance of the clock and clockn lines, which is estimated to be about 30 pF, about half of which comes from the lines and about half from the devices connected to the lines. There also are no termination resistors for the clock5 and clockn5 lines, because the 74AHC04 that drives these lines is not hard switched and hence has a high and poorly defined output impedance.

The transmission lines on the four-layer PCB were designed assuming that the dielectric between the outer and inner metal layers is 360 μm thick and has a relative dielectric constant of 4.35, which should more or less match the Eurocircuits PCBProto technology. The dielectric thickness between the two inner layers is not critical, as long as the dielectric can easily handle 300 V.

Even though there are SMD components placed on the bottom, no bottom legend layer was used, simply because the Eurocircuits PCBProto process doesn't feature a bottom legend. Some letters and dots in bottom copper indicate what the correct place and orientation for the SRC4392, DIX4192 and TE0630 module are, for all other components one has to look in the KiCAD database. There is solder mask on both sides, which helps with soldering but also improves insulation.

The boards were designed using the KiCAD open-source PCB design program. It is a very powerful program, although it has some peculiarities. For example, via stitching is quite clumsy; you have to start a wire at a component pad associated with the correct net and then switch back and forth between the layers of the planes that need stitching. Normal wires can sometimes lose their connectivity and never get it back anymore when you accidentally touch them with a component pad that belongs to another net, for example while moving or rotating this unrelated component. Often the undo-button or the "Rebuild board connectivity" function can solve the problem, but sometimes you just have to delete the track and put in a new one.

One KiCAD feature that I rather like is the fact that schematic components can be associated with different footprints. By assigning a connector footprint or no footprint at all, you can also include off-board components in your schematic if you wish. I have used this trick for the switches and the mains transformer.

Regarding the choice of mains transformer, I did most experiments and measurements with an old Nordmende transformer (bought at Radio Twenthe in The Hague) that supplied all the required voltages. As it was quite large and made an audible humming sound, I later switched to using two modern toroids, a Noratel AA-53129 (50 VA, 2 times 115 V) for the -300 V supply and a Noratel AA-58023 (50 VA, 2 times 7 V) for the heaters and the +5 V supply.

Theoretically a 30 VA transformer should be sufficient for the high voltage supply, but using a slightly oversized 50 VA transformer brings the filtered supply voltage closer to the desired -300 V and reduces the overshoot you get when the valves have not warmed up yet. As there are many components used that are rated for 400 V, it is essential that the filtered supply voltage never gets more negative than -400 V, not even when the valves are still cold and the mains voltage is 10 % above nominal.

NTC resistor TH₁, an Epcos B57153S0479M0, acts as an inrush current limiter for the heaters. As there is about half a volt of voltage drop across it when heated up, it should only be used if the voltage from the transformer is a bit higher than the nominal 6.3 V, for example because you use a transformer meant for 220 V with 230 V mains or because you use a 7 V

instead of a 6.3 V toriod. In fact I had to add an extra 0.1 Ω resistor in the heater circuit to get to the correct voltage with the Noratel AA-58023.

A suitable connector for the connection to the transformer(s) is Würth Elektronik type 691311500006 from the WR-TBL / 311 series, or any other well-insulated 6-pin connector with 5.08 mm pitch. The extension connector and the connector for the SPDIF/AES3 inputs can be standard DIN 41651 flat cable connectors. Most of the other connectors are shrouded headers with 2.54 mm pitch. I used Amphenol FCI Dubox headers, crimp sockets and socket housings; the sockets and socket housings officially require a special crimping tool, but they are easily mounted without one once you realize that the barb of each contact has to point to the side of the connector. The FPGA module requires two special SMD connectors, as indicated on the schematic.

If you want to connect the neon lamps by connectors, these need to have 5.08 mm pitch, yet have thin enough pins to fit in 1.3 mm holes. A normal 2.54 mm-pitch header is not suitable because when the lamps are switched on, the voltage across them can be over 300 V for a few microseconds. With 2.54 mm pitch, the pads on the PCB would have been too close for 300 V. A terminal block with modest current rating will probably work, or maybe one can take a three-pin shrouded header and remove the centre pin.

Nothing needs to be mounted at the indicated test points (TP).

Decoupling capacitors C_{10} , C_{19} , C_{23} , C_{30} , C_{44} , C_{45} and C_{52} should be film capacitors with a very low inductance, such as stacked rather than wound capacitors, or wound capacitors that have all windings shorted.

Capacitors C_{35} must be an ordinary, run-of-the-mill aluminium electrolytic capacitor rather than a low-ESR / low $\tan(\delta)$ device. It is used as a frequency compensation capacitor for the LM317 and its losses are an essential part of the compensation.

I specified a working voltage of at least 1000 V (DC) for C_{119} and C_{120} because the voltage across them has a large AC component. Large AC voltages are more difficult for the capacitor's insulation than a DC voltage equal to the peak value of the AC voltage, see <http://www.wima.de/EN/pulseselection.htm>. On top of that, there is only limited filtering between the mains and the voltage across these capacitors and there can sometimes be nasty spikes on the mains. Class X2 capacitors meant for at least 300 V AC would be even better.

The inductors $L_1...L_5$ in the crystal oscillator and clock buffer are small colour-coded axial inductors. The self-resonant frequencies of L_1 , L_2 , L_4 and L_5 obviously need to be much greater than 27 MHz. I used inductors of the Bourns 78F series. Trimmer capacitor C_7 is a ceramic trimmer from the Multicomp CV05 series. It has two pins with a pitch of 5 mm.

The DIX4192 datasheet contains many suggestions for possible SPDIF or AES3 input application circuits. If you should choose a transformer-coupled version, it makes sense EMC-wise to mount the transformer straight behind the panel of the enclosure of the DAC. I therefore did not reserve any space on the PCB for SPDIF or AES3 transformers. Similarly, I did not put an optical receiver on the board, only a 5 V-compatible level shifter you can use with 5 V as well as 3.3 V optical receivers. To prevent overlap currents, don't forget to ground the level shifter's input when there is no optical receiver connected to it.

Reconstruction filters

Except for the first filter capacitors and the DC blocking capacitors, the reconstruction filters are mounted on a separate single-layer board. As the filter components are quite large, have only few connections between them, carry no high voltages and are not exceedingly sensitive to board leakage currents, a simple single-layer board with no solder mask works fine and is much cheaper than using a larger four-layer board.

The schematic diagram and the PCB design can be found in the KiCAD archive ThevalveDAC_PCB_FPGA_files\filterPCBfiles\filter_KiCAD\filter.zip. A version meant for a single channel or for an unbalanced version is ThevalveDAC_PCB_FPGA_files\filter_singlechannelPCBfiles\filter_singlechannel_KiCAD\filter_singlechannel.zip. A PDF version of the schematics and a Gerber version of the PCB layout are placed in directories with rather self-explanatory names.

Reconstruction filter topologies

As suggested in the valve DAC article, there are three options for the filter topology: balanced with transformer, balanced without transformer and single-ended. Figure 1 shows the balanced version with transformer. Resistors R_3 and R_7 are to be chosen such that the parallel connection of the resistor and the expected load is $600\ \Omega$. The indicated value of $634\ \Omega$ is meant for loads around $10\ \text{k}\Omega$ ($634\ \Omega$ being the E96 value closest to the theoretical $638.3\ \Omega$). For $600\ \Omega$ loads these resistors are obviously not mounted. Wire resistance of the transformer brings the differential-mode impedance at the primary side to around $750\ \Omega$. The RC-networks R_1 - C_1 - C_4 , R_4 - C_5 - C_8 , R_5 - C_9 - C_{12} and R_8 - C_{13} - C_{16} compensate for the spreading inductance of the transformers. R_2 and R_6 are the common-mode terminations of the filters, these need not be accurate.

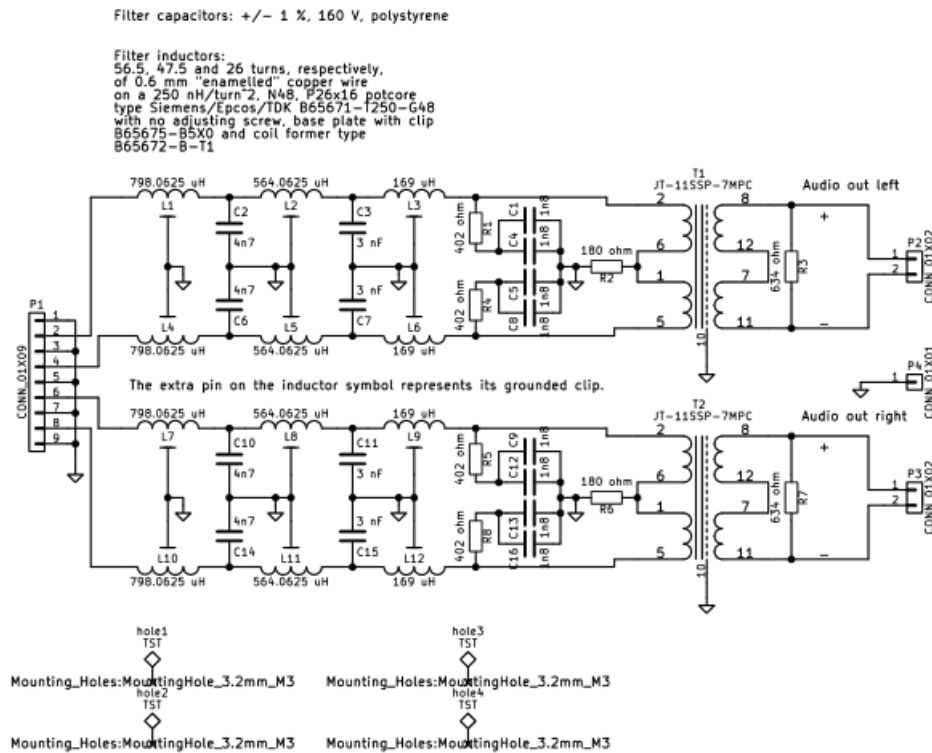


Figure 1 Balanced filter with transformer. Please keep in mind that pins 2 and 8 of the transformers are diagonally opposite to each other. There are two possibilities for the inductances and capacitances, see Table 1.

Figure 2 shows the transformerless balanced version. It is much cheaper than the version with transformer, but has the disadvantage that it can only drive balanced inputs. In fact it can only drive relatively high-impedance ("bridging") balanced inputs because of the need to have a reasonable common-mode termination. R_1 , R_4 , R_5 and R_8 are to be chosen such that the parallel connection of the load and the series connection of two of these resistors is $750\ \Omega$. This boils down to $405.4\ \Omega$ resistors (nearest E96 value: $402\ \Omega$) for $10\ \text{k}\Omega$ load. As suggested by the way the schematic is drawn, I haven't designed a PCB for this variant, but it can easily be mounted on the normal filter board by leaving out some components and using some wire bridges.

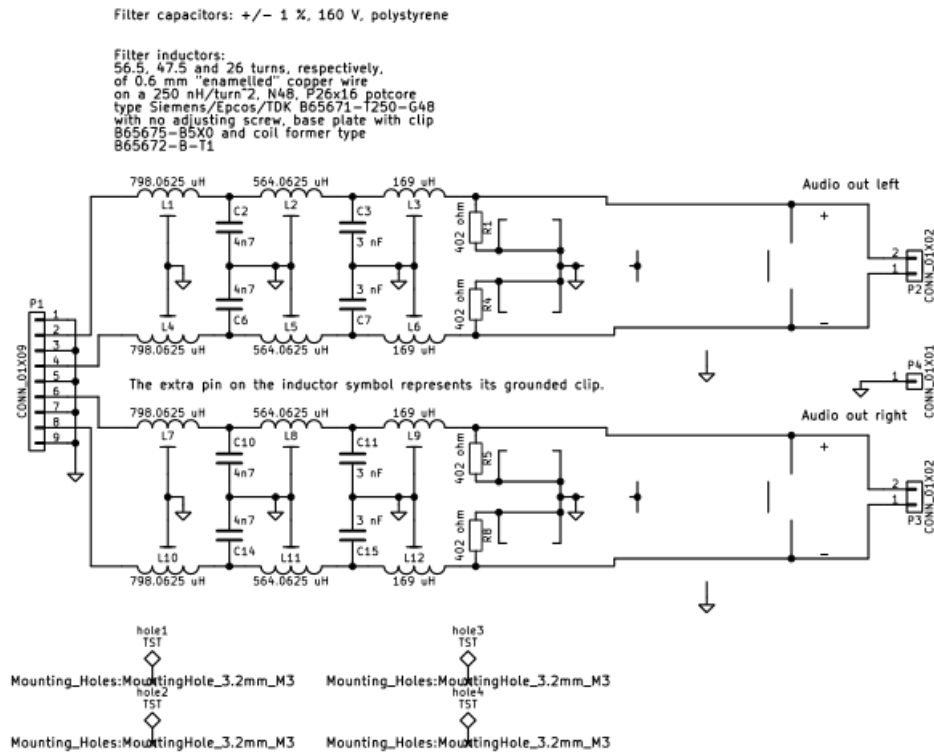


Figure 2 Balanced filter without transformer; note the reversed polarity of the output compared to Figure 1! There are again two possibilities for the inductances and capacitances, see Table 1.

Figure 3 shows a transformerless version for single-ended loads. This is the cheapest version, but as there is no suppression of common-mode noise, the ratio of the maximum signal to the noise floor is less than with the other versions (the difference was measured to be about 3 dB). The output signal level is also halved compared to Figure 2. R_1 or R_5 in parallel with the load impedance must be 375 Ω . For 10 k Ω loads, this boils down to 389.6 Ω , or 392 Ω after rounding to an E96 value. The 10 μ F DC blocking capacitor of the negative output on the main PCB must be left in place, it is only grounded on its output side now.

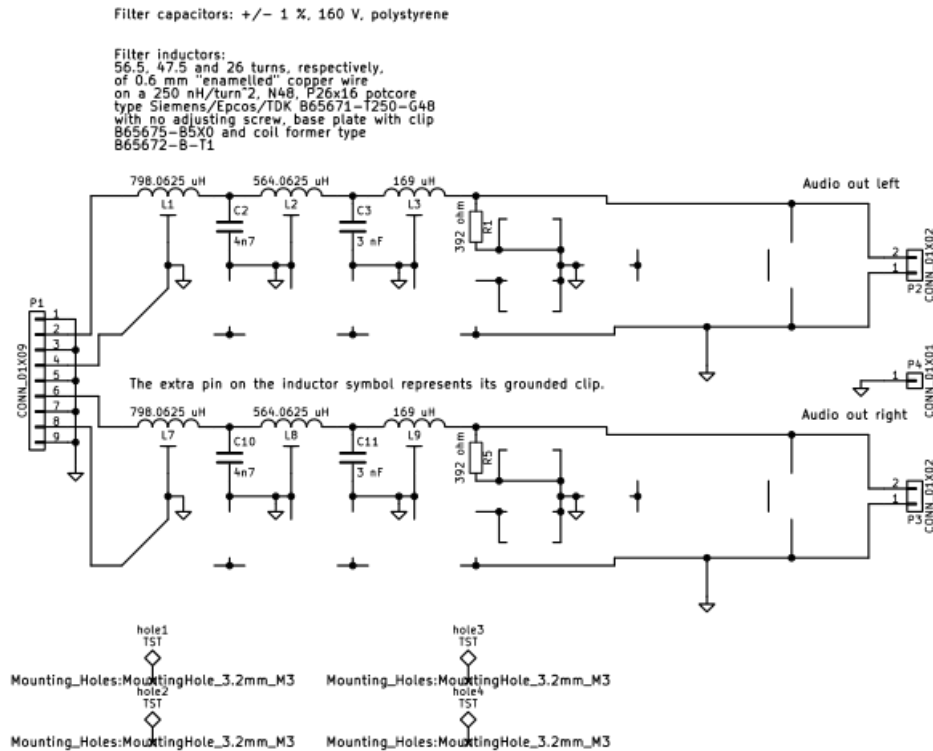


Figure 3 Single-ended version. The negative output of the main PCB is grounded now. There are still two possibilities for the inductances and capacitances, see Table 1.

The main PCB and filter PCB together are just too wide to be mounted in a 19-inch enclosure. It may therefore be handy to split the filter PCB up into two parts that can be mounted on top of each other, hence the "filter_singlechannel_*" files. If you choose the unbalanced filter version, one single-channel PCB suffices.

Reconstruction filter component values

As mentioned in the article, the valve DAC performs best in PWM8 mode, but the out-of-band noise in PWM8 mode is a bit high if you should want it to drive a solid-state amplifier with poor slew rate. A possible solution is to go for other analogue filter coefficients.

In the 1960's, DeVerl S. Humpherys invented the so-called transitional Gaussian-Chebyshev or transitional Gaussian-to- x dB filters. Theoretically, a filter with Gaussian impulse response would have linear phase and no overshoot or ringing. All well-known filter types with almost linear phase response and little overshoot, like Bessel filters and 0.05° linear phase filters, approximate the theoretical Gaussian filter in some sense or other.

By making a Chebyshev approximation to the magnitude response of a theoretical Gaussian filter up to the point where this response has dropped to $-x$ dB, Humpherys obtained filters that have smooth roll-off to the $-x$ dB point and then become steep. Their passband phase response is quite good. The step response has some ringing, but at frequencies that are already filtered out by the digital filters and that should therefore not be excited much by the

signal. Anyone who wants to know the ins and outs of these filters is referred to DeVerl S. Humpherys, "Equiripple network approximations using iteration techniques", *Proceedings of the National Electronics Conference*, vol. 20, 1964, pages 753...758 or to DeVerl S. Humpherys, *The analysis, design and synthesis of electrical filters*, Prentice-Hall, Englewood Cliffs, N. J., 1970. The LC prototype values can be found in Anatol I. Zverev, *Handbook of filter synthesis*, Wiley, New York, 1967.

The alternative values are shown in Table 1. To get a better match with the passband response of the original filter and to further improve phase response, C_2 , C_6 , C_{10} and C_{14} are increased a bit compared to their theoretical value (theoretically they should be 6.07413 nF, assuming ideal current source drive). Use at your own risk, as I haven't tried this variant.

Table 1 Alternative filter component values. The termination resistors are not changed, neither are the extra RC networks on the version with transformer.

Components	Original value (equiripple phase)	Alternative value with better out-of-band noise rejection (transitional Gaussian to 6 dB)
C_{28} , C_{29} , C_{50} and C_{51} on main PCB	8.2 nF	8.2 nF
L_1 , L_4 , L_7 and L_{10}	798.0625 μ H, 56.5 turns	915.0625 μ H, 60.5 turns
C_2 , C_6 , C_{10} and C_{14}	4.7 nF	6.8 nF
L_2 , L_5 , L_8 and L_{11}	564.0625 μ H, 47.5 turns	826.5625 μ H, 57.5 turns
C_3 , C_7 , C_{11} and C_{15}	3 nF	4.7 nF
L_3 , L_6 , L_9 and L_{12}	169 μ H, 26 turns	306.25 μ H, 35 turns

Build up

It is handy to first solder the SMDs, particularly the fine-pitched ones, before starting with the through-hole components. Once there are through-hole components mounted, the PCB can not be laid flat on a table anymore. Besides, if soldering the fine-pitched SMDs should go wrong and the PCB should become unusable, at least all larger components are salvaged for a new attempt on a new PCB.

Although it is much easier to solder TQFP-48 packages with solder paste and hot air, it is not entirely undoable with a good soldering station and conventional solder wire². Properly align the package and solder two opposite corner pins. After checking again that everything is really well-aligned on all of the four sides and adjusting if needed, try soldering the other pins as neatly as you can, making sure that every pin indeed gets soldered (on a board with HAL finish, it is generally easier to see shorts than to see pins without solder). Inevitably there will be quite some pins shorted by excess solder, but this can easily be solved with fine desoldering braid. Check with a magnifying glass (or with the naked eye when you are short-sighted) and/or with an ohm meter and two fine needles that all pins are correctly soldered and are not shorted to their neighbours. Soldering the 80-pin connectors for the FPGA module is done in a similar way, but the alignment is much easier because they have plastic pins that just fit in holes in the PCB.

² At least not with good-quality tin-lead solder, I didn't dare to try it with lead-free solder. The larger components can be hand-soldered with lead-free tin-silver-copper solder without any problems.

Some resistors are best mounted a few millimetres above the surface of the PCB, either because they get quite hot or to improve insulation to a track that passes under them. I did this with R₃₆, R₄₁, R₄₃, R₄₄, R₄₈, R₈₆, R₁₀₀, R₁₂₈, R₁₃₁, R₁₃₂, R₁₃₆, R₁₃₉, R₁₄₀ and R₁₄₁.

The polystyrene capacitors have the lowest melting point of all the through-hole components, so I mounted them last. These capacitors are all grounded on one side. They usually have a ring indicating which side is connected to the outer foil. This side must preferably be connected to ground, so that the outer foil works as an electrostatic shield.

It is advisable to remove flux residues, particularly from small DC blocking capacitors that handle high voltages, such as C₁, C₂, C₃, C₂₆, C₃₃, their right-channel colleagues C₂₀, C₂₁, C₂₂, C₄₈, C₅₅, and C₄ and C₁₅. With a bit of moisture, flux residues become electrolytic conductors. I usually gently scratch off most of the residue, making sure not to damage the solder mask, and try to remove the remains with a cotton bud soaked in isopropyl alcohol and a dry cotton bud to dry the PCB again. Mind you, some IC type number prints dissolve in isopropyl alcohol.

Valve sockets for PCB mounting are available with many incompatible footprints. Two relatively common types are ceramic noval (B9A) sockets with flat 1.6 mm-wide pins placed on a circle with a diameter of approximately 21 mm, and ceramic sockets with flat 1 mm-wide pins placed on a circle with a diameter of approximately 19 mm. The footprint I drew for the noval valve sockets is more or less compatible with both, as it has pads with a hole of 2 mm placed on a circle with a diameter of 20 mm. The 85A2 requires a miniature (B7G) socket, these often have pins of about 1 mm wide on a circle with a diameter between 16 mm and 19 mm in diameter. The PCB footprint has pads with 2 mm holes placed on a circle with a diameter of 17 mm. As many valve sockets also have a centre pin, the footprints have a 3.5 mm hole in the middle connected to the -300 V plane.

The PCB design on the Linear Audio website has many improvements compared to the first PCB. Besides the changes discussed in chapter 13 of the article, some clumsy footprints were replaced with better versions. For example, the 100 µF, 450 V electrolytic capacitors originally had footprints with 1.1 mm holes and 10 mm pitch, while most radial 100 µF, 450 V electrolytic capacitors either have wires with 7.5 mm pitch, or snap-in pins that require 2 mm holes at 10 mm pitch. The modified PCB design is compatible with both.

There are many footprint-incompatible variants of the BAS70 diodes used in the DAC core. The variant I've used is a single diode in a SOT-23 package that is simply called BAS70 and not BAS70-04, BAS70-05, BAS70-06, BAS70W, BAS70Z, BAS70J, BAS70K or BAS70-with-cream-on-top.

Availability of the components

Good sources for valves and associated components are Radio Twenthe in The Hague and the NVHR fairs, or similar antique electronics fairs in other countries. The polystyrene capacitors, the trimmer capacitor and the connectors for the FPGA module came from Farnell, the potcores from Distrelec. The SO-14 version of the 74AHCU04 is, unfortunately, not available from Farnell anymore, but Mouser and Digikey still supply it. Most of the other components are available from any of the larger electronic component distributors. High capacitance polypropylene capacitors are readily available at loudspeaker DIY shops such as Speaker & Co. The FPGA module and its JTAG programming cable come from Trenz Electronic. The optional output transformers came straight from Jensen Transformers; in three day's time they travelled all the way from the west of the USA to the Netherlands before getting stuck at the Dutch customs for twelve days.

As some inverters of the SN74AHCU04 (U₁₇) are used in a rather analogue fashion and as U₁₇ and the flip-flops U₄ and U₂₇ (74AHCT74) directly impact the timing of the data signal at the valves with respect to the clock, I recommend using the exact same brand and type of chips as I used, if possible. That is, a TI SN74AHCU04DE4 and NXP 74AHCT74D,112 devices.

Preparing the FPGA module

Before putting the FPGA module in its sockets, check R₁₀₂ and R₁₀₃ on the module and program the flash memory as explained in the TE0630 user manual. My configuration file needs R₁₀₂ mounted and R₁₀₃ not mounted (which connects I/O bank 0 to the FPGA's 3.3 V supply). Instead of using an official 0 Ω resistor for R₁₀₂, you can just as well put a blob of solder between its pads. As explained on the Trenz website under TE0630 known issues, switch S_{1A} must be off and S_{1B} on (unless you want to write your own FPGA configuration that uses the module's USB interface).

Regarding programming, if you want to use the exact same configuration file I used, you only need to program the flash memory with the MCS file using the free Xilinx software and a JTAG programming cable, using the procedure explained in the TE0630 manual. If you wish to change the configuration, you will need the Verilog files, the constraints (UCF) file, the FIR filter coefficient files (.coe), the overview of settings of the FIR compiler and the overview of synthesis and implementation settings. All valveDAC FPGA-related files are in ThevalveDAC_PCB_FPGA_files\FPGAfiles.

I chose the Verilog hardware description language because of its popularity among the digital designers I know; however, the more I use it, the less I understand its popularity. For one of the main logic circuit design tools, it is a remarkably illogical and inconsistent language full of stupid quirks. For example, there are two ways to define a constant, but only one works properly for constants longer than 31 bits unsigned or 32 bits signed, the other results in an incorrect value. Net types must be declared but also have a default, so you won't get a sensible error message when you forget to declare one. One of the silliest features is the way additions of signed and unsigned numbers are handled; when at least one of the variables to be added is unsigned, the whole addition becomes unsigned. People new to Verilog can save enormous amounts of time by first reading Stuart Sutherland and Don Mills, *Standard gotchas – subtleties in the Verilog and System Verilog standards that every engineer should know!*, 2006, available on http://www.sutherland-hdl.com/papers/2006-SNUG-Boston_standard_gotchas_presentation.pdf.

I have used switches S_{1C} and S_{1D} on the FPGA module to control how the clipping lamp U₂₂ behaves. A switched-on S_{1D} switches off the clipping lamp when you are in surprise mode. With S_{1C} switched on, the clipping neon lamp does not react to quantizer clipping, only to integrator and adder clipping.

With both switches off, the clipping neon lamp is powered for about 593 ns (actually less than its ignition delay) whenever the quantizer clips, for about 1 s whenever an integrator clips and for about 5 s whenever an adder clips (which indicates that there is a bug in the FPGA configuration, as adder clipping is supposed to be impossible; in any case, I haven't seen this happen so far). Quantizer clipping is fairly normal in a sigma-delta modulator, so with S_{1C} switched off, the lamp seems to light up continuously and only shines a bit brighter when the music is loud. Occasional quantizer clipping does little harm because the resulting error gets suppressed quite effectively by the noise shaping loop. Integrator clipping, on the other hand, indicates that the sigma-delta is really working at higher levels than it can properly handle.

Trimming procedure

The tuning procedure of the oscillator is as follows. Put the trimming potmeter R_{42} somewhere halfway its range. Attach a DC voltage meter between the testpoints (across R_{52}) and adjust C_7 until testpoint P_3 is as far as possible below the negative supply rail (P_4). Use a well-insulated adjuster rather than an ordinary metal screwdriver for safety reasons as well as to prevent detuning. Then adjust the trimming potmeter R_{42} for -0.5 V between P_3 and the negative supply rail, readjust C_7 and readjust the potmeter for -0.5 V between P_3 and the negative supply rail. If the testpoint doesn't go below the supply rail at all, reduce the potmeter's resistance and try again. Assuming ideal grid rectification and taking into account the maximum specified crystal resistance and the measured typical value of C_0 , a DC voltage of -0.5 V should roughly correspond to the allowable 500 μ W drive level.

In the DAC cores, the trimming potmeters RV_1 (left) and RV_2 (right) that set the balance between the DAC anode resistors must be adjusted for minimum idle-channel noise. This can be done by listening to the noise without any input signal while adjusting the trimming potmeters, trying to ignore the scratching sounds that occur when you change the wiper position.