CZECH TECHNICAL
UNIVERSITY
IN PRAGUE

F3

Faculty of Electrical Engineering
Department of Microelectronics

Bachelor's Thesis

Modular Analog Synthesizer with Digital Sequencer

Ing. Jan Onderka

2022

Acknowledgement / Declaration

I would like to thank my supervisor doc. Ing. Stanislav Vítek, Ph.D., for his supervision and inputs.

In addition, I would like to thank my family for supporting me throughout the course of my studies.

I dedicate this thesis to my late cat Blecha. She brought me many joys during the time we were together.

Prohlašuji, že jsem předloženou práci vypracoval samostatně a že jsem uvedl veškeré použité informační zdroje v souladu s Metodickým pokynem o dodržování etických principů při přípravě vysokoškolských závěrečných prací.

V Praze dne 20. 5. 2022

....

Abstrakt / Abstract

V této práci jsou shrnuty často používané techniky pro konstrukci modulárních zvukových syntetizátorů. Je zhotoven modulární syntetizátor s digitálním sekvencerem a analogovými moduly navržený zejména pro účely demonstrace a studia elektrotechniky.

Pro kompatibilitu s profesionálními syntetizátory jsou následovány technické parametry běžného standardu modulárních syntetizátorů Eurorack. Místo upevnění modulů ve skříni jsou nicméně použity malé plastové konstrukční krabičky pro snížení ceny a jednodušší výrobu.

Syntetizátor je zkonstruován a vyhodnocen s důrazem na cenu, užitné vlastnosti a složitost. I přes některé problémy, které by mohly být opraveny v další revizi, je syntetizátor plně funkční a hodí se pro užití jako demonstrační a výuková pomůcka.

Klíčová slova: analogový syntetizátor, digitální sekvencer, rozdílová syntéza, nástroj pro demonstraci elektroniky

In this thesis, commonly-used techniques for construction of modular audio synthesizers are summarized. A modular synthesizer with a digital sequencer and analog modules is constructed, designed especially for electronics demonstration and learning purposes.

For compatibility with professional synthesizers, the electrical aspects of a common modular synthesizer standard Eurorack are followed. However, instead of rack mounting, small plastic box enclosures are used for cost-effectiveness and ease of construction.

The synthesizer is constructed and evaluated with focus on cost, performance and complexity. Despite some problems which could be fixed in a later revision, the synthesizer is fully functional and well-suited for usage as a demonstration and learning tool.

Keywords: analog synthesizer, digital sequencer, subtractive synthesis, electronics demonstration tool

Contents

Tables / Figures

Chapter 1
Introduction

Synthesizers are musical instruments in which audio signals are generated and altered electronically. Unlike classic instruments such as pianos or guitars, no conversion from mechanical domain (e.g. vibration of strings) to electrical domain is performed. This has distinct advantages in the practically unlimited audio possibilities. In practice, however, synthesizers for production of music (as opposed to sound effects) are typically inspired by classic instruments, allowing for basic emulation of their sound while providing additional possibilities.

Musical analog synthesizers first started to reach musical prominence in mid-1960s, especially due to the synthesizer makers Robert Moog and Don Buchla. With increasing prominence of digital electronics, the main commercial focus later switched to digital synthesizers. Each approach has its distinct advantages and disadvantages [1, p. 1-2].

I have been intrigued by audio synthesizers and closely related effect units for a long time. It is a specific area of electronics with low entry requirements but high complexity. A rudimentary synthesizer can be made on a breadboard with few components, but professional-grade audio synthesizers can easily cost thousands of dollars or even more. Due to the sensitivity of human hearing to undesirable sound artifacts, such as bad tuning or non-harmonically related frequencies in the sound spectrum, designing a practically usable and cost-effective synthesizer is a major engineering challenge.

The goal of this thesis is to build (and evaluate) a modular synthesizer for electronics demonstration purposes. This means that the synthesizer characteristics should lie between "toy" synthesizers and costly professional synthesizers. That way, in a university setting for example, students could use and try to understand the various modules, which would use circuit techniques similar to professional modules. Furthermore, they could also design their own module and interface it with others without much fear of destroying costly equipment.

I wanted to keep the modules compatible with professional synthesizers to allow for expansion and interplay. After weighing my options, I decided to make the synthesizer electrically compatible with the Eurorack specification [2], but discard the mechanical specifications of rack-mounted units in favour of small, self-contained printed circuit boards enclosed in small construction boxes. This choice reduces cost and build complexity, but also structural rigidity present in rack units.

During the course of my work, I designed six modules for my modular synthesizer (not counting a few trivial support boards):

- **Power module**, which converts power from 12 V AC power adapter to ±12 V DC power rails from which the other modules are powered.
- **Digital sequencer** comprised of an STMicroelectronics development kit and a custom daughterboard, which provides control capabilities via buttons, potentiometers, and LEDs. The daughterboard also contains circuitry for translation of signals between the ARM microcontroller unit (MCU) / codec levels and Eurorack levels.
- **Voltage-Controlled Oscillator (VCO)** with a 1 volt per octave (V/Oct) Control Voltage (CV) input and triangle, pulse and ramp outputs. to avoid the need for manual

trimming, I implemented a digital trimming solution using an AVR MCU and digital potentiometers.

- **Voltage-Controlled Amplifier (VCA)** with largely linear CV gain control.
- **Voltage-Controlled Filter (VCF)** with CV-controlled filter cutoff.
- **Attack-Release Envelope Generator (AR EG)** with CV-controlled Attack and Release intensities.

After designing and building the modules, I performed objective measurements and tested the usability of the synthesizer. There are some problems in certain modules due to design oversights, and output characteristics are lackluster due to noise levels and some signal contamination by mains frequency harmonics. These problems could be alleviated in a future revision.

The form factor is not optimal for professional use due to lack of structural rigidity, but it seems well-suited for its intended purpose as an electronics demonstration tool. I plan to continue the development of the synthesizer in the future so it can be, hopefully, used by others as a demonstration and learning aid.

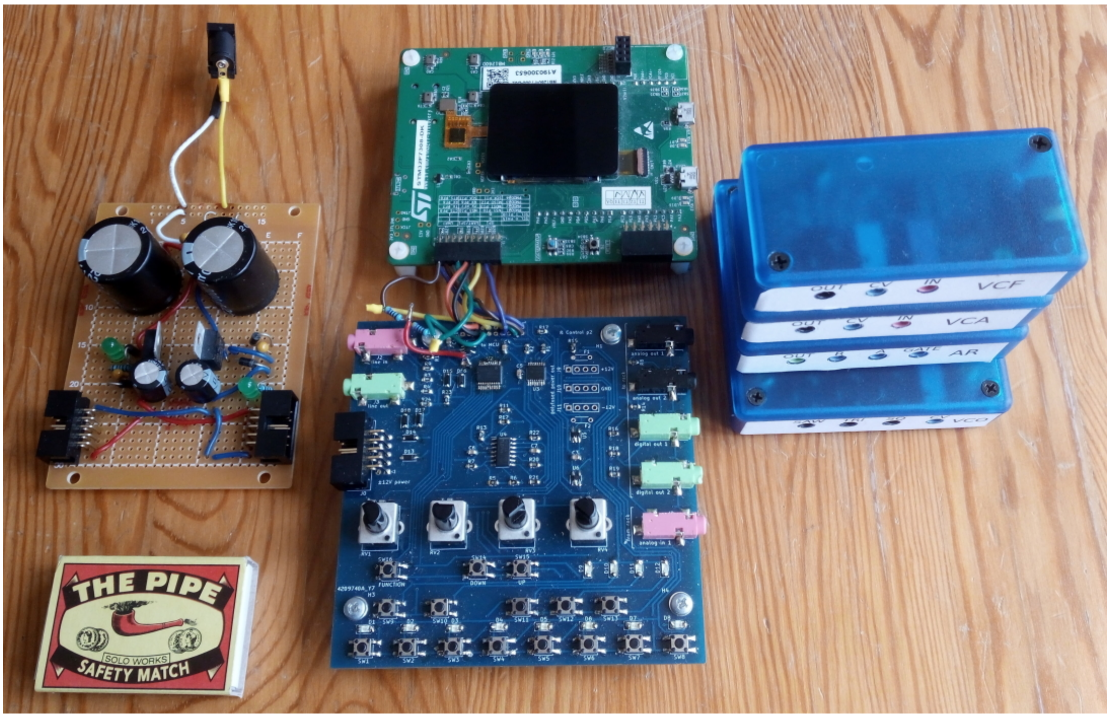

Figure 1.1. Photo of all modules, not including support boards. From left to right: power supply, sequencer (STM32F7308 Discovery Kit with daughterboard), the other modules which are enclosed in Hammond Manufacturing Corporation 1551 Series enclosures. A standard size matchbox is included for size reference. The daughterboard is connected to the Discovery Kit manually using wires instead of a proper connector due to a design error discussed in Section 5.2.

Chapter 2
Synthesizer design considerations and approaches

In this chapter, I present signal considerations that should be taken into account for synthesizer construction and high-level classifications of various approaches to synthesizer design. I discuss my choices of approaches most suited to my goals.

2.1 Signal considerations

In modular synthesizers, four major types of signals are present [3, p. 130], presenting slightly different challenges. Between modules, signals are ubiquitously carried as voltages, similarly to almost all common signal interconnects.

- **Audio signals.** The importance of various audio signal properties is determined by human hearing capability. From a high-level perspective, the hearing process can be viewed as a complex time-frequency analyzer [4, ch. 2]. The audible frequencies are contained in a band spanning approx. from 20 Hz to 20 kHz, the exact range depending on the listener. For audio equipment construction, this has several consequences for audio signal handling:

  - The stationary component may be discarded.
  - Components designed for almost-stationary signals (i.e. some operational amplifiers) cannot be used.
  - Radio-frequency design techniques are unnecessary as the wavelength of a 20 kHz wave in free space is approximately 15 kilometres, well above synthesizer size. Circuit analysis techniques are sufficient in almost all cases.

  Perceived loudness and pitch of the signal are related logarithmically to the excitation strength and frequency, respectively [4, ch. 3]. The possible variation of excitation strength (corresponding to the signal power) is extreme: amplified rock music listened to from close by may be 120 dB above the threshold of hearing (that is, having 10^{12} times more power than the threshold), even though such levels are considered unsafe due to potential hearing damage [3, ch. 2]. In practice, this means that the noise floor of audio signal processing circuits must be very low. Audio CD format features signal-to-noise ratio (SNR) above 90 dB and increasing the resolution (and SNR, unless the sound system is limited otherwise) does not lead to appreciable differences [5]. Thus, ideally, the synthesizer should reach such SNR levels, although this is tricky to accomplish in practice.

  A periodic sound with multiple harmonics is perceived as having the pitch of the fundamental (first) harmonic [6, p. 1143]. The ratio of the harmonics is perceived as the timbre of the sound. For sounds with a single fundamental (e.g. a single note being played), nonlinear distortion only produces harmonics, only changing the perceived timbre. Nonlinear distortion thus is not especially problematic for synthesizers, as it becomes "part of their sound" [7, subsec. 5.1]. Obviously, this only applies in moderation.

- **Control signals.** Normal control signals typically contain a stationary component that largely determines the behaviour of the module input it is routed to, in addition to other higher-frequency components up to the upper end of audio frequency range. Additional frequency components will correspond to the module behaviour being modulated: for example, using a low-frequency oscillator (LFO) to modulate Voltage-Controlled Oscillator (VCO) pitch results in vibrato (continuous change of pitch), while using it to modulate Voltage-Controlled Amplifer (VCA) gain results in tremolo (continuous change of amplitude). An audio signal may be used as CV as well, producing potentially interesting sounds.

  The considerations for control signals are therefore similar to considerations for audio signals, with differences largely arising from the fact that the stationary component is important and cannot be discarded, while higher audio-frequency components may not be as important. The stationary component is especially important for consideration in control signals as usage of various components and circuits (especially op-amps and unmatched resistor dividers) may change the offset or gain of the stationary component considerably, resulting in reduced performance. Noise is undesirable as well, as it may result in artifacts being present in the subsequently affected audio signal.
- **Gate signals.** Gate signals are typically distinguished from normal control signals by being switched between two signal levels due to an action such as note being played [3, p. 130], conventionally acting as an input of a voltage-controlled switch or comparator. However, a gate signal can also be used as a classic control signal, so it should ideally be noise-free similarly to control signals.
- **Trigger signals.** Trigger signals are switched between two signal levels like gate signals, but only change their level momentarily as an action is performed [3, p. 130]. As I do not use any trigger signals in my modules, I will not discuss them in detail.

Ideally, all signals should be treated similarly by the modules, with frequency range of interest ranging from stationary to approx. 20 kHz, so that different types of signals can be mixed and matched for more sonic possibilities. The most important requirement is high SNR, distortion being secondary.

Offset and gain errors are of varying importance depending on the context in which they are used. VCO pitch control (as well as their handling in the VCO itself) is especially sensitive. Typically, the 1 V/Oct scheme is used: with each +1 V CV change, the VCO frequency is reduced by half, corresponding to the note played moving exactly one octave up [1, p. 38]. This scheme has many nice properties: a large amount of notes can be represented by a small range of octaves, note offsets can be "stacked" just by summing them, and symmetric vibrato can span a certain semitone fraction exactly, both up and down.

Unfortunately, when the 1 V/Oct scheme is used, even small offset and gain errors result in perceivable offsets. A pitch difference of c cents can be converted to the corresponding voltage offset as $2^{(c/1200)}-1$. For worst-case consideration, negative differences result in (very slightly) smaller voltage offset amplitude. A -5 cent difference corresponds to an approx. -2.88 mV voltage offset, while -20 cent difference corresponds to an approx. -11.49 mV voltage offset. Therefore, even relatively small offset or gain errors may affect the listening experience considerably, and pitch control signal handling should be very careful. Furthermore, as all signals may be used for pitch control, reduction of gain and offset output errors is still of importance in other modules.

2.2 Synthesis domains

For synthesizer implementation, the most important consideration is how the audio signal is generated. Considering only purely electronic synthesizers (in contrast to Hammond organs or Rhodes and Wurlitzer electric pianos, all of which feature mechanical components in addition to electronics), there are two very distinct approaches: analog and digital synthesis, each with their advantages and drawbacks.

- **Analog synthesis.** The signal is generated or changed using traditional circuit components such as passives, diodes, transistors, operational amplifiers (op-amps), manual or voltage-controlled switches, etc., without any digital or computer-based components used [8]. Analog synthesis has a distinct advantage over its digital counterpart in the fact that the signal has direct correspondence to the produced sound. This allows for easy interaction between distinct modules and straightforward usage of electrical properties (e.g. two signals can be mixed by connecting them together). The absence of analog-digital conversion circuitry or digital interfaces contributes to lower initial cost of any separate module.

  Unfortunately, electronic components exhibit many non-ideal characteristics, which may contribute to undesirable artifacts such as control voltage and signal offset and gain errors, low noise floor, distortion, oscillator instability and drift, etc. Careful design and component selection is required. Unfortunately, while manufacturing techniques have improved continually, many specialized analog audio ICs have been supplanted by their digital counterparts [9].
- **Digital synthesis.** Theoretically, using digital synthesis, signals can be constructed and transformed perfectly in the digital domain. However, the signals must be converted between the digital and analog worlds by the means of analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) [10, p. 60-63]. Performance in the analog domain is then governed by the parameters of the converter chip. Together with the cost of the processing unit and required support circuitry, this results in significant initial cost of a digital module.

  The actual performance of digital synthesizers is also a bit more complicated in practice. In addition to latency incurred during conversion between the domains [10, p. 60-63], *aliasing* artifacts may be produced. For mitigation of such artifacts, higher digital processing capabilities are required [11].

In my case, the choice of a digital sequencer was both required by the assignment and obvious, as the sheer amount of control available is unmatched by analog techniques. For other modules, I decided to use a purely analog signal path, to demonstrate classic analog design techniques and reduce module cost. In the VCO module, I also added digital trimming circuitry to alleviate oscillator tuning problems.

2.3 Synthesis types

There are many possibilities for production of sounds using a synthesizer, and correspondingly a number of different categorizations. One of the popular categorizations is by the method complex waveforms are attained. I have compiled a non-exhaustive list:

- **Subtractive synthesis.** A complex waveform containing multiple harmonics (e.g. triangle, ramp or pulse wave) is produced by the oscillator, after which various frequencies are attenuated by filters to achieve the desired timbre. This type of synthesis is typical for Moog-style "East Coast" synthesizers [12].

- **Waveshaping.** The initial waveform is transformed using a waveshaper circuit, resulting in a different timbre. This type of synthesis is typical for Buchla-style "West Coast" synthesizers [13].
- **Additive synthesis.** The initial waveforms are multiple sine waves or similar waves with low amplitude of higher harmonics and the desired effect is achieved by adding them together. This type of synthesis is famously used in electromechanical Hammond organs, transistor organs, and their emulations.
- **Frequency modulation (FM).** The frequency of the oscillator is modulated by another oscillator, the oscillator frequencies usually being harmonically related. This is problematic in the analog domain due to oscillator inaccuracies. The most famous use of FM synthesis is in digital instruments, as originally devised by John Chowning [3, ch. 2].

All of the listed methods can be (and are) combined in practice. For my own synthesizer, I opted to use traditional subtractive synthesis, which is easy to understand, well-suited for analog implementation, and the audio signals produced sound interesting even without many modules used. The traditional implementation of subtractive synthesis features one or more voltage-controlled oscillators, filters, and amplifiers, as well as one or more Attack-Decay-Sustain-Release (ADSR) envelope generators (EGs). For my implementation, I opted to build one of each of these modules, only replacing the ADSR EG with a simpler Attack-Release (AR) EG.

2.4 Modular synthesizer formats

A modular synthesizer format defines parameters needed for physical and electrical compatibility between various modules, to allow combination of modules from different manufacturers. There are three major categories of formats [14]:

- **Small format.** This category prominently includes the Eurorack format. Small knobs, jacks etc. are used, which increases value compared to large formats but may be uncomfortable to use for some users [14].
- **Large format.** This category includes the Dotcom and MOTM formats inspired by the old Moog modular format, with 1/4 inch jacks used for interconnections and ± 15 V dual supply voltages, as well as the Modcan-A format where banana jacks are used [14, 15].
- **Brand-unique.** The other formats are brand-unique, resulting in incompatibility with modules from other brands.

Since the Eurorack format is dominant and cheaper than the large formats [14], I decided to make my modules electrically compatible with it. The Eurorack electrical specification is available in [2]. In particular, two-row 16-pin or 10-pin headers and corresponding Insulation-Displacement Contact (IDC) connectors are used as power connectors. Dual-supply ± 12 voltages are provided, i.e. $+12$ V, -12 V, and 0 V rails are carried.

For signal inputs and outputs, mono 3.5mm phone jacks are used. The usual signal voltage ranges are [2]:

- **Audio signal.** $-5..+5$ V (10 Vpp).
- **Control voltage.** $-2.5..+2.5$ V (5 Vpp) for low-frequency oscillators, and $0..+8$ V for envelope generators.
- **Gate.** 0 V for the low state, 5 V for the high state. Triggers are normally triggered on rising edge.

Of course, in practice, the signal voltage may be anywhere between the rails. For my purposes, I chose to largely follow the specification, although I used 0..+5 V nominal range for envelope generators so that gates could be used in the voltage-controlled amplifier (VCA) as control voltages, for triggering between zero and unity gain.

2.5 Part selection

From Sections 2.1 and 2.4, it is apparent that the choice of parts may have a significant impact on the performance of the synthesizer. Therefore, in this section, I list some common choices that have to be made.

Passives. Good passive part tolerance is essential for prevention of voltage-divider offset and gain errors. Fortunately, 1% tolerance resistors can be sourced easily enough nowadays for prices similar to the 5% tolerance resistors. Lower-tolerance resistors are available with higher pricing, but there are few places where they would offer a significant advantage. Therefore, 1% tolerance resistors can be used almost universally. For capacitors, the role of the part in the circuit is important. Most commonly used filtering capacitors can be bought with high tolerances.

Operational amplifiers. The most important parts of the synthesizer are definitely op-amps, as they are used ubiquitously and a poor op-amp selection can ruin the characteristics of the synthesizer. Unlike ideal op-amps, real op-amps have several important parameters that must be considered, as well as many less important ones (depending on the application).

In practice, for Eurorack-style designs, the TL07x series of JFET input stage op-amps is used almost universally, due to their low cost, high availability, and relatively good parameters. Crucially, the slew rate of TL07x is plenty enough for all audio-frequency signals, preventing their distortion due to slew rate limits. Furthermore, in JFET input stage op-amps, almost no current flows through the inputs [16], so there are no current flow considerations when designing circuits using these op-amps. The harmonic distortion is also a non-issue for synthesizer design.

The typical equivalent input noise voltage of TL07x between 10 Hz and 10 kHz is 4 uV (-108 dBV) at 25 degrees Celsius [16], almost the same as the Johnson noise of a 100 kOhm resistor considering the same temperature and bandwidth, which seems to be the reason why the largest signal-path resistors present in modular synthesizer circuits are typically 100 kOhm. While reasonable by itself, as the signals are attenuated by factors of up to 1000 (-60 dB gain) in some circuits, the noise floor can rise significantly.

Additionally, there is a significant input offset voltage in TL07x, 3 mV typical and 10 mV maximum (varying slightly between models) [16]. This is especially problematic when considering offset-induced pitch errors discussed in Section 2.1. Furthermore, TL07x is susceptible to *phase reversal*, where the polarity of the output voltage changes when one of the op-amp inputs is near the negative rail [17]. However, op-amps with better parameters are typically sold for more than double the price of TL07x and there are sourcing issues, making them infeasible for low-cost synthesizer design.

Voltage-controlled switches. Their selection is easy, as the amount of available "high-voltage" switches for ±12 V dual-supply voltages or higher is largely limited to some switches from the Vishay DG family. However, the switches are costly and sourcing the exact type in the exact package needed is still problematic, so they are best used sparingly.

Chapter 3
Typical subtractive synthesizer modules

As disscused in Chapter 2, I decided to create a largely analog modular subtractive synthesizer. In this chapter, I will detail the various modules and common circuit techniques used in construction of such synthesizers.

Unfortunately, there does not seem to be a commonly available textbook-style handbook for synthesizer module design, so I was forced to piece together the commonly used circuits from various sources I found on the Internet. This seems to stem from the fact that synthesizer construction is done either by professionals who do not share their techniques, or by hobbyists. Academic interest in the matter is limited.

3.1 Power supply

While not a "module" in the sense of audio or control signal generator / processor, the power supply is the most important part of a modular synthesizer. After all, the synthesizer does not work without it.

For Eurorack-style modules, ±12 V dual-supply voltages are used, that is, there are three voltage rails in total: +12 V, ground (0 V), and −12 V. This allows for signals to be centered around 0 V, simplifying circuits. However, providing the three rails is not easy, since classic wall power adapters provide only two rails. There are three popular ways to resolve this problem:

- Directly building a converter from the mains voltage to the modular system voltages. This presents a large amount of engineering and formal difficulty due to the dangerous voltage levels involved, but offers the best energy conversion efficiency. This method is used in commercial products, e.g. the Doepfer power supply A-100PSU3 [18].
- Using a wall power adapter with a DC output, generating the negative voltage using a charge pump IC. This method typically has low maximum output current, which limits its use to settings where only low current is needed, such as in the Klon Centaur guitar pedal [19].
- Using a wall power adapter with an AC output, performing rectification to ±12 V in the power module. This method is preferred by hobbyists.

The direct power converter construction was clearly not desirable for my goals. Negative voltage conversion is possible, but inadvisable for usage with modular synthesizers where current requirements are high. The AC power rectification approach seems to be the best for non-professional use-cases.

The traditional power module rectification circuit for Eurorack-style modular synthesizers is visualised in Figure 3.1. The 12 V AC inputs are galvanically isolated from the mains voltage, so there is no connector polarity issue with the barrel connector: either the outside or inside contact may be used as a ground reference. The other one serves as a 12 V AC input. The input is half-wave rectified by diodes D1 and D2. Large electrolytic output capacitors C3 and C4 are used for smoothing, the small bypass capacitors C1 and C2 ensuring stability. After the voltage has been regulated by the LM7812 and LM7912 linear regulators, output smoothing is performed by C5 and C6. The LEDs D5 and D6 provide information about the state of the power supply and bleed the capacitors when the power module is turned off. The diodes D3, D4, D7, and D8 are used for protection purposes [20, sec. 2a].

The half-wave rectifier power source seems to be popular for many reasons: it is inherently safe due to the low voltages involved (mains-level conversion being done in a certified 12 V AC power adapter), the component cost is low, overcurrent and thermal protections are present in the LM7812 and LM7912 regulators, and the output is reasonably suited for low-cost audio-level rails. The major disadvantages are halving of maximum output power compared to the power adapter, low efficiency of the linear regulators and consequent thermal dissipation, requiring heatsinks if large output current is required [20, sec. 2a].

I built the half-wave rectifier power module for my synthesizer (as will be further detailed in Section 5.1), but the evaluation results in Chapter 8 suggest that mains harmonics bleed into module output signals, unacceptably lowering SNR. Therefore, proper mains harmonics filtering should be a major consideration when building a synthesizer.

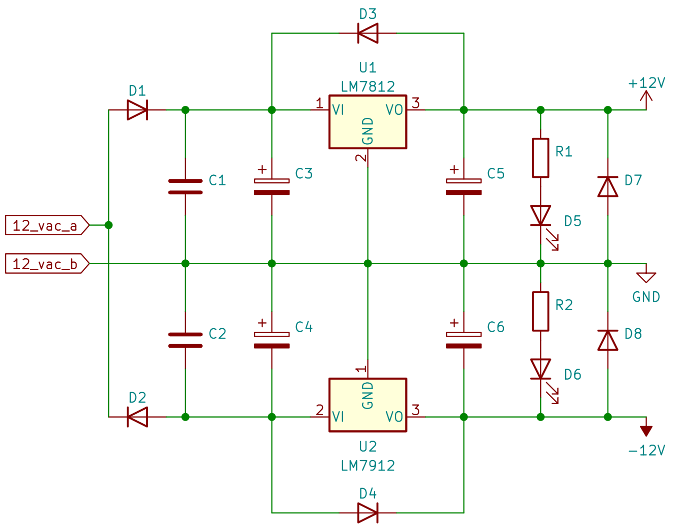

Figure 3.1. Simplified synthesizer power module half-wave rectifier schematic, remade schematic from [20, sec. 2a].

3.2 Control module

The user should be able to control what is being played by the synthesizer. Many various methods have been devised, their development being an integral part of synthesizer history [3, ch. 1]. In conservative designs, keyboards similar to piano and organ keyboards are used, sometimes joined by other note pitch input methods (e.g. ribbon controller) or other parameter inputs (e.g. modulation wheel).

Another popular design of control module, typical for electronic music, is a sequencer, into which a sequence of notes (and other control inputs) can be programmed. This sequence is then played repeatedly, which allows the user to manipulate other parameters in the meantime [3, p. 172]. The major adavantage of sequencer-based control modules is that there are no major limitations to the form factor or mode of user input. Therefore, I will not discuss them further in this section. The details of my implementation are discussed in Section 5.2.

3.3 Voltage-Controlled Oscillator (VCO)

A voltage-controlled oscillator is the heart of subtractive synthesizers. Typically, the pitch is controlled using 1 V/Oct standard CV, i.e. the frequency of the output waveform is doubled (corresponding to one octave up) when the CV increases by 1 volt. The frequency corresponding to 0 V is typically not fixed in classic designs, allowing the user to tune it to arbitrary frequency. This presents some difficulty when playing together with other instruments.

The main problem with 1 V/Oct CVs is that they must be converted to voltage or current that corresponds linearly to the signal period before the result can be integrated. Changing the direction of integration or resetting the integrator output after reaching a predetermined output value results in an oscillating waveform at the VCO output, and may be designed using standard op-amp and switch techniques.

A classic matched pair BJT exponential converter circuit is studied in many sources, such as [21, 22, 23]. In particular, [23] gives a thorough examination of the "basic exponential core", which almost exactly corresponds to Figure 3.2.

In short, the operational amplifier U1B acts as a current source: the amplifier operates in a feedback configuration through R3 and Q2. Due to the feedback configuration, there is a virtual ground on its negative input. Since no current flows into the negative input of U1B, the reference current through R4 must be balanced by current of equal magnitude and inverse direction through Q2. Therefore, Q2 collector current is constant.

Supposing that the transistor Q1 is held at ground and considering Ebers-Moll-style model where BJTs are voltage-controlled devices and collector current depends on base-emitter voltage exponentially, Q1 and Q2 collector currents must be the same as the base-emitter voltages are the same. This should largely cancel a device-specific current-magnitude constant [23, p. 1-2]. The remaining undesirable constant which impacts V/Oct scaling is compensated by the thermistor R2.

After the undesirable scaling constants are compensated, a small signal applied to the base of Q1 should result in the emitter current of Q1 being almost exactly exponentially dependent on it. Of course, this all depends on proper component matching. In practice, some temperature and CV amplitude dependence will be retained no matter how well the circuit is designed and built.

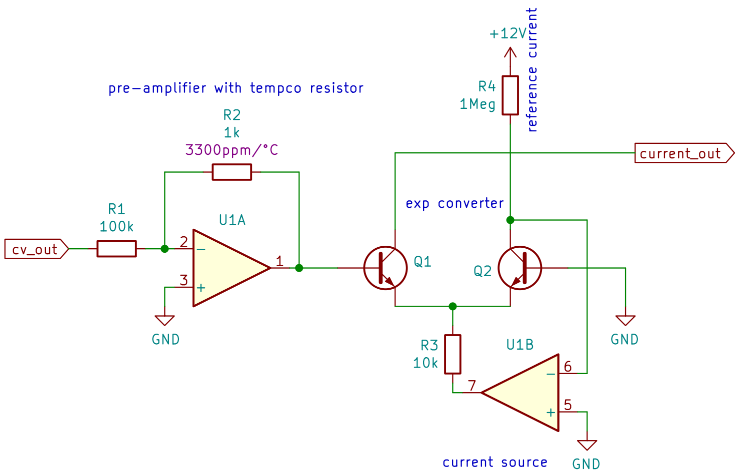

Figure 3.2. Simplified schematic of an exponential converter core.

3.4 Voltage-Controlled Amplifier (VCA)

A Voltage-Controlled Amplifier amplifies an incoming (typically audio) signal by gain determined by its Control Voltage (CV) input [3, p. 151]. In essence, it is a special analog multiplier with an audio signal as its one input and the control signal as the other one. The exact function of gain determined by CV might not be linear, and will determine the characteristics of the resulting sound.

There are various techniques for construction of analog VCAs which differ in various parameters [7]. Crucially, they may introduce unwanted nonlinear distortion, and the maximum attenuation may be too low, so the amplified signal never "turns off" completely. While usage in synthesizers is not as demanding as in high-fidelity audio, as the non-ideal parameters become the part of the sound of the synthesizer, care must be taken.

For a general-purpose synthetizer VCA for educational purposes, the requirements are different than for compressor/limiter VCAs chiefly considered in [7]. Thus, my evaluation of various techniques differs from [7] in some cases:

- Blackmer circuit VCAs provide excellent performance, but due to the amount of well-matched parts needed, only IC-based implementations are reasonable. This leads to sourcing and cost issues and is not very educational, as the circuit functionality is hidden inside the IC. Output gain is exponentially related to input CV.
- Analog multipliers can provide excellent performance as well if the proper part for audio use is chosen. Their output gain is linearly related to input CV. The cost, sourcing and functionality-hiding issues are similar as for the Blackmer circuit.
- Long-tailed pair VCAs provide relatively good performance with tradeoffs between distortion and noise, as input signals must be heavily attenuated for linearity purposes. Only transistors, op-amps, and passives are required for construction.
- VCAs based on an Operational Transconductance Amplifier (OTA) are essentially long-tailed pair VCAs integrated inside the OTA. Parameters may be better than for a discrete long-tailed pair, but only a few types of OTAs with mediocre parameters are widely available today. The functionality is hidden inside the IC.
- Optocoupler-based implementations are infeasible for use as general-purpose VCAs due to high recovery times.
- Single-transistor VCAs are problematic due to temperature concerns.
- Diode-based VCAs (conceptually similar to a diode-ladder VCF circuit that is explored in Section 3.6) cannot be used due to low maximum attenuation, in addition to other problems.
- Switch-based VCAs require careful selection of components and are against the spirit of analog signal path, providing little insight into classic techniques.
- Digital VCAs are against the spirit of analog signal path, in addition to latency problems.

It is apparent that the only reasonable classic solution that does not offload the problem to ICs is the long-tailed pair-based solution.

The long-tailed pair core in a cheap "resistor-based" form with a differential amplifier on the outputs is visualised in Figure 3.3. While the long-tailed pair circuit is described in many sources, the conceptual idea of gain change via CV manipulation of the long tail voltage is problematic to understand. I will try to do so concisely.

For simplicity, one can consider the resistor R3 as a current source. This allows for a simplification where the emitter voltage V_E is fixed to approx. -0.5 V (depending on the transistor pair used) compared to ground. The approximation works well if the resistor R3 is very large (approx. 100 kΩ in practice) and large positive CV (over 1 V) is applied.

If the transistors Q1 and Q2 are matched and have approximately the same base voltage, the current flowing through R3 is approx. $\frac{V_{\mathrm{CV}}-0.5}{R_3}$. This current is split between the transistors Q1 and Q2. An imbalance caused by a small signal going into the base of Q1 will result in a difference between the emitter currents of Q1 and Q2. As emitter voltages must remain the same due to the feedback action of the differential amplifier, the current difference will be transformed into voltage difference on the output of U1A. Thus, the small signal will be almost unchanged, but scaled proportionally to $V_{\mathrm{CV}} - 0.5$.

Of course, the discussed view only holds for larger positive CV voltages. If the CV is negative, there is almost no current flowing through the transistors and the VCA has almost zero gain. Between these cases, there is a soft knee due to the actual exponential dependence of BJT collector current on V_{BE}.

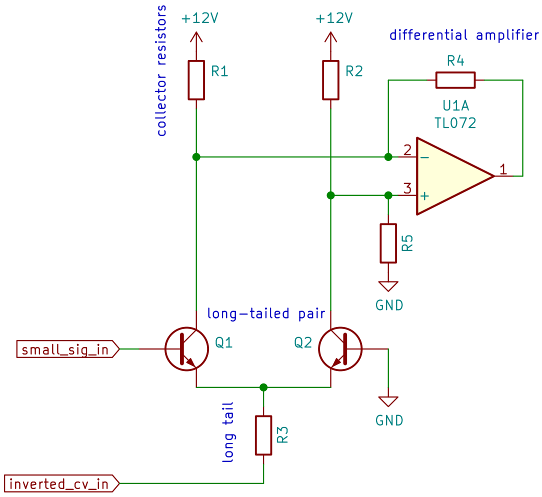

Figure 3.3. Schematic of a long-tailed pair Voltage Controlled Amplifer core. Resistors R1 and R2 have the same resistance R_{c}. Resistors R4 and R5 have the same resistance R_{a}.

3.5 Envelope Generator (EG)

An audio signal produced by a classic instrument such a piano can be considered to be composed of two parts: its audio-frequency waveform that determines the pitch and timbre, and its *envelope* which contains slow changes of signal amplitude. The piano-note envelope rises sharply as the note is pressed and the hammer strikes the piano strings, after which it slowly falls as the strings lose energy. In synthesizers, such behaviour can be approximated by generating the envelope using an Envelope Generator and using it as a CV input to the VCA. Of course, different uses are also possible, for example, modulating the filter cutoff frequency [3, p. 151-152].

A classic ADSR envelope generator has four control parameters: attack, decay, sustain, and release. When the note is pressed (or, more abstractly, the gate input of ADSR is turned on), the EG output rises to the maximum value during attack phase, then falls to the sustain level during the decay phase. When the note is released (the gate input is turned off), the EG output falls to zero during the release phase [3, p. 151-152]. Using this scheme, many interesting sounds can be generated.

A simpler envelope generator is the Attack-Release Envelope Generator (AR EG), where there are only the attack and release parameters available. During the sustain phase, the envelope remains at its maximum.

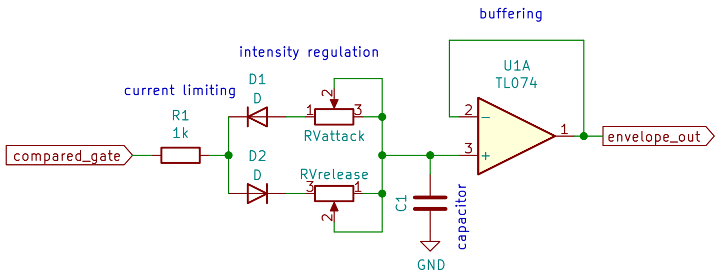

Figure 3.4. Simplified schematic of low-pass-filter Attack-Release Envelope Generator from [24].

In Figure 3.4, the core of a very simple AR EG is shown. The gate input is first rescaled to full rail range (e.g. by an op-amp in comparator configuration). It is then fed to a low-pass RC filter. Only one diode conducts, based on the rescaled gate voltage, which determines the potentiometer that sets the time constant of the RC filter.

There are some problems with such a simple EG: the signal rises to the rails, produces an exponential curve, and the potentiometers must be operated manually. As such, the intensity of attack or release cannot be manipulated automatically in interesting ways. This is not consistent with my goals.

I did not find any widely-used EG circuits with voltage-controlled intensities. Fortunately, design of such circuits is only a matter of combining traditional op-amp and switch techniques.

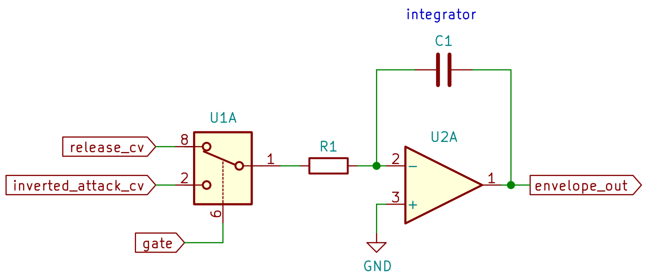

Figure 3.5. Simplified schematic of an integrator-based Attack-Release Envelope Generator with voltage control of attack and release intensity.

A simple AR EG with attack and release intensities controlled by voltage inputs can be constructed using a simple Single-Pole Double-Throw (SPDT) switch and an integration stage, as shown in Figure 3.5. Switch-based selection of the appropriate CV can also be used for creation of ADSR envelope generators [25].

Integration using an op-amp in integrator configuration produces linear slopes, which may or may not be desirable. Obviously, the problem of integration ending at rails (or, worse, with phase reversal) must be tackled, as Eurorack-style envelope voltages should be contained within 0..+5 V range. The trivial solution is to use diode clamping, however, there are diode-drop errors using such solution. I discuss my implementation, which resolves the diode-drop errors, in Section 5.5.

3.6 Voltage-Controlled Filter (VCF)

There are four main types of filters depending on the frequencies they pass through: low-pass, high-pass, band-pass, and band-reject / notch. The cutoff frequency is selected in voltage-controlled filters by the cutoff CV. The filters have a certain number of poles, which determines the rolloff slope and consequently their sonic characteristics.

In synthesizers, the most common VCF designs are two-pole (12dB/Oct attenuation) and four-pole (24dB/Oct attenuation) [3, p. 157].

Classic synthetizer VCFs also feature a resonance control. As the output of a filter is fed back to its input, a peak appears in a narrow part of the frequency spectrum, its location dependent on the cutoff frequency. This allows for interesting sounds. In some cases, the VCF may even self-oscillate using resonance [3, p. 158]. In modular synthesizers, resonance can be also achieved by routing the output of the VCF to its input after attenuating it by an attenuator or a VCA.

Typical VCFs are based on transistor or diode ladders. They are typically complex, especially commercial ones (a detailed analysis can be found in [26]), but the core idea is quite simple: the transistors or diodes are used alongside capacitors, and they function as variable resistance elements, their resistance set using large-signal cutoff CV.

In Moog-inspired filters, a long-tailed pair stage is used together with transistor/diode stages [26], similarly to the long-tailed pair VCA discussed previously in Section 3.4. Unfortunately, these filters are rather complex and many parts are required. However, a very simple diode ladder filter was implemented in [27], operation of which is easy to understand even with the additional complexity of frequency-dependent gain.

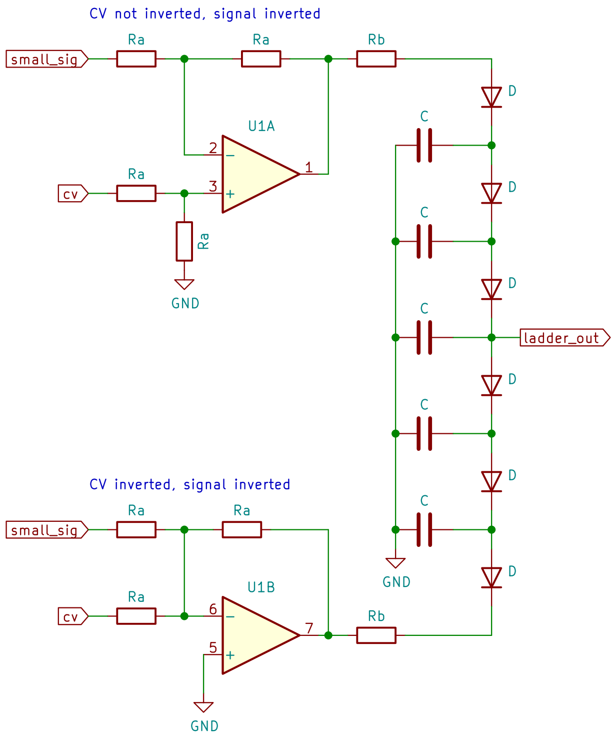

Figure 3.6. Simplified schematic of the core of the VCF from [27]. All components with the same reference in this schematic are assumed to have the same value. However, in theory, if symmetry is retained, various capacitor values and diode types could be used to obtain a different frequency response shape.

The core of the VCF is visualised in Figure 3.6. First, small-amplitude signal is superimposed onto large-amplitude CV using operational amplifiers. U1A is in unity-gain differential configuration, inverting the small signal but not the cutoff CV. U1B, on the other hand, is in unity-gain inverting summing amplifier configuration, inverting both the small signal and the cutoff CV. The outputs are sent to a ladder comprised of diodes and capacitors. As the ladder output is balanced, the CV should cancel out, leaving only the appropriately modified small-amplitude signal at the output. The resistors Rb are present to avoid excessive currents when the CV is large.

The effect of the ladder on small signals can be seen clearly by considering the diodes as resistive elements, the values of which are set by the cutoff CV. It is then immediately apparent that a low-pass filter is formed. The desired behaviour of the filter is best designed by simulation as direct analysis is rather unwieldy. However, from the topologic point of view, the circuit is certainly very elegant.

Chapter 4
Implementation choices common to multiple modules

In this chapter, I detail the choices I made for the design and implementation of my synthetizer that are common to multiple modules. In general, I aimed to make the synthesizer cost-effective, but opted for conservative solutions where I suspected potential problems.

4.1 Form factor

The fundamental difference of my synthesizer from common Eurorack synthesizer modules was the form factor. I decided early in design phase that the form factor should be different from standard rackmount approach for easier and more cost-effective construction of modules.

I wanted to use standard small construction enclosures for each module. In this market, the Hammond Manufacturing Corporation is probably the most reputable manufacturer, with their 1590 series of aluminium alloy enclosures used ubiquitously in guitar pedals. However, aluminium is not easy to work with. Therefore, I opted for ABS enclosures from the 1551 series instead. These enclosures contain PCB mounting holes and are easy to modify using a Dremel-style rotary tool. The "front panel" of the module can be marked with self-adhesive paper to aid module type and input identification.

In practice, the 1551 series enclosures are a bit problematic, as only a few sizes are large enough to contain the module PCBs. After considering the options available, I used 80x40 mm 1551KTBU [28] for most modules in an enclosure (AR EG, VCA, VCF) and 80x80 mm 1551XTBU [29] for the VCO module which ended up bigger than others due to use of more components.

For the sequencer module and its daughterboard, I did not plan for any enclosure as I used a STM32F7308 Discovery Kit for the sequencer core which has no enclosure anyway. Due to time constraints, I also did not build a version of the power module in an enclosure.

In Figure 4.1, all non-support modules are shown. Originally, I planned to arrange the enclosed modules so that the front panel containing the jacks would face upwards and the module power connector would face backwards from the user. However, this configuration was unstable due to the small weight of the modules and the fact that the enclosure sides are not perfectly right-angled. I therefore decided to label and use the modules in the shown configuration. This results in the module inputs placed to the right of outputs, counterintuitively for Western users (although guitar pedals, unlike typical synthesizers, have the same configuration). This should be rectified in the future by reversing the order of the jacks.

Figure 4.1. Photo of all modules that are in enclosures. A standard size matchbox is included for size reference.

4.2 Printed Circuit Board (PCB) layouts and parts

To design schematics and layouts of the Printed Circuit Board (PCB) modules (all except power supply), I used the KiCad electronics design automation software. For reasonable costs and assembly complexity, I used double-layer PCBs with components placed only on the top layer. I designed the board outlines and screw holes according to the enclosure datasheets [28, 29], shortening them in comparison to datasheet maximum PCB sizes by 2.25 mm and 3.04 mm, respectively, at the side opposite to 3.5mm sockets, so the PCB could be placed into the enclosure with sockets installed before being pushed to seat the sockets properly in drilled holes. Routing was difficult, but I managed to route all modules with a combination of 0.3 mm and 0.5 mm track widths, with bottom layer ground pour.

I used CUI Devices SJ-352X-SMT sockets due to their relatively low cost, high availability and multitude of colours for easy distinguishment between inputs and outputs, as well between audio and CV/gate signals. I used pink for audio inputs, black for signal outputs, blue for CV/gate inputs, and green for CV/gate outputs. In retrospect, as pink and green are typically associated with audio inputs and outputs, it would be much better to keep that association and use blue for CV/gate inputs and black for CV/gate outputs.

I used SMT components where possible. As for footprints, I used the TL07xAC opamps in Small Outline Integrated Circuit (SOIC) package for easier soldering. Interestingly, the TL07x series is also made in similarly named Small Outline (SO) package among others, which is bigger than the SOIC package. I used the Vishay DG series switches in Thin Shrink Small Outline Package (TSSOP) due to their arguably better availability than in SOIC. As for other ICs, I used what was available in reasonably large yet cheap packages. For passives, I chose the imperial 0805 (metric 2012) form factor for cost and footprint reduction and easy soldering. As discussed in Section 8.2, due to the choices made, the modules were easy enough to solder.

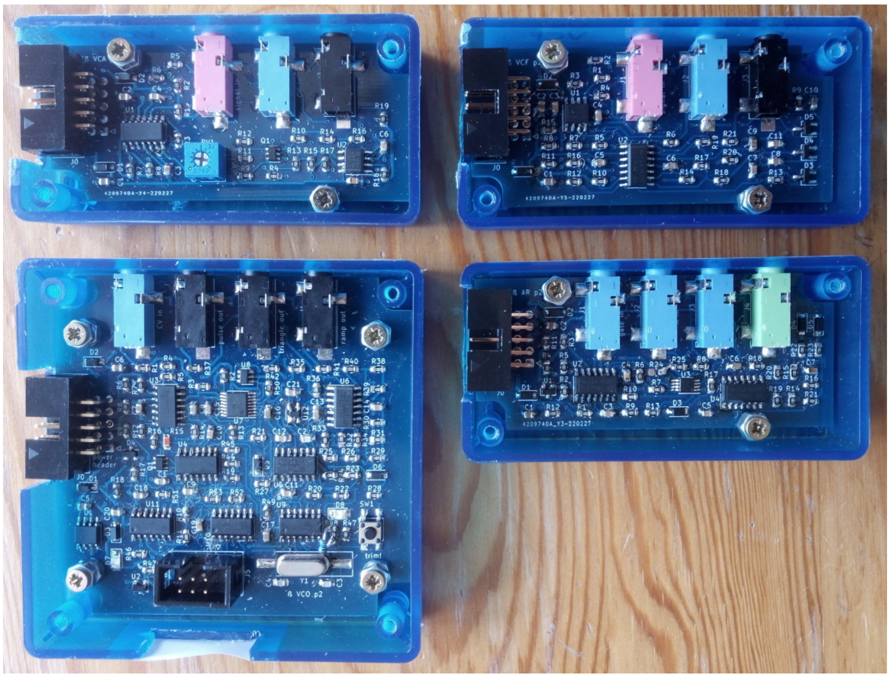

Figure 4.2. Photo of all enclosed modules with their lids removed. Clockwise from top right: VCF, AR EG, VCO, VCA. For PCB mounting, I used 2.5x10 mm self-tapping screws with two M3 nuts used as standoff due to problems with sourcing smaller and shorter screws.

4.3 Module power input

For Eurorack-style modules, +12V and −12V are conventionally used [2]. In some older Eurorack modules, a +5V rail is required to be supplied as well through a 16-pin connector. Given the ubiquity and low cost of 5V regulators, this was gradually replaced by on-module regulation from the +12V rail. As such, I opted to use the popular alternative 10-pin connector scheme for the power supply as well as modules, forgoing the +5V rail generated on the power supply (as well as system-wide CV and gate inputs).

The power connector format is essentially based on 2.54 mm male pin headers with 2 rows and 5 columns (which can also be expressed as as 2 rows and 10 positions) being present on modules. An Insulation-Displacement Contact (IDC) connector with the matching number of rows and positions can be used to connect two male headers. While bare male headers are used by Doepfer, shrouded headers are also available, typically with a notch used for keying the IDC connector so that it cannot be plugged in the opposite direction (which would result in opposite power supply polarity on the Eurorack power supply format and thus destruction of unprotected modules).

For my modules, I chose to use the 10-pin connector with protection diodes, as visualised in Figure 4.3. to reduce the amount of work, I used a schematic part from the Eurocad KiCad part library [30]. Unfortunately, the ordering of pins in the schematic is opposite from the Eurorack specification when used with a conventional KiCad header footprint, resulting in polarity reversal. In the Eurorack specification, the coloured (typically red) wire of IDC connectors that signifies position of pin 1 is specified to route the −12V rail [2].

I did not catch the polarity error until the PCBs were manufactured. Fortunately, the only change necessary was in the power module, where I reversed the wires going to the power output headers. However, in a future revision, this should definitely be changed to comply with the specification.

As for the protection diodes, I used 1N4148W diodes for very conservative and cheap protection. As a result, there is approx. 0.7 V voltage drop on the rails. In a future revision, Schottky diodes with at least 20 V peak reverse voltage or a FET-based solution should be considered for a lesser voltage drop.

As for power conditioning, I did not have enough space to add electrolytic capacitors to modules, so I used two ceramic 10 uF bypass capacitors, each between one rail and ground. I also added 0.1 uF capacitors near op-amps and other ICs as standard practice, one near each voltage rail pin for proper decoupling (with some exceptions where 2 ICs shared the capacitors to save space). The choice not to use large ($\geq$ 100 uF) electrolytic capacitors might have influenced signal contamination by mains frequency harmonics, shown and discussed in Chapter 7.

In modules where smaller-voltage rails were necessary, they were generated by common linear voltage references or regulators (depending on the output current needed), with circuits based on their respective datasheets. I will not discuss them in Chapter 5 except for where they are important.

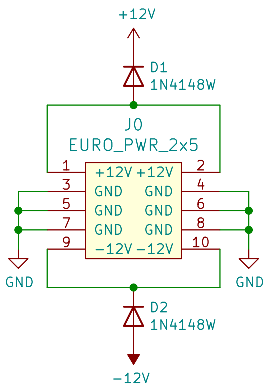

Figure 4.3. Schematic of power module input used in the synthesizer modules. Due to a design error, the module pins are incorrectly ordered, resulting in polarity reversal compared to Eurorack specification.

4.4 Signal inputs and outputs

The choice of an input stage circuit for the synthesizer modules is not trivial. While an inverting amplifier configuration is extremely common in Eurorack module schematics found on the Internet, I chose to use a non-inverting buffer for all inputs instead. I will discuss my reasoning in this section.

In Figure 4.4, a common Eurorack-style input is shown. An inverting amplifier configuration is used. There are some distinct advantages to this style of input over non-inverting amplifier configurations. The gain of the signal may be immediately reduced, reduction of signal gain being necessary for many module implementations (which can be seen in Chapter 5).

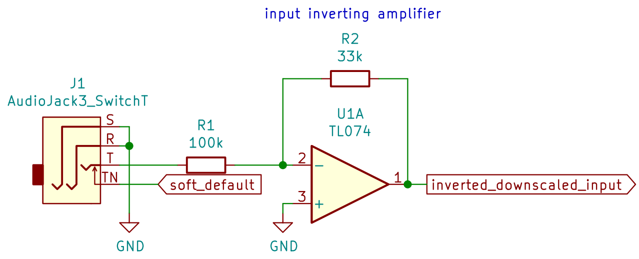

Figure 4.4. Schematic of Eurorack-style signal input.

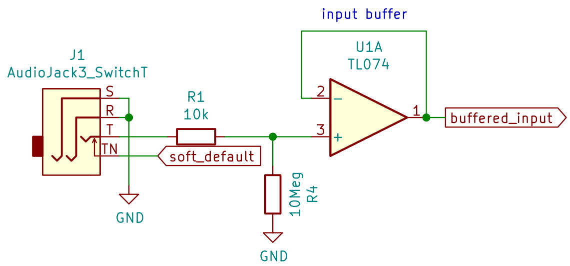

Figure 4.5. Schematic of buffered signal input used on my synthesizer modules.

The inverting amplifier configuration also effectively removes the possibility of TL07x phase inversion in normal operation: inverting input voltage always stays close to 0 V due to negative feedback. Furthermore, if there are multiple outputs connected to the input, each with the same effective impedance, their voltages are almost exactly summed, which is more useful than the voltage divider (or, more generally, a star network) arising from a non-inverting amplifier configuration.

However, there is a problem with the inverting amplifier configuration: the gain is not just $-\frac{R_2}{R_1}$, but $-\frac{R_2}{R_1+Z_{\text{input}}}$, where impedance Z_{input} will typically be given as as sum of the output resistor impedance, cable impedance and connector impedances. In Eurorack-like modules, it should typically be dominated by a 1 kΩ resistor after the output buffer (the tolerance of which is not well-specified). If R_1 is 100 kΩ, the typical value for this configuration, there is an approx. 1% gain difference from $-\frac{R_2}{R_1}$. If R_1 is smaller, the gain difference becomes larger. R_1 larger than 100 kΩ adversely affects noise, as discussed in 2.5.

In my synthesizer modules, I opted to use a buffer arrangement visualised in Figure 4.5. The gain dependence on output impedance is almost completely eliminated. The resistor R1 now serves the sole purpose of input protection against overvoltage (limiting input clamp current). I have also added the resistor R2 to reduce the noise picked up when the tip switch is not engaged but the input is floating (example, when a patch cable is inserted into the input but floating on the other end). This does not work very well in practice, however, and forms a voltage divider. In the future, I would opt to eliminate these resistors.

Looking back at the input design, I would probably not use the buffered approach on most modules as the gain dependence of the inverting amplifier configuration is not important for most uses. The gained precision would probably not be significant anywhere except for VCO V/Oct CV. The rather high input offset voltage of the TL072 also decreases precision significantly, which is not countered, but in fact worsened if more op-amps become necessary. Therefore, in my current opinion, the buffered approach would be preferable if low-offset and/or low-noise op-amps were used on the VCO CV path. The inverting amplifier approach can be used for other modules without noticeable loss of precision, reducing cost and complexity.

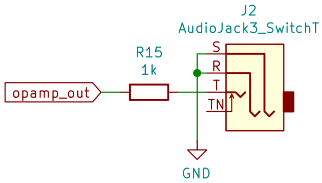

Figure 4.6. Schematic of signal output used on my synthesizer modules.

As for the input jack connections, the tip (T) carries the signal in Eurorack. For compatibility with some mono sockets that connect a TRS cable ring to the ground instead of sleeve, both the ring (R) and sleeve (S) connections should be connected to the ground if a TRS 3.5mm jack socket is used. The tip switch (TN) can used to provide the default input if no cable is connected. Since all pins of the 3.5mm jack socket can be shorted together in some situations, if the default input is not ground (to which audio signals carried on the cables can be shorted indefinitely), the corresponding rail should be connected via a current-limiting resistor.

For outputs, the situation is much easier, as visualised in Figure 4.6. As the TL07x op-amps can easily drive 1 kΩ loads (which happens in the worst case when the output is short-circuited to the ground), it is only needed to route the corresponding op-amp output to the socket tip (T) contact through the output resistor. The ring (R) and sleeve (S) are connected to the ground, and the tip switch (TN) is left unconnected. This also means that variants of the 3.5mm jack socket without the tip switch can be used for outputs, which can reduce costs and alleviate potential sourcing issues.

Chapter 5
Implementation of synthesizer modules

In this chapter, I outline my synthesizer module implementation subject to considered goals and constraints. For purposes of conciseness, only important sections of schematic design are shown and discussed. Printed Circuit Board (PCB) layouts are not discussed due to their design similarity. Support boards (two power connector hubs and a potentiometer input board) are also not discussed due to their triviality. Full schematics and layouts are attached to this thesis, as discussed in Appendix B.

5.1 Power supply

As discussed in Section 3.1, there are two main options for building a ± 12 dual-supply power module without unsafe voltages involved: either using a DC power adapter and generating the negative voltage using a charge pump IC, or using an AC power adapter and rectifying the AC voltage to negative and positive DC rails.

At first, I experimented with the former method, using a +12VDC barrel jack power adapter and generating the negative supply rail using a charge pump IC TC7662ACPA. Unfortunately, this method quickly became unusable as more modules were added, resulting in unwanted 10 mVpp signals at approx. 5 kHz in signal outputs, unacceptable for listening experience.

Supply current requirements of the synthesizer are significant, especially due to the TL07x op-amps used, which have a JFET differential pair at their inputs but the rest of their architecture is based on BJT transistors. This results in 1.4 mA typical and 2.5 mA maximum supply current per each op-amp channel. Since each module features 4 to 12 op-amp channels, supply current requirements rise quickly as more modules are added.

After deliberation, I built a half-wave rectifier power supply module instead, visualised in Figure 5.1, which is used in the final synthesizer. Specifically, I used the circuit from [20, sec. 2a] which features required capacitors and protections some of other variations of this circuit on the Internet lack.

As for component values, large capacitors are needed for input and output smoothing. I used 2200 uF electrolytic input smoothing capacitors ([20, sec. 2a] suggests 4700 uF, but these were too large for my perfboard), and 100 uF output smoothing capacitors, as suggested by [20, sec. 2a].

As the prototype of the half-wave rectifier power supply built on breadboard was largely sufficient for my use, I did not design a PCB layout and instead focused on other modules. In the future, I am planning to design a PCB for it as well. Contrary to other modules, I expect thermal dissipation to be a major concern, as improper thermal design could severely limit the maximum output current and, consequently, the maximum size of the modular system.

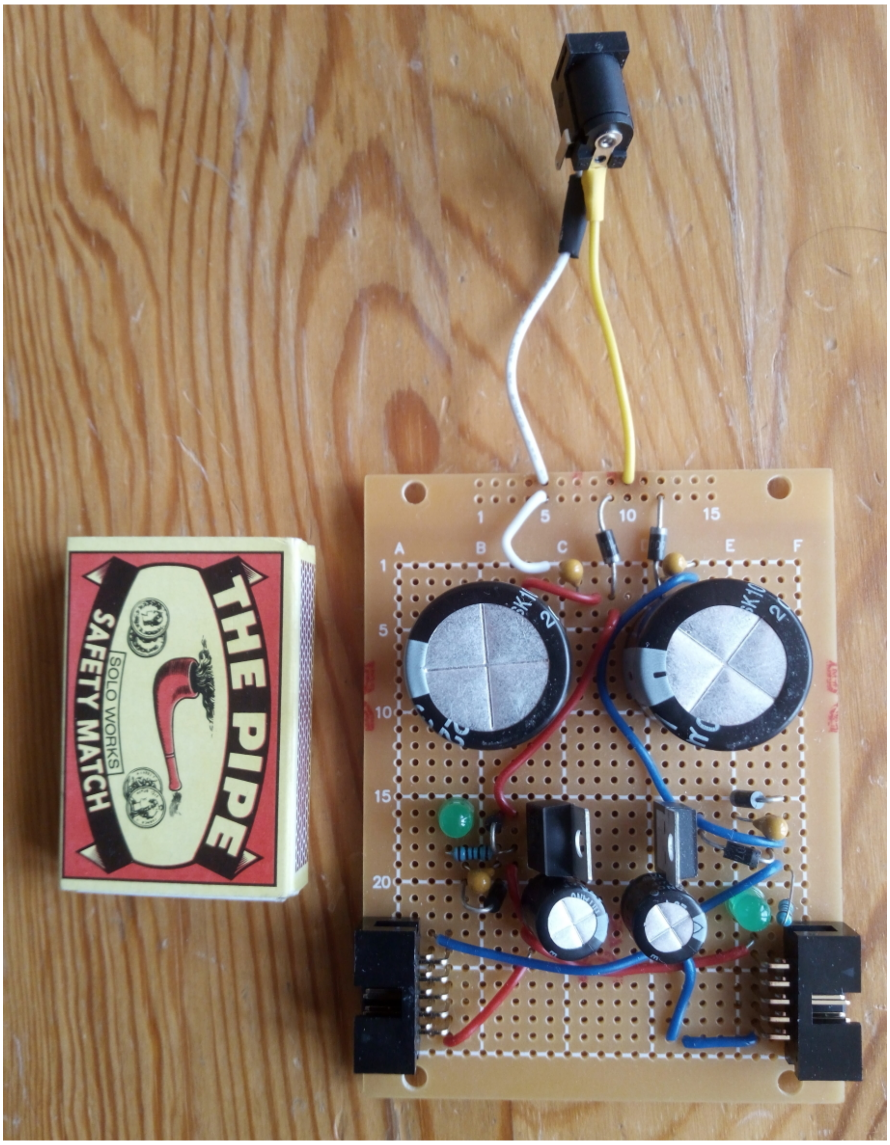

Figure 5.1. Power supply module used in the synthesizer. Module power connector polarities are reversed compared to the Eurorack specification, as discussed in Section 4.3.

Unfortunately, during objective evaluation of the synthesizer, I found that there is some mains frequency harmonics bleeding into module output signals, as discussed in Chapter 8. Better design of the power distribution system including the power module and power rail smoothing and/or filtering in other modules should alleviate these problems in the future.

5.2 Sequencer module

Designing a digital sequencer module from scratch for purposes of this thesis would be time-prohibitive. Instead, I decided to use the STM32F7308 Discovery Kit as a basis for the module. Importantly, the kit includes an STM32 microcontroller with ARM Cortex M7 core, a WM8994ECS/R codec with a stereo microphone input and a headphone output connected to separate 3.5mm jacks, and a 240x240 pixel LCD touchscreen. While it is no longer in production, it was replaced by the STM32F723E Discovery Kit which includes a similar microcontroller with more Flash memory and seems to be identical otherwise. The cost is very reasonable, approx. 40 USD per unit.

I had to solve two problems, however. First, the kit does not feature much input capability (only a touchscreen and a single user button). Second, the codec cannot output Eurorack-level 10 Vpp signals and only two audio output channels are severely limiting as pitch control CV output and note-on gate output leave no channels for other use. To resolve these problems, I designed a daughterboard for the kit.

In Figure 5.2, the sequencer module is shown. The daughterboard is mainly responsible for enhancing the user input / output (I/O) capabilities and analog voltage level translation from and to Eurorack level. I also added ground and polyfused ±12 V headers to the daughterboard to make it easier to make a breadboarded module, but I realized the danger of inadvertently touching the sequencer or a sensitive part of the daughterboard with a power rail jumper wire, so I decided against emplacing the headers and polyfuses.

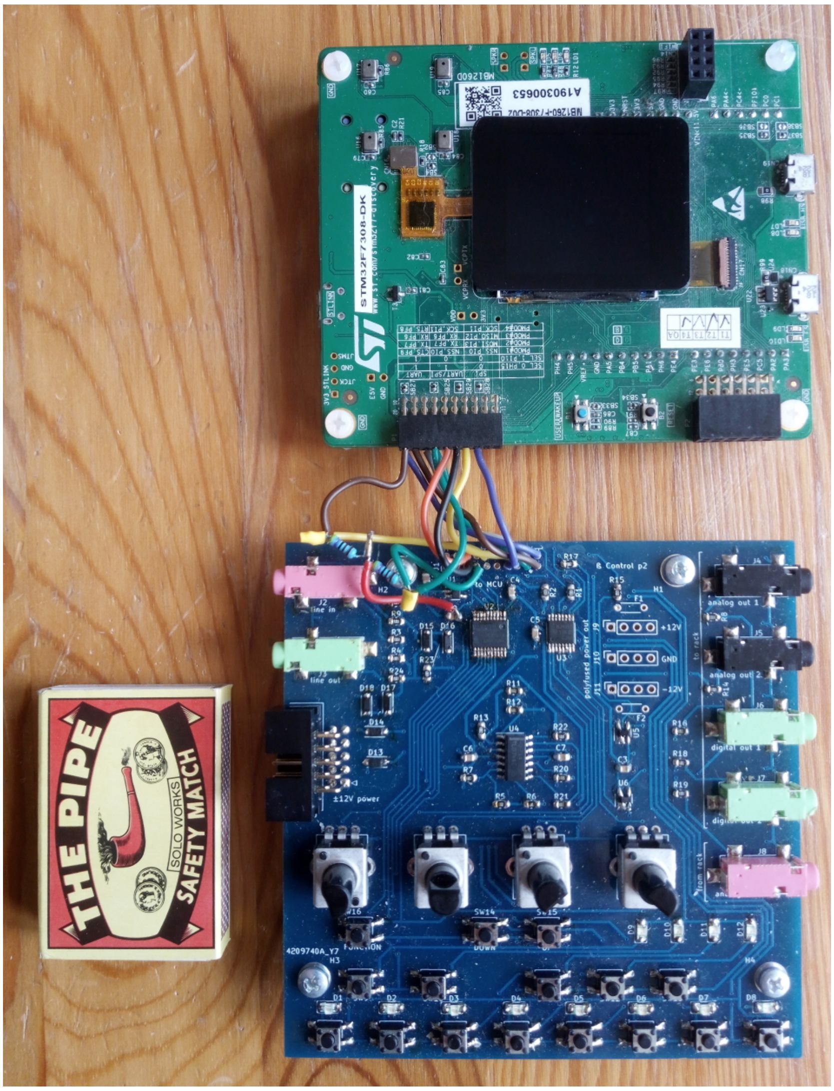

Figure 5.2. Sequencer module, comprising of the STM32F7308 Discovery Kit with a custom daughterboard. A standard size matchbox is included for size reference. The daughterboard is connected to the Discovery Kit manually using wires instead of a proper connector due to a design error discussed in Section 5.2.

Unfortunately, while I correctly connected the STMod+ connector pins in KiCad schema using a schematic symbol from the KiCad standard library, the only matching standard library connector footprint had different pin numbering. I did not realize this until the PCBs were made, forcing me to route the connections manually using wires. I also forgot to include I2C pullup resistors in the schematic, so I added them manually as well. Other than that, there were no significant problems.

5.2.1 User input and output

I had to augment the input capabilities of the STM32F7308 Discovery Kit, since touchscreen-only input would not provide suitable musical control experience. Therefore, I included 16 switches and 4 potentiometers on the daughterboard. I also added 12 LED diodes for output, although the Discovery Kit LCD display would remain as the main user interface output due to its capability to display significant amounts of information. As can be seen in Figure 5.2, 13 switches are strategically placed so that they form a piano-like one-octave-and-note keyboard. The functions of the inputs and outputs are defined by sequencer firmware and are thus discussed in Section 6.1.

In Figure 5.3, the I/O component schematic is shown. As there are 32 I/O components in total and the STMod+ connector used to connect the daughterboard to the Discovery Kit has only 20 pins, 4 of which are used for power, I arranged the LEDs and switches into a matrix configuration for multiplexing. The potentiometers are set up in a voltage divider configuration.

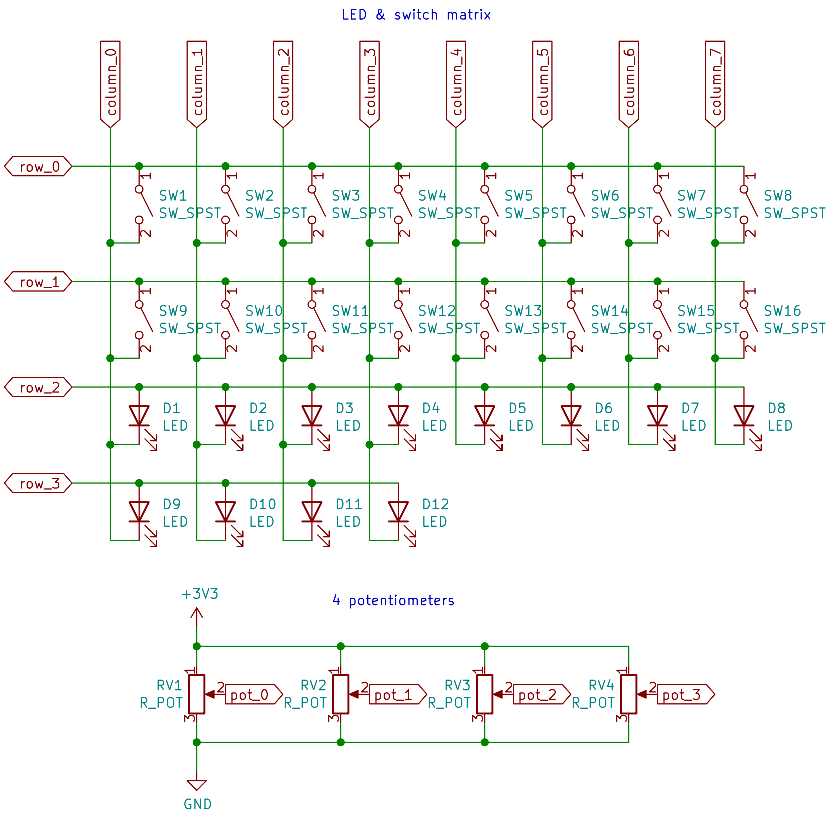

Figure 5.3. Sequencer daughterboard schematic, part 1: User I/O components.

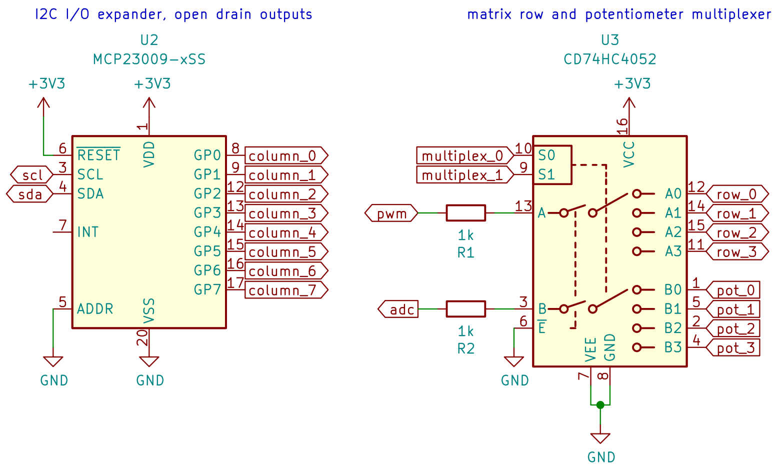

Figure 5.4. Sequencer daughterboard schematic, part 2: User I/O expander and multiplexer.

Multiplexing is facilitated by a digital I/O expander on the column side and an analog multiplexer on the row side, as shown in Figure 5.4. The multiplexer is a simple 74HC4052 IC, its selection inputs and multiplexed pins `pwm` and `adc` connected to the STM32F7308 microcontroller via the STMod+ connector.

For potentiometer value measurement, the appropriate potentiometer is selected using the muliplexer, after which voltage on the `adc` pin can be measured by the Analog-to-Digital Converter (ADC) peripheral on the STM32F7308 microcontroller. The resistor R2 is used to avoid potential damage if the `adc` pin is mistakenly configured as microcontroller output.

The switch / LED matrix is controlled by the `pwm` pin from the row side and a MCP23009 expander from the column side. The expander communicates with the microcontroller using `scl` and `sda` I2C interface pins. These pins also must be pulled up to +3V3 by approx. 10 kΩ resistors, something that I forgot about during PCB design, forcing me to later add them manually.

After selecting the row, pulling the `pwm` pin low, and setting the expander to input mode with pullups enabled, the status of all switches in the row can be read from the expander. Similarly, each LED can be lit by selecting the row, setting the appropriate column as open-drain output pulled low, and controlling the `pwm` output using Pulse Width Modulation (PWM), with duty cycle set according to the desired relative brightness. Using Time-Division Multiplexing (TDM), all of the inputs and outputs can be controlled almost simultaneously, as discussed in Section 6.1.

5.2.2 Voltage level translation and digital outputs

In addition to user I/O, the other responsibility of the daughterboard is translation of codec output voltage levels to Eurorack levels, as well as translation of the final sequencer output signal to line level so it can be recorded by a standard audio interface or played through active speakers safely and without clipping. In addition, I also implemented two digital gate outputs controlled by the STM32F7308 microcontroller so that note-on gate can be sent through one of these outputs, freeing a valuable analog output channel for a more interesting output instead.

The WM8994 amplifier has its headphone outputs, connected to a stereo 3.5mm jack on the STM32F7308 Discovery Kit, ground-referenced using a charge pump [31, p. 143]. Therefore, there is no need for an offset. DC component is not suppressed on the codec (unless such an option is selected in codec registers) nor on the Discovery Kit. Therefore, both audio and control signals can be output. Headphone output levels are not well-specified in [31]. According to my measurements, the codec outputs 2.7 Vpp signals at most with internal Programmable Gain Amplifier (PGA) 0 dB gain selected. With maximum +6 dB gain, the signals are at most 5.4 Vpp, which is still insufficient for Eurorack levels. As Eurorack levels specify 10 Vpp reference signal level, I chose to apply a gain of 3 to be sure. In Figure 5.5, the circuit is shown, comprising of a single non-inverting amplifier for each channel. In addition, while the headphone output is stereo, the Eurorack-style modules accept mono inputs, so each channel is output on the tip of its own jack.

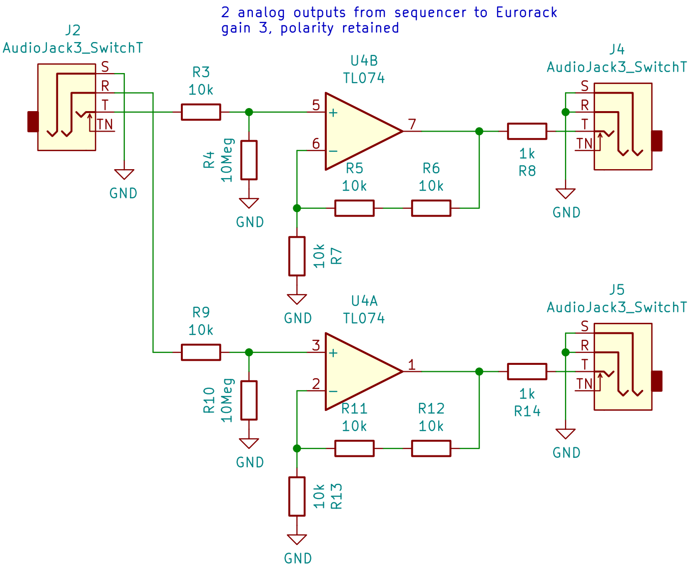

Figure 5.5. Sequencer daughterboard schematic, part 3: Sequencer analog output level translation.

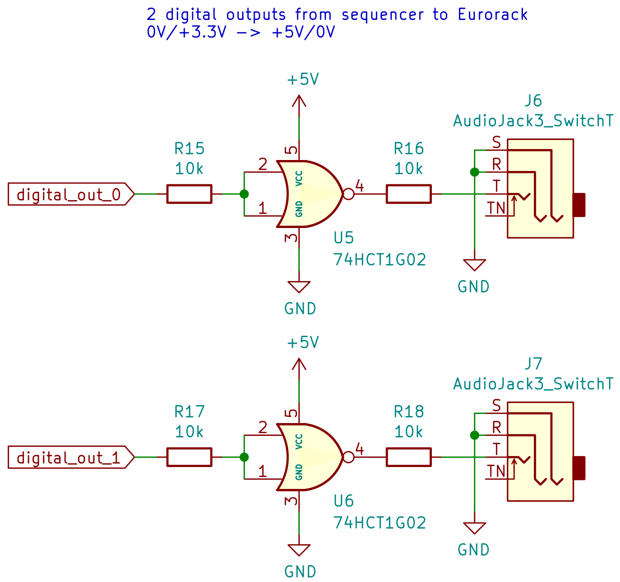

Figure 5.6. Sequencer daughterboard schematic, part 4: Sequencer digital output level translation.

The digital output signals are supplied from the STM32F7308 microcontroller via the STMod+ connector. As the microcontroller outputs are +3.3 V level and Eurorack-style gates are +5 V level, I performed level conversion using a TTL-input 74HCT gate, as can be seen in Figure 5.6. As the +5V rail was supplied by the Discovery Kit via the STMod+ connector, I used 10 kΩ output resistors to avoid potentially severely loading the rail in case another output is connected to the digital outputs.

Unfortunately, the +5V rail supplied by the Discovery Kit is rather noisy, as discussed in Section 7.3. While the digital outputs can be used as gates, usage as control voltage is very problematic. Furthermore, output resistances are larger than typical, which may result in unexpected voltage divider behaviour when connecting two outputs together. In retrospect, I would generate the +5V rail used for digital output gates from the +12V power supply rail instead.

I also included a single translation circuit for the other direction, so that the synthesizer output can be safely recorded with consumer-grade line level equipment. Eurorack reference-level 10 Vpp signals are scaled down to 1 Vpp. Diode clamping is used to ensure that the levels remain at safe values. While this solution may introduce some distortion in normal operation due to reverse-bias diode current, it works well enough in practice.

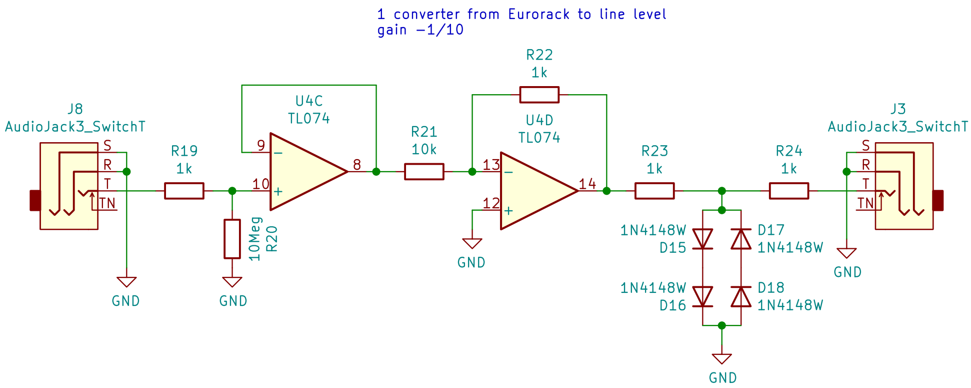

Figure 5.7. Sequencer daughterboard schematic, part 5: Eurorack level to line level translation.

5.3 Voltage-Controlled Oscillator (VCO)

For the construction of the VCO, I was inspired mainly by [21], but made extensive changes during design so that the end result is very different.

First, I wanted the VCO to produce all of the classic output shapes: triangle, ramp (also called saw due to the shape), and pulse (or square, as the duty cycle is constant 50% as implemented). Triangle waveform generation from ramp produced in [21] may pose problems due to discontinuities. Therefore, I used triangle as basic waveform, as a ramp can be generated easily from it. Second, I did not want to use an Operational Transconductance Amplifier (OTA), preferring a operational amplifier-only solution due to costs and sourcing. Third, and most crucially, due to the amount of frequency drift in cheap synthesizer VCOs (both in offset and V/Oct), I decided to implement a digital trimming solution for the VCO. Unfortunately, this made the VCO rather complex.

The core design of the VCO can be roughly split into four parts:

- Input stage, where a choice is made between the input CV and input reference for trimming.

- Trimming stage, where 1 V/Oct trim and 0 V frequency trim is done.
- Exponential converter and inversion stage, where the processed 1 V/Oct CV is converted to linear CV for processing.
- Oscillator core, where the linear CV is integrated and integrated value overflow results in integration direction change, resulting in triangle-wave oscillation.
- Output stage, where the pulse output is brought to reference level and ramp output is generated.

5.3.1 Input stage

The input stage is visualised in Figure 5.8. The main difference from a regular input stage is that reference voltage `trim_ref` can be selected instead of the input CV for purposes of trimming. This choice is made before the input buffer so that voltage offset difference between normal input and reference input is minimized. The `trim_ref` voltage is provided by an MCP1501-10 voltage reference. 1.024 V is supplied by the reference in normal operation. When the reference is tristated using a shutdown pin, 0 V is supplied via a 10 kΩ resistor. The choice of reference level is necessary for V/Oct trimming.

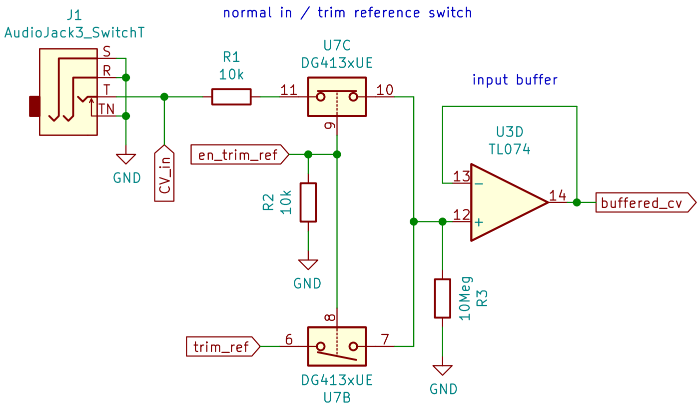

Figure 5.8. Voltage Controlled Oscillator schematic, part 1: Input stage.

5.3.2 Trimming stage

V/Oct trimming is problematic as the V/Oct CV must be multiplied by the trimming factor, i.e. a controlled-gain amplifier is needed. Using a VCA is not possible, as it would be subject to inaccuracies, further reducing VCO precision. Programmable Gain Amplifier (PGA) ICs typically feature coarse gain control either for measuring equipment usage or audio usage. For V/Oct trimming, I needed a way to control the gain in small increments, though not necessary precisely, as monotonicity would be sufficient for feedback control.

After considering various approaches, I devised a solution using the MCP4251 dual digital potentiometer (digipot) IC, as seen in Figure 5.9. While a single digital potentiometer typically has either 128 or 256 stops which is not precise enough, I combined the two potentiometers together in a coarse-fine configuration inspired by [32]. Using this configuration, the coarse digipot can be adjusted first, the fine digipot second.

Unfortunately, as the digipots are +5V-only, voltage preprocessing and postprocessing must be done. In the inverting op-amp U3A, the level of incoming CV is scaled down and referenced to +2.5 V. The output current is then limited and additional Schottky diode protection is added. In retrospect, this is not strictly needed as MCP4251 has a generous ±20 mA maximum input clamp current.

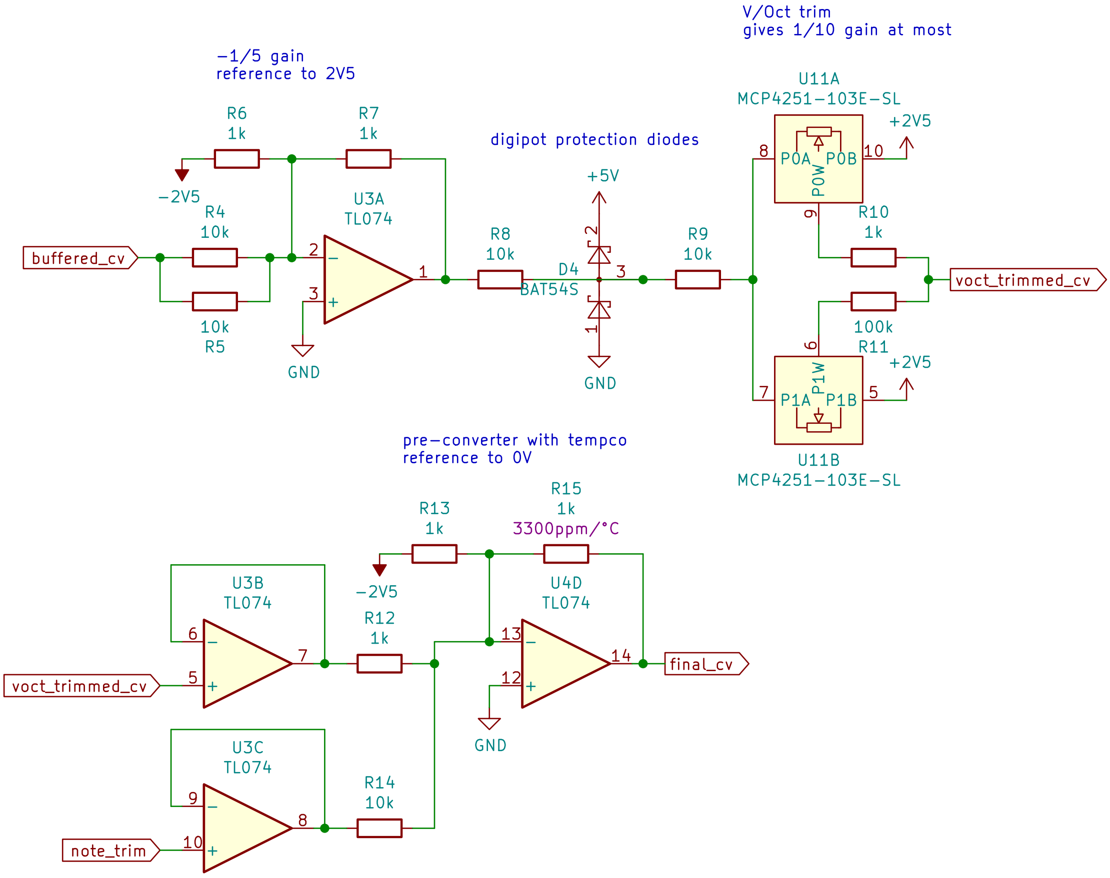

Figure 5.9. Voltage Controlled Oscillator schematic, part 2: Trimming stage.

After V/Oct trimming is performed, the result is buffered by U3B and referenced back to 0 V by U4D. At this point, the attenuation of input CV is 50 at least and 100 when both digipots are centered. Middle note (0 V CV) frequency trimming is also added at this point. While the frequency trimming voltage could be generated by a Digital-Analog converter (DAC), I used another MC4251 in the coarse-fine configuration due to the very fine increments possible and Bill of Materials (BOM) simplification purposes. While I could use a quad-digipot IC, I opted not to due to their lesser availablity.

The resistor R15 is a Positive Temperature Coefficient (PTC) thermistor required by the next stage for smaller temperature dependence, as discussed in Section 3.3. Unfortunately, I was unable to source a 3300 ppm / degree Celsius thermistor in SMD configuration, so I used a 3900 ppm / degree Celsius thermistor instead. This could have contributed to lackluster temperature stability of the VCO, as detailed in Section 7.4.

In addition to the thermistor problem, I also overloaded the +2V5 rail (produced by an LM4040 voltage reference) during design, so the actual +2V5 rail voltage is roughly +1.3 V. While this did not render the VCO inoperable, the triangle and ramp wave amplitude was affected in particular, leading to the digipot middle frequency rising as well to approx. 1.3 kHz. As remaking the already completed circuit on the PCB would be problematic, I only replaced R23 and R24 in Figure 5.11 with 10 kΩ resistors to counteract the increased middle frequency with higher integration time.

5.3.3 Exponential converter and oscillator core stage

The exponential converter stage visualised in Figure 5.10 is based on the exponential converter circuit explained in Section 3.3. I used a matched pair of NPN transistors in one package BCM56DS. The op-amp U4B is used as a current-to-voltage converter. I added a very conservative current-limiting resistor R16 for additional protection of the transistor from excessive base currents. However, the resistor is not necessary upon further inspection as the U4B can only sink approx. 1.2 mA through the emitters at most, well below limiting values.

The oscillator core in Figure 5.11 is comprised of four building blocks. First, the linear CV from the exponential converter is inverted if deemed necessary. After that, integration occurs. During design, I targeted 220 Hz middle frequency, resulting in component choices shown.

While I could use a single op-amp in positive feedback with hysteresis configuration for inversion of the `go_down` voltage, I considered the dependence on voltage extremes achievable on amplifier output problematic, especially since TL07x is not a "rail-to-rail" operational amplifier. Instead, I chose to use two op-amps in comparator configuration. Each of them can detect overflow in one direction. The integrator value is divided by 2 before the comparators to avoid phase inversion. When the integrator value overflows, the corresponding diode conducts and U6C is flipped. I consider this a good solution that does not depend on rails, though the component cost is quite significant.

When the `go_down` voltage polarity is flipped, integration polarity flips as well due to U7D. This leads to oscillation where `integrator_val` contains a triangle wave and `go_down` contains a square wave of the same frequency.

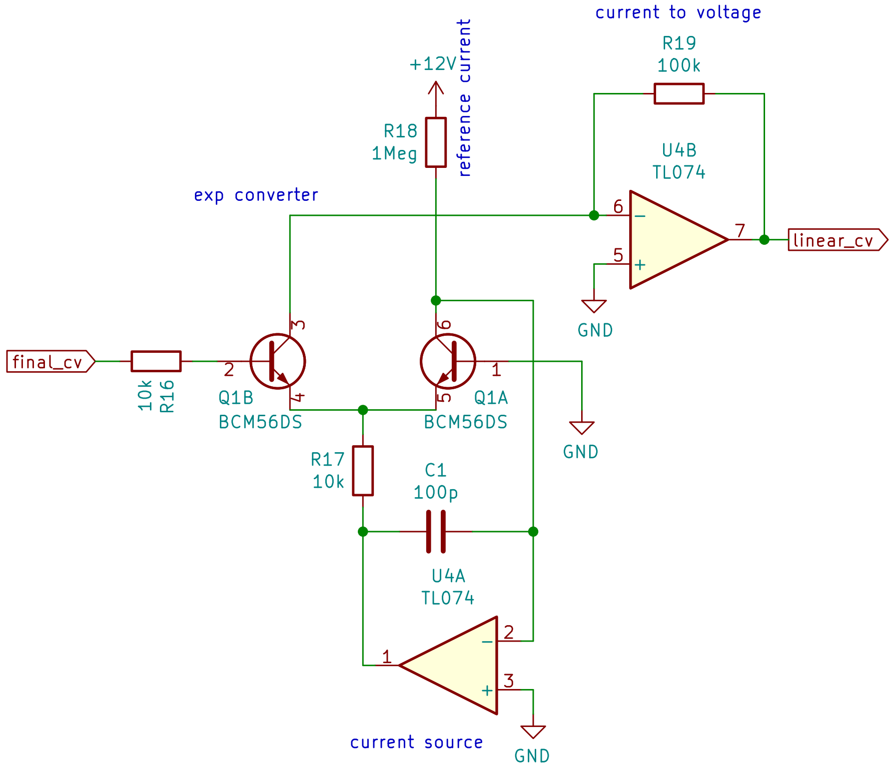

Figure 5.10. Voltage Controlled Oscillator schematic, part 3: Exponential converter.

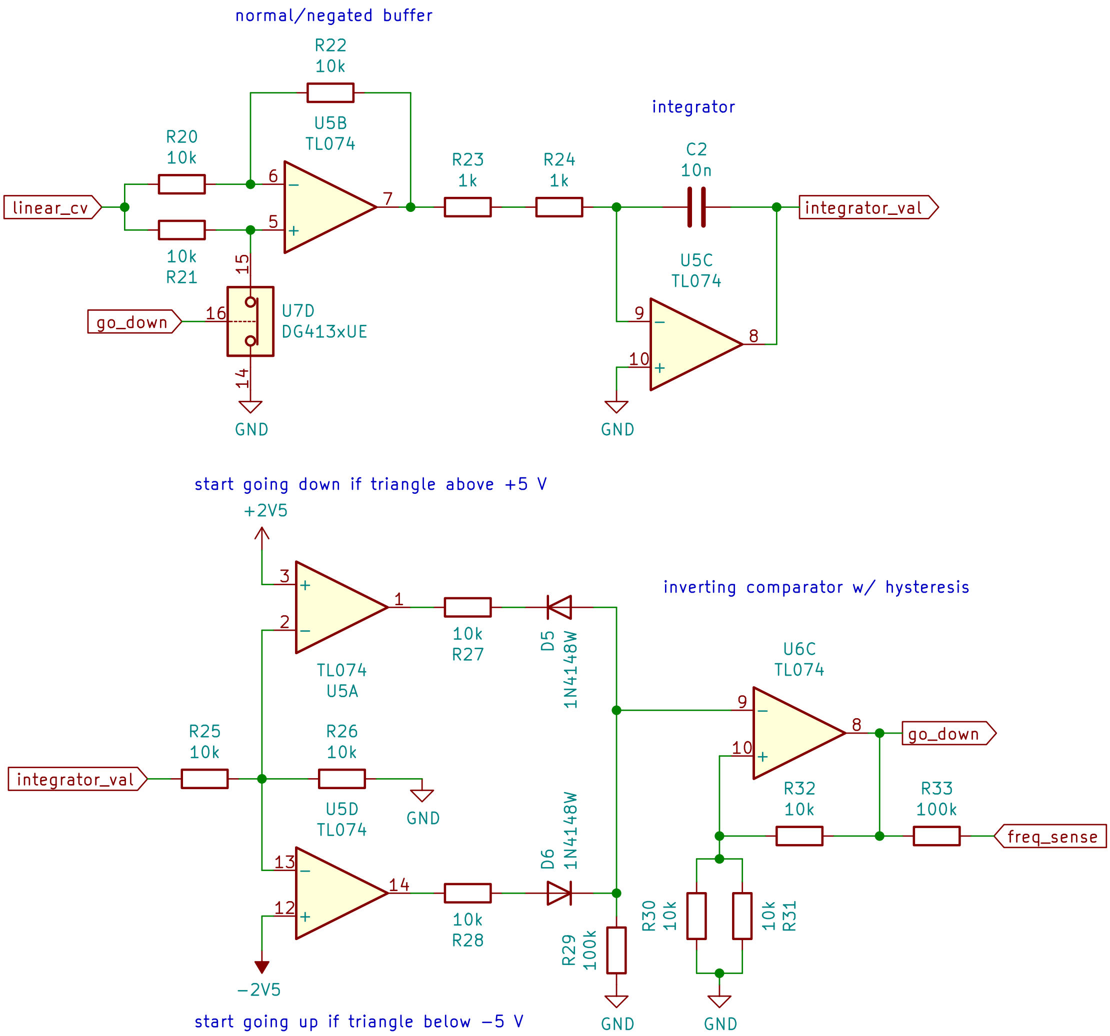

Figure 5.11. Voltage Controlled Oscillator schematic, part 4: Oscillator core.

5.3.4 Output stage

The VCO output generation in Figure 5.12 is rather straightforward. The pulse output is generated by feeding the `go_down` signal (which is close to positive or negative rail) to the NOR gate U12, which has either 0 V or +5 V at its output. Amplified by U6A in non-inverting amplifier configuration with gain 2 referenced to +5 V, the final pulse wave output voltage is either +5 V or −5 V.

The ramp output is produced simply by inverting the integrated value using the `go_down` signal. The resulting ramp wave fundamental is an octave higher than the triangle and pulse wave fundamentals.

Since one op-amp channel was left unused in the VCO, I used it to buffer the triangle wave output. In theory, buffering could prevent loading the integrator output and result in a more consistent performance, but I doubt that the difference would be noticeable in practice.

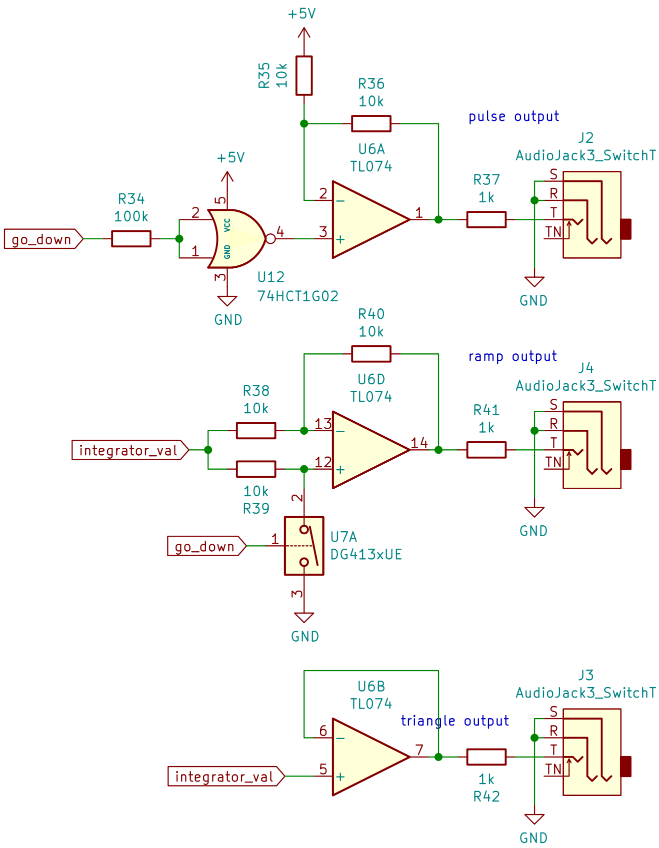

Figure 5.12. Voltage Controlled Oscillator schematic, part 5: Outputs.

5.4 Voltage-Controlled Amplifier (VCA)

I implemented a Long-Tailed Pair VCA discussed in Section 3.4. Specifically, I was inspired by [33], tweaking some values and parts for reduction of BOM. In Figure 5.13, the input stage of the VCA is visualised. The only item of note is my gain reduction solution via U1B. In [33], gain is lowered by a voltage divider in front of the transistor input. I chose to use an operational amplifier to lower the gain, and chose to use higher values of voltage divider resistors to simplify BOM. Unfortunately, the noise performance of the resulting VCA is very lacking, as discussed in Section 7.5. Johnson noise of the 1 MΩ resistor R3 is badly matched to TL07x, as discussed in Section 2.5. The 100 Ω / 100 kΩ resistor divider in [33] seems to be a superior choice.

The long-tailed pair stage in Figure 5.14 is the core of the VCA. I already explained the basic idea of long-tailed pair VCA in Section 3.4. Compared to the idealized circuit, resistors R13 and R14 were added. This result in slight output voltage scaling with no other significant differences. As scaling is also adjusted by R11, R12, and RV1, the resistors R13 and R14 could be removed without significant problems in a future revision. Compared to [33], I adjusted some values for BOM reduction.

In the output stage of the VCA, visualised in Figure 5.15, further amplification by 10 is performed. U2A is mainly present so that an even number of op-amp channels is used, as the further amplification could be performed using U2B as well. In retrospect, I would probably eliminate U2A and U1B from Figure 5.15, decreasing the number of op-amp channels needed to 4, so that only a single TL074 IC is needed.

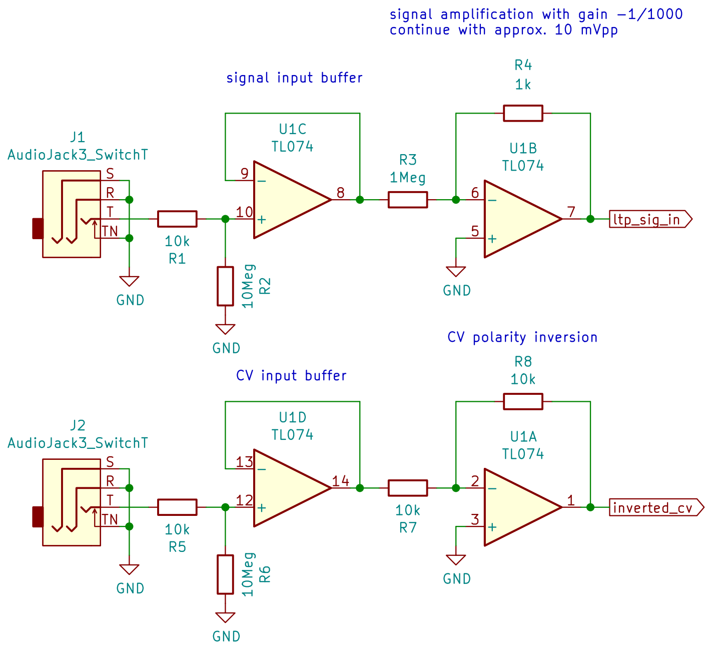

Figure 5.13. Voltage Controlled Amplifer schematic, part 1: Input stage.

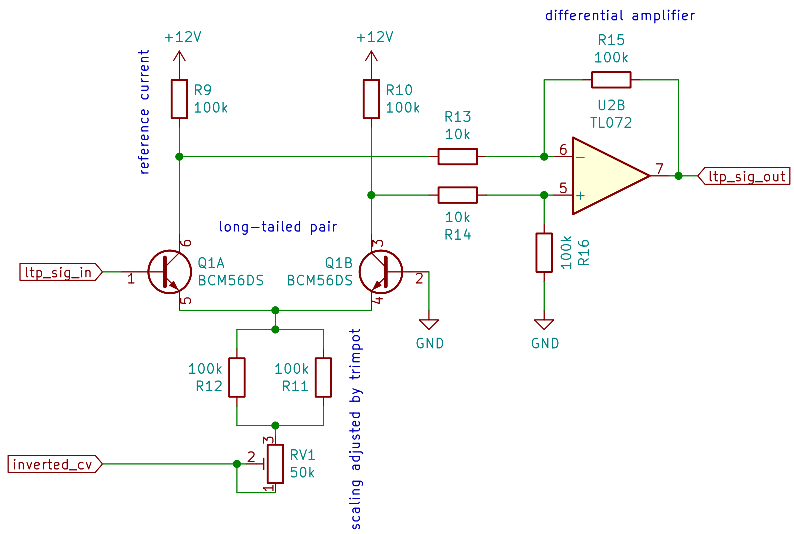

Figure 5.14. Voltage Controlled Amplifer schematic, part 2: Long-tailed pair stage.

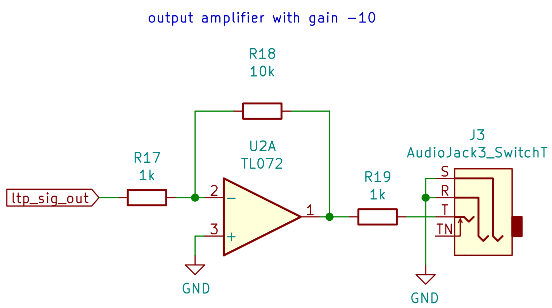

Figure 5.15. Voltage Controlled Amplifer schematic, part 3: Output stage.

5.5 Attack-Release Envelope Generator (AR EG)

As discussed in Section 3.5, I did not find any widely-used EG circuits with CV-controlled intensities. Therefore, I decided to design my AR envelope generator myself, with inputs for gate, attack intensity, and release intensity, and a single envelope output.

In the first part of the EG, visualised in Figure 5.16, the appropriate intensity is selected using a voltage-controlled single-pole double-throw (SPDT) switch. The gate input buffer is not needed. I included it to even out the number of op-amp channels.

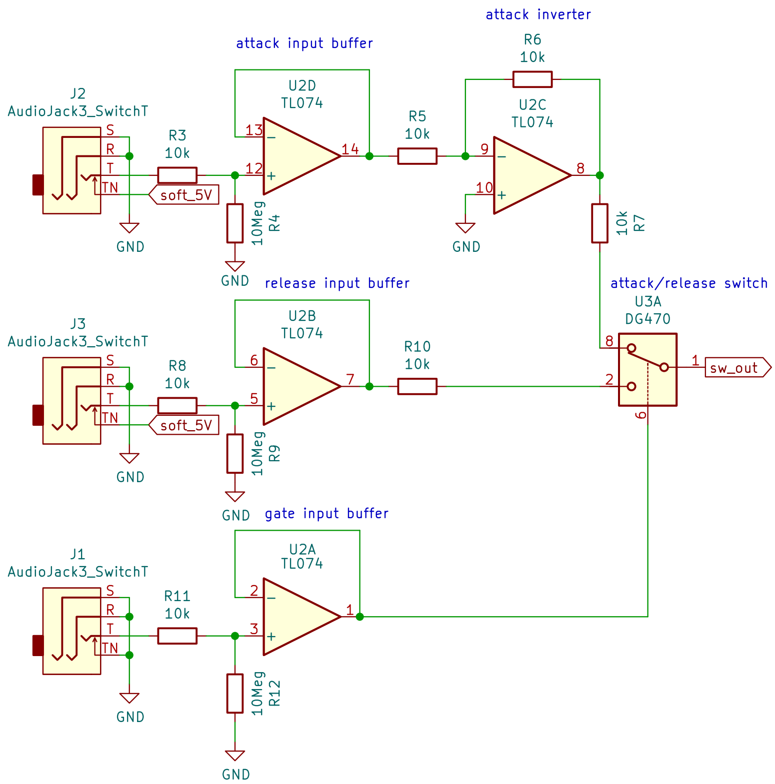

Figure 5.16. Attack-Release Envelope Generator schematic, part 1: Integration intensity CV switching on based on the gate input.

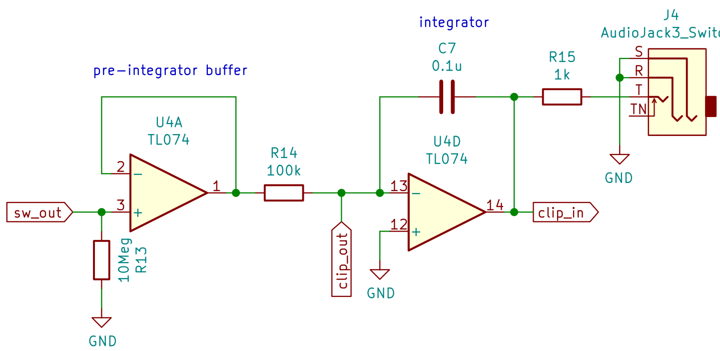

Figure 5.17. Attack-Release Envelope Generator schematic, part 2: Integration of selected intensity voltage

After the intensity is selected, it is integrated, as visualised in Figure 5.17. In my module implementation, the selected intensity is buffered first to eliminate the dependence of the integration constant on switch resistance, which can vary. The buffer is prevented from floating on break-before-make action by R13. The pre-integrator buffer could be omitted, especially if the switches with low input resistance and its flatness are used. I chose the integration values to allow fast integration times, measured in Section 7.6. In retrospect, I would raise the integration time as most of the intensity CV range results in envelope changes too fast to notice.

To keep the envelope output in 0..+5 V range, I used a parallel feedback loop to the integrator, visualised in Figure 5.18. The operational amplifiers U4A and U4B operate as non-inverting amplifiers with gain of 1001, but can be thought of as comparators for the purposes of explanation.

The top op-amp U4C compares the current envelope output voltage with a +5V reference rail. In normal operation, the envelope output is below +5V, the comparator output is negative, D4 does not conduct and the parallel feedback loop branch is inactive. Once the envelope output is above +5V, the comparator output is positive and D4 starts conducting. Since the integrator input value polarity is negated, this results in the envelope output being pulled down. As the resistance of R22 is 10 times lower than the resistance of R14 (in Figure 5.17), it dominates over the selected intensity and the envelope output voltage is clipped only slightly above +5 V.

The bottom op-amp U4B behaves similarly to U4C, the only difference being that the envelope output voltage settles slightly below 0 V. I included R18, R19, R20, R21 to reduce amplifier gain for prevention of possible undesirable behaviour (especially overshoot), but they could be removed if care is taken.

I also added a phase reversal prevention diode D3. When the envelope output is only slightly above the −12V rail (which could happen, for example, during startup), the output of U4C could be inverted, as if envelope voltage was above +5 V. This would lead to the envelope output being aggressively pulled further down, precluding normal operation. Such error mode cannot occur when D3 is present.

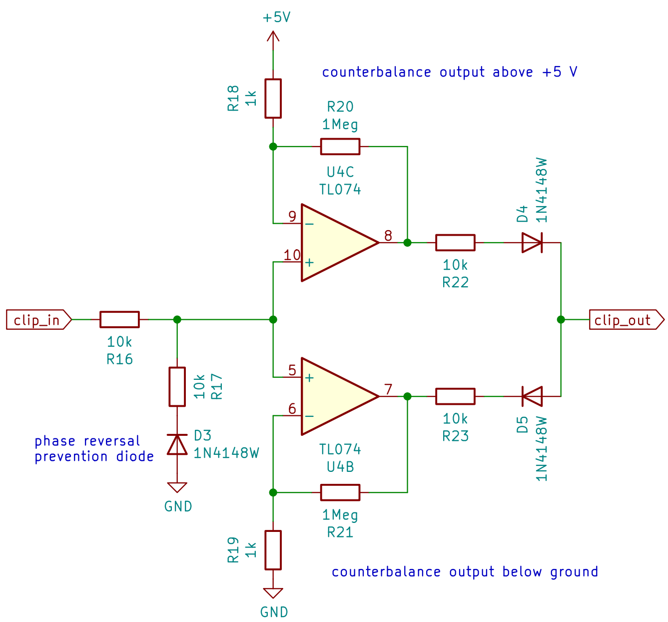

Figure 5.18. Attack-Release Envelope Generator schematic, part 3: Feedback loop for limiting integrated value to 0V..+5V

5.6 Voltage-Controlled Filter (VCF)

As discussed in Section 3.6, I wanted to use a filter that would require only a small number of easily-sourced parts, and the very simple diode ladder filter from [27] seemed to be ideally suited for my purposes. I adapted it for my uses in a simple version without resonance control.

The input stage of the VCF in Figure 5.19 serves to buffer the audio signal and cutoff CV. The signal gain is reduced by 101 in U2A. The choice of this gain factor will become clear in the following stages.

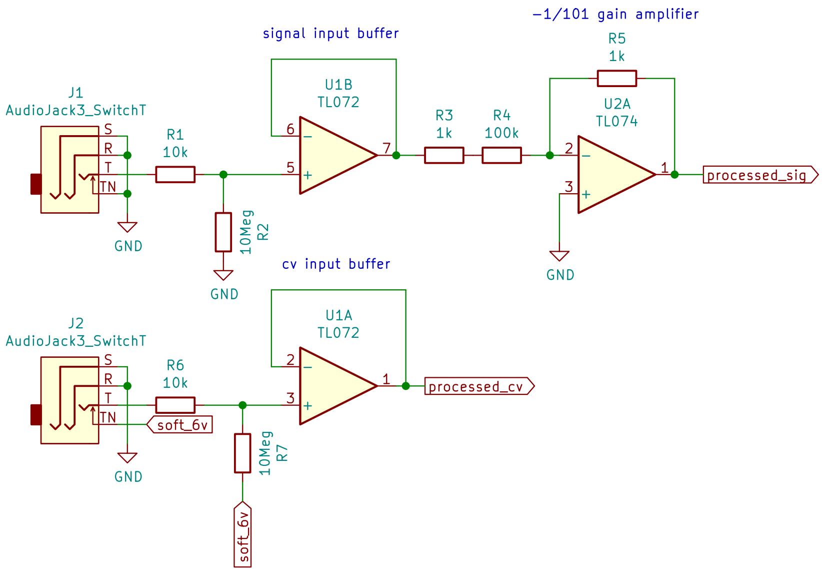

Figure 5.19. Voltage-Controlled Filter schematic, input stage.

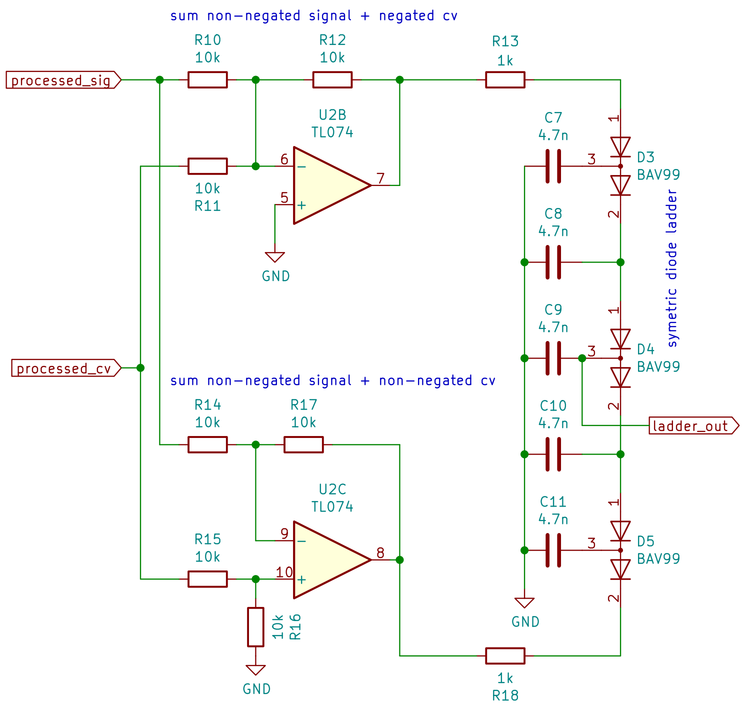

Figure 5.20. Voltage-Controlled Filter schematic, filtering stage.

The core of the VCF, as I implemented it, is visualised in Figure 5.20. Considering that cutoff CV was not inverted in the previous stage, and comparing it to Figure 3.6 and [27] or analysing the circuit directly, it is apparent I have made a design error: for positive CV, the diodes are reverse-biased and the ladder does not work as intended, the filter instead working when CV is negative.

I traced the source of the error to the fact that I verified the circuit using KiCad-native Ngspice simulation, where circuit netlists are generated with proper component pin numbering. As manufacturers of dual diodes and transistors typically only provide single-diode SPICE models, I had to create the package model with correct pin numbering myself. Unfortunately, I used wrong polarity for the BAV99 dual diodes, not realizing my mistake until after PCBs were made. I plan to fix this oversight in a future revision.

I chose the values of capacitors to be 4.7nF so that the frequency response is almost linear up to 20 kHz when the CV is 5 V (−5 V when the unintended CV negation is taken into account), designing by simulation. As discussed in Section 7.7, this is not an optimal choice, as it results in interesting cutoff range being clustered around 0..+1 V (−1..0 V when the unintended CV negation is taken into account).

The last part of VCF is the output stage, which is extremely simple. Non-inverting amplifier U2D is used to amplify the ladder output without any loading (which could change the frequency response of the VCF considerably). As the gain of U2D in this configuration is 101, I used an additional resistor in the input stage gain-reduction op-amp U2A to reduce the gain by 101 there, as shown in Figure 5.19. While the actual gain in built modules may vary due to component tolerances, I avoided consistent gain error by including the resistor R3.

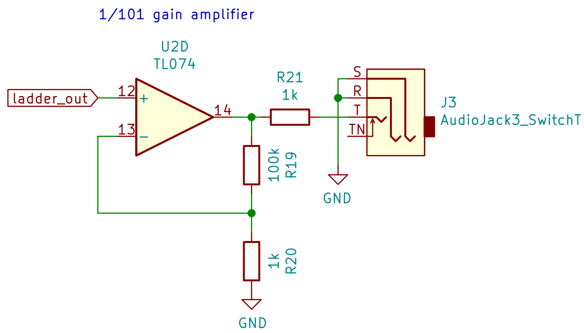

Figure 5.21. Voltage-Controlled Filter schematic, output stage.

Chapter 6
Synthesizer firmware

In this chapter, I discuss the firmware I developed for the microcontrollers used in the sequencer and VCO modules. The full firmware, not including open-source libraries used, is available attached to this thesis, as discussed in Appendix B.

6.1 Sequencer firmware

The high-performance STM32F7308 ARM Cortex-M7 microcontroller is used in the sequencer module. In addition to high-level sequencer control behaviour which determines the ability of the user to create interesting and musical sequences, the other two major implementation sections were audio output and low-level I/O control.

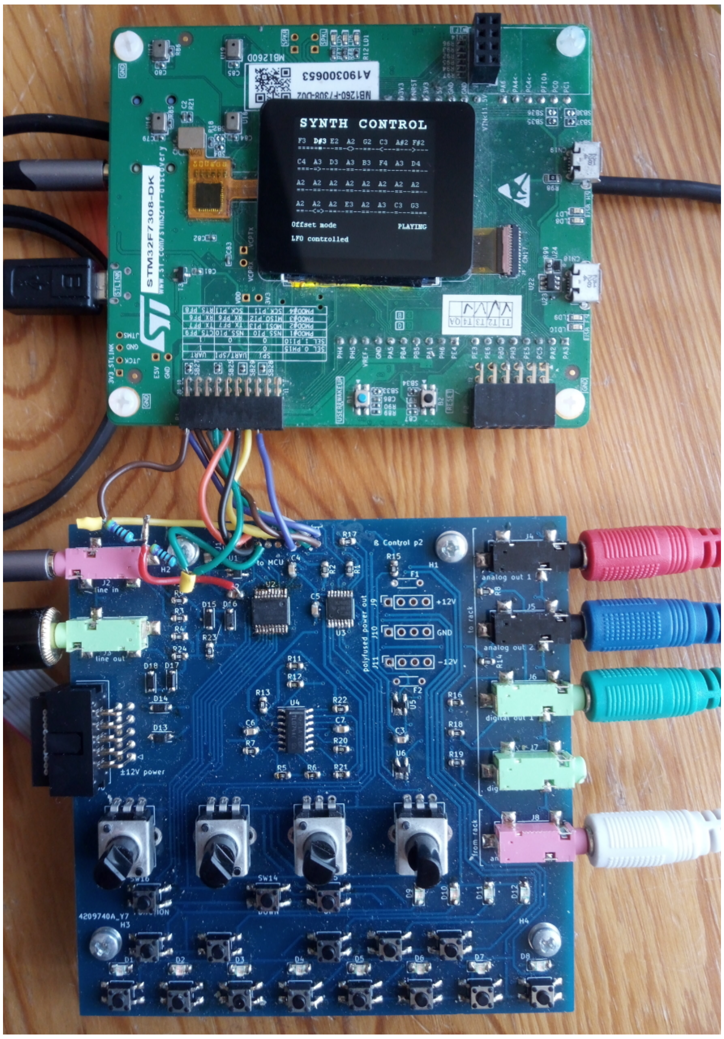

Figure 6.1. Sequencer module during normal operation. The turned-on LEDs are almost invisible in the photo due to ambient light.

6.1.1 Control behaviour

During layout design, I was inspired by classic sequencer-based synthesizers, especially the Roland TB-303 synthesizer, leading me to use a one-octave-and-note piano-style keyboard for note control. I made the other decisions during firmware development. By changing the firmware control behaviour, the sequencer could be tailored to the tastes of each user, with hardware as the only limitation.

The sequencer during normal operation is shown in Figure 6.1. Conceptually, there are 4 sequences comprised of 8 beats. One of the sequences and its beats is selected, and the selected sequence is played through on repeat unless the sequencer is stopped. Most of the outputs depend on the current beat and its fraction:

- **Analog output 1.** Current beat note pitch CV.
- **Analog output 2.** Sine-wave low-frequency oscillator (LFO) output.
- **Digital output 1.** Gate status, determined by current beat fraction.
- **Digital output 2.** Current beat accent.

From the user perspective, there are three rows of switches. There are 4 *operation modes* that change the behaviour of the lower and middle row of the switches:

- **Note mode.** In this mode, the lower and middle row switches act together as a piano-style keyboard and change the current beat note when pressed. To alleviate the small number of notes available, I implemented "octaveless" change, i.e. the note octave is selected to minimize the absolute semitone change.
- **Offset mode.** In this mode, the current beat fraction where the gate opens (is pulled high) can be set by pressing one of the lower row switches, with the leftmost switch corresponding to $\frac{1}{8}$, the second to $\frac{2}{8}$, and so on. In addition, pressing one of the middle row switches sets the fraction to 0.
- **Length mode.** This mode is similar to the offset mode, but the length of beat fraction during which the gate is being held open is changed by pressing the switches.
- **Accent mode.** Pressing one of the lower row switches changes whether the corresponding beat in the current sequence is accentuated.

The two rightmost switches in the upper row change the octave of the current beat note up or down in every mode. The leftmost upper switch acts as a function modifier. Pressing switches while holding the function modifier changes their function. Current mode, sequence, and beat can be changed, the sequencer can be stopped (or restarted, if already stopped), and the function of the rightmost potentiometer can be toggled.

Beat speed is controlled by the leftmost potentiometer. By default, the other three potentiometers control frequency, amplitude, and offset of the LFO, which is digitally synthesized as a sine wave on the second analog output. The function of the rightmost potentiometer also can be toggled to control note glide, the other value being retained.

In the lower LED row, the beat currently played is lit with high brighness. In addition, the beats in which the gate opens are lit with low brightness (except in accent mode, where the accented beats are lit with low brightness). In the higher LED row, the LED corresponding to current mode is lit with high brightness, while the LED corresponding to current sequence is lit with low brightness. In the future, I would consider using 8 LEDs in the higher row instead of 4 so that both the mode and sequence could be shown easily without interfering with each other.

The sequences and sequencer status are shown on the LCD display as visualised in Figure 6.2. In practice, this provides information much better than the LEDs, with all information about the sequencer readily available.

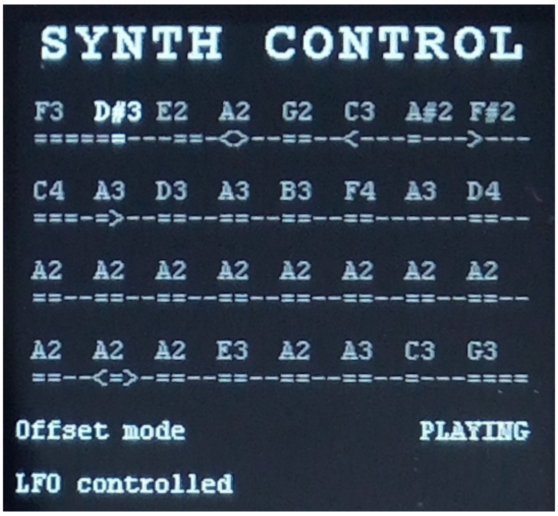

Figure 6.2. Sequencer module LCD display during normal operation. Four sequences comprising of eight beats are shown. Under each beat note, beat fractions during which the gate is open are displayed. The currently played beat and its fraction is highlighted. In the lower part of the display, sequencer status is shown.

The sequencer is rather capable and sequences can be created quickly, although the need for switching between various modes, combined with other sequencer handling issues that I discuss in Section 8.3, makes it rather awkward and unsuitable for professional music production. However, it is still quite fun to use.

6.1.2 Audio output

My choice to include digital output pins was problematic as synchronicity between analog and digital outputs is needed. However, on the STM32F7308, the Serial Audio Interface (SAI) peripheral has its clock separate from the microcontroller core clock and is filled using an internal First-In-First-Out (FIFO) queue. Furthermore, the Direct Memory Access (DMA) peripheral can be configured to feed the queue automatically, typically from a circular buffer with multiple parts that are filled in the main program loop. DMA is used in STM32F7308 Discovery Kit Board Support Package audio codec initialization. However, output synchronicity becomes problematic, as I did not find any way to use the audio clock for digital output control purposes.

For a clean solution, I wrote audio initialization and handling myself, selecting FIFO mode, an interrupt occurring when the FIFO is no longer full (i.e. ideally every sample). In addition to the interrupt filling the FIFO, it changes the states of digital outputs correspondingly. After resolving implementation problems, the concept works nicely in practice. Overhead is incurred due to one interrupt per audio sample, but the chosen microcontroller is powerful enough that it is not a problem with optimized interrupts.

6.1.3 Input and output control

I discussed the I/O schematic in Section 5.2.1 and will only explain the firmware aspect here. The Discovery Kit LCD is periodically written to on per-character basis when no interrupts are pending and there are no audio buffers to process in the main program loop. For daughterboard I/O, a Time-Division Multiplexing (TDM) scheme is used. A free-running timer is set up to generate an interrupt every 0.2 milliseconds. The interrupt handler performs the following actions depending on the current timeslot row type and prolog status:

- **Row prolog action.** a command is issued to the I/O expander to force all pins into input-with-pullup mode, the PWM pin is set as always low (zero duty cycle), and multiplexer pins are updated with the row value. This assures that the row is switched properly for the non-prolog action in the next timeslot.
- **Switch row non-prolog action.** a command is issued to the I/O expander to read the switch values. After the read completes, the current switch values are updated according to the result. Next row prolog follows in the next timeslot.
- **LED row non-prolog action.** The PWM pin duty cycle is set to zero. A command is issued to the I/O expander to force the current LED column pin into output mode. After the write completes, the PWM pin duty cycle is set according to the desired brightness. Next LED follows in the next timeslot, next row if current LED was the last in the row.

The scheme makes it possible to process all switches, LEDs, and potentiometers. Switches are debounced and potentiometer values filtered. Overhead is minimized as all I2C expander communication is done asynchronously and the timer and I2C transfer/receive complete interrupt handlers are optimized.

The choice of base time-division multiplexing frequency is crucial, especially for LEDs: the higher the frequency, the higher the overhead, but too low frequency will result in LED flickering. As 16-bit communication over 100 kHz clock I2C interface occurs during each timeslot and there are 18 timeslots, I have chosen a 5 kHz slot timer, resulting in LED refresh rate of 278 Hz. This seems to work well, although I would choose a higher refresh rate on the order of kHz for professional equipment to prevent strain in individuals sensitive to flicker.

Overall, the scheme works well. However, due to the small total LED duty cycle (at most $\frac{1}{18}$), the LEDs are not very bright with the selected PWM resistor value of 1 kΩ (resistor R1 seen in Figure 5.4). A lower value should be used for higher brightness, subject to absolute maximum ratings.

6.2 Voltage Controlled Oscillator (VCO) trimming firmware

For digital trimming, I used an ATtiny24A microcontroller due to its low cost, sufficient number of pins and ability to use an external crystal with minimal external circuitry. This is crucial as the internal RC oscillator is accurate only to 10% from the factory and 1% after user calibration [34, p. 181].

As implemented, trimming takes about 10 seconds and is performed both after system reset and when a switch is pushed (this was mainly done for testing, as the the switch is inside the enclosure and impossible to reach normally). As discussed in Section 7.4, in the final firmware, trimming is performed four times after system reset, with one-minute periods of normal operation between them, so that the VCO reaches working temperature before the final automatic trim.

The trimming algorithm is based on the Successive-Approximation (SAR) ADC principle. The digipot acts as a DAC. Unlike SAR ADC, where the DAC and reference voltages are compared, the frequency of the VCO is measured and compared to the desired frequency and the digipot value is successively refined to achieve the minimum error.

The frequencies are measured indirectly, via the number of microcontroller clock cycles in 8 periods (for increased accuracy). When the rising edge of the pin connected to the VCO square wave signal (through a large resistor for input clamp current limiting) is detected, an interrupt is called that takes a snapshot of a free-running 16-bit timer. By adding the 8 snapshot differences, the final value can be computed and compared, either to the corresponding expected period for middle note trimming, or to the other measured frequency for V/Oct trimming.

For full trimming, the following procedure is followed:

1. All digipots are set to the middle position.
2. The coarse middle note digipot is trimmed so that 0 V frequency is 110 Hz.
3. The fine middle note digipot is trimmed similarly to the coarse one.
4. The coarse V/Oct digipot is trimmed so the 1.024 V CV frequency corresponds to the 0 V CV frequency times $2^{1.024} \approx 2.033$.
5. The fine V/Oct digipot is trimmed similarly to the coarse one.

In practice, digital trimming as designed works very well (apart from the fact that VCO tuning is rather unstable by itself, discussed in Section 7.4). However, I made some errors and oversights in the trimming and support circuitry during design.

First, I erroneously thought that pins annotated as Master-In-Slave-Out (MISO) and Master-Out-Slave-In (MOSI) in the ATtiny24A datasheet perform that function when the ATtiny24A functions as Serial Peripheral Interface (SPI) master. However, these annotations apply when it is being programmed in slave mode. Thus, the pins were reversed, forcing me to bit-bang the SPI interface instead.

Second, while using the interrupt for timer capture works, it is non-ideal. Using the Input Capture peripheral which snapshots the value of the timer would be ideal, but I connected the VCO square wave signal output to a pin that does not feature this peripheral. Of course, these problems could be fully fixed in a later revision.

Chapter 7
Objective measurements

For the purposes of a proper evaluation of the synthetizer, objective measurements of quantifiable parameters are required. I was able to verify that all modules perform their intended function. However, I found significant areas for future improvement through the interpretation of the measurements as well.

7.1 Methodology

To obtain objective measurements of module performance, I recorded or measured the outputs of the modules in response to various test inputs. I also measured the amount of noise superimposed on important power rails.

7.1.1 Used measurement instruments

I used the following instruments for measurement:

- Metex M-3860M True-RMS multimeter for measurement of DC and RMS AC voltages and currents.
- Rigol MSO5074 mixed signal oscilloscope for taking rail and signal oscillograms that are visualised in the figures in this section. Unfortunately, the oscilloscope seemed to have an offset problem, making it unusable for taking DC measurements.
- Focusrite Scarlett 2i2 2nd Generation audio interface, used for recording audio signal outputs. Only the left input and output channels were used.

7.1.2 Input signal generation

As I did not have a signal generator for input CVs, I generated them using a potentiometer voltage divider, measuring the output with the multimeter to achieve the correct voltage. This is a decidedly low-precision approach, with approx. ± 5 mV practical accuracy due to potentiometer drift, but it worked well enough for my purposes.

I generated the noise, 1 kHz sine wave, and constant zero test audio signals using the Audacity digital audio editor and played them using the audio interface. Unfortunately, audio interface output levels are low compared to Eurorack reference signals, complicating evaluation. I alleviated the problem by rescaling the output signals to standard units before evaluation, as discussed in Section 7.1.4.

7.1.3 Signal measurement

I performed rail measurements using the oscilloscope in AC mode. As for module configuration, all modules were powered and the following "standard" mode was used: sequencer CV output routed into VCO input, sequencer gate output routed into AR EG gate input, VCO output routed to VCF input, VCF cutoff set to fully open, VCF and AR EG outputs routed into VCA input and CV respectively. Sequencer pitch CV was kept at 0 V, while gate was toggled at low frequency.

I used the oscilloscope for measurement of the CV signal output of the AR EG module as well. As the most important module parameter was the time difference between integration start and end, the offset problems were not an issue.

I recorded the audio signal outputs using the audio interface in line mode, with the gain knob adjusted so that test signal loopback would not be clipped. I used the Audacity digital audio editor for playback and recording. As expected, the audio interface noise floor was significantly (10 dB or more) lower than the noise floor of synthesizer modules.

7.1.4 Measurement post-processing

I performed post-processing of the recorded output audio signals in the MATLAB numeric computing environment. First, to obtain the recording sample values in standard units, I rescaled the sample units to volts using a 1 kHz test signal recorded with input looped back to output using a patch cable, with the same audio interface settings as the output recordings. I measured the test signal RMS voltage to be approx. 0.459 Vrms and rescaled the signals accordingly.

I performed the evaluation of recordings using power spectra. The spectra were calculated from 9 seconds of output signal (the leading second skipped to avoid discontinuities at the beginning) using the Welch method with a $\frac{7}{8}$ window overlap. I used a Dolph-Chebyshev window with 2^{14} samples and 140 dB sidelobe attenuation factor. Other traditionally used windows are typically inappropriate due to insufficient sidelobe attenuation, which would be especially apparent in Figure 7.9 due to very high signal-to-noise ratio. In addition to the Dolph-Chebyshev window, the Blackman-Harris window is also usable in practice, with only slight sidelobe artifacts that taper off.

7.2 Power supply

The quality of power supply was measured near the power supply using an oscilloscope in AC mode. In Figure 7.1, an oscillogram of the +12V rail is presented (the -12V rail behaves very similarly). There are some apparent oscillations on the rail. The peak-to-peak voltage is below 5 mV and RMS voltage is approx. 1.5 mV (approx. 67.4 dB below 10 Vpp sine wave RMS voltage). This is problematic, as our ideal target for signal-to-noise ratio is approx. 90 dB or more, as discussed in Section 2.1. While minimal supply-voltage rejection ratio for TL07x is typically at least 80 dB (depending on the model) [16], gain reduction in modules (up to 60 dB attenuation) makes this problematic.

Unfortunately, as can be seen from the evaluation of modules where gain is reduced in Sections 7.5, 7.6, and 7.7, power supply ripple rejection is definitely not sufficient, as mains frequency harmonics contamination is apparent. This was surprising to me as I did not expect a well-known synthesizer DYI power supply to be so problematic to use with my modules. In the future, more thorough design of the entire power distribution system is needed.

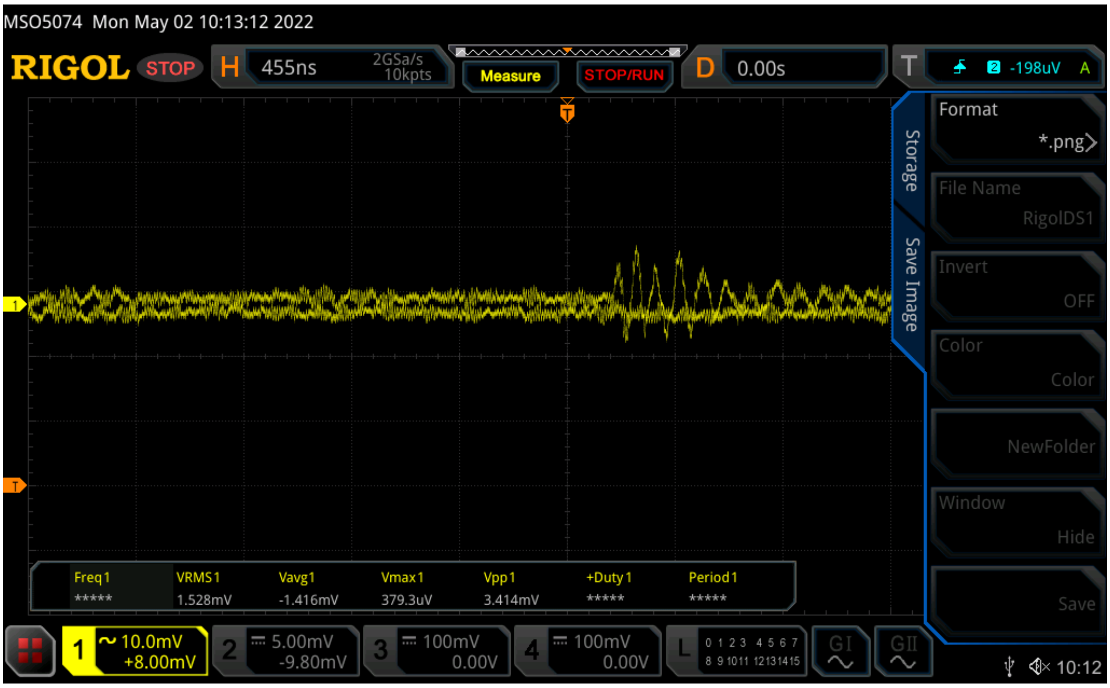

Figure 7.1. +12V supply rail AC oscillogram, measured between +12V and supply ground. The modules were operated in standard mode during measurement.

7.3 Sequencer module

The analog audio outputs of the sequencer do not present significant problems due to their straightforward implementation. However, in addition to the fact that absolute gain of the codecs is not defined, as discussed in Section 5.2, there are offset errors of up to ±6 mV on the daughterboard outputs. This is presumably due to the TL072 input voltage offsets. I have largely compensated the offsets in firmware, though this requires recompilation for every different board and does not resolve drift. But a better solution would be to use an op-amp with smaller maximum voltage offset, as the V/Oct note CV from the sequencer is the one where maximum precision is important.

The sequencer module daughterboard is specific in the fact that it features both ±12V power supply rails, which supply the TL074 used for analog level translation circuits, and a +5V rail supplied directly from the STMod+ connector. I used the +5V rail for digital output level translation purposes as discussed in Section 5.2.2, so its noise level may directly impact audio quality if the digital outputs are used as CVs instead of as gates.

Unfortunately, as can be seen in Figure 7.2, the +5V rail is extremely noisy and suitable only for digital use. In the future, if the daughterboard-style expansion is used, a separate +5V rail should be generated from the +12V rail for usage with digital outputs instead.

As seen in Figure 7.3, the +3V3 supply rail is noisy in the vicinity of 100 kHz as this is the frequency of I2C clock by which the MCP23009 expander IC is controlled. This is not problematic as no audio signal is affected by this rail.

Figure 7.2. +5V supply rail AC oscillogram, measured between the STMod+ 5V and ground pin. The modules were operated in standard mode during measurement. The STM32F7308 Discovery Kit was powered by the STLink USB connector connected to a battery-powered notebook USB port.

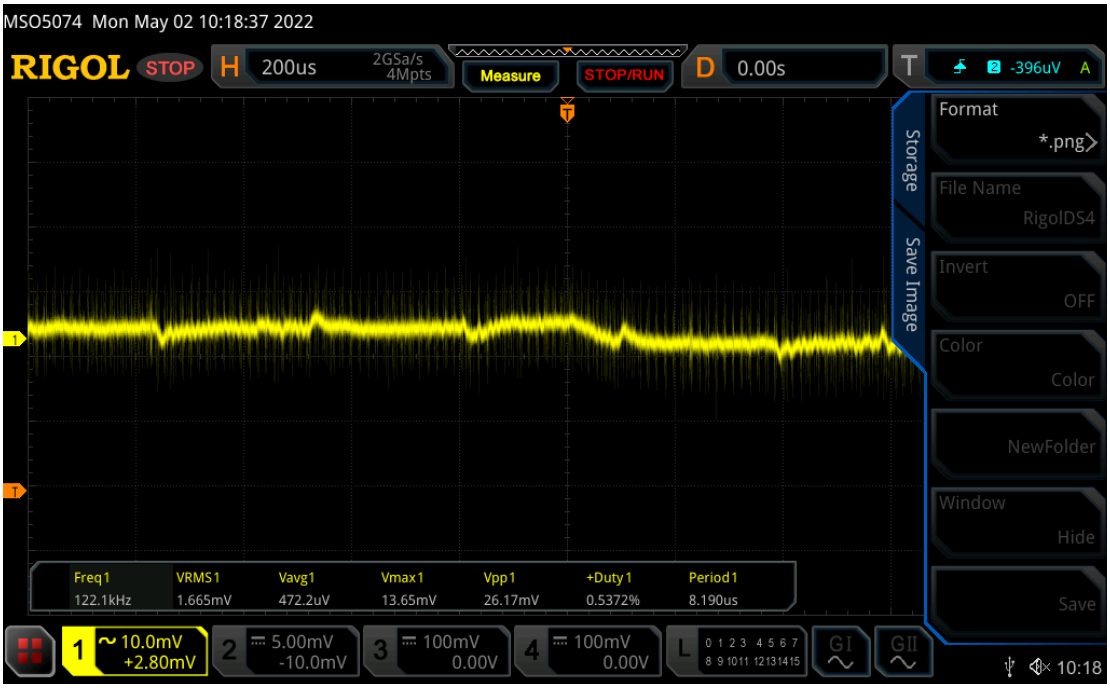

Figure 7.3. +3V3 supply rail AC oscillogram, measured between the sequencer daughter-board TC1015 regulator output and ground pin. The modules were operated in standard mode during measurement. The STM32F7308 Discovery Kit was powered by the STLink USB connector connected to a battery-powered notebook USB port.

7.4 Voltage-Controlled Oscillator (VCO)

VCOs are very problematic in general due to the fact that precision frequency control is needed, as discussed in Section , but they are prone to tuning instability and drift. In addition to using a matched-transistor exponential converter with thermistor compensation (unfortunately, with a different temperature coefficient than recommended due to sourcing issues, as discussed in Subsection 5.3.2), I implemented a digital trimming method.

My original idea was to perform a single trimming pass when the module is turned on. Unfortunately, this was not viable. From Figure 7.4, it is apparent that while the digital trimming works correctly, the parameters of the VCO change during the first few minutes as it heats up, resulting in change of the twentieth harmonic from 2.2 kHz to 2.296 kHz. This corresponds to tuning up by 74 cents. Clearly, the precision gain from digital tuning is irrelevant in this scenario. Therefore, I implemented a scheme in which the trimming microcontroller waits for one minute before trimming is performed again, four times in total.

The multishot trimming scheme is shown to alleviate the effect of temperature-based detuning in Figure 7.4. As the temperature fluctuates, there are still some changes in tuning even after last trimming action. The user also must wait for 3.4 minutes before using the VCO or ignore the brief trimming actions before the automatic trimming is completed.

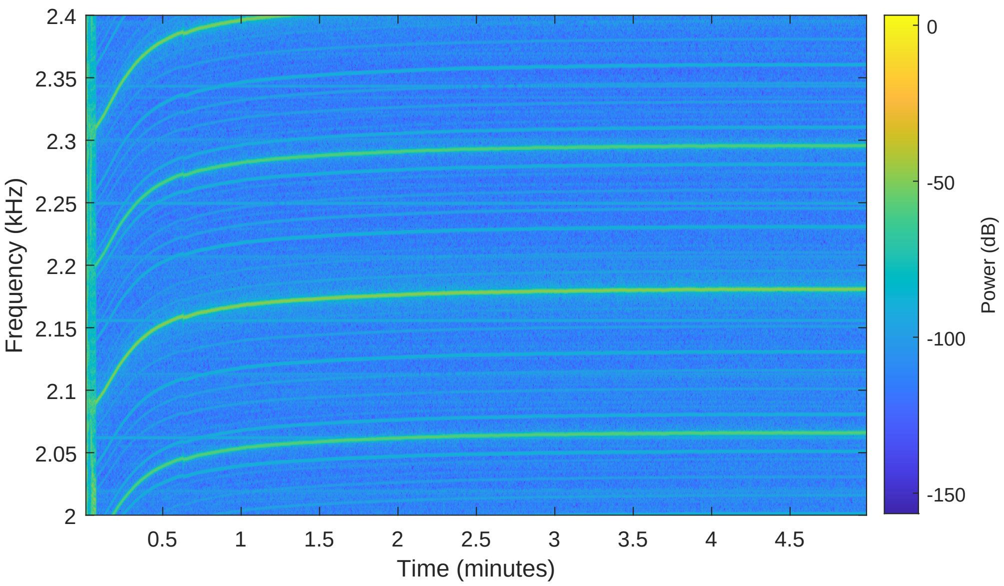

Figure 7.4. Spectrogram of single-shot VCO trimming when the module is turned on. The twentieth harmonic (2.2 kHz) of a 0 V CV, 110 Hz triangle output is shown in the middle. Using the first harmonic would result in insufficient time or frequency precision. The broad-spectrum part in the first 10 seconds is the digital trimming action. The measurement was performed in an interior room with approximate temperature of 25 degrees Celsius.

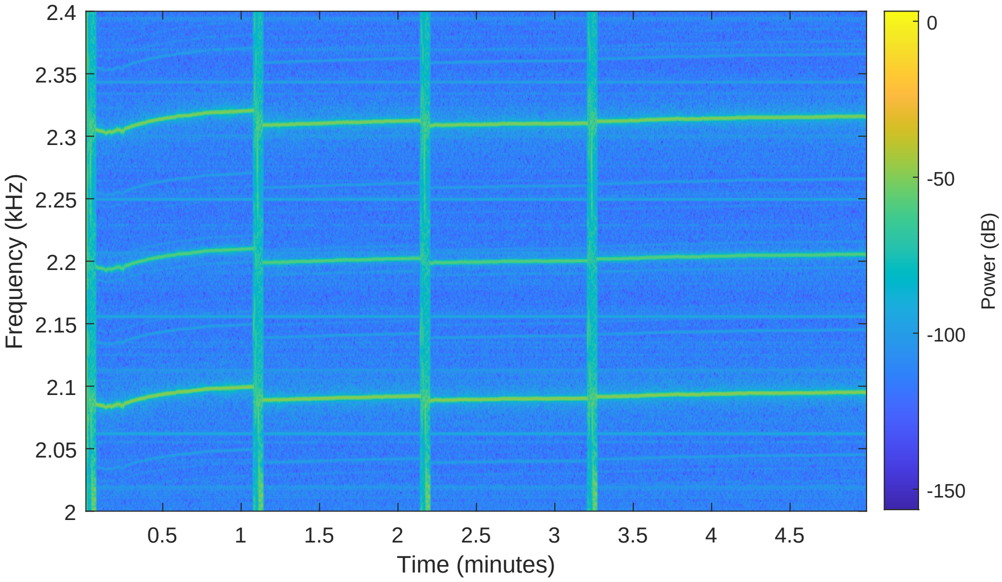

Figure 7.5. Spectrogram of multishot VCO trimming when the module is turned on. The twentieth harmonic (2.2 kHz) of a 0 V CV, 110 Hz triangle output is shown in the middle. Using the first harmonic would result in insufficient time or frequency precision. The broad-spectrum part in the first 10 seconds, around 1.1, 2.2 and 3.3 minutes, is the digital trimming action. The measurement was performed in an interior room with approximate temperature of 25 degrees Celsius.

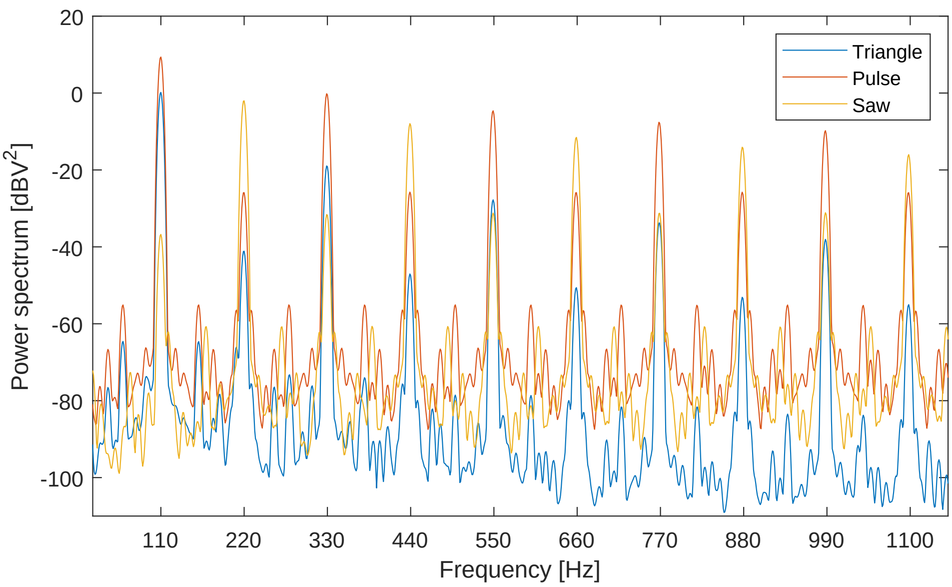

Figure 7.6. Power spectrum of VCO harmonics with 0 V CV, corresponding to 110 Hz. All outputs are shown. The frequency axis is linear in this plot so that all harmonics are evenly spaced. The intended fundamental of the constructed saw wave output is an octave higher than the fundamentals of the triangle and pulse outputs.

In Figure 7.6, the power spectrum of all outputs is shown. Due to the design error forcing the 2V5 rail to approx. +1.3 V (as discussed in Subsection 5.3.2), the power of triangle and saw outputs is severely reduced compared to the pulse output.

Harmonics that are not present in the ideal versions of waveforms are only reduced in the VCO module outputs: the triangle and pulse waves contain some even harmonics, and the saw wave contains a 110 Hz subharmonic, even though it should be ideally missing. Furthermore, some contamination of the pitch CV by the mains frequency is apparent from the creation of 110 Hz fundamental sidebands at 60 Hz and 160 Hz.

While a clean VCO is desired in non-audio engineering, the output signals are quite “interesting”, unlike a precisely generated wave. Therefore, in this situation, these characteristics are not necessarily problematic. The mains frequency contamination should be resolved, however, for the possibility of cleaner outputs.

The tuning errors seen in Figure 7.7 are definitely problematic. As discussed in Section 2.1, the just noticeable error is approx. 5..20 cents. Using the reasonable value of 10 cent maximum absolute error, the useful range is only approx. 1.5 octaves.

All in all, analog VCOs have been supplemented by digitally-controlled oscillators for a good reason. Clearly, better temperature compensation is needed, as is tuning error reduction. While I am happy with the performance of my trimming circuit, I feel that its costs do not justify it in this case. In the future, I would consider replacing the VCO by a digital equivalent that would sample the control voltage and create the output waveform digitally. With currently available microcontrollers and codec ICs, the end product would probably be much better for same or lower costs than its analog counterpart. Analog VCOs could be used for low-frequency oscillators where precise CV tuning is not necessary.

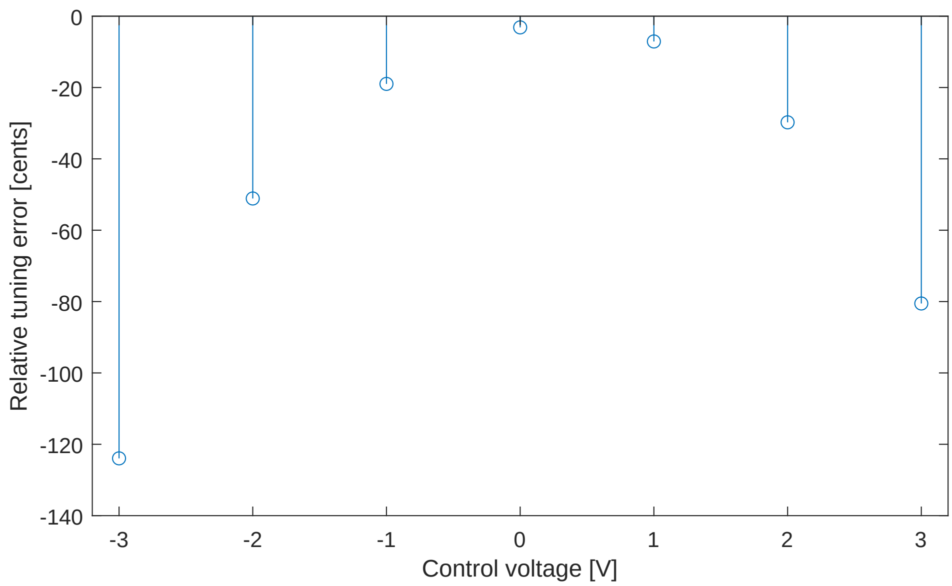

Figure 7.7. VCO tuning error in ±3 octave range. 100 cents = 1 semitone.

7.5 Voltage-Controlled Amplifier (VCA)

For evaluation of the VCA, I used a 1 kHz sine test signal as audio input. Previously, I adjusted the gain of the VCA via its gain trimming potentiometer, so that the input and output amplitude would be the same with +5 V CV.

From Figure 7.8, it is immediately apparent that the VCA output is heavily contaminated by mains frequency harmonics with +5 V CV, with wideband noise also rising to much higher levels than with 0 V CV. While the unwanted spectral components are somewhat reasonable when 0 V CV is applied, the behaviour with +5 V CV is clearly unacceptable.

As evidenced by Figure 7.9, signal gain matching between loopback and +5 V CV output is almost perfect. In fact, they are matched to 0.1 dB, which is a professional-level result. However, as matching is performed manually by adjusting the trimmer present on the VCA, this parameter may vary. When 0 V CV is applied, attenuation is 75 dB, which is somewhat reasonable but not impressive.

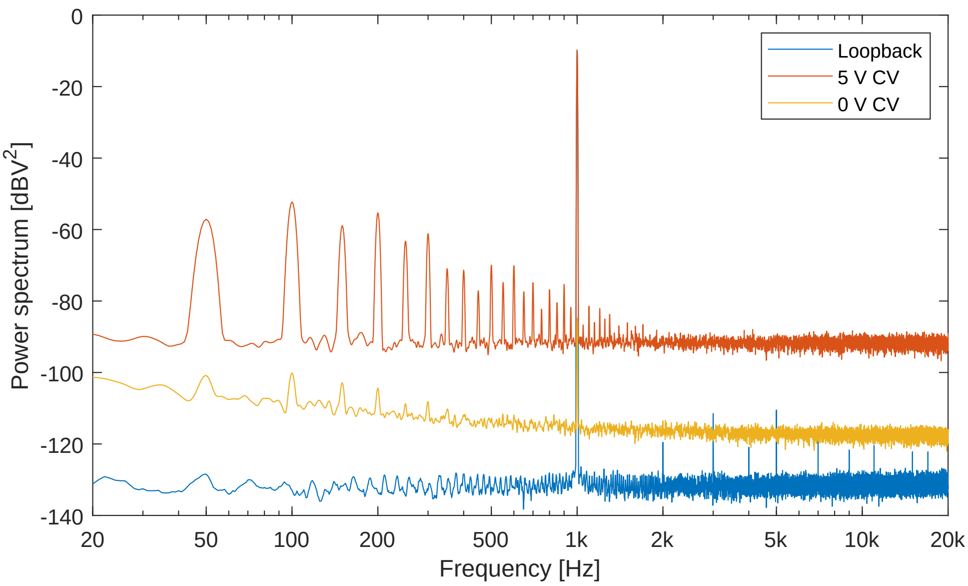

Figure 7.8. Overview of the power spectrum of the VCA with the test audio signal applied. Design extremes +5 V and 0 V are shown, as is loopback spectrum for comparison.

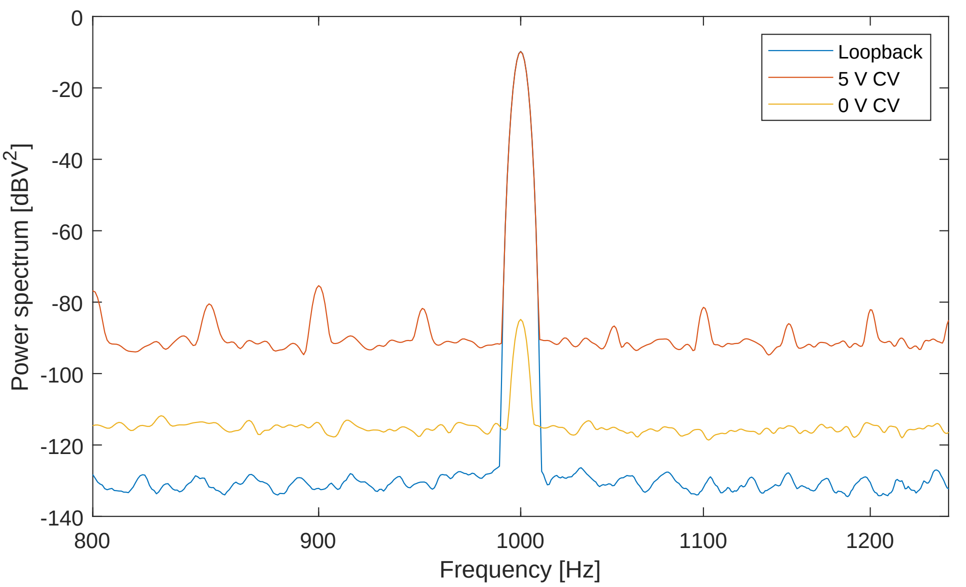

Figure 7.9. Power spectrum of the VCA for the test signal, zoomed in around the test signal frequency. Design extremes +5 V and 0 V are shown, as is loopback spectrum for comparison.

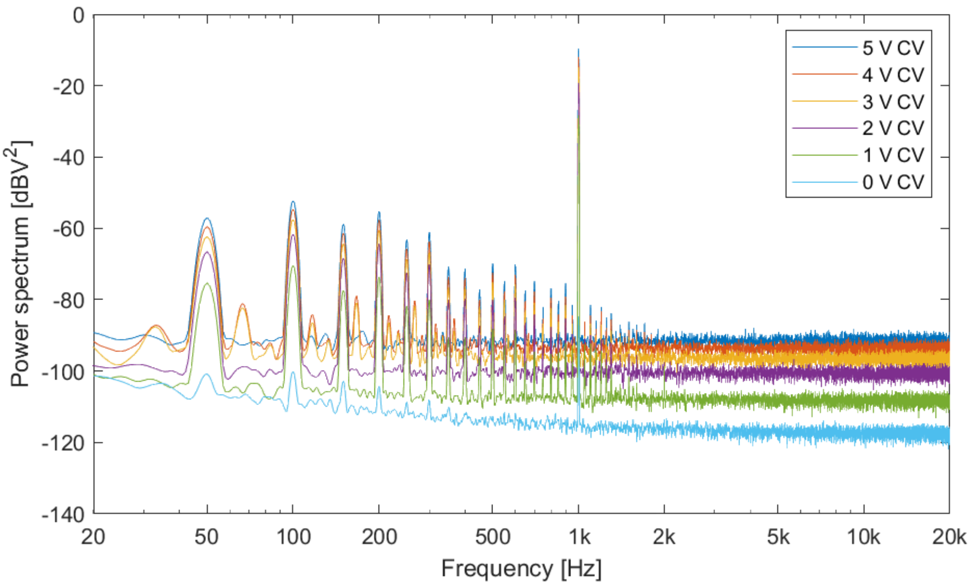

Figure 7.10. Power spectrum of the VCA for the test signal, control voltages between 0 V and 5 V.

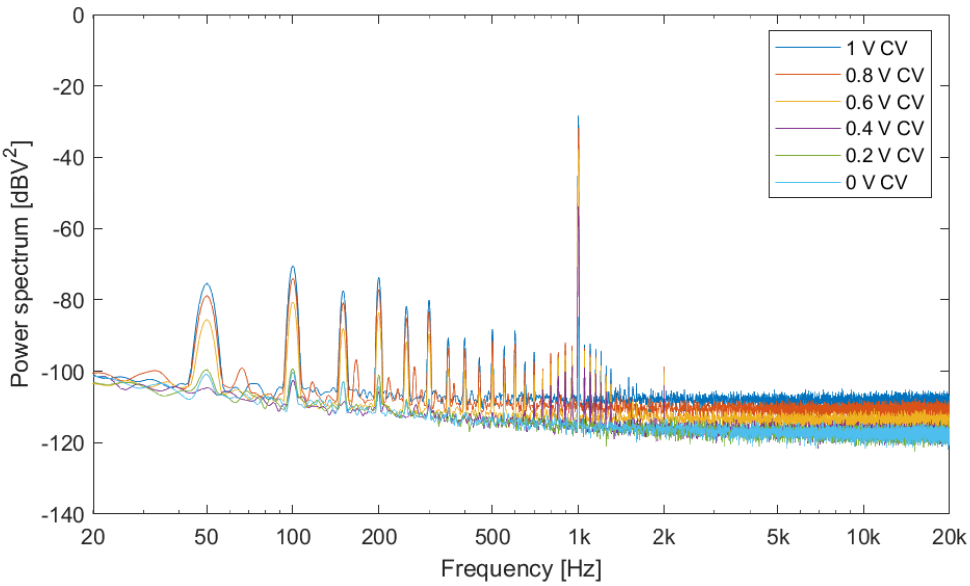

Figure 7.11. Power spectrum of the VCA for the test signal, fractional control voltages between 0 V and 1 V.

From Figures 7.10 and 7.11, it is apparent that both mains frequency harmonic contamination and broadband noise level is heavily dependent on CV. Therefore, the input of the long-tailed pair or that the long-tailed pair stage itself is suspect.

It can be seen in Figure 7.12 that for larger control voltages, amplitude gain exhibits almost linear dependence on CV. However, as perceived signal intensity is based on signal power rather than amplitude, the dependence of perceived intensity on CV is square-root in this region. In the 0..1 V region, there is a soft exponential knee followed by the amplifier closing between 0.2 V and 0.4 V, achieving the maximum 75 dB attenuation.

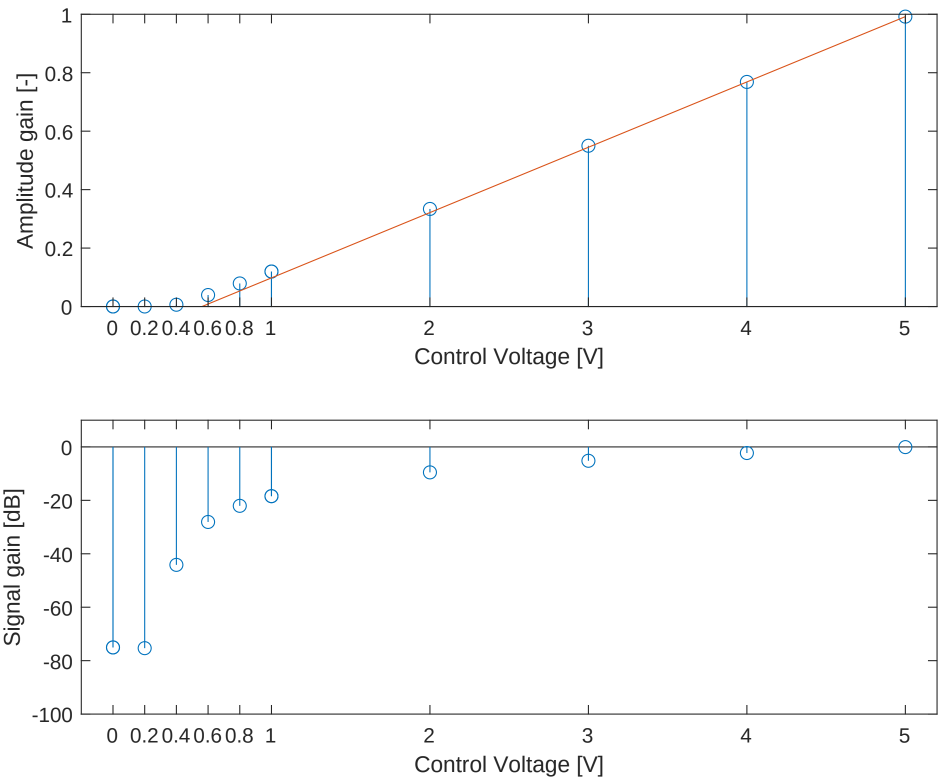

Figure 7.12. VCA gain at various control voltages, measured using the power spectrum at test signal frequency, referenced to loopback. Both linear amplitude gain and logarithmic gains are shown. A straight line is drawn through the two rightmost points in the amplitude plot to show that amplitude gain is almost linear for larger CV values.

As-is, the mains frequency harmonics artifacts are the most problematic, which I confirmed as present by listening tests. While working with Eurorack reference level signals alleviates this problem somewhat, the artifacts are definitely present. In addition to power supply improvements, further power tailoring on this module is advisable, as its susceptibility to power rail ripple is severe due to high attenuation of the input signal.

7.6 Attack-Release Envelope Generator (AR EG)

There are no major problems with AR EG. The generated CV signal features clean linear slopes, as expected. The attack / release time with +1 V CV is approx. 56 ms. In practice, this means the slopes are too fast to notice unless the CV is a fraction of a volt, and the integration should be slowed down in a future redesign. Using the multimeter, I measured the high level to be +5.01 V and low level to be 0.00 V, an excellent adherence to the design interval.

In Figure 7.14, the fastest slope achievable can be seen as well. The high level is reached in 5.7 ms, corresponding to an approx. +10 V CV. Full +12 V CV cannot be used as phase reversal occurs in attack intensity inversion op-amp. All in all, the AR EG performs as expected.

Figure 7.13. Example oscillogram showing the envelope generator gate triggered periodically, with attack CV only slightly above 0 V for low intensity. Release was CV left unconnected, resulting in fast release time. The envelope generator output is shown on Channel 1. The voltage offset of low level is non-indicative as the used oscilloscope had an offset problem. Channel 2 is non-indicative.

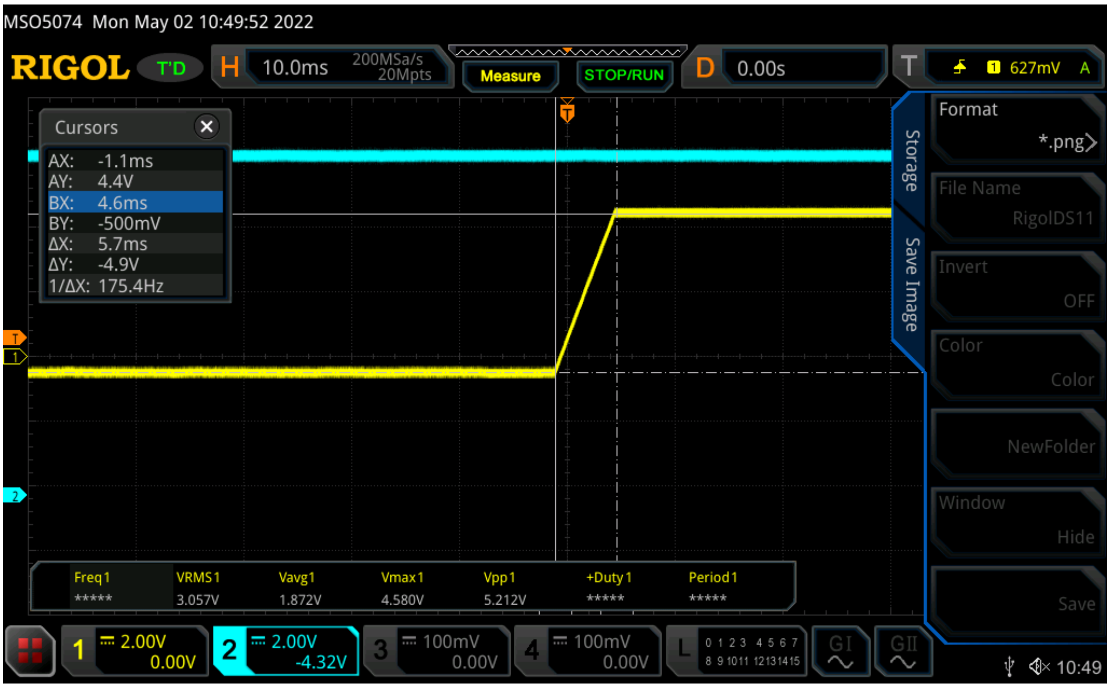

Figure 7.14. Oscillogram of fastest attack slope of the envelope generator, achievable with an approx. +10 V CV. The envelope generator output is shown on Channel 1. The voltage offset of low level is non-indicative as the used oscilloscope had an offset problem. Channel 2 is non-indicative.

7.7 Voltage-Controlled Filter (VCF)

As discussed in Section 5.6, due to a design error, the polarity of the CV cutoff input is inverted. There are severe stationary component changes at the extremes of −8.3..+0.2 V range, presumably due to effects of TL07x phase reversal and/or circuit operation with back-biased diodes. To evaluate the filter, I have measured its output gain using white-noise signal, computing the output gain with respect to a loopback recording of the white noise.

It is apparent from Figure 7.15 that in the −5..−2 V range, there is barely any change in filter response. The filter exhibits a small, approx. 0.6 dB gain reduction in its passband. The −3dB point relative to the passband is approx. 7.5 kHz for −5..−3 V CV and 6.5 kHz for −2 V CV. Clearly, the low-pass behaviour in this region is largely CV-independent.

In Figure 7.16 a much more interesting CV range is shown. Clearly, the filter cutoff is highly dependent on CV in this range. The sloping behaviour is monotonic and gentle, which is expected and desired for an audio synthesizer to prevent successive notes having noticeably different output power, which would be jarring. The contamination of output by the mains frequency and its multiples can be clearly seen in 0 V and −0.2 V cutoff CV responses.

I also tested the noise floor of the VCF by sending constant zero signal to the signal input, as visualised in Figure 7.17. While there is mains frequency harmonics contamination in the output power spectrum, it is not as prevalent as in the VCA by far. However, unlike the VCA, where the mains harmonics are largely attenuated when the amplifier is “off”, the low-pass filter in the VCF attenuates lower harmonics less and they are thus still significant even when the cutoff frequency is low.

The measurements confirm my experience from listening tests where small negative CV had to be used for an interesting filter effect. A redesign should alleviate this problem. In addition, mains harmonics bleeding and the problems with stationary component changes should be further investigated and fixed. Other than that, the filter frequency response is very pleasant, considering its simplicity.

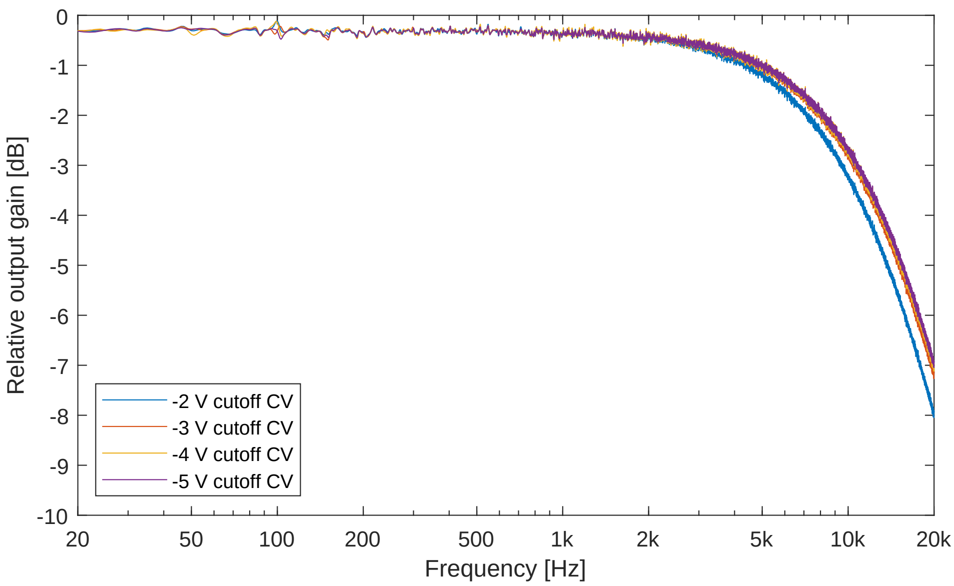

Figure 7.15. Output gain of the VCF with respect to loopback response, large cutoff CV magnitudes.

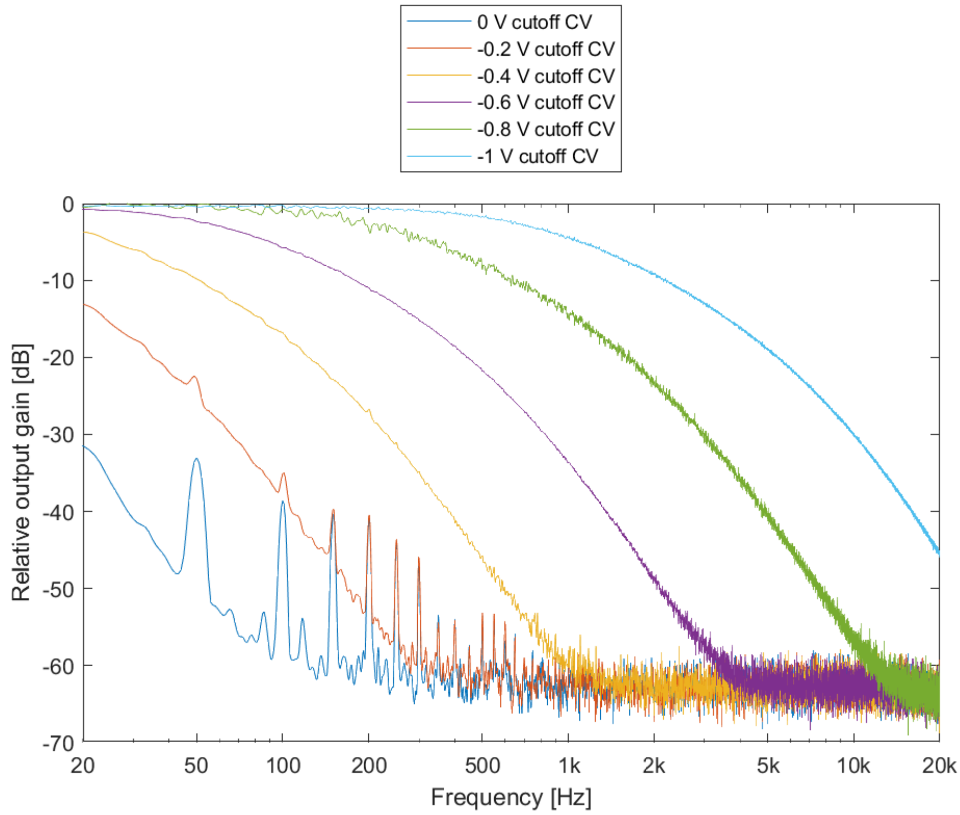

Figure 7.16. Output gain of the VCF with respect to loopback response, small cutoff CV magnitudes.

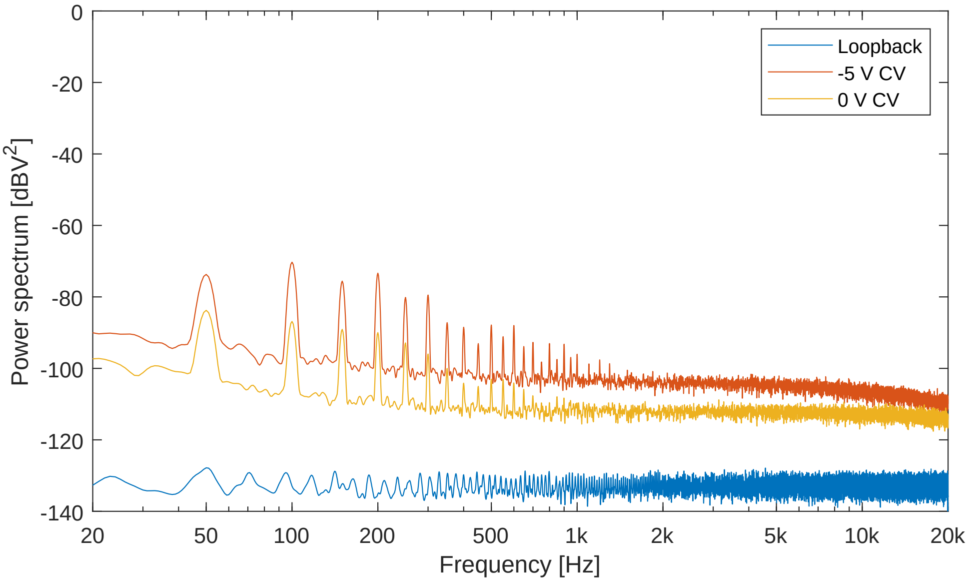

Figure 7.17. Output power spectrum of the VCF with constant zero signal applied.

Chapter 8
Evaluation of the synthesizer

In this chapter, I will evaluate the synthesizer with respect to its cost, design and build complexity, and performance.

8.1 Cost

The synthesizer must be evaluated, above all, with respects to its cost: problems that are forgivable in a toy synthesizer which costs 50 USD are not in a professional synthesizer which costs thousands of USD.

	Enclos.	Connector	Op-amp	Signal	Misc.	**Total**
Power	0	5.64	0	0	9.37	15.01
Daughterb.	0	9.40	0.81	3.19	9.34	22.75
VCO	4.10	6.33	3.25	5.80	5.87	25.35
VCA	3.07	4.35	2.09	1.46	1.88	12.85
AR	3.07	5.62	1.64	3.08	1.69	15.09
VCF	3.07	4.35	2.09	0.37	1.01	10.89
Total	10.24	30.08	8.25	10.82	27.47	**86.86**

Table 8.1. Part costs for every module in US dollars. Converted from CZK using conversion rate 1 USD = 23.825 CZK (Czech National Bank foreign exchange rate on 13-05-2022). “Signal” parts include ICs, matched pair transistors, and diodes in the signal path, excluding op-amps, voltage references or regulators. Cost for 10 Eurorack-style power headers used on power expansion boards was included in the power module connector cost.

I have computed the Bill of Materials (BOM) costs of the modules using the actual costs of parts I bought, as shown in Table 8.1. The costs are for reference purposes only, as they might vary significantly based on the distributor and location, as well as sourcing considerations (if a part is out of stock, a replacement may be significantly more costly, or, in some cases, less costly but with worse parameters). I purchased most of the parts through Mouser Electronics with delivery to the Czech Republic. I built the power module with components purchased previously from Farnell and GM Electronic.

The part costs in Table 8.1 are very optimistic, as it is usually a good idea to purchase one or two additional parts of every type for protection against an unintentional destruction of a part when emplacing it. For small passives, 100 units or more typically have to be bought for a reasonable price per unit (fortunately, my aim to simplify the BOMs considerably mean that not many types of frequently-used passives are needed). This is not taken into account in Table 8.1. The discussed additional costs can raise the total cost of parts for an initial set of modules up to 170 USD, although the costs should taper off if more units of the same module are made.

As for the distribution of costs, the enclosures and connectors amount to 46.4% of total costs. This is despite the comparatively cheap enclosures and connectors used.

Depending on the interpretation, this might lead to considerations that more expensive circuitry could be used in modules if it gives better performance, since the enclosure and connector costs are already paid for.

Additional costs are due to the PCBs, power module and support module perfboards, and the sequencer development board. The boards were manufactured by JLC-PCB. The final price with shipping for 5 units of all module boards was 42.80 USD (generally, only multiple-of-5 unit quantities are possible to order at Chinese PCB manufacturers).

The STM32F723E Discovery Kit (almost exact successor to the STM32F7308 Discovery Kit I used) currently costs approximately 40 USD. Furthermore, some perfboards and miscellaneous parts are needed for the power board and support boards. Thus, the expected total cost for the first set of modules is approx. 250 USD, with some parts and PCBs left over. The cost could be reduced up to approx. 150 USD per set when building multiple sets of modules. The synthesizer is definitely more expensive than “toy” synthesizers, but the cost is not severe enough to discourage experimentation.

Additional costs I did not consider in the total cost computation are the costs of tools required for the construction of modules. A good temperature-controlled soldering iron with thin solder and desoldering braid is necessary, as is a rotary tool and some miscellaneous tools (screwdriver, tweezers, pliers, wire strippers. . .). Additionally, an AVR ISP programmer is required to program the VCO trimming microcontroller (the Discovery Kit is programmable using standard USB connection). Of course, with the exception of the solder and desoldering braid, these items are a non-consumable investment.

8.2 Complexity

There are two kinds of “complexity” present in the synthesizer, design complexity and build complexity.

Design complexity. Varies significantly between the modules. The small modules (VCA, AR EG, and VCF) are based on standard op-amp arrangements with a single core technique (long-tailed pair, integration with limiting, diode ladder). Understanding of their function is therefore based on understanding of the core technique. Hopefully, I have written enough about the techniques in Chapter 3 that they are at least a bit understandable. I feel that it is important that the techniques are not hidden in ICs. Instead, properties of diodes, transistors and standard op-amps are used.

The design complexity of the VCO is large, especially due to my decision to include a digital trimming function. While the digital trimming performs well (nothwithstanding the very limited tuning stability of the VCO as implemented), I feel that the complexity introduced is not worth it. In future designs, I would probably use the designs based on exponential converter for non-critical low-frequency oscillators (LFOs) only, with audio VCOs replaced by digitally controlled oscillators.

As for the sequencer daughterboard, there are two functionalities mixed together: I/O expansion and input level translation. Both of these requirements were born out of necessity and are largely self-explanatory. The most complex part is the sequencer firmware, which has to be written carefully and with a high level of understanding of the STM32F7 microcontroller. However, this simply is an unavoidable part of digital audio processing using microcontrollers: the amount of data to be processed is high, so a relatively performant microcontroller that includes an audio interface peripheral must be used. While an ARM-based Cortex-M platform is rather complex by itself, it it a best choice for the sequencer core nowadays due to its excellent price-value proposition.

Build complexity. The modules can be soldered by hand easily. I used a Weller WE1010 soldering station with good results. Thin Shrink Small Outline Package

(TSSOP) form factor ICs are a bit problematic, best drag-soldered with solder bridges removed using a desoldering braid. Unfortunately, hard-to-spot bridges may still be present afterwards. To ensure that there are no inadvertent shorts, continuity testing should be used, which can become tiresome. Small Outline Integrated Circuit (SOIC) form factor ICs, on the other hand, are extremely easy to solder correctly. Unfortunately, their larger size may prevent the designer from adding more features to the space-limited PCBs.

Small low pin count packages were less problematic to solder than I initially expected. However, care must be taken to handle them carefully as their small size makes their pins more prone to bending and/or separating. This is especially problematic if few replacements are available. Using a larger package should be considered where possible.

I chose the imperial 0805 (metric 2012) form factor for resistors and regular capacitors. I was wary about this, but it did not pose a problem in practice. The imperial 0603 (metric 1608) form factor also could be considered if more integration is needed. Below that size, I think that the soldering difficulty would be too high for a regular hobbyist, even with good tools.

Drilling and cutting the enclosures is not problematic in practice. The results are not perfect, but they are good enough. Labeling the “front panel” with self-adhesive paper that has holes punched through with a classic paper hole punch hides irregularities in the drilled 3.5 mm phone socket holes quite well.

The modules as designed were largely interesting and fun to construct. The small modules (VCA, AR EG, VCF) in particular feature few BOM parts, making it possible to emplace the whole board in an hour and immediately test the module.

8.3 Performance

As discussed in Chapter 7, design errors were made in most of the modules, which could be rectified in a later revision. Otherwise, the modules fulfill their purposes. Thus, in this section, I will mostly document my subjective opinions about the modules.

The chosen form factor, modules inside small plastic enclosures, is ideal for easy construction of simple modules without advanced tools, but the low rigidity of the completed synthesizer makes it hard to use it as a serious music production tool. Modules placed at the top of each other tend to slide away when inserting patch cables. I partly alleviated this problem by taping the modules together with masking tape, but this is not a complete solution. As the module power connectors are regularly spaced, a custom PCB could be made with female headers so that they would replace the IDC cables, resulting in much better (although probably not total) stability. This would also reduce the number of ungainly IDC cables and power connector hubs.

The modules not in enclosures (power module, sequencer, support boards) are awkward to use. For serious use, the sequencer in particular should definitely be redesigned as a single enclosed module, or a professional control module should be used. In particular, there is a chance of inadvertently touching the development kit with a patch cable, which could result in its destruction.

However, after overcoming the awkwardness of the control module, I was pleasantly surprised by the interesting sounds that can be made using the synthesizer. Recordings of some of them are attached to this thesis, as discussed in Appendix B. The more traditional examples are:

- **Pitch control.** Sequencer pitch CV output is routed to VCO CV input. One of the VCO outputs routed through the sequencer level translator to a sound card with headphone monitoring. Continuous sound can be heard, the pitch controlled by the sequencer.

- **Vibrato.** Sequencer LFO output with additional offset is routed into the VCO input. The result is vibrato (continuous change of pitch).
- **Organ-style envelope.** VCO output is routed through the VCA. Sequencer gate output is routed to VCA CV input. The result is an on-and-off envelope, similarly to an organ.
- **Tremolo.** Sequencer LFO output with additional offset is routed to VCA input. The result is tremolo (continuous change of amplitude).
- **Attack-release envelope.** Sequencer gate output is routed to AR EG gate input. AR EG output is routed to VCA level input. Attack and release intensities are controlled externally, resulting in different shapes of the generated envelope that is applied by VCA.
- **Low-pass filtering.** VCO output is routed through the VCF. VCF gate voltage is controlled externally. The result is low-pass filtering with selectable cutoff.

The amount of possibilities is large. For even more interesting results, more modules would be needed (in my opinion, at least a second set of the VCA, AR EG, and VCF modules). In addition, control modules in enclosures (voltage output selectable by potentiometer, summation modules, potentiometer-controlled gain change...) would definitely improve sonic possibilities, as well as handling. Lastly, a breakout board with integrated power outputs and signal inputs / outputs which could be inserted into a breadboard would allow an unlimited amount of experimentation and demonstration of electronics by building new modules right before the audience.

8.4 Overall

The synthesizer as designed is not a “toy” synthesizer, neither a professional-level one. It seems at its best as an electronics demonstration and learning aid, which was the design goal. The modules are easy to assemble by hand and can be even used without enclosures. The single-purpose nature of each module makes it relatively easy to understand its function and perhaps even design others.

The major drawbacks of the synthesizer in its current form are definitely its low-quality output, contaminated by mains frequency harmonics, and the handling awkwardness. A future redesign should alleviate or eliminate these problems. If the drawbacks were to be rectified, I would consider the cost of the synthesizer, around 250 USD for all parts and PCBs necessary, very justifiable.

Chapter 9
Conclusion

My task was to research the techiques used for construction of audio synthesizers and their modules, design my own modular synthesizer for electronics demonstration purposes and evaluate its properties. In Chapter 1, I introduced the theme of the thesis and my goals. In Chapter 2, I discussed the considerations for synthesizer construction and common categorizations of synthesizers, and discussed my reasoning for building an analog subtractive synthesizer with a digital control module. In Chapter 3, I presented common subtractive synthesizer modules and circuits commonly used for their construction. In Chapter 4, I discussed my choices for implementation of common module sections, and discussed my implementation of the specific modules in Chapter 5, augmenting common techniques with my own design ideas. In Chapter 6, I discussed firmware I developed for microcontrollers present in some of the modules. In Chapter 7, I measured objective parameters of the synthesizer. Finally, in Chapter 8, I evaluated the synthesizer according to various criteria, taking into account the costs of the synthesizer modules, objective measurement results, and my subjective opinions about the synthesizer and its usability.

The task of constructing a modular synthesizer from the ground up proved to be rather complex and time-consuming, as each module has a different function and design requirements, resulting in a multitude of similar yet different circuit techniques used to resolve them. Unfortunately, many of these techniques are largely confined to DIY audio synthesizer usage and there are few reference sources. Understanding is achieved largely by piecing together various sources on the Internet with common circuit analysis techniques.

Despite the problems encountered, I was able to design and succesfully construct a modular synthesizer of my own, electrically largely compatible with the common Eurorack standard for modular synthesizers. All of the built modules are fully functional, though some of their parameters are problematic due to errors and oversights I made during construction. I plan to fix them in the future. I will provide links to future work on the synthesizer format introduced in this thesis in [35], so it hopefully can be used by others in the future as well.

Appendix A
Thesis assignment

CTU CZECH TECHNICAL UNIVERSITY IN PRAGUE

BACHELOR'S THESIS ASSIGNMENT

I. Personal and study details

Student's name: **Onderka Jan** Personal ID number: **457971**

Faculty / Institute: **Faculty of Electrical Engineering**

Department / Institute: **Department of Microelectronics**

Study program: **Electrical Engineering, Electronics and Communications**

II. Bachelor's thesis details

Bachelor's thesis title in English:

Analog Modular Music Synthesizer with a Digital Control

Bachelor's thesis title in Czech:

Analogový modulární hudební syntezátor s digitálním kontrolérem

Guidelines:

Research and summarize commonly used techniques for constructing music synthesizer instruments and their classic modules. Design and implement an analog modular synthesizer with a digital control module that may be used for electronics demonstration purposes. Evaluate the synthesizer's performance, cost, and complexity of the design.

Bibliography / sources:

[1] WILSON, Ray. Make: Analog Synthesizers: Make Electronic Sounds the Synth-DIY Way. Maker Media, Inc., 2013.
[2] HAYES, Thomas C.; HOROWITZ, Paul. Learning the art of electronics. Learning the Art of Electronics, 2016.
[3] HORN, Delton T. Digital Electronic Music Synthesizers. Tab Books, 1988.

Name and workplace of bachelor's thesis supervisor:

doc. Ing. Stanislav Vítek, Ph.D., Department of Radioelectronics, FEE

Name and workplace of second bachelor's thesis supervisor or consultant:

Date of bachelor's thesis assignment: **31.01.2022** Deadline for bachelor thesis submission: __________

Assignment valid until: **30.09.2023**

[redacted] [redacted] [redacted]

doc. Ing. Stanislav Vítek, Ph.D. — Supervisor's signature

prof. Ing. Pavel Hazdra, CSc. — Head of department's signature

prof. Mgr. Petr Páta, Ph.D. — Dean's signature

III. Assignment receipt

The student acknowledges that the bachelor's thesis is an individual work. The student must produce his thesis without the assistance of others, with the exception of provided consultations. Within the bachelor's thesis, the author must state the names of consultants and include a list of references.

9.5.2022 — Date of assignment receipt

[redacted] — Student's signature

Appendix B
List of attachments

Complete module schematics and layouts, as well as the firmware, are available as attachments to this thesis. On physical copies, they are contained on a Secure Digital card attached to the inside back cover. They should also be available for electronic download where the electronic copy of this thesis is published.

In addition to the above required locations where the attachments are fixed and cannot be changed, I also published them in [35] where I can make updates in the future, e.g. to inform about a next version of the modules. That attachment location should therefore be preferred by all readers except the ones evaluating this thesis.

The directory structure of the attachments is as follows:

- `firmware` .... Firmware folder.
  - `sequencer` .... Sequencer firmware.
  - `vco_trim` .... Firmware for VCO trimming.
- `modules` .... Module schemas and layouts.
- `recordings` .... Sample synthesizer recordings.
- `thesis` .... Source files necessary for creation of this thesis.

Appendix C
List of acronyms and abbreviations

AC	Alternating Current
ADC	Analog-to-Digital Converter
ADSR	Attack-Decay-Sustain-Release (envelope)
AR	Attack-Release (envelope)
BJT	Bipolar Junction Transistor
BOM	Bill of Materials
CV	Control Voltage
CZK	Czech koruna
DAC	Digital-to-Analog Converter
DC	Direct Current
Digipot	Digital Potentiometer
DMA	Direct Memory Access
DYI	Do It Yourself
EG	Envelope Generator
FET	Field-Effect Transistor
FIFO	First-In-First-Out
FM	Frequency Modulation
I/O	Input / Output
I2C	Inter-Integrated Circuit
IC	Integrated Circuit
IDC	Insulation-Displacement Contact
JFET	Junction Field-Effect Transistor
LCD	Liquid Crystal Display
LED	Light-Emitting Diode
LFO	Low-Frequency Oscillator
MCU	Microcontroller Unit
Op-amp	Operational Amplifier
OTA	Operational Transconductance Amplifier
PCB	Printed Circuit Board
PGA	Programmable Gain Amplifier
ppm	Parts Per Million
PWM	Pulse-Width Modulation
RC	Resistor-Capacitor
RMS	Root Mean Square
SAI	Serial Audio Interface
SAR	Successive-Approximation
SD	Secure Digital
SNR	Signal-to-Noise Ratio
SO	Small Outline
SOIC	Small Outline Integrated Circuit
SPDT	Single-Pole Double-Throw (switch)

SPI	Serial Peripheral Interface
TDM	Time-Division Multiplexing
TRS	Tip-Ring-Sleeve
TSSOP	Thin Shrink Small Outline Package
USB	Universal Serial Bus
USD	United States Dollar
V/Oct	Volts per octave
VCA	Voltage-Controlled Amplifier
VCF	Voltage-Controlled Filter
VCO	Voltage-Controlled Oscillator
Vpp	Peak-to-peak voltage
Vrms	Root Mean Square voltage

References

[1] R. Wilson. *Make: Analog Synthesizers: Make Electronic Sounds the Synth-DIY Way*. Maker Media, Inc.„ 2013. ISBN 9781449345198.

[2] Doepfer Musikelektronik GmbH. *Technical Details A-100*. `https://doepfer.de/a100_man/a100t_e.htm`. Retrieved 2022-03-09.

[3] Mark Vail. *The Synthesizer: A Comprehensive Guide to Understanding, Programming, Playing, and Recording the Ultimate Electronic Music Instrument*. Cary, United States: Oxford University Press, Incorporated, 2014. ISBN 0195394895.

[4] Mendel Kleiner. *Acoustics and Audio Technology*. 3rd edition. J. Ross Publishing, Inc., 2012. ISBN 978-1-60427-052-5.

[5] Archimago. *24-Bit vs. 16-Bit Audio Test - Part II: RESULTS & CONCLUSIONS*. `http://archimago.blogspot.com/2014/06/24-bit-vs-16-bit-audio-test-part-ii.html`. Retrieved 2022-05-15.

[6] Malcolm J. Crocker. *Handbook of Acoustics*. John Wiley & Sons, 1998. ISBN 978-0-471-25293-1.

[7] Rod Elliott (Elliott Sound Products). *VCA Techniques Investigated*. 2012. `https://sound-au.com/articles/vca-techniques.html`. Retrieved 2022-03-18.

[8] Chris Meyer. *Analog*. 2016. `https://learningmodular.com/glossary/analog/`. Retrieved 2022-05-19.

[9] Don Tillman. *Last of the OTA's*. 2005. `http://www.till.com/blog/archives/2005/06/last_of_the_ota.html`. Retrieved 2022-05-19.

[10] Peter Elsea. *The Art and Technique of Electroacoustic Music*. Middleton (WI), United States: A-R Editions, Inc., 2013.

[11] Adam Kagan. *Should I Be Oversampling?* 2021. `https://www.sonarworks.com/blog/learn/should-i-be-oversampling`. Retrieved 2022-05-19.

[12] Chris Meyer. *East Coast Synthesis*. 2016. `https://learningmodular.com/glossary/east-coast-synthesis/`. Retrieved 2022-05-19.

[13] Chris Meyer. *West Coast Synthesis*. 2016. `https://learningmodular.com/glossary/west-coast-synthesis/`. Retrieved 2022-05-19.

[14] Electronic Music Wiki. *Format*. `https://electronicmusic.fandom.com/wiki/Format`. Retrieved 2022-05-15.

[15] Electronic Music Wiki. *Dotcom*. `https://electronicmusic.fandom.com/wiki/Dotcom`. Retrieved 2022-05-15.

[16] Texas Instruments. *TL07xx Low-Noise FET-Input Operational Amplifiers*. 2021. `https://www.ti.com/lit/ds/symlink/tl074.pdf`. Retrieved 2022-05-13.

[17] North Coast Synthesis Ltd. *What's the deal with phase reversal?* 2019. https://northcoastsynthesis.com/news/whats-the-deal-with-phase-reversal/. Retrieved 2022-05-13.

[18] Doepfer Musikelektronik GmbH. *A-100 Analog Modular System - Accessories and Spare Parts*. https://doepfer.de/a100z_e.htm. Retrieved 2022-03-15.

[19] Electrosmash. *Klon Centaur Analysis*. https://www.electrosmash.com/klon-centaur-analysis. Retrieved 2022-03-15.

[20] chillibasket: ARobotics, and Technology Blog. *Modular Synth – Dual 12V Power Supply*. https://wired.chillibasket.com/2020/06/dual-power-supply/. Retrieved 2022-05-03.

[21] Thomas Henry. *VCO-1*. https://www.birthofasynth.com/Thomas_Henry/Pages/VCO-1.html. Retrieved 2022-05-07.

[22] North Coast Synthesis Ltd. *Exponential converters and how they work*. https://northcoastsynthesis.com/news/exponential-converters-and-how-they-work/. Retrieved 2022-05-07.

[23] openmusiclabs. *Thermal Compensation of Analog Exponential Converters*. http://www.openmusiclabs.com/files/expotemp.pdf. Retrieved 2022-05-07.

[24] synthnerd. *Envelope Circuits: a simple AR design using op amps*. 2016. https://synthnerd.wordpress.com/2016/04/06/envelope-circuits-a-simple-ar-design-using-op-amps/. Retrieved 2022-05-14.

[25] René Schmitz. *Voltage controlled envelope generator*. https://www.schmitzbits.de/vcadsr.html. Retrieved 2022-05-14.

[26] Timothy E. Stinchcombe. *Analysis of the Moog Transistor Ladder and Derivative Filters*. 2019. https://www.timstinchcombe.co.uk/synth/Moog_ladder_tf.pdf. Retrieved 2022-05-14.

[27] Moritz Klein. *Designing a diode ladder filter from scratch*. https://youtu.be/jvNNgUl3al0. Retrieved 2022-05-05.

[28] Hammond Manufacturing Corporation. *1551KTBU*. https://www.hammfg.com/part/1551KTBU. Retrieved 2022-05-12.

[29] Hammond Manufacturing Corporation. *1551XTBU*. https://www.hammfg.com/part/1551XTBU. Retrieved 2022-05-12.

[30] Nebojsa Petrovic. *Eurocad*. https://github.com/nebs/eurocad. A library of Eurorack components and footprints for KiCad. Retrieved 2022-05-12.

[31] Cirrus Logic. *WM8994 Product Datasheet, Revision 4.6*. https://statics.cirrus.com/pubs/proDatasheet/WM8994_Rev4.6.pdf. Retrieved 2022-05-07.

[32] Spehro Pefhany. *Circuit for a coarse and fine setting potentiometer?* Electrical Engineering Stack Exchange. https://electronics.stackexchange.com/q/144540. Version 2014-12-19, retrieved 2022-05-07.

[33] Moritz Klein. *Designing a classic transistor-VCA from scratch*. https://youtu.be/yMrCCx6uqcE. Retrieved 2022-05-06.

[34] Microchip. *ATtiny24A/44A/84A tinyAVR Data Sheet*. `https://ww1.microchip.com/downloads/en/DeviceDoc/ATtiny24A-44A-84A-DataSheet-DS40002269A.pdf`. Retrieved 2022-05-07.

[35] Jan Onderka. *Modular Analog Synthesizer with Digital Sequencer: Attachments to the Bachelor Thesis*. `https://doi.org/10.6084/m9.figshare.19729501`. Retrieved 2022-05-07.