
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Paschou, Effrosyni; Esqueda Flores, Fabian; Välimäki, Vesa; Mourjopoulos, John
Modeling and measuring a Moog voltage-controlled filter

Published in:

Proceedings of the Ninth Annual Conference of the Asia-Pacific Signal and Information Processing Association

DOI:

[10.1109/APSIPA.2017.8282295](https://doi.org/10.1109/APSIPA.2017.8282295)

Published: 01/01/2017

Document Version

Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:

Paschou, E., Esqueda Flores, F., Välimäki, V., & Mourjopoulos, J. (2017). Modeling and measuring a Moog voltage-controlled filter. In *Proceedings of the Ninth Annual Conference of the Asia-Pacific Signal and Information Processing Association* (pp. 1641 - 1647). IEEE. <https://doi.org/10.1109/APSIPA.2017.8282295>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Modeling and Measuring a Moog Voltage-Controlled Filter

Effrosyni Paschou^{*†}, Fabián Esqueda^{*}, Vesa Välimäki^{*}, and John Mourjopoulos[†]

^{*}Acoustics Lab, Dept. Signal Processing and Acoustics, Aalto University, Espoo, Finland

[†]Wire Communications Laboratory, University of Patras, Patras, Greece

Abstract—The analog voltage-controlled filter used in historical music synthesizers by Moog is modeled using a digital system, which is then compared in terms of audio measurements with the original analog filter. The analog model is mainly borrowed from D’Angelo’s previous work. The digital implementation of the filter incorporates a recently proposed antialiasing method. This method enhances the clarity of output signals in the case of large-level input signals, which cause harmonic distortion. The combination of these two ideas leads to a novel digital model, which represents the state of the art in virtual analog musical filters. It is shown that without the antialiasing, the output signals in the nonlinear regime may be contaminated by undesirable spectral components, which are the consequence of aliasing, but that the antialiasing technique suppresses these components sufficiently. Comparison of measurements of the analog and digital filters show that the digital model is accurate within a few dB in the linear regime and has very similar behavior in the nonlinear regime in terms of distortion. The proposed digital filter model can be used as a building block in virtual analog music synthesizers.

I. INTRODUCTION

Robert Moog developed for his music synthesizers a voltage-control analog filter (VCF) [1], which is one of the most influential innovations in music technology. The Moog VCF is a fourth-order resonant lowpass filter with self-oscillating capabilities built around a transistor “ladder” topology [1]. The filter can be made time-varying using two control parameters that independently adjust the two main parameters of the filter’s response, the cutoff frequency and the resonance, i.e., the Q value. This made it possible to generate interesting sounds using subtractive synthesis. Another striking feature of the Moog ladder filter is its distinctive “warm” sound, which can be attributed to the inherent nonlinear behavior of the circuit. The time-variant and distorting characteristics of the Moog filter have become a crucial part of electronic music.

In the era of digital audio, a desire to replace old analog technology with software started to evolve, which led to the development of virtual analog (VA) modeling [2], [3]. For the case of analog VCFs, the digitizing started gradually several decades ago [4], [5] and has continued ever since. More recently, audio researchers have devised digital models for many other analog musical devices, such as synthesizer oscillators [6], guitar pedals [7]–[9], distortion circuits [10]–[12], audio transformers [13] and ring modulators [14].

Stilson and Smith were the first to propose a linear digital model of the Moog filter [5]. Huovilainen derived a different digital model, which accounts for the nonlinear characteristics

of the transistor ladder [15]. Huovilainen and Välimäki simplified the nonlinear Moog filter model so that it only uses one memoryless mapping function [6]. Other authors have derived digital variations of both the linear Moog filter [16] and the nonlinear one [17]–[21]. Recently, also other analog synthesizer filters have been the subject of digital modeling [22], [23].

Nonlinear audio processing is prone to aliasing artifacts. They are caused by the fact that the spectral components caused by nonlinear distortion can appear at frequencies higher than the Nyquist limit (i.e., half of the sampling frequency) and they are then reflected down to the audio band leading to inharmonic disturbances. A traditional method to suppress aliasing is to increase the Nyquist limit by using oversampling. However, oversampling is a very inefficient method, as it requires extra filtering (decimation and interpolation filters) and it multiplies the number of operations per sample. As an alternative, sophisticated signal processing methods have been developed for antialiasing. Esqueda et al. have shown how an approximation of a bandlimited ramp function can be applied to the reduction of aliasing in hard clipping and rectification of signals [24], [25]. Recently, antialiasing methods based on the antiderivatives of the nonlinear mapping function have been proposed [26], [27].

In this paper we investigate how well a digital model of the Moog ladder filter, which incorporates the nonlinear behavior of transistors, corresponds to the analog ladder filter. The analytical analog model is based on the work of D’Angelo and Välimäki [20], [21]. Two digital implementations are then compared: a trivial (i.e. non-compensated for aliasing) version using the implicit midpoint rule [18], [28] and an antialiased form based on the method proposed in [26] and [27]. Measurements of the frequency response of the system in the linear regime and the distortion in the nonlinear regime allow to evaluate the similarities and differences.

The rest of this paper is organized as follows: Section II describes the analog Moog filter and introduces the digital models used in this work; Section III presents measurement results from the analog and digital filters, comparing their spectral characteristics; Section IV concludes this paper.

II. THE MOOG LADDER FILTER

The Moog transistor ladder utilized the base-emitter resistance of bipolar junction transistors to implement buffered

voltage-controlled RC sections [15]. By cascading four identical filtering stages, the circuit is capable of implementing a lowpass filter with a 24-dB/octave rolloff. To control the resonance of the filter, part of the output is fed back to the input stage.

A. Analog Large-Signal Model

One of the most desirable features of the Moog ladder filter is its distortion, which can be attributed to the inherent nonlinear behavior of the circuit. Specifically, this nonlinear behavior is caused by the exponential current-voltage (I-V) relationship of p-n junctions, which causes the signal at the input of each filter stage to saturate. Saturating distortion introduces harmonic components that modify the timbre of the output sound. In the case of the Moog, the level of the introduced harmonics is dependent on the amplitude of the input signal.

In order for digital simulations of the Moog filter to faithfully preserve the nonlinear behavior of the original circuit, the exponential I-V curve of the transistors must be incorporated into the model [15]. In this work, we will use the large-signal continuous-time model presented by D'Angelo and Välimäki [16]. This model is given by the following system of ordinary differential equations (ODEs):

$$\frac{1}{w_c} \frac{dV_1}{dt} = -\tanh(AV_1) - \tanh[A(V_{in} + kV_4)], \quad (1)$$

$$\frac{1}{w_c} \frac{dV_2}{dt} = \tanh(AV_1) - \tanh(AV_2), \quad (2)$$

$$\frac{1}{w_c} \frac{dV_3}{dt} = \tanh(AV_2) - \tanh(AV_3), \quad (3)$$

$$\frac{1}{w_c} \frac{dV_4}{dt} = \tanh(AV_3) - \tanh(AV_4), \quad (4)$$

where V_1 , V_2 , V_3 , and V_4 are the states of the filter, V_{in} (in V) is the input signal, $A = 1/2V_T$, $V_T \approx 26$ mV is thermal voltage, and parameter $k \in [0, 4]$ is the feedback coefficient. The output of the system is extracted from the last filter state, i.e. V_4 . Parameter w_c determines the cutoff frequency of the system:

$$w_c = 4\pi V_T f_c, \quad (5)$$

where f_c is the cutoff frequency in Hz.

The system of equations (1)–(4) can be represented by the block diagram shown in Fig. 1a. Blocks H_{1-4} represent the filtering stages. Fig. 1b and 1c show the structure of the first three stages (H_{1-3}) and the last one (H_4), respectively. It must be noted that all four filtering stages are identical, but the building blocks of H_4 are arranged differently from those of H_{1-3} in order to be in accordance with the general block diagram of Fig. 1a and thus, with the ODEs of the model.

B. Digital Model

Equations (1)–(4) describe a set of implicit relations between signal values V_{1-4} . In the digital domain, implicit relations are known as “delay-free loops” and are quite common

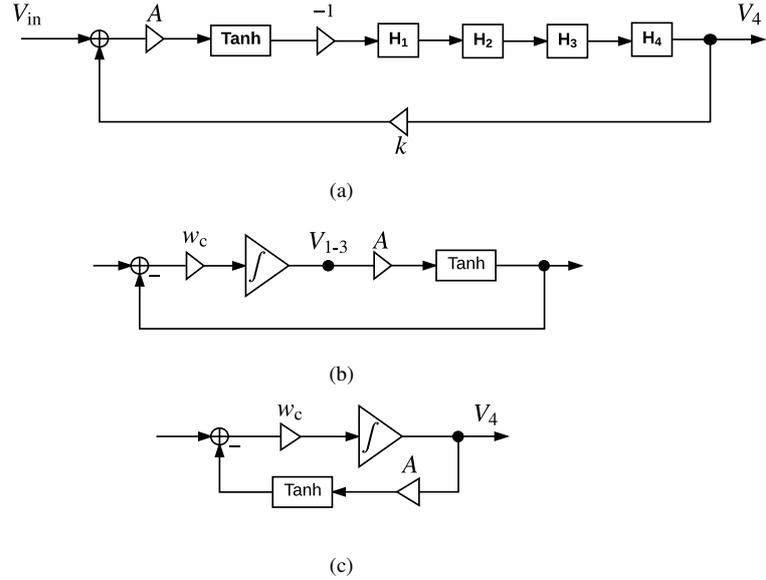


Fig. 1: Block diagram representation of (a) the Moog ladder filter, and (b) the first three (H_{1-3}) and (c) last (H_4) filtering stages.

in VA applications. Delay-free loops can be solved trivially for the case of linear systems using standard algebraic techniques without affecting the response of the system. For the case of nonlinear systems like the one described by (1)–(4) this process becomes nontrivial. A straightforward approach is to use an explicit discretization scheme such as Euler’s method [28]. This process is presented in [15] and is equivalent to inserting fictitious unit delays in the system feedback paths. Unfortunately, this approach does not preserve the linear response of the system and imposes high constraints on its sampling rate. Recently, D’Angelo and Välimäki proposed the use of a network of linear filters to account for the effects of using explicit methods [21].

In this work we first derive a digital implementation of the Moog ladder filter without antialiasing using the implicit midpoint rule, also known as the midpoint method, and approximate the solution of the multiple filter states using an iterative solver. The midpoint method is A -stable which means its region of absolute stability comprises the entirety of the left half-plane [28]. This approach was originally proposed within the context of the Moog VCF in [18] and expanded for real-time implementation in [19]. Iterative solvers such as Newton-Raphson have been previously used in the context of VA modeling [8]. Fontana and Civolani presented a similar treatment for the filter inside the EMS VCS3 synthesizer [22], [29]. In this work we use superscripts to denote the sample index n , i.e. $V_{in}^n \equiv V_{in}[n]$.

Applying the midpoint rule to (1)–(4) gives us

$$V_1^n = V_1^{n-1} - w_c T [\tanh(S_1) + \tanh(S_0)], \quad (6)$$

$$V_2^n = V_2^{n-1} + w_c T [\tanh(S_1) - \tanh(S_2)], \quad (7)$$

$$V_3^n = V_3^{n-1} + w_c T [\tanh(S_2) - \tanh(S_3)], \quad (8)$$

$$V_4^n = V_4^{n-1} + w_c T [\tanh(S_3) - \tanh(S_4)], \quad (9)$$

where T is the sampling period (i.e. the inverse of the sampling rate f_s), and

$$S_0 = A [V_{\text{in}}^n + V_{\text{in}}^{n-1} + k(V_4^n + V_4^{n-1})] / 2, \quad (10)$$

$$S_1 = A (V_1^n + V_1^{n-1}) / 2, \quad (11)$$

$$S_2 = A (V_2^n + V_2^{n-1}) / 2, \quad (12)$$

$$S_3 = A (V_3^n + V_3^{n-1}) / 2, \quad (13)$$

$$S_4 = A (V_4^n + V_4^{n-1}) / 2. \quad (14)$$

The value of the states of the filter at time step n can be approximated using e.g. the Newton-Raphson method [18], [28]. To do so, we rewrite (6)–(9) in the form $\mathbf{F}(\mathbf{V}) = \mathbf{0}$, where $\mathbf{V} = [V_1^n \ V_2^n \ V_3^n \ V_4^n]^T$ is the states vector we wish to find and

$$\mathbf{F} = \begin{cases} V_1^n - V_1^{n-1} + w_c T [\tanh(S_1) + \tanh(S_0)] = 0 \\ V_2^n - V_2^{n-1} - w_c T [\tanh(S_1) - \tanh(S_2)] = 0 \\ V_3^n - V_3^{n-1} - w_c T [\tanh(S_2) - \tanh(S_3)] = 0 \\ V_4^n - V_4^{n-1} - w_c T [\tanh(S_3) - \tanh(S_4)] = 0. \end{cases} \quad (15)$$

We can then approximate the value of \mathbf{V} by iterating over

$$\mathbf{V}_{j+1} = \mathbf{V}_j - (\mathbf{D}\mathbf{F}(\mathbf{V}_j))^{-1} \mathbf{F}(\mathbf{V}_j), \quad (16)$$

where \mathbf{V}_0 is an initial guess and $\mathbf{D}\mathbf{F}$ is the Jacobian matrix

$$\mathbf{D}\mathbf{F} = \begin{bmatrix} 1 + gJ_1 & 0 & 0 & kgJ_0 \\ -gJ_1 & 1 + gJ_2 & 0 & 0 \\ 0 & -gJ_2 & 1 + gJ_3 & 0 \\ 0 & 0 & -gJ_3 & 1 + gJ_4 \end{bmatrix}, \quad (17)$$

where $g = w_c T A / 2$ and $J_i = 1 - \tanh^2(S_i)$ for $i \in [0, 4]$.

C. Antialiasing Method

The discrete-time form proposed in the previous section will suffer from aliasing distortion when driven by high input levels due to the harmonics introduced by the \tanh functions [15]. This condition could be ameliorated using oversampling. However, this approach is far from ideal as it will introduce redundant operations which will increase the computational costs of the system. In this work we proposed incorporating the first-order antiderivative antialiasing method proposed by Parker et al. [26].

The antiderivative antialiasing method can be applied to arbitrary memoryless mappings of the form

$$y^n = f(x^n), \quad (18)$$

where $f()$ is a nonlinear function. The antialiased form is then given by

$$\tilde{y}^n = \frac{F_0(x^n) - F_0(x^{n-1})}{x^n - x^{n-1}}, \quad (19)$$

where $F_0()$ is the antiderivative of $f()$. As explained in [26], this antialiased form is derived from the analytical convolution of a continuous-time representation of the input signal with a simple two-point lowpass kernel.

When $x^n \approx x^{n-1}$, (19) becomes ill-conditioned. When this happens the antialiased form must be replaced by

$$\tilde{y}^n = f\left(\frac{x^n + x^{n-1}}{2}\right), \quad (20)$$

which is used to bypass the method while accounting for the half-sample delay introduced [26].

We once again discretize (1)–(4) using (19) on the right-hand side and rearrange in the form $\mathbf{F}_{\text{AA}}(\mathbf{V}) = \mathbf{0}$, where \mathbf{F}_{AA} is the antialiased system to solve:

$$\mathbf{F}_{\text{AA}} = \begin{cases} V_1^n - V_1^{n-1} + w_c T \left[\frac{I_1^n - I_1^{n-1}}{d_1} + \frac{I_0^n - I_0^{n-1}}{d_0} \right] = 0 \\ V_2^n - V_2^{n-1} - w_c T \left[\frac{I_1^n - I_1^{n-1}}{d_1} - \frac{I_2^n - I_2^{n-1}}{d_2} \right] = 0 \\ V_3^n - V_3^{n-1} - w_c T \left[\frac{I_2^n - I_2^{n-1}}{d_2} - \frac{I_3^n - I_3^{n-1}}{d_3} \right] = 0 \\ V_4^n - V_4^{n-1} - w_c T \left[\frac{I_3^n - I_3^{n-1}}{d_3} - \frac{I_4^n - I_4^{n-1}}{d_4} \right] = 0, \end{cases} \quad (21)$$

where

$$d_0 = A (V_{\text{in}}^n + kV_4^n - V_{\text{in}}^{n-1} - kV_4^{n-1}) \quad (22)$$

$$I_0^n = \ln [\cosh (AV_{\text{in}}^n + kAV_4^n)] \quad (23)$$

and

$$I_i^n = \ln [\cosh (AV_i^n)] \quad (24)$$

$$d_i = A (V_i^n - V_i^{n-1}) \quad (25)$$

for $i \in [1, 4]$. The new Jacobian matrix required to implement Newton's method is now given by

$$\mathbf{D}\mathbf{F}_{\text{AA}} = \begin{bmatrix} 1 + Z_1 & 0 & 0 & kZ_0 \\ -Z_1 & 1 + Z_2 & 0 & 0 \\ 0 & -Z_2 & 1 + Z_3 & 0 \\ 0 & 0 & -Z_3 & 1 + Z_4 \end{bmatrix}, \quad (26)$$

where

$$Z_0 = P_0 [d_0 \tanh (AV_{\text{in}}^n + kAV_4^n) - (I_0^n - I_0^{n-1})] \quad (27)$$

$$P_0 = kw_c T A / d_0^2 \quad (28)$$

and

$$Z_i = P_i [d_i \tanh (AV_i^n) - (I_i^n - I_i^{n-1})] \quad (29)$$

$$P_i = w_c T A / d_i^2 \quad (30)$$

for $i \in [1, 4]$.

When the denominator of the antialiased form becomes ill-conditioned and is replaced with the form (20), the method becomes equivalent to the implicit midpoint method described by (15) and (17).

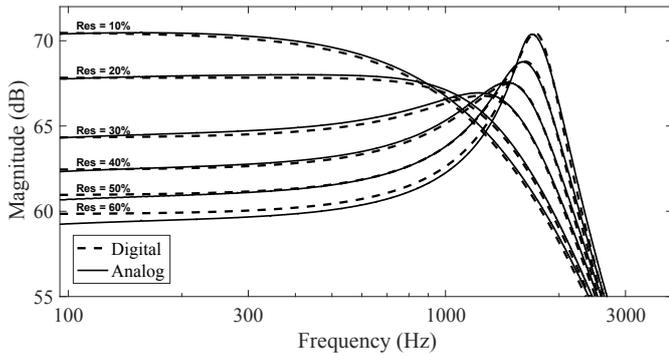


Fig. 2: Linear response of the analog and digital filters for $f_c = 2$ kHz and resonance in knob positions from 10% to 60%.

III. MEASUREMENTS AND COMPARISON

For the validation of the digital model described in Sec. II, measurements on an analog Moog synthesizer (Mother-32) were made. The measuring process was divided into two parts. In the first part, the linear response of the ladder filter, for various cutoff frequencies and knob positions for the resonance, was obtained by using a sweep-sine as input. In the second part, the nonlinear behavior of the filter was examined and high-voltage waveforms were selected as inputs. A sample rate of 44.1 kHz was used in all measurements and simulations.

In the following subsections the measurement process for both parts is explained. Comparisons between the measurements and the digital model are discussed, and the results of applying the antialiasing technique to the nonlinear model are presented.

A. Linear Response

In order to measure the linear response of the analog filter an exponential sine sweep (ESS) was created following the method described in [30]. A 6-second ESS ranging from 10 Hz – 22 kHz was synthesized for the measurements. This signal was used as input to the analog filter and the output of the VCF was recorded. In order to retrieve the impulse response of the system from the recorded signal, the method relies on deconvolution. As explained in [30] the spectrum of the ESS signal is not flat, and the same problem can be observed in the inverse signal that is used for the deconvolution and is the time reversed version of ESS signal. To compensate for this, the inverse filter undergoes an amplitude modulation. This process is detailed in [31].

The resulting signal, after the deconvolution with the proper inverse signal, consists of the linear response of the filter and the harmonic distortion responses. Given that our aim is to study the linear response only, a time window was applied so as to extract only the linear part of the deconvolved signal.

Next, the linear response of the digital model presented in Sec. II, was measured. As explained in [15], when the level of the input signal $V_{in} \ll 26$ mV the nonlinear effects of the tanh functions become negligible and the linear response of the model can be measured. To achieve this, a low-level impulse was chosen as the input V_{in} to the trivial model.

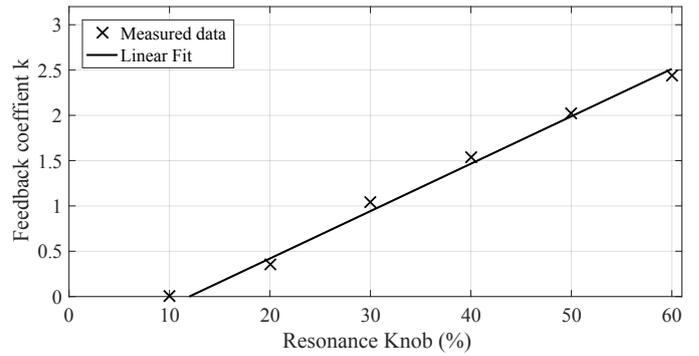


Fig. 3: Mapping between resonance's control knob position in the analog filter and feedback coefficient k in the digital model.

When comparing the measurements with regards to the resonance, in contrast with the frequency-to-frequency mapping between the cutoff frequency control knob and the ω_c value of the digital model, the position of the resonance control knob has no direct mapping to a digital value. For this reason, throughout this paper we use percentage values to indicate the position of the resonance knob in our measurements. The far left position corresponds to 0% resonance and the far right to 100%.

Figure 2 compares the linear magnitude of the frequency response of the measured analog device and its trivial digital implementation for different resonance values. A fixed cutoff parameter $f_c = 2$ kHz was used on both filters. The resonance parameter k in the digital model was chosen empirically based on the analog measurements. From these curves it can be observed that the differences between the response of the measured device and the simulated response are minimal. It should be noted that a DC offset was added to the digital measurements in order to match the low-frequency gain levels of the analog measurements. The 0% measurement is exactly the same as the 10% one. At 70% resonance the analog filter starts behaving nonlinearly, right before it begins to self-oscillate at approx. 80% and the measurement becomes illegible due to noise.

Based on the results in Fig. 2, we can detect a mapping between the resonance knob positions on the analog filter and parameter k in the digital. Fig. 3 shows this mapping. The measured data represents the knob positions of the resonance in our measurements and the respective k value in the digital model. It can be seen that the relation between the analog and digital resonance is linear. It should be noted that positions 0% and 10% both correspond to $k = 0$.

During the measurement it was observed that the labels on the synthesizer's cutoff frequency knob do not map directly to the measured cutoff frequency of the filter. For that reason, some additional measurements were made in order to identify the correct mapping between the knob positions and the actual cutoff frequency of the filter. Fig. 4 illustrates this mapping. For the cutoff frequency knob in positions 0–100% and the resonance knob in the 0% position, the -12 -dB drop in the magnitude response of the filter was detected. This point marks

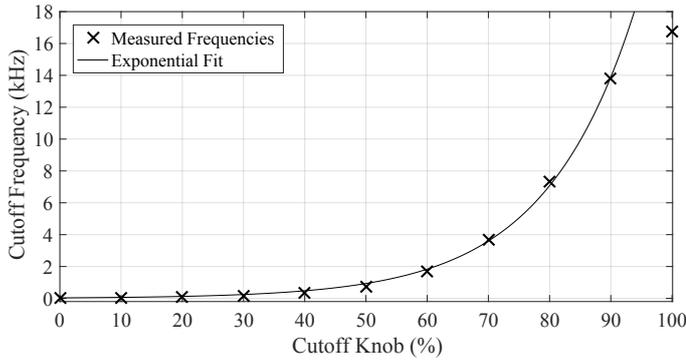


Fig. 4: Mapping between the cutoff frequency control knob position and the measured cutoff frequency of the analog filter. The data point at 100% is considered an outlier and has not been included in the exponential fit.

the cutoff frequency of four cascaded lowpass filters in the absence of resonance. As shown in Fig. 4, the relationship between the physical knob positions and the actual cutoff frequency fits an exponential curve, which indicates that the scale for the control knob is logarithmic.

B. Nonlinear Response

As explained in Sec. II A, in order to excite the nonlinear behavior of the analog filter, a high-level input signal is needed. Two different waveforms, an analog sinewave and a sawtooth, were selected as inputs for our measurements,

with their amplitude set at 12 V peak-to-peak and fundamental frequency 500 Hz. The same waveforms were recorded and used as input to the trivial digital model. The cutoff frequency and resonance parameters were set on both the analog system and digital model to approx. 3.7 kHz (i.e., knob positions at 70%) and 0, respectively.

After performing a series of measurements with these high-level waveforms, it became evident that the analog filter was preceded by a passive attenuation stage, designed to bring the input signal down to operating level. For this reason, we tried to incorporate this stage to the digital model by attenuating the amplitude of the recorded input signal. The attenuation value was chosen empirically to match the analog measurements.

Figures 5 and 6 compare the magnitude spectra of the analog and digital model, when driven by the same input waveform. The analog measurements with the sinewave and the sawtooth input can be seen in Fig. 5a and 6a, respectively. From Fig. 5a we can observe how the system introduces odd harmonic distortion components. The low-level even harmonics in the magnitude spectrum can be attributed to noise artifacts in the input signal and to the imperfect nature of the transistors, whose saturating behavior is slightly asymmetric. However, due to their low level these harmonics can be neglected. Figures 5b and 6b show the result of driving the trivial digital model using the same input waveforms and filter parameters. It is observed that the digital model approximates the analog filter's nonlinear behavior to a great extent.

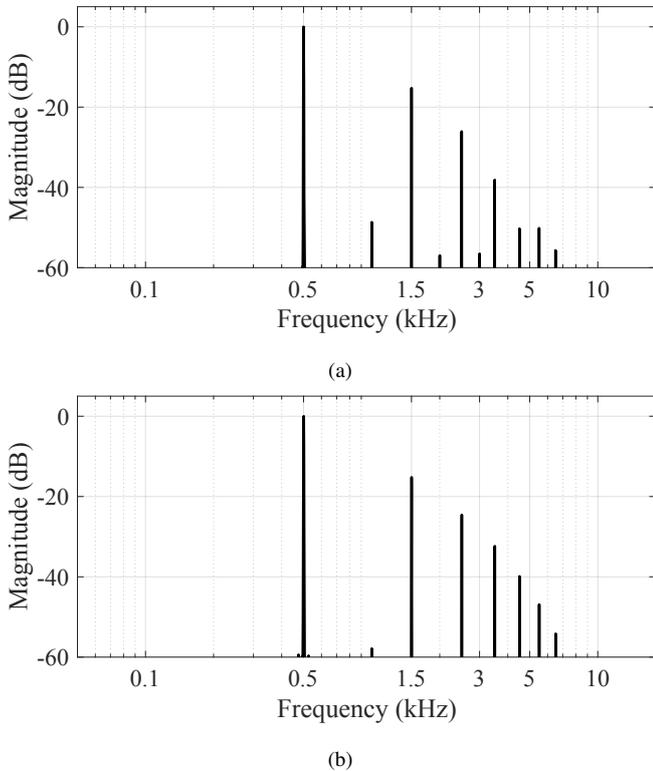


Fig. 5: Comparison between (a) the analog and (b) the digital filter measurements with an analog sinewave input signal.

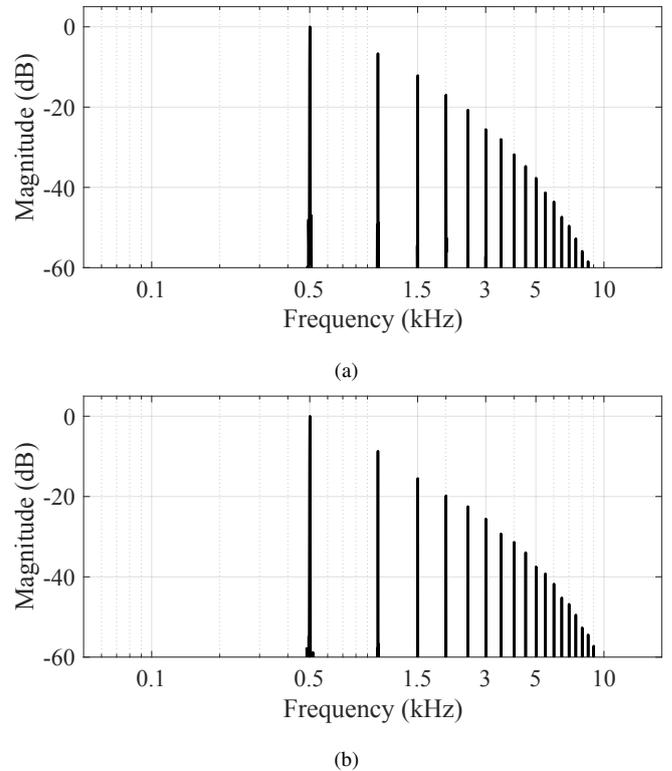


Fig. 6: Comparison between (a) the analog and (b) the digital filter measurements with an analog sawtooth input signal.

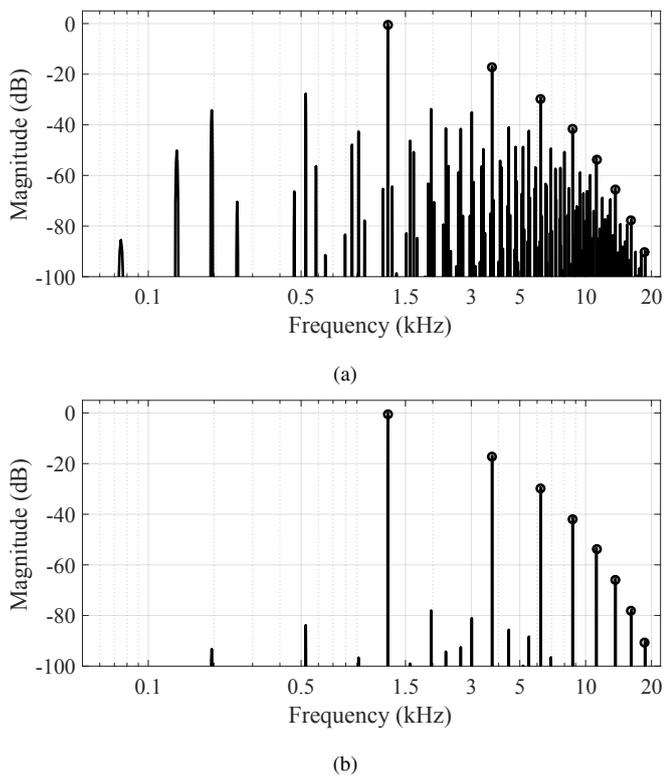


Fig. 7: Effects of aliasing distortion on (a) the trivial and (b) the antialiased digital filter model. The circled spectral peaks are the desired harmonics whereas the remaining spectral components are undesirable aliasing.

C. Antialiasing

The trivial digital model suffers from aliasing distortion when driven by high input frequencies and large input levels. The effects of aliasing can be observed in the magnitude spectrum shown in Fig. 7a, where the input of the trivial digital model is a 1245-Hz sinewave with a peak amplitude of 5 V. The circles on some of the peaks mark the fundamental frequency of the input and the harmonics introduced by the filter’s distortion. All the other peaks represent the aliasing components that seem to be in relevant magnitude with the harmonics and therefore, able to spoil the timbre of the output.

To reduce these artifacts we apply the antialiasing method described in Sec. II. The magnitude spectrum at the output of the antialiased system, when driven by the same input signal can be seen in Fig. 7b. The application of the antialiasing technique reduces the aliasing components significantly with up to 50-dB suppression. This aliasing reduction is, also, clearly audible.

IV. CONCLUSIONS

This paper has reported on how the analog resonant filter proposed by Robert Moog can be modeled as a digital system and how the properties of the digital model relate to the analog filter. Both the linear (small-signal) and the nonlinear (large-signal) behavior were investigated. The small-level characteristics were tested by feeding a sinusoidal sweep signal to the analog filter and by comparing the response to that of the

digital model. It was shown that the digital and analog filter have a very similar frequency response with respect to the two main parameters, the cutoff and resonance. Differences of only 1-dB are below the threshold of audibility.

In the nonlinear regime, it becomes important to account for the harmonic distortion caused by transistors, since this contributes to the unique sound of Moog synthesizers. The nonlinear behavior was tested using high-level sinusoidal input signals. The measurement results show similar nonlinear behavior, such as distortion components in the output spectrum. An antialiasing method based on the use of the antiderivative of the nonlinear shaping function was incorporated in the digital model, reducing the aliasing to a negligible level.

ACKNOWLEDGMENT

This work was conducted in spring 2017, when Ms. E. Paschou spent three months as an ERASMUS+ exchange student at the Aalto Acoustics Lab, Espoo, Finland.

REFERENCES

- [1] R. A. Moog. “A voltage-controlled low-pass high-pass filter for audio signal processing”. In *Proc. 7th Conv. Audio Eng. Soc.*, New York, USA, Oct. 1965.
- [2] J. O. Smith. “Physical modeling synthesis update”. *Computer Music J.*, 20(2):44–56, Summer 1996.
- [3] V. Välimäki, S. Bilbao, J.O. Smith, J.S. Abel, J. Pakarinen, and D. Berners. “Virtual analog effects”. *DAFX: Digital Audio Effects, Second Edition*, U. Zölzer (ed.), pages 473–522, 2011.
- [4] D. Rossum. “Making digital filters sound “analog””. In *Proc. Intl. Computer Music Conf. (ICMC 1992)*, pages 30–34, San Jose, CA, USA, Oct. 1992.
- [5] T. Stilson and J. O. Smith. “Analyzing the Moog VCF with considerations for digital implementation”. In *Proc. Intl. Computer Music Conf.*, Hong Kong, Aug. 1996.
- [6] V. Välimäki and A. Huovilainen. “Oscillator and filter algorithms for virtual analog synthesis”. *Computer Music J.*, 30(2):19–31, 2006.
- [7] A. Huovilainen. “Enhanced digital models for analog modulation effects”. In *Proc. Intl. Conf. Digital Audio Effects (DAFx-05)*, pages 155–160, Madrid, Spain, Sept. 2005.
- [8] M. Holters and U. Zölzer. “Physical modelling of a wah-wah effect pedal as a case study for application of the nodal DK method to circuits with variable parts”. In *Proc. Intl. Conf. Digital Audio Effects (DAFx-11)*, pages 287–290, Paris, France, Sept. 2011.
- [9] F. Eichas, M. Fink, M. Holters, and U. Zölzer. “Physical modeling of the MXR Phase 90 guitar effect pedal”. In *Proc. Intl. Conf. Digital Audio Effects (DAFx-14)*, pages 153–158, Erlangen, Germany, 2014.
- [10] D. T. Yeh, J. Abel, and J. O. Smith III. “Simulation of the diode limiter in guitar distortion circuits by numerical solution of ordinary differential equations”. In *Proc. Intl. Conf. Digital Audio Effects (DAFx-07)*, pages 197–204, Bordeaux, France, Sept. 2007.
- [11] R. C. D. de Paiva, S. D’Angelo, J. Pakarinen, and V. Välimäki. “Emulation of operational amplifiers and diodes in audio distortion circuits”. *IEEE Trans. Circuits Syst. II, Exp. Briefs*, 59(10):688–692, Oct. 2012.
- [12] K. J. Werner, V. Nangia, A. Bernardini, J. O. Smith III, and A. Sarti. “An improved and generalized diode clipper model for wave digital filters”. In *Proc. 139th Conv. Audio Eng. Soc.*, New York, USA, Oct.–Nov. 2015.
- [13] R. C. D. de Paiva, J. Pakarinen, V. Välimäki, and M. Tikander. “Real-time audio transformer emulation for virtual tube amplifiers”. *EURASIP J. Adv. Signal Process*, 2011, Feb. 2011.
- [14] J. D. Parker. “A simple digital model of the diode-based ring modulator”. In *Proc. Intl. Conf. Digital Audio Effects (DAFx-11)*, pages 163–166, Paris, France, Sept. 2011.
- [15] A. Huovilainen. “Non-linear digital implementation of the Moog ladder filter”. In *Proc. Intl. Conf. Digital Audio Effects (DAFx-04)*, pages 61–164, Naples, Italy, Oct. 2004.

- [16] S. D'Angelo and V. Välimäki. "Generalized Moog ladder filter: Part I—Linear analysis and parameterization". *IEEE Trans. Audio, Speech, and Lang. Process.*, 22(12):1825–1832, Dec. 2014.
- [17] T. Hélie. "Volterra series and state transformation for real-time simulations of audio circuits including saturations: Application to the Moog ladder filter". *IEEE Trans. Audio Speech Lang. Process.*, 18(4):747–759, May 2010.
- [18] P. Daly. "A comparison of virtual analogue Moog VCF models". Master's thesis, University of Edinburgh, Edinburgh, UK, 2012.
- [19] F. Esqueda. "Real-Time iOS implementation of a virtual analogue synthesiser model". Master's thesis, University of Edinburgh, Edinburgh, UK, 2013.
- [20] S. D'Angelo and V. Välimäki. "An improved virtual analog model of the Moog ladder filter". In *Proc. IEEE Intl. Conf. Acoustics, Speech and Signal Processing (ICASSP-13)*, pages 729–733, Vancouver, Canada, May 2013.
- [21] S. D'Angelo and V. Välimäki. "Generalized Moog ladder filter: Part II—Explicit nonlinear model through a novel delay-free loop implementation method". *IEEE Trans. Audio, Speech, and Lang. Process.*, 22(12):1873–1883, Dec. 2014.
- [22] F. Fontana and M. Civolani. "Modeling of the EMS VCS3 voltage-controlled filter as a nonlinear filter network". *IEEE Trans. Audio, Speech, Language Process.*, 18(4):760–772, Apr. 2010.
- [23] J.D. Parker and S. D'Angelo. "A digital model of the Buchla lowpass-gate". In *Proc. Intl. Conf. Digital Audio Effects (DAFx-13)*, pages 278–285, Maynooth, Ireland, Sept. 2013.
- [24] F. Esqueda, S. Bilbao, and V. Välimäki. "Aliasing reduction in clipped signals". *IEEE Trans. Signal Process.*, 60(20):5255–5267, Oct. 2016.
- [25] F. Esqueda, V. Välimäki, and S. Bilbao. "Rounding corners with BLAMP". In *Proc. Intl. Conf. Digital Audio Effects (DAFx-16)*, pages 121–128, Brno, Czech Republic, Sept. 2016.
- [26] J. D. Parker, V. Zavalishin, and E. Le Bivic. "Reducing the aliasing of nonlinear waveshaping using continuous-time convolution". In *Proc. 19th Intl. Conf. Digital Audio Effects (DAFx-16)*, pages 137–144, Brno, Czech Republic, Sept. 2016.
- [27] S. Bilbao, F. Esqueda, J. D. Parker, and V. Välimäki. "Antiderivative antialiasing for memoryless nonlinearities". *IEEE Signal Process. Lett.*, 24(7):1049–1053, July 2017.
- [28] D. Griffiths and D. Higham. *Numerical Methods for Ordinary Differential Equations*. Springer-Verlag, London, UK, 2010.
- [29] S. Zambon and F. Fontana. "Efficient polynomial implementation of the EMS VCS3 filter model". In *Proc. Intl. Conf. Digital Audio Effects (DAFx-11)*, pages 31–35, Paris, France, Sept. 2011.
- [30] A. Farina. "Advancements in impulse response measurements by sine sweeps". In *Proc. 122nd Audio Eng. Soc. Conv.*, Vienna, Austria, May 2007.
- [31] Q. Meng, D. Sen, S. Wang, and L. Hayes. "Impulse response measurement with sine sweeps and amplitude modulation schemes". In *Proc. 2nd Intl. Conf. Signal Processing and Communication Systems*, pages 1–5, Gold Coast, Australia, Dec. 2008.