

NON-LINEAR DIGITAL IMPLEMENTATION OF THE MOOG LADDER FILTER

Antti Huovilainen

Laboratory of Acoustics and Audio Signal Processing
Helsinki University of Technology, P.O. Box 3000, FIN-02015 HUT, Espoo,
Finland
ajhuovil@acoustics.hut.fi

ABSTRACT

This paper presents a non-linear digital implementation of the Moog ladder filter. The implementation is relatively efficient and suitable for inclusion into real-time systems, for example virtual analog synthesizers. The analog circuit is analyzed to produce a differential equation. This equation is solved using Euler's method, and the result is shown to be equivalent to a cascade of first order IIR sections with embedded non-linearities. Finally, the filter structure is modified to improve tuning.

1. INTRODUCTION

Time varying filters are used in many musical applications - synthesizers, effects units and samplers. They are especially important nowadays in virtual analog synthesizers. The voltage-controlled filter published by Robert Moog in 1965 [1] is perhaps the most famous of them.

There exist several published digital filters suitable for musical applications, such as the State-Variable Filter [2]. The Moog filter itself has also been converted to digital form by Stilson and Smith [3]. These filters are linear and some people feel that they sound "digital" and lack the "warmth" that is characteristic of analog filters - especially the Moog filter. Rossum [4] has published a filter that claims to have the same "warm" sound by embedding a non-linearity within the filter.

In this paper the Moog filter circuit is analyzed and a digital implementation of it is presented. The implementation models the inherent non-linearities of the original circuit and should thus give more accurate results compared to linear digital filters. As the digital filter is directly based on the analog circuit, there are no extra coefficients that would need to be tuned by ear - the "warmth" is determined by the input amplitude. While the resulting filter requires more computation than traditional linear filters, modern advances in computing power allow it to be used in real-time systems. Since the large signal behavior of the circuit is analyzed, traditional pole-zero analysis is not directly usable. Not using pole-zero domain can also offer otherwise hidden insights.

Additional material and audio examples are available at <http://www.acoustics.hut.fi/publications/papers/dafx2004-moog/>

2. THE MOOG LADDER

Moog implemented his filter using an innovative transistor ladder circuit shown in Figure 1. It employs the base-to-emitter

resistance of bipolar transistors to realize voltage-controlled RC-sections. The transistors also act as buffers between each stage.

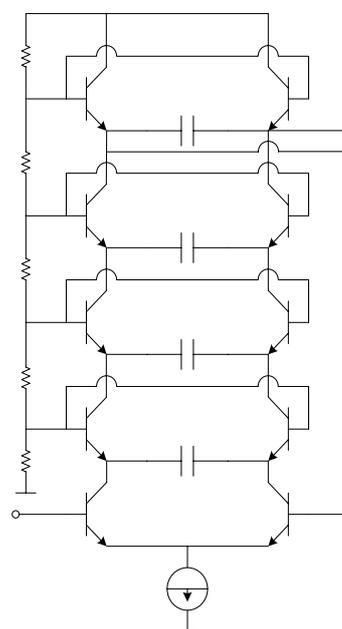


Figure 1: *The Moog ladder circuit* [1].

The circuit has four stages, each of which consists of two transistors and a capacitor. Each stage is driven by the output current of the previous stage apart from the first stage, which is driven by a differential transistor pair. The bottom of the ladder is driven by control current I_c . Differential output is taken from the emitters of the transistors in the last stage to a high impedance differential amplifier. To produce resonance, a portion of the output is routed to the other side of the input differential amplifier.

The current gain, or β , of the transistors is high (values of > 200 are typical) and the resistors feeding the transistor bases are small valued, so transistor base voltages are effectively constant for all stages. It is a reasonable approximation to assume that the transistor β is infinite and therefore the base current is zero. Thus the stages are buffered from each other. Further, since the base-emitter voltage is logarithmically dependant on the collector current, the emitter voltage varies very little and the Early effect

[6] can be neglected. Therefore each stage depends only on the current state and the input current coming from the previous stage.

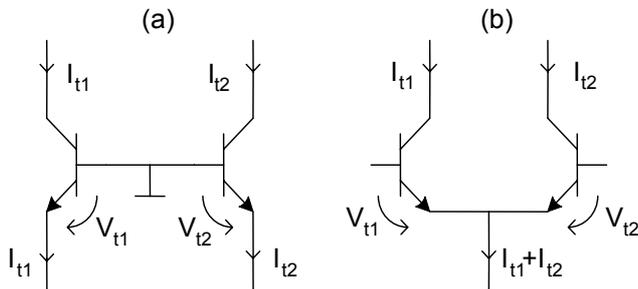


Figure 2: (a) Differential pair used in Moog ladder, (b) Traditional differential pair.

2.1. Differential pair

Figure 2a shows the differential pair used in the ladder. V_{t1} and V_{t2} are the base-emitter voltages of the transistors while I_{t1} and I_{t2} are the collector currents. Since we assume that transistor beta is infinite and no base current flows, the emitter current is the same as the collector current.

Since the signal current is differential, we are interested in finding out the relationship between V_{t1} , V_{t2} and $I_{t1}-I_{t2}$. The equation for $I_{t1}-I_{t2}$ in the circuit of Figure 2a is same as for the circuit in Figure 2b.

Equation for the differential current in a transistor pair [5] is

$$I_{t1} - I_{t2} = I_{diff} = (I_{t1} + I_{t2}) \tanh\left(\frac{V_{t1} - V_{t2}}{2V_t}\right), \quad (1)$$

where V_t is the so-called thermal voltage of a transistor [6]. Equation (1) holds when the transistors are assumed to be perfectly matched, beta is infinite and the Early effect is neglected.

3. DIFFERENTIAL EQUATION

We can now derive a differential equation for the filter. Since the stages are buffered from each other, it makes sense to first derive a differential equation for a single stage.

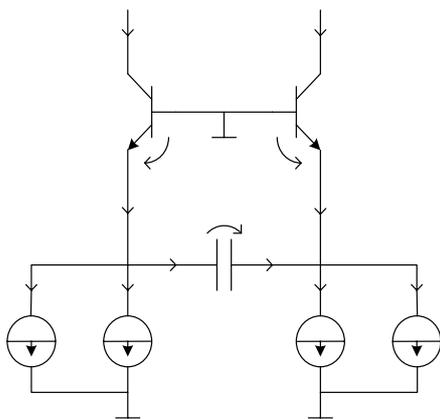


Figure 3: Single stage of the Moog ladder.

3.1. Single stage

Figure 3 shows the equivalent circuit for a single stage of the Moog ladder. V_{t1} , V_{t2} , I_{t1} and I_{t2} are the same as in Figure 2. I_{ctl} is half of the ladder control current and I_s is half of the signal current. It can be seen that

$$I_{t1} = I_{ctl} + I_s - I_c, \quad (2)$$

$$I_{t2} = I_{ctl} - I_s + I_c, \quad (3)$$

$$I_{t1} + I_{t2} = 2I_{ctl}, \quad (4)$$

$$I_{t1} - I_{t2} = 2I_s - 2I_c, \quad (5)$$

where I_{t1} and I_{t2} are the base-emitter voltages of the transistors, I_{ctl} is half of the control current, I_s is half of the signal current and I_c is the current through capacitor C .

From equation (4) it follows that

$$I_{t1} - I_{t2} = 2I_{ctl} \tanh\left(\frac{V_c}{2V_t}\right), \quad (6)$$

where V_c is the voltage over the capacitor C . Equation (1) can now be written as

$$2I_c = 2I_s - 2I_{ctl} \tanh\left(\frac{V_c}{2V_t}\right), \quad (7)$$

which will prove useful in the next Section.

3.2. Differential equation for a single stage

We can now write the differential equation for a single Moog ladder stage. The equation for current I_c through a capacitor is

$$I_c = C \frac{dV_c}{dt}, \quad (8)$$

where V is voltage over the capacitor and C is its capacitance. Inserting equation (8) into equation (7) then gives

$$2C \frac{dV_c}{dt} = 2I_s - 2I_{ctl} \tanh\left(\frac{V_c}{2V_t}\right). \quad (9)$$

As each stage is driven either by the previous stage or by the differential input amplifier, equation (9) can be written as

$$\frac{dV_c}{dt} = \frac{I_{ctl}}{C} \left(\tanh\left(\frac{V_{in}}{2V_t}\right) - \tanh\left(\frac{V_c}{2V_t}\right) \right). \quad (10)$$

4. DISCRETIZING THE FILTER

4.1. Difference equation for a single stage

To be useful for a digital implementation, the differential equation must now be solved. The easiest way is to use Euler's method. Although Euler's method has some drawbacks, it is very useful for this particular case. As the differential equation is of first order, the solution is inherently stable. Since there is a non-linearity, oversampling must be used and this brings the Euler solution closer to the ideal solution. Runge-Kutta or some other higher order method could also be used, but it would require evaluation of the equation between samples. This is problematic as it is equivalent to having a higher sample rate for the input than the output. It also poses a problem for resonance, since that is achieved by feeding back some of the ladder output. Euler's

method also has the advantage that the resulting difference equation is similar to a normal one-pole IIR lowpass filter as seen in Section 5.

Euler solution for equation (10) is

$$V_c(n) = V_c(n-1) + T_s \frac{I_{ctl}}{C} \left(\tanh\left(\frac{V_{in}(n)}{2V_t}\right) - \tanh\left(\frac{V_c(n-1)}{2V_t}\right) \right), \quad (11)$$

where T_s is the time interval between samples. For sample rate F_s ,

$$T_s = \frac{1}{F_s}. \quad (12)$$

4.2. Difference equation for the complete filter

Difference equations can now be written for the full ladder filter.

$$y_a(n) = y_a(n-1) + \frac{I_{ctl}}{CF_s} \left(\tanh\left(\frac{x(n) - 4ry_d(n-1)}{2V_t}\right) - W_a(n-1) \right) \quad (13)$$

$$y_b(n) = y_b(n-1) + \frac{I_{ctl}}{CF_s} (W_a(n) - W_b(n-1)) \quad (14)$$

$$y_c(n) = y_c(n-1) + \frac{I_{ctl}}{CF_s} (W_b(n) - W_c(n-1)) \quad (15)$$

$$y_d(n) = y_d(n-1) + \frac{I_{ctl}}{CF_s} \left(W_c(n) - \tanh\left(\frac{y_d(n-1)}{2V_t}\right) \right) \quad (16)$$

where $x(n)$ is the input, $y_a(n)$, $y_b(n)$, $y_c(n)$ and $y_d(n)$ are the outputs of individual filter stages, r is the resonance amount ($0 < r \leq 1$) and

$$W_{\{a,b,c\}}(n) = \tanh\left(\frac{y_{\{a,b,c\}}(n)}{2V_t}\right) \quad (17)$$

It can be seen that each stage uses as input the tanh of the output of the previous stage. This is also used by the previous stage during the next sample. The calculation result can be stored and thus only five tanh calculations per sample are required. These can be implemented efficiently with table lookups or polynomial approximations.

5. IMPROVING TUNING

5.1. Tuning of a single stage

While this paper is concerned with the large signal model of the Moog ladder filter, it is interesting to see what equation (11) is for low signal amplitudes. For small inputs ($-0.5 < x < 0.5$), \tanh function is almost linear. Equation (11) then becomes

$$V_k(n) = V_k(n-1) + \frac{I_{ctl}}{CF_s} \left(\frac{V_{in}(n)}{2V_t} - \frac{V_k(n-1)}{2V_t} \right), \quad (18)$$

where F_s is the sample rate. Equation (18) is similar to that of a normal digital one-pole lowpass filter.

Scaled impulse invariant transform [7] is impulse invariant transform scaled so that the dc gain is one. The difference equation for a one-pole lowpass filter transformed with scaled impulse invariant transform is

$$y(n) = y(n-1) + g(x - y(n-1)), \quad (19)$$

which is the same as equation (18) with

$$g = \frac{I_{ctl}}{2V_t CF_s} \quad (20)$$

substitution. As can be seen, I_{ctl} , C and F_s determine the tuning. Since I_{ctl} and C do not affect anything else, their exact values are irrelevant. Coefficient g can therefore be computed the same way as with a normal scaled impulse invariant transformed one-pole filter

$$g = 1 - e^{-2\pi F_c / F_s}, \quad (21)$$

where F_c is the cutoff frequency.

Making this substitution to equation (13) and substituting x for input and y for output gives

$$y(n) = y(n-1) + 2V_t g \left(\tanh\left(\frac{x(n)}{2V_t}\right) - \tanh\left(\frac{y(n-1)}{2V_t}\right) \right) \quad (22)$$

Making the substitution to equations (14)–(17) gives the difference equations for the complete filter.

5.2. Resonance

To produce resonance, the filter output is fed back inverted. In the analog filter, each stage causes a 45 degree phase shift at the cutoff frequency, producing a combined phase shift of 180 degrees at the cutoff frequency. This phase shift, combined with inverting, causes the feedback to be positive at the cutoff frequency and thus it emphasizes frequencies around the cutoff. The attenuation for a single stage is 3 dB at cutoff, producing a total attenuation of 12 dB for the complete filter at the cutoff point.

In the digital implementation, the unit delay in the feedback path causes an additional phase shift, making the combined phase shift to be

$$p = 4p_{stage}(f, F_c) + 180 \frac{f}{F_s}, \quad (23)$$

where p is the total phase shift and p_{stage} is the phase shift of a single filter stage. This additional phase shift causes the resonance frequency to vary from the cutoff frequency. Another effect is that the attenuation at resonance frequency is no longer exactly 3 dB. This means that the feedback amount required to produce the desired resonance varies with frequency.

5.3. Compensation

Stilson and Smith show some methods to compensate this shifting of resonance frequency and amplitude [3]. However, these methods require the use of tuning and resonance amount compensation tables. This has an unfortunate side effect owing to the difference between tuning for zero resonance and tuning for self oscillation. Here, the tuning table must be two-dimensional, or some form of interpolation is needed between no tuning and full tuning (depending on the resonance amount). Further, methods where transfer function zeroes are introduced require two different coefficients to be used and interpolated in the filter loop.

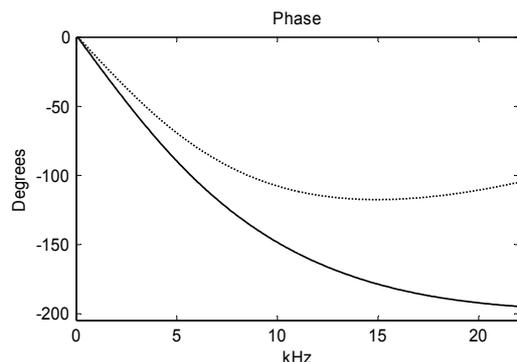


Figure 4: Phase shift of the filter with (solid) and without (dashed) feedback path delay.

Another approach is to stay with the scaled impulse invariant transformed filter stage and see how the filter structure might possibly be modified to compensate in order to get feedback phase as close to 180 degrees as possible. Figure 4 shows the phase shift for the filter both alone and with the feedback path unit delay using two times oversampling (88.2 kHz sample rate). It can be seen that the phase shift of the filter starts falling back to zero at higher frequencies and this somewhat compensates for the phase shift introduced by the unit delay.

The resulting phase shift is now slightly too small at the cutoff frequency. Addition of half-unit delay causes the phase shift to be almost exactly 180 degrees at cutoff up to about $F_s/4$. The half-unit delay can be realized by averaging two samples. At very high frequencies ($f > F_s/4$) the situation is now worse than without the extra delay, but as some oversampling is required because of the nonlinearities, this does not matter in practice. Figure 5 shows the tuning and amplitude error with and without the extra delay. With the extra delay added, the error in tuning is less than 10% for $f < F_s/4$ (or $f < 22$ kHz for F_s of 88.2 kHz).

With two times oversampling, the remaining tuning error can be eliminated by using a tuning table. Since the error is so small, the resonance tuning compensation can be combined with the tuning for scaled impulse invariant transformed one-pole filter into a single table. The small error in the frequency response is unlikely to be audible.

6. CONCLUSIONS

A digital implementation of the Moog ladder filter has been presented with non-linearities of the circuit correctly modelled. The implementation is similar to a normal IIR filter made of cascaded first-order sections, but the first order sections have non-linearities embedded within them. The filter is directly based on the Moog transistor ladder circuit and thus requires no user tunable parameters other than the cutoff frequency and input amplitude.

While more computationally intensive than traditional IIR filters, the filter is still suitable for real-time and DSP implementations. As some of the calculations are shared between stages, the implementation requires only five tanh-function evaluations. These can be implemented efficiently with table lookups or polynomial approximations.

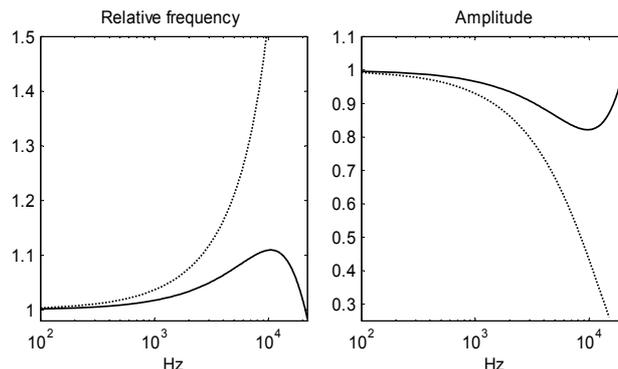


Figure 5: Tuning and amplitude error for compensated (solid) and uncompensated (dashed) filter.

Some oversampling is required to avoid aliasing. Slight modification of the filter structure with use of oversampling also makes the resonance frequency almost the same as the cutoff frequency, thus requiring only modest tuning and resonance gain compensation.

Additional material and audio examples are available at <http://www.acoustics.hut.fi/publications/papers/dafx2004-moog/>

7. ACKNOWLEDGEMENTS

The author would like to thank Professor Vesa Välimäki for his valuable support and feedback, as well as Merlijn Blaauw, Martin Fay, John McMonagle, Andreas Tell and George Yohng for comments and proofreading.

This work has been supported by the Academy of Finland (project #104934).

8. REFERENCES

- [1] R. A. Moog, "A Voltage-Controlled Low-Pass High-Pass Filter for Audio Signal Processing," *17th AES Convention*, New York, Oct. 1965.
- [2] H. Chamberlin, *Musical Applications of Microprocessors*, Hayden Book Company, 1980, pp. 489–492.
- [3] T. Stilson and J. O. Smith, "Analyzing the Moog VCF with Considerations for Digital Implementation", in *Proc. Int. Computer Music Conf.*, Hong Kong, China, 1996, pp. 398–401.
- [4] D. Rossum, "Making Digital Filters Sound Analog," *Proc. Int'l Computer Music Conf.*, San Jose, California, Oct. 1992, pp. 30–34.
- [5] A. Gratz, *Operational Transconductance Amplifiers*, <http://www.compuserve.de/stromeko/diy/OTA.pdf>
- [6] A. S. Sedra and K. C. Smith, *Microelectronic Circuits*, Fourth International Edition, Oxford University Press, 1998.
- [7] T. Thong and J. McNames, "Transforms for Continuous Time System Modeling," *45th Midwest Symp. on Circuits and Systems*, Tulsa, Oklahoma, Aug. 2002, pp. 408–411.