An Efficient Three-State Model of the EMS VCS3 Voltage-Controlled Filter

GriffinBoy In-House Paper, July 2026

Abstract

In this paper we present a fast virtual-analog model of the voltage-controlled filter used in the EMS VCS3 synthesizer. The VCS3 filter is a diode ladder with nonlinear feedback, and earlier circuit-derived models have shown that its characteristic behaviour can be simulated from the underlying circuit equations. The remaining problem is runtime cost: a faithful four-state implementation consumes too much realtime budget to serve as a standard filter in a software instrument.

The proposed filter is a source-derived three-state version of the same EMS model. In the main runtime benchmark, median CPU falls from 11.77% to 2.90% at 8x, and from 6.03% to 1.45% at 4x. These are speedups of 4.06x and 4.17x. On the calibrated Razer Blade 15 this corresponds approximately to 20% versus 6% at 8x and 10% versus 3% at 4x. Sweep and driven-saw measurements show that the main resonant behaviour is retained across the tested mid and high cutoff settings.

1. Introduction

The Voltage Controlled for Studio with 3 Oscillators, better known as the EMS VCS3, was designed by David Cockerell in 1969 and produced by Electronic Music Studios in London. Its voltage-controlled filter is a diode ladder: the cutoff is set by current-controlled junction diodes, and a feedback amplifier pushes the circuit into the unstable, resonant behaviour associated with the instrument.

In 2008, Civolani and Fontana derived a nonlinear differential-equation model of the EMS VCS3 filter from the circuit [1]. In 2010, Fontana and Civolani reformulated the same model as a nonlinear filter network with delay-free-loop solution [2]. In 2011, Zambon and Fontana reduced the cost of that model by reformulating the equations in polynomial form [3]. These papers established the EMS filter as a serious virtual-analog modelling problem: the diode ladder can be modelled from circuit equations, but the solver has to cope with nonlinear feedback, parameter motion, and high sample rates.

That last point is the runtime problem. A four-state nonlinear circuit model is a useful accuracy reference, but an instrument plugin may need the filter at 4x or 8x oversampling, with sample-rate modulation, several voices, and enough CPU left for oscillators, envelopes, effects, and the host. The aim of this work is therefore practical: keep the EMS circuit equations, remove unnecessary realtime work, measure the error, and report the CPU gain plainly.

The reference throughout this paper is the full four-state EMS circuit model from the Civolani-Fontana formulation, advanced with an implicit trapezoidal/Newton solve. The proposed model is a fast three-state EMS filter derived here. It removes one upstream capacitor state by local current balance, keeps the three downstream ladder states, and uses a fixed retune:

```
internal cutoff   = 0.50 * requested cutoff
internal feedback = 2.00 * requested feedback
```

Figure 1 shows the Ryzen 9 7900 CPU result; Section 7 gives the Razer Blade 15 estimate.

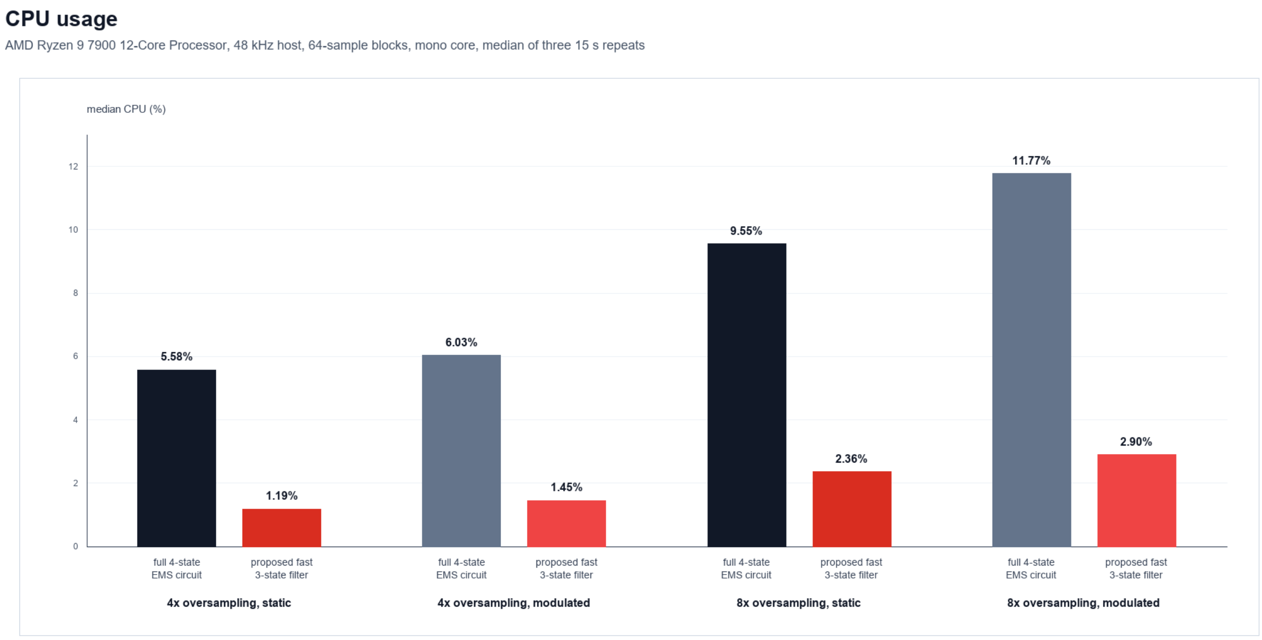

Figure 1. Median CPU usage for the full four-state EMS circuit model and the proposed fast three-state EMS filter. The benchmark used an AMD Ryzen 9 7900, 48 kHz host rate, 64-sample blocks, mono processing, and 4x or 8x oversampling.

2. Full Four-State EMS Model

The starting point is the Civolani-Fontana VCS3 diode-ladder model. Let the capacitor voltages be:

```
v = [v1, v2, v3, v4]
```

The nonlinear terms are:

```
x1 = tanh((v2 - v1) / (2 gamma))
x2 = tanh((v3 - v2) / (2 gamma))
x3 = tanh((v4 - v3) / (2 gamma))
x5 = tanh(v4 / (6 gamma))
u  = tanh((vin - (K + 1/2) v4) / (2 VT))
s  = I0 / (2 C)
```

where VT is thermal voltage, `gamma = eta VT`, `I0` is the cutoff-control current, `C` is the ladder capacitance, and `K` is the feedback gain. The state equations are:

```
dv1/dt = s * (u + x1)
dv2/dt = s * (x2 - x1)
dv3/dt = s * (x3 - x2)
dv4/dt = s * (-x5 - x3)
```

The output is proportional to the final ladder state:

```
y = (K + 1/2) v4
```

This model is the accuracy reference used below. Its cost comes from preserving the coupled nonlinear state update at audio rate and then repeating that work at each oversampled substep.

3. Route to the Reduction

The final model was found by source-derived reduction, not by fitting a black-box response. Five routes were tested.

First, we tried keeping all four states and reducing only the numerical work around them: transformed polynomial coordinates, fixed-point network evaluation, and stable full-state formulations. These routes kept the reference structure, but they did not remove enough work from the realtime path.

Second, we derived a two-state quasi-static reduction by eliminating both upstream capacitor states. It was fast, but it removed too much ladder memory and moved the resonance map too far.

Third, we derived an unretuned three-state reduction by eliminating only the first upstream state. This preserved more of the ladder, but the peak was not in the right place. At requested cutoff 4 kHz, feedback 6, and drive -42 dB, the full four-state EMS circuit model peaked near 3840 Hz, while the unretuned three-state reduction peaked near 7111 Hz and 13.3 dB.

Fourth, we tried feedback-only compensation. Doubling feedback recovered resonance height, but it did not recover cutoff position. The peak moved to about 8016 Hz at 21.1 dB.

Fifth, we applied cutoff and feedback compensation together. The chosen map was cutoff x0.50 and feedback x2.0. That brought the important resonance peak back into line without adding audio-rate work beyond a control scaling. Figure 2 shows the route from the full model through the failed reductions to the proposed filter.

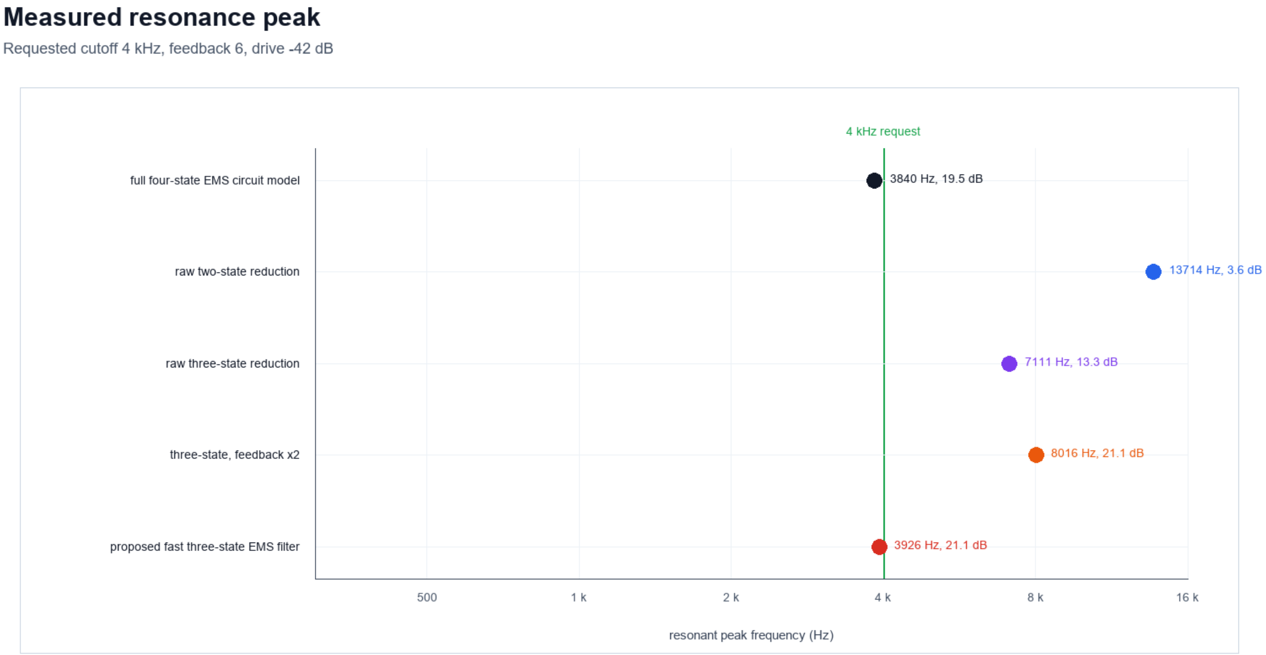
Measured resonance peak

Requested cutoff 4 kHz, feedback 6, drive -42 dB

Figure 2. Measured resonance peak for the full four-state EMS circuit model, the two-state reduction, the unretuned three-state reduction, feedback-only compensation, and the proposed fast three-state EMS filter.

4. Three-State Equations

The useful approximation is local current balance at the first ladder state:

```
dv1/dt = 0
0 = u + x1
x1 = -u
```

Substitution into the remaining equations gives the three-state system:

```
dv2/dt = s * (x2 + u)
dv3/dt = s * (x3 - x2)
dv4/dt = s * (-x5 - x3)
```

with:

```
x2 = tanh((v3 - v2) / (2 gamma))
x3 = tanh((v4 - v3) / (2 gamma))
x5 = tanh(v4 / (6 gamma))
u  = tanh((vin - (K + 1/2) v4) / (2 VT))
```

The runtime state is:

```
z = [v2, v3, v4]
```

The eliminated voltage can be reconstructed from:

```
v2 - v1 = 2 gamma atanh(-u)
```

but the audio output depends on `v4`, so the runtime filter does not need that reconstruction.

The discrete-time update uses a non-iterative midpoint step:

```
k1          = F(z[n], input[n], params[n])
z_mid       = z[n] + (h / 2) * k1
input_mid   = (input[n] + input[n - 1]) / 2
params_mid  = (params[n] + params[n - 1]) / 2
k2          = F(z_mid, input_mid, params_mid)
z[n + 1]    = z[n] + h * k2
```

Cutoff and feedback are evaluated per sample and interpolated into the midpoint calculation, so the filter remains usable under fast modulation.

5. Transfer Accuracy

Sweep-probe measurements compare the proposed fast three-state EMS filter against the full four-state EMS circuit model. After cutoff x0.50 and feedback x2.0 compensation, the resonant peak position and peak shape are close through the useful mid and high cutoff range.

Representative peak measurements are:

Requested setting	Full four-state EMS circuit model	Proposed fast three-state EMS filter
0.5 kHz, fb 6, -42 dB	480.5 Hz, 19.2 dB	492.2 Hz, 21.1 dB
2 kHz, fb 6, -42 dB	1921.9 Hz, 19.5 dB	2003.9 Hz, 20.7 dB
4 kHz, fb 6, -42 dB	3840 Hz, 19.5 dB	3926 Hz, 21.1 dB
4 kHz, fb 9, -42 dB	3972.7 Hz, 23.4 dB	4031.2 Hz, 25.3 dB
10 kHz, fb 9, -42 dB	9867.2 Hz, 23.5 dB	10113.3 Hz, 25.2 dB

The sweep-probe residuals for the proposed fast three-state EMS filter were:

Requested setting	Median residual	p90 residual	Max residual
0.5 kHz, fb 6, -42 dB	0.7567 dB	11.98 dB	19.30 dB
2 kHz, fb 6, -42 dB	0.1353 dB	3.466 dB	9.504 dB
4 kHz, fb 6, -42 dB	0.1321 dB	1.985 dB	9.706 dB
4 kHz, fb 9, -42 dB	0.1751 dB	2.404 dB	19.19 dB
10 kHz, fb 9, -42 dB	0.1687 dB	1.127 dB	14.41 dB

The 0.5 kHz row is the weakest sweep match. From 2 kHz through 10 kHz the residuals are much smaller, and the peak lands close to the full model. Figure 3 shows the 4 kHz, feedback 6 transfer overlay.

Figure 3. Sweep-probe transfer for requested cutoff 4 kHz, feedback 6, and drive -42 dB. The lower panel is proposed fast three-state EMS filter minus full four-state EMS circuit model.

6. Driven Harmonic Detail

A linear sweep does not show how a nonlinear filter redistributes a saw waveform. Figure 4 therefore uses a 30 Hz bandlimited-saw input at -12 dB. The plot shows the harmonic levels directly, with no per-trace normalization, so the visible height differences are part of the result.

At requested cutoff 4 kHz, feedback 9, and input -12 dB, the proposed fast three-state EMS filter follows the same resonant buildup as the full model, but the 4x and 8x versions put slightly different weight around the resonance and upper harmonics.

Saw harmonic detail: 4 kHz

30 Hz bandlimited saw, requested cutoff 4 kHz, feedback 9, input -12 dB

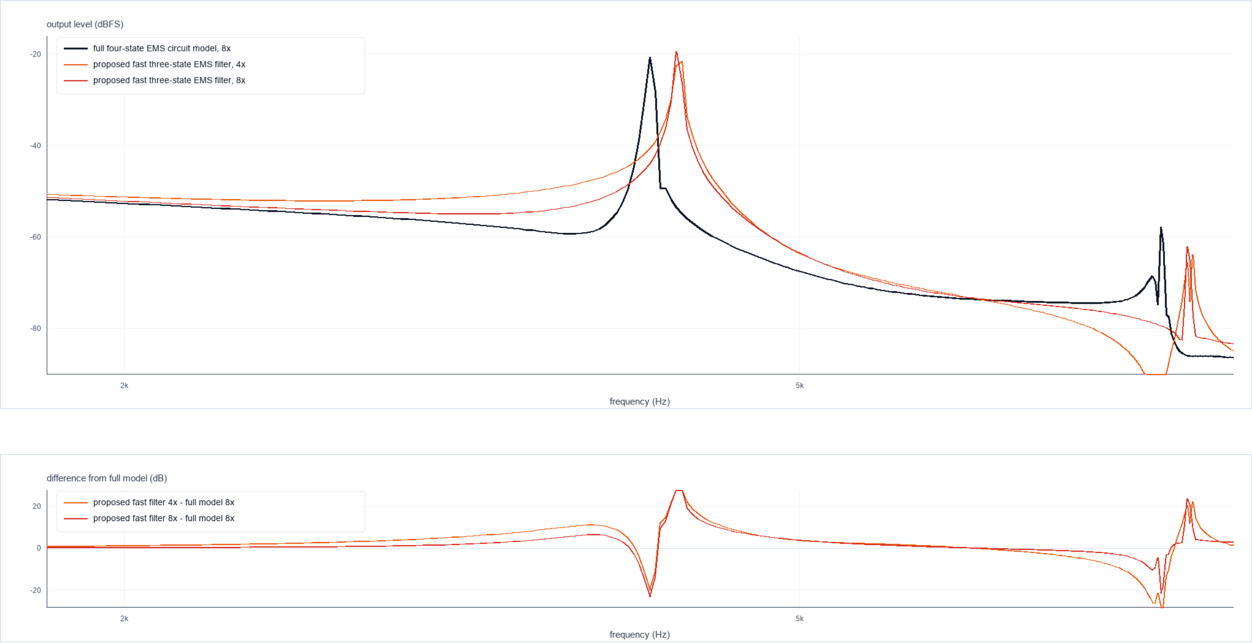

Figure 4. Saw harmonic-detail comparison for requested cutoff 4 kHz, feedback 9, and input -12 dB.

7. CPU Results

All CPU results in this paper use the AMD Ryzen 9 7900 benchmark machine. The benchmark ran at 48 kHz, 64-sample blocks, mono processing, Release build, and SIMD enabled. Four-state and three-state rows use the same patch, input level, channel count, oversampling factor, and modulation stream. Each table entry is the median of three 15-second repeats.

Model and condition	Median CPU	p95 CPU	p99 CPU	Peak CPU
full four-state EMS circuit model, static mono 4x	5.5800%	5.8275%	6.0750%	7.8075%
proposed fast three-state EMS filter, static mono 4x	1.1850%	1.2675%	1.3763%	2.9025%
full four-state EMS circuit model, modulated mono 4x	6.0300%	7.5450%	7.8300%	10.9725%
proposed fast three-state EMS filter, modulated mono 4x	1.4475%	1.5300%	1.6500%	3.2550%
full four-state EMS circuit model, static mono 8x	9.5513%	10.8600%	11.3400%	14.2575%
proposed fast three-state EMS filter, static mono 8x	2.3625%	2.5350%	2.7675%	4.5675%
full four-state EMS circuit model, modulated mono 8x	11.7675%	12.7575%	13.6088%	18.3975%
proposed fast three-state EMS filter, modulated mono 8x	2.8950%	3.1275%	4.1775%	5.7150%

The median speedups are:

```
static mono 4x:    5.5800 / 1.1850 = 4.71x
modulated mono 4x: 6.0300 / 1.4475 = 4.17x
static mono 8x:    9.5513 / 2.3625 = 4.04x
modulated mono 8x: 11.7675 / 2.8950 = 4.06x
```

Earlier calibration against the Razer Blade 15 found the Ryzen 9 7900 roughly 1.7x faster for measurements using the full four-state model and about 2.0x faster for measurements using the proposed fast filter. Approximate Razer Blade 15 values are:

Row	Ryzen 9 7900 measured	Razer Blade 15 estimate
full four-state EMS circuit model, modulated mono 4x	6.03%	about 10%
proposed fast three-state EMS filter, modulated mono 4x	1.45%	about 3%
full four-state EMS circuit model, modulated mono 8x	11.77%	about 20%
proposed fast three-state EMS filter, modulated mono 8x	2.90%	about 6%

8. Oversampling

At 48 kHz, 2x oversampling was not robust enough at high cutoff and high feedback. In the listening patch it could produce audible glitches and simulation-stability problems. At 4x, the proposed fast filter was stable in the tested patch. At 8x, informal listening preferred the result against the full four-state EMS circuit model: the low-frequency response felt more solid and less damped than 4x, and the high-resonance behaviour sounded closer to the full model.

The transfer measurement supports that preference in a high-cutoff row. At requested cutoff 10 kHz, feedback 9, and drive -42 dB, the full four-state model at 8x peaked at 9738 Hz and 25.1 dB. The proposed fast filter at 4x peaked at 11016 Hz and 24.8 dB. The proposed fast filter at 8x peaked at 9984 Hz and 25.4 dB.

Figure 5. Oversampling transfer check at requested cutoff 10 kHz, feedback 9, and drive -42 dB. The full four-state EMS circuit model is run at 8x. The proposed fast three-state EMS filter is shown at 4x and 8x.

9. Scope of the Method

This EMS reduction is not a port-Hamiltonian formulation, a neural model, a trajectory fit, or a hybrid fallback to the four-state solver. The process was:

1. start from the full four-state EMS ODE;
2. remove one upstream capacitor state by local current balance;
3. keep the downstream ladder states dynamic;
4. use a non-iterative midpoint update;
5. compensate the changed cutoff and feedback map;
6. measure transfer, driven harmonic detail, and realtime CPU.

The method is likely useful only for topologies where a state can be removed with a clear circuit meaning. It worked here because the eliminated EMS state is upstream of the output and the three remaining states still carry the resonant ladder memory. The same idea should be re-derived for each filter rather than copied as a generic shortcut.

10. Conclusion

The proposed fast three-state EMS filter keeps the source nonlinear diode laws, removes the first upstream capacitor state, and applies a static control retune of cutoff x0.50 and feedback x2.0. It is not numerically identical to the full four-state EMS circuit model, but it preserves the main resonant sweep shape in the range that matters most for the tested patches.

The CPU reduction is large. In the preferred 8x oversampled modulated case on the AMD Ryzen 9 7900, median CPU fell from 11.7675% for the full four-state model to 2.8950% for the proposed filter, a 4.06x speedup. In the lower-cost 4x oversampled modulated case, median CPU fell from 6.0300% to 1.4475%, a 4.17x speedup. On the Razer Blade 15, these rows are estimated at about 20% versus 6% for 8x and about 10% versus 3% for 4x.

The practical recommendation is 8x oversampling for the closest sound in demanding patches, with 4x as the stable lower-cost mode.

References

[1] M. Civolani and F. Fontana, "A nonlinear digital model of the EMS VCS3 Voltage-Controlled Filter," Proceedings of the 11th International Conference on Digital Audio Effects, Espoo, Finland, 2008.

[2] F. Fontana and M. Civolani, "Modeling of the EMS VCS3 Voltage Controlled Filter as a Nonlinear Filter Network," IEEE Transactions on Audio, Speech and Language Processing, 2010.

[3] S. Zambon and F. Fontana, "Efficient Polynomial Implementation of the EMS VCS3 Filter Model," Proceedings of the 14th International Conference on Digital Audio Effects, Paris, France, 2011.

[4] F. Fontana and E. Bozzo, "Explicit Fixed-Point Computation of Nonlinear Delay-Free Loop Filter Networks," IEEE/ACM Transactions on Audio, Speech, and Language Processing, DOI 10.1109/TASLP.2018.2847442.

[5] F. Fontana and E. Bozzo, "Newton-Raphson Solution of Nonlinear Delay-Free Loop Filter Networks," IEEE/ACM Transactions on Audio, Speech, and Language Processing, DOI 10.1109/TASLP.2019.2924842.