DSP: Gain, the Feedback Loop, also the Noise Gate

Author Maximilian Ardalan Bernhard

Institutions FH Joanneum, KUG

Course Design & Research 2

Date 2026-08-31

1. Noise Gate and Software Gain

The raw microphone signal carries both the sound I want, the tone reflected back from the room, and a noise floor that comes from the ADC, the microphone electronics, and the room itself. That noise floor matters here because the system is a feedback loop: left to itself, the loop would amplify the noise round after round until it saturated the signal.

The fix for this issue was implementing a noise gate before the mic's signal passes on. The gate's job is to attenuate the signal whenever it falls below a level threshold, so that the otherwise-silent loop does not accumulate noise.

The gate begins with an envelope follower, which tracks the signal's amplitude using asymmetric attack and decay, following the technique Zölzer describes in DAFX (U. Zölzer 2011). Attack is fast, so the envelope snaps upward the moment a transient appears and catches it. Decay is slow, so the envelope relaxes gradually instead of possibly clicking between sounds.

The gate opens and closes using two thresholds rather than one, a hysteresis that stops the gate from chattering at the boundary. The envelope must rise above the open threshold of three hundred to open the gate, and must fall below the much lower close threshold of one hundred and fifty to close it. Keeping the two thresholds apart prevents the gate from flickering open and shut on every tick of noise near the edge.

Switching the gate in a single step would make an audible click, so the gain it applies ramps smoothly instead. When the gate opens, it fades in by stepping toward full scale; when it closes, it fades out by stepping back toward silence. The result is a smooth fade rather than an abrupt switch, basically an equivalent of an AR[1] amplitude envelope.

Only after the gate does the signal meet the software gain, which is implemented in Q15 fixed-point arithmetic, where the value 32768 represents unity, a gain of exactly 1.0 times. The computation is a single multiply followed by a shift right of fifteen bits, with clamping to keep the result inside the sixteen-bit range. Q15 being a standard format to represent fractional values in fixed-point DSP.

[1] Attack & Release.

The gain is adjustable in small steps of about 0.01 times, from zero, silence, up to unity. That fine resolution is essential, because coarse steps are useless when tuning a feedback loop: too much gain and the loop howls, too little and the loop is lost entirely. The software gain is applied on the microphone side, in the `MicGimbal` A2DP source where the samples are read, but it must be balanced against the loudspeaker side of the loop as well. Both ends work together, and the balance between them is what determines whether the loop stabilizes or runs away.

2. Tuning the Feedback Loop

With the gain and the gate in place, the whole system becomes a positive feedback loop. The microphone captures the sound reflected from the wall, the processed signal is played back through the loudspeaker toward the same wall, and the wall reflects it back to the microphone again.

Whether that loop stabilizes or swells infinitely is governed by one simple condition: if the total loop gain exceeds one at any frequency, the signal at that frequency grows without bound and the system howls; whereas if the gain stays below one, the sound decays. Tuning the system is therefore mostly the task of keeping the loop gain just below the point of instability.

The wall sits at the heart of this loop and acts as a filter within it. A reflective wall does not reflect all frequencies equally: its geometry, its material, and the room it bounds emphasize certain frequencies, the room modes, and attenuate others, phenomena Beranek treats in depth (Beranek and Mellow 2019). When Lucier re-records his speech again and again, it is precisely this selective emphasis that morphs his words until leaving only the room resonances (Lucier 1969). Here the same filtering happens continuously, and moving the gimbals changes which frequencies are favored, effectively creating the live, steerable version of Lucier's process.

Fine gain tuning finds the sweet spot. A gain just below unity at the dominant frequency keeps the loop alive at that one resonant tone, producing a sustained, evolving resonance rather than uncontrolled howling. This is the equilibrium the installation is meant to inhabit: loud enough to shape the sound, and quiet enough to stay on this side of instability.

The noise gate is what makes that equilibrium reachable. Without it, the loop would amplify the noise floor continuously, and there would be no way to tell the useful signal from the noise. By silencing the loop whenever the signal is quiet, the gate keeps the noise from accumulating and hands the ear a clean starting point.

In the end, the gain must be adjusted by ear. I set the gain through the serial console, turn it up slowly, listen to the output, and correct course by what I hear. The ears are the primary debugging tool, and they remain faster than any measurement for finding the edge of instability in a live room.

Bibliography

Beranek, Leo Leroy, and Tim J. Mellow. 2019. *Acoustics: Sound Fields, Transducers and Vibration.* Second edition. Academic Press.

Lucier, Alvin. 1969. "I Am Sitting in a Room.".

U. Zölzer. 2011. *DAFX: Digital Audio Effects.* 2nd ed. Wiley. https://doi.org/10.1002/9781119991298.