Notes On Soft Mute And Analog Tuning In DSP Radios

On the surface, soft mute in the modern DSP-chipped receiver seems a bit mysterious. What is it? How does it work? Why do we have it? In this article we'll explore how soft mute works and explain its technical details.

As a side topic, but somewhat related to soft mute, we'll also tackle the pseudo-analog tuning of the Silabs 483x DSP chip and see what's up there. This is the chip used in many of the cheap portable and pocket radios found today. They are mechanically tuned with a dial knob and have a traditional AM band scale with dial indicator. But they aren't the pocket radio you remember from the old days.

So let's get right to it.

SOFT MUTE

What is soft mute?

Soft-mute is a further lowering of the audio level of the received signal when it drops below a prescribed signal-to-noise ratio. It was implemented in consumer grade DSP radios to provide a more "comfortable listening experience" for the casual listener and not the DXer. The idea is to relieve the listener from all that nasty low level "static" and "interference", or as Silicon Labs states: "....to attenuate the audio outputs and minimize audible noise in compromised signal conditions."

Soft mute attenuation is available in the Si473x digitally-tuned series of chips as well as the Si483x analog-tuned series of chips. The soft mute feature is triggered by the SNR (signal-to-noise) metric. The SNR value is directly readable by the chip's software when you tune to a station. The software reads the quality of the signal through its SNR value and makes soft mute changes accordingly. The SNR threshold for activating soft mute is programmable, as are soft mute attenuation levels, attack/release rates and attenuation slope.

The Tecsun PL-380, PL-310, PL-330, and other radios all may set different soft mute values than the chip's default values shown below. Settings for soft mute are initialized during the power up sequence.

The 4 soft mute parameters: Rate, Slope, Max Attenuation, Threshold.

Rate (default): 278 dB/second (range 1-255, actual figure 278 = setting * 4.35)

Determines how quickly the soft mute is applied/released when soft mute is allowed (enabled).

Slope (default): 2 dB (range 1-5 dB per dB below SNR threshold)

The attenuation slope for soft mute application - in dB of attenuation per dB SNR below the soft mute SNR threshold. Translated: how much audio attenuation is applied as the SNR and signal quality decreases. A setting of 2 will lower the audio by 2 dB for each 1 dB reduction of SNR below the starting threshold at which soft mute kicks in. An example: soft mute starts to kick in when the SNR decreases to 10 dB. At 10 dB, there is 0 dB of soft mute. When the SNR decreases to 9 dB, soft mute reduces the audio level by 2 dB. When the SNR decreases to 8 dB, soft mute reduces the audio level by another 2 dB (4 dB total). By the time the SNR hits 2 dB, the soft mute has reduced the audio level to a max of 16 dB. It will go no lower as the max soft mute has been applied. Note that every 6 dB of audio reduction is a halving of the audio voltage level. 12 dB of reduction is then 1/4 of the original audio voltage level. 16 dB (max soft mute) is a reduction of 84.2% (0.158).

Max Attenuation (default): 16 dB (range 0-63 dB, max attenuation of soft mute)

If set to 0, soft mute is disabled entirely.

Threshold (default): 10 dB (range 0-63 dB, SNR at which soft mute starts to engage). Silabs states, "for a tuned frequency".

Note that the Threshold setting is applicable only "for a tuned frequency". I take this to mean that soft mute is dis-engaged totally when not tuned to an exact 9 or 10 KHz channel, which is apparently why the 1 KHz off-tuning hack works.

What you're hearing when a signal's SNR lowers below the threshold and the soft mute kicks in is the Slope factor in action. The Slope factor is lowering the audio volume accordingly.

.....

How to defeat soft mute?

Soft mute can be somewhat minimized by increasing signal strengths to the radio by using a directly-coupled loop, passive loop or other inductively coupled antenna. What happens is you are increasing signal levels, thus improving the SNR, making the signal exceed the threshold where soft mute is engaged.

The other (original) hack is to tune off the channel by 1 KHz and raise the volume on the radio. Being off-channel disables soft mute.

.....

Two other interesting parameters effecting tuning and seeking, not related to soft mute.

AM Seek/Tune SNR Threshold.

SNR Threshold which determines if a valid channel has been found during Seek/Tune.

Specified in units of dB in 1 dB steps (0–63). Default threshold is 5 dB.

This tells us that when you do a scan, only stations with >5 dB SNR are eligible to be stored.

AM Seek/Tune Received Signal Strength Threshold (RSSI).

RSSI Threshold which determines if a valid channel has been found during Seek/Tune.

Specified in units of dBµV in 1 dBµV steps (0–63). Default threshold is 25 dBµV.

This tells us that when you do a scan, only stations with >25 dBµV RSSI are eligible to be stored.

.....

THE ANALOG TUNING ALGORITHM

The Silabs 483x series of chips are analog tuned and they have no digital LCD display. Tuning is accomplished through a tuning knob connected to a 100K ohm potentiometer. They attempt to mimic the analog tuning of the old traditional analog superhet radios when you "sweep" through a station's carrier. Silabs has developed a special tuning formula in software to simulate this. From the DXer's point of view it doesn't work. I've given a lot of thought to how their algorithm works in software.

Over the summer here in North America I have bought quite a few of the cheap Chinese analog-tuned DSP Ultralights. Though I have found some can be quite sensitive (like the Sangean SR-35 and the ultra cheap Dreamsky Pocket Radio), the SiLabs tuning algorithm is still wonky and masks a lot of weaker adjacent channel signals. It becomes tedious for serious DXing. Poor selectivity and overload problems can also be evident on these units, depending on the unit.

As stated, the problem with the current analog-tuned theory is that a weaker adjacent channel signal is masked deliberately if next to a more overwhelming signal.

A typical tuning scenario goes like this. Find a strong station where you know a weaker station sits right next to it on the adjacent channel. The weaker station would be strong enough to be received on a normal superhet radio. With the Silabs 483x radio, tune to the strong station's channel. Now tune to the adjacent channel (the weaker station). The strong station is still there, only at a slightly reduced volume. The radio is attempting to mimic tuning "through" a station like in the old days, increasing the strength of the station as you approach its channel center, then decreasing the strength as you depart. But where is the weaker station?

Here is what is happening in software (I think), preventing you from receiving the weaker adjacent channel.

Let's say the following numbers below, 0 | 5 | 20, represent frequencies 1020, 1030, and 1040 KHz. 1020 KHz has no signal on channel. 1030 KHz has a weak signal of SNR 5 dB. 1040 KHz has a strong signal of SNR 20 dB.

```
FREQ  1020   1030   1040
SNR =   0  |   5  |  20
```

In these DSP radios, hardware generates a tuning interrupt in software when changing the tuning knob. It causes the software to take over and analyze what just happened.

Initially, tune to 1040 KHz from somewhere above in frequency and begin receiving the strong station.

Now tune to 1030 KHz. Software then does this:

1. The tuning interrupt is generated.

2. Hard mute the audio.

3. With audio off, electronically retune to the new channel (1030) and test the new channel's SNR. If valid (SNR >= 5 dB), remain on this new channel and unmute the audio. If not valid (SNR < 5 dB), electronically retune back to the original channel (1040) and reduce the audio 6 dB and unmute. The dial will point to 1030 even though we're hearing 1040.

Now tune to 1020 KHz. Software then does this:

1. The tuning interrupt is generated. Remember, though the radio dial shows 1030 KHz, the radio is still electronically tuned to 1040 KHz.

2. Hard mute the audio.

3. With audio off, electronically retune to the new channel (1020 this time) and test the new channel's SNR. If valid (SNR >= 5 dB), remain on this new channel and unmute the audio. If not valid (SNR < 5 dB), electronically retune back to the original channel (1040) and reduce the audio an additional 6 dB and unmute. The dial will point to 1020 even though we're hearing 1040.

Additionally, for each of the two scenarios above, we must also be sure in step 3 that the original 1040 channel maintains a SNR above the SNR of the newly tuned channel or we force-tune to the new channel.

Electronically retuning the DSP chip is simply a matter of electronically setting the proper internal capacitance to resonate with the ferrite coil at the desired frequency. It's done with a single software command.

It's complicated.

If you start at 1020 KHz then approach 1030 from below the situation changes, as we are comparing 1030 to 1020 now, 1020 having no signal at all. If 1030 is a valid channel (SNR >= 5 dB) then the DSP chip tuning remains at 1030, the hard audio mute is unmuted, and the station is received. Drawing from this scenario, we can conclude that if we approach a weak signal from the right tuning direction that we might be able to hear it.

Compounding the problem, these 483x chips also generally have soft mute enabled, which may mask very weak stations. The weak station will still need to overcome the soft mute threshold to some degree.

According to Silabs, this new tuning algorithm has been "audience tested" to a positive level of acceptance. The best approach for the DXer would be to have a radio where soft mute is disabled altogether and no tuning algorithm so that when you move the tuning dial it always changes the frequency.

Surprisingly, this wonky tuning algorithm can be somewhat minimized by increasing signal strengths to the radio by using a directly-coupled loop, passive loop or other inductively coupled antenna. What happens is you are increasing signal levels, thus improving the SNR, so the signal meets the threshold requirements for a valid signal. The radio then tunes to the proper signal and frequency.

A description of even weirder analog tuning anomalies can be read here:

Notes On The XHData D-219 Analog DSP Radio

I hope this analysis of soft mute and the DSP analog tuning mechanic has proven useful and interesting. All technical data has been gleaned directly from Silabs data sheets for the respective 473x and 483x DSP chips. The programming guide for these chips was particularly helpful in understanding the operation of soft mute.

Posted by RADIO-TIMETRAVELLER at 6:20 AM

2 comments:

Subs said...

Hi

I got myself a SONY ICF19. I live in Abu Dhabi, UAE.

In the lower end of AM scale the tuning is bit off. 810kHz All India Radio, Rajkot comes in clear and strong but at a little below 800kHz mark. In an analogue radio it is as simple as turning the Osc coil to adjust. On DSP radios like SONY ICF19 how do I do it? It is new and came in only yesterday but not worth sending back to Amazon. So thinking of retune if I can.

Otherwise it is a lovely radio, Right size, very good pleasing SONY sound, tuning seems rock steady (I get 1476kHz from Dubai but it drifts in my other radios even ATS909). FW is very strong reception - was receiving stations as far as 100 miles or more without extending the antenna!!

October 13, 2022 at 12:49 AM

RADIO-TIMETRAVELLER said...

Hi Subs. The only DSP radio I know which allows user-accessible tweaking of tuning to correspond to dial setting is the digital Tecsun PL-880. You definitely won't find this in less expensive models. The DSP chips also aren't traditional superheterodynes with a local oscillator, not in the traditional sense.

The so-called newer "analog tuned" DSP models only mimic analog tuning, as you already know. The dial scales on these type radios are a linear scale, not logarithmic, so it should be easy for the manufacturer to lay them out accurately. The DSP chip is actually tuned using a linear 100K resistor potentiometer. Problems may occur with the linearity of the pot or the layout or positioning of the scale itself. Possibly if you pull the radio's case off you could reposition the scale.

My best analog-tuned DSP radio is the CC Radio EP Pro from C. Crane Company. It goes for about $100 US, and is a very sensitive unit, using an 8 inch ferrite. It also includes a twin-coil antenna tuning circuit. But even this radio suffers a little from dial misalignment.

If you have a sensitive radio in the Sony '19, I would keep it.

Thanks for your comments.

Bill