Understanding Controlled Envelope Single Sideband

CESSB increases your average transmitted power without making you sound more "processed."

Dave Hershberger, W9GR

Controlled Envelope Single Sideband (CESSB) is a new system that allows your rig to output more average power while keeping peak envelope power (PEP) the same. Single sideband (SSB) modulation produces large envelope peaks. To avoid distortion and splatter, these peaks have been traditionally controlled by Automatic Level Control (ALC) systems, which reduce the transmitted power level so that the envelope peaks do not get clipped in the RF power amplifier. CESSB avoids the clipping problem by not generating the envelope peaks in the first place. CESSB is generated using a modified RF clipper system that allows the average transmitted power to nearly double. CESSB is not an audio processor or something that goes between your microphone and your rig.

Introduction

Most RF power amplifiers are peak power limited. So if the peaks are big, we must turn down the amplitude of the entire signal to keep it within the limitation of the power amplifier. Conversely, if we could control the peaks and reduce them, then we could turn up the amplitude of the overall signal, increasing our average power.

You might think that by accurately limiting the audio level going into an SSB transmitter, the output envelope amplitude would also be accurately limited. It's not. You might also think that RF clipping would accurately control the output envelope amplitude of an SSB transmitter. It doesn't. If we accurately peak-limit audio into an AM or FM transmitter, the modulation *is* accurately limited. This is definitely not the case for SSB.

Overshoots in SSB Modulators

Figure 1 shows a reference audio tone preceding speech. The peaks of the audio tone and of the peak-limited speech input are the same amplitude, but overshoots of conventional SSB modulation upset that relationship. If you input a single tone into an SSB transmitter, you will always get a predictable output amplitude. If you set your audio tone to a level producing the maximum PEP rating of the RF power amplifier, then you might think that is the correct peak level for speech audio. It is not. If you accurately compress and limit your audio signal, as shown in Figure 2, then apply that well-limited audio signal to a conventional SSB modulator, it will overshoot like crazy (see Figure 3).

SSB modulators overshoot because of the Hilbert transform operation. All SSB modulators perform a Hilbert transform, some indirectly. A Hilbert transform applies an audio phase shift of 90 degrees for all frequencies within its range. The phasing method of SSB performs the Hilbert transform directly using a phase shift network. See the sidebar, "The Hilbert Transform."

The initial tone bursts in Figures 2 and 3 have an RF amplitude of 1.0, corresponding to maximum PEP of the subsequent RF amplifier. We end up with a 61% overshoot in Figure 3. That means the PEP is 2.59 times what it should be. So we must turn down our RF output power to avoid flat-topping caused mostly by overshoots generated by the Hilbert transform operation. Shifting the phase of a single frequency does not change its amplitude. However, when the phases of all the many different frequencies in speech are changed, new points of constructive and destructive interference create peaks that overshoot.

ALC has traditionally dealt with SSB envelope overshoot. We sense RF peaks, then reduce the gain of the transmit path to avoid flat-topping and splatter. But traditional ALC acts too late in the process. The first overshooting RF envelope peak gets clipped, and that causes some splatter. ALC then turns down the power of your SSB transmission.

With current SSB technology, we may compress and limit the audio speech. We then apply that to an SSB modulator. It overshoots like crazy, so we use ALC to reduce our transmitted power and avoid flat-topping of the signal.

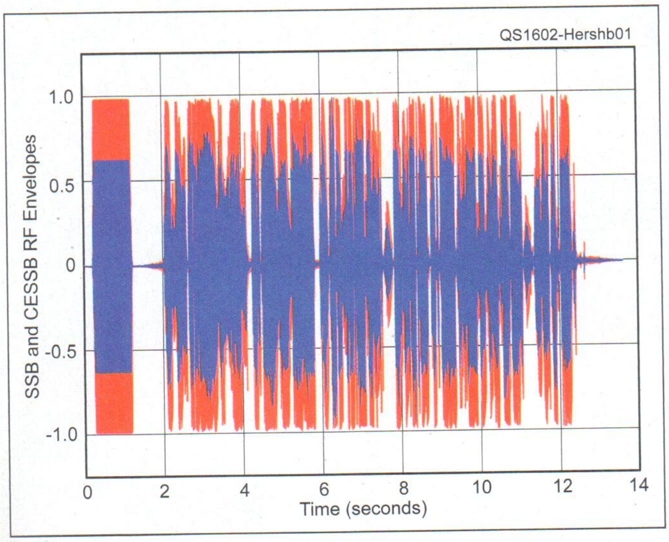

Figure 1 — Conventional SSB (blue) compared to CESSB (red) have the same PEP, but the average power of CESSB is 2.28 times greater.

About RF Clipping

RF clipping helps, but does not solve the problem. Figure 4 shows how RF clipping is done in conventional analog circuitry. First, a double sideband balanced modulator generates an IF signal, which is applied to the SSB band-pass filter. The filter method will overshoot just as the phasing method does. We then clip the IF signal.

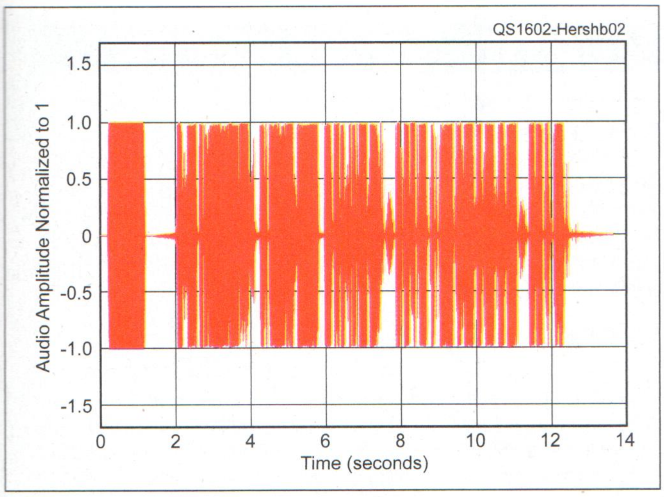

Figure 2 — Accurately peak-limited audio.

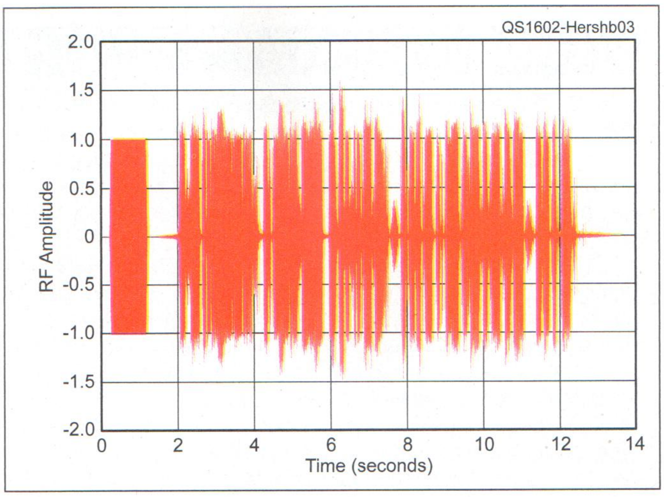

Figure 3 — SSB RF waveform when the signal of Figure 1 is applied to an SSB modulator.

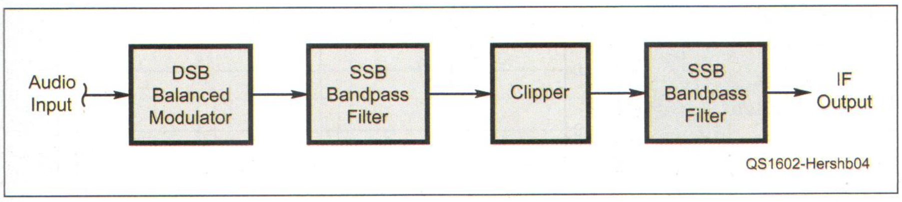

Figure 4 — Conventional analog RF clipper.

Clipping produces RF or IF harmonics, and it produces out-of-band intermodulation distortion components that would cause splatter if they were not removed. So, we must run the clipped SSB signal through another SSB band-pass filter. The problem is that this second SSB filter also overshoots. The more you clip, the more overshoot you get from the second filter.

Figure 5 shows the output of the RF clipper. The overshoot is now 37% instead of 61%. But PEP is still 188% of what we want — almost double the desired power. My simulation used a linear phase filter, but a practical analog crystal or mechanical filter would have produced even more overshoot because of group delay.

Spectral truncation and RF cycle clipping are two mechanisms that cause the overshoot in a conventional RF clipper. Spectral truncation is the removing of the out of band clipping distortion components. A conventional RF clipper does not clip the envelope. Rather, it clips individual RF cycles. There is a difference. RF cycle clipping creates, in the limit, square waves. After band-pass filtering, only the fundamental component remains. But the fundamental component peak of a square wave is bigger than the square wave amplitude by $4/\pi$. That would be 27% overshoot, or somewhat less in practice because we are not producing perfect square waves.

Figure 6 shows short bursts of RF sine waves with 0, 3, 6, and 10 dB of clipping,

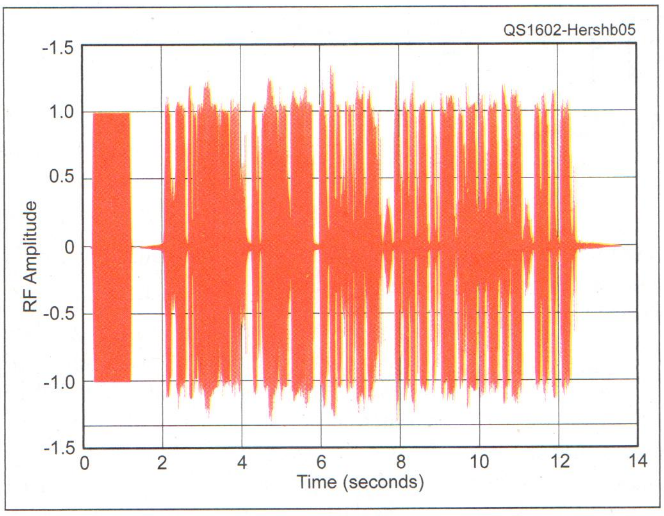

Figure 5 — SSB RF waveform from a conventional analog RF clipper.

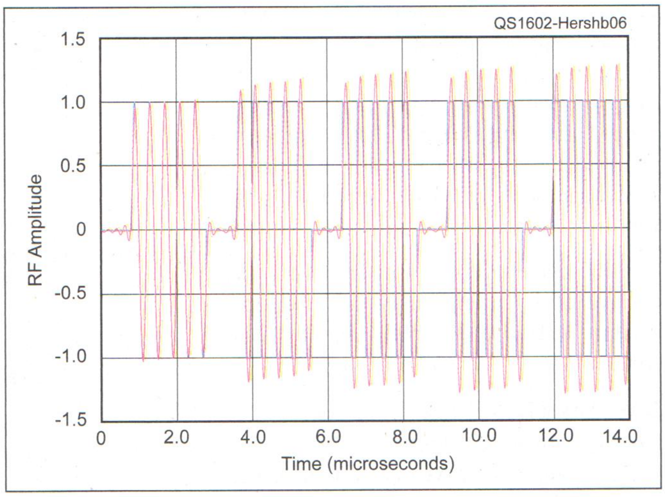

Figure 6 — Effect of clipping RF cycles at 0, 3, 6, 10 dB, and "infinite clipping."

The Hilbert Transform

SSB modulators overshoot because of the Hilbert transform operation (an audio phase shift of 90 degrees for all audio frequencies in its range). A Hilbert transform occurs in every SSB transmitter. A Hilbert transform takes nicely peak-limited audio and turns it into a peaky mess with lots of overshoot. The RF envelope will be likewise peaky.

Figure A shows the phasing method of generating SSB, as used by some of the earliest SSB transmitters, like the classic Hallicrafters HT-37. Some modern DSP rigs also use the phasing method. In Figure A, we split the input audio, which has a phase versus frequency of $\phi(\omega)$, into two paths. One path remains $\phi(\omega)$, while the second path passes through a $\phi(\omega)$ + 90-degree phase shift network. The output of this second path is the Hilbert transform of the first path signal. We then apply these two audio signals to two mixers driven in quadrature (RF or IF phased 0 and 90 degrees to each other), and combine them to get single sideband.

A filter-type SSB transmitter also performs a Hilbert transform indirectly. Think about the case where you are receiving an SSB signal with a BFO at the correct frequency. Now if you shift the BFO phase by 90 degrees, it shifts all of the demodulated audio frequencies by 90 degrees too. There's your Hilbert transform. The I and Q components of every SSB signal are always related by a Hilbert transform.

A Hilbert transform has a perfectly flat frequency response. Its phase response causes the problem by changing how the different frequencies add and subtract with one another. Where different frequency components might have subtracted before, they might add after passing through the Hilbert transform, creating new peaks.

SSB Envelope of a Square Wave

About the worst thing you could put into a Hilbert transform is a square wave. Figure B shows a band-limited 100 Hz square wave (blue) applied to a Hilbert transform. A square wave is the summation of sine harmonics having coinciding zero crossings. After the Hilbert transform, the sines

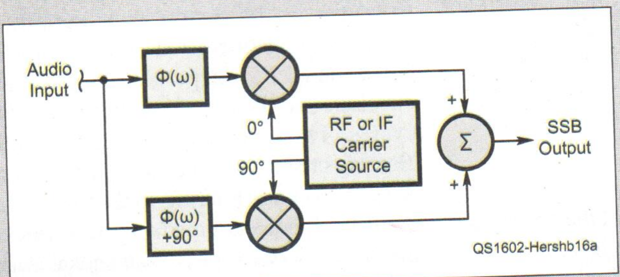

Figure A — SSB using the phasing method.

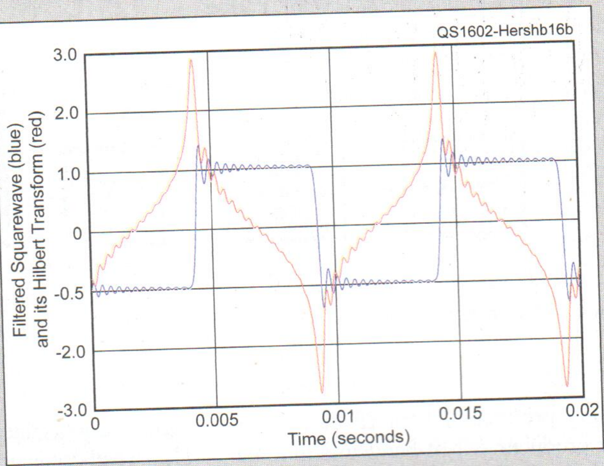

Figure B — A band-limited square wave (blue) and its Hilbert transform (red).

followed by a square wave, which is effectively infinite clipping of a sine wave. The red trace shows the output of a low-pass filter that removes the harmonics. As you can see, there is overshoot, because the fundamental harmonic component (red) of a clipped sine wave is bigger than the (blue) clipped sine wave. When the harmonics are removed, the amplitude is no longer limited. This is a reason why RF clippers don't work very well.

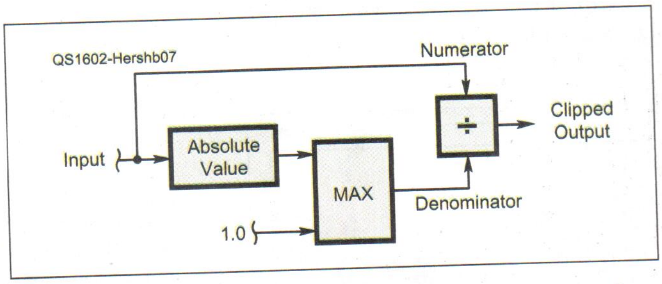

Figure 7 — Clipping by division.

The RF cycle overshoot problem can be eliminated by doing baseband envelope limiting instead of RF-cycle clipping. But some overshoot will still remain.

Baseband Envelope Clipping

Even though an RF clipper does not fix the problem, baseband envelope clipping is the first step in CESSB. Baseband envelope clipping is an improved version of RF clipping, but it is done at audio frequencies instead of radio frequencies. Baseband envelope clipping is performed on the two baseband audio signals. We limit the envelope amplitude directly, rather than trying to clip RF cycles.

Figure 7 shows how to make a clipper using a divider rather than diodes. This is more complicated than a diode circuit, but it can be modified to process the complex (in the mathematical sense) dual-path signals used for SSB baseband signals.

We apply an input signal to an *absolute value* circuit, which can be just a full-wave rectifier. The output of that is applied to a function whose output is the greater of its two inputs. You can think of that *max function* as an analog *diode OR gate*. That makes the denominator input to our divider. The input signal is the numerator.

become cosines (90-degree phase shift) and now waveform peaks coincide. That is why the red waveform in Figure B looks so peaky. The overshoot from the square wave is almost 200%. This is important because it shows us that doing speech clipping before an SSB modulator will produce huge peaks. The rms value (average power) of the two waveforms is the same, but the peaks are very different.

Figure C shows the envelope of an SSB signal with a filtered square wave as its input. Anything over an amplitude of 1.0 is an overshoot. This envelope spends almost all of its time over 1.0, and overshoot is more than 200%.

SSB Envelope Overshoots with Speech

We don't normally transmit square waves. Figure D shows what happens with speech. The blue trace is our input audio from an audio compressor and peak limiter. The initial block shows a tone burst at 100% modulation; that is, the audio level corresponding to maximum PEP of the RF power amplifier. Notice that the blue speech audio does not exceed the level of our 100% tone burst.

The red trace is the output of a Hilbert transform filter. All of the audio phases have been shifted 90 degrees relative to the blue trace. Our nicely peak-limited audio isn't peak-limited anymore. In fact, the overshoot is 59% for this speech sample. There is a lot of "fuzz" on the Hilbert transform output. With some voice waveforms, asymmetry may also appear.

Notice that the tone burst at the beginning doesn't produce any overshoot. That's because it is just a single frequency. Overshoot happens only when there are multiple frequencies present, and where the phase shift causes new points of constructive and destructive interference among the various frequency components. Shifting the phase of a single frequency doesn't cause any overshoot, whereas shifting the phases of multiple frequencies does.

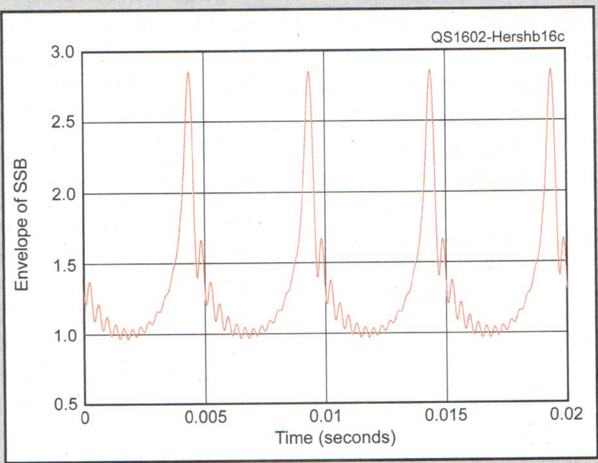

Figure C — SSB envelope resulting from the band-limited square wave of Figure B.

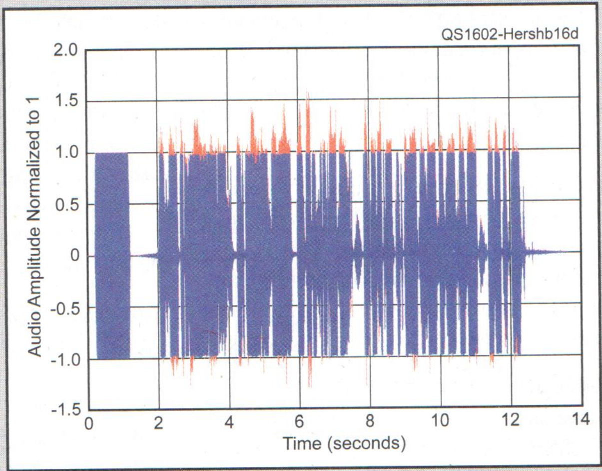

Figure D — Peak-limited speech audio (blue) and its Hilbert transform (red).

So, an input signal that is between –1 and +1 gets divided by 1 and there is no change to the signal. But if the input signal goes to 1.5 for example, then the denominator becomes 1.5, and 1.5 divided by 1.5 is 1. The signal is clipped.

This works for single input signals. But unlike a diode clipper, this method can be expanded to work with the dual (mathematically) complex signal pairs we use to generate SSB. These two audio signals (α and β) can come from a phasing-type 90 degree phase difference network.

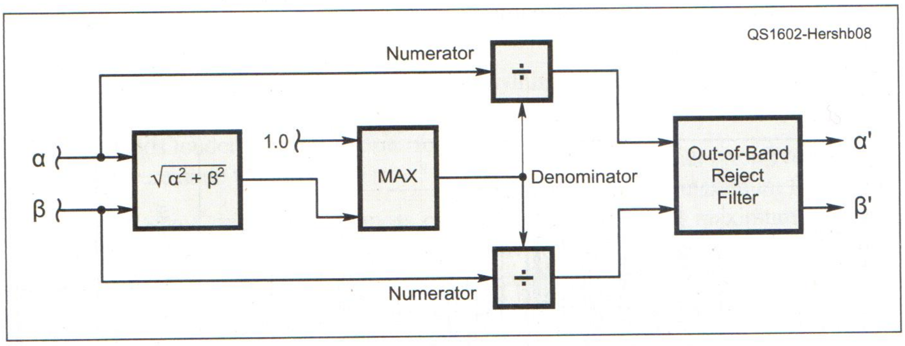

Figure 8 — Baseband envelope clipper for complex signals.

Figure 8 shows the baseband envelope clipper. We want to limit the SSB signal envelope amplitude. The envelope signal is the square root of the sum of the squares of the two inputs (a *modulus function*). So we have a function that generates the envelope from the α and β signals. Next we have the *max function*, which selects either the envelope signal or 1.0, whichever is greater. That becomes the denominator signal for both dividers. When we divide the α and β signals by the same denominator, we change only the length of the resulting vector, but not its phase.

After this clipping, we have to clean up the spectral mess using a filter that rejects the out-of-band components. This method does not clip any RF cycles, because there are no RF cycles, so it avoids the 4/π overshoot shown earlier, but it still overshoots.

overheat, or worse. A power supply designed for full power key-down CW should be adequate.

(5) CESSB will have systems implications. Current SSB technology doesn't do very much baseband audio processing, because conventional SSB overshoots like crazy. So today's SSB radios rely on ALC and sometimes RF clipping. As we have shown, those techniques are really not very effective. With CESSB the audio processing can and should be moved backwards into the baseband audio domain. Modern broadcast-type multi-band limiting and distortion-canceled smart clipping can be applied, with much better results than a simple wideband RF clipper. I think we will see the development of some very good speech processing algorithms, which can be combined with CESSB for no overshoot.

(6) Your SSB signal will be louder without sounding more "processed."

CESSB is something that needs to be built into the radio (or at least allowed for) from the beginning. If the radio is software defined, then new code can be written to implement CESSB. CESSB has been included in the firmware for the FlexRadio 6000 series since the summer of 2014. FlexRadio has made both laboratory and on-air tests to verify the usefulness of CESSB.

TAPR makes a family of boards that you can use to build your own software defined radio.[4] Warren Pratt, NRØV, has written CESSB code for the openHPSDR library for the TAPR boards.[5]

CESSB should be coming to new radios from other manufacturers as well. I am placing this technology in the public domain and, in particular, to the "ham domain," royalty free.

Retrofitting CESSB

While it is theoretically possible to retrofit existing rigs for CESSB, it would be quite difficult in practice. CESSB is normally tightly integrated into the SSB modulator. However, in some special situations, CESSB processing can be done externally. Older analog radios are generally not linear phase, so I don't think that there will be any Heathkit HW-101 radios running CESSB anytime soon.

See also, "External Processing for Controlled Envelope Single Sideband" in the January – February 2016 issue of *QEX* for details on the external processor approach.[6]

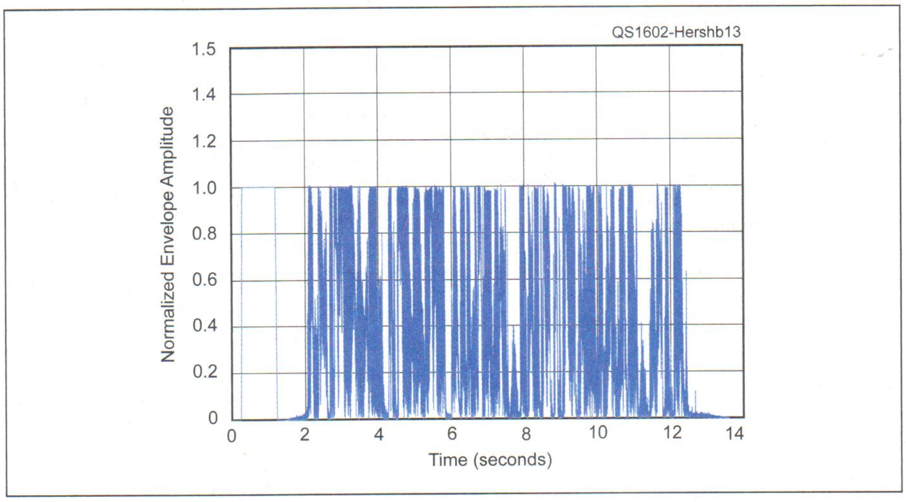

Figure 13 — CESSB RF envelope at the output of the baseband envelope "more than a clipper."

Conclusion

SSB can be generated without producing envelope overshoots beyond a defined maximum PEP level. This allows higher average power while controlling PEP. The technique is nonlinear, but virtually inaudible, and does not make your signal sound more processed. The technique consists of a conventional baseband envelope clipper, followed by a second envelope clipper that reduces its instantaneous gain more than a conventional clipper would. This is equivalent to a cascade of seven or more conventional envelope clippers.

CESSB is intended for use with voice. Digital modes, especially those with non-constant envelope functions, will be transmitted with less nonlinear distortion if CESSB is turned off.

Notes

1 The two audio files are available from **www.arrl.org/qst-in-depth**.

2 Information and the latest version of *GNU Octave* is available from **www.gnu.org/software/octave**.

3 D.L. Hershberger, W9GR, "Controlled Envelope Single Sideband," *QEX*, Nov/Dec 2014, pp 3 – 13, **www.arrl.org/files/file/QEX_Next_Issue/2014/Nov-Dec_2014/Hershberger_QEX_11_14.pdf**.

4 S. Cowling, WA2DFI, "The High Performance Software Defined Radio Project," *QEX*, May/June 2014, pp 3 – 13.

5 TAPR openHPSDR library project (**openhpsdr.org/**).

6 D.L. Hershberger, W9GR, "External Processing for Controlled Envelope Single Sideband," *QEX*, Jan/Feb 2016, pp 11 – 14.

ARRL Life Member Dave Hershberger, W9GR, was first licensed in 1965 at age 14 as WN9QCH. Dave holds a Bachelor's degree in math from Goshen College and Bachelor's and Master's degrees in Electrical Engineering from the University of Illinois. He has been awarded 19 US patents. Dave is Senior Scientist at Continental Electronics. His recent projects include two ATSC digital television broadcast exciters with adaptive linear and nonlinear precorrection, a DSP based FM/HD Radio® exciter with adaptive precorrection, new high-power uplink transmitters for the JPL/NASA Deep Space Network, and a 2.4 MW VLF transmitter system. You can reach Dave at 10373 Pine Flat Way, Nevada City, CA 95959; **w9gr@arrl.net**.

For updates to this article, see the *QST* Feedback page at www.arrl.org/feedback.

VOTE

Did you enjoy this article? Cast your vote at www.arrl.org/cover-plaque-poll