

Controlled Envelope Single Sideband

Introducing Controlled Envelope SSB; greatly increase your SSB “talk power” by accurately limiting envelope peaks in the SSB modulator. Generate SSB without the big overshoot peaks that make ALC necessary with conventional SSB modulators. Watch your wattmeter read higher than before.

Abstract

Achieving simultaneous accurate control of both amplitude and bandwidth is a difficult problem. When amplitude-limited audio is filtered to limit its bandwidth, the filter may overshoot substantially. It loses its amplitude limiting ability. If the resulting overshoots are clipped, the amplitude is controlled but the signal’s bandwidth increases because of the clipping distortion. The signal loses its bandwidth limiting. Systems exist for correcting audio low-pass filter overshoot. But single sideband (SSB) is a more difficult problem, because of the inevitable Hilbert transform regardless of the method of SSB generation. ALC systems reduce the amplitude of an SSB signal in response to overshooting envelope peaks. Fast ALC may result in clipping and splatter. Slow ALC will significantly reduce transmitted power. This paper presents a method for generating SSB without system overshoots. The result is higher transmitted power without audible distortion.

Objective

One of the reasons for the existence of the Amateur Radio Service is the development of new techniques for radio communication. I hope that this paper will present a useful method for improving the effectiveness of SSB transmitters for both amateur and commercial applications. This technique is being placed into the public domain and in particular, the “ham domain.”

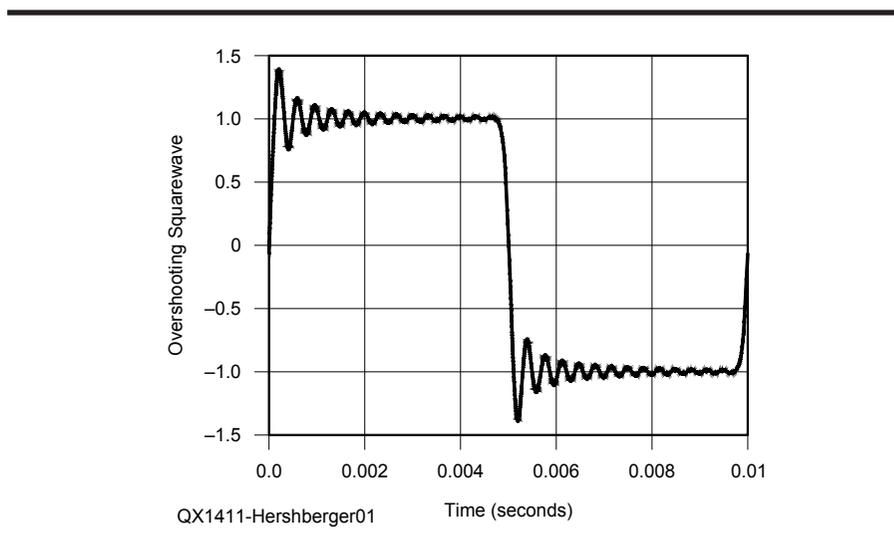


Figure 1 — 100 Hz square wave filtered by 3 kHz elliptic low-pass filter.

Benjamin Franklin expressed it well: “As we enjoy great advantages from the inventions of others, we should be glad of an opportunity to serve others by any invention of ours; and this we should do freely and generously.”

Background

In the 1970s, FM stereo broadcasters’ on-air loudness was affected by overshoot in the sharp cutoff 15 kHz low-pass filters (or 19 kHz notch filters) used in stereo generators of the time. These filters were

necessary to prevent crosstalk between the left plus right and left minus right subchannels. As the stereo generator’s filters would overshoot, the modulation would have to be reduced to keep infrequent peaks from exceeding 100% modulation. A number of manufacturers responded with systems that controlled overshoot accurately, allowing a loudness increase of 2 to 4 dB without any audible increase in distortion. The overshoots themselves contained relatively low energy. Eliminating them resulted in a substantial loudness

increase without any perceptible increase in distortion. Such overshoot control systems are still in widespread use today.

In this paper a classic method for real signal low-pass filter overshoot control is presented. Then the technique is extended to the complex baseband signals used to generate SSB.

Overshoot Control for Low-pass Filters

The problem with low-pass filters is that they overshoot on amplitude-limited signals

such as square waves. Figure 1 shows what happens when a 100 Hz square wave signal, accurately limited to a value of 1.0, is filtered by a sharp cutoff 3 kHz elliptic low-pass filter. It overshoots substantially.

There are several classes of non-overshooting or low-overshoot low-pass filters. These include Gaussian, Bessel, transitional, and parabolic filters. These are all linear filters. All of them have a roll-off characteristic that is far too slow to be useful for FM broadcasting or SSB generation, however. Therefore, an effective system for

overshoot control must be nonlinear.

My system for overshoot control of low-pass filters (expired US patent #4,134,074) started from the notion that overshoot could be reduced, but not eliminated, by clipping and refiltering. Clipping controlled the overshoots, but created harmonic distortion beyond the cutoff frequency of the filter. A filter after the clipper would remove the harmonic distortion — but it would overshoot. The overshoot of the second filter was not as high as the first filter, however. Theoretically, the clip, filter, clip, filter process could be repeated *ad nauseum* until the overshoot was insignificant. Figure 2 shows what happens with seven stages of filter, clip, filter, clip, and so on.

One cycle of a square wave is shown after the first, second, third stage, and so on. Even after the seventh pass, the overshoot is still over 4%. This method does not converge very quickly. Such a system would be impractical in analog circuitry. So, I came up with a method to make the system converge in “one fell swoop.” What was needed was something that did “more than clipping.”

A clipper can be visualized as a circuit that creates a series of peaks, which are then subtracted from the input signal. That is, it creates “clippings” that are the “tops” (and bottoms) of the input waveform. These are subtracted from the input to create a clipped waveform. The gain of the “clippings” in a simple clipper is unity.

Figure 3 shows what happens when the overshoots are simply clipped off. The overshoot “clippings” (lower amplitude trace) may be separated by subtracting the output from the clipper from its input. When the square wave with clipped overshoots is filtered, the second waveform in Figure 2 above is the result. The overshoot is reduced but not eliminated.

Something that would do “more than clipping” could apply a gain factor to the “clippings” — such as a gain of 2. The original overshooting signal shown in Figure 1 with such processing is shown in Figure 4.

Note that the overshoots that would have exceeded 1.0 are turned around and made to go the other way. The signal of Figure 4 contains out of band distortion and should be low-pass filtered with linear phase. When the “more than a clipper” processes the signal, and the result is linear phase low-pass filtered, the result is a large reduction in overshoot. Almost all of the overshoot is removed in this single more-than-clip and filtering process. The result is shown in Figure 5. The block diagram for this system is shown in Figure 6.

Audio is assumed to come from a peak limiter device, which ensures that nothing over 100% is being applied to the system.

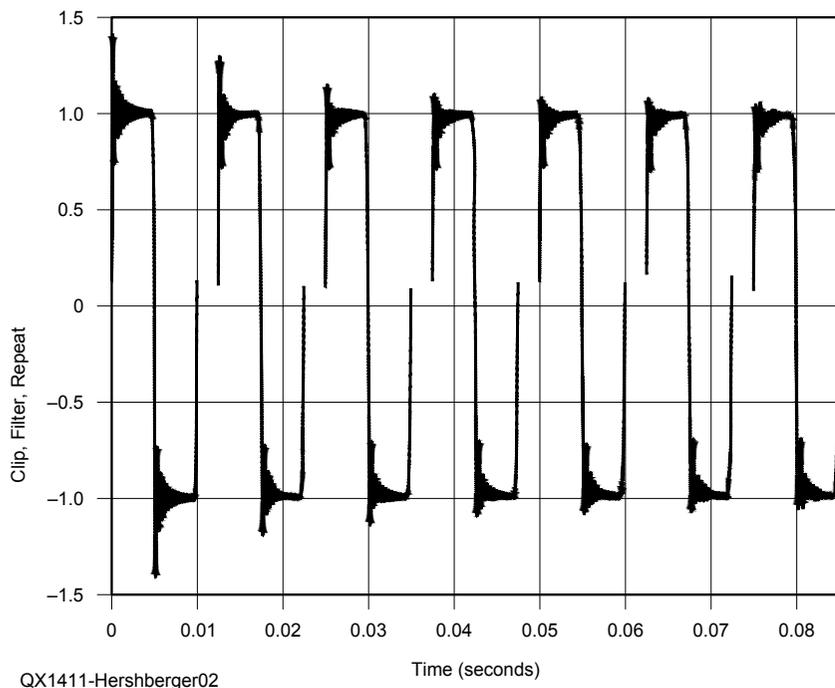


Figure 2 — Effects of repeated clipping and filtering.

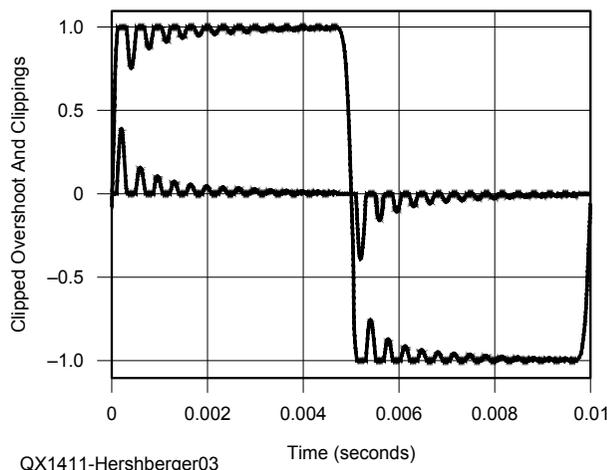


Figure 3 — Overshoots clipped off and separated.

This signal is low-pass filtered, and the low-pass filter will overshoot. The output of the low-pass filter is clipped. The output of the clipper is subtracted from the clipper's input, to obtain the "clippings." The clippings are then amplified by a factor of 2. They are then subtracted from the low-pass filter's output to create the waveform of Figure 4. That waveform is then low-pass filtered again with a linear phase low-pass filter. The result is simultaneous control of both peak amplitude and bandwidth.

There are two possible enhancements to this system. First, the simple gain factor of 2 generally results in insufficient overshoot correction at lower amplitudes, and too much correction at full amplitude. So rather than applying a simple gain factor of 2 in the "gain" element shown in Figure 6, the clippings are divided by the greater of 1.0 or the absolute value of the clipper input. Then a slightly lower gain factor of 1.9 is applied. The signal shown in Figure 4 has been processed this way. As an expression, the correction is:

$$corr(t) = \frac{1.9 \cdot clippings(t)}{\max[1, |lpf(t)|]} \quad [Eq 1]$$

Where:

corr is the additive correction signal

clippings is the input of the clipper minus its output

max is the function that returns the greater of its two inputs

lpf is the output of the first low-pass filter

The vertical bars operator is absolute value

Second, there is a way to reduce the signal processing complexity. Rather than passing the entire signal through the second low-pass filter, just the "clippings" may be passed through the filter if the rest of the signal is delayed to match the delay of the second low-pass filter. This modification is shown later in Figure 13.

When we generate SSB, the system as described above will not work for overshoot

control. In DSP, single sideband is usually generated by creating two orthogonal baseband audio signals, such as the well-known phasing method. The phasing method

as implemented in analog hardware required very close matching and phasing of the two signal paths. This was difficult to do in analog

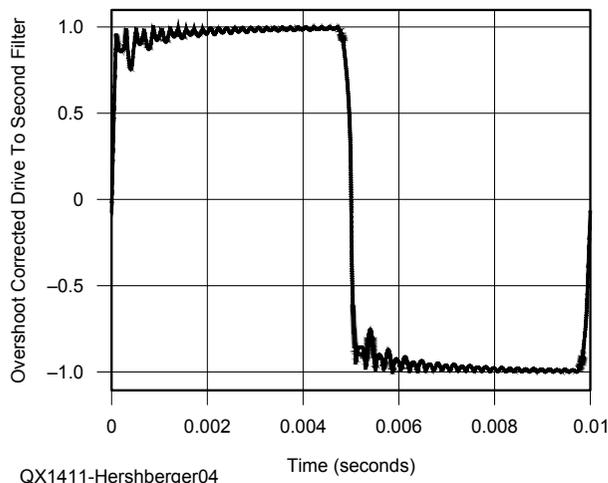


Figure 4 — 100 Hz filtered square wave processed with "more than clipping."

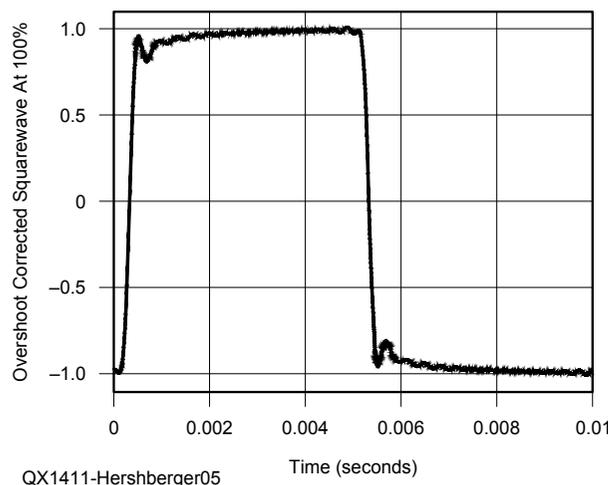


Figure 5 — Square wave filtered with overshoot compensation.

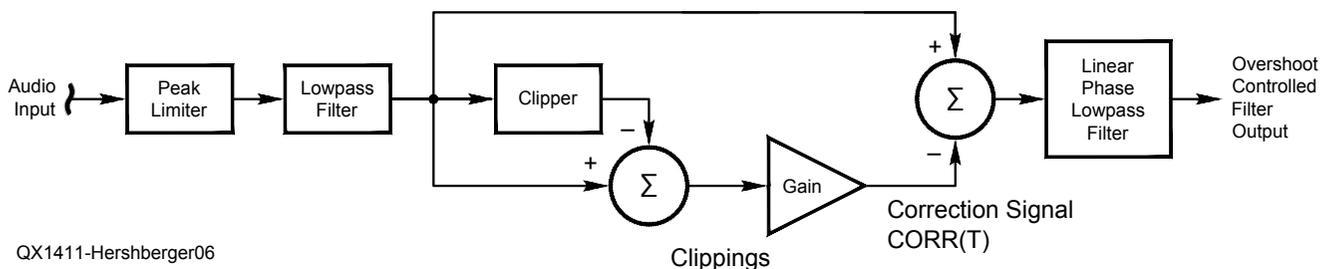


Figure 6 — Overshoot compensation system for real signals.

Figure 7 — First three Fourier components of a 100 Hz square wave. Notice that zero crossings are coincident at 0, 0.005, and 0.01 seconds.

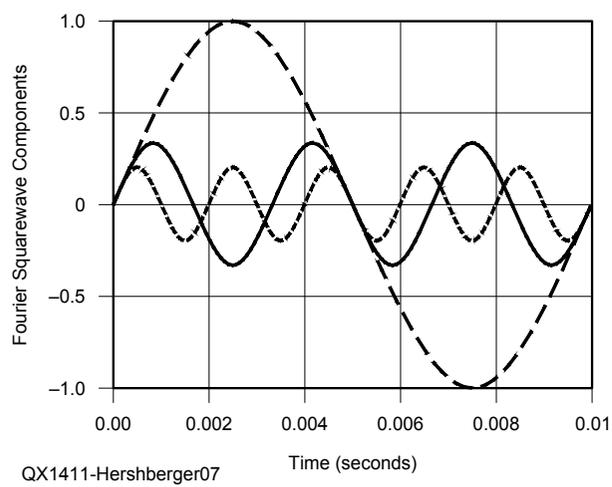


Figure 8 — First three Fourier components of a 100 Hz square wave, all shifted 90° (Hilbert transformed). Notice that waveform peaks, rather than zero crossings, are coincident at 0, 0.005, and 0.1 seconds.

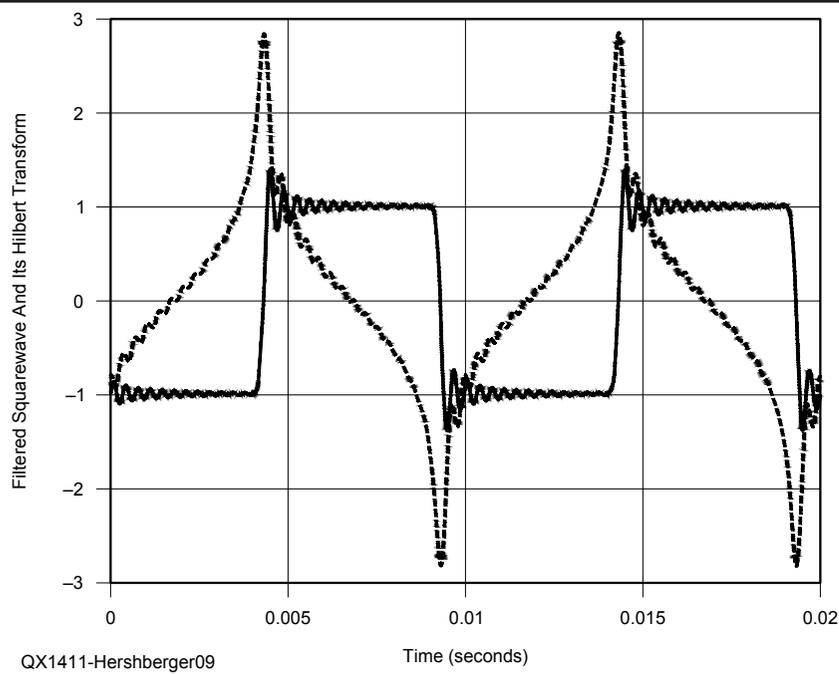
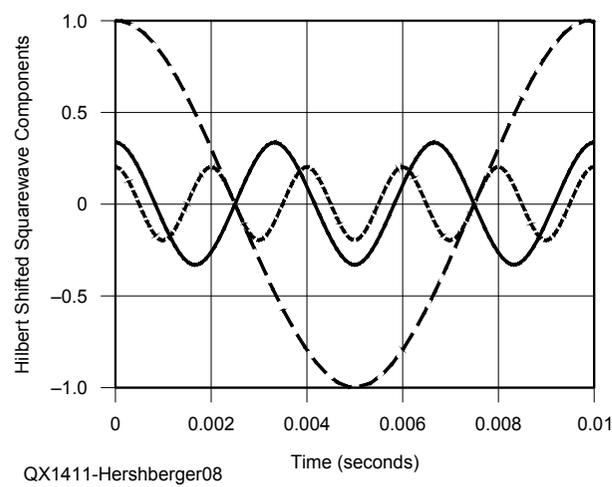


Figure 9 — Filtered square wave and its Hilbert transform.

hardware but exact matching is trivial in DSP. If we try to control overshoot independently in the two paths, however, the system will not work. A different approach is required for complex signals.

Overshoot Control for Complex Signals

Now let us extend the overshoot correction system for real signals to complex signals, so that we can generate SSB without overshoots.

What is a complex signal? It is a signal with an in-phase (I) part and a quadrature (Q) part, which can also be treated as real and imaginary parts. Physically, there are two signal paths. Complex signals are also known as analytic signals.

The phasing method for generating single sideband is an example of creating a complex signal in physical hardware. We take a modulating signal (speech audio) and shift all of its frequency components by 90° (a Hilbert transform). The original audio is applied to an RF mixer operating at one phase, and the phase shifted audio is applied to another RF mixer operating with a 90° RF phase shift. When the two double sideband signals are added, one sideband cancels and the other adds. We have single sideband.

$$ssb(t) = m(t) \cos(\omega_c t) + H[m(t)] \sin(\omega_c t) \quad [Eq 2]$$

Where:
 $ssb(t)$ is our single sideband signal
 $m(t)$ is the modulating function
 $H(\)$ indicates the Hilbert transform
 ω_c is the radian carrier frequency, and
 t is time.

In modern DSP based modulators for many signal types, including 8 level vestigial sideband digital television, orthogonal frequency division multiplexing (OFDM) and other types, the signals are generated at baseband and then up converted to the operating frequency in the later stages. Overshoot compensation of SSB signals is also best done at baseband.

Regardless of how we generate SSB (filter method, phasing method, Hilbert transform method, or Weaver method), we always get a signal that has a Hilbert transform in it.¹ Choose any demodulation phase for an SSB signal, and if you shift that RF phase 90°, the demodulated audio phase will also shift 90° for all frequencies. That audio phase shift constitutes a Hilbert transform.

A square wave causes large peaks when Hilbert transformed. A square wave is a summation of a fundamental sine wave plus odd harmonic sine waves, each diminished by the harmonic order (third harmonic at

¹Notes appear on page 13.

1/3, fifth harmonic at 1/5, and so on). All of the sine waves have their zero crossings coincident. This is shown in Figure 7.

Now if we shift all of the audio frequencies by 90°, those sine waves turn into cosines. While the sine wave components of the square wave share zero crossings, the cosine components are all coincident at their peak amplitudes! This is shown in Figure 8.

When we apply a Hilbert transform to baseband audio, we can get some really nasty waveforms. Beginning with our low-pass filtered 100 Hz square wave example, Figure 9 shows what the Hilbert transform looks like. The Hilbert transform overshoots are much more severe than the peaks of the filtered square wave. The envelope of an SSB signal generated from this Hilbert transform pair will spend most of its time overshooting. This is why the Hilbert transform of a square wave is so spiky looking. And for that reason, clipped speech applied to an SSB modulator will also produce some large spikes in the RF envelope.

In Figure 10, 1.0 is “full modulation,” corresponding to an input signal amplitude of 1.0. The envelope is much larger, creating some nasty peaks. In fact, the envelope rarely dips below 1.0 in Figure 10!

Again, this envelope (or something even worse) will be produced regardless of the manner of SSB generation (filter, phasing, Hilbert, or Weaver), because I and Q are always related by a Hilbert transform in an SSB signal.

For a complex signal, what we are

ultimately trying to control is not the amplitude of I and Q individually. We are trying to control the envelope of the resulting signal, which is the modulus, or

$$\sqrt{I^2(t) + Q^2(t)}$$

To convert the real overshoot control system to complex, we must form the envelope signal and use it to control the gain of both I and Q. For this conversion, we should note that a clipping function can also be done by a divider. If we divide a real

signal by the maximum of 1 or the absolute value of the real signal, the real signal will be clipped. It is this approach (division) that will extend the real technique to complex signals. Instead of dividing by the absolute value of the real signal, we will divide by the modulus of the complex signal.

The complex signal overshoot control system with a Weaver SSB modulator is shown in Figure 11. Most of the signal paths shown in this control system are complex. The wide lines represent complex signals and the narrow lines are real signals.

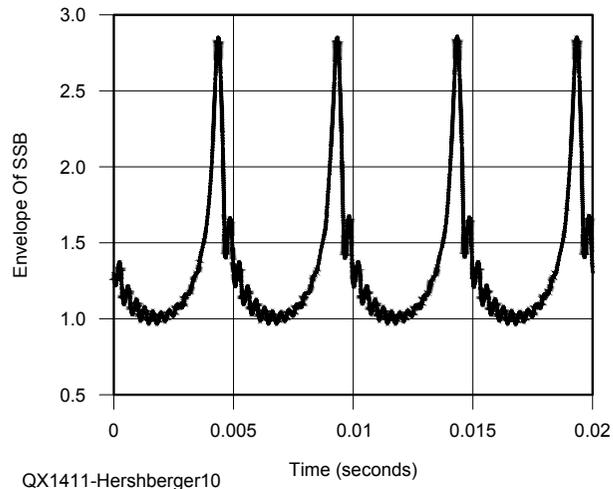


Figure 10 — Envelope of square wave SSB from I and Q signals of Figure 9.

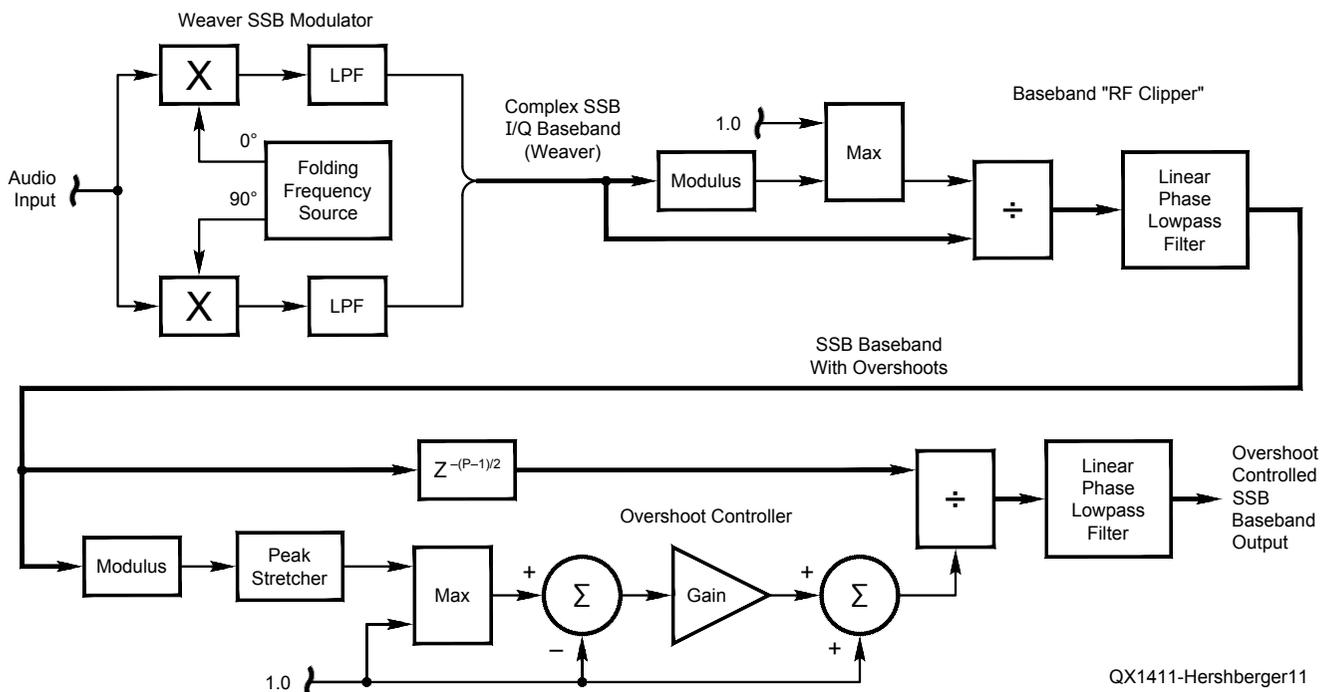


Figure 11 — Overshoot control for single sideband.

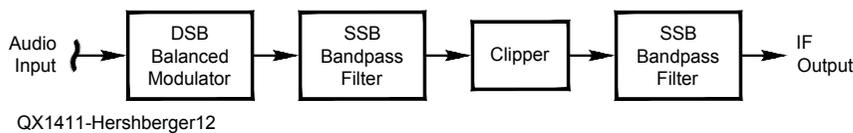


Figure 12 — Conventional analog RF clipper. RF cycle overshoot, spectral truncation, and filter group delay effects reduce its effectiveness.

Overshoot control for single sideband is simplest if we use the Weaver system (see the Generating SSB with Complex Math sidebar) because the filters operating on the real I and Q signals are all real (no complex coefficients). A further advantage of the Weaver system is that the bandwidth of the baseband I and Q signals is less than half what it would be with the Hilbert transform method. This allows use of a lower sampling rate. If the sampling rate is cut in half, the number of filter coefficients will also be cut in half for the same shape factor. This results in a Weaver SSB computation rate $\frac{1}{4}$ of what it would be with a Hilbert transform or phasing method SSB modulator.

Figure 11 shows a Weaver modulator for generation of SSB baseband signals from incoming peak-limited audio. The folding frequency for a 300 to 3000 Hz system would operate at 1650 Hz. The two low-pass filters have a bandwidth of 1350 Hz. The resulting baseband spectra are folded through DC.

The subsequent overshoot control system is divided into two subsystems. First there is a baseband “RF clipper.” It operates at audio baseband, and not at RF, but its operation is similar in function to a conventional “RF clipper.” See Figure 12. The purpose of this function is to remove most of the overshoot from the process of generating SSB. *This is not intended to be an audio processor.* Clipping threshold is set to 0 dB, which is referred to an audio tone at maximum input level.

“RF clipping” is interesting because it results in no harmonic distortion — only intermodulation distortion. As such it is usually preferable to simple audio baseband clipping, which produces both harmonic and intermodulation distortion. “RF clipping” of a single sine wave tone results in no distortion at all.

A conventional analog “RF clipper” generates single sideband at some nonzero frequency, clips the RF cycles, then band-pass filters the result to eliminate out of band distortion components. The result is an SSB signal with greater average power. A conventional RF clipper cannot control peaks very well, however. It controls instantaneous peak RF amplitude, which is not the same thing as peak RF *envelope* amplitude. As each RF cycle is clipped, the clipped sine

waves may be approximated as trapezoidal waves. RF harmonics are produced. In the limit (a large amount of clipping) the waveform will approach a square wave. When a square wave is filtered back to sine wave, the result will be overshoots. The peak amplitude of the fundamental component of a square wave is larger than the square wave peak by $4/\pi$. That would be 27% overshoot. Additional overshoot results from spectral truncation and group delay distortion in the second filter.

The baseband “RF clipper” shown in Figure 11 does not limit the instantaneous amplitude of RF cycles, because there are no RF cycles. Instead it limits the RF envelope. The RF envelope is the modulus of the complex time domain signal. There will be no $4/\pi$ overshoots. There will, however, still be some overshoot from spectrum truncation.

The modulus (envelope) signal is first generated. The larger of either the envelope or 1.0 is applied to the denominator input of a complex divider. The I and Q signals are divided by this inverse gain value. The modulus of the divider output will be perfectly limited to 1.0. The nonlinear processing will result in some out of band distortion products, however. The distortion is removed by a complex low-pass filter (which is really just two real low-pass filters for the Weaver system). The output of the low-pass filter will have some overshoots.

The overshoots from the RF clipper are then processed in the next subsystem shown in Figure 11. Again, the modulus is formed and the maximum of the modulus or 1.0 is generated.

The modulus value is next applied to a “peak stretcher.” This simply takes the maximum of the current sample and several previous samples. The number of samples will depend on the sampling rate. The recommended length of the peak stretch window is given by Equation 3.

$$Pwin \approx \frac{0.3}{BW} \quad [Eq\ 3]$$

Where:
 $Pwin$ is the length of the peak stretch window in seconds ($Pwin$ is approximately 111 μ s), and BW is the bandwidth of the SSB signal (2700 Hz in this case).

Equation 4 converts this to samples.

Generating SSB with Complex Math

Let’s look at the different ways to generate single sideband from a complex math perspective. There are three commonly used ways to generate SSB: the filter method, the Hilbert transform method (a variant of which is known as the “phasing method”) and the Weaver method (also known as the “third method”). In analog circuitry the filter method is most common. In DSP, the Hilbert and Weaver methods are more common. In DSP, obtaining the two perfectly matched signal paths is trivial for Hilbert and Weaver, but generating a high Q band-pass filter for the filter method at high sampling rates is difficult.

First we look at the filter method. The filter method uses no complex math at all. All of the calculations are real. Nevertheless, we will show the mirrored positive and negative spectral components as part of Figure A.

The top line of Figure A shows the spectrum of the baseband modulating signal. Next is the carrier signal. The third spectrum shows the output of a balanced modulator, which multiplies the audio and the carrier. We see a double sideband signal at positive and negative frequencies. The next step is to apply a band-pass filter that passes the upper sideband and rejects the lower sideband. The final spectrum is our real SSB signal.

Next we will generate SSB with the Hilbert transform method. This will get us into some complex math. Figure B shows the progression of the generation of the signal.

The top spectrum is the modulating audio. The next spectrum is the response of a filter derived from a Hilbert transform. H is a Hilbert transform multiplied by $-j$, plus a unit real impulse response with a time delay of one half of the Hilbert transform delay. This sum of a Hilbert transform in the imaginary coefficients plus the delayed real unit impulse will suppress the negative frequencies, but allow the positive frequencies to remain.

The next spectrum shows the result of filtering $m(t)$ with G . Only positive baseband frequencies remain. The next spectrum is a complex carrier sinusoid, at a positive frequency only. The next spectrum shows what happens when the complex positive frequency audio is multiplied by the positive carrier. Now we have complex SSB, at a positive frequency only. Finally, we discard the imaginary part of the signal. The result is a real SSB signal with mirrored positive and negative frequency components.

A variation on the Hilbert transform

method uses phase difference networks. Analog phasing method SSB exciter used phase difference networks rather than Hilbert transforms. A set of phase difference networks maintains a 90° audio phase shift between its outputs, but it incurs additional phase distortion beyond that of a Hilbert transform system. Phase difference networks are easier to implement in analog circuitry. In DSP, fewer calculations are required for phase difference networks if an infinite impulse response (IIR) implementation is used.⁵

Finally we will generate SSB using the Weaver method. Figure C shows the spectrum at each step in the process of generating the signal.

As before, the top spectrum is the modulating baseband audio. The next spectrum is a complex frequency-shifting or “folding” carrier, selected to be in the middle of the audio passband. The third spectrum shows what happens when the baseband audio is multiplied by the shifting carrier. The mirrored spectrum is shifted downwards. The next spectrum shows the response of a real low-pass filter. It is symmetrical around DC. In actual implementation, it consists of two identical real filters, one filtering the real part and the other filtering the imaginary part. After this filter is applied to the shifted baseband audio, what remains is the single audio spectrum folded through DC. Next we can multiply that complex folded audio spectrum by a complex carrier. The result is a complex (positive frequency only) single sideband signal. Finally, we discard the imaginary part, leaving a real single sideband signal with mirrored positive and negative spectral components.

In this article, we assume that the Weaver method is being used to produce single sideband. Since the overshoot correction process requires nonlinear processing followed by additional filtering, that filtering will be simplest to implement with real filters (no complex coefficients) of the type used in the Weaver method. Overshoot control can be done with the Hilbert method, but the filters must all be complex instead of real. Alternatively, since the Hilbert or phasing method baseband signals are already in analytic form, they may be frequency shifted downward with a complex multiplication by a complex tone to put them in spectrally folded Weaver format. After the overshoot correction is completed (using real filters) the baseband signals can be shifted back to Hilbert format by multiplying by the conjugate of the same complex tone. Note that with analytic signals, no filtering is required because there are no unwanted spectral images.

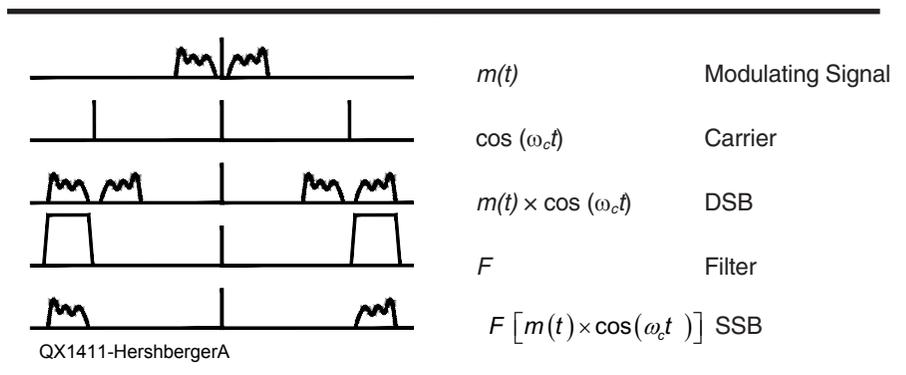


Figure A

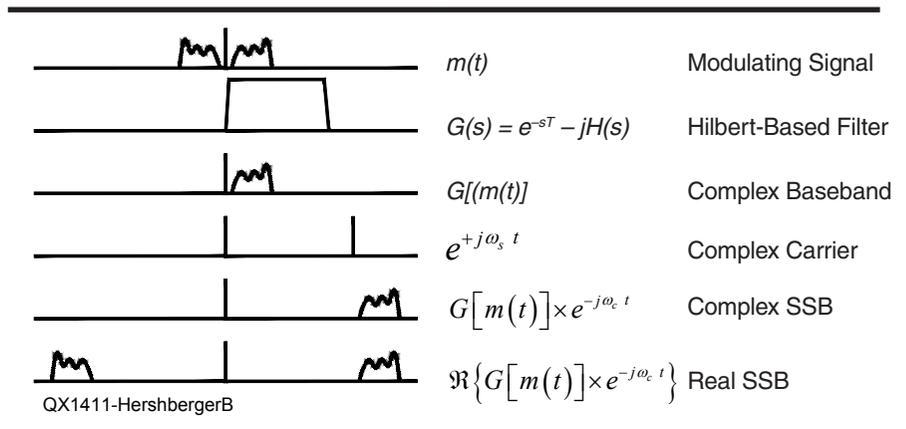


Figure B

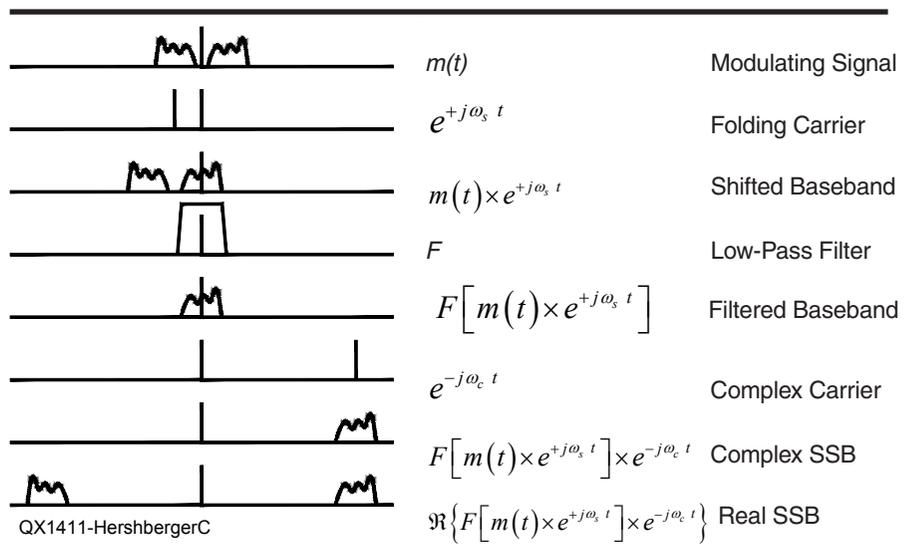


Figure C

$$Pn = \max\left[3, \text{round}(Pwin \cdot Fs)\right] \quad [\text{Eq } 4]$$

Where:

Pn is the number of samples

$\max()$ selects the largest value from its arguments

$\text{round}()$ rounds to the closest integer, and Fs is the sampling rate in Hz.

Pn should be an odd number and at least 3. For example, for a 2.7 kHz bandwidth and a 48 kHz sampling rate, the closest integer value for Pn would be 5. Note that there is a compensating delay in the signal path such that the peak is stretched equally to both preceding and subsequent samples. The compensating delay for $p = 5$ would be 2 samples. Odd values of Pn are suggested so that the compensating delay will have an integer number of samples.

The peak stretcher may be omitted. Without it the overshoot will be about 5% on voice peaks. With it, the overshoot is reduced to about 1.5%.

The output from the peak stretcher has unity subtracted from it to produce output only when there is overshoot. The resulting signal is applied to a gain factor of approximately 2.0. With low sampling rates and a peak stretch of 3 samples, the optimum gain factor could be less than 2.0. After 1 is added to return the quiescent gain to unity, it is applied to the denominator input of a complex divider. The divider will perform a “more than clipping” function similar to what was done in the real overshoot compensation system. The difference is that the correction is applied in a divider rather than additively, and the gain control affects the real and imaginary components equally.

Finally, the overshoot compensated complex baseband signal is low-pass filtered to produce CESSB. If the final filter is made slightly wider in bandwidth than the incoming signal, there will be less overshoot. The Weaver filters used for SSB generation

had a cutoff frequency of 1350 Hz. The final filter has a bandwidth of 1450 Hz. The same 1350 Hz filter could have been used with slightly increased overshoot (a small fraction of 1%).

Note that the entire system of Figure 11 has no effect on signals below 100% modulation. Applying a steady state tone below 100% will not invoke operation of any of the clippers. The system only affects transients that would produce overshoot. Audio compression to increase density, compression, peak limiting, and so on should be applied before the audio input to Figure 11.

DSP Economy

It is possible to reduce the amount of DSP horsepower required to control overshoot by implementing the filtering method shown in Figure 13. The method is shown for the baseband “RF clipper” but it may also be applied to the overshoot correction processing block as well.

In Figure 13, the input to the divider is subtracted from its output. The result is the nonlinear correction signal that has been added to reduce the envelope peaks. Rather than pass the entire signal through the low-pass filter, only the correction signal is passed through the filter. Since the correction signal has low energy, the specifications of the linear phase low-pass filter may be relaxed in both the passband and the stopband. With fewer filter coefficients than would otherwise be required, the overall calculation rate is reduced. Notice that the main signal path only incurs delay and no filtering. If the low-pass filter is a finite impulse response (FIR) filter, it should have an odd number of coefficients so that the delay line may have an integer number of sample delays.

Sampling rates should be mentioned at this point. Generating the initial SSB signal requires only that the sampling rate should be higher than twice the bandwidth of the baseband signals. The Weaver SSB

generation method has a big advantage here since a typical 2.7 kHz wide SSB signal only requires a 1350 Hz baseband bandwidth. (The Weaver baseband spectrum is folded through DC.) The sampling rate could be as low as 3 to 4 kHz. The “RF clipper” and the overshoot compensator both generate wider bandwidth correction signals, however. To avoid aliasing, the sampling rate at that point should be higher. The simulations for this article used a sampling rate of 48 kHz for all processing, which is overkill. Inspecting the spectra of the nonlinearly processed Weaver baseband signals show that the bandwidth grows to only about 5 kHz at the -60 dB point. So the sampling rate for the nonlinear processing could be as low as 10 to 12 kHz, provided that the Weaver method is used.

In practice, input audio filtering and Weaver SSB generation could be done at a sampling rate of 8 kHz. Then a half band interpolator could raise the sampling rate to 16 kHz where all subsequent nonlinear processing could be done.

Results

A test audio file was prepared using SOX, an open-source command line audio processing program.² SOX is an acronym for “SOund eXchange.” The SOX code used applies filtering, fast gain control, and peak limiting to a test audio file of my speech. The processing is aggressive, producing 100% peaks frequently.

Simulation was done using GNU *Octave*, an open-source matrix math program.³ *Octave* reads in the processed audio file, and appends to it a 1 kHz tone at 100% modulation as an amplitude reference. (Here “100%” corresponds to the maximum PEP capability of the downstream RF power amplifier.) SSB is generated using the Weaver method. The envelope of the SSB generator is shown in Figure 14.

The envelope is normalized to 1.0, which corresponds to the peak of the audio signal.

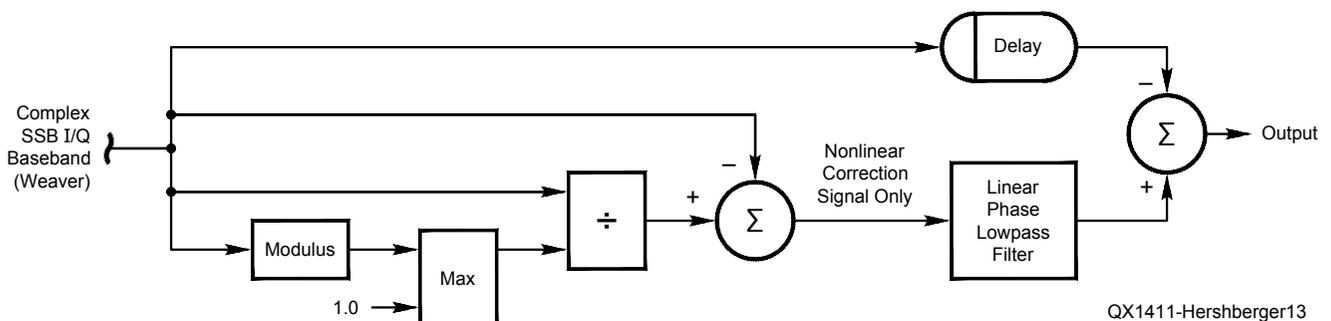


Figure 13 — Method for reducing calculation rate.

Overshoot is 59%. The tone burst at the end has overshoots at its beginning and end, but the steady state portion is at the nominal 1.0 level.

Figure 15 shows the envelope after the baseband “RF clipper” stage. Because of spectral truncation (filtering), the RF envelope is still not controlled. It is better than the conventional SSB generator, but there is still 21% peak overshoot with this audio sample. (A conventional analog RF clipper such as the one shown in Figure 12 would have even worse overshoot, because of the $4/\pi$ problem and nonlinear phase filters.)

Figure 16 shows the final output from the system. The peaks are well controlled. Overshoot is only 1.6%. The overshoots at the beginning and end of the tone burst are gone too.

Several supporting files have been placed on the ARRL *QEX* files website.⁴ Included are:

1) The original unprocessed audio WAV file ([SSB-test-wideband.wav](#)).

2) The SOX command script that filters the input audio to SSB bandwidth, then compresses, and limits the audio (Sox-ssb-process.bat).

3) The FIR filter that restricts the audio bandwidth to 300 to 3000 Hz (SSBBPF.TXT). This filter is used by SOX.

4) The WAV file of the SOX-processed audio ([SSBaudioprocessed.wav](#)). This file was used to create Figures 14, 15, and 16.

5) GNU *Octave* source code which generates SSB, “RF clipped” SSB, and overshoot compensated SSB (cessb.m). A version compatible with *Matlab*TM is also included (cessb_matlab.m).

6. The demodulated audio output WAV file from the Weaver SSB generator signal of Figure 14 ([ssbmodem.wav](#)).

7. The demodulated audio output WAV file from Figure 16 ([cessbmodem.wav](#)).

The last two audio files can be compared to hear how using SSB peak control has very little effect on audio quality.

Implementation

Implementation can be done efficiently with a DSP chip and/or FPGA code. This method of SSB generation is something best integrated into a transmitter or transceiver, as opposed to an external add-on box. An add-on box would have to generate SSB at the radio’s intermediate frequency and inject it in place of the radio’s original SSB modulator.

System setup is best done with a single reference sine wave tone at 100% modulation. The clippers should be set to operate on anything that exceeds the envelope of the reference tone. If the incoming signal is low

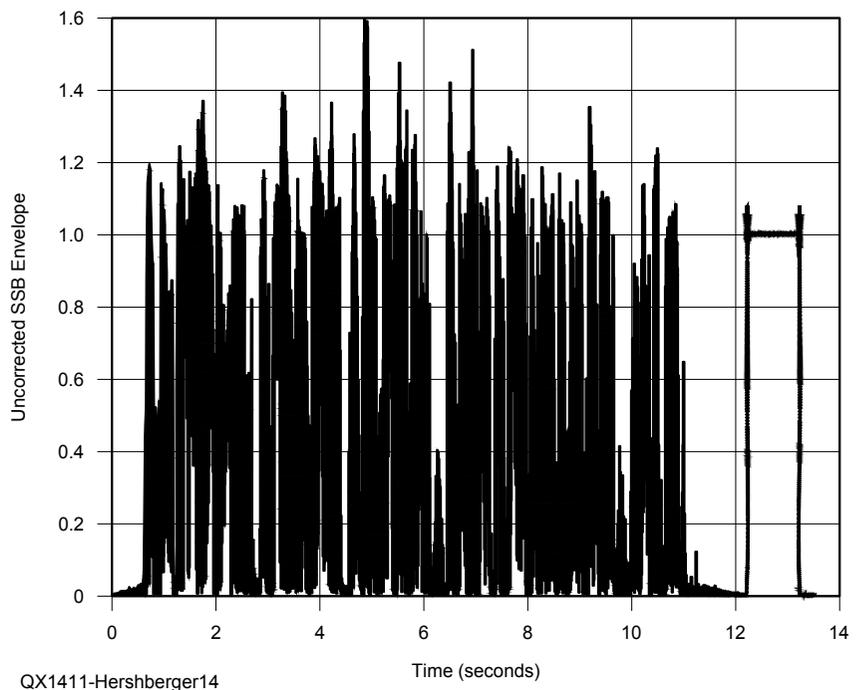


Figure 14 — Envelope of conventional SSB generator. Overshoot is 59%.

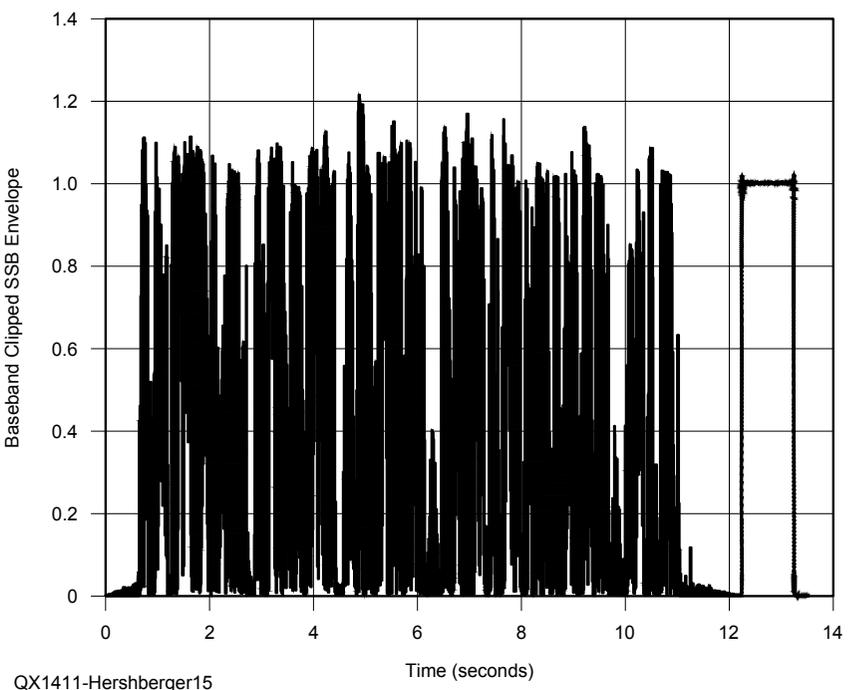


Figure 15 — Envelope of conventional SSB Generator with Baseband RF Clipping. Overshoot is 21%.

enough to produce no overshoots, the system does nothing to the SSB baseband signal. Only when the envelope exceeds that of the 100% tone will there be overshoot control processing.

It is possible to overdrive the algorithm by driving the baseband “RF clipper” into a few dB of clipping beyond what is necessary to remove the Hilbert transform overshoots.

This will provide a further increase in average power. I think it is better to do baseband audio processing, however, and let the CESSB system only remove overshoots. Historically, baseband audio processing was not considered particularly effective for increasing average SSB power — because of envelope overshoots. This is no longer true, with the introduction of CESSB. With

baseband audio processing, sophisticated multiband compression and clipping is possible, with better results than a simple single-band RF clipper.

Application Results

FlexRadio made the first hardware implementation of CESSB and reports a 2.56 dB increase in average power over fast look-ahead ALC. Figure 17 shows average power as a function of time for CESSB versus ALC.

CESSB is intended for use with speech signals. Although it is a nonlinear process, the nonlinearity has a negligible effect upon speech. For non-speech SSB applications, however, such as digital modes (PSK, JT65, SSTV, and others), CESSB should be tested to be sure that any average power increase is not offset by nonlinear distortion.

Conclusions

SSB generation generally results in large envelope peaks well above the reference level set by a steady-state tone, even when the input audio is accurately peak limited. SSB envelope overshoot is caused by spectrum truncation and nonlinear phase shifts (particularly from the Hilbert transform). Use of “RF clipping” reduces, but does not eliminate the overshoots. With overshoot compensation, about 3.8 dB of peak reduction is possible using this test speech sample. If power is set to keep peak envelope power the same, this results in an average power increase of about 140%!

Comparisons to a well-designed fast look-ahead ALC system have resulted in an average power increase of about 80% —

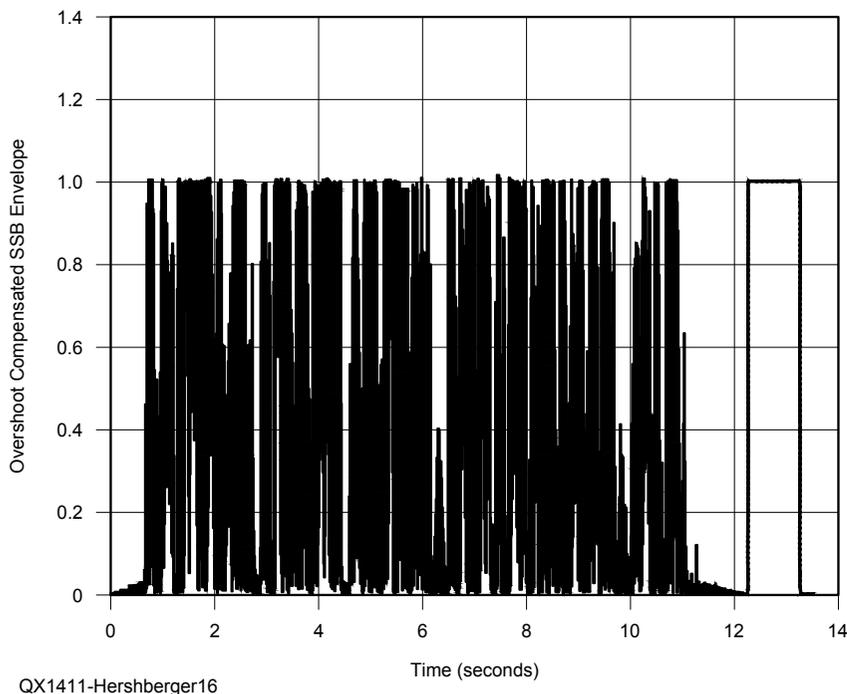


Figure 16 — Envelope of Controlled Envelope SSB Generator. Overshoot is 1.6%

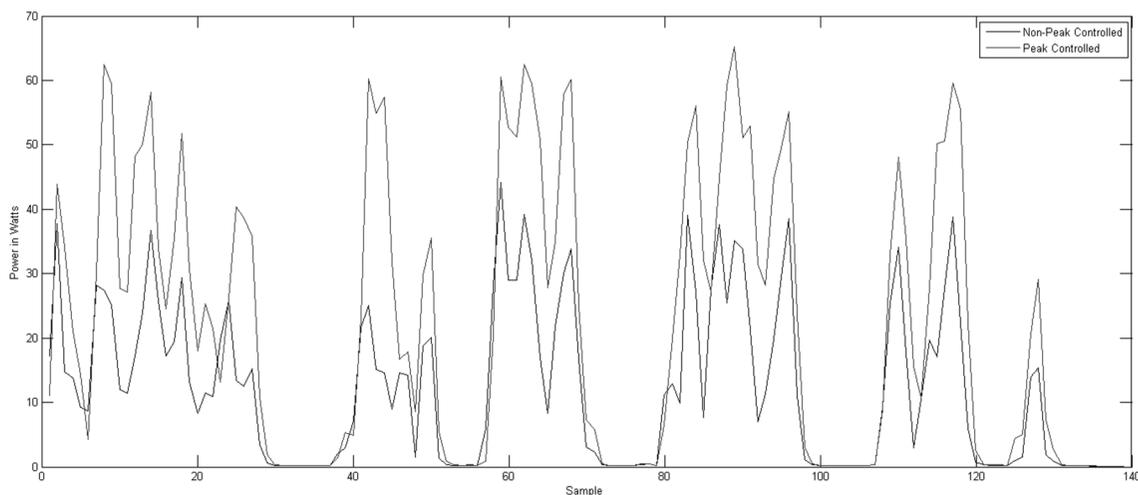


Figure 17 — Average SSB power with fast look-ahead ALC (lower trace) and CESSB (upper trace). (Graphic supplied by FlexRadio)

about 2.56 dB. (That is about the same as converting a single beam antenna to a stacked array.) Your wattmeter should read the same PEP as before, but the average power will be higher.

The processing used does not produce any significant audible artifacts. It may be used in conjunction with speech processing.

The final step in the Controlled Envelope SSB process is a sharp band-pass filter — guaranteeing bandwidth limiting in addition to accurate amplitude control.

In summary, Controlled Envelope SSB envelope control is accomplished by:

1) Prefiltering and peak limiting the audio input signal.

2) Baseband “RF clipping” of the SSB signal to reduce Hilbert transform overshoots.

3) Overshoot compensating the remaining envelope peaks resulting from baseband “RF clipping.”

By accurately controlling SSB envelope peaks at the point where the SSB is generated, ALC is unnecessary. ALC, even with look-ahead, reduces transmitted power when it does not have to — before and after an envelope peak. Average transmitted

power can be significantly higher, without introducing “speech processor” type artifacts. An SSB signal with well-controlled envelope peaks makes more efficient use of the RF power amplifier, and produces higher average power for a given peak envelope power.

Acknowledgment

I would like to thank FlexRadio for reducing CESSB to practice, and for verifying its usefulness both with lab tests and real on-air testing.

Notes

¹Ward Silver, NØAX, Editor, *The 2014 ARRL Handbook*, ARRL, 2013, Chapter 13, Section 1.4, Figure 16.13B, p 13.6. ISBN: 978-1-06259-001-7; ARRL Publication Order No. 0007, \$49.95. ARRL publications are available from your local ARRL dealer or from the ARRL Bookstore. Telephone toll free in the US: 888-277-5289, or call 860-594-0355, fax 860-594-0303; www.arrl.org/shop; pubsales@arrl.org.

²You can find more information about SOX and download a copy of the SOX software at: sox.sourceforge.net/.

³There is more information about GNU *Octave* on the *Octave* home page at www.gnu.org/software/octave. You can also download the latest version of GNU *Octave* from that website.

⁴The supporting files that accompany this article are available for download from the ARRL QEX files website. Go to www.arrl.org/qexfiles and look for the file [11x14_Hershberger.zip](#).

⁵Theodore A Prosch, DL8PT, “A Minimalist Approximation of the Hilbert Transform,” *QEX*, Sep/Oct 2012, pp 25 – 31.

Dave Hershberger, W9GR, was first licensed in 1965 at age 14 as WN9QCH. He is an ARRL Life Member. 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 megawatt VLF transmitter system.