By Dave Hershberger, W9GR

DSP—An Intuitive Approach

Learn how those new-generation digital signal processors can make your ham life easier!

Today's digital signal processors, such as my own DSP-3,[1 2] do a lot to enhance ham-radio operation They remove both man-made and natural noise, and make operation more pleasant during difficult band conditions. DSP-based CW filters add razor-sharp selectivity to receivers that lack it, and they provide extra-narrow filtering (100 Hz or less) that is not generally available in analog IF filters. Custom band-pass filters are available for special modes such as SSTV and data communications. A few new functions keep things interesting: DTMF and CTCSS tone decoding.

The real strength of DSPs is their adaptive filtering capability, which makes auto- ic notch filtering and noise reduction a ity. DSPs frequently make it much more enjoyable to work stations with weak signals Sometimes, when atmospheric noise disappears, DSP can seem like pure magic!

Well, it's not really magic. It's all done with digital technology, which is now inexpensive enough to bring some advanced signal-processing techniques to the radio amateur.

I promise not to use any equations or mathematics in this article. Equations and math are fine when you need to quantify or accurately model something, but I want to give you an intuitive understanding of what goes on in these wonderful DSP boxes. I will explain DSP with words, pictures and a little metaphorical hand waving.

DSP Chips

DSP is usually performed by specially developed microprocessors (DSP chips). These differ from general-purpose microprocessors in several ways: Most importantly, DSP chips can perform certain mathematical operations *very* quickly. They have *multipliers* that form the product of two 16-bit or larger numbers in *one* instruction cycle. (Even those general-purpose microprocessors that have "multiply" instructions take *many* instruction cycles to perform a multiplication.) DSP usually involves a lot of multiplication, which makes a fast multiplier necessary. DSP chips also quickly perform the repetitive *multiply, accumulate, and data shift* instruction sequence, which is common in DSP.

Although DSP technique has existed for several decades, *inexpensive* DSP has become a reality only recently. About a decade ago, first-generation DSP chips cost $200 each. Now, powerful DSP chips are available in the $5 to $20 range, even in small quantities This makes them ideal candidates for low-cost amateur applications. Thanks to low-cost DSP chips, digital approaches to signal filtering and processing are now more than competitive with traditional analog methods.

DSP has several advantages over analog circuits. DSP units need no alignment or adjustments. They are free from drift or performance change with component aging and temperature variations and offer perfectly repeatable and predictable results. DSP allows easy implementation of complex signal-processing structures where analog component tolerances would be critical.

DSP also allows a filter to *change* itself, under its own control, in response to different conditions. This is the basis of adaptive filtering.

Differences between Analog and Digital Processing

In analog signal processing, a varying voltage or current represents a sound or a radio signal. We process that voltage or current using resistors, capacitors, inductors, amplifiers, mixers, crystal filters and other components. These circuits enable us to perform various signal-processing tasks such as band-pass, low-pass, and high-pass filtering, modulation, demodulation, frequency conversion, etc

DSP first converts an analog waveform to a sequence of numbers called *samples*. From that point we perform similar signal-processing tasks, but our toolbox is different. We manipulate the numerical data with digital circuits such as multipliers, adders, barrel shifters, delays and other digital-logic operations Then we convert the resulting numbers back to an analog waveform, which may drive a loudspeaker or headphones.

The digital domain offers many new operations that are difficult or impractical to implement with analog circuitry Modern DSP chips can easily perform signal delays, multiplication, nonlinear processing, phase-linear filtering, time-varying, signal-dependent or adaptive filtering and more.

The DSP toolbox offers capabilities, economies and limitations Don't expect to do straightforward receiver filtering at a high IF, such as 9 MHz, anytime soon. In addition to requiring a lot of DSP horsepower, it would also require a wide dynamic range analog-to-digital converter (with at least 16 bits of resolution). Such parts are not practical—at least not yet

Therefore, DSP applications that exist today operate at AF or very low IFs (25 to 50 kHz). Even though we are restricted to mostly low frequencies, a mathematical principle called the *frequency shifting theorem* tells us that anything we do to achieve *audio* selectivity in a linear receiver (CW and SSB receivers are linear; FM is not) is equivalent to filtering or signal processing at RF or IF. Although A/D converters limit DSP dynamic range and AGC derived before the DSP may pump from signals removed by the filter, it is basically true that audio selectivity is

[1]Notes appear on page 42

equivalent to RF or IF selectivity for SSB and CW reception

DSP CW Filters

Consequently, we can make DSP audio filters that are much sharper and narrower than the best crystal CW filters. We can design for phase linearity, too. This means that the time delay is equal across the passband, which reduces ringing and distortion of the CW envelope.

In DSP, we have no parts-tolerance problem. If a DSP instruction says, "multiply the present signal value by 23751," we know that the signal will be multiplied by exactly 23751, no more and no less. We can make highly selective filters without worries about accurate tuning, tracking and matching.

DSP easily performs time delays. We need only stuff signal samples into a memory, wait a while, and recall them. In analog circuits, a delay line requires a huge LC ladder structure with precise mutual coupling between the inductors.

A *transversal filter* can be designed for a phase-linear characteristic. This filter is very difficult to build with analog techniques, because it uses many delay lines or a single delay line with many taps. Because time delay and multiplication are easy in DSP, DSP transversal filters are common. In DSP, however, we call transversal filters *finite-impulse-response* (FIR) filters.

Figure 1 shows the frequency response and time delay curves of a typical (nonlinear phase) analog CW filter. This five-pole, elliptic, band-pass filter has a –3 dB bandwidth of 200 Hz and a –60 dB bandwidth of 350 Hz. Time delay is least at the center frequency, and it increases sharply as frequency approaches the passband edges.

Figure 1 also shows curves for a 100th-order FIR filter designed to the same specifications. The passband shape differs from that of the analog filter; but more importantly, the time delay is a *flat line*, constant with varying frequency.

The analog filter's phase distortion effectively delays the edges more than the bodies of the dots and dashes. Figure 2 shows keying envelopes filtered with the analog and phase-linear digital filters. Notice that the analog filter smears the keying transients by delaying them, while the digital filter distributes them equally before and after the edges and reduces their peak amplitudes (ringing).

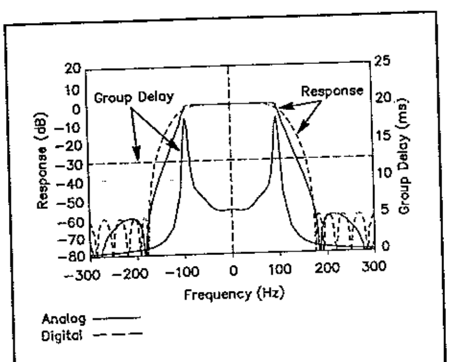

Figure 1—Response and group-delay curves for analog and digital (FIR) CW filters.

Ringing in CW Filters

There is a lot of misinformation on this subject. All band-pass filters ring; they *must* in order to work! There are two causes of ringing, however. The first is reduced bandwidth. All band-pass filters ring due to this mechanism. The second is nonlinear phase, also known as time delay distortion or phase distortion. Conventional analog filters incur this additional ringing due to phase distortion. Since a FIR filter is phase linear, there is no additional ringing from this mechanism. For a given bandwidth, a FIR filter rings less than a conventional analog filter, but it still rings. Many DSPs use FIR CW filters, so they present only minimal ringing.

FSK reception also benefits from phase-linear filters because phase distortion tends to smear adjacent bits so that they overlap in time.

Adaptive Filtering for Voice Signals

Perhaps the most wonderful benefit of DSP to Amateur Radio today is noise reduction. Many amateurs still have not heard a demonstration of DSP noise reduction. When they do, a common initial reaction is a dropped jaw and wide-open eyes.

DSP noise reduction removes most of the hiss from weak signals and makes such operation a lot more pleasant and less fatiguing. One ham told me that a few years ago he was calling CQ on one of the ham satellites. Even though he heard his signal loud and clear, nobody would answer him. Upon checking his equipment, he found that his DSP noise reducer was on. He switched it off and heard his signal fall into the noise! Apparently only he was listening to the satellite with DSP!

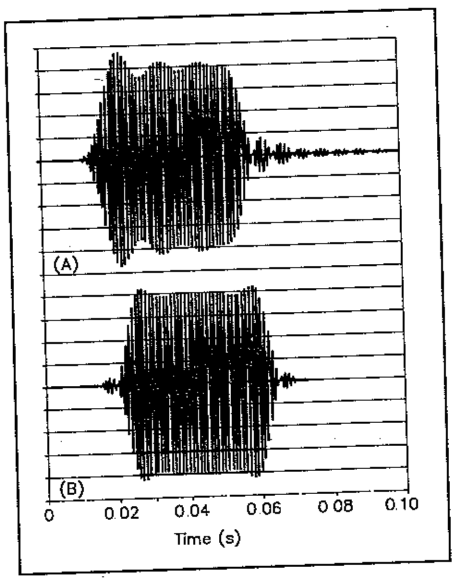

Figure 2—Analog filter phase distortion smears the edges of a CW pulse and increases ringing at A. At B, the result is much cleaner with a phase-linear (FIR) DSP filter.

Regardless of whether the noise reduction algorithm is the Widrow-Hoff Least Mean Square (LMS) algorithm,[3,4,5,6] or the Discrete Fourier Transform (DFT) based "spectral subtraction" algorithm,[7,8] DSP noise reduction works by finding the most significant spectral lines in a signal and then forming band-pass filters around the strongest energy concentrations. The particular algorithm only determines *how* this is accomplished. To understand adaptive noise reduction, it helps to think about filters in a different way.

Frequency-Domain Filtering

Conventional filters discriminate among signals on the basis of frequency. For example, a 2.4-kHz-wide crystal filter for SSB passes frequencies from 300 to 2700 Hz. It will reject an RF signal that produces a 4000-Hz tone or one in the opposite sideband because they are the wrong frequencies. Similarly, that filter would pass a 1000-Hz audio tone in the correct sideband. We might say that conventional filters operate in the *frequency domain*.

Autocorrelation Filtering

DSP noise-reduction filters and automatic notch filters operate by discriminating for or against a signal based on its degree of *autocorrelation*. We might call this kind of signal classification the *autocorrelation domain*.

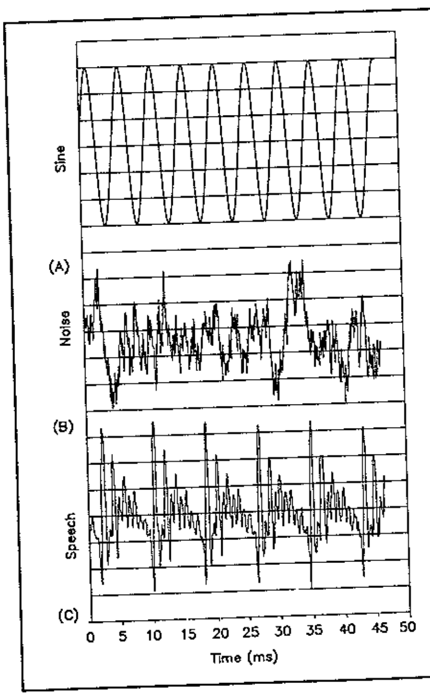

Figure 3—Time-domain displays of autocorrelation (see text) within common signals. Sine waves (A) are highly autocorrelated; noise (B) has no autocorrelation; speech (C) has partial autocorrelation.

What is "autocorrelation"? It has to do with the degree of repetitiveness of a signal Refer to Figure 3 as we discuss three important signal groups.

Consider a sine wave (Figure 3A) A sine wave has the highest possible degree of autocorrelation because every cycle in the sine wave is *exactly the same* as the next cycle. A DSP filter that rejects signals with the highest degree of autocorrelation is an automatic notch filter, because interfering carriers appear as pure sine-wave audio tones.

At the other extreme are thermal and atmospheric noise, which sound like hiss. Figure 3B also shows an example of band-limited noise. (That is, noise contained in some arbitrarily narrow band of frequencies.) There is no repetitive pattern in the noise A DSP filter that rejects signals with autocorrelation is a noise reducer

What about speech signals? Figure 3C shows the diphthong[9] in the word "nine" Note that this speech waveform has a certain degree of repetitiveness Each cycle is not exactly the same as its predecessors but there is a lot of similarity. This is generally true of voice signals

To summarize, pure sine waves have maximum autocorrelation; speech has some; background noise has none at all.

So a DSP filter that rejects signals with high autocorrelation and signals with zero autocorrelation is a simultaneous automatic notch and noise-reduction filter This filter will allow semiautocorrelated speech signals to pass, while rejecting both noise and audio tones!

How do we filter signals based on their autocorrelation? The LMS algorithm provides a way to do this.

Any signal can be represented as either a time-domain signal, which you can observe on an oscilloscope, or as a frequency domain signal, which you can observe on a spectrum analyzer.

More repetitive signals show more of their energy in discrete spectral lines (frequency domain). For example, a pure 400-Hz sine wave has all of its energy concentrated in a 400-Hz spectral line A nonrepetitive noise signal has its energy distributed throughout the spectrum, with no spectral energy concentrations. Speech is somewhere in between. It has a few strong spectral lines, and the rest of its energy is spread throughout the audio spectrum in weaker frequency components.

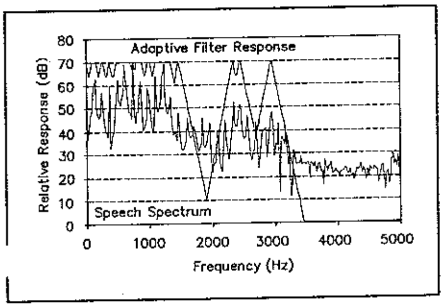

Figure 4—Adaptive filter response to a speech signal

Figure 4 shows the spectrum of the diphthong in Figure 3C. There are several strong spectral lines in the frequencies below 1500 Hz, and a few more in the 2500 to 3000-Hz range.

Figure 4 also shows a possible adaptive-filter response. We can derive this response from the voice spectrum by looking for the most significant spectral lines, and forming band-pass filters around them.

If the voice signal also contains noise, the adaptive filter will reduce the noise in the stop bands. In this case, noise around 2000 Hz and above 3000 Hz would be reduced.

Noise that is very close to voice signal components is not significantly reduced *by the DSP* But, due to a phenomenon known as "noise masking," another signal processor—the human brain—eliminates close-in noise components![10,11] Noise masking in the human auditory system makes the weaker of two closely spaced frequencies essentially inaudible

So the DSP attenuates noise that is not close to the voice-signal frequency components. The ear-brain combination removes the remaining noise, which is close to the strong spectral lines.

As the voice signal spectrum changes from syllable to syllable, the adaptive filter response will follow it, continuously changing its shape with the input signal's variations

Figure 5 represents common adaptive filters in the autocorrelation domain. The horizontal scale of this graph has noise to the far left, pure tones to the far right, and somewhat repetitive speech signals in the middle.

A DSP noise-reduction filter rejects signals with little or no autocorrelation, and passes everything else

If the adaptive filter forms notches instead of passbands around the strongest spectral lines, it's an automatic notch filter. It rejects highly repetitive signals and allows everything else to pass.

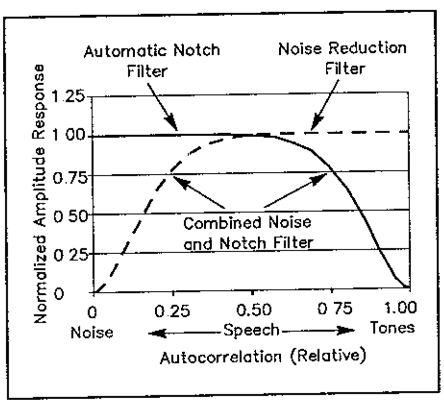

Figure 5—Response versus autocorrelation plots for noise-reduction and automatic notch filters. A speech filter applies both curves simultaneously to reduce white noise and carriers

An adaptive filter for speech provides simultaneous noise reduction and automatic notching. This filter rejects nonrepetitive noise and pure tones, but allows everything else (somewhat repetitive speech signals) to pass

Spectral Subtraction

Spectral subtraction[7,8] is another way to reduce the noise in voice signals. This technique accomplishes much the same thing as the LMS algorithm, but in a different way.

Up to this point, all of the DSP algorithms we have discussed work by processing a series of numbers that represent the signal waveform as a function of *time*. Spectral subtraction, on the other hand, works by processing a series of numbers that represent the *frequency content* of our input signal.

To do this, DSPs use a relatively complex mathematical operation (called a *transform*) to change the signal representation from the *time domain* to the *frequency domain*.

For example, what comes out of an analog to digital converter is a series of numbers which represent the audio voltage in *time* increments at 0 µs, 100 µs, 200 µs, etc. The transformation operation yields a series of numbers that indicate signal energy in *frequency* increments at 300 Hz, 320 Hz, 340 Hz, etc, up through 3000 Hz or more

A complementary inverse transform returns the frequency data to a time-domain signal. If we do the time-to-frequency transform and follow it immediately with the frequency-to-time (inverse) transform, we get our original signal back

Spectral subtraction is a three-step process:

1. Transform signal to frequency domain
2. Process frequency domain data
3. Inverse transform back to time domain.

This process repeats for successive short segments (a fraction of a second) of audio.

Spectral subtraction relies on two assumptions:

1. Voice-frequency energy is concentrated in a small number of frequencies, and
2. Noise energy is uniformly distributed throughout the audio spectrum

Spectral subtraction algorithms try to determine the "noise floor" of a signal. The process assumes that any frequency-domain value at or below the noise floor is noise and sets the energy at that frequency to zero. Conversely, it considers signals above the noise floor to be voice components and allows them to pass

Figure 6A shows the spectrum of Figure 4, but with some noise added. The spectral-subtraction threshold is a horizontal line at

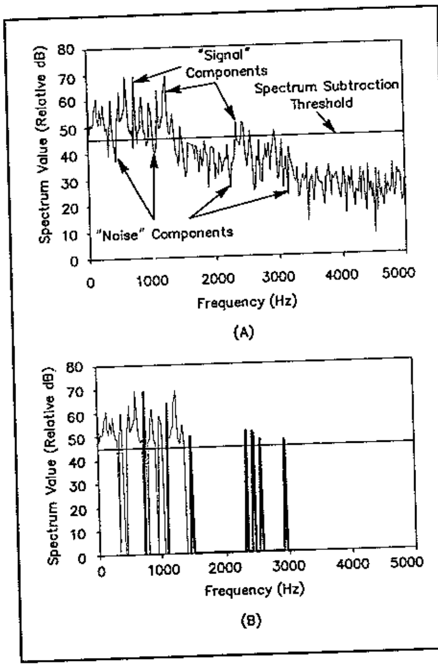

Figure 6—Spectral subtraction removes signal component frequencies with energy below some threshold (A). B shows the result

+45 dB. We consider anything below this line to be "noise," and anything above it to be "signal." When signals below the threshold are set to zero, the spectrum of Figure 6B results.

Some amateur DSPs use spectral subtraction for noise reduction, but the algorithm has several disadvantages:

1. It takes a substantial amount of time to perform the forward and inverse transforms. The resulting delay through the DSP (a fraction of a second) can create an annoying "electrical backlash" condition. The delay makes it difficult to rapidly tune receivers, because the audio and dial position are not synchronized. The delay also makes very rapid contest exchanges impossible.

2. Spurious audio "tones" result when processing noisy signals in some implementations. These appear as seemingly random beeps at random frequencies. They are caused by the algorithm's imperfect ability to distinguish between signal and noise in the frequency domain.[12]

3. Spectral subtraction requires much more DSP computing power than the LMS algorithm.

All algorithms have their disadvantages; yet under some conditions, spectral subtraction may provide the best performance. The JPS NIR-12 DSP, for example, provides *both* spectral subtraction and LMS noise reduction (which JPS refers to as *adaptive peaking*). The user can choose the best noise reduction method for each situation.

To summarize, spectral subtraction allows concentrated spectral energy (speech) to pass, and rejects low-amplitude signals over a wide range of frequencies. Spectral subtraction does this by transforming the signal to the frequency domain, throwing away the smaller spectral-energy values, then performing the inverse transform to the time domain.

Different Kinds of Noise

What kind of noise can DSP remove best? DSP is most effective with noise that has a basically constant level (eg, hiss) rather than spiky noise (the kind conventional noise blankers remove). DSP noise reduction algorithms work best with moderate to weak noise; conventional noise blankers work best when noise is very strong. If the noise comes in bursts that remain at least 6 dB below the instantaneous signal strength, a DSP filter will probably help. When line noise is strong, a conventional noise blanker works better than DSP because it has access to the wideband RF signal prior to IF filtering.

Line noise can mean a lot of things. If the noise is strong 120-Hz impulses, DSP will not help much, but a conventional noise blanker will. When the noise is closer to white noise or a background "frying" noise (as opposed to a buzz), DSP will provide some relief.

Denoisers or Fixed-Frequency CW DSP Filters?

Although some amateurs use adaptive noise-reduction filters for CW work, nonadaptive CW filters are better in most applications. Remember that adaptive filters work by automatically forming little band-pass filters around the most significant spectral lines in the signal. For CW operation, the filter must create the band-pass filters with each dot and dash. Also, we have advance knowledge of the desired signal's spectrum, so we can design a fixed-frequency filter that is optimal and need not adapt itself. Furthermore, DSP fixed-frequency CW filters are phase linear, while adaptive filters formed by noise-reduction algorithms usually are not phase linear.

It makes sense to use adaptive noise-reduction filters for CW when tuning across the band; they provide a wide (SSB) bandwidth that passes more signals. When you want to hear only one signal, it is probably best to use dedicated CW filters.

Conclusion

Today, DSP provides a lot of value at low cost by reducing noise through advanced adaptive filtering algorithms. It also supplements IF selectivity with additional filters for a variety of modes. Additional functions—such as tone decoders—can be added where software space permits. The flexibility of general-purpose DSP chips allows designers to include a lot of functions in a single box.

As DSP hardware costs continue to fall and algorithm (software) development improves, DSPs will perform more and more functions in our radios. Multiple-bandwidth IF filtering, SSB generation, speech processing, passband tuning, FSK and FM generation and detection, fine tuning and bandwidth compression are all feasible as DSP functions.

Acknowledgments

Thanks and credit go to Dr Steven Reyer, WA9VNJ, who reviewed this article for gross oversimplifications and inaccuracies; Curt Holsopple, K9CH, for expunging incomprehensible "technobabble"; and to my wife Sandy Hershberger, N6SMF for her support and insight.

Notes

[1]Glenn Swanson, KB1GW, 'Product Review: Quantics W9GR DSP-3 Audio Filter," August 1995 *QST*, pp 73 to 75.

[2]Peter Bertini, K1ZJH, "Quarterly Review," Summer 1995 *Communications Quarterly* pp 85 to 90.

[3]Steven Reyer, WA9VNJ, and Dave Hershberger, W9GR, "Using the LMS Algorithm for QRM and QRN Reduction," September 1992 *QEX*, pp 3 to 8.

[4]Bernard Widrow *et al*, 'Adaptive Noise Cancelling: Principles and Applications,' *Proceedings of the IEEE*, Vol. 63, No. 12, pp 1692 to 1716, December 1975.

[5]Dave Hershberger, W9GR, 'Low Cost Digital Signal Processing for the Radio Amateur," September 1992 *QST*, pp 43 to 51.

[6]Dave Hershberger, W9GR, "A Few Words about DSP,' Summer 1995 *Communications Quarterly*, pp 80 to 84.

[7]S. Boll, "A Spectral Subtraction Algorithm for Suppression of Acoustic Noise in Speech,' *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-29, pp 113 to 120, April 1979.

[8]D. Hall, KF4KL, 'Spectral Subtraction Eliminates Noise from Speech in Real Time," May 1995 *Personal Engineering*, pp 51 to 54.

[9]A dipthong is a gliding, monosyllabic speech sound.

[10]D. M. Green, *An Introduction to Hearing*, Lawrence Erlbaum Associates, Inc. Hillsdale, NJ, 1976.

[11]J. J. Zwislocki, 'Masking: Experimental and Theoretical Aspects of Simultaneous, Forward, Backward, and Central Masking,' in *The Handbook of Perception*, Volume IV, Academic, New York, 1978.

[12]R. Preuss, "A Frequency Domain Noise Cancelling Preprocessor for Narrowband Speech Communications Systems," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Volume ASSP-29, pp 212 to 215, April 1979.

Dave Hershberger, W9GR, PO Box 2163, Nevada City, CA 95959. Dave was first licensed in 1965 at age 14 as WN9QCH. He presently holds an Extra Class license, and is a life member of ARRL. Dave has a BA in mathematics from Goshen College and BS and MS degrees in electrical engineering from the University of Illinois. Dave is a Principal Engineer for Continental Electronics Corporation, headquartered in Dallas, Texas. Continental is involved in the research, design and manufacture of special purpose super-power RF energy systems as well as broadcast and communication transmitters. In his "spare time" Dave produces DSP kits for the amateur market. Dave is active on HF SSB and CW and VHF FM and packet. Dave's e-mail address is w9gr@psyber.com