Design Considerations for the KDSP2

By Lyle Johnson, KK7P
May 6th, 2003

Background

I am more of an SSB operator than a CW operator. I found my K2 to be a good radio for SSB, but I wanted a radio that was as desirable for SSB operation as it is for CW.

The missing features for reception included a denoiser to reduce listening fatigue, a notch filter to kill heterodynes as other stations tuned up nearby, and more versatile filtering. An outboard DSP unit could be used, but that is inconvenient at best. And for portable operation, it is clumsy.

What about integrating a DSP unit into an EC2 enclosure along with the KPA100 and KAT100? Hmmmm…that wouldn't help for portable operation. Examination of the K2 design revealed a DSP unit could be installed in place of the KAF2 analog audio filter.

Why DSP?

DSP is not new to the Amateur community. In fact, almost every current transceiver either offers it included, or as an optional accessory. Throughout the 1990s, the transition was made from external DSP to internal DSP in most radios. Today's high-performance radios incorporate DSP to implement more and more of their functionality.

Digital Signal Processing allows the creation of multiple filters, including adaptive filters, at low cost and with consistent performance. Dave Hershberger, W9GR, led the way with low-cost DSP with his landmark article "Low Cost Digital Signal Processing for the Radio Amateur," published in the September 1992 issue of *QST*. He explained some of the mysteries of its operation in "Using the LMS Algorithm for QRM and QRN Reduction" in the September, 1992 *QEX*. Johan Forrer, KC7WW, expanded on Hershberger's work when he published "A DSP-Based Audio Signal Processor" in the September, 1996 issue of *QEX*. Finally, Bob Larkin, W7PUA, presented his homebrew Software Defined Radio as the DSP-10 article series starting in the September, 1999 issue of *QST*, using DSP to create most of the radio's feature set.

While some of the performance of DSP can be realized with analog circuits, complexity increases dramatically when more than a few filters are required, or if adaptive noise canceling is desired. The Datong Frequency Agile Audio Filter was an early design that used analog circuitry to suppress heterodynes - an automatic notch filter (see *QST* for August, 1979, page 45).

AF or IF Processing?

The earliest Amateur DSP units were audio (AF) based. They were connected to the headphone or speaker jack, and usually included a speaker amplifier. Impressive performance was possible, but usually required careful setting and monitoring of audio levels to get enough audio signal, yet not overdrive the DSP.

Processing at IF often allows better functionality, particularly if the DSP is used to implement the radio's automatic gain control (AGC). IF processing also enables some interesting features, such as binaural reception. Due to limitations of the circuits used to convert analog to digital signals, most DSP IF processing uses a really low frequency IF of 10 to 40 kHz. This means multiple conversions, popular in many of today's radios, usually at the expense of phase noise, birdies, images and other performance-robbing artifacts.

Design Tradeoffs

The Elecraft K2 transceiver already implements many of the features that dictate the use of IF DSP in other radios. The AGC system is excellent, and the variable width IF filter allows rejection of most signals that might otherwise "pump" the radio's AGC. For the most part, the power of DSP can be harnessed to the K2 using AF DSP.

Finally, the K2 uses single conversion to obtain its excellent performance. Adding another conversion stage didn't seem like a good idea.   Or even a very practical one.

Balanced Audio

When the design started in earnest, it was apparent that the DSP would need to accept the audio signal output directly from the product detector. This is a differential audio signal. The output of the DSP would also have to be differential to preserve the balance and external audio noise rejection characteristics of the radio.

AGC Considerations

You may have noticed that many radios with internal DSP do not allow AGC to be turned OFF in order to protect the DSP's analog-to-digital converter input circuit from overdrive.

Since some K2 operators operate with AGC OFF, and others never turn it off, the dynamic range of audio signals presented to the KDSP2 is much wider than is common in most AF DSP systems. The AF GAIN control of the K2 is **after** the audio filter connector, so allowing the operator to adjust the audio level going into the DSP is not realistic.

During field test it was determined that an audio limiter, driven only in the case of a very strong signal with AGC OFF, is necessary to prevent overdriving the DSP. The alternative would be an internal AGC loop within the DSP. Adding another AGC loop to a radio of the caliber of the K2 is a step that must be taken only after a lot of careful thought. After all, if the operator wants AGC OFF, he might not be happy with a DSP AGC system that could not be defeated!

Power Miser

The K2 is, at its heart, a radio designed to be taken into the field for portable operation, often using an internal battery pack. In order to be compatible with this philosophy, the DSP unit had to be designed to be very power efficient.

This immediately dictated the use of a 16-bit DSP chip. 32-bit DSPs generally use a lot of power, often measured at a watt or more! And, they tend to be quite expensive. 16-bit DSPs, on the other hand, are often used in portable audio equipment, like MP3 players. The Analog Devices ADSP-2184N was chosen for the KDSP2. It consumes very little power, has more processing muscle than is needed for the application at hand, is low in cost, and is reasonably compatible with DSP code already developed for some of the projects mentioned above.

At the heart of any DSP system are the analog-to-digital (ADC) and digital-to-analog (DAC) converters. Since the K2 uses balanced audio, dual ADCs and DACs were indicated, or an analog circuit to convert between balanced and unbalanced circuits.

A search of suitable components brought the Texas Instruments TLV320AIC23B CODEC (coder-decoder, DSP-speak for a combined ADC and DAC component) to light. This is a stereo device; hence it has dual ADCs and DACs. It was designed for applications like pocket MP3 players. This meant it had very low power consumption, and very good ADC and DAC performance. It is a 16-bit device, so it is a perfect match to a 16-bit DSP. No other readily available part seemed to combine the features of dynamic range and low power.

But, there was a catch. The TI part was designed for TI DSPs. Its processor interface was incompatible with the Analog Devices DSP - and the Analog Devices' compatible CODECs consumed more power. What to do?

Xilinx is a company which makes a wide range of programmable logic devices, called CPLDs. They offer a very low power series of economical CPLDs, the XPLA3 family. A Xilinx CPLD was pressed into service to adapt the DSP to the CODEC.

Lastly, Flash memory was needed. The Flash had to store the program for the DSP, as well as the filter parameters (called coefficients) for the multitude of filters that would be desired. It would mostly be idle, though.

Modern cellular phones use lots of Flash memory, and the technology has evolved to be low power, with almost zero power drain when not active. Perfect for our application!

In the end, the KDSP2 that resulted from these decisions draws only about 60 mA when operating, and can be placed in a bypass mode that typically consumes less than 10 mA.

This met the need for battery-powered field operation. If interference is bad, 60mA will get rid of it. When the bands are friendly, 50 mA can be saved.

Power Regulator

Even more power could be saved if a switching regulator were employed rather than a linear regulator to power the KDSP2. But, switchers using through-hole technology are bulky and may create a lot of nearby, noisy magnetic and electrical fields. Since the KDSP2 would be located in the vicinity of the PLL, it seemed prudent to not use a switcher.

Dynamic Range

With the latest 3 kilobuck radios extolling 24 bit ADCs and 32-bit DSPs, can 16 bits do the job?

The answer is yes. As with any engineering problem, there are tradeoffs. 16 bits provide a theoretical 96 dB dynamic range - that is, the ratio from the weakest discernable signal to the strongest one that can be tolerated. The ADC used in the KDSP2 has an actual dynamic range of about 90 dB.

The K2's AGC keeps the signals at the output of the product detector well within this range. Use of the DSP in no way reduces the dynamic range of the **radio**.

A typical signal of moderate-to-strong strength will be about 20 dB below the peak capability of the ADC. This means the "noise floor" is about 70 dB below that. Real world considerations, at least at this QTH, make that noise floor more like 30 to 50 dB below the peak signal.

The CW and SSB filters in the KDSP2 are designed to provide no more than 2 dB of ripple and at least 65 dB of attenuation to signals outside the design audio passband. The data filters are flat within 0.1 dB at the cost of reduced ultimate attenuation on the order of 40 to 45 dB. This provides great **practical** selectivity with little delay.

What is Filter Delay?

When a signal passes through a filter, it is delayed in time between the input and the output. If the filter is narrower than the signal trying to get through it, there may also be other issues, such as ringing.

Whether the filter is analog using op amps, or digital using DSP, the issues are the same.

Most DSP filters use an approach called Finite Impulse Response (FIR), while op-amp and switched capacitor filters are almost universally Infinite Impulse Response (IIR) types. Before your eyes glaze over, this is just engineer-speak for two very different ways to make a filter.

IIR filters can become oscillators. Sharp filters will ring. But they are simple and practical and widely used.

FIR filters can't oscillate, and ringing is usually less. They aren't very realistic for analog circuits, but are quite practical for DSP.

When people talk about FIR filters, they start discussing taps and filter length. What this really comes down to is that to have steep skirt selectivity and good rejection, an FIR filter needs to be long, with a lot of taps. Many audio DSP units have filters with 200 or more taps. Long filters with lots of taps mean greater signal delay from input to output.

After studying the effect of various lengths and types of FIR filters, the KDSP2 filters were designed to use 128 taps. The DSP can easily handle more, and the Flash memory can hold more. But 128 provided good selectivity while keeping the delay through the filter reasonably short. In fact, the signal only "lives" in the filter for 16 milliseconds, and has its peak effect at 8 milliseconds. Many other DSP filters have twice the delay. With high speed CW, the shorter delay may be beneficial.

Full Break-In Supported

One complaint often voiced by CW operators using outboard AF DSP relates to the delay in the filters. Full break-in operation is often not practical with these systems.

The KDSP2 is in the K2 audio chain **prior** to sidetone injection. Since the sidetone doesn't pass through the DSP, filter delays don't affect the ability of the CW operator to hear his own keying in real-time. Thus, the excellent full break-in ability of the K2 is preserved when using the KDSP2.

Integration

Getting into the audio path, and tapping into the radio's power system, were easy. The KAF2 paved the way, and the K2 was designed from the outset to support the KAF2, so connectors were available.

The user interface, on the other hand, required significant effort. The first paper design required jumpers, LEDs and knobs hidden inside the K2, accessible with the top cover removed, or by drilling holes in a side panel. Yeech!

It was clear that the K2's existing front panel controls had to be used, along with the display, or the project would be one-of-a-kind.

There was no room in the K2 MCU for additional code to support the KDSP2, so another way had to be found. This meant interfacing to the auxbus.

A little more study made it clear that the KDSP2 could pre-empt the KAF2 and its real-time clock interfaces and warp them into the buttons and display resources required.

Once that was decided, a user interface had to be designed in the spirit of the K2. Operation had to be simple and fast. Control functions had to be clear. Menus had to be consistent. Every feature had to be accessible from the front panel without opening the radio. Often-used features had to be accessible quickly, with minimal button pressing.

After some effort, a menuing scheme was developed that was practical and consistent with the K2. It has since evolved into a very flexible, yet convenient, interface.

For example, four filters are available for each radio mode (CW, SSB and RTTY). You simply press the AFIL button to cycle through them. If you change modes, the last filter selection you made in that mode will be retrieved.

Each of the four filters can be set up from a menu with up to 512 filters to choose from.

You can toggle the denoiser (automatic noise reduction) on and off from the first DSP menu, by tapping the DISPLAY button. In SSB mode, you can also toggle the automatic notch filter from the same first menu.

Flexibility

K2 owners vary widely in their use of the radio. Some operate exclusively CW; others SSB; still others HF data modes. Some use their radios for microwave work with transverters. Some are serious contesters. Some are casual ragchewers.

It is hard to make "one size fits all." The K2 uses a modular approach to satisfy the varying needs of the community. The KDSP2 uses a similar approach in its design.

Many operators will be satisfied with the KDSP2 out-of-the-box defaults. They'll just set the clock, and use the radio, selecting among the filters and toggling the denoiser on and off. If they go camping, they may put the DSP in bypass mode when the band is friendly to save current consumption.

But others will want to tweak things. They like to take the covers off of their radios. They drive stick shift automobiles.

The KDSP2 may be the most configurable DSP available today.

You can select from hundreds of filters - as with many DSPs.

You can turn the denoiser on and off - as with many DSPs.

You can select from various levels of denoiser performance - as with many DSPs.

You can get under the hood, so to speak, and play with the internal LMS algorithm Beta and Decay factors. Huh? This means you can tune the automatic noise reduction and notch filter performance to suit your tastes and the need for the band conditions at the time.

You can balance the gains of the various mode filters as well as the gain of the system when the CW or SSB denoisers are activated.

Every adjustment or setting is remembered in EEPROM, so it is there the next time you turn the radio on.

You can truly customize this DSP to your personal tastes.

Distortion and Artifacts

Any time you pass a signal through a filter, you alter it. If you don't, you didn't need the filter!

DSP filters are no exception.

If you use an automatic notch filter, the filter will be constantly seeking something to notch. Whether done at IF or AF, you will hear it. If you activate a denoiser, it will look for noise to suppress. Like the notch, you will hear it. Some may call these effects distortion. Others may refer to them as artifacts. All filters have these effects.

While it is not practical, or perhaps even possible, to completely eliminate these effects, you can adjust them. Less suppression means less audible effect on the signals you want to hear.

Different DSP algorithms have different-sounding qualities of these effects. The KDSP2 uses least-mean-squares (LMS) algorithms for its notch and denoiser, and produces audio artifacts typical for this type of algorithm.

Noise Reduction

There are all sorts of sounds that we Amateurs collectively term "noise." Static crashes, ignition noise, interference from appliance motors, power line noise, and background "hiss" all fall under the definition.

High amplitude impulse noise, such as that from an ignition system, is readily handled by IF noise blankers that operate on the signal before the selectivity of the IF filter can affect them. For the K2, the KNB2 does a very good job.

The noise reduction for which DSP is most suited, is that of background hiss. It's the sort of noise you hear when you attach an antenna, tune to an unoccupied part of the band, and crank up the gain.

Or when you are trying dig a weak station "out of the noise."

It is non-repetitive, and random in nature. It is sometimes referred to as thermal noise because it is the result of random motion of electrons in a conductor, like an antenna.

This is the type of noise which the KDSP2 is designed to mitigate.

In the world of DSP, there are two major approaches (algorithms) used to reducing this type of noise: LMS (least-mean-squares) and spectral subtraction.

LMS algorithms analyze a signal on the fly, looking at waveforms in time. This is the sort of display you might see on an oscilloscope. Using LMS, the DSP attempts to distinguish between signals that look like noise, and signals that look like, well, signals. It then filters out the noise.

Spectral subtraction algorithms look at a signal based on frequency components, like a spectrum analyzer display, or a "bandscope." If a decision is made that the energy in a particular frequency area is noise, it is so marked. The frequency information is then converted back to a waveform in time, which is passed on to the audio stages of the receiver.

LMS, used in the KDSP2, is usually faster than spectral subtraction. It has its adherents, as well as its detractors. Like most sounds we intentionally listen to, some like it one way and some prefer the other.

Kitability

Early on, it was apparent that modern DSP technology dictated the use of surface mount technology. This is anathema to many kit builders.

The solution was to make a two-board system.

The DSP portion is a pre-assembled and tested module. It plugs into the interface board like a large IC.

The interface to the K2 - MCU for auxbus transactions, retrieving button presses, displaying menus, audio levels and muting - is a through-hole technology kit in the Elecraft tradition.

The result is a high performance DSP that is available to kit builders. All it really lacks are toroids to wind…

Open Source

One of the reasons for choosing the ADSP-2184N over the TI DSP chips was the availability of code done by other Amateurs, using Windows and DOS software tools that were free. By using proven algorithms, the work they had already done could be applied. And by putting it into the KDSP2, their work could assist many more people.

A very real benefit to this is that by using open source code, the resulting code is also open source. This means that you can study it, use it, improve it -- it is great for education and self-training.

Another advantage is that if you do improve it, you must contribute it back to the community, so others can benefit from your efforts, too.

Field Testing

When the KDSP2 project was first revealed on the Elecraft reflector, an overwhelming number of people responded to the request for volunteers to participate in field testing the design. Limiting the number of participants was a difficult task. The people selected were and remain a tremendous asset to the project.

About this time, Elecraft and I agreed that the project would benefit if it were adopted by Elecraft. The Field Testing phase was hosted by Elecraft as the effort changed from a personal one to a team. And what a team!

Some weaknesses in the design were uncovered early on. Circuit changes were made, new firmware made and distributed. Features were added, user interfaces were altered,

menuing systems were improved. Everyone had suggestions and made valuable contributions.

Being a sanctioned project provided access to some internal Elecraft design information that helped make a good product better.

The initial production release of the KDSP2 represents literally hundreds of man-hours of testing and evaluation by a patient, enthusiastic group of fellow Elecrafters.

I am indebted to each of them.

Conclusion

The KDSP2 provides high-performance AF Digital Signal Processing for the K2.   It works well within the system design constraints imposed by the K2. Flash-based memory systems ensure field upgradability.

It also provides an opportunity for users of the product to directly contribute to its enhancement -- and education on the way DSP is used in a modern radio -- by employing the open source model for the DSP software.

The KDSP2 is in a very real sense a product of the K2 community.

The KDSP2 Field Test Team

Dale Boresz, WA8SRA
Wayne Burdick, N6KR
Dave Gropper, KK7SS
Lyle Johnson, KK7P
Carl Morris, WN3DUG
Mike Murray, W0AG
Bob Nielsen, N7XY
Steve Russell, N7SR
Gary Surrency, AB7MY
Eric Swartz, WA6HHQ
Bob Tellefsen, N6WG