WA4WDL describes a homebrew double balanced beat indicator that he uses for precise frequency measurements from 5 kHz to 20 MHz. At the heart of John's design is a pair of North Hills model 1301 baluns and a quad array of Schottky diodes.

QEX

In order to ensure prompt delivery, we ask that you periodically check the address information on your mailing label. If you find any inaccuracies, please contact the Circulation Department immediately. Thank you for your assistance.

About the Cover

John Franke, WA4WDL, was inspired by an article in the September/October 2009 issue of *QEX* magazine titled "Atomic Frequency Reference for Your Shack" by Bob Miller, KE6F. The design included a double balanced beat indicator as a means for quickly comparing two oscillators. John determined that while the project used a regular single-polarity meter, a zero-center meter would have been better. John's solution was to build his own using two North Hills baluns, the associated circuitry and a surplus enclosure. Read about it in this issue!

In This Issue

Features

Index of Advertisers

Larry Wolfgang, WR1B

Empirical Outlook

I have written a number of times on this page about how much I enjoy building electronics projects. I believe many readers also share my enthusiasm for "melting solder." The most popular *QEX* articles seem to be those that describe the construction of some piece of radio gear, whether it is a transmitter, receiver, transceiver, or some piece of test gear. Readers seem to like any project that includes full details as well as the availability of circuit boards and even full or partial parts kits.

In my case, this goes back to junior high school, before I even had an Amateur Radio license. I remember building several crystal radio sets with my dad, and then experimenting with different circuit configurations on those receivers. I built a 2 tube receiver as a science fair project one year, and although it didn't receive as many short wave stations as I had hoped, experimenting with the circuit and "spare" parts I could find proved fun and educational. That project was "bread-boarded" on a piece of pine, with terminal-strip tie points and "open frame" wiring. Today most people would cringe at the sight, and insist on some protective cabinet and other electrical safety features. That's the way nearly everyone built projects in the 1960s, though.

In more recent years I have stuffed many circuit boards and even ventured into the realm of surface mount components. I have talked with many hams, both old timers and new recruits, who lament that people can't build anymore. I have heard that it is too hard to find parts, because many of the neighborhood electronics emporiums have disappeared. Of course we can still find many parts at hamfests, and there are mail-order companies that will sell small quantities to individuals. Club group purchases are a good way to work around minimum purchase requirements, and can even reduce the price if you are ordering in larger quantities. It does take a little more planning, however, than just a quick trip to a RadioShack or other local supplier.

I like to encourage people to try building with surface mount components. I really believe nearly anyone could enjoy the thrill of building some new radio or accessory and using it on the air if they would just try some of the tools and techniques available to us. Of course, it is true that these parts get smaller and smaller, with more and more leads per inch on the ICs. It is becoming more difficult to find some traditional parts with leads, and some transistors and ICs are being phased out by the manufacturers. I think this points to the fact that for those of us who like to build, we will have to adapt our procedures and revise some of our expectations.

Certainly the availability of kits makes it easier to assemble that new project. I recently began working on a new kit that will end up as a *QST* Product Review. While it is quite an extensive project, and the end result will be a piece of test equipment with some very impressive features, the construction involves installing "old" leaded parts on through-hole circuit boards. The PIC processor board and several other boards that have surface mount components come fully assembled. This helps illustrate the fact that circuit assembly has become an automated process. It is fairly simple to have a robotic arm select, pick up and place a surface mount part in the exact right place. You and I will struggle to hold that part, need some large magnification to verify the component value or ID, and will be challenged to place it exactly. Machines can also pick and place those leaded components very easily.

In those "good old days" we often repaired a piece of equipment by troubleshooting to find the defective component, and then removing that capacitor, resistor or whatever it was, and replacing it with a new part. This was really pretty simple when all of the components were soldered to a terminal strip fastened to a metal chassis. We have also removed and replaced defective components on through-hole circuit boards. Today, the component density of the typical circuit board — even those we build with surface mount parts — make it much more difficult to remove and replace defective parts. Consumer electronics are truly throw-away products. We may try a little harder with a piece of our Amateur gear, but more often the best we can hope for is to swap a circuit board rather than a component-level repair.

I was thinking about all of this because I am writing this editorial as I rehab from a total knee replacement. Isn't it interesting that as our electronics become more difficult to repair, even just thrown away when it breaks, medical science has advanced to the point that replacing worn out body parts is a routine procedure? While I learn to use my new knee, another member of our ARRL HQ staff is out having a second hip replacement. I am aware of several other current or former HQ Staff members who have artificial knees, and I imagine that just about every reader knows someone who has had a knee or hip replacement. I hope we can continue to repair at least some of our electronics, but I am sure thankful that we can so easily replace some worn out body parts!

James Ahlstrom, N2ADR

221 Pleasant Plains Rd, Stirling, NJ 07980 **jahlstr@gmail.com**

Using a Digital Transceiver as a Vector Network Analyzer

The author describes how his digital transceiver hardware can be used as a vector network analyzer just by changing the software.

In a previous article I described a transceiver based entirely on digital electronics.[1] The transmit signal was generated by a digital-to-analog converter (DAC) and the received signal was sampled by an analog-to-digital converter (ADC). Computation was performed in a field programmable gate array (FPGA), and Ethernet connected the hardware to a personal computer that ran my *Quisk* software.[2] Another article by Detlef Rohde, DL7IY and Günter Richter, DL7LA appeared in Funk Amateur.[3] This led to an improved hardware version called the HiQSDR by Helmut Göbkes, DB1CC, and there are PC boards and completed units available.[4] Stefan, DL2STG, has contributed improved software,[5] including *QSDR*, an alternative to *Quisk*, *QVNA*, a VNA program that can plot Smith charts, and *HiQScope*, a program to return a block of samples so you can display the whole 61 MHz bandwidth. There are two groups available for more information, one on Yahoo[6] and one on Nabble.[7] It is a real honor for me to work with other talented hams around the world on our joint homebrew project.

In my 2010 article, I noted that the hardware was quite generic, and that the software made it a transceiver. I suggested that with different software it could just as well be a spectrum analyzer and tracking generator. After a little thought I realized that if I could measure the phase, I would have a vector network analyzer (VNA), an instrument that generates an RF signal as an output and measures the amplitude and phase of a return signal at its input. This VNA can be used to measure the response of filters, calculate the impedance of antennas and

calculate S parameters. This article describes the changes needed in the software for the FPGA and the PC. No changes were needed in the hardware.

The Basics

We need to generate an RF signal at the output of the hardware and receive the return signal while measuring its amplitude and phase relative to the output. We also need to do this at multiple points over some frequency range so that we can plot a graph. For example, we may want a graph of the SWR of a 40 meter dipole at one thousand points between 7.0 and 7.3 MHz, or one point every 300 Hertz. This is not something that a transceiver would normally do, so we need to change the FPGA and PC software. The changed software works for the original hardware described in *QST*, and for the newer HiQSDR. We are going to make the transmit and receive frequencies the same so that the received signal is a quadrature I/Q signal at dc. That could cause drift problems as an analog design, but is not a problem for a digital mixer.

Changes to the FPGA Software

The transceiver has a digital VFO for receive and another independent digital local oscillator for transmit. That enables the transmit and receive frequencies to be different, and enables both to operate at the same time for full duplex operation. But there is no reason that these two oscillators should have any phase relation, especially as the frequency changes. So the first change is to send the receive VFO directly to the DAC for transmit; that is, the same oscillator is used for transmit and receive. This means that

there is a constant phase between the two, but not necessarily a zero phase difference. As we will see, we will need a calibration run to measure the phase difference. Next, we need to collect about 1000 data points at different frequencies for the scan. There are about 1000 points because we are measuring one point for each pixel in the graph display, and a modern display is roughly 1000 pixels wide. The frequency scan could be done in the PC software, but it is easier to do it in the FPGA. We are receiving a continuous stream of data samples from the hardware, and if we change the frequency by PC command, we would need to figure out which samples came from the old frequency and which came from the new. And a PC scan would be slow because even for a one millisecond Ethernet command, that would be one second for 1000 points.

So we change the FPGA software as follows: We send the number of data points desired as an addition datum to the FPGA. We send the receive frequency as usual and use it for both receive and transmit. We send the frequency increment between each point (as a phase increment) in the field that used to be the transmit frequency. When the number of data points is zero, the FPGA operates as usual as a transceiver. Otherwise, the FPGA operates as follows: The frequency is set to the starting receive frequency. Then the FPGA discards the next 8000 points while the external circuit adjusts to the new frequency. At a 122.88 MHz clock, this is a delay of 65 microseconds. Then, the next 4096 data points are added up. The data points are I/Q signals at dc; that is, they are complex numbers. So, the averaging acts like a narrow low-pass filter that tunes in only the exact transmit frequency. This pre-

vents spurious signals, for example, from an antenna, from distorting the measurement. This averaged point is sent to the PC, the frequency is incremented, and the process repeats. After all the points are sent and the scan is complete, the FPGA sends a zero data point. The PC looks for this zero data point so it can figure out when the scan starts. The PC can check that it has exactly the number of data points specified, and knows they are in correct order.

Changes to the PC Software

The *Quisk* software is a transceiver and we are trying to do something completely different, so the solution is to write a completely new VNA program *quisk_vna.py* to replace the transceiver program *quisk.py*. There are a few other changes to enable reception of the scan data blocks, but nothing major. This enables us to write new code to solve the problem at hand, namely plotting each data point as a pixel on a graph. The VNA program still uses features of the old code that creates buttons, draws graphs, etc.

The new VNA software supports a transmission mode and a reflection mode. Transmission mode means that the RF output is connected to an attenuator, then to a test device like a filter, then to the RF input. Reflection mode is used with a resistive return loss or SWR bridge and is a bit more complicated.

Hardware Protection

Note that if you directly connect the hardware RF output to the input, you will overload the RF input, the level will cause the ADC to clip and destroy the data validity, and you may damage the preamp and even the ADC. Never directly connect the RF output to the input without an attenuator that reduces the RF to a safe level. For some measurements RF is connected through an amplifier to the input and the extra power may certainly cause damage. What I mean is you must *always use attenuators* to reduce the output power and input power to a safe level that does not clip the ADC. I use the HAT series of attenuators from Minicircuits,[8] which are available as an inexpensive kit. The measurement accuracy depends on an accurate 50 Ω output and input impedance in the hardware and the attenuators will improve the accuracy too.

Transmission Mode

In transmission mode we connect the RF output to an attenuator and then connect a test device such as a filter, amplifier or cable to the input. The graph will plot the response of the device. If an amplifier is in-line add additional attenuation at the input.

First we must have a hardware calibra-

tion. We connect the output to the input with just an attenuator in line. We are trying to measure and correct for any variation in amplitude and phase due to the Quisk or HiQSDR hardware absent a device under test (DUT). To do this, we measure the amplitude and phase every 15 kilohertz from zero to a little over 60 MHz, or 4003 data points. These measurements are saved and used as corrections with linear interpolation for frequencies between these points. The calibration data are shown on the graph. Note that transmission mode needs its own calibration, and if you change to reflection mode the calibration must be repeated. Calibrations are not saved, and if you exit the program, you must perform the calibration again. It is not necessary to repeat the calibration if you just change the frequency span of the graph. Note that it is up to you to connect the cables as

appropriate for the calibration; the software cannot "see" what you are doing.

To perform a calibration, first select transmission mode and press **Cal Remove**. Then, connect the output to the input with an attenuator. Push the **Cal Short** button. If you want, you can remove the cable and press **Cal Open** to measure and subtract background, but this is optional. The result is shown in Figure 1. The **Trans** and **Refl** buttons select transmission or reflection mode, the **Cal** buttons start a calibration run, the **Run** button is used for a measurement and the **Help** button provides help. The red vertical line is a tuning line and causes the lower status line to show data for the indicated frequency. The black line starting at -10 dB and decreasing gradually is the amplitude, and it decreases due to the filters in my hardware that are designed to attenuate signals above
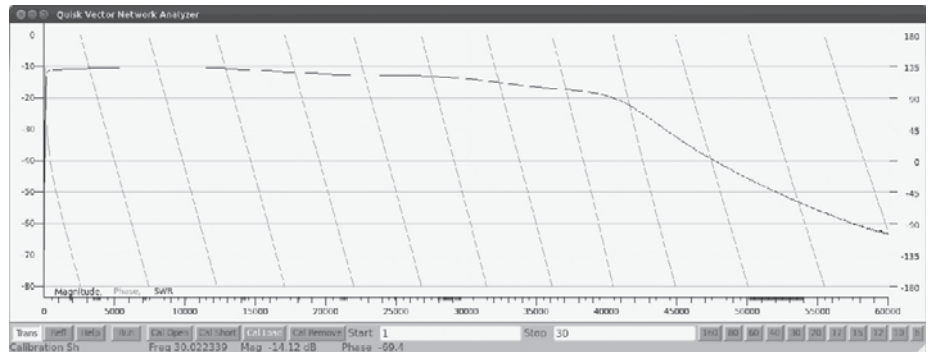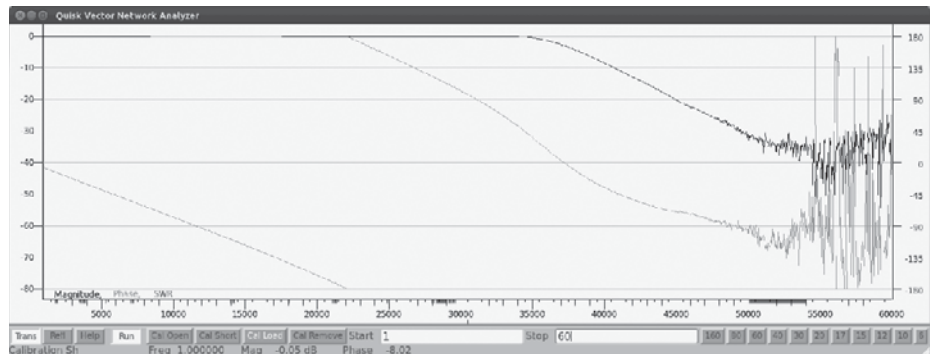

Figure 1 — Transmission mode calibration graph.

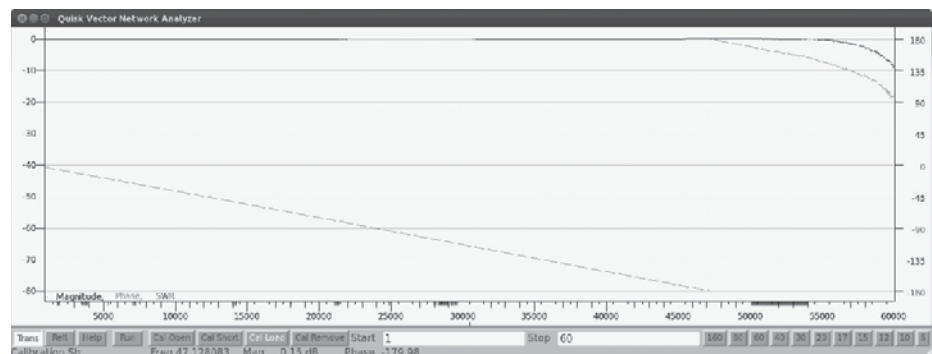Figure 2 — Transmission response of 35 MHz filter.

Figure 3 — Transmission response of 205 cm cable.

30 MHz. Despite the reduced level, the VNA works well to over 50 MHz. I was using a 20 dB attenuator and the maximum level was comfortably below the maximum level of zero dB. I could have used 12 dB; generally you should use just enough attenuation to be below 0 dB, but not too much more.

The diagonal green lines are the phase, and we see that it varies with frequency on a regular basis. The phase is plotted from minus 180 to plus 180 degrees, and it folds over to the bottom when it goes over 180 degrees. This can result in some ragged and confusing phase graphs if this point is not kept in mind. The phase repeats every 4.7 MHz and is caused by a time delay between transmit and receive of 211 nanoseconds, or 26 clocks. Some of this delay is in the pipelined ADC and the rest is in the registers. Although a constant phase would have been nice, this doesn't cause problems because the phenomenon is linear and the linear interpolation of the calibration data corrects for it. If you now press **Cal Short** again to end calibration and press **Run** to take a measurement, two nice flat lines at zero dB and zero degrees result.

Figure 2 shows the result of inserting a 35 MHz low pass filter. Note the ragged display above 55 MHz, which is caused by the low signal levels.

Figure 3 shows the result of inserting a cable. The phase changes by 180 degrees, or a half wavelength, at a frequency of 47.128 MHz. The speed of light is 299.8 million meters per second, so the electrical length is 318 centimeters. The physical length is 208 centimeters, so the velocity factor is 65%. We can see that this is useful for measuring the electrical length of cables. But the other side of the coin is that adding a cable distorts the phase when we are trying to measure something else. Adding a device generally adds more cable, and that changes the phase. This is not too troublesome at
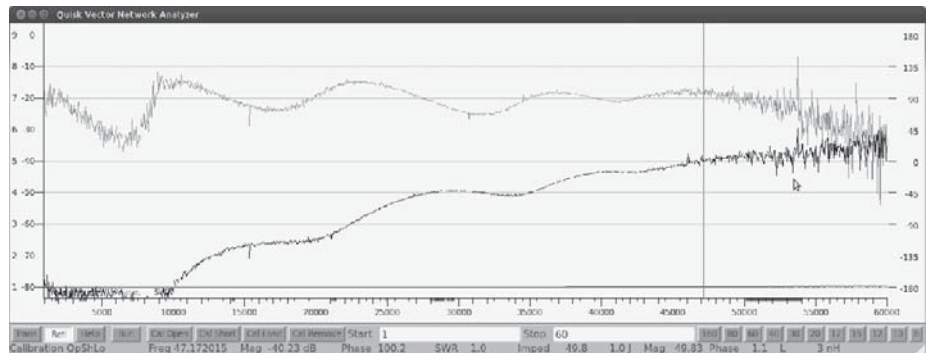
Figure 5 — Return loss of a 50 Ω standard.
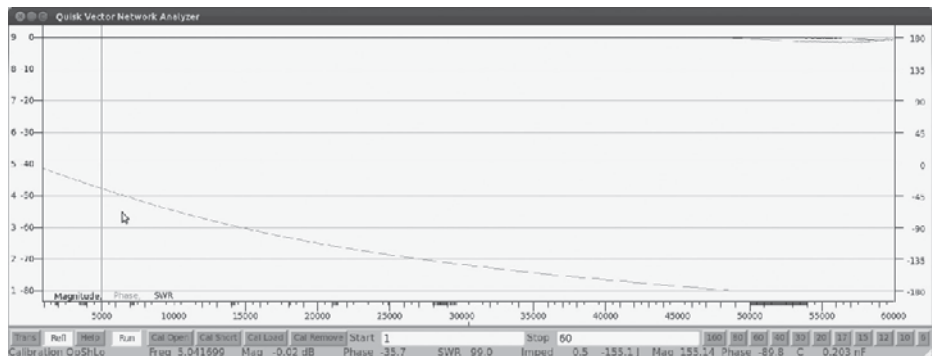
Figure 6 — Return loss of a 200 pF capacitor.
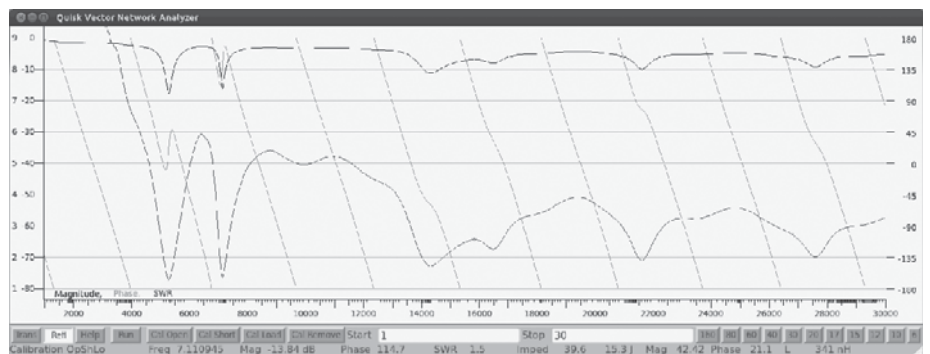
Figure 7 — Return loss of my fan dipole.

HF, but it is a major headache at microwave frequencies, and engineers must specify the exact position along the connections that the phase is measured.

Reflection Mode

Reflection mode requires a resistive return loss bridge as shown in Figure 4. Each bridge arm is 50 Ω and, if the unknown impedance is also 50 Ω, the bridge is in balance and the detector output is zero. Directions for constructing such a bridge are in the ever-useful *Experimental Methods in RF Design*.[9] Here we need to know exactly what output voltage V we can expect for an arbitrary unknown impedance. This can be solved by a bit of algebra, as shown in *Radio Frequency Design*,[10] and the result is:

$$V = (Z - 50) / (Z + 50)$$

We can solve this for the complex impedance Z given V, but more importantly, V is the definition of the complex reflection coefficient. So we are directly measuring the reflection coefficient, a very useful quantity. Note that this assumes that the generator and detector both have a 50 Ω impedance.

For reflection mode, first select reflection and press **Cal Remove**. Then connect the RF output to an attenuator and then to the generator terminal of the bridge; connect the RF input to the detector terminal. Now connect an open connector to the unknown terminal, or leave it unconnected, and press **Cal Open**. This provides a minimal calibration and you can now take data.

To do better, connect a shorted connec-

Figure 4 — Schematic of a resistive return loss bridge.

tor to the unknown terminal and press **Cal Short**. To do better still, connect a 50 Ω standard to the unknown terminal and press **Cal Load**. Leaving the 50 Ω standard in place, press **Run**, and the result is Figure 5. The point here is the low baseline amplitude. The phase does not mean much due to the low amplitude. Figure 6 shows the result of connecting a 200 pF capacitor across the unknown terminal. The capacitance can be read in the status bar. You can also measure inductors and resistors. Try to keep the impedance around 50 Ω, perhaps in the range 5 to 500 Ω for best accuracy. The impedance will seem to change with frequency due to stray inductance and capacitance. Think about what you are doing to interpret your measurements correctly.

The real point of this VNA is antenna measurement. Figure 7 shows the result of connecting my fan dipole to the unknown terminal. It shows a graph of the SWR from 1 to 30 MHz. The tuning line at 7.111 MHz shows an SWR of 1.5 in the status bar, and an impedance of $39.6 + 15.3j$ Ω. This fan dipole was designed for the 60, 40 and 20 meter bands. Apparently the 80 meter band needs a longer antenna. At 3.5 MHz the SWR is 6.7 and the impedance is $8.2 - 32.6 j$.

Conclusions

I am not proposing that this is a good design for making a VNA since the frequency response is limited to the 1 to 55 MHz range. There are many other designs that are better.[11] But there are two points to be made: First, those users who have my hardware or the HiQSDR now have a powerful VNA at zero cost. The frequency range is limited, but it covers all the frequencies of the transceiver and, therefore, all of the relevant antennas. Second, it shows the power of digital electronics in general and Software Defined Radio in particular. We are almost used to the idea that a new software version can update our DSP radios with new features, such as a better noise blanker or AGC. This project demonstrates that sometimes a software update can convert an apple into an orange, and a free orange at that!

References

[1] James C. Ahlstrom, "An All-Digital Transceiver for HF," *QEX* January/February 2011, pp 3-8.

[2] **http://james.ahlstrom.name**

[3] Detlef Rohde and Günter Richter, "Direktabtast-Transceiver für *Linux*", *Funk Amateur*, August 2010, pp 814-817.

[4] **www.hiqsdr.org/**

[5] **http://dl2stg.de/stefan/hiqsdr/index.html**

[6] **http://groups.yahoo.com/group/n2adr-sdr/**

[7] **http://quisk.973856.n3.nabble.com/**

[8] **www.minicircuits.com**

[9] Wes Hayward, Rick Campbell and Bob Larkin, *Experimental Methods in RF Design*, ARRL, 2003, ISBN: 0-87259-879-9.

[10] Wes Hayward, *Radio Frequency Design*, ARRL, 2004, ISBN: 0-87259-492-0

[11] Prof. Dr. Thomas C. Baier, "A Small, Simple, USB-Powered Vector Network Analyzer Covering 1 kHz to 1.3 GHz," *QEX* January/February 2009, pp32-36.

James Ahlstrom, N2ADR, was first licensed as KN3MXU in 1960. He received a BS in physics from Villanova University in 1967 and a PhD in physics from Cornell University in 1972. He then moved to New York to work in the financial business. He is now retired. His license lapsed while raising his family, and he was re-licensed in 2006. He currently holds an Amateur Extra class license. Besides Amateur Radio, he enjoys bird watching, skiing, music and working out at the gym. **QEX-**

Bob Dildine, W6SFH

2321 Sycamore Ave, Santa Rosa, CA 95404 **w6sfh@sonic.net**

Buck Converter Basics

Learn how switching power supplies perform their unique "magic."

Several years ago I was asked to come out of retirement and take over the design of several high efficiency switching power supplies. I had almost no experience with switchers and considered them to be mysterious and full of black magic. But after getting my feet wet and working with them for a while, the mystique faded and I realized they were based on a few simple concepts that I hope to share here.

How It Works

In case you are unfamiliar with the term, a buck converter is a step-down dc-to-dc converter. Let's see how the converter works using the diagram shown in Figure 1.

When switch S1 closes, current flows through the inductor and into the load, charging the inductor by increasing its magnetic field and increasing V_{out}. When V_{out} reaches the desired value, we open S1 and close S2. Current continues to flow in the inductor as the magnetic field collapses and the inductor discharges. Before the inductor completely discharges, we open S2 and close S1 and the cycle repeats. We can adjust the ratio of V_{out} to V_{in} by varying the duty cycle of S1. The longer S1 is turned on, the greater V_{out} will be. The duty cycle of S1 is usually called the converter's duty cycle. If the switches and the inductor are lossless, V_{in} is converted to V_{out} with no loss of power and the conversion is 100% efficient.

The Math, Choosing an Inductor

Now let's look at the math that describes what's going on. The current through the inductor, I_L is made up of the average current which is equal to the load current, and the ripple current which is the change in current as the switches are opened and closed. See Figure 2.

The peak to peak ripple current is usually set to be about 30% to 40% of the load current at the start of the design and it is related to the voltage across the inductor by the familiar expression...

Figure 1 — A basic buck converter.
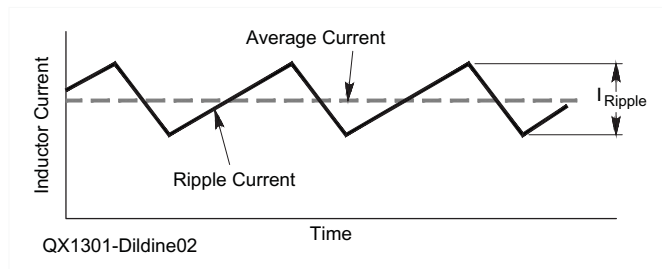
Figure 2 — Inductor current. Current increases when S1 is closed and decreases when S2 is closed.

$$v = L\frac{di}{dt}$$

Integrating and solving for i, we have

$$i = \frac{1}{L}\int vdt$$

or simply

$$I_{Ripple} = \frac{1}{L}V\Delta t$$

When the inductor is charging, V is the voltage across the inductor, $V_{in} - V_{out}$ and Δt is the time that S1 is closed. When the inductor is discharging, V is simply V_{out} and Δt is the time S2 is closed. (Of course this assumes there is no voltage drop across the switches, which in practice is not the case but we'll deal with that later.) It should be obvious that

I_{Ripple} is the same for the two cases, when S1 is closed and when S2 is closed, so

$$I_{Ripple} = \frac{1}{L}\left(V_{in} - V_{out}\right)\Delta t_{S1} = \frac{1}{L}V_{out}\Delta t_{S2}$$

Solving this equation for the ratio of V_{out} to V_{in}

$$\frac{V_{out}}{V_{in}} = \frac{\Delta t_{S1}}{\Delta t_{S1} + \Delta t_{S2}} = S1's\, duty\, cycle$$

which is what we said in the first paragraph.

From the expression for ripple current we can choose values for the inductor and the switching frequency. Notice that for a given current and voltage difference, the inductor value is proportional to Δt. In other words, the greater the switching frequency, the lower the required inductor value. But losses in the

inductor and switches increase with frequency so there's a practical limit. With modern parts this limit is around 1 MHz.

Let's take a simple example based on the circuit in figure 1. Say we have a 12 volt source and we need to supply the load with 5 volts at 0.6 amps. Let's choose a switching frequency of 300 kHz and assume a ripple current of 33% of the load current. What value inductor is needed? Taking the above equation for ripple current and solving for L we have...

$$L = \frac{(V_{in} - V_{out})\Delta t_{S1}}{I_{Ripple}}$$

S1's duty cycle (the converter duty cycle) is the ratio of the input voltage to the output voltage.

$$\frac{5 Volts}{12 Volts} = 0.417$$

The period of the switching cycle is

$$\frac{1}{300 kHz}$$

or 3.33 microseconds so

$\Delta t_{s1} = 0.417 * 3.33$ 1.388 microseconds. The ripple current is 0.33*0.6 = 0.2 amps. So the required inductance is

$$L = \frac{(12V - 5V)*1.388\mu s}{0.2 amps} = 48.6\mu H$$

As a check, we'll use the equation for ripple current during the time when S2 is closed:

$$I_{Ripple} = \frac{1}{L}V_{out}\Delta t_{S2} = \frac{1}{48.6\mu H} * 5V *$$
$$(3.33\mu s - 1.388\mu s) = 0.2 amps$$

which is what we started out with.

Notice a couple of things about the inductor value. It's proportional to the time that S1 is on which implies that it's inversely proportional to the switching frequency. And it's inversely proportional to the ripple current which implies that it's inversely proportional to the load current. So as the switching frequency increases, you need less inductance and as the load current increases you need less inductance. But this also means that for a certain value of inductor there is minimum practical load current. It's sort of analogous to the critical inductance in a choke input filter for a traditional linear power supply.

Besides the inductance value, there are two other important things to consider when choosing an inductor: dc resistance of the winding and core saturation. The dc resistance must be low enough that resistive losses are not a problem and the core must be big enough that it doesn't saturate with the peak ripple current. Of course the core size will depend on the material. In Linear

Technology's Application Note 35, Jim Williams gives some advice on empirically choosing an inductor from the junk box[1]. For the maximum performance, it's best to use a more analytical approach as outlined in most switching converter data sheets.

The Output Capacitor

Now what about the output capacitor? Its purpose is to filter the inductor's ripple current and make a nice clean dc voltage on the output. So it needs to store enough charge to hold the output voltage constant over the inductor's charge-discharge cycle. Another way of looking at the output capacitor is that it bypasses the inductor's ripple current to ground, so its impedance (its reactance X_c plus its equivalent series resistance, ESR and its equivalent series inductance, ESL) must be low enough that the ripple current times the impedance is less than the desired output voltage ripple:

$$V_{Ripple} = I_{Ripple}[ESR - j\frac{1}{8fC} + j4f(ESL)]$$

where f is the switching frequency. Note that the factors 8 and 4 in the above equation would normally be 2π for sinusoidal waveforms but the ripple current here is triangular[2]. When the switching frequency is in the 10s or 100s of kilohertz or higher, relatively small values of capacitors can be used, but they need to have low ESR. For example, the reactance of a 22μF capacitor at 300 kHz is about 25 mΩ. If the ESR is also about 25 mΩ, the total impedance is about 35 mΩ. Monolithic ceramic capacitors are a good choice, as they often have ESRs of 10 mΩ or less. But be careful to choose capacitors with X7R or X5R dielectrics, as the high dielectric ceramics such as Y5V or Z5U have very poor temperature and voltage coefficients and can lose a significant portion of their rated capacitance over temperature and at their rated voltage. Even the X7R and X5R capacitors have poor voltage coefficients and can lose up to 80% of their capacitance at their rated voltage. So

[1]Notes appear on page 9.

it's a good idea to use a capacitor with a rated voltage several times higher than the supply's output voltage. This also improves the capacitor's reliability.

The Input Capacitor

In Figure 1 the input voltage is shown as an ideal voltage source with infinite current capacity and no resistance. All real voltage sources have a finite current capacity and some finite resistance as shown in Figure 3 so an input capacitor is used to store charge and smooth out the current drawn from the source. When S1 is closed, the input current is the inductor current that we've already seen is an increasing ramp as shown in Figure 2. The average value during the time S1 is closed is the same as the output current. When S1 is open, the input current is zero. So the input current to the supply can be approximated by a rectangular pulse of magnitude equal to the load current and duty cycle equal to the duty cycle of S1. The dc component of the input current is the average of this rectangular pulse and is simply S1's duty cycle times the load current[3]. The capacitor supplies the input current's ac component and the RMS value of the input capacitor current is $I_{RMS} = I_{LOAD}\sqrt{D(1-D)}$ where D is the converter's duty cycle (S1's duty cycle). (This expression or variations of it[4] shows up without explanation in most switching regulator data sheets and application notes. See the sidebar for the derivation.) Strange as it may seem, the input capacitor is chosen based on its ESR and current rating rather than its capacitance value. As with the output capacitor, low impedance at the switching frequency is what's important to minimize the input ripple current ripple that the source sees. Since monolithic ceramic capacitors have ESRs on the order of 10 milliohms or less, choosing a capacitance that yields a reactance of about that value or less is a good start. To further minimize input ripple current an inductor can be inserted before the input capacitor.

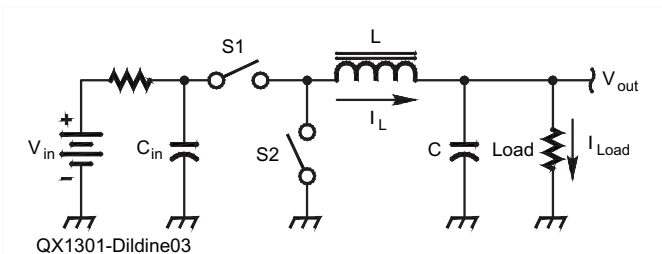## The Switches

Well this is all fine, but where do we

Figure 3 — Including an input capacitor.

get a couple of switches for S1 and S2 that are lossless and can be thrown at a several hundred kilohertz rate? The answer is modern power MOSFETs and Schottky diodes. While they're not completely lossless, modern power FETs come close with typical on resistances on the order of 10 mΩ or less. Replacing S1 in Figure 1 with a MOSFET and S2 with a Schottky diode we have the circuit of Figure 4.

Q1 takes the place of the switch S1 and D2 simply steers the inductor's discharge current to the load. D2 is analogous to a rectifier diode. Q1 is turned on and off at the switching frequency by a genie that lives in the magic box and controls the duty cycle while watching a voltmeter connected to *Vout*. But notice that now during the inductor discharge time the diode is in series with the load so the voltage drop across the diode represents power lost. In high current applications or low voltage applications where the diode drop is a significant portion of the output voltage, we may not be able to tolerate the power lost in the diode. If we're willing to add a little complexity and give the genie in the magic box more work to do, we can replace the diode with another MOSFET as shown in Figure 5.

Here switch S2 is replaced by Q2 which is turned on and off by our genie synchronously with Q1. It is analogous to the synchronous rectifier sometimes used in the old vibrator power supplies and so this circuit is referred to as a synchronous converter. (*Synchronous* refers to the two MOSFETs being synchronized to each other, not to some external frequency reference.) Since modern MOSFETs have very low on resistance, minimal power is lost through Q2 during the inductor discharge period.

The Controller

That's all well and good, but where do we find a magic box with a genie inside to do our bidding? The answer to this one comes in the form of integrated switching supply controllers available from all the major IC companies. Even the simplest controllers include an oscillator to generate the switching frequency, all the necessary FET drive circuitry, and a feedback loop to regulate the output voltage by controlling the FET duty cycles. Some also include the FETs and some even include the inductor.

Real-World Components

So far we've assumed either perfect ideal components or at least they're pretty good. How do real-world components affect efficiency? As mentioned before, there will be loss in the inductor due to its dc resistance. The inductor will also dissipate power due to core losses which will increase with higher switching frequency. In a non-synchronous
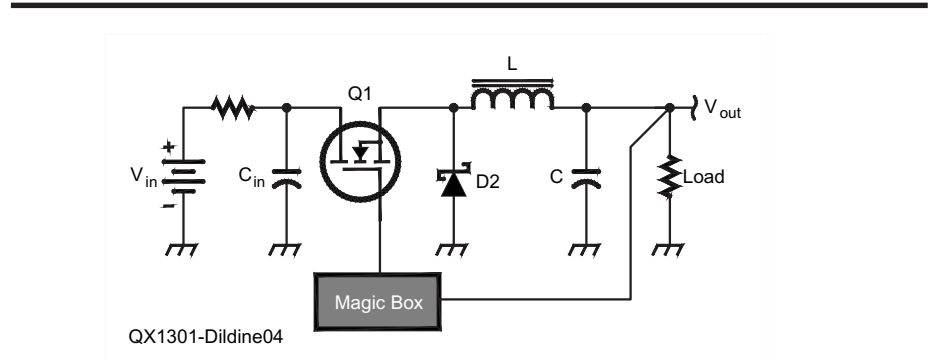
Figure 4 — A buck converter with a MOSFET and diode rectifier.
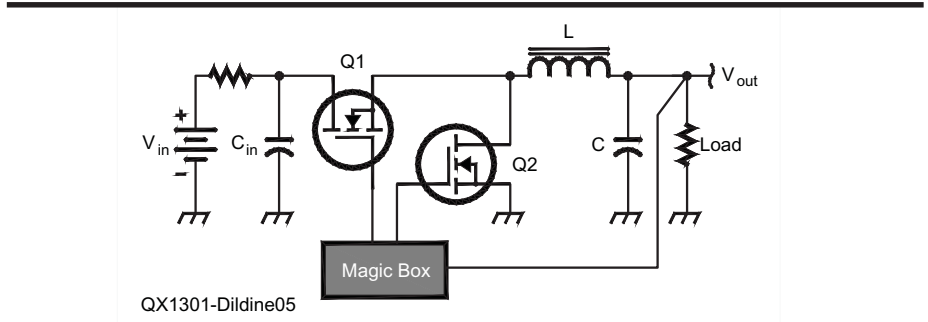
Figure 5 — Adding a MOSFET synchronous rectifier.

converter using a Schottky diode in the place of switch S2, power is lost due to the voltage drop in the diode. As mentioned above, this can be mitigated at the cost of extra complexity by replacing the diode with a MOSFET in a synchronous converter.

The power lost through the series FET, Q1 is due not only to its dc on resistance, but also to transition losses caused by pumping current into and out of its Miller capacitance. So while it's best to select a low resistance FET for Q2, you need to balance the FET's Miller capacitance against its on resistance for Q1. Most manufacturers of integrated switching controllers give a good explanation of FET losses, including the effects of Miller capacitance, in their data sheets and application notes.

Further Resources

I've only scratched the surface of buck converters, but hopefully some of the mystery has been removed and I've shown that the basic operation isn't rocket science. But as always, the devil is in the details and although it's relatively easy to get a simple buck converter working with the modern controllers, optimizing the design for maximum efficiency, minimum ripple, and all the other parameters that make a high performance circuit can be complex.

Fortunately there are a lot of resources

that can help the designer. First and foremost are the data sheets of the various controllers that are available. Even carefully reading the data sheets for parts you may not consider can yield a lot of useful information. Various manufacturers offer detailed application notes describing both specific parts and generic examples of switching supply design. Particularly well written are those by Jim Williams of Linear Technology Corporation[5].

There are several simulation tools that are quite useful for switching supplies and circuit analysis in general. *LTSpice*[6] is a free general purpose SPICE simulator with particular emphasis on efficient simulation of switching supplies built around Linear Technology's controllers. Analog Devices also offers several design and simulation tools for their controllers[7] although they are more specialized. And Texas Instruments offers *TINA*[8], another free SPICE simulator.

Notes

[1]Williams, Jim, *Step-Down Switching Regulators*, Linear Technology Corporation, Application Note 35, p. 22.
[2]Analog Devices Inc., *ADP1829 Data Sheet*, p. 16. (**www.analog.com/static/imported-files/data_sheets/ADP1829.pdf**)
[3]In this regard, a buck converter can be thought of as a "dc transformer" where the turns ratio is converter's duty cycle. In the

Input Capacitor RMS Current

So where does that expression for input capacitor current come from? Let's take a look. The input capacitor current can be approximated by the trapezoidal pulse shown in Figure A1.

Here we assume the current drawn from the input source is constant and equal to I_{AV}. This is especially true if the input capacitor is preceded by a small inductance to smooth the current variations from the input source. During the time S1 is closed and the inductor is charging, the inductor current is supplied from both the capacitor and the input source. The average inductor current is equal to the load current, I_{Load} so the current drawn from the input capacitor is the difference between the load current and the current from the input supply, I_{AV}. Because current is flowing out of the capacitor, the sign is negative: $I_{CAP} = -(I_{Load} - I_{AV})$. Remember, the input current to the buck converter is simply the load current times the duty cycle, D so,

$$I_{CAP} = -\left(I_{Load} - DI_{Load} \right) = -I_{Load}\left(1 - D \right)$$

When S1 is off, the input capacitor regains the charge lost. No current is flowing into S1 so all of the current from the input source is flowing into the input capacitor: $I_{CAP} = I_{AV} = DI_{Load}$. Now current is flowing into the capacitor so the sign is positive.

The RMS value of the waveform in figure A1 can be calculated by remembering the definition of RMS: the (square) Root of the Mean of the Squares. All this means is that we take the area under each part of figure A1's waveform and square it, then add the results and divide the total by the total time covered by each part of the waveform. Finally we take the square root of the whole thing. It's easier to do it than explain it:

$$I_{CAP}RMS = \sqrt{\frac{\left[-I_{Load}(1-D) \right]^2 D + \left[DI_{Load} \right]^2 (1-D)}{D + (1-D)}}$$

The $[-I_{Load}(1-D)]^2 D$ term represents the area under the waveform when S1 is on and the $[DI_{Load}]^2(1-D)$ term represents the area under the waveform when S1 is off. The denominator $D + (1-D)$, is the time of the whole cycle and is of course just 1. So simplifying the expression for RMS current,

$$I_{CAP}RMS = I_{Load}\sqrt{D(1-D)}$$

Where D is the converter's duty cycle.
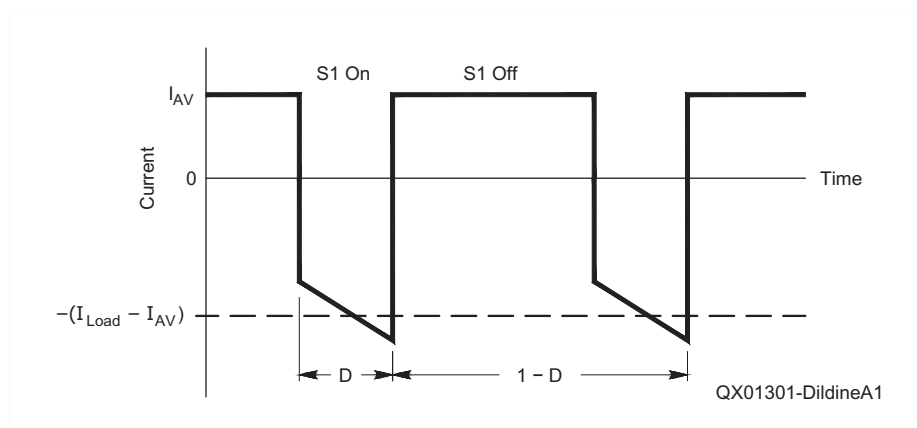
Figure A1 — Input capacitor current.

ideal case, the output voltage is the input voltage times the duty cycle and the output current is the input current divided by the duty cycle.

[4]Seemingly widely differing expressions for the input capacitor's RMS current are given in various data sheets and application notes. By remembering that the converter's duty cycle is the ratio of output voltage to input voltage, $D = V_{out}/V_{IN}$, these differences can usually be reconciled.

[5]See for example Application Notes 25 and 35 by Jim Williams as well as Application Notes 19 and 44 by Carl Nelson, all published by Linear Technology Corporation.

[6]**www.linear.com/designtools/ software/#LTspice**

[7]**www.analog.com/en/power-management/ products/cu_eval_pmp_tools/fca.html**

[8]**www.ti.com/tool/tina-ti**

First licensed in 1957 as WN6SFH, Bob Dildine has been primarily interested in the technical aspects of Amateur Radio. He received a BSEE degree from San Diego State College in 1967 and an MSEE from the University of California at Berkeley in 1973. He then joined Hewlett-Packard as a research and development engineer where he has done precision analog design on microwave synthesizers and network analyzers. Bob retired from Agilent Technologies in 2003 and continues to enjoy designing and building Amateur Radio projects, as well as restoring vintage electronic equipment. **QEX-**

John M. Franke, WA4WDL

4500 Ibis Court, Portsmouth, VA 23703 **jmfranke@cox.net**

Double Balanced Beat Indicator

Build a better indicator for precision measurements.

An article in the September/October 2009 issue of *QEX* magazine, "Atomic Frequency Reference for Your Shack," by Bob Miller, KE6F, contained an interesting feature best described as a double balanced beat indicator. It seemed to be an ideal inexpensive passive instrument for quickly comparing two oscillators. I had long considered trying such a circuit. On close examination, I believe I discovered an error in that the author used a regular single-polarity meter versus a zero-center scale meter to monitor the output signal.

The meter in the article should work well over a span of 180° of phase error change, but for the second 180° the meter needle would be pinned against the zero stop. I did a search on the ARRL website but could not find reference to any errata or feedback concerning the article. So, I sent an e-mail to the author asking about the meter. Bob agreed a zero-center meter for the beat indicator would have been best. He had done what I often do and used what he had available in his junk box.

Building a Better Indicator

I breadboarded a beat indicator using a Mini-Circuits model ZAD-8 double balanced mixer and a 50-0-50 microammeter. The ZAD-8 is rated for 500 Hz to 10 MHz.

I compared a repackaged military surplus O-1267A/USM-207 1.000000 MHz timebase oscillator to a Hewlett-Packard model 3336B signal generator. I quickly discovered the time base inside the signal generator was 14 Hz high. I next used a HP model 10811 crystal oscillator as an external reference to drive the signal generator. The difference between the HP 10811 and the O-1267A/USM-207 was only 0.08 PPM. The measurements were quick and fun. Then and there I knew I needed a beat indicator as a
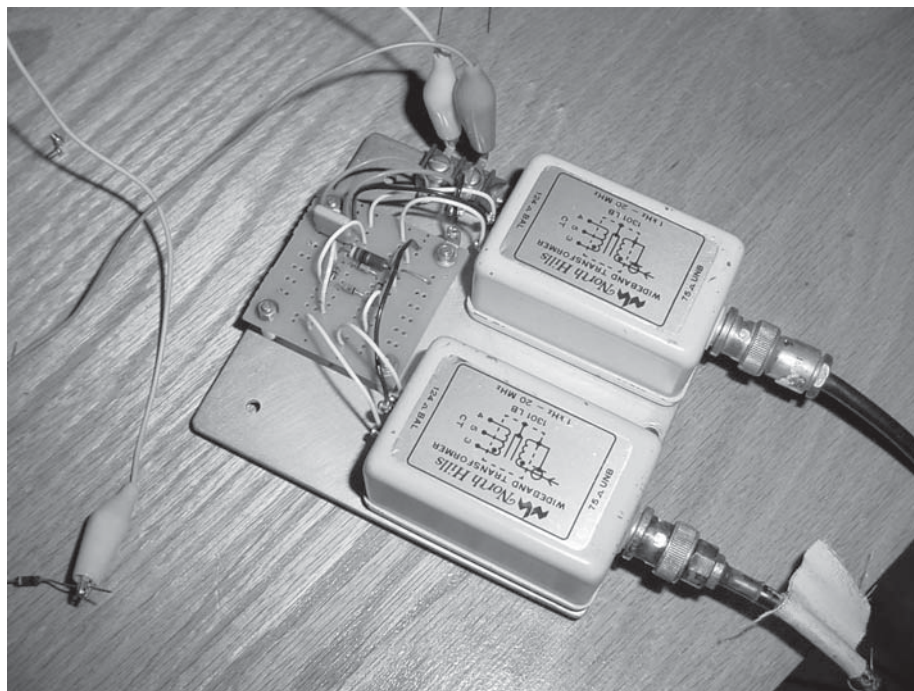
The prototype double balanced beat indicator with North Hills baluns.

permanent part of my work station.

At first, I considered making two beat indicators; one covering 5 kHz to 1 MHz and the other covering 1 MHz to 20 MHz. Then, I found and took advantage of a buy-it-now option on an auction website for four North Hills' model 1301LB baluns covering 1 kHz to 20 MHz for less than $25, including shipping. The baluns came in shielded cases with BNC unbalanced input and terminal post balanced output connections. The baluns are unique in that the balanced output windings have a center tap. Using just two of the

baluns would allow me to make one double balanced mixer that would cover both frequency ranges.

I prototyped a double balanced mixer using the North Hills baluns and a quad array of 1N5711 Schottky diodes (see Figure 1). The double-balanced mixer output port is terminated with a 47 Ω resistor in series with a non-polarized 2.2 µF capacitor in parallel with a 0.001 µF capacitor. I added a RC low-pass filter consisting of a 150 Ω resistor in series and a shunt capacitor of 0.0012 µF between the output port and the
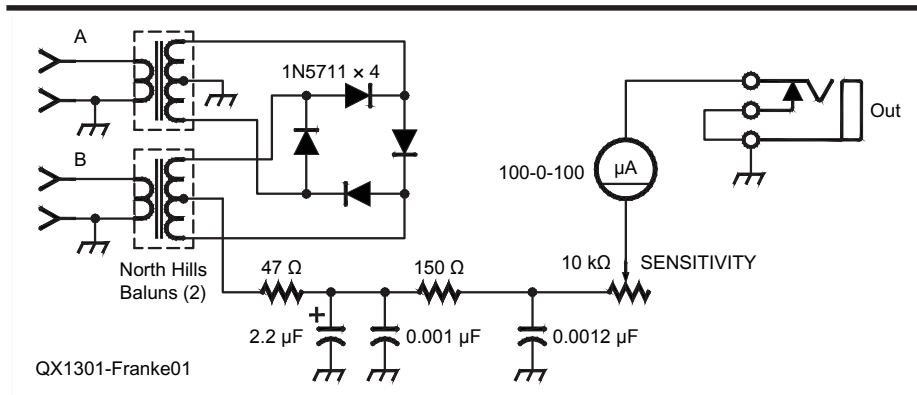
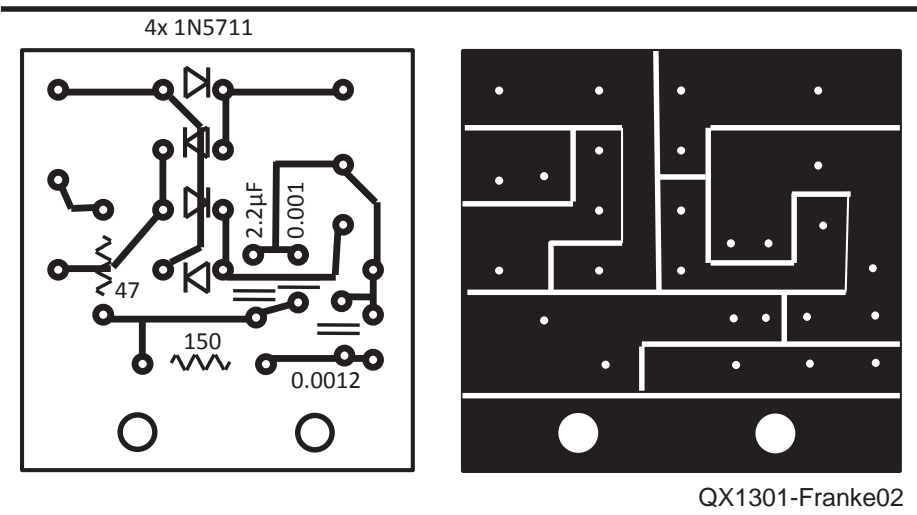Figure 1 — Schematic of the Double Balanced Mixer Beat Indicator.

Figure 2 – The evolution of the circuit board layout from contour to block design.

meter. I mounted everything, including a prototyping circuit board and a two-terminal barrier strip, on an aluminum plate, approximately 4 inches square. I assembled everything together in less than an hour. I did try a quad array of 1N4148 diodes and the results were essentially indistinguishable from the 1N5711 diodes.

If you do use a commercial double balanced mixer, make sure the IF port is dc coupled. For maximum sensitivity, the meter should have a low internal resistance. Most 100-0-100 microammeters have an internal resistance of approximately 1 kΩ. For the same reference input level, the measured full-scale sensitivities with meters having internal resistances of 1350 Ω and 65 Ω were -7 dBm and -20 dBm respectively. I stayed with the 1350 Ω movement because the 65 Ω meter was already being used in another project.

I tested the final circuit at 10 kHz and at 10 MHz. For both tests, the reference frequency was generated with the Hewlett-Packard model 3336B signal generator. For 10 kHz, I used the 10,000.000 Hz output from a Jupiter OEM GPS receiver. For 10 MHz, I used the 10.000000 MHz output from a FE-5860A rubidium frequency standard. Everything worked fine. The tests verified that the system works over the full range of anticipated frequencies.

Finishing Touches

I laid out the design for a dedicated printed circuit board for the final version of the beat indicator by placing the parts on a piece of perforated phenolic. Removing the parts one by one, I marked the holes with a marker pin. I cut out a 2-inch square piece of 1/16-inch thick double-sided G-10 glass-epoxy printed circuit board stock. Then I clamped it to the perforated template and drilled the

holes for the parts.

With the holes drilled, I evolved the trace layout from contoured lines into a design using simple block areas of copper separated by straight line borders (Figure 2). I laid down strips of drafting tape to mask the board where I wanted to etch away the copper and then spray painted both sides of the board. Masking the ground plane side not only provided improved shielding, it saved on etchant and reduced the time required for etching. After the paint dried, I removed the drafting tape strips and etched the board in ferric chloride. I populated the circuit board and used it to replace the prototype circuit board, fastened it on the test frame using two small right angle brackets, and completed the wiring to the baluns and barrier terminal strip.

To house the instrument, I salvaged the cabinet from a defunct Hewlett-Packard 350D attenuator. I mounted the plate carrying the two North Hills baluns and the double balanced mixer inside the enclosure. I found a nice 10 kΩ sensitivity control and dial plate from a military surplus TS-117/GP wavemeter. I drilled more than 70 holes in the front panel of the case and then used a half round file to smooth the edges of the main hole for the panel meter.

I added two handles on the front panel and four rubber feet to the bottom of the case. I mounted two BNC female-to-female feedthrough connectors on the front panel and used two BNC pig tail jumper cables to connect the panel connectors to the baluns.

Next, I filed round two extra D-shaped holes that were on the front panel and are now filled with a plug and a closed-circuit 1/4-inch phone jack. The jack provides access to the output signal for observation on an oscilloscope or to drive a chart recorder.

I cleaned and spray painted the front panel antique white and the two handles flat black. Dry transfer label decals were added to the beat indicator faceplate and given a clear protective overcoat of lacquer. I printed the beat indicator schematic on a small piece of sticky back paper and pasted it inside the enclosure. This is an action I am trying to take more often because it is very helpful if I ever have to troubleshoot the instrument, or if I pass it along to a new owner.

Conclusion

The instrument has proven itself to be very helpful for testing FE-5680A rubidium frequency standards, as well as Jupiter OEM GPS satellite receivers and experimental GPS disciplined crystal oscillators. I saw an e-mail on a newsgroup wherein John C. Roos, K6IQL, mentioned the fact that he uses a double balanced mixer to check his systems against a rubidium standard prior to participating in on-the-air frequency measure-
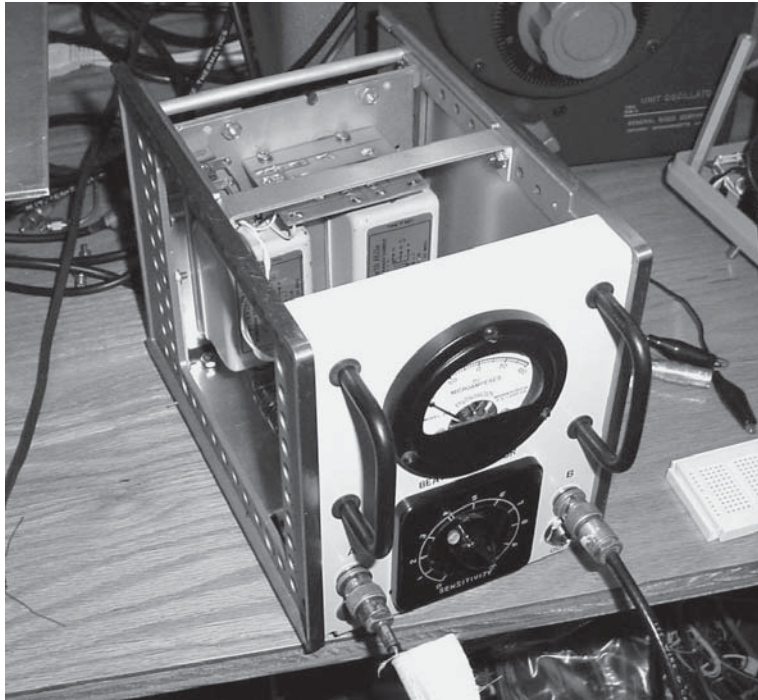
The completed 1 kHz to 20 MHz double balanced beat indicator.

ment tests (FMTs). Bob Miller, KE6F, used his embodiment to adjust his rubidium standard against a GPS-steered rubidium for a 100 second beat. Using 10 MHz sources, one beat in 100 seconds corresponds to a frequency difference of 1 PPB (one part per *billion*).

John earned his Novice license in the early '60s and currently holds an Amateur Extra with the call sign WA4WDL. He also earned his First Class and now General Radiotelephone Operator License with Ship Radar Endorsement. While in college John worked as a transmitter engineer at two AM broadcast stations.

His educational history includes AAS, BSEE and MS, Physics degrees. John retired from NASA after more than 30 years of service.

John currently works as an export control consultant and volunteers as a docent at the Virginia Air and Space Center. His interests include electronic warfare, microwaves, VLF and precision timing. He is the inventor or co-inventor on three US patents. John has also authored or co-authored 125 professional and Amateur Radio publications. QEX–

Richard Harris, G3OTK

10 South St, South Petherton TA13 5AD, United Kingdom **rjharris@iee.org**

The Drive Level Sensitivity
of Quartz Crystals

G3OTK investigates why the motional resistance of some quartz crystals changes with drive level and may be the cause of intermodulation in crystal filters.

Until recently I had thought of quartz crystals as being linear devices, at least until driven to levels at which damage could occur. From time to time I have read articles that mention intermodulation products produced by crystal filters but such products can only be generated if there is non-linearity within one or more of the crystals that comprise a filter. The question therefore arises as to the source of such non-linearity and whether every crystal exhibits this effect. If it is measureable, then crystals could be screened and those with minimum non-linearity selected for use in a filter.

Research has been undertaken over the past few decades into non-linearity effects in quartz crystals and the cause has been found to be variations in the "motional" (or series) resistance of crystals with drive level. The motional resistance is the result of losses within the quartz element itself, through the supporting structure and to the air (unless in an evacuated holder).

Many experiments have been described over the years. The motional resistance of some crystals may be low at low drive levels, but as the current increases it switches to a much higher value before falling back as the current increases still further. There may also be hysteresis, with changes in the motional resistance being dependent upon whether the current is increasing or decreasing.

The motional resistance may take another form, being high at low currents but decreasing with increasing drive level, finally stabilizing at a much lower level. Such a characteristic may prevent an oscillator from starting, although it may work satisfactorily if shocked into oscillation, perhaps by a voltage transient.

These motional resistance characteristics have been reported to be strongly temperature dependent and unpredictable, with some crystals from the same batch exhibiting the
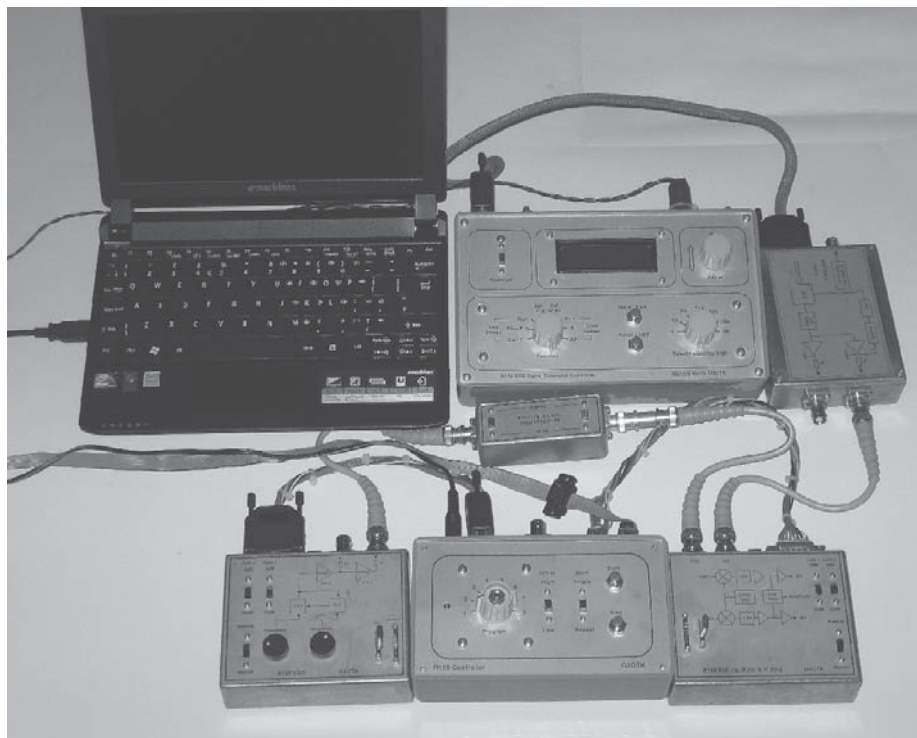
The equipment used to make the motional resistance measurements. Left to right along the bottom row is the VXO/phase detector, controller and mixer/low frequency IF amplifiers.

effect while others do not. It may lie dormant only to appear some time later. This phenomenon has been given various names, such as "starting resistance", "sleeping sickness", "second level of drive" and (the term I will use) "drive level sensitivity" or DLS. An Internet search using these phrases as search strings soon finds references to this effect.

Very little information about the DLS of crystals has been published in Amateur Radio magazines. This article reviews some of the papers published in professional journals, describes the equipment that I made to measure this effect and presents some of the results that I obtained. All of the crystals were AT-cut parts, which operate in the thickness-shear mode, and were purchased new from component distributors in the UK.

Review of Published Papers

One of the first investigators was Bernstein[1], who used a modified crystal

[1]Notes appear on page 21.

impedance meter to deduce the motional resistance characteristics of crystals. He concluded that those that showed drive level sensitivity had particles bound to the surface of the resonator by a contaminating film of oil or other sticky material. He found that changes in the motional resistance exhibited hysteresis because these changes depended upon whether the drive level was increasing or decreasing. His explanation was that very small particles of quartz or other material were closely coupled to the resonator at low crystal currents. As the drive level increased the contaminating film was stretched beyond its elastic limit and the particles were then no longer coupled to the quartz surface. When the drive level was reduced, a level was reached at which the film became effective again. Bernstein found that washing the crystal element in acetone or alcohol removed the particles and other contaminants and the motional resistance became independent of the drive level.

Nonaka, Yuuki and Hara[2] found that the motional resistance of some crystals switched between widely different values at a critical current, which depended upon whether it was increasing or decreasing. Their explanation was that at low drive levels the static friction kept the particles bound to

the surface of the crystal. However, at a critical level of drive the momentum of particles exceeded the static friction and moved relative to the surface of crystal, absorbing additional energy and causing an increase in the apparent motional resistance. At this critical drive level, the mass loading of the crystal due to the particles decreased and gave a step increase in the series resonant frequency. As the drive level increased further, these frictional losses decreased as a proportion of the total losses and so the motional resistance decreased.

Dworsky and Kinsman[3] proposed a model that considered the crystal as a spring-mass system with a particle free to move in a crevice, perhaps a scratch on the surface of the crystal. The particle would be expected to experience inelastic collisions with the crevice walls at high drive levels. This situation could be described by an equivalent electrical circuit with some non-linear components. They showed that approximate solutions to the equations for this circuit could yield both forms of motional resistance variations described above, depending on the values used in their model.

Quartz crystal manufacturing techniques have evolved to such an extent that crystals at popular frequencies can now be purchased

for less than 25 cents in quantity, a remarkably low price for a precision component. Has the drive level sensitivity of motional resistance been eliminated by improved production methods? To find out, I made a manually operated test jig to measure the motional resistance and series resonant frequency of crystals at various drive levels[4]. I could only make a very limited number of measurements over an 80 dB range of crystal currents and consequently I may have missed important detail but I did find crystals that showed anomalous motional resistance. Making these measurements was time consuming and this manual method was not suitable for assessing batches of crystals or for making measurements under temperature controlled conditions.

The next stage of my investigation was to design and build equipment that could make a large number of measurements automatically. This article describes this equipment and some of the results that I obtained.

The Measuring Equipment

The block diagram of the measuring equipment is shown in Figure 1. The measurement method is sensitive to the harmonic content of the test signal and several circuit
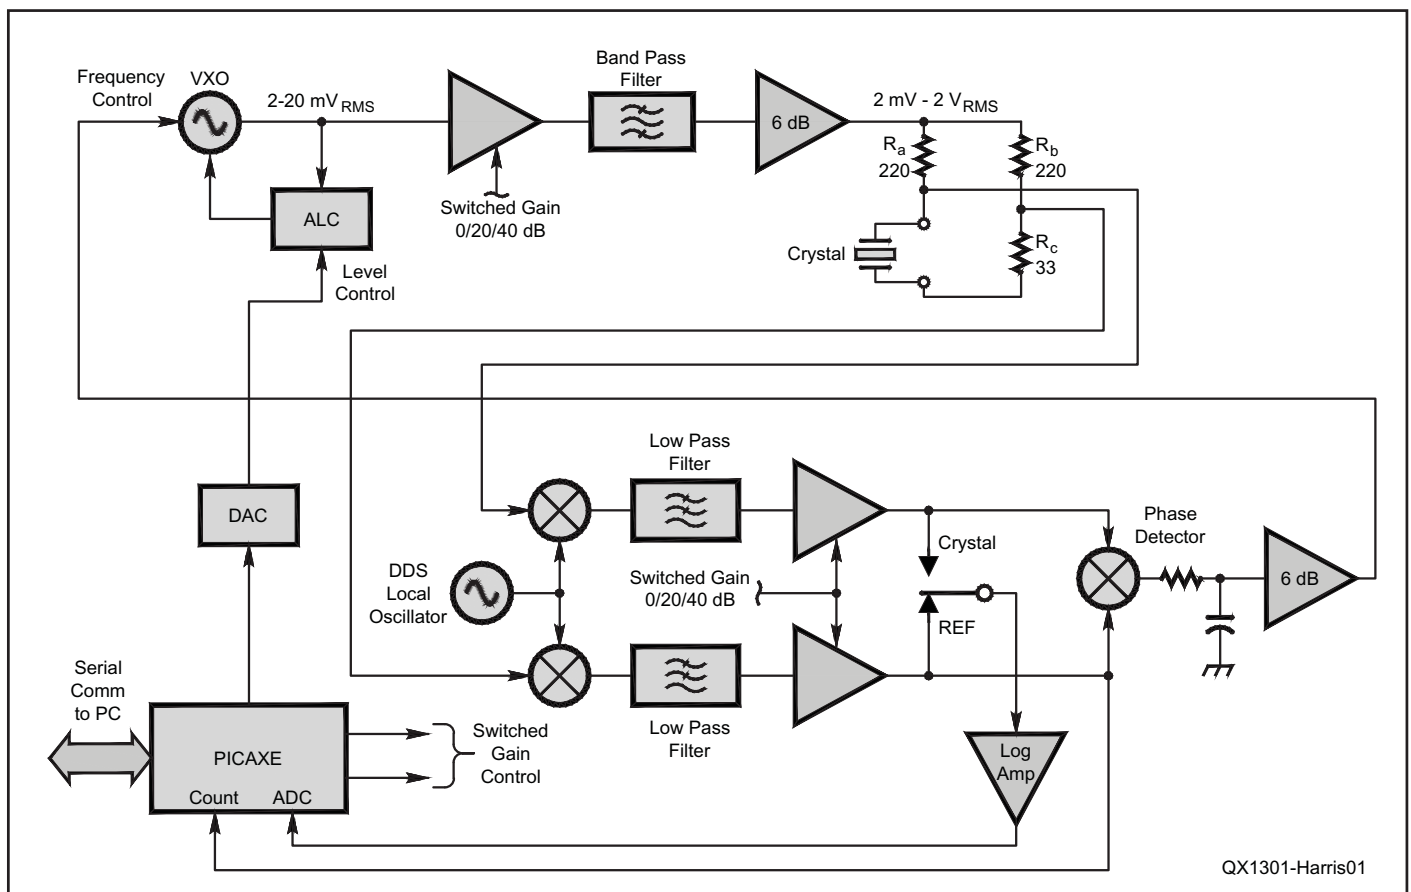
Figure 1 — Measurement equipment block diagram.
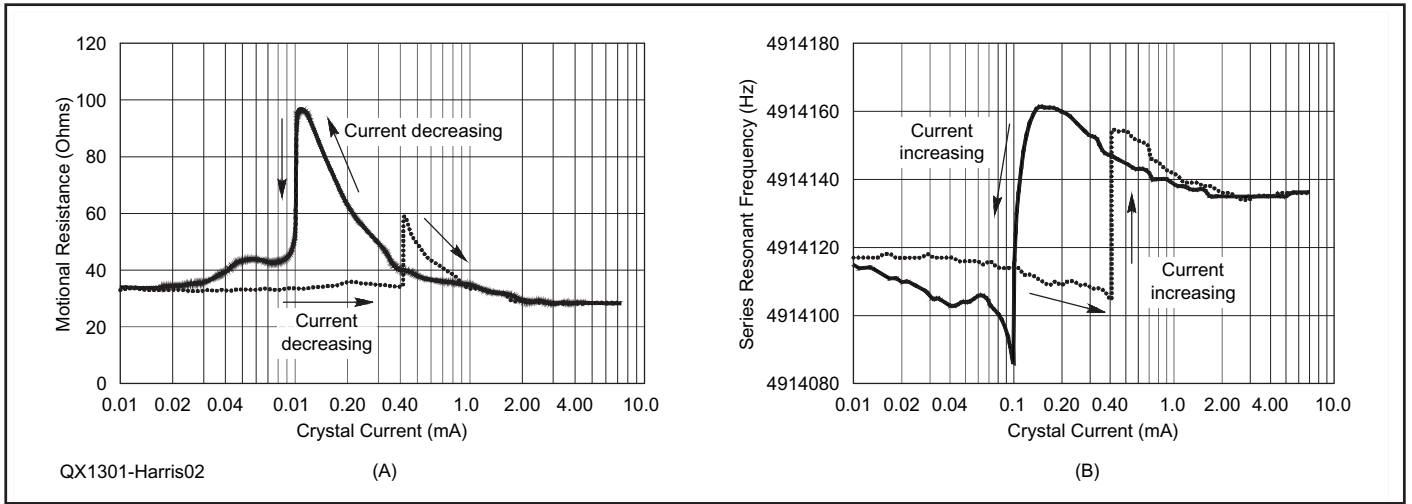
Figure 2 — (A) The motional resistance and (B) series resonant frequency of crystal #6 at 25 C. These show abrupt switching and hysteresis.

techniques are used to keep the harmonics to an acceptably low level. The test signal is generated by a VXO that is locked to the series resonant frequency of the crystal under test. Automatic Level Control (ALC), using an AD8307 log amplifier as a detector, limits the amplitude of oscillation to a maximum of 20 mV, so that the oscillator transistor is operating in a substantially linear fashion with a low harmonic content. By controlling the comparison voltage of the ALC circuit, the amplitude can be reduced to 2 mV, giving 20 dB of variation.

The output of the oscillator is amplified by two LM6181 wide band current mode amplifiers in series, each with selectable gains of 0 and 20 dB. The signal is passed through a
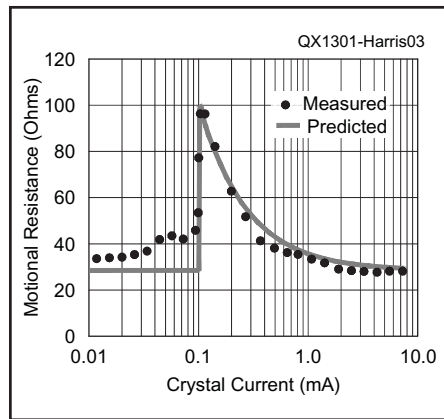
Figure 3 — Comparison of measured and predicted response of crystal #6 at 25 C.

band pass filter to reduce the harmonic content still further. The filter is buffered by a third LM6181 amplifier, which has a gain of 6 dB to compensate for the insertion loss of the filter. The controllable amplitude of the oscillator and the switchable gain stages give a total test signal variation of 60 dB.

The current through the crystal under test is determined by the voltage at the output of the 6 dB amplifier, the series resistor Ra and the motional resistance. It is adjustable over the range of about 8 μA to 8 mA. At the series resonant frequency of the crystal, the voltage at the junction of Ra and the crystal will be in phase with the voltage at the output of the 6 dB amplifier. The potentiometer network Rb and Rc provides a sample of the
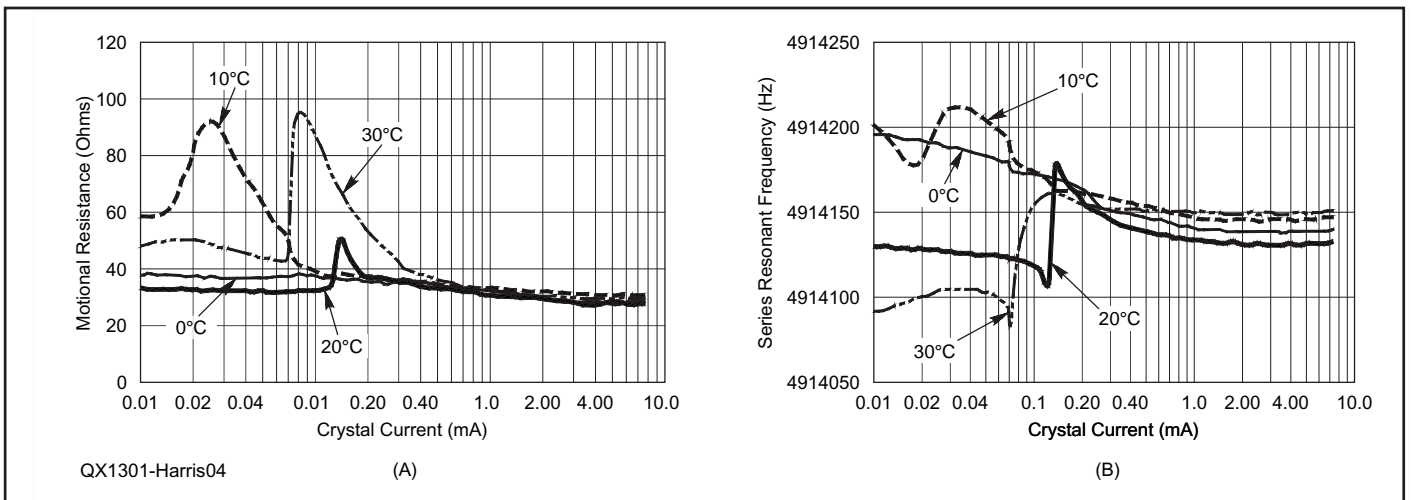
Figure 4 — (A) The motional resistance and (B) series resonant frequency of crystal #6 measured at various temperatures in an environmental chamber. These are for the current sweeping from high to low only.

amplifier output that is at a level similar to the voltage across the crystal at resonance. These two voltages are heterodyned down to a low intermediate frequency of one or two kHz by two MC1496 mixers, operating in multiplicative rather than switching mode. The local oscillator is a home-made DDS signal generator.

The outputs of the two mixers are filtered by identical three pole RC low pass filters and amplified by op-amps with selectable gains of 0, 20 or 40 dB. The phases of the two signals are compared by a phase detector consisting of two zero crossing detectors, using a LM393 dual comparator, connected to the D and CLOCK inputs of a D-type flip-flop. The output of the phase detector will be either high or low depending on which signal leads the other. An RC low pass filter averages the signal from the phase detector and controls the frequency of the VXO via a buffer amplifier.

The outputs of the two intermediate frequency amplifiers are sampled in turn by an AD8307 log amplifier. The motional resistance of the crystal is calculated from the ratio of these voltages and the values of Ra, Rb and Rc,

The controller uses a PICAXE micro-processor, which is a PIC processor with an integral "PICAXE Basic" interpreter and is very easy to program. The comparison voltage for the VXO ALC is generated by an 8-bit DAC connected to the processor using the I²C bus. The test signal can be increased or decreased in steps as small as 0.8 dB (about 10%), giving up to 78 measurements over a 60 dB range. The output of the log amplifier is digitized by one of the 10 bit ADCs of the processor. The PICAXE is not capable of measuring the frequency of the VXO directly but it does have a timed counting facility that is usable up to frequencies of a few tens of kilohertz. The output frequency of one of the intermediate frequency stages is counted and subtracted from the local oscillator frequency to give the frequency of the VXO. The counting period can be programmed to be 1 second (1 Hz resolution) or 10 seconds (0.1 Hz resolution).

The processor sends data consisting of gain settings, VXO ALC voltage, log amplifier output voltages and frequency in CSV format to a laptop. The data is loaded directly into an *Excel* spread sheet. Additional columns in the spread sheet are used to calculate the crystal current, motional resistance and series resonant frequency. The formulae for these calculations are sent as part of the formatted data set so that the calculated values automatically appear during the test. The controller can be set to sweep the current up and down continuously to record changes that may occur over periods of hours.

The log amplifier has a nominal conversion factor of 25 mV/dB. As the level of the signal changes, this factor varies very slightly and appears as small ripples on the graphs of motional resistance versus crystal current. These ripples repeat every 20 dB during a sweep of crystal current. For crystals that have a constant motional resistance, an error correction function can be derived by making a sweep with the nearest value fixed resistor fitted instead of the crystal.

The accompanying photo shows the equipment that was used to make the measurements. To the right of the laptop is the controller for the DDS signal generator, which is used as the local oscillator, with the DDS RF unit next to it. Bottom left is the VXO, RF amplifiers and phase detector module, with the controller to the right. On the far right is the crystal measurement and dual low frequency IF module. In the middle is a switchable 5 or 10 MHz band pass filter, which reduces the harmonic content of the signal to a negligible level.

The crystal measurement module has two crystal holders made from strips of turned-pin IC socket, one fitted to the outside of the unit for ease of changing crystals, the other inside for use when the equipment is used in the environmental chamber. Crystals must be protected from air currents when small changes of frequency are to be measured.

This equipment has much in common with a vector network analyzer, but with a crucial difference. A VNA sets the frequency and measures amplitude and phase, whereas this equipment sets the phase (to zero degrees) and measures the frequency and amplitude.

An Example of Motional Resistance that Switches

The first crystal that we will examine is a low cost 4.9152 MHz unit in a HC49/U package. It is important that crystals can be individually identified and this item was #6 from a batch of 25 crystals purchased from a component distributor. The current sweep started at about 8 mA and was reduced in steps of about 10% until a current of 8 μA was reached. The current was then stepped back up to the starting value, giving a total of 156 measurements. These measurements were made in an environmental chamber with the temperature set to 25 C. The motional resistance and series resonant frequency plots are shown in Figures 2A and 2B. These show many of the characteristics reported in the literature.

As the current is reduced the motional resistance increases until a peak is reached, at which point the resistance switches to a much lower value, similar to that measured at high currents. At the same time the series resonant frequency switches to a lower value. The motional resistance then remains reasonably
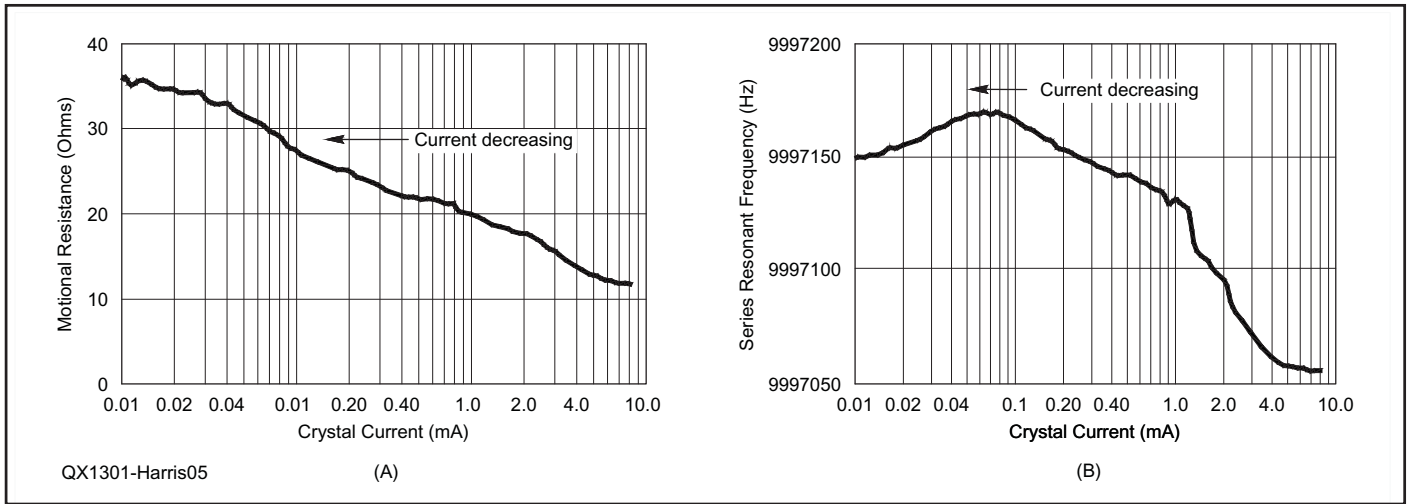
Figure 5 — (A) The motional resistance and (B) series resonant frequency of crystal #20 with weak coupling between particles and the surface of the crystal.

constant as the current is further reduced. When the current is increased from low levels we see similar resistance and resonant frequency plots but shifted to higher currents. These show the hysteresis effect that has been reported by other experimenters.

According to Dybwad[5], the change of frequency with current may be positive or negative depending upon whether a particle is strongly or weakly coupled to the surface of the resonator. This crystal had particles that were strongly coupled to the surface at low currents but less so at high currents.

My explanation is as follows. At low currents the particles are bound to the surface of the crystal by a sticky coating. These particles increase the mass of the crystal and lower the series resonant frequency. As the drive to the crystal increases, a critical current is reached at which the sticky substance is unable to hold the particles at peak acceleration and there is relative movement between the particles and the coating. This movement generates localized heating, reducing the viscosity of the coating and leading to increased relative movement, more heating and a fur-

ther reduction in viscosity. Energy is being dissipated and so this is analogous to an increase in the motional resistance. There can be a big increase in the motional resistance for little or no increase in current above this critical value. The particle is no longer tightly bound to the surface of the crystal and so the mass loading of the crystal decreases and the series resonant frequency increases.

Now consider what happens when the current is well above this critical value and starts to reduce. Less heat is generated but some is still being produced when the critical current is reached that we noted when increasing the current from a low level. The current can be reduced still further until a level is reached at which the sticky coating can once again hold the particle and so the heating ceases, the viscosity increases and the motional resistance switches to a lower value. This is the source of the hysteresis effect.

We can make an estimate of how the
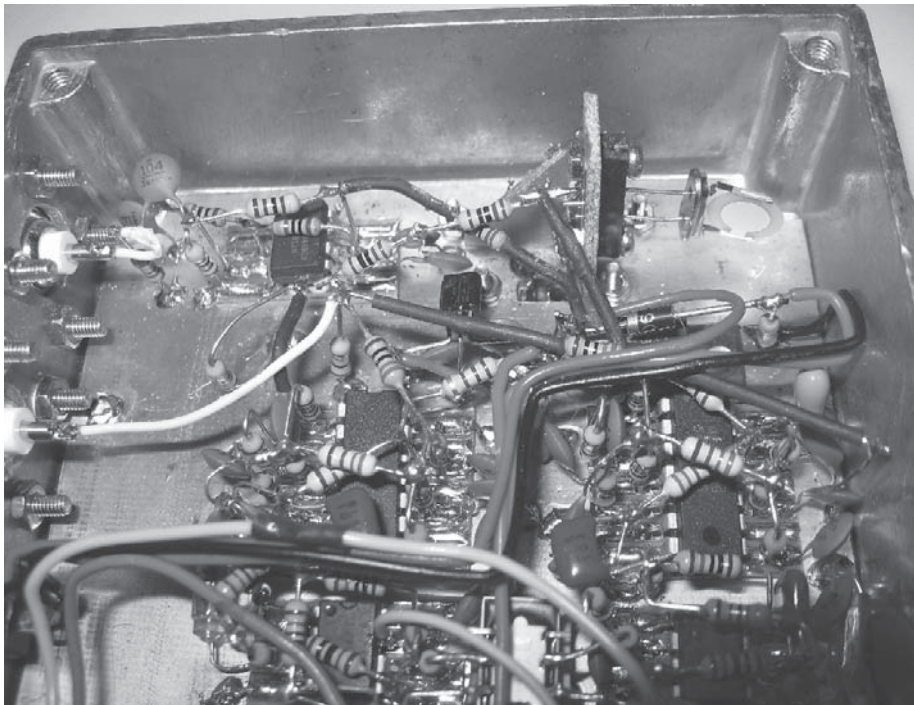
"Manhattan" type construction is used, with dual-in-line ICs surface mounted onto small pieces of strip board fixed to a ground plane. The crystal element is the circular disc at the top right of the photograph.

Figure 6 — The motional resistance after blowing talcum powder at an opened crystal and after a final washing with IPA.

motional resistance varies with current. Let the motional resistance be Rm when the crystal current Ix is greater than the critical current Ic and let it be a constant value, $Rm0$, when Ix is less than Ic. Above the critical current there will be relative movement between the sticky coating and the particle. We will assume that the frictional force on a particle is proportional to its mass $M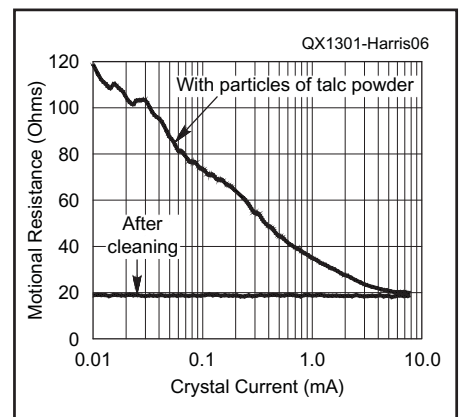$ and that the average relative velocity between the particle and the surface of the crystal is proportional to the current. The average power that is dissipated is given by

$$Ix^2 Rm = Ix^2 Rm0 + K M Ix \quad for\ Ix > Ic$$

where K is the constant of proportionality. Dividing this equation by Ix^2 we obtain

$$Rm = Rm0 + KM / Ix \quad for\ Ix > Ic$$

Below the critical current the second term is zero because there is no relative movement between the particle and the surface of the resonator and so

$$Rm = Rm0 \quad for\ Ix < Ic$$

Nonaka et al[3] give a similar equation but without a derivation. The empirical best fit values for crystal #6 at 25 C is

$$Rm = 25.8 + 7.2 / Ix \quad for\ Ix > 0.1\ mA$$

and

$$Rm = 25.8 \quad or\ Ix < 0.1\ mA$$

where Ix is in milliamps and Rm is in ohms. Figure 3 compares these formulae with the measured motional resistance of crystal #6 at 25 C and shows good agreement.

I found that the motional resistance of such crystals varied over short time spans and that successive current sweeps were seldom the same. It seems probable that particles can migrate over the surface of a crystal, perhaps due to the effects of gravity or to small differential forces caused by variations in the surface tension of the contaminating film.

The Effects of Temperature

The anomalous motional resistance effect has been reported to be very temperature dependent and unpredictable. I wanted to investigate this aspect and fortunately I had access to a temperature controlled environmental chamber. I placed the VXO, crystal measurement and controller modules in the chamber but kept the DDS local oscillator outside to minimize its frequency drift with temperature. I set a series of temperatures from 0 to 35 C in 5 C increments,

allowing at least 30 minutes at each temperature for the equipment temperature to stabilize. I made most measurements with the 4.9152 MHz crystal #6. Some results for decreasing crystal currents are shown in Figures 4A and 4B for the motional resistance and series resonant frequency respectively. At 0 C there is little variation in the motional resistance, although there is a significant change in resonant frequency, showing that the effect is not dormant.

The behavior of motional resistance with temperature is unpredictable. At 10 C and 30 C there is a large increase in the motional resistance, whereas at the intermediate temperature of 20 C there is a much smaller increase, which is at a higher current than at the other two temperatures. I have not shown the motional resistance and resonant frequency plots for increasing crystal currents because these show similar changes but at higher currents.

An Example of Weakly Coupled Particles

Another characteristic type of motional resistance with crystal current is shown in Figure 5A, with the series resonant frequency shown in Figure 5B. This was measured using a 10 MHz crystal designated #20. As the current reduces, the motional resistance rises. In this case the motional resistance depends much less upon whether the current is increasing or decreasing and so only the decreasing current sweep is shown. The resonant frequency increases as the current reduces and this is characteristic of particles that are weakly coupled to the surface of the resonator. One possibility, to be discussed in the next section, is that the particles are bound to the surface by the intermolecular van der Waals force, rather than a surface contaminant.

Dybwad's explanation is that the resonator can be thought of as a mass attached to a spring. A particle can be considered to be a small second mass coupled to the first mass by a second spring. His analysis shows that the dividing line between a resonant frequency that increases and one that decreases with changes of crystal current is when the particle and its motional spring has a resonant frequency that is the same as the crystal.

Inducing DLS

If DLS is caused by particles attached to the surface of the resonator, can we induce it? Virgil Bottom[6] did so and I replicated his experiment. I selected a 4.9152 MHz crystal in an HC49/U package (designated as #49), which showed no significant variation in motional resistance or series resonant frequency, and removed most of the housing
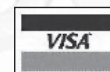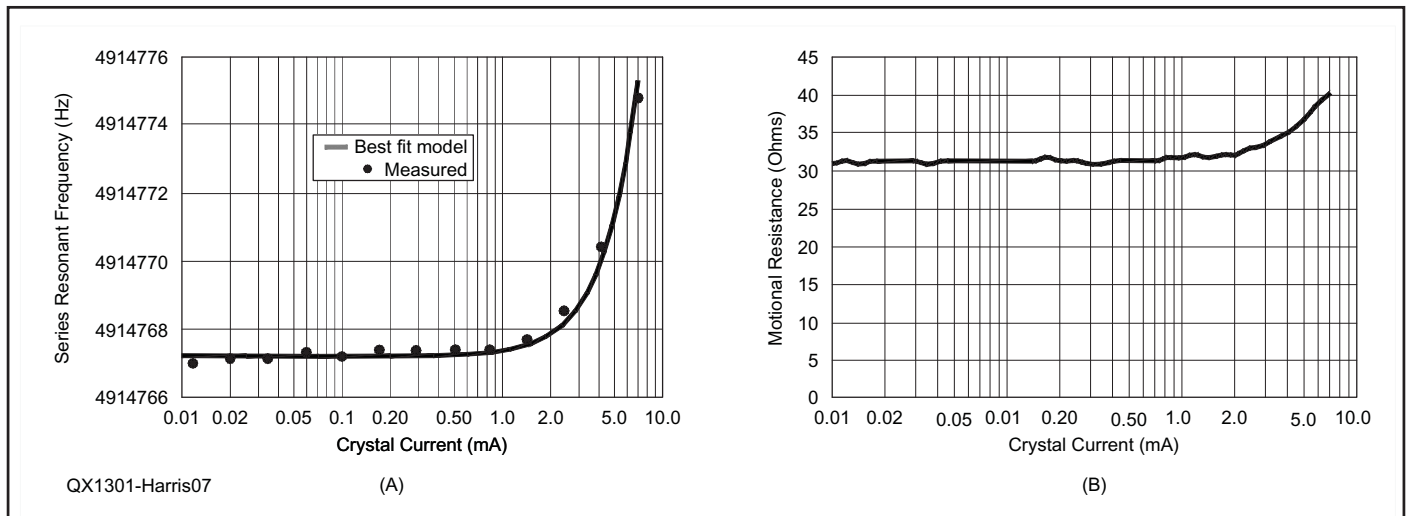
Figure 7 — (A) The series resonant frequency and (B) motional resistance of crystal #5 in a low profile HC49/4H package showing the effects of high drive levels.

using a small grinding wheel. I washed the crystal element with an isopropyl alcohol (IPA) spray to remove grinding debris and allowed it to air-dry because wiping might have contaminated the surface. This crystal was fitted to the measurement module, as shown in the accompanying photo. I checked that it did not show any drive level sensitivity that could have been caused by residual particles from the grinding operation and that the series resonant frequency was unchanged. The motional resistance remained constant at about 18 Ω.

I blew a small amount of talcum powder at the crystal from a distance of about half of a meter and then measured it again. I found that it exhibited strong drive level sensitivity, with the motional resistance increasing from 20 Ω to 80 Ω as the current decreased. After several hours sweeping the current over the 8 μA to 8 mA range, the motional resistance stabilized at about 120 Ω at 10 μA.

I wanted to find out whether the motional resistance was sensitive to temperature and ran a number of tests under temperature controlled conditions between 10 C and 40 C. There was no significant variation and only a small amount of hysteresis, which seems to be a characteristic of weakly coupled particles. To make a comparison sweep without the talc, I washed the crystal element with IPA from a dispenser on my bench. Another sweep at 25 C showed that the motional resistance was now constant at about 22 Ω but the series resonant frequency had dropped by 200 Hz. My conclusion was that the IPA from the dispenser may have been contaminated by prior use and that there was now a coating on the crystal element that increased

the mass very slightly and reduced the resonant frequency. Several more washings with clean IPA raised the resonant frequency closer to the original value and the motional resistance stabilized at about 18 Ω. The motional resistance with the particles of talc and after the final clean is shown in Figure 6. With particles of talc, the resonant frequency increased as the current decreased, much as shown in Figure 5B.

When we looked at how the series resonant frequency changed with current when there was a sticky contaminant on the surface of the crystal, we found that the frequency switched to a lower value as the current reduced. However, in this experiment the frequency increased. This difference is significant. Washing with IPA removed any sticky coating on the surface of the crystal and so the small particles of talc were held by the van der Waals force. The surface of the crystal moved back and forth under shear motion and the bounded particle acted as a small oscillatory system as described by Dybwad.

So how strong is the van der Waals force? It can be very strong and it is the force that allows geckos to climb smooth vertical surfaces. Virgil Bottom calculated that it could be between one and one thousand times the mass of a particle, depending on its size. The smaller the particle, the greater will be the force relative to its weight. He calculated that an excitation voltage of one volt at the series resonant frequency of the crystal would cause peak accelerations of the surface of 1000 G and this would be sufficient to dislodge the smallest particles. Passing a high current through the crystal to shed particles and eliminate DLS has been suggested by others. However, the particles are still within

the housing and so the DLS may return later.

High Crystal Current Effects

I found that the series resonant frequency of crystals almost always increased at the highest currents. According to Parzen[7], the series resonant frequency of an AT-cut crystal increases with current. The change of frequency is not due to the heating of the crystal element because AT-cut crystals have a negative temperature coefficient at moderate temperatures. I saw this effect for both the standard HC49/U and the low profile HC49/4H crystals. According to Parzen:

$$\Delta f / Fo = a\, Ix^2$$

where Δf is the incremental frequency above the series resonant frequency Fo measured at low currents, Ix is the crystal current in Amps and a is the constant of proportionality which, according Parzen, is between 0.02 A^{-2} and 0.05 A^{-2} for AT-cut crystals. I measured several crystals from batches at 4.9152 MHz and 10 MHz and found good agreement with the above formula. The series resonant frequency of a 4.9152 MHz crystal (designated #5) is shown in Figure 7A, together with the empirically determined best fit constants, as per Parzen's formula, of Fo = 4,914,791.1 Hz and a = 0.029 A^{-2}. The small discontinuities at about 0.08 and 0.8 mA coincided with gain switching of the two intermediate frequency amplifier chains and are due to small phase response mismatches of the order of 1 or 2 degrees.

The series resistance increases as the drive increases, as shown in Figure 7B. I have

not been able to find any information about whether there is an established relationship between the motional resistance and high crystal currents.

Conclusions

Anomalous changes in the motional resistance, with associated changes in the series resonant frequency, can still be found in quartz crystals manufactured today. Changes in motional resistance can be large and can occur with a small increase in current above a critical value. Such changes show hysteresis with drive level and may be very temperature dependent. There is evidence that particles may migrate over the surface of the crystal and the motional resistance can change significantly over a short period of time.

Changes in the motional resistance of rogue crystals may be sufficient to prevent an oscillator from starting or may limit the amplitude of oscillation. Such a situation may be puzzling because an oscillator using such a crystal may operate perfectly satisfactorily if shocked into oscillation or may work at some temperatures but not at others.

Crystals in the low profile HC49/4H holder exhibited less anomalous resistance and series resonant frequency variations than crystals in the standard HC49/U holder. The low profile package uses a small rectangular strip of quartz, whereas the standard package uses a quartz disc. This finding probably applies to surface mount crystals. Over-drive effects were observed with low profile crystals but were much as predicted and can be avoided by limiting the crystal current to less than 1 mA.

Selection of crystals for use in filters can be undertaken but this is not an absolute guarantee that the crystal will not suffer from this phenomenon because loose particles in the holder may become attached later.

Notes

[1]M. Bernstein, "Increased Crystal Unit Resistance at Oscillator Noise Levels", *Proc. 21st Annual Frequency Control Symposium*, April 1967, pp. 244-258.
[2]S. Nonaka, T. Yuuki, K. Hara, "The Current Dependency of Crystal Unit Resistance at Low Drive Level", *Proc. 25th Annual Frequency Control Symposium*, April 1971, pp. 139-146.
[3]L. Dworsky, R. G. Kinsman, "A Simple Single Model for Quartz Crystal Resonator Low Drive Level Sensitivity and Monolithic Filter Intermodulation", *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, Vol. 41 No 2, March 1994, pp. 261-268.
[4]R. J. Harris G3OTK, "Variations in the Motional Resistance of Quartz Crystals with Drive Level", *Radio Communication*, Vol. 87 No 12, December 2011, pp. 42-45.
[5]G. L. Dybwad, "A Sensitive New Method for the Determination of Adhesive Bonding between a Particle and a Substrate", *Journal of Applied Physics*, 58(7), 1 October 1985, pp. 2789-2790.
[6]V. E. Bottom, "The Behaviour of Quartz Resonators at Low Drive Levels", *Proc. 5th Crystal Conference*, 1983, pp. 245-249.
[7]Benjamin Parzen, "Design of Crystal and Other Harmonic Oscillators", John Wiley and Sons, 1983.

Richard Harris was licensed as G3OTK in 1961. He received Bachelor and Master Degrees in Electrical Engineering from the University of Bath in the UK. Although he has spent much of his professional life undertaking electronic design, for the last ten years he has been responsible for Quality Assurance, health & safety and environmental management. He is a member of the Itchen Valley Amateur Radio Club.

QEX

Richard C. Jaeger, K4IQJ

2160 Estate Dr, Auburn, AL 36830 **k4iqj@mindspring.com**

Multi-Element End-Fire Arrays of K9AY Loops

An array of switched loop antennas can achieve good performance in environments that do not favor traditional four-square vertical arrays.

In the past I have tried a four-square array of short verticals, but in this area of Alabama we have poor ground conditions (red clay over rock shelves) and the array never really performed up to my expectations. (The short vertical array did seem to perform better after a hard rain.) My best receiving antenna was most often a pair of K9AY loops located at the center of the four-square. The loops appear to be less dependent upon ground conditions.

I wanted to try a four-square array of K9AY loops, but had a great deal of trouble placing the array on my lot because of the position of the house and other esthetic considerations. After reading papers by Breed[1] and Lankford[2], I realized I could achieve similar performance with three-element end-fire arrays. A binomial weighted (1:2:1) three-element array has a pattern that is nearly identical to that of a four-square. In essence, the two side elements of the four-square are collapsed into a single center element in the three-element co-linear array. The center element must then deliver approximately twice the current of each end element.

[1]Notes appear on page 31.
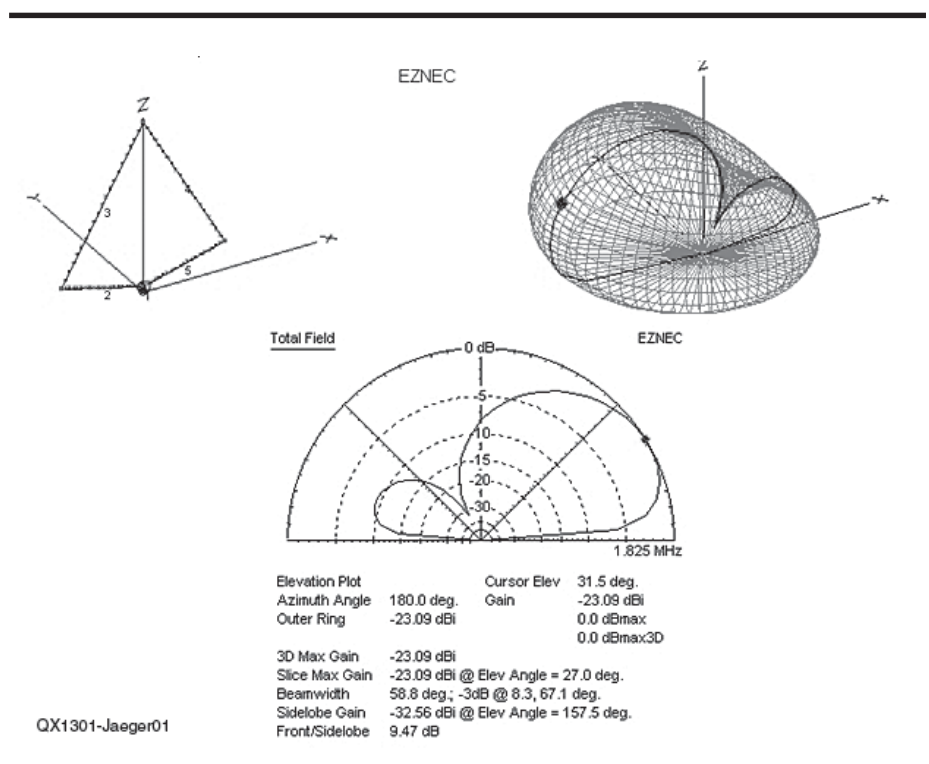
Figure 1 — K9AY Loop - 1.825 MHz. Average Gain: -30.4 dB, RDF 7.4 dB. The loop termination resistor is 510 Ω.

Table 1

Comparison of End-Fire Arrays with a Single Loop – 80 Foot Spacing

	160 Meter RDF (dB)	80 Meter RDF (dB)	160 Meter / 80 Meter Crossfire Phasing (deg)
Single Loop	7.4	7.4	---
2-Element Array	10.5	10.0	-200 / -220
3-Element Array	12.1	11.3	(0,-200,-400) / (0,-220,-440)
4-Element Array	14.6	12.9	(0,-195,-390,-585) / (0,-210, -420,-630)

In the antenna comparisons that follow, I desired maximum rejection of noise, so I chose to optimize the Receiver Directivity Factor (or RDF) proposed by W8JI and discussed in ON4UN's books[3]. Atmospheric noise generally comes in from all directions, and the RDF compares the maximum forward gain of the main antenna lobe to the average gain of the antenna over the above ground hemisphere in dB:

$$RDF_{dB} = G_{for}(dB) - G_{avg}(dB)$$

EZNEC automatically calculates the average gain from its 3D plots, and the forward gain is obtained from the 2D elevation or horizontal patterns.

Basic Array Simulation Results

Before finally deciding upon the three-element array, I performed a number of

simulations with *EZNEC* to compare 2, 3 and 4-element arrays with a single loop, and the results appear in Table 1 and the plots in Figures 1-5. Wires 2 and 5 of the basic loop in Figure 1 are 15.5 feet long, and wires 3 and 4 are 27.5 feet in length. The top of the loop is at 27.5 feet, and the feed point of the loop is one foot above the ground. The loops are spaced 80 feet apart in the multi-element arrays and ground is modeled using *EZNEC* ground with conductivity of 3 mS/m and a relative dielectric constant of 20[4].

The two-element array is easy to implement (See the sidebar "Two-Element Designs") and yields a 3 dB improvement in RDF compared to a single loop. Addition of the third element increases the RDF to above 12 dB, and a four-element (1:3:3:1) array achieves an RDF in excess of 14 dB[5]. However, this latter array will not be discussed further here because of its narrow spatial bandwidth and somewhat greater difficulty in realizing the required current ratios (optimal 1:2.6:2.6:1).

Each additional element tends to narrow the beam width and reduces the response at high angles to achieve the improvements in RDF. From the plots you'll notice that the take-off angle and beam width decrease as the number of elements increases, and the rejection off the back increases significantly. By using the cross-fire feed[3] technique, the patterns that are similar on 80 and 160 meters, as are the simulated RDF values[6].
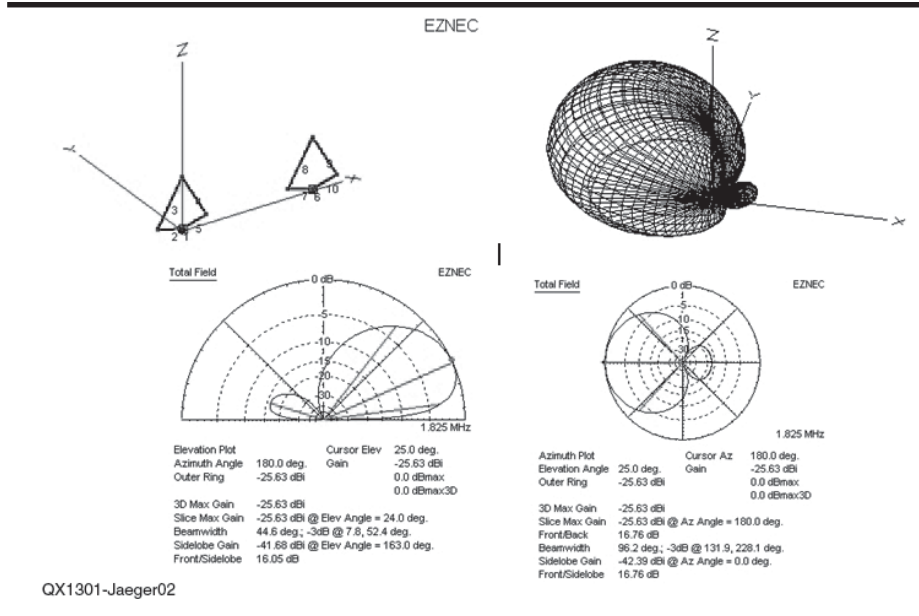
Figure 2 – Two element K9AY Array - 1.825 MHz, 80 foot spacing, -200° phasing. Average Gain: -36.2 dB, RDF 10.5 dB. The loop termination resistor is 510 Ω.

Three-Element Array Optimization

After studying of the initial simulation results, I decided to continue with the three-element array project to achieve high RDF with acceptable bandwidth. Although the four-element antenna can achieve higher RDF, the beamwidth becomes too narrow for use with only two arrays. Discussions in *ON4UN's Low-Band DXing* books indicated that higher performance could be achieved with a slight modification of the 1:2:1 current distribution. Extensive simulation work with *EZNEC* for various current distributions and array phasings yielded the results in Table 2 and Figure 6.

The RDF can actually be pushed above 13 dB, but the beam width again narrows more than I desired. My final design choice was a 1:1.8:1 current distribution with 0/-200°/-400° element phasing. On top band, the simulations predict an RDF of 12.7 dB, a take-off angle of 21°, and a beam width of 69°. The back/side lobes are down by at least 20 dB. Interestingly enough, this is essentially the same choice made independently by John Kaufman, W1FV, for the three-element core of his 9-circle vertical array recently published in *NCJ*[7].
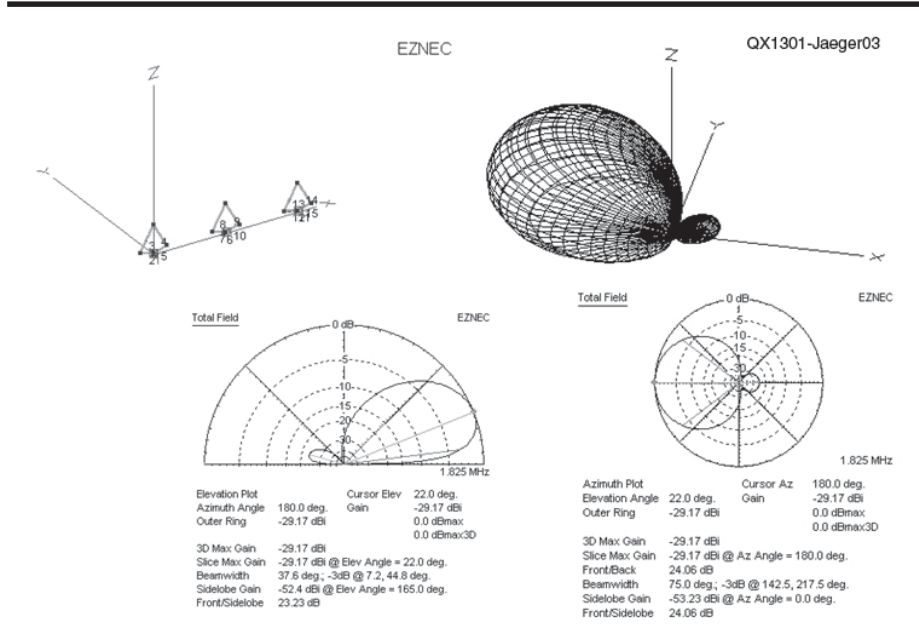

Figure 3 – Three element K9AY Array - 1.825 MHz, 80 foot spacing, -200°, -400° (-40°) degree phasing. Average Gain: -41.2 dB, RDF 12.1 dB. The loop termination resistor is 510 Ω.

Array Implementation

The layout of the two arrays on our lot appears in Figure 7. One three-element array is oriented NE/SW and a second array is used for the NW/SE directions. The two arrays require 6 loops in contrast to the eight loops needed for a four-square of K9AY antennas. Because of factors related to my lot and house, an 80-foot spacing is used between loops in the NE/SW direction, and a 57.5-foot spacing is used in the NW/SE antenna. The actual default array headings are 60° and 130° as depicted in the azimuthal plot in Figure 8.

Care was taken to try to make the loops in each array as physically identical as possible using the dimensions given earlier. Since the K9AY loops are nominally 500 Ω antennas, the loops are connected directly to Hi-Z Amplifier[8] inputs without the matching transformers that are typically used to match the K9AY loops to 50 or 75 Ω coax. See the sidebar "Loop Direction Switching." The loops in the NE/SW array are supported by 26 feet 6-inch fiberglass poles from Max-Gain Systems Inc.[9], whereas the SE/NW loops are supported by aluminum tubing that is grounded at the bottom. I see no interaction with the aluminum supports, and they are more rigid. However, simulation indicates that the aluminum poles should not be grounded.

The loops are separated from the top of the aluminum supports by a 12-inch length of PVC that slides half way into the tubing. The supports are guyed with black Nylon twine from the sides perpendicular to the loops, whereas the loops themselves stabilize the structure in the plane of the loop. The corner pair of crossed loops does not require separate guys.

A system diagram appears in Figure 9. The Hi-Z amplifiers for the two outer elements are set to have an output resistance of 75 Ω, whereas the output resistance of the center amplifier is set to approximately 37.5 Ω in order to drive two coax lines that are connected to the side element inputs of the four-square controller. To achieve the 1:1.8:1 distribution, the gain of the center amplifier is actually reduced by a factor of 0.9 by increasing the output resistance to 46 Ω. The controllers implement the cross-fire phasing method in which a transformer provides 180° of the phase shift required for the center element. Thus the short phasing line is nominally 20° (200° - 180°). However, the short phasing lines are actually cut for a 17° phase shift in order to compensate for an added 3-4° phase shift that exists in the short phasing path of the controllers. The long phasing line is set for 40° (400° - 360°).

I happened to already have two four-square controllers, so I used one for each
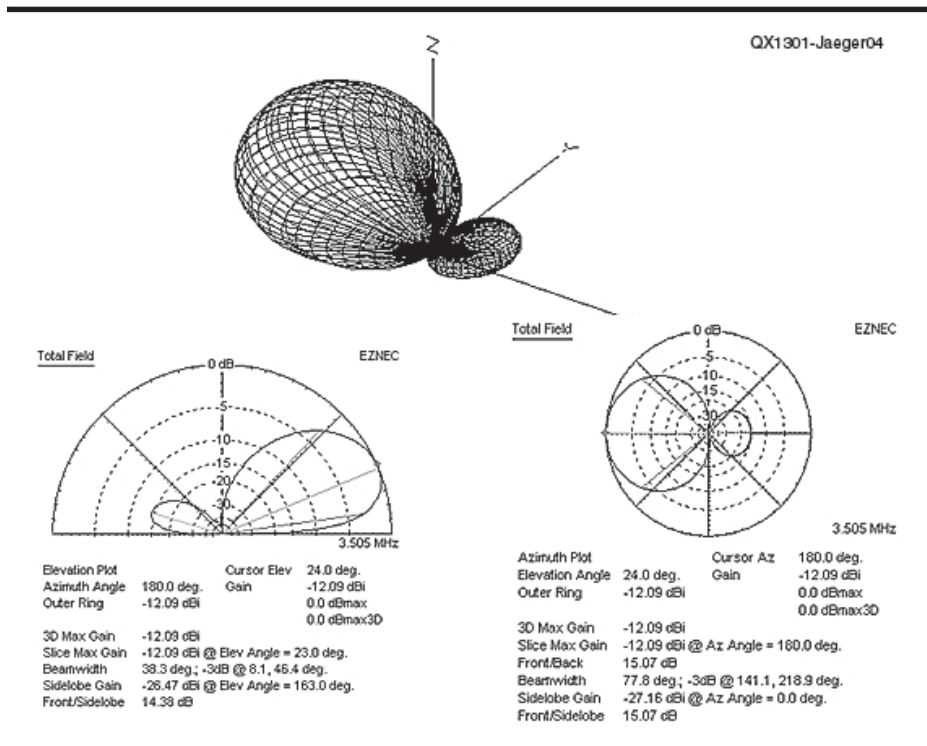
Figure 4 – Three element K9AY Array - 3.505 MHz, 80 foot spacing, -200°, -400° (-40°) degree phasing. Average Gain: -24.0 dB   RDF 11.9 dB. The loop termination resistor is 510 Ω.
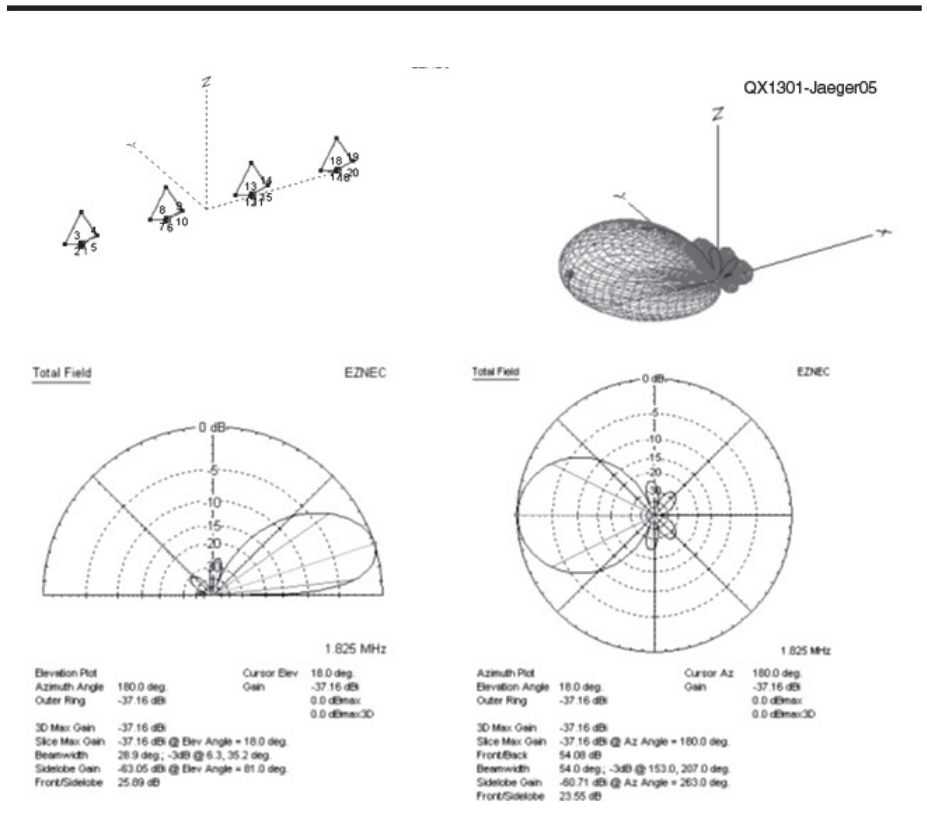
Figure 5 – Four element K9AY Array - 1.825 MHz, 80 foot spacing, -195°, -390° (-585°) degree phasing. Average Gain: -51.6 dB, RDF 14.5 dB. The loop termination resistor is 510 Ω.

QX1301-Jaeger06

Figure 6 — RDF vs phase shift of the short phasing line. The final design choice is an RDF of 12.6 dB for a center element phase of -200° and a relative amplitude of 1.8.

QX1301-Jaeger07

Figure 7 — Array orientation on the lot.

QX1301-Jaeger08
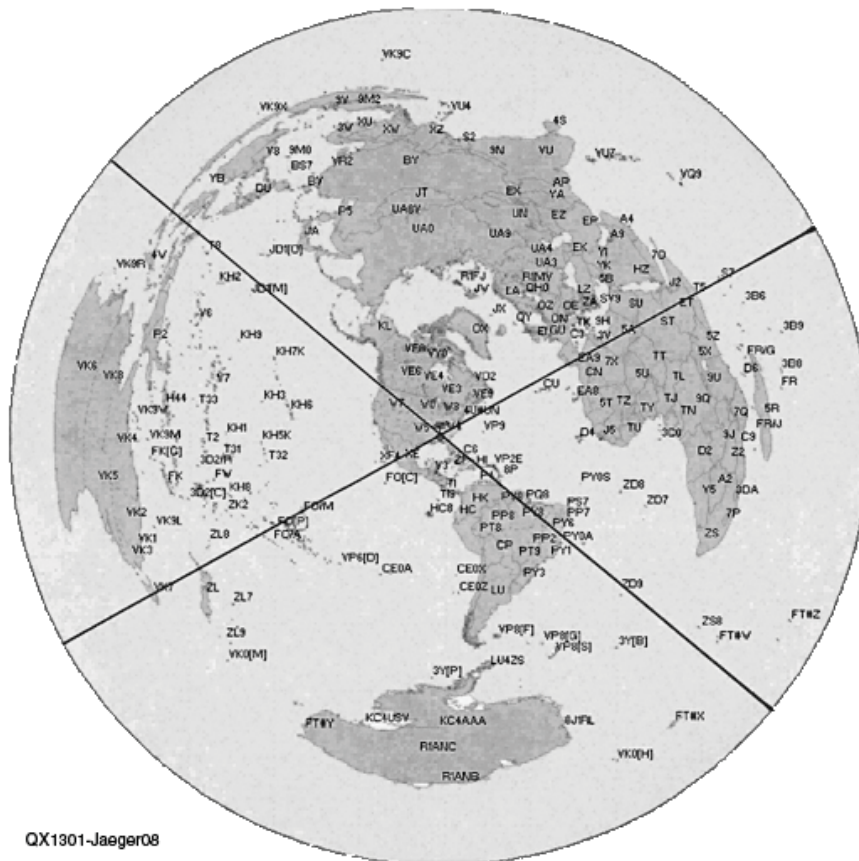
Figure 8 — Azimuthal plot showing actual array directions from K4IQJ's station. (Plotted with *DXAtlas***.)**

Table 2
Simulation Results
1:2:1 Binomial End-fire Arrays* - 1.825 MHz - RDF versus Crossfire Phasing

Phasing	−190/−20	−195/−30	−200/−40	−205/−50	−210/−60
RDF (dB)	12.6	12.3	12.1	11.8	11.6
Angle	21°	21°	22°	22°	23°
Beam Width	69°	72°	75°	77°	80°
Back Lobes (dB)	−16	−20	−24	−28	< −30
Gain (dBi)	−31.7	−30.4	−29.3	−28.3	−27.3

1:1.9:1 Binomial End-fire Arrays* - 1.825 MHz - RDF versus Crossfire Phasing

Phasing	−190/−20	−195/−30	−200/−40	−205/−50	−210/−60
RDF (dB)	13.0	12.7	12.4	12.1	11.8
Angle	20°	21°	21°	21°	22°
Beam Width	66°	69°	72°	75°	78°
Back Lobes (dB)	−17	−21	−26	−28	−29
Gain (dBi)	−32.3	−30.9	−29.7	−28.6	−27.6

1:1.8:1 Binomial End-fire Arrays* - 1.825 MHz - RDF versus Crossfire Phasing

Phasing	−190/−20	−195/−30	−200/−40	−205/−50	−210/−60
RDF (dB)	13.3	13.0	**12.7**	12.4	12.0
Angle	20°	20°	**21°**	21°	22°
Beam Width	62°	66°	**69°**	73°	76°
Back Lobes (dB)	−16	−18	**−20**	−21	−23
Gain (dBi)	−33.0	−31.5	**−30.1**	−28.9	−27.8

*Arrays with both 80 ft. and 57.5 ft. spacing are essentially the same.

Loop Direction Switching

An illustration of single and double termination switching. See Table A1.

Table A1
Comparison of K9AY 3-Element End-Fire Arrays and Four Square
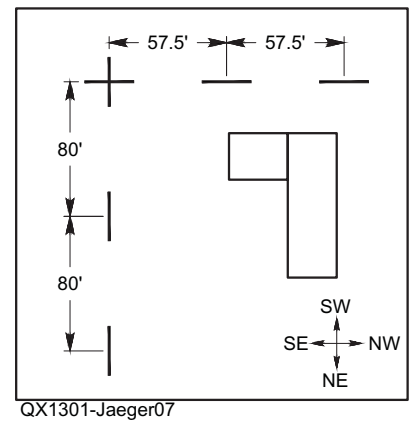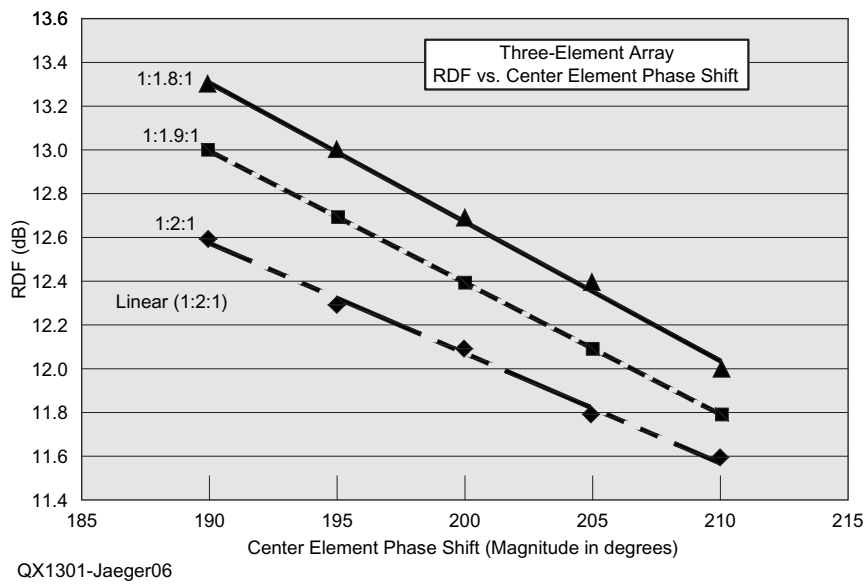
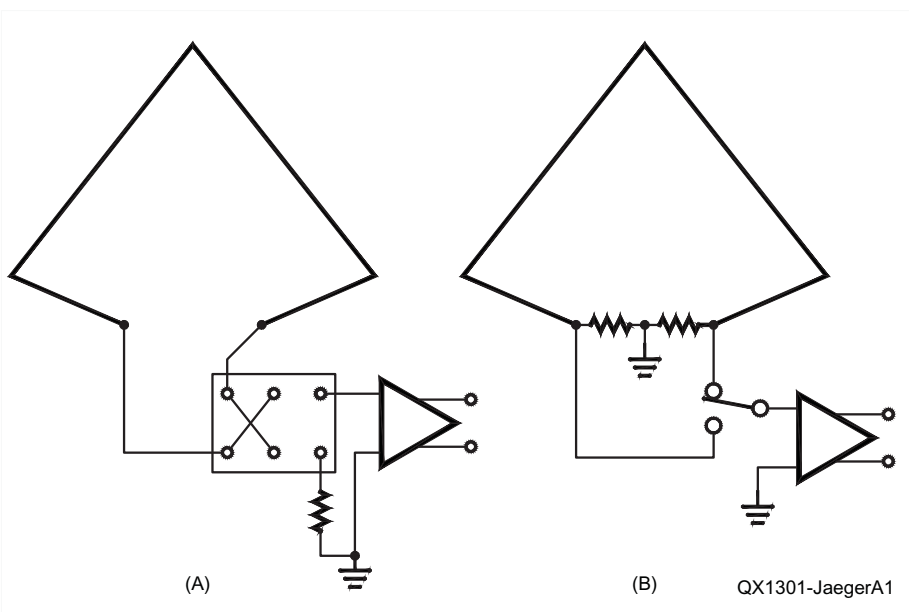|	Orthogonal End-Fire Arrays	Four Square
Number of Loops	6	8
Number of Amplifiers	5 or 6	4

(A)   (B)   QX1301-JaegerA1

Figure A1 — Single termination with double-pole double throw relay (A) and double termination with single-pole double throw relay (B). In the double termination output is reduced by about 6 dB. The loop termination resistor is 510 Ω.

array and did not develop the switching matrix needed to use one controller for both arrays. One array uses the four-square controller from Hi-Z Antennas to implement the phasing and array control, and the second utilizes the four-square controller from DX Engineering[10]. The use of two controllers allows additional potential advantages that will be mentioned later. Note that in order to reverse the array directions, the loop directions must be changed at the same time that the controller phasing is switched. Circuits for single and double terminated loop switching appear in the sidebar "Loop Direction Switching."

Figure 10 presents photos looking down array axis from the end and corner. The antennas are well hidden in the trees and are not noticable even though they are in our front yard. The poles are mounted on wooden supports cut from 6-foot treated 2 x 4s as in Figure 11. The 2 x 4 is cut into 2 feet 6 inch and 3 feet 6 inch pieces and nailed together, thereby providing a shelf plus a 1 foot region to which the poles are attached with stainless steel U-bolts. The wooden support is buried 18 inches in the ground. The amplifiers, terminating resistors and direction switching relay are mounted in weatherproof boxes (from Lowes) as in Figures 11 and 12. Stainless steel hardware is used for all connections. The loops connect to the top of the boxes, the coax and ground connections exit the bottom of the box, and the direction control comes into the side. I would probably mount the amplifier against the back wall of the box if I did it again, in order to facilitate access to the gain adjustment pot on the circuit board, and I would use two separate boxes at the array corner.

All feed and phasing lines utilize flooded RG-6 coax, and the lines from each amplifier to the controller were chosen to be 90 feet. Each feed line goes through a 50-Ω beaded coax choke (from The Wireman[10]) at the output of the amplifier. Pairs of similar chokes are placed in the array output lines approximately half way between the controllers and the shack as in Figure 13. The two beaded chokes are connected together with a barrel connector with the coax shield grounded at the barrel connector. The mismatch with RG-6 can be tolerated because of the short length of the 50-Ω line. "Braid breakers" were later added to the NE/SW system, but no difference in performance was noted.[3]

Adjustment and Characterization

Phasing lines were carefully adjusted using a VNA-2180 network analyzer by measuring the resonant frequency of the open-circuit line (Note that this can also be done with any of the single-port antenna

Table 3

Phase Shift of Phasing Lines – Calculated from Open Circuit Measurements

Phasing Line	NE/SW	NE/SW	NW/SE	NW/SE
1	21.5 feet / **17.2°**	33.0°	**16.7°**	32.0°
2	21.5 feet / **17.2°**	33.0°	—	—
3	49.0 feet / **39.3°**	75.5°	**41.3°**	79.2°

Table 4

Amplifiers Plus Controller Gain and Phase Matching

1.827 MHz Results

Loop	NE/SW Gain	NE/SW Phase	NW/SE Gain	NW/SE Phase
Front	+0.13 dB (1.01)	0° (ref)	+0.05 dB (1.01)	0° (ref)
Middle	+5.19 dB (1.82)	**−200.3°**	+5.30 dB (1.84)	−200.1°
Back	−0.13 dB (0.99)	−399.8° (**−39.8°**)	−0.05 dB (0.99)	−401.5° (−41.5°)

3.505 MHz Results

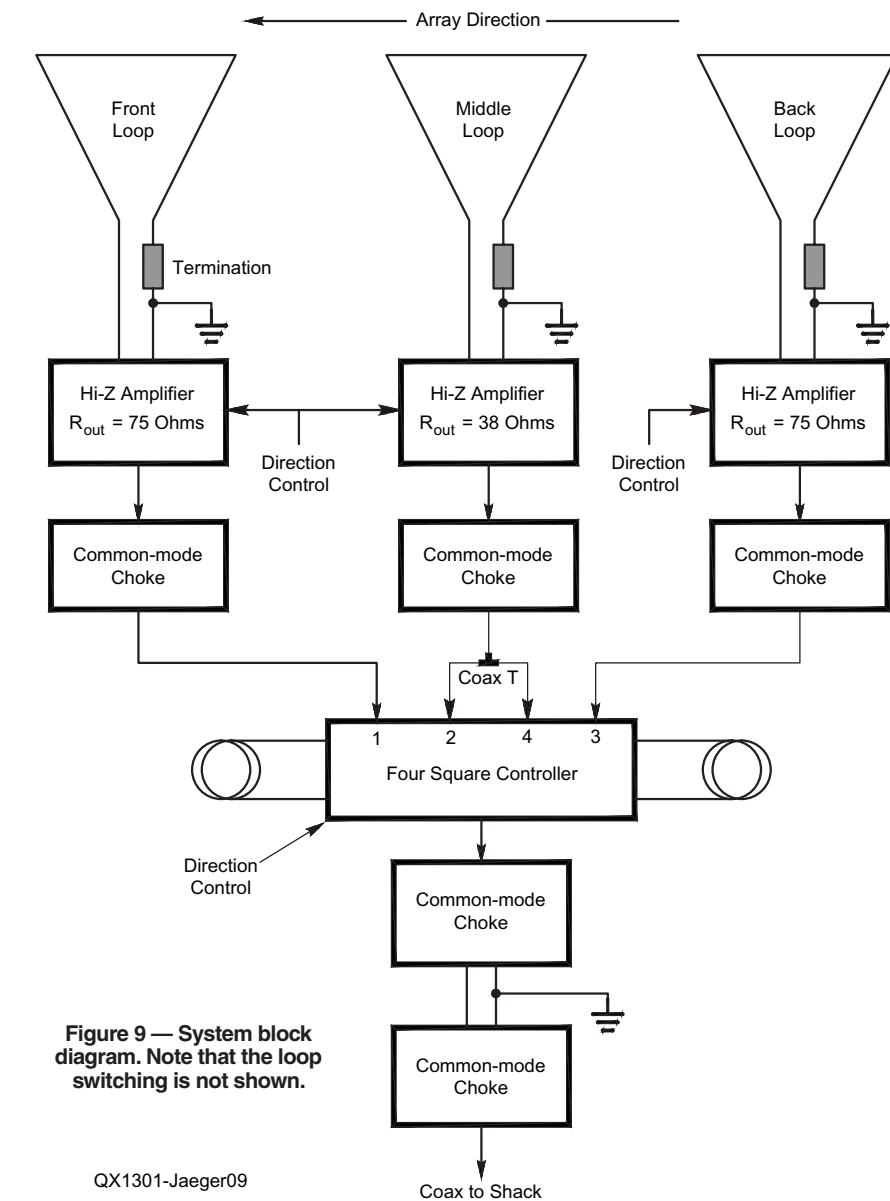| Loop	NE/SW Gain	NE/SW Phase	NW/SE Gain	NW/SE Phase
Front	+0.14 dB (1.02)	0° (ref)	−0.01 dB (1.00)	0
Middle	+5.07 dB (1.79)	**−218.6°**	5.26 dB (1.83)	−220.8
Back	−0.14 dB (0.98)	**−434.0° (−74.0°)**	0.01 dB (1.00)	−440.8 (−80.8°)

Figure 9 — System block diagram. Note that the loop switching is not shown.

QX1301-Jaeger09

analyzers.), and the results appear in Table 3. After one step of trimming, the phases were within about 1 degree of the design values.

I first built a wide-spaced two-element NE/SW array and then added the third loop. The observed changes in front-to-back ratio were quite noticeable with each added loop. Other than reduced gain in the shorter array, I see little difference between the 57.5 foot and 80 foot three-element arrays on either 160 or 80 meters.

The VNA-2180 was used to characterize the path from the input of each Hi-Z amplifier to the output of the corresponding array controller. A 50 Ω coax feeds the signal from the VNA to input of the Hi-Z amplifier and is terminated in a 51 Ω resistor at the amplifier input. A 75 Ω coax section brings the signal from the output of the controller back to the input of the VNA, and a 75-to-50 Ω pad is used for matching at the VNA input. After the first set of measurements, the Hi-Z amplifier gains were tweaked a bit and the final results appear in Table 4. On 160 meters, the relative gains and phase shifts are very close to the design values. The element phase shifts agree well with the phase shift of the phasing lines in Table 3. The phase shifts for the center elements exhibit the additional 3-4 degrees of phase shift in the controller paths as mentioned above. Similar results are obtained on 80 meters, although the excess phase shift appears to be slightly larger.

Since the same phasing lines are used on 160 meters and 80 meters, the phase shift contribution of the line is approximately twice as much on 80 meters as compared to 160 meters, and results in reduced RDF on 80 meters. Improved performance could be achieved by switching phasing lines. However, because the signal strengths are generally much larger on 80 meters, the reduction in RDF is not a major issue, and switching phasing lines does not seem worth the effort.

Simulation results for the two arrays using the measured amplitude and phase information and appear in Table 5 and Figure 14 below. The lower value of RDF on 80 meters results from the larger phase shifts through the delay lines (0/-220°/-440°) versus (0/-200°/-400°) on 160 meters. As mentioned above, the reduction in RDF is usually not an issue because of the stronger signals on 80 meters. In addition, the reduced RDF is offset by a wider beamwidth.

The Bottom Line

The receiving arrays are in use most of the time on the low bands to help save my ears. In fact, I use them on any frequency where there is an advantage, including 40 meters, 30 meters and higher, although

the patterns may be quite distorted. Generally there are lobes pointing somewhere, and I even use the arrays on 17 meters and 12 meters[12]. With the three-element arrays, I am now hearing the best yet on 160 meters and 80 meters. The output of the longer array is greater than that of the shorter array as expected, and the observed front/back ratios are also consistent with the simulation results. Signals often disappear when switching the direction by 180°. I frequently hear stations that I cannot hear on my inverted-L transmitting antenna. For instance, S79GM, FR/DJ7RJ and 5R8RJ would probably not have been worked on either 160 or 80 without these antennas. I could hear the FR and 5R8 signals nightly for a week on both bands (obviously propagation also favored me). When I first listened to the first PJ4 DXpedition signals, they were very weak on 160 meters, but I could copy them clearly on the SE array. Other recent successes include XU7ACY, 4L/UUØJM, 4L5O, 9L5MS, 5M2TT, 3D2C and PTØS. I had more than a dozen contacts with VK3ZL through the summer noise during June, July and August of 2010-2012. I missed 9Q5ON on 160 meters because they couldn't hear me, but I did work them for a new one on 80 meters. It is important to realize, however, that these antennas will not work well all the time, particularly since the high angle response is surpressed in order to achieve high RDF. The old adage that one can never

Table 5
Simulation Results Using Measured Data

Band	160 foot Array (NE/SW)		115 foot Array (NW/SE)	
	160 Meters	*80 Meters*	*160 Meters*	*80 Meters*
RDF (dB)	12.6	11.5	12.4	11.5
Angle	21°	25°	21°	24°
Beam Width	70°	84°	71°	84°
Gain (dBi)	-30.1	-10.7	-33.8	-12.8
F/B (dB)	32.2	22.3	25.6	54.1
Front/Side (dB)	29.4	20.6	18.4	27.7

Figure 10 — End of NE/SW array (left) and the corner of the two arrays (right).

(A)　　　　(B)　　　　(C)

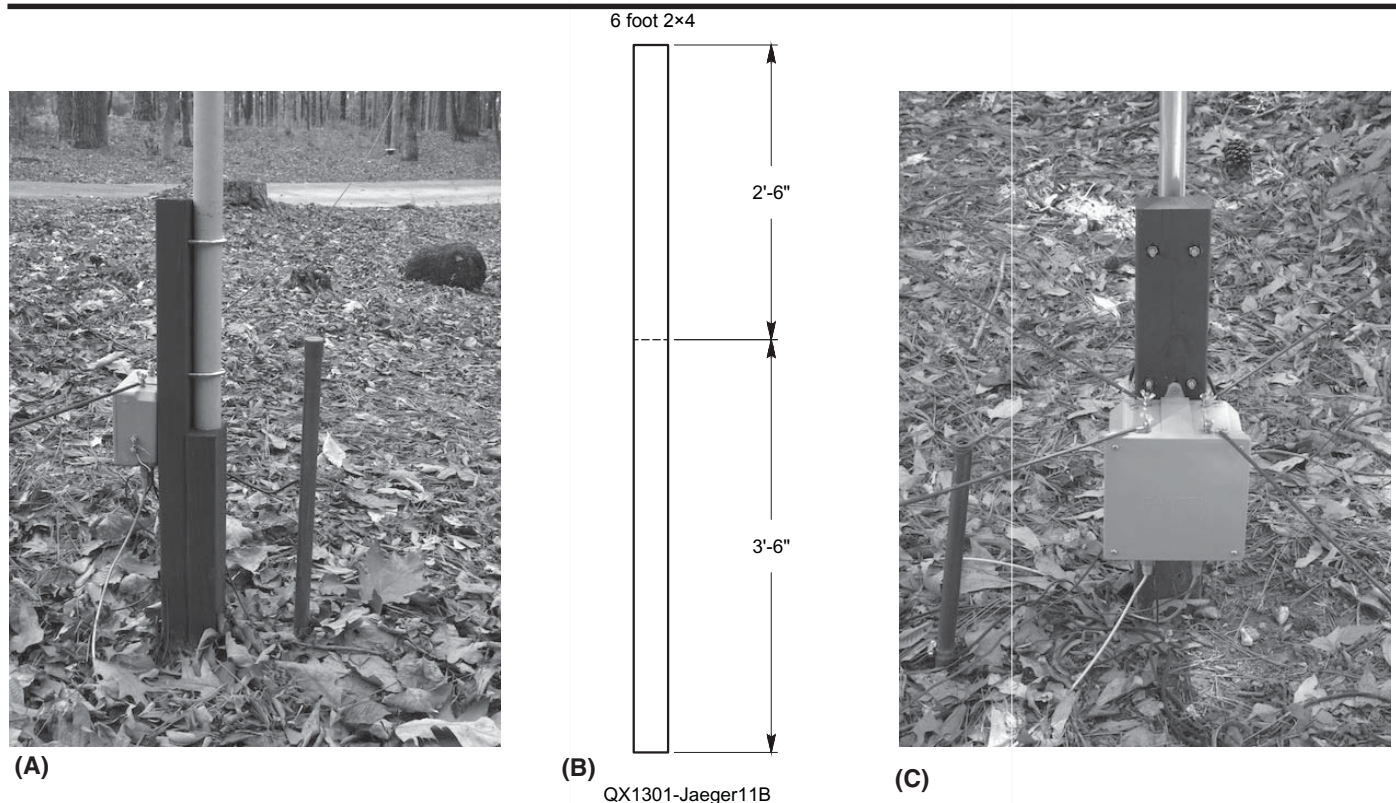QX1301-Jaeger11B

6 foot 2×4

2'-6"

3'-6"

Figure 11 — Loop and amplifier (A), the loop support (B) and the corner amplifier box (C).
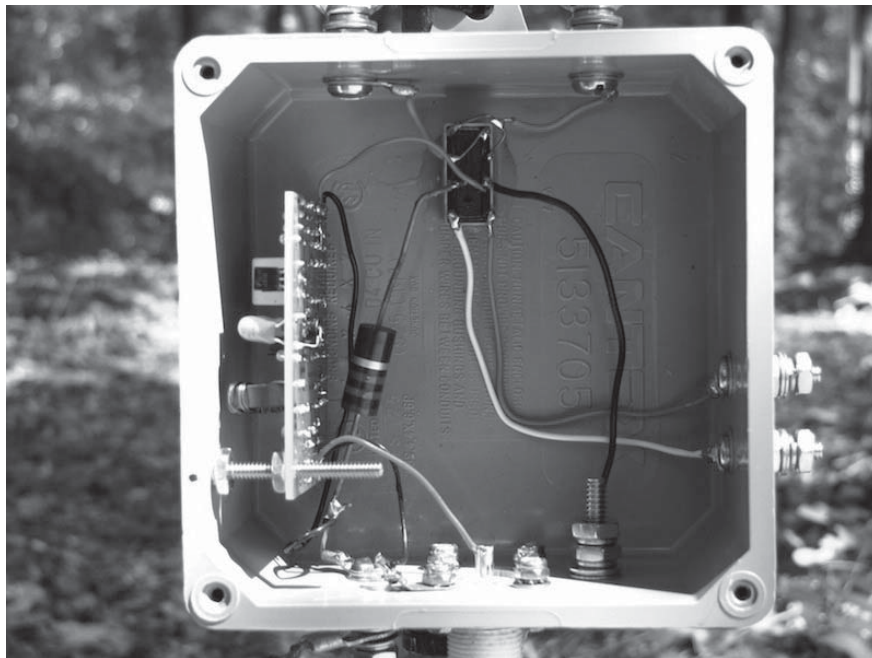
Figure 12 — Exterior and interior views of amplifier/control box.

have too many low-band antennas is certainly true.

For comparison purposes, I have put up a three-element short vertical array in parallel with NE/SW loop array and separated from it by approximately 20 feet. In several months of testing, the loop array has always outperformed the vertical array except for one instance when they were about the same.

Further Ideas

Having two separate arrays has some additional potential uses that are being explored. The arrays have very high rejection off the sides, so one can use the second array as a noise antenna with one of the commercial phasing controllers to eliminate pesky noise sources or strong unwanted signals without substantialy reducing the desired signal. Or, the output of the two arrays can be combined to provide a pattern that is switchable in 8 directions. The combined patterns have an RDF of about 9.7 dB in the ±45° directions, approximately 2 dB better than either array at ±45°.

Array additions have been recently implemented using a hybrid combiner with a pair of switches. It is now in routine use. Note that the output phase and amplitude of the two arrays need to be balanced for best results.

The aluminum tubing supports for the NW/SE array were chosen to permit a direct comparison between the loop array and a
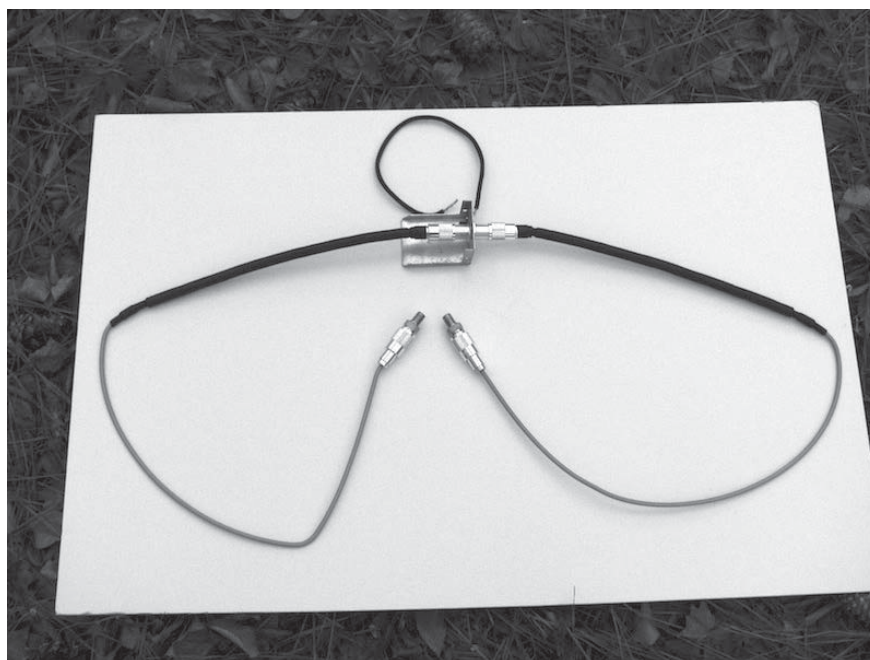

Figure 13 – Feed line chokes from The Wireman.

Two-Element Designs

Figure B1 depicts an implementation of one of a pair of two-element arrays using the Hi-Z 2-3 element controller. The 200° phase shift from Table 1 is achieved by an inverting transformer in the controller plus a 20° phasing line. The feed lines from the loop antennas to the controller should have identical lengths. I have designed one of these systems with a narrow 50 foot spacing so that the pattern will be maintained up through 40 meters. A second pair of loops is used with the controller to get the orthogonal directions. Note that the phasing of the two-element array can also be done with the front element plus one of the side-element inputs of a four-square controller.

Alternate Two-Element Implementation

Figure B2 depicts an implementation of a passive two-element array with transformers in place of the amplifiers. A 200° phase shift from Table 1 is achieved by inverting the transformer winding of the front element and increasing the length of the back coax by the equivalent of 20°. The coaxes are joined with a 0° hybrid followed by a matching transformer to drive the output line (see Chapter 7 of *ON4UN's Low Band DXing*, 5th edition). The hybrid is designed to match the coaxes so that the phase shift is controlled by the line length. Direction control is not shown in this figure. Additional information can be found in "Enhancing the Performance of the K9AY Receiving Loop Array" presented at the 2012 Dayton Antenna Forum and available at **www.k3lr.com** or **www.k4iqj.com**.
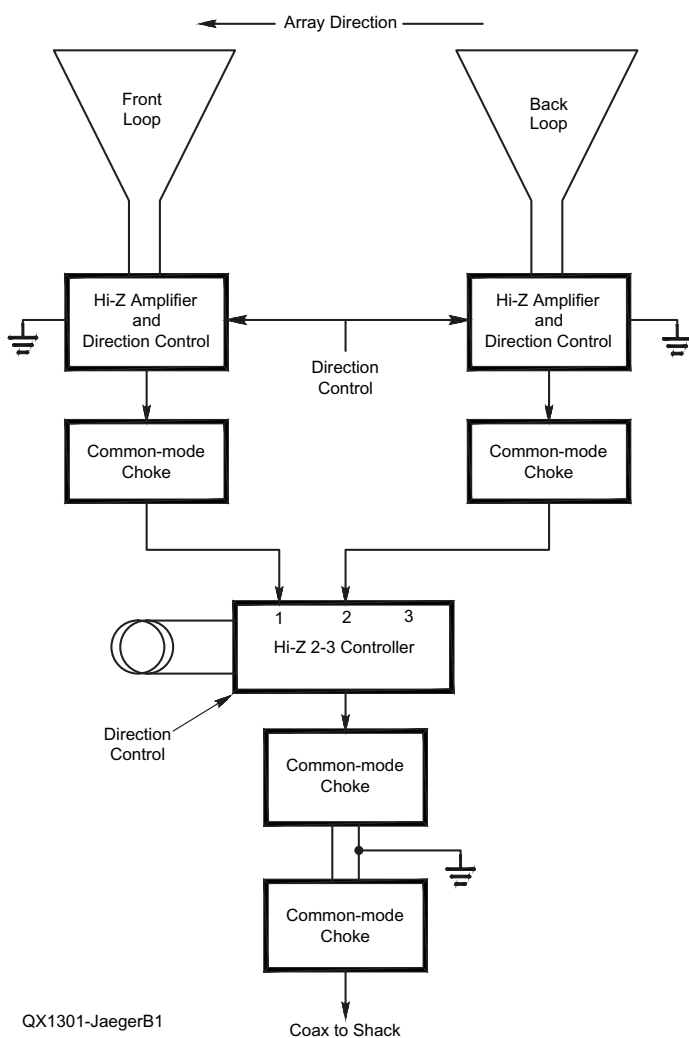

QX1301-JaegerB1

Figure B1 — A pair of two-element arrays using the Hi-Z 2-3 element controller.
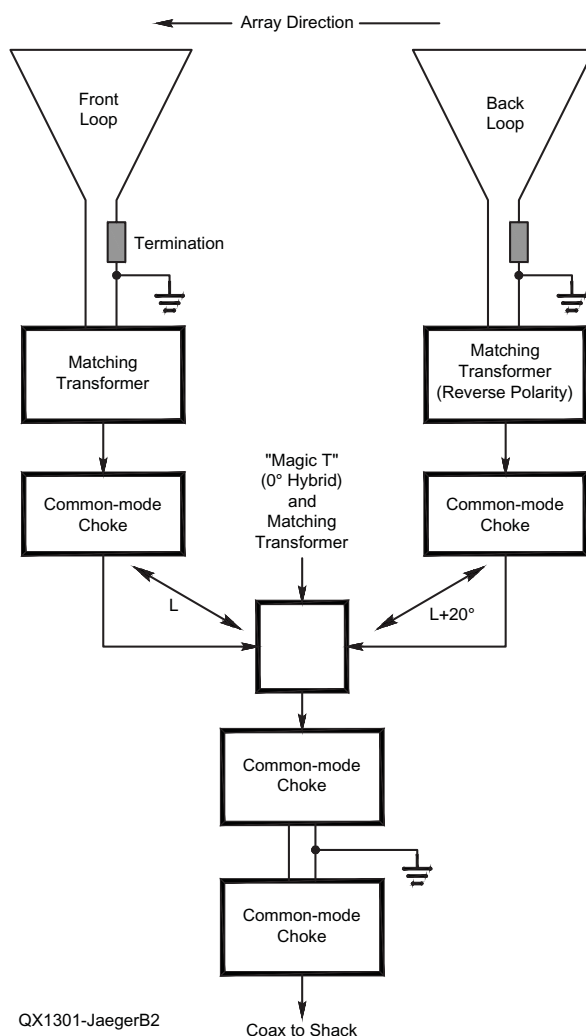

QX1301-JaegerB2

Figure B2 — A passive two-element array with transformers in place of the amplifiers.

similar array of short verticals. Simulation has indicated that there is essentially no interaction between the loops and the ungrounded aluminum supports. I plan to add a set of relays to enable switching between the loops and short verticals and make direct comparisons between the two types of arrays. I will post results in the future as they become available at **www.k4iqj.com**.

Notes

[1]Gary Breed, K9AY, "Arrays of K9AY Loops: 'Medium-sized' Low Band RX Antenna Solutions," expanded version from the W9DXCC Convention, Sept. 15, 2007. **www.aytechnologies.com**

[2]Dallas Lankford, **http://groups.yahoo.com/group/thedallasfiles**.

[3]John Devoldere, ON4UN's *Low-Band DXing*, 4th and 5th Editions, ARRL, Newington, CT.

[4]**www.fcc.gov/mb/audio/m3/index.html**

[5]In fact certain four-element designs can achieve RDF of over 15 dB.

[6]I have recently realized that the four-element array can also be implemented using existing controllers and amplifiers with some additional switching. However, the beam width is too narrow for my present application.

[7]John Kaufman, W1FV, "A Compact Dual-Band 9 Circle Receiving Array – Part 1," *National Contest Journal*, vol. 39, no. 5, pp. 14-18, September/October 2011. Part 2 appears in the November/December 2011 issue of *NCJ*.

[8]Hi-Z Antennas by Lee Strahan, K7TJR, and Dick Ewing, KO7N, Culver, OR: **www.hizantennas.com**.

[9]Max-Gain Systems Inc: **www.mgs4u.com**.

[10]DX Engineering: **www.DXEngineering.com**.

[11]The Wireman: **www.thewireman.com**.

[12]An advantage of the loops is an essentially constant impedance to above 20 MHz.

Dick was first licensed as KN4IQJ in 1959 and upgraded a few months later. He is primarily a CW operator who enjoys DX and contest operation. Dick is on top of the SSB/Mixed Honor Rolls, needs only P5 on CW and has competed 5BWAZ. He has confirmed 246 countries and 38 zones on 160 meters. Over the years he has operated from PJ4, KH6, ZL, TA, KL7, BV, VS6 and UA3.

In his other life, Dr Jaeger is Professor Emeritus of Electrical and Computer Engineering at Auburn University and founded the Alabama Microelectronics Center. He is author/co-author of Introduction to Microelectronic Fabrication, 2nd edition *and* Microelectronic Circuit Design, 4th edition*, and received the 2004 IEEE Undergraduate Teaching Award.*

He is an IEEE Fellow, Past President of the IEEE Solid-State Circuits Society, former Editor of the IEEE Journal of Solid-State Circuits*, and chaired both the International Solid-State Circuits Conference and the International VLSI Circuits Symposium. He was the Founding Editor of* IEEE Micro Magazine.
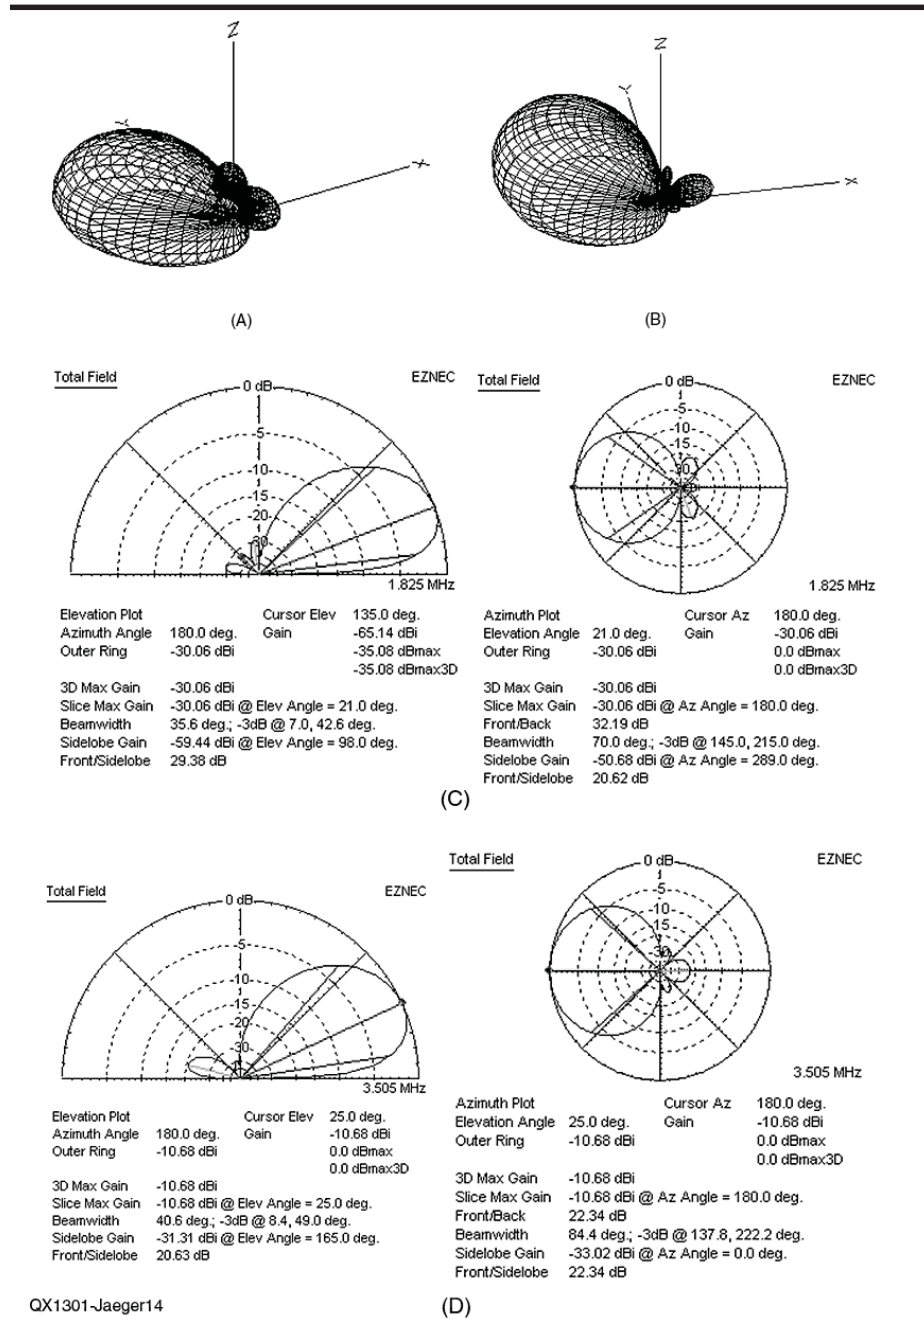
Figure 14 — 3D simulations of final NE/SW 160 foot array on (A) 160 meters and (B) 80 meters. A simulation of the final NE/SW 160 foot array on 160 meters is shown at (C) and a simulation of the final NE/SW 160 foot array on 80 meters appears at (D).

QX1301-Jaeger14

Clark Sell, N9APK

2418 Lydia Ave, Zion, IL 60099

Use Arduino Technology to Create a Clock/10 Minute Timer

And put it to work bringing your SB-630 back to life!

I had an old Heathkit SB-630 station console in my shack that was just taking up space. If you're familiar with the SB-630, you know one of its most familiar features was its mechanical digital clock. My SB-630 was a great piece of gear in its day, but now its clock had become terribly inaccurate.

The condition of my SB-630's clock provided the inspiration I needed to create a replacement. My new timepiece would be highly accurate and include a battery backup. It would also feature a 10-minute legal identification timer that would automatically send my call sign in CW and activate a visual alarm. My new clock would be designed to reset the 10-minute timer whenever it sensed my transmitted signal — or whenever I chose to reset it manually using any old infrared remote control.

My goal was to do all this in such a way that it would preserve the classic appearance of the SB-630 as much as possible, despite the fact that the mechanical clock display would be replaced with an LCD. Of course, this clock project can also be a stand-alone device that's useful in any station, even if an SB-630 is not available.

The brain of this project is the popular Arduino microprocessor. The Arduino project was born in Italy in 2005 and by 2010 over 120,000 units had been sold. This is an open source platform, simple in design and low in cost. The typical Arduino processor and project board costs $30 US. The software is a form of *C++* and there are quite a few online resources available to explore, along with a wealth of hardware resources worldwide.

Arduino systems have expansion cards that stack on top of the microprocessor and each other. They call these expansion boards *shields* and there are many to choose from, or you can build your own.

Easy prototyping is why "Arduino fever" has captured the electronic hobbyist world. With the new parts that are now on the market, the average ham can create devices that might have been well beyond his reach a few years ago.

Getting Started

A very good source on the web for getting started with Arduino projects is **www.ladyada.net/learn/arduino/index.html**. This site will show you how to get up and running with simple project ideas. It will also link to software you'll need to
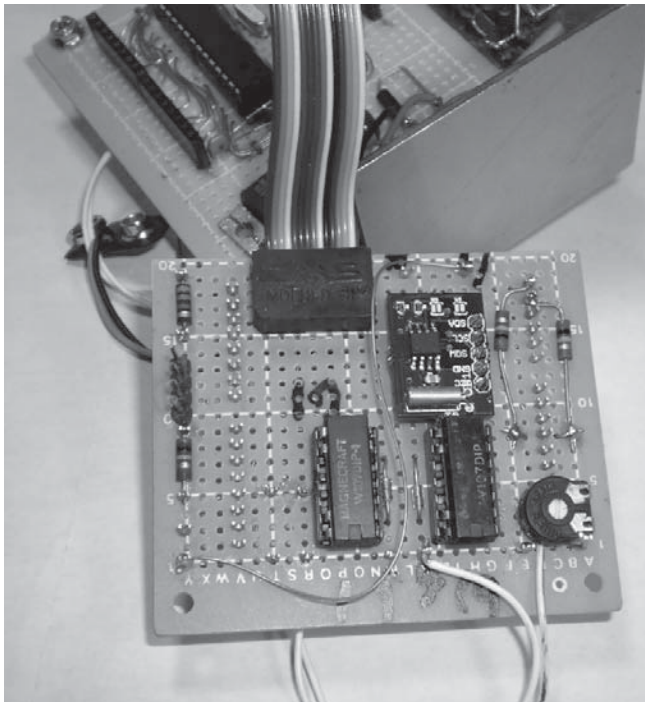

Original (circa 1960) SB-630 station console with a slight mod to switch audio from another radio into the phone patch.
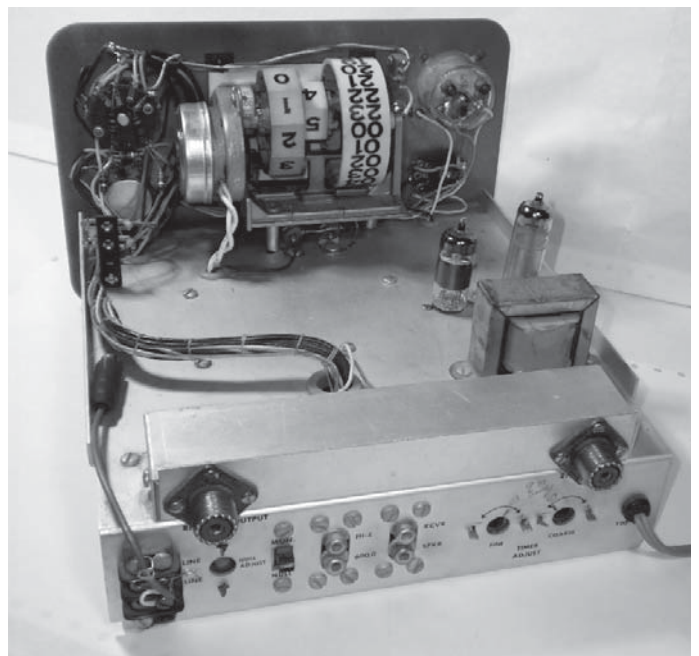
Rear view of the console showing the old clock.
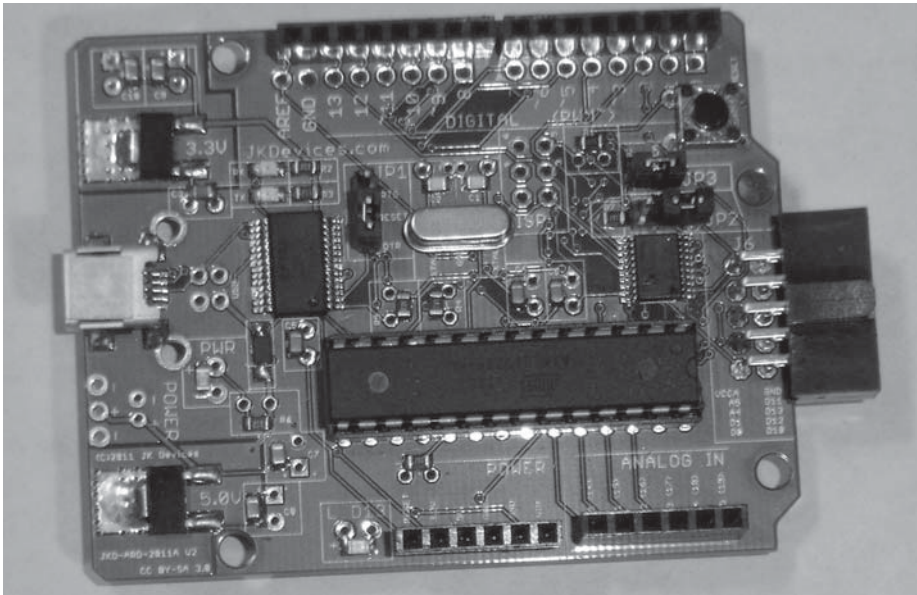
Arduino board can be used as the main board or development of software for other projects.

The Real Time Clock with battery backup shown with the two relays and LCD contrast control.

download such as a free IDE (Integrated Development Environment). Even if you are not going to write software you'll need the IDE to load the program to the microprocessor to run this project.

Background

The open source approach to software has many benefits, among them the fact that people throughout the world are using the same tools to create, build and test ideas. Open source software can be freely downloaded and modified as long as you give credit to the original author and do not attempt to sell it. The example of an open source code that I used in this project was written by Maurice Ribble in 2008 and was modified in **tronix-stuff.com/tutorials**, Example 7.4. I downloaded and examined the code, but it didn't have a 10 minute countdown timer, much

less the means to reset it.

One of my ideas was to have the microprocessor read the trigger switch and then execute the code based on the position of switch. So, I took the Example 7.4 code and moved all of the lines having to do with Date and Day to its own function (in *BASIC* that would be a known as a *sub-routine*) and placed a call to it as "daydate."

The main loop of the Example 7.4 code reads the DS1307 (a real time clock) and prints the time to the LCD every second. It then prints the day and date, and repeats. I added a switch input for the code to check. When the hardware switch closes, it places 5 V on pin 8 of the processor. When the processor checks this pin and finds it is high "(if(digitalRead(8))==HIGH)" it will execute the code that follows the IF statement. The code that follows the IF statement is the 10 minute countdown timer. The routine simply subtracts 1 from the countdown seconds, compares for zero, subtracts one from countdown minutes and checks for zero, flashes the LEDs and redisplays the date and day.

I didn't like the idea of a simple buzzer sounding at the end of the 10 minute countdown. Instead, I found a small program that sends CW known as *Simple Arduino Morse Beacon* written by Mark VandeWettering, K6HX. I added it to my modified clock code and changed a few things. In addition to adding my call sign, I increased the sending speed to 30 WPM to avoid causing issues with the main timing loop. I also attached a connection from pin 6 of the processor, through a transistor, to a monitor speaker.

Future Expansion

I was looking for an atomic clock module to add to the project so that the clock could be regularly calibrated to WWV. Unfortunately, my search wasn't successful. An alternative is to add an EtherShield with a memory card, which would add a web server to the SB-603 so that you could control your station remotely, as well as set the time from the Internet. There are Bluetooth and Wi-Fi shields available as well for wireless access.

The SB-603 has SWR and VU metering that could also become part of this project. In fact, you could display the SWR and VU levels on the clock LCD. Should anyone achieve this before I do, please let me know.

Sketch Options

In the Arduino world a *sketch* is a unit of code that is uploaded to, and run on, the Arduino board. *SB630RTCW.PDE* is the sketch (source code name) for this project and it is included in the *SELL.ZIP* file that you can download from the *QEX* website at **www.arrl.org/qexfiles**. Of course, you will need to change a few lines to add your call sign.
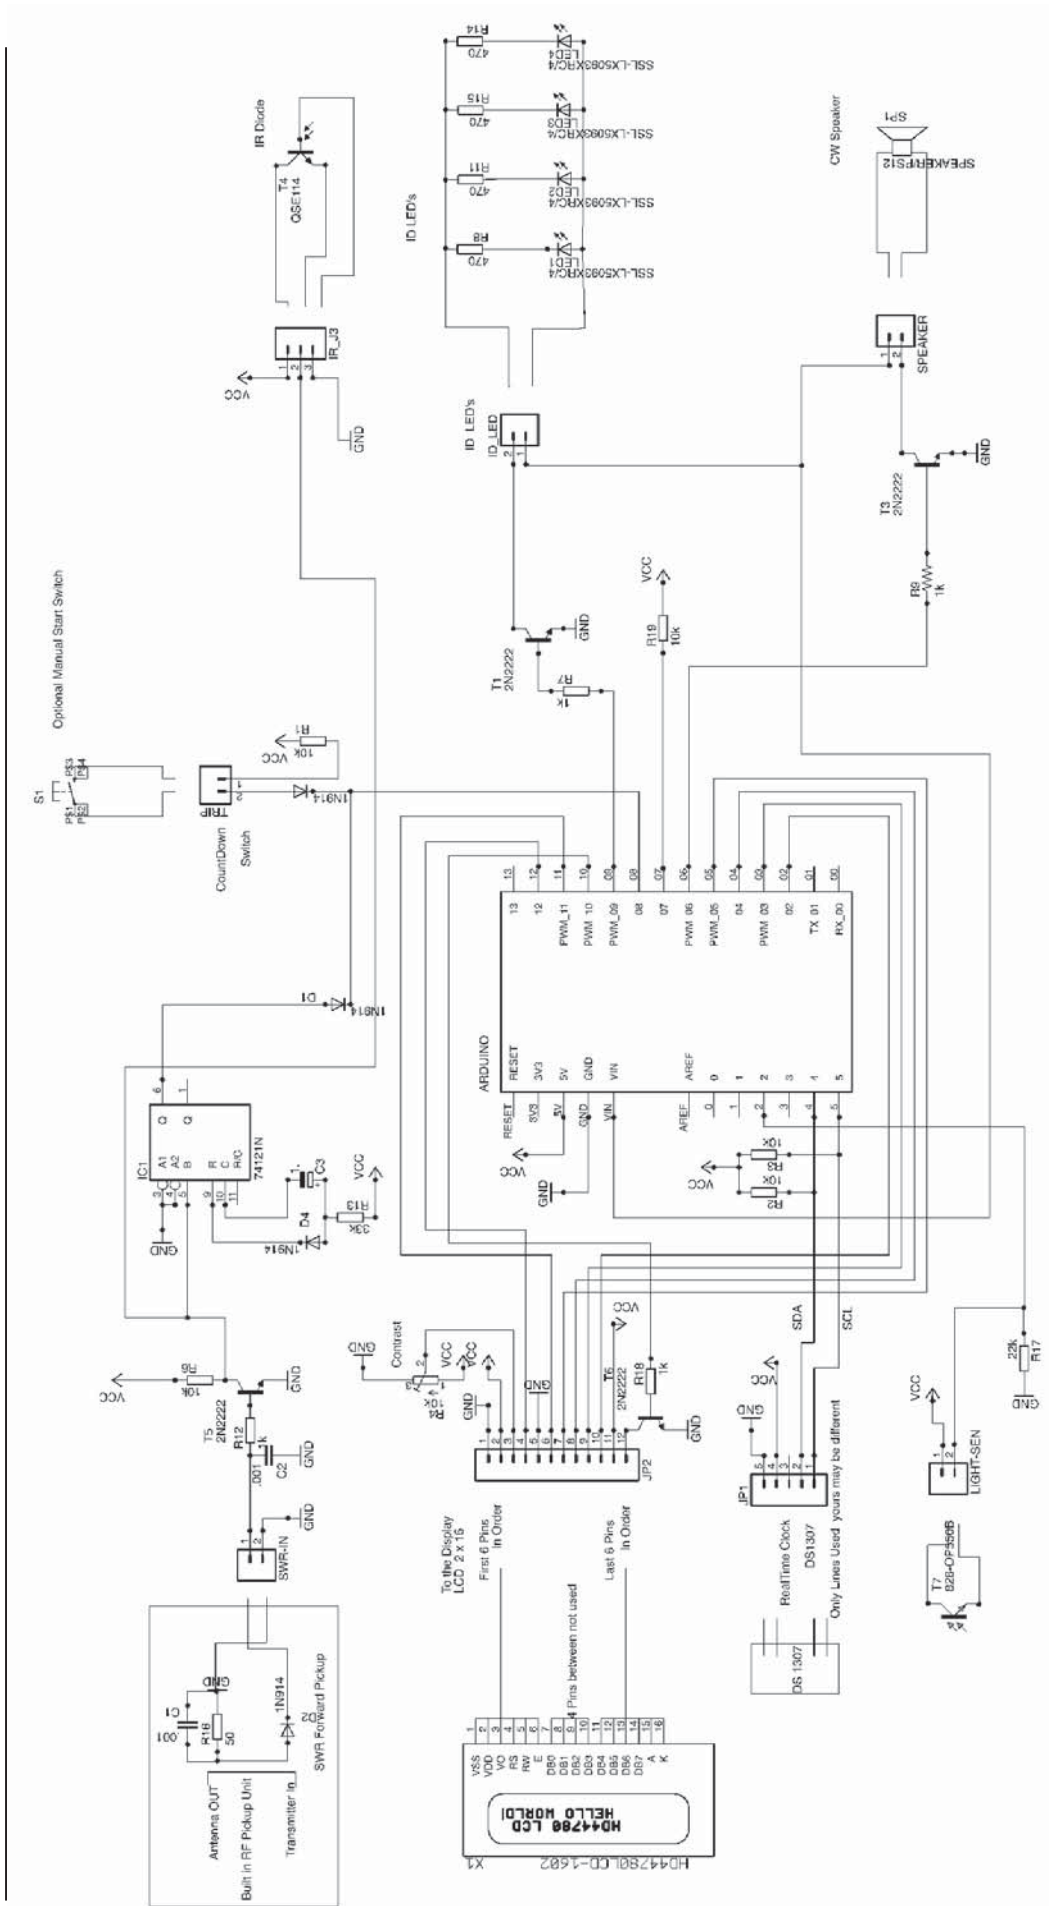
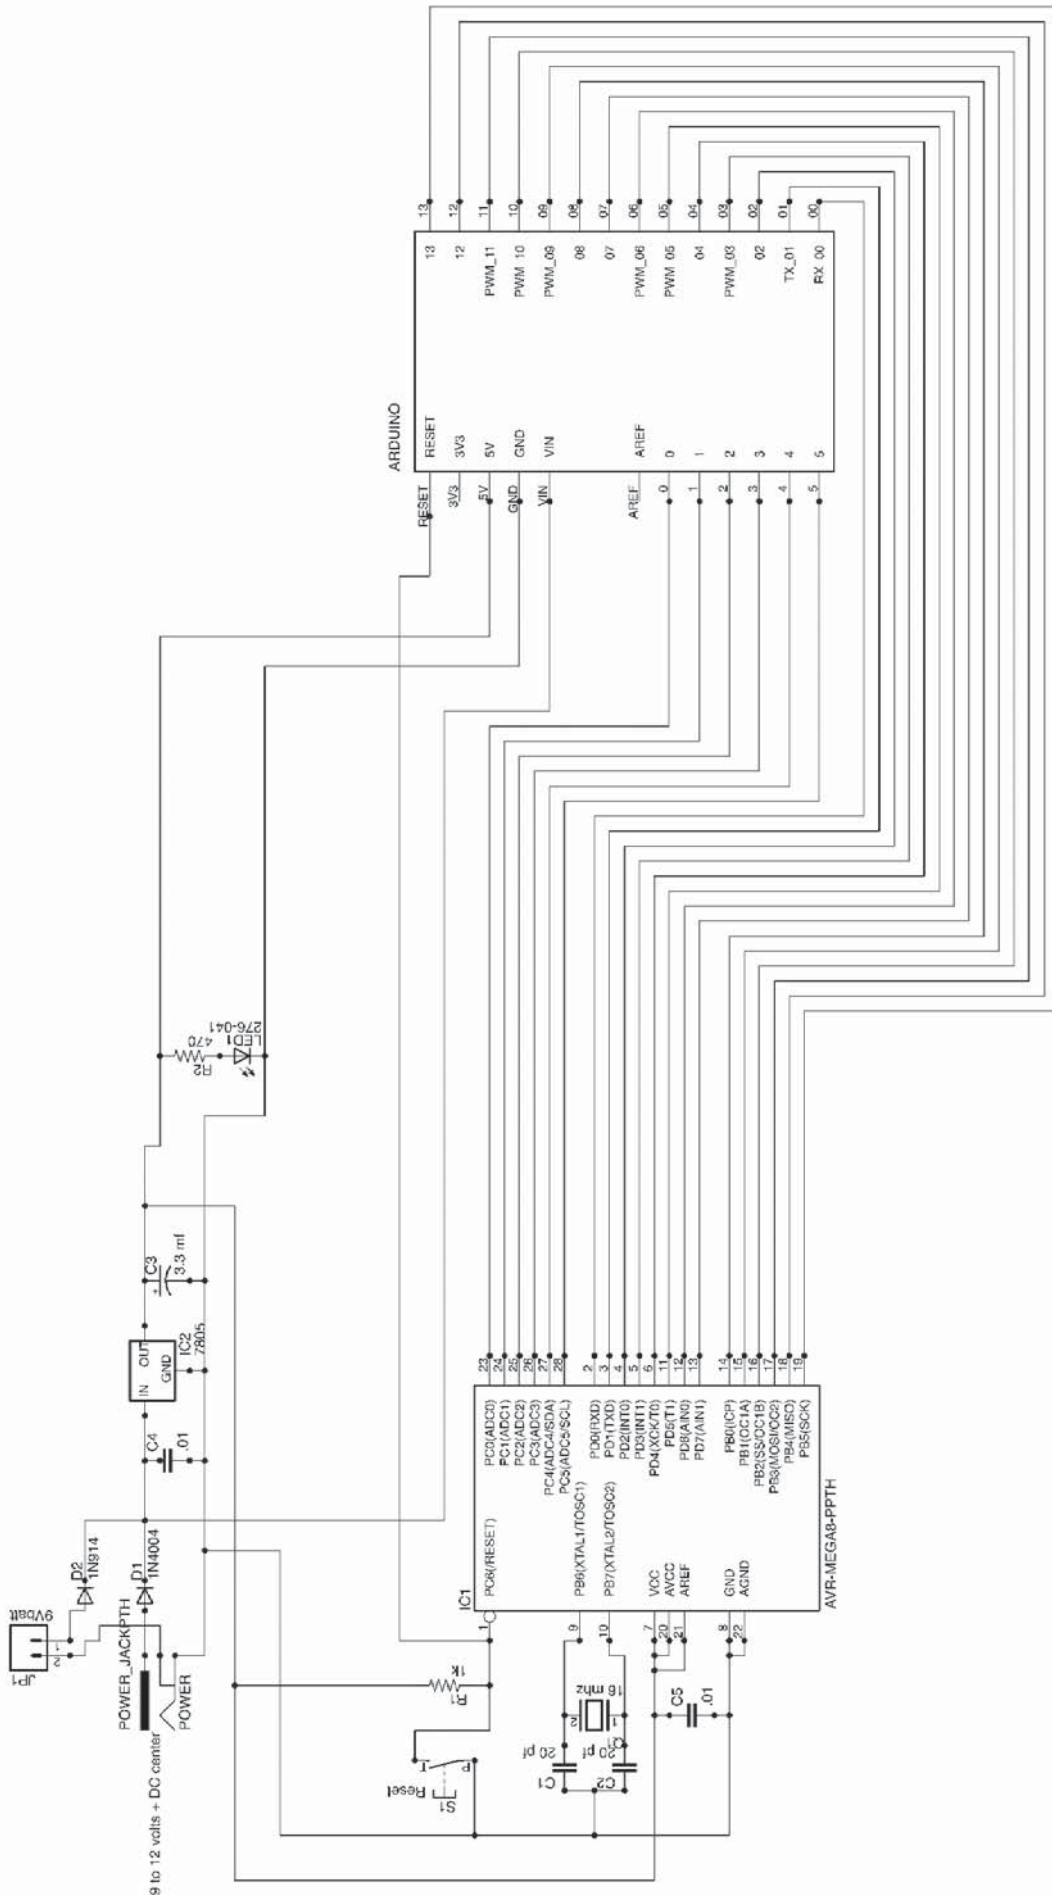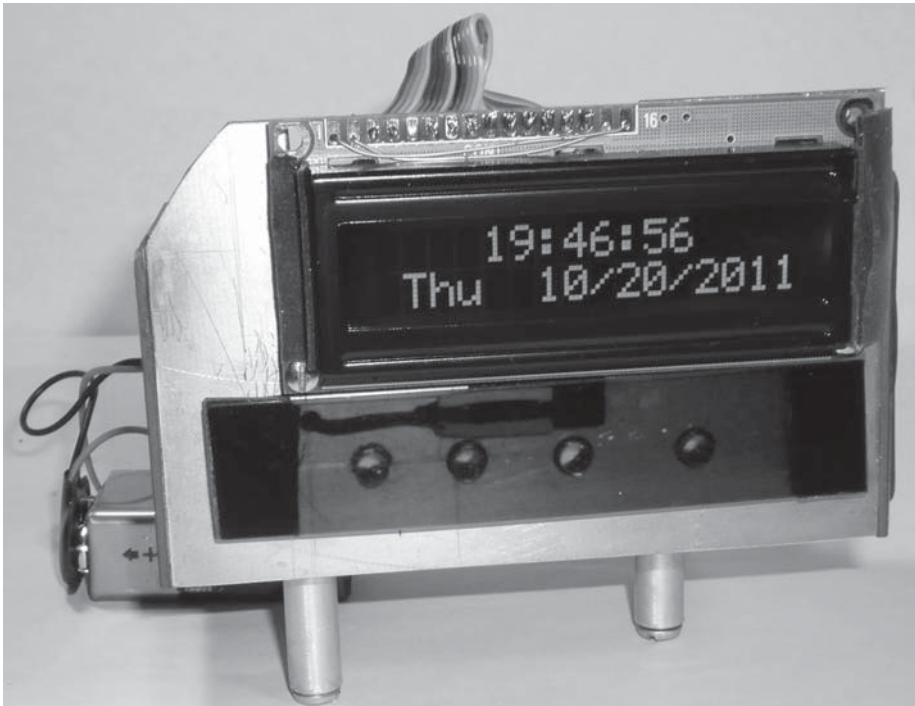Figure 1 – Schematic diagram of the LCD/RTC/Relay shield board.

Figure 2 – Schematic diagram of the You Build It CPU.

The assembled digital clock showing the LEDs and IR diode under the LCD, all running on a 9 V battery.

Suggested software alterations to the code include the following four items (look for *******CHANGE LINES ***** callouts in the code)

```
// **************** CHANGE LINES ***************************
//         Change to Your Call sign
//
// NOTE ************************* CHANGE LINES ******************* //
// Change these values to what you want to set your clock to for the first time.
//
second = 10;
minute = 8;
hour = 22;
dayOfWeek = 2;
dayOfMonth = 14;
month = 11;
year = 11;
// *************************** CHANGE LINES *************************
// You only want to set your clock once and then put the // back in front of the next //   line
//  setDateDs1307(second, minute, hour, dayOfWeek, dayOfMonth, month, year);
//
```

Note: You only want the line of code "setDateDS1307" to run once, and then you must comment it out and reload the software back to the project or the clock will reset to the values above each time the unit is powered up. I didn't add anything in the code to set the clock after it is running.

```
//**********Change if don't want to display Zulu (GMT)************************
lcd.print("Zulu ");  //Default
//lcd.print("      "); // If you don't want to Display Zulu put // if front of the line
above and remove the two at the beginning of this line
```

If you don't like my call at startup change the lines with my call in line 328

```
lcd.print("SB630 RTC w/CW"); // print text and move cursor to start of next line
lcd.setCursor(0,1);
lcd.print("      DE N9APK");
sendmsg("DE N9APK");
```

The Arduino

I used a Duemilanove & UNO Arduino compatible board with an ATmega328 microprocessor (with a bootloader) that I got at **www.jkdevices.com**. The bootloader is code that resides in the processor. Its purpose is to communicate with your computer so that you can load your software.

You'll need an LCD for the display. I started with blue on white and changed to a white on black. The selection you make will be based upon your preferences, as there are a number of combinations available.

The Shield

The microprocessor and software interface to the outside world through hardware shields that sit atop the Arduino board. For this project I used a single shield that I label the LCD/RTC/Relay board. The schematic is shown in Figure 1. I tested the hardware design on a proto jumper board to ensure that it would work with the software. When I was confident that it would work, I assembled the shield using a RadioShack perfboard and wire wrap pins. The pins are long enough to use as male pins to plug into the processor board and allow the components and wire to be soldered to them.

This board connects the LCD panel to the microprocessor and provides a location to mount the Real Time Clock and the logic to control the bright LEDs, as well as the CW speaker, infrared interfacing and two 5-V DIP relays. Point-to-point wiring was used just like the old days, except that I used 30-gauge wire because it saves time when you are only building one unit. You can purchase a supply of wire at RadioShack. It is great stuff to use, but extremely challenging to strip.

I used 5-V DIP relays because they are small. They can switch either ac or dc and isolate the loads from the microprocessor.

If you don't like the perfboard idea you can purchase an "Arduino Protoshield Kit" from many sources on the web. This is a PCB with connectors in the correct locations for the busses, but it is a perfboard just the same, so you will have to use point-to-point wiring. You are working with tiny components, so you'll need a small soldering iron, a sharp tip and small-strand solder.

I utilized a DS1307 Real Time Clock with a battery backup module. It is pre-assembled and offers a small parts count and low cost. The $4 price tag isn't much for a working module compared to the time it takes to round up the separate parts. I found mine on eBay. Make sure that you identify where the seller is located as he might be overseas and the module could take three weeks or more to arrive.
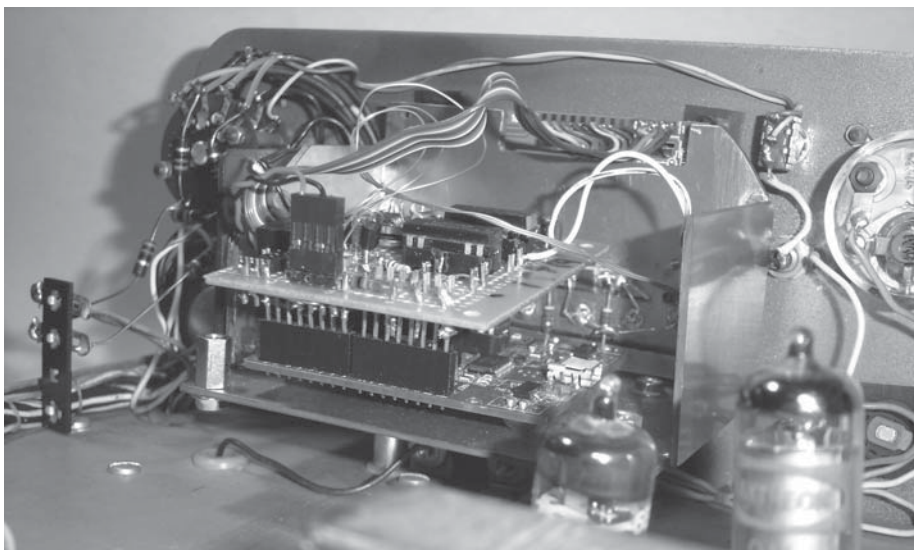
An option added to the second design
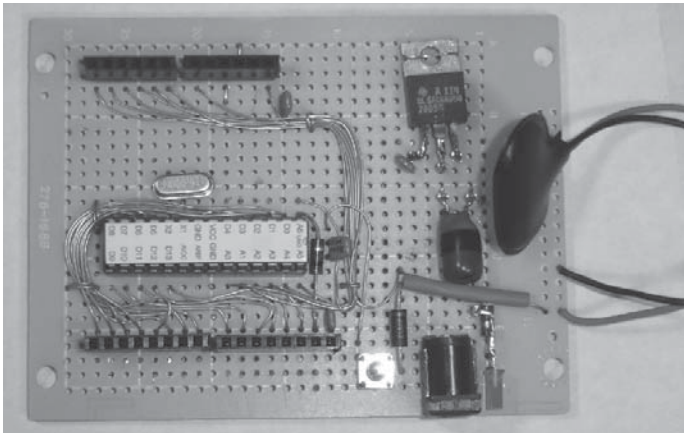
My no-frills You Build It CPU.
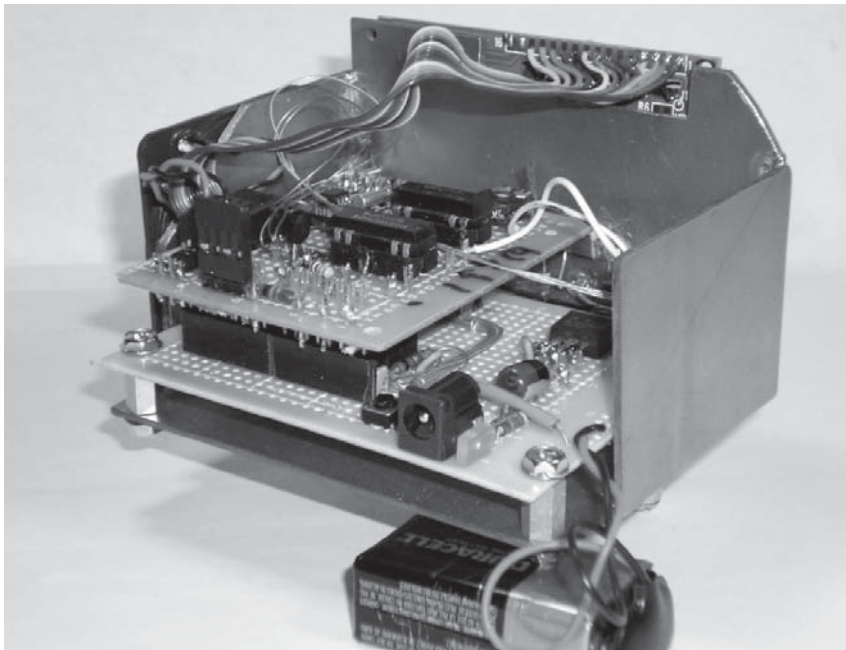
Rear view of the clock showing the relay board on top and the You Build CPU on the bottom.

The completed unit using the full Arduino board for the CPU.

was an ambient light sensor to adjust the backlight of the LCD display. I only pick two levels of brightness to set the LCD. The software shows you how to have the CPU monitor the light intensity in the room. I didn't find a clever place to hide the sensor, so I drilled two very small holes and just stuck it at the top of the bezel (see the photo). It bothers me that I haven't got a better place to hide the sensor. It looks like one would have to cut in to the bezel and I think that wouldn't look better than it does now.

The schematic for the LCD/RTC/Relay board was created with *Eagle CAD* software (a free version is available at **www.cadsoftusa.com/downloads/**). This software can generate a PCB based on the schematic, although I didn't use that option. I've included several *Eagle CAD* files in *Sell.ZIP* for you to try.

An Advanced Concept

I also built a perfboard CPU in the Arduino style – a design I call the You Build It CPU. The Arduino headers and spacing allow the LCD/RTC/Relay board to be plugged in easily. This design has just the minimum parts needed: CPU socket, clock oscillators, power regulator and buss connectors. See the schematic in Figure 2.

This board was also laid out with *Eagle CAD* and you could make a PCB from the file if you wish. There is one jumper (blue, look for the file) on the single sided PCB to be installed.
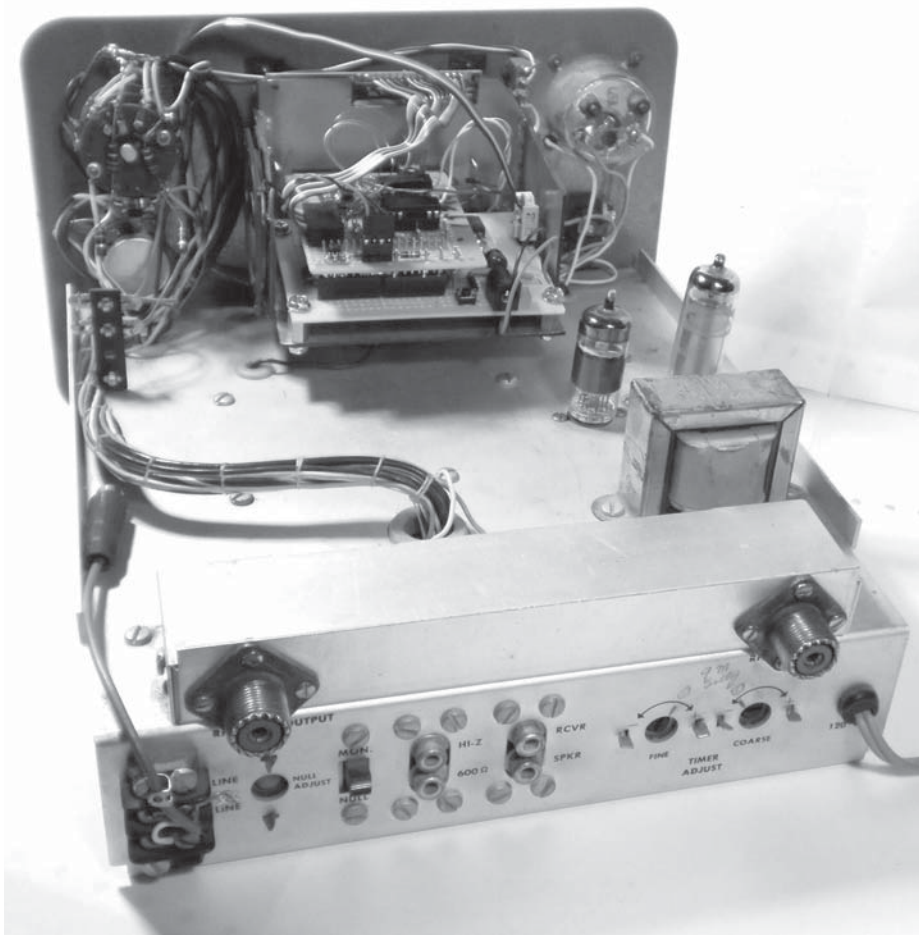
Stand-Alone Version

As I mentioned previously, the project was designed to be an add-on to the SB-630, but this in no way limits it to the Heathkit. You could package it into many other enclosures.

SB-630 Installation

I constructed the housing for the clock/timer out of a simple 4-sided box made of double sided copper-clad board (easy to cut and assemble). I cut one 3 × 4-inch front, one 4 × 2.5-inch bottom and two 2.5 × 2-inch sides. Use a triangle and large soldering gun/iron to assemble the pieces, using the triangle to keep them at right angles to each other where needed.

My goal was to keep the project compact and self-contained, much like the mechanical clock it was going to replace. To that end I wasn't going to use the #47 lights that are original equipment on the SB-630. Instead, I installed four high-bright LEDs under the LCD display to light up the red **IDENTIFY** prompt in the front of the unit. I installed the LEDs by carefully drilling holes on the front of the box. I added the infrared transistor just

The new clock installed.

Copper strips held on with double stick tape isolating the connection to each LED and resistor on the backside of the LCD display board.

above the LEDs and below the LCD cover with the red filter to add more color to the bright LEDs.

The backside of the support that holds the LCD is where the power bus and LED current limiting resistors are located. They are mounted on strips of double sided copper cut in ½-inch pieces and held in place with double stick tape.

Double stick tape will also hold the LCD and the red filter to the copper box that houses the unit. You will need to cut a 1½ × ½ inch notch at the top of the front piece of copper to feed the cable from the LCD to the LCD/RTC/Relay board. Drill two holes on the bottom of the new unit to mount it to the old clock spacers (use a paper template to mark your hole spacings).

Remove the power wires to the clock. You have the option of installing an ac receptacle at the point where the old clock was located so you can plug in the ac adapter for the new clock. I also removed the clear plastic lens that was in front of the old clock to provide a larger opening for the LCD. The LCD will fill the entire opening.

A jumper from SB-630's SWR circuit needs to be routed to the LCD/RTC/Relay board for the automatic ID trigger to work. Connect the jumper to the FWD pin of the SB-630's **FWD/REV SWR** meter switch.

Of course, you must have a transmitter connected to the coax connection on the back of the SB-630 if you want the auto ID function to operate.

Connection Options for the SB-630

In *Sell.ZIP* file on the *QEX* website you'll find a marked-up SB-603 schematic diagram that indicates the tap to sample RF for the Auto ID function.

In *Sell.ZIP* I also include a parts list.

Conclusion

I am quite happy with the results. It brought a purpose back to my old SB-630 and provided me with an identification reminder in a timely manner – not to mention a reliable clock for the shack. Being fully automatic, it is one less thing I have to worry about when I am on the air.

Clark was an active SWL before the age of 14 (WPE3GQF)and was a radio teletype and CW operator in the US Army. His career included 30 + years in computer system and network design, building and repairing systems. Clark also worked for more than 10 years as an electronics engineer/technician for industrial controls companies. Licensed in 1978, Clark is an Amateur Extra operator active in HF, UHF/VHF, computer, digital, satellite, antennas, remote control systems and Arduino projects.
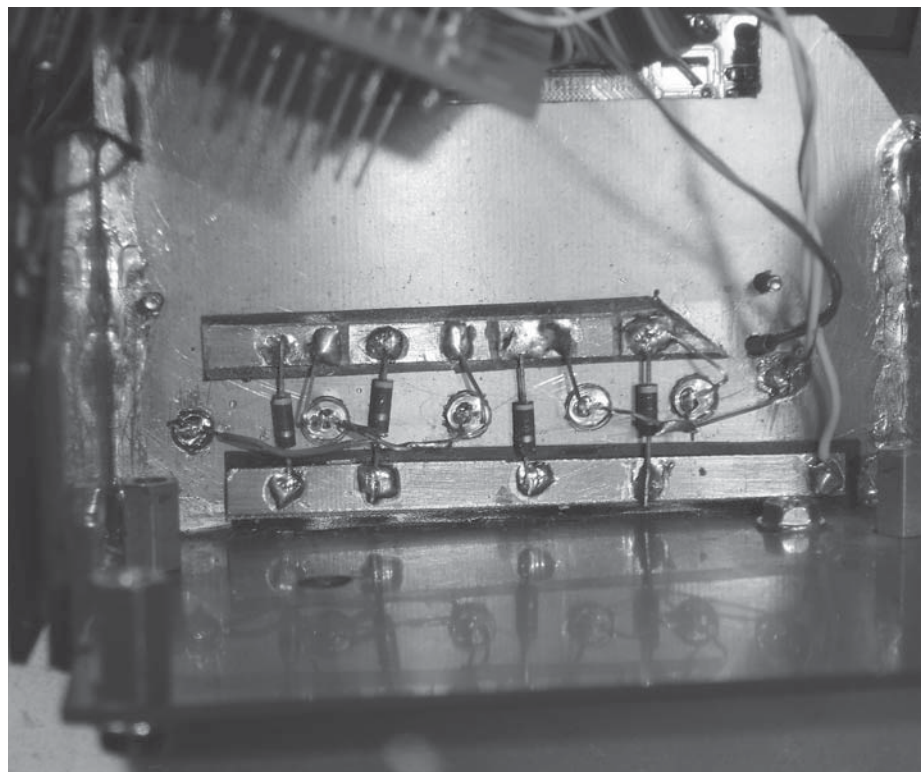
Here is the finished SB-630 showing the new ambient light sensor at the top right of the bezel. Also, the LCD is white on blue. Note that the RCA jack and toggle switch to the left of the bezel are for a modification to the phone patch for a different receiver input, which is not included in this Arduino project.

My complete Heathkit SB station with the newly modified station console.

Maynard A. Wright, W6PAP

6930 Enright Dr, Citrus Heights, CA 95621 **w6pap@arrl.net**

Octave for T- and Pi- Networks

Bring the power of software to matching network design.

In "*Octave* for L-Networks"[1] and "More *Octave* for L-Networks,"[2] we developed methods for designing L-networks that can directly match complex source and load impedances. Can we do the same for T- and pi-networks? Although an L-network will suffice for most impedance matching applications, there are times when we would like to adjust parameters of the network that are not independently controllable in an L-network. Before we proceed, we need to think a bit about what we actually did when we matched complex impedances with an L-network.

In "*Octave* for L-Networks," we set up a system of two equations in two unknowns, specified the impedances to be matched, and solved for the two unknowns, the series and shunt reactive elements of the L-network. Using two equations gave us "two degrees of freedom," or the ability to solve the system so as to control two variables. We chose to specify the real and imaginary (resistive and reactive) components of the input impedance of the L-network and to solve for the L-network elements that would produce that complex input impedance.

When the load impedance is resistive (the imaginary, or reactive, part of the load impedance equals zero), the equations for the elements of the L-network given in "*Octave* for L-networks" reduce to the equations of Sections 5.6.1 or 20.4.2 of the *2012 ARRL Handbook*[3].

The design procedure for L-networks described in Section 5.6.1 of the *2012 ARRL Handbook* notes that the procedure is "bilateral."[4] That is, reversal of the network will allow the source impedance to be transformed into the load impedance. This is true when the source and load impedances are resistive. When either or both impedances are complex, then we must match both the real and reactive components of each impedance to provide a match on both sides. An L-network does not provide sufficient degrees of freedom in its mathematical model

to do that.

The design procedures in "*Octave* for L-Networks" and "More *Octave* for L-Networks," therefore produce an exact match on the input side of the network, which is usually of most interest, and provide matches as specified in Table 1 on the load side of the network, that is, the output impedance of the network seen looking back into the network from the load with respect to the load impedance.

Note that these are fundamental characteristics of the networks involved and other design procedures that take into account complex source and load impedances, such as the circle diagram techniques of B. Whitfield Griffith, Jr., *Radio-Electronic Transmission Fundamentals*,[5] and the Smith Chart design techniques of Phillip H. Smith, *Electronic Applications of the Smith Chart*,[6] will also produce networks that are subject to the impedance limitations of Table 1.

Going Beyond the L-Network

By adding reactive elements to our network, we could design a network that will exactly match complex impedances on each side. Since we must match two resistive components and two reactive components, a network consisting of at least four reactive elements and specified by four simultaneous equations will be required. We really don't have much need, though, in most cases, to exactly match the output impedance of a network and we rarely see matching networks with four elements. Why, then, would we ever use more than two elements in any impedance matching network?

T- and Pi-Networks

The provision of three independent reactances may increase the versatility of a matching network. The use of a pi-network in an amplifier tank circuit allows us independent adjustment of the circuit's resonant frequency and its impedance transformation ratio. The use of a pi-network or T-network in an antenna tuner allows for more versatile adjustment of impedance ratios, often with the use of less expensive components, than does an L-network.

In addition, there are characteristics of a matching network other than impedance ratios and resonance that we might want to consider. Sections 5.6.2 and 5.6.3 of the *2012 ARRL Handbook*,[3] for instance, include design procedures for pi- and T-networks that provide a specified Q for each network. Such a design is desirable if the network must exhibit a particular performance over a range of frequencies, to provide, for example, some minimum degree of harmonic suppression, as well as matching the source impedance.[7]

In the *Radio Engineers' Handbook*[8], Terman takes a different approach and specifies a pi- or T-network that exhibits a specified phase shift at a particular frequency of interest. Such a design is required when the phase shift is important, such as when feeding signals to multiple elements of a directive antenna system.

Note that both the *ARRL Handbook* and Terman provide design procedures that will handle only resistive source and load impedances. The *ARRL Handbook* recommends using series reactances to counter the effects of any reactive components of the source or

Table 1
Output Impedance Matching Condition

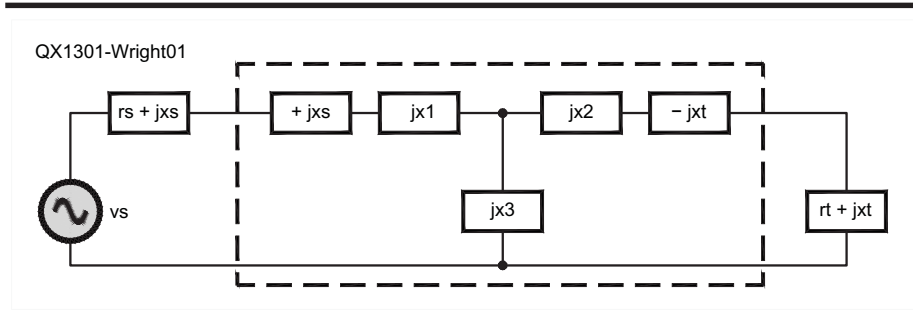| Input Z	Output Z	Output Impedance
REAL	REAL	DIRECT MATCH
REAL	COMPLEX	CONJUGATE MATCH
COMPLEX	REAL OR COMPLEX	NO MATCH

[1]Notes appear on page 42.

Figure 1 – T network diagram. Note that +*j*xs and *j*x1, as well as -*j*xt and *j*x2, are combined in the output of the code in Table 2 as we don't need separate capacitors or inductors in the physical network to handle the functions of these series components.
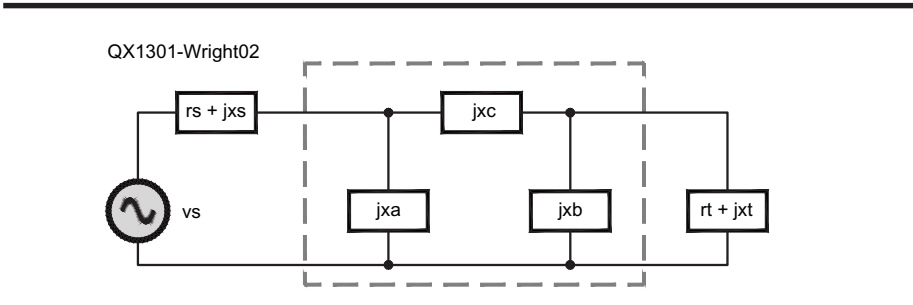

Figure 2 – Pi network diagram.

load impedance after designing the network for the resistive component only. As we saw earlier, this will provide a conjugate match, rather than a direct impedance match and, if used, will change somewhat the Q and/or the phase shift of the network. In most cases where Q is important, it is sufficient to design a network that will exhibit some minimum Q, so the resistive design procedures with added reactances will normally be sufficient to produce such networks by ensuring that the Q is sufficiently high to allow the additional reactances or by iteratively adjusting the design until a satisfactory Q is achieved.

In the case of the network where phase shift is important, however, we are generally interested in accurately producing the desired shift. Adding extra reactances to the legs of such a network may affect its phase shift to an unacceptable extent. If we simply add extra compensating reactances, we'll need to loop through multiple iterations of the design to home in on the desired phase shift. We'll see whether we can instead use *Octave* to produce a network that will provide a complex impedance match while producing exactly the phase shift we need.

The usual convention for networks such as those we are dealing with is to label the elements as impedances, such as "Z1, Z2, and Z3." Here we must consider how the various complex impedances of the network, the source, and the load interface with each other so, in Figures 1 and 2, we have labeled the elements of the networks as reactances ("*j*x1,

*j*x2, and *j*x3") and have included the *j* operator to emphasize that they are not real. Note that +*j*xs and *j*x1, as well as -*j*xt and *j*x2, in Figure 1 are combined in the output of the code in Table 2 as we don't need separate capacitors or inductors in the physical network to handle the functions of these series components.

Using *Octave* to Design T- and Pi-Networks

We'll begin by designing a T-network. Generating the T-network before the pi-network is easier as we can simply add the additional reactances algebraically to the series arms of the T-network and then perform a T- to pi-network conversion to obtain the corresponding pi-network.

Since we are interested in matching the source impedance exactly, we'll do that with our added reactance on the source side. On the load side we'll maximize power transfer and make sure that the resistive design using Terman's equation will still match our input impedance by using a conjugate match.

It might be tempting to think that we can ignore the phase shift of the added reactance on the load side of the network because it will just offset the reactance of the load. We are, though, interested in the phase shift of just the network itself, so our reference point for the output of the network is at the input to the load impedance and not somewhere inside the load so as to include the reactive part of the load impedance in the calculation.

The formulas for converting the T-network to a pi-network are from *Reference Data for Radio Engineers*.[9] Note that these conversions preserve the phase shift of a network as well as its impedance transformations. They will not, however, preserve the Q of a network during transformation as the conversion is valid at only the frequency at which the reactances of the T-network are specified[11].

The *Octave* code that generates the T- and pi-networks is listed in Table 2. We'll first accept the real and imaginary parts (resistance and reactance) of the source and load impedances, and our desired phase shift, from the keyboard. Phase shifts in this code can be positive or negative. We'll convert the phase shift in degrees to radians as *Octave's* trig functions expect arguments in radians:

ph *= pi / 180;

Next, we'll calculate the additional phase shift that will be added when we add our matching reactances to the series legs of the T-network:

phin = –arg(rs / (rs + 1j * xs));
phout = arg(rt / (rt + 1j * xt));

Now we'll subtract the phase shifts of the added reactances from the desired phase shift so that when we design the T-network and add those elements in, the phase shift will have the correct value:

ph –= phin + phout;

We'll then design the T-network using formulas from Terman[8]:

x1 = -(rs * cos(ph) - sqrt(rs * rt)) / sin(ph);
x2 = -(rt * cos(ph) - sqrt(rs * rt)) / sin(ph);
x3 = -sqrt(rs * rt) / sin(ph);

Then we'll add the matching reactances to the series legs of the T-network so as to provide a direct impedance match on the input side and a match as shown in Table 1 on the output side:

x1 + = xs;
x2 – = xt;

We'll then print out the T-network elements. The T- to pi-network conversion equations from *Reference Data for Radio Engineers*[9] are designed to accept impedance values rather than reactances, so we'll convert the real numbers representing the reactances of our T-network to complex impedances whose real parts are zero:

z1 = 0 + 1j * x1;
z2 = 0 + 1j * x2;
z3 = 0 + 1j * x3;

We'll then make the conversion to the pi-network:

zc = (z1 * z2 + z1 * z3 + z2 * z3) / z3;
za = (z1 * z2 + z1 * z3 + z2 * z3) / z2;
zb = (z1 * z2 + z1 * z3 + z2 * z3) / z1;

Note that, when we print out the results, we use imag(zc), imag(za), and imag(zb) to print out only the imaginary part of the complex number that represents the impedance of each pi-network element. Limited precision in *Octave's* representation of the numbers

involved in these calculations make it likely that each complex number will have a small residual real part, somewhere around 1e-14, that we'll discard.

Alternative Design Considerations

The T-network or pi-network resulting from this design procedure will exhibit a matched impedance toward the source, an impedance as listed in Table 1 toward the load, and a phase shift as specified. With some minor changes to the code of Table 2, we could use the same techniques to design a network of specified phase shift that will provide a conjugate match on each side.

Note that we have no control over the polarities of the elements in the networks we are designing when we use the code in Table 2 to design for a particular phase shift. We may end up with any combination of positive and negative reactive values that must be implemented by various, sometimes "odd," combinations of inductors and capacitors.

Although there is nothing fundamentally wrong with such networks, we might want to use a particular configuration to minimize component expense or to provide some low-pass filtering function while matching impedances. Figure 5.58 of the *2012 ARRL Handbook*[3] shows fourteen network configurations along with equations to specify their elements.

The code in Table 2 cannot implement those equations, but we can use our code to force a particular network configuration if we simply want impedance matching and are not concerned about the value of the phase shift through the network. Figure 74 of the *Radio Engineers' Handbook*[8] shows eight T- and pi-network configurations, along with phase shifts that will force the elements to assume particular polarities. If, for instance, we want a pi-network with a series inductor and parallel capacitors, we might specify a large negative phase angle. If we want a T-network with series capacitors and a parallel inductor, we would use a large positive phase angle. Other forced configurations are possible, but those we have considered here are the most common.

"Large" and "small" phase angles as specified in Figure 74 of the *Radio Engineers' Handbook* are somewhat subjective and the best value in each case is dependent on the ratio of impedances to be matched and the ratios of real to imaginary components of both source and load impedances. Using the Table 2 code in this way is somewhat analogous to the use of the parameter N in the older *ARRL Handbooks*[10] to optimize element values.

Conclusions

Using GNU *Octave*, we are able to design a T- or pi-network that matches complex source and load impedances and that produces a specified phase shift or a particular network

Table 2
GNU *Octave* Code for T- and PI-Networks

```
#!/usr/bin/octave
printf("\n\n    OCTAVE T- AND PI-NETWORK DESIGN SCRIPT");
rs = input("\n\n        ENTER SOURCE RESISTANCE: ");
xs = input("        ENTER SOURCE REACTANCE: ");
rt = input("  ENTER TERMINATING RESISTANCE: ");
xt = input("    ENTER TERMINATING REACTANCE: ");
ph = input("  ENTER PHASE SHIFT IN DEGREES:  ");
ph *= pi / 180;
phin = -arg(rs / (rs + 1j * xs));
phout = arg(rt / (rt + 1j * xt));
ph -= (phin + phout);
x1 = -(rs * cos(ph) - sqrt(rs * rt)) / sin(ph);
x2 = -(rt * cos(ph) - sqrt(rs * rt)) / sin(ph);
x3 = -sqrt(rs * rt) / sin(ph);
x1 += xs;
x2 -= xt;
printf("\nT-NETWORK ELEMENTS:\n");
printf("X1 = %g\n", x1);
printf("X2 = %g\n", x2);
printf("X3 = %g\n", x3);
z1 = 0 + 1j * x1;
z2 = 0 + 1j * x2;
z3 = 0 + 1j * x3;
zc = (z1 * z2 + z1 * z3 + z2 * z3) / z3;
za = (z1 * z2 + z1 * z3 + z2 * z3) / z2;
zb = (z1 * z2 + z1 * z3 + z2 * z3) / z1;
printf("\nPI-NETWORK ELEMENTS:\n");
printf("XA = %g\n", imag(za));
printf("XB = %g\n", imag(zb));
printf("XC = %g\n\n", imag(zc));
```

configuration. We have a code base in *Octave* that will allow us, with modifications, to optimize other T- or pi-network parameters that might be of interest.

Notes

[1]M. Wright, "*Octave* for L-Networks," *QEX*, Mar/Apr, 2011.

[2]M. Wright, "More *Octave* for L-Networks," *QEX*, May/June 2012.

[3]*The ARRL Handbook for Radio Communications*, 2012, The American Radio Relay League, Inc. 2011, Sections 5.6.1 and 20.4.2.

[4]Note that the bilateral constraint of Section 5.6.1 of the *2012 ARRL Handbook* applies when the impedances on both sides of the L-network are real, but not when one or both are complex. The reciprocity theorem as stated in the *Radio Engineers' Handbook*,[8] Page 198, does apply to the L-network regardless of the impedances being transformed. We can verify this by replacing the source voltage by a short and by inserting an equivalent voltage in series with the load. Regardless of which configuration of Table 1 is in use, the current flowing in the load before the substitution and the current flowing in the source impedance afterward will be identical, verifying compliance of the network with the reciprocity theorem.

[5]B. Whitfield Griffith, Jr., *Radio-Electronic Transmission Fundamentals*, McGraw-Hill, 1962: reprint available from ARRL at **www.arrl.org**.

[6]Phillip H. Smith, *Electronic Applications of the Smith Chart*, McGraw-Hill, 1969: reprint available from ARRL at **www.arrl.org**.

[7]Some earlier editions of the *ARRL Handbook*[10] feature a design for a pi-net-

work using parameter N. The text advises choosing N to "optimize circuit component values." Rearrangement of the equations for the elements of the pi-network in one edition or the other will show that N in the earlier *Handbooks* is the same parameter as Q in the more recent editions.

[8]Frederick E. Terman, *Radio Engineers' Handbook*, McGraw-Hill, 1943, Section 3, Paragraph 25.

[9]International Telephone and Telegraph Corporation, *Reference Data for Radio Engineers*, Fourth Edition, International Telephone and Telegraph Corporation, 1956, page 142.

[10]*The 1978 Radio Amateur's Handbook*, ARRL, 1977, page 54.

[11]William Littell Everitt, E.E., Ph.D., *Communication Engineering*, McGraw-Hill, 1937. Note, however, that for small phase shifts and some impedance ratios, the T- to pi-network conversion, or the inverse, may provide a network that approximates closely the Q of the source network.

Maynard Wright, W6PAP, was first licensed in 1957 as WN6PAP. He holds an FCC General Radiotelephone Operator's License with Ship Radar Endorsement, is a Registered Professional Electrical Engineer in California and is a Life Senior Member of IEEE. Maynard has been involved in the telecommunications industry for more than 49 years. He has served as technical editor of several telecommunications standards and holds several patents. He is a Past Chairman of the Sacramento Section of the IEEE. Maynard is Secretary/Treasurer and Past President of the North Hills Radio Club in Sacramento, California.

QEX~

Ray Mack, W5IFS

17060 Conway Springs Ct, Austin, TX 78717; **w5ifs@arrl.net**

# SDR Simplified

A look at High Speed Conversion

A Recap

Until now we have looked at projects that we could buy with pin money (Blackfin Stamp, PC sound card, TI DSP stick). The Blackfin Stamp at $235 was the most expensive of the development platforms, but the most capable. None of the easy implementations of those systems could process much above about 1 Msps. The holy grail of SDR is to hook the antenna relay to the DAC output and the ADC input. A decade ago, such a system was just barely possible with government money. You still cannot achieve the goal with your lunch money, but it is available at the price of recent very high end analog amateur transceivers.

Commercial High Speed Hardware Implementations

Flex Radio has just announced the Flex 6000 series of radios that range from $4500 to $7500. For comparison, some of the high end Yaesu radios can cost over $10,000. The Flex 6000 series radios cover 300 kHz to 77 MHz (to 165 MHz for the 6700) completely in the digital domain. The ADC operates at 246 MHz and the DAC operates at 492 MHz. The specifications for this radio are very impressive. The engineering and craftsmanship is top-notch. Much of the technology in the radio is a spinoff of government projects. This radio is a closed system and not amenable to direct experimentation by the owner, but it shows what is possible at the very high end of technology. Although it is a closed system, Flex is working with other developers to expand the capabilities of the system. They are also working on an interface specification and a software interface to allow users and third party developers to control some of the proprietary portions of the design. This will allow software applications and client tools to control the radio using various aspects of the internal hardware and software as building blocks.

The FPGA that Flex uses is the Xilinx Virtex-6 part number XC6VLX130T. This part has 480 DSP blocks which are multiply accumulate elements for FIR or IIR imple-
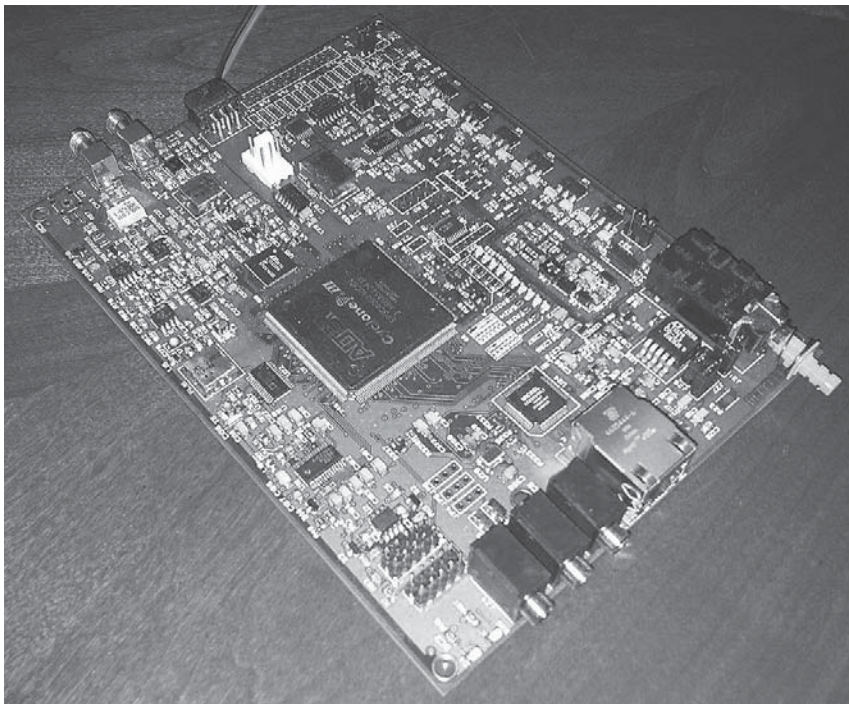
(A)
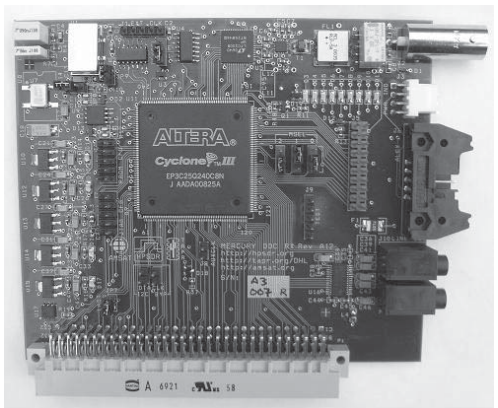
(B)
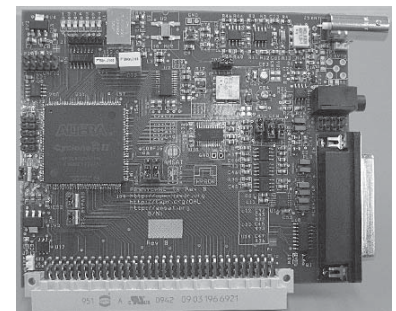
(C)

Figure 1 – Photos of the available boards from the HPSDR project and TAPR. Photo A is the Hermes transceiver board using a 3C40 FPGA. Photo B is the Mercury receiver using a 3C25 FPGA. Photo C is the Penny Lane transmitter using a 2C8 FPGA.

mentation. Flex partitions those between eight receive channels that operate totally in parallel. A big part of what makes the Flex implementation unique is the extensive engineering work to produce those parallel receivers. A few of the DSP blocks are dedicated to the preselecter and down converter tasks to reduce the data rate from 246 Msps to several Msps. The rest of the DSP functionality is time multiplexed to use all of the time available in the FPGA resources to reduce the data rate to 192 ksps for each of the receivers. The result is that the 6500 implements up to 191 billion MAC operations per second and 78 billion floating point operations per second. The $4500 price for the radio is pretty reasonable when you see that the FPGA is $900 in small quantities and the ADC is $130 in small quantities.

A more economical and open system is the HPSDR project which is developed by a group of amateurs in association with TAPR. It is not a part of TAPR, but shares resources with them. The Flex system is a fully functioning radio that lets you enjoy the DSP advantages of a radio without your technical involvement; the system from the HPSDR group gives you the DSP functionality with all of the resources exposed and requires you to be intimately involved in its operation and configuration. It is intended as an experimenter system, and you are expected to be your own support system if you buy HPSDR hardware. The HPSDR system is aimed specifically at amateurs who want to experiment with high end DSP in radios. The latest board (Hermes) combines both a Digital Up Converter for transmit and a Digital Down Converter for receive in a single FPGA. It combines the functionality of Mercury (Digital Down Converter receiver), Penny Lane (Digital Up Converter transmitter), Metis (high speed PC interface), and Excaliber (frequency stability card). Fig-ure 1 shows photos of Hermes, Penny Lane, and Mercury.[1] The frequency range for all of the HPSDR systems is currently 50 kHz to 55 MHz. The Hermes board is no longer available as an assembled and tested product. However, you can buy a bare board for $40 and build your own system. It is composed of mostly surface mount parts, so assembly on your own might be daunting. At least the FPGA used is a PQFP gull wing package, so you don't have to deal with the issues of a ball grid array. Digi-Key carries the part in stock (at least as of this article) for $73. The FPGA is the Altera EP3C40 device from the Cyclone III family that has 126 multipliers available. One implements FIR filters in this FPGA using the multiplier elements and additional logic to create multiply accumulator blocks.

You can still buy an assembled and tested Penny Lane board for $399 and a Mercury board for $469. You will also need a few extra boards to round out the system as well as a PC for control. So, if you want to start with something reasonably assured of working, it will cost around $1000. The Mercury receiver uses an Altera 3C25 series FPGA which has 66 multipliers. The Mercury with appropriate firmware will only support two simultaneous receivers where Hermes can support up to eight. Penny Lane uses an EP2C8 from the Cyclone II family which has more than enough resources to handle up converting to 55 MHz using CIC filters and a numerically controlled oscillator. If you are comfortable ordering internationally from India, Apache Labs sells a complete Hermes based transceiver for $895.

The next level down is a selection of FPGA evaluation kits. Xilinx and Avnet have partnered to develop an evaluation kit based on the Zynq 7000 series FPGA. Fig-ure 2 shows a photo of this board.[2] The Zed Board is available for $399 ($315 for qualifying student or academic use). This is an incredible price since the Zynq 7020 sells for about $245 in single quantity. The Zynq chip contains a full FPGA as well as a dual core ARM 9 and associated MCU peripherals (1 Msps ADC, SPI, I2C, USB, etc.). The FPGA contains regular I/O and gates as well as dedicated DSP elements. The Zed Board uses the XC7Z020-1CLG484CES part. The part contains 220 of the Xilinx DSP slices (presumably the same as used in the Virtex parts). The Zed Board package contains the free web version of the Xilinx *ISE FPGA* development software. You can download the free version from the web as well. The support model for the Zed Board is similar to the Blackfin Stamp and HPSDR models. Primary support is through the user community at **zedboard.org** using *Linux* and *Android* as the operating systems. The user community is just coming up to speed
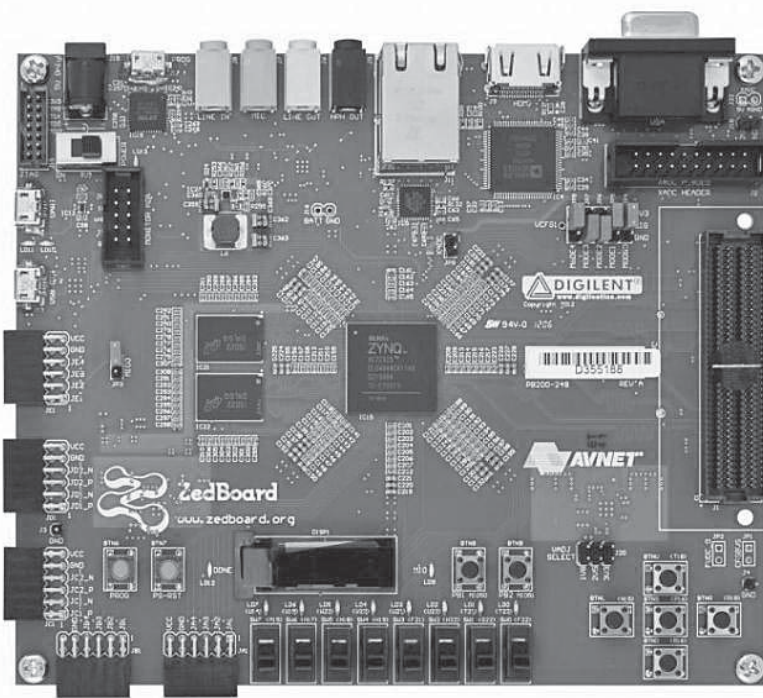
Figure 2 – Photo of the Zed Board Zynq FPGA/MCU evaluation board from ZedBoard.org
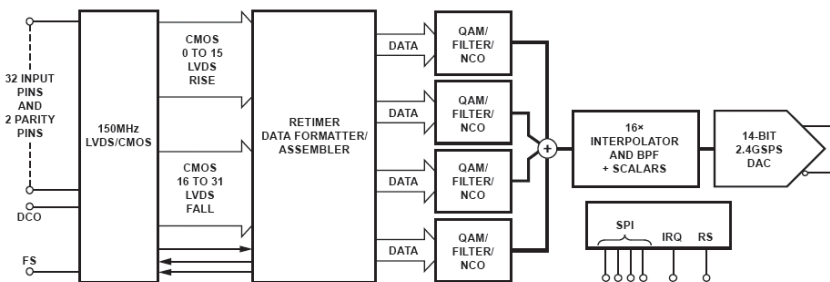
Figure 3 – The block diagram of the Analog Devices AD9789 Digital Up Converter that operates up to 2.4 Gsps. This part is capable of RF output from DC to 3 GHz.
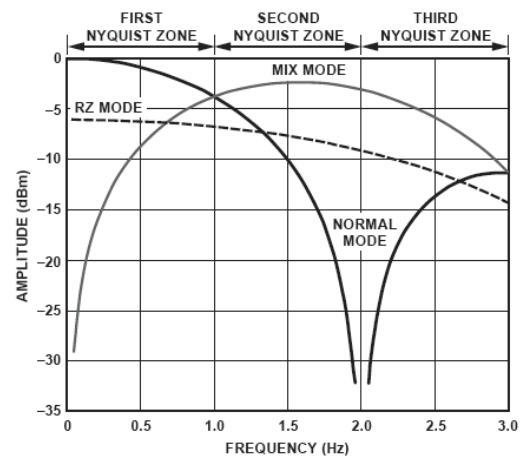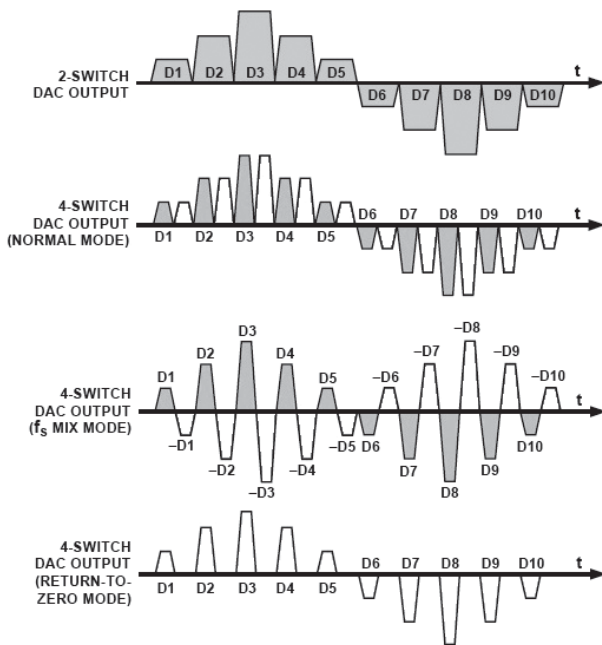
[1]Notes appear on page 47.

Figure 4 – Waveforms comparing an ordinary two switch DAC with the three modes available in the AD9789. The graph shows the frequency ranges and relative output for each of the three modes.
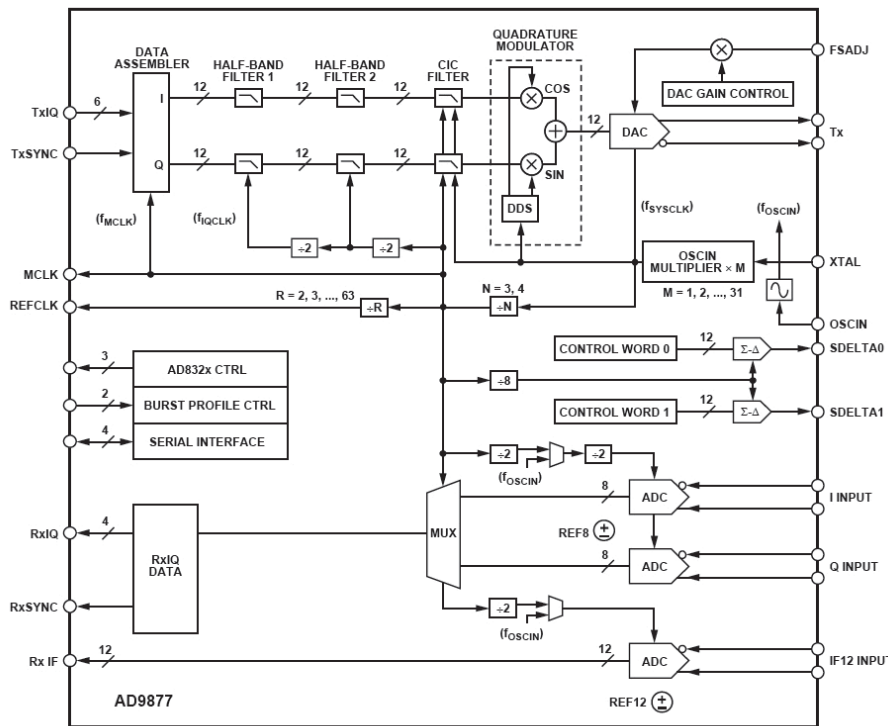

Figure 5 – The block diagram of the Analog Devices AD9877 Cable Modem IC. The transmit chain has Digital Up Conversion. The receive section has very capable 33 Msps ADCs, but no digital down conversion.

since the system has only been readily available since October 2012. I expect the learning curve for the Zed Board *Linux* to be equivalent to the adventures we experienced with the Analog Devices Blackfin Stamp.

The Zed Board is a full computer with DDR memory, HDMI video, Ethernet, sound card I/O, and flash "hard disk" among others. The board also has a set of FMC connectors for attaching accessory boards. You can see those connectors on the right side of the photo in Figure 2. Once again, if you have a large

budget you can do high speed SDR work using this platform and one of two SDR accessory cards. The VITA 57.1 industry standard for FMC connectors allows you to connect a large variety of third party accessory boards to the Zed Board.

Analog Devices is producing an SDR kit that plugs into the FMC connectors on the Zed Board. There are two options for using the Analog Devices board. The first is to buy a complete kit from Avnet for $1450 that includes the Zed Board, the Analog Devices FMCOMMS1-EBZ, and associated pieces. The second is to buy the Zed Board from Avnet for $399 and buy the FMCOMMS1-EBZ from Analog Devices for $750. It is unclear to me why the combined kit from Avnet would be $300 more expensive. The FMCOMMS1-EBZ is just available for order in early November with a six week lead time. The FMCOMMS1-EBZ is intended to demonstrate wireless capabilities from 400 MHz to 4 GHz and is not specific to the Zed Board.

Another SDR kit for use with the Zed Board is available from 4DSP (Austin, TX). This system is also intended to support SDR for wireless digital applications. One model of the FMC30RF operates from 400 MHz to 1.2 GHz and the second operates from 1.2 GHz to 3 GHz. 4DSP has a large number of other DSP boards that will also connect to the Zed Board. Their focus is on commercial applications of DSP, so the $1195 cost of the board reflects that focus. They are pleased to sell in single unit quantities, though. 4DSP supplies firmware for both Xilinx and Altera FPGAs to implement digital up conversion or digital down conversion in addition to its hardware offerings. The FMC30RF uses TI wire-

less parts for digital up conversion and digital down conversion. 4DSP sells both hardware and software for SDR and DSP. They have IP cores for FFT, video, and SDR for implementation in FPGAs.

You can experiment with the same series of FPGA as used in the Flex 6000 series with the SP 601 evaluation kit for the Virtex 6 family. This board is $295 from Xilinx and uses the same FMC connector as the Zed Board. The free *ISE Web Pack* software supports this board. The SP 605 board is even more capable, but it is $495. It includes a node locked version of the higher performance *ISE Logic Edition* and other debug/measuring tools.

If you would like to experiment with the Altera EP3C25 series that is used in the HPSDR Mercury board, you can buy the Cyclone III FPGA Starter Development Kit for $199. This board uses a connector similar to the FMC system from Xilinx, but they are not the same. The intent for high speed interconnect is the same, though. The free version of Altera development software is called *Quartus II Web Edition*. Altera has a software package called *DSP Builder* that works in combination with *MATLab* and *Simulink* from MathWorks, but all three packages must be purchased, and they are quite expensive. The FPGA code from HPSDR is a good starting point for any experimentation you might want to do. All of the FPGA code for Mercury, Penny Lane, and Hermes was generated using the free version of *Quartus II*.

Dedicated High Speed Hardware

TI and Analog Devices each produce parts that will perform the digital up or down conversion. The parts are all aimed primarily at the wireless industry.

Analog Devices has three digital up converter ICs that differ primarily in the internal clock rate with corresponding differences in price. The AD9856 runs at 200 Msps, has a 12 bit DAC, can create signals up to 80 MHz and costs $35. The AD9857 is an improved version that has a 14 bit DAC and costs just $34. The next step up is the AD9957 that runs at 1 Gsps, has a 14 bit DAC and costs $36. Since I do my own board layout and assembly, these parts are very appealing because they have TQFP gull wing packages. I am not ready to tackle ball grid array (BGA) packages just yet.

The fourth member of the Analog Devices digital up converter family is the AD9789 that runs at 2.4 Gsps, has a 14 bit DAC, and costs $73. The part can produce signals up to 800 MHz in the primary Nyquist zone. This part is different from the other three in that it has provisions for producing output in the second and third Nyquist zones. That means the part can produce output from DC all the way through UHF! The down side is that it is only available in a BGA package. That is not an insurmount-

able problem, but it effectively increases its unit cost due to the need for specialized assembly tools. Figure 3 shows the block diagram of the part.

The DAC output uses a four switch driver rather than the two switch driver that most DACs use. One can select one of three DAC drive modes to obtain the desired type of output. Figure 4 shows representative output waveforms for the three modes and the resulting output spectrum. Normal Mode works in a manner similar to an ordinary two switch DAC as a zero order hold circuit. In a normal DAC the switches are supposed to change at exactly the same time. However, if there is any delay between the channels (a certainty and

dependent on chip design, etc.), one switch can change before the other and produce a glitch that is dependent on the initial state of the switches. The glitch energy varies based on the codes and increases noise in the output. The four switch method used in the AD9789 overlaps the switching in the middle of a sample time and at the edge of a sample time. This eliminates the code dependent glitches, but creates a constant "glitch" at twice the sample frequency. The part implements a Mix Mode output that chops the output at twice the DAC frequency with a pulse train of alternating values of 1 and -1 as we would see in a double balanced mixer. The waveform looks exactly like a DSB-SC waveform from a mixer driven
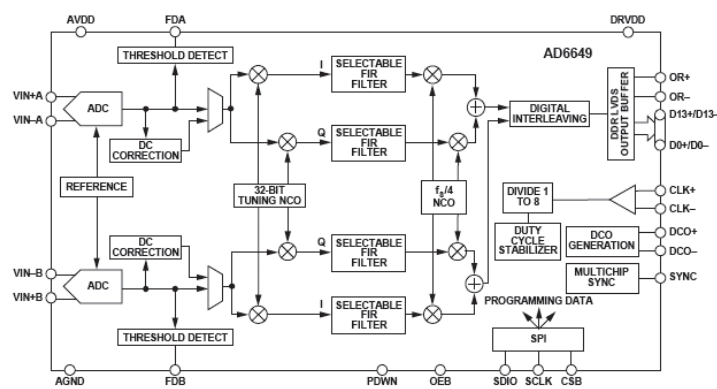
Figure 6 – The block diagram of the Analog Devices AD6649 Digital Down Converter analog front end IC. It has two independent I/Q channels with digital mixing for use in diversity reception applications.
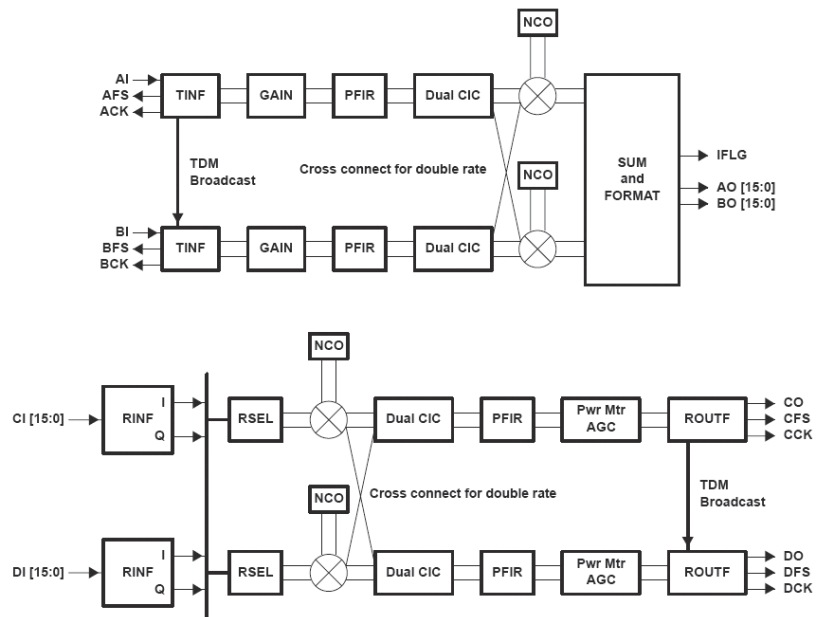
Figure 7 – The block diagram of the TI GC5016 when configured to use two channels for transmit and two channels for receive. Channels A and B are used for transmit and Channels C and D are used for receive.

by a square wave. The third mode is called Return to Zero mode and operates as if the output were multiplied by a pulse waveform with values of one and zero much like a single balanced mixer.

The AD9879 and AD9877 are from the DOCSIS cable modem product family and cost $21. The DOCSIS transmit channel is generally in the low VHF range (up to 65 MHz) with the receive channel occupying one of the normal cable TV channels. The receive channel has two 33 Msps ADCs but no digital down conversion. The transmit channel does include a fixed 16X digital up conversion in the AD9879 and programmable 12X or 16X in the AD9877. Each part contains a DDS oscillator and quadrature mixer for frequency translation. Figure 5 shows the block diagram of the AD9877.

The Analog Devices web site is a bit frustrating when trying to find Digital Down Conversion ICs. A search for digital down converter brings you to a product matrix that shows the AD6620 through AD6636, but all of those parts are not recommended for new designs. Being more creative in searching leads you to the AD6649 which appears to be the only production digital receiver IC in the Analog Devices portfolio. The block diagram is shown in Figure 6. The ADC can sample up to 250 Msps. It is interesting that one of the signal to noise specifications is measured at 245.76 Msps which is the exact frequency used by Flex. That frequency is 1280 times 192 kHz, so there is a good engineering reason for that sample frequency. The part is useful from DC to 400 MHz input frequency. Of course, frequencies above 125 MHz operate the part in sub-sample mode and require analog filtering to eliminate image frequencies. The part is available from Digi-Key in single quantities for $118. The package is a 64 pin lead-less chip scale package, so soldering should not be especially difficult.

Analog Devices has a part that is a hybrid between an analog and fully digital part in the AD9874 which costs $28. It contains an analog down converter (LNA, mixer, and PLL oscillator) to reduce the input signal from the input range (10 MHz to 400 MHz) down to approximately 3 MHz. The 3 MHz signal is converted to digital in a Σ-Δ ADC. The digital data is decimated in a CIC filter which can have a value of 60 x M or 48 x M (where M is a value between one and 16). Such large decimation values will provide very manageable sample rates for final filtering and demodulation.

TI has a large portfolio of digital RF parts. Like Analog Devices, some of their digital down/up converter parts are in production but not recommended for new designs. My opinion is one of the nicest parts they have is the GC5016. It is not supported for new designs but is a full production part. The evaluation module is no longer available, so you would
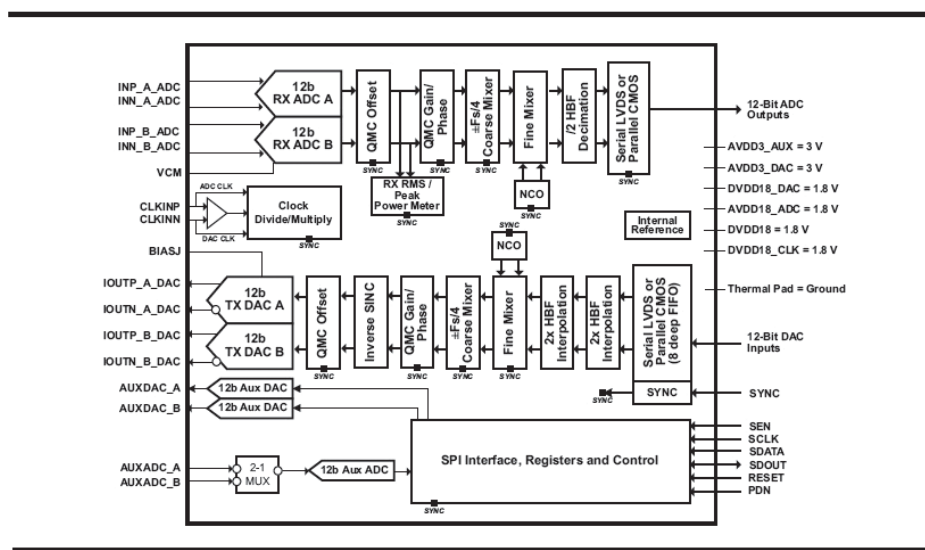
Figure 8 – The block diagram of the TI AFE7225 transceiver IC.
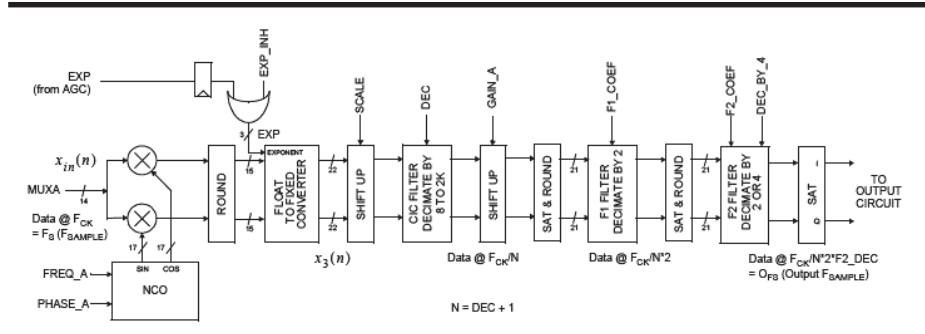
Figure 9 – The block diagram of one of the two channels of the TI CLC5903 Digital Down Converter IC. Each of the channels is identical.

need to create your own board from the Gerber files from TI (which I can supply) in order to experiment with the part. The reason I like the part so much is that it does both down and up conversion in one part and costs just $39. It is a BGA package, so that is a small issue. The GC4016 is a similar part that is only a digital down converter and is $64. The GC6016 is a fully supported part with significantly more capabilities that supersedes the GC5016, but it is $104. The GC5016 block diagram in transceiver mode is shown in Figure 7. It is capable of 320 Msps when set up using double rate mode. Probably the biggest task in experimenting with these parts is the substantial software effort to get all of the configuration registers set.

4DSP uses another set of TI digital radio ICs on the FMC30RF board for use with Xilinx FPGA evaluation modules. The AFE7225 is the heart of the system. It contains the ADCs and DACs for data conversion as well as numerically controlled oscillators for digital frequency translation. It provides 2X decimation in the receive path and 4X interpo-

lation on the transmit path. Figure 8 shows its block diagram. The sample rate on receive is 250 Msps and 125 Msps on transmit. It is a $54 part in a QFN package (similar to a chip scale package), so soldering to a bard is not terribly difficult. The AFE7225 requires an FPGA to further reduce the receive data rate for use by a general purpose DSP.

The CLC5903 is another interesting part. It is a digital mixer, down converter, and filter IC that is available in a PQFP package for $33. Its input sample rate is 78 Msps for each of its two receive channels. Figure 9 shows the block diagram of one of the channels of the part. This is a very capable part that can decimate by significant amounts in the CIC section and then capable FIR decimation by up to an additional 8X. Depending on your needs it can decimate down to 192 kHz sample rate or lower to reduce the speed requirements of the final DSP.

Notes

[1] Photos courtesy of TAPR.
[2] Photo courtesy of ZedBoard.org

2012 *QEX* Index

Features

About the Cover

Empirical Outlook (Wolfgang)

Letters to the Editor

Out of the Box (Mack)

SDR: Simplified (Mack)

Tech Notes

Upcoming Conferences

QEX