QEX

$5

March/April 2013

www.arrl.org

A Forum for Communications Experimenters

Issue No. 277

K6IQL presents A Precision DDS for the Frequency Measuring Tests. A bit more involved than many of the DDS projects we have seen, this one provides resolution accuracy to the nearest 0.1 millihertz.

ARRL The national association for AMATEUR RADIO™
225 Main Street
Newington, CT USA 06111-1494

HF/50MHz ALL MODE TRANSCEIVER TS-590

Nothing But Performance

The TS-590S

Kenwood has essentially redefined HF performance with the TS-590S compact HF transceiver. The TS-590S RX section sports IMD (intermodulation distortion) characteristics that are on par with those "top of the line" transceivers, not to mention having the best dynamic range in its class when handling unwanted, adjacent off-frequency signals.*

- **HF-50MHz 100W**
- **Digital IF Filters**
- **Built-in Antenna Tuner**
- **Advanced DSP from the IF stage forward**
- **500Hz and 2.7KHz roofing filters included**
- **Heavy duty TX section**

- **2 Color LCD**

KENWOOD

Customer Support: (310) 639-4200
Fax: (310) 537-8235

Scan with your phone to download TS-590S brochure.

www.kenwoodusa.com

ISO9001 Registered
Professional Systems Business Group
JVCKENWOOD Corporation

ADS#43712

*** For 1.8/3.5/7/14/21 MHz Amateur bands, when receiving in CW/FSK/SSB modes, down conversion is automatically selected if the final passband is 2.7KHz or less.**

QEX (ISSN: 0886-8093) is published bimonthly in January, March, May, July, September, and November by the American Radio Relay League, 225 Main Street, Newington, CT 06111-1494. Periodicals postage paid at Hartford, CT and at additional mailing offices.

POSTMASTER: Send address changes to: QEX, 225 Main St, Newington, CT 06111-1494 Issue No 277

Offices

225 Main St, Newington, CT 06111-1494 USA
Telephone: 860-594-0200
Fax: 860-594-0259 (24 hour direct line)
e-mail: **qex@arrl.org**

Subscription rate for 6 issues:

In the US: ARRL Member $24,
nonmember $36;

US by First Class Mail:
ARRL member $37, nonmember $49;

International and Canada by Airmail: ARRL member $31, nonmember $43;

Members are asked to include their membership control number or a label from their QST when applying.

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 Roos, K6IQL, presents A Precision DDS for the Frequency Measuring Tests. Unlike many of the DDS projects we have seen, this one provides resolution accuracy to the nearest 0.1 millihertz, with an RF output between 100 kHz and 40 MHz. This meets John's requirements for an FMT DDS.

In This Issue

Features

Index of Advertisers

The American Radio Relay League

The American Radio Relay League, Inc, is a noncommercial association of radio amateurs, organized for the promotion of interest in Amateur Radio communication and experimentation, for the establishment of networks to provide communications in the event of disasters or other emergencies, for the advancement of the radio art and of the public welfare, for the representation of the radio amateur in legislative matters, and for the maintenance of fraternalism and a high standard of conduct.

ARRL is an incorporated association without capital stock chartered under the laws of the state of Connecticut, and is an exempt organization under Section 501(c)(3) of the Internal Revenue Code of 1986. Its affairs are governed by a Board of Directors, whose voting members are elected every three years by the general membership. The officers are elected or appointed by the Directors. The League is noncommercial, and no one who could gain financially from the shaping of its affairs is eligible for membership on its Board.

"Of, by, and for the radio amateur," ARRL numbers within its ranks the vast majority of active amateurs in the nation and has a proud history of achievement as the standard-bearer in amateur affairs.

A *bona fide* interest in Amateur Radio is the only essential qualification of membership; an Amateur Radio license is not a prerequisite, although full voting membership is granted only to licensed amateurs in the US.

Membership inquiries and general correspondence should be addressed to the administrative headquarters:

ARRL
225 Main Street
Newington, CT 06111 USA
Telephone: 860-594-0200
FAX: 860-594-0259 (24-hour direct line)

Officers

President: KAY C. CRAIGIE, N3KN
570 Brush Mountain Rd, Blacksburg, VA 24060

Chief Executive Officer: DAVID SUMNER, K1ZZ

The purpose of *QEX* is to:

1) provide a medium for the exchange of ideas and information among Amateur Radio experimenters,

2) document advanced technical work in the Amateur Radio field, and

3) support efforts to advance the state of the Amateur Radio art.

All correspondence concerning *QEX* should be addressed to the American Radio Relay League, 225 Main Street, Newington, CT 06111 USA. Envelopes containing manuscripts and letters for publication in *QEX* should be marked Editor, *QEX*.

Both theoretical and practical technical articles are welcomed. Manuscripts should be submitted in word-processor format, if possible. We can redraw any figures as long as their content is clear. Photos should be glossy, color or black-and-white prints of at least the size they are to appear in *QEX* or high-resolution digital images (300 dots per inch or higher at the printed size). Further information for authors can be found on the Web at **www.arrl.org/qex/** or by e-mail to **qex@arrl.org**.

Any opinions expressed in *QEX* are those of the authors, not necessarily those of the Editor or the League. While we strive to ensure all material is technically correct, authors are expected to defend their own assertions. Products mentioned are included for your information only; no endorsement is implied. Readers are cautioned to verify the availability of products before sending money to vendors.

Larry Wolfgang, WR1B

Empirical Outlook

Experimenting for Fun and Necessity

QEX: A Forum for Communications Experimenters. There, it says it right in the title of our magazine: *QEX* is for experimenters. What kind of experimenting do you do? Why do you experiment? What do you learn from or gain from your experiments?

Most of us will be quick to say that experimenting is fun. We build projects, test new circuits, even try new communications modes, because it we enjoy it. Along the way we learn something new, whether it is a personal accomplishment such as a deeper understanding of some electronics principle, or perhaps we help advance the state of the electronics/communications art.

When you share the results of your experimentation you help your fellow Amateurs gain insight and understanding. You may spur some to continue your research and provide even greater benefits to the Amateur Radio community or mankind.

That may seem like quite a lot, just from one *QEX* article, but who knows where that information may lead someone else. By sharing the joy of building projects and experimenting with electronics you might be inspiring the next generation of scientists and engineers.

We share some of those "stories" in every issue of *QEX*, and this one is no exception.

Wayne Openlander, W9NZB, describes his efforts to build a small magnetic loop antenna for 20 meters. Wayne's article is not a construction project, but a story about some of the experiments he conducted along his journey. Using an infrared thermometer to search for hot spots along his loop is an interesting idea. Perhaps others will find this technique useful in other ways.

Sam Green, W0PCE, describes some of the fun he had experimenting with voltage to frequency converter ICs. Sam describes several ways he used these converters to create an audio signal that he could hear as he adjusted various circuits or equipment. There are several ideas here that readers may find helpful in their own experimentation. Sam even describes one circuit that lead to a patent application. Who knows where your experiments might lead?

John Roos, K6IQL, describes a direct digital synthesizer circuit that he developed with one of his particular interests in mind — ARRL frequency measuring tests (FMT). John points out that you would need a reference oscillator accuracy of 7 parts in one billion to accurately measure a 40 meter signal in one of these tests. A DDS that would prove perfectly adequate as a transceiver LO or VFO will not fulfill the needs for an FMT station. John describes his efforts to build a suitable DDS While his circuitry is a bit more involved than the typical DDS that we have looked at, he did achieve a synthesizer that provides mHz resolution, suitable for his FMT application. Might this also lead to new performance in more general applications? Where might it take us?

Ray Mack, W5IFS, continues to lead us through some DSP topics in our study of software defined radio concepts. In this issue we learn more about noise analysis and adaptive filtering DSP techniques. I know I have been learning a lot from Ray's tutorials, and gaining a better understanding of many topics I've read about but never quite "gotten" over the years. It is fun to learn about new technology.

In the earliest days of radio, amateur experimenters probably felt this same sense of pride in extending the body of knowledge for future generations. They were discovering new ways to make use of this new thing called radio. Those must have been exciting times to be an Amateur Radio operator!

One of the things that strikes me about those early experimenters is that they couldn't just go to the local parts emporium (or on-line store) and order the capacitors, tubes, transformers and other parts they needed. Often they had to build these elements from materials found in a literal junk heap or salvaged from some equipment that was no longer useful. I am in awe of the You Tube videos I've seen of some amateur experimenter building a power tube in his home workshop or crafting a fine Morse paddle from a block of brass or other metal. I could not do it, but some of these expert machinists, with fairly sophisticated workshops, produce some fine equipment. But in those early days of radio, the amateur experimenters were building tubes, capacitors, transformers and more with rather crude workshops and no step by step instructions! What they accomplished is truly amazing.

The construction of components went well beyond just the joy of crafting your own parts. Then it was absolute necessity, because suitable parts were not commercially available. In this issue we have an article by Luiz Amaral, PY1LL/AC2BR. Luiz has been experimenting on the 500 kHz band, and had to build his own transmitter. When he realized he needed an output filter for his transmitter he built the capacitors and inductors needed for the filter. (See his article in the Jan/Feb 2012 issue of *QEX* for a description of his home made capacitors.)

With the prospect of a new Amateur Radio allocation in the 475 kHz region (See "It Seems to Us," Feb 2013 *QST* by Dave Sumner, K1ZZ), more of us may be taking an interest in this region of the spectrum. While an official allocation in that frequency range probably means commercial equipment will become available, we may find ourselves experimenting with home built components of necessity again, at least for a while. Have fun!

Wayne Openlander, W9NZB

2909 N Wolcott Ct, Chicago, IL 60657, **engineering@microwavemagician.com**

A Two Turn Magnetic Loop Antenna for 30 through 10 Meters

The journey to success with this antenna included several interesting detours.

I've been working on the construction of an efficient small magnetic loop antenna for Amateur Radio on 20 meters. This is the story of what I did, not a construction project.

My home is in a restricted area and I cannot have a highly visible antenna, nor run much power without inviting problems from my neighbors and the homeowners association.

I started with a 42 × 42 inch square loop using ¾ inch ID copper pipe with an L match and a trombone tuning capacitor following a design published by John Portune, W6NBC.[1,2]

I then spent a number of hours operating the loop in my basement lab at low transmitter power levels. Using an established measurement criterion – comparing the calculated Q from the radiation resistance[3a and 3b] and loop inductance[4] to the measured Q — I calculated an overall efficiency of 5%.

I looked for power-sucking hot spots with an infrared thermometer. Operating at 25 W input, both the tuning capacitor and the choke balun became warm. I had made the capacitor insulator from Delrin and this was getting hot. I replaced all the Delrin insulators with Teflon insulators and then replaced the RG58 coax in the choke balun with RG142/u. This cooled down the hot spots but with no appreciable improvement in efficiency.

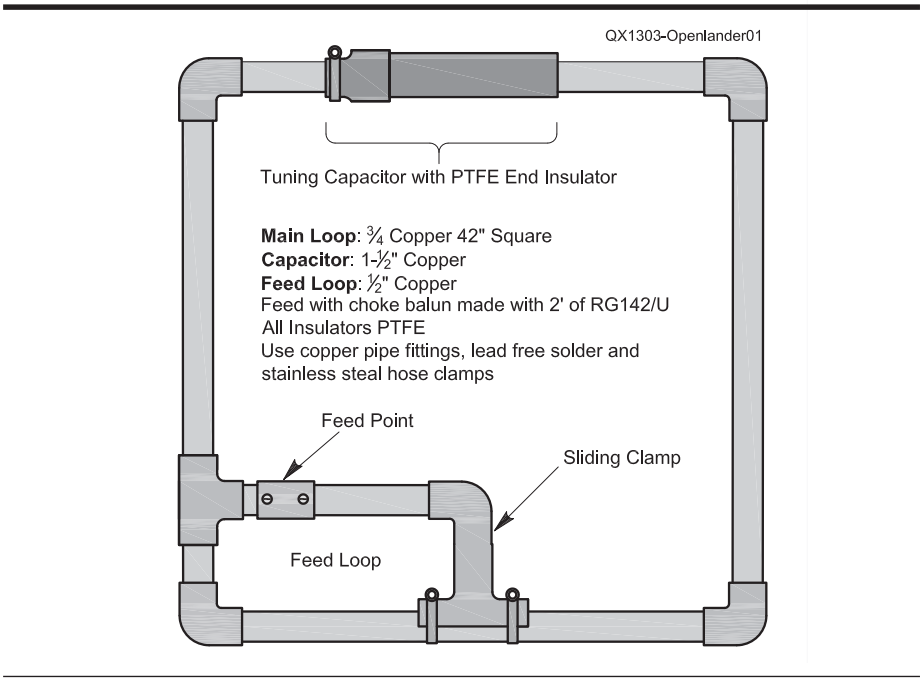

Figure 1 – John Portune's 20 meter loop.

Tests and Tweaks

I put the loop outside on my deck. Using the concept of reciprocity[5] and averaging several days of WSPR operation (a digital mode used to study HF propagation), my observations appeared to confirm the 5% efficiency estimate.

Given the remarkably low efficiency of the loop and, aside from WSPR, my inability to make any contacts with it, I rebuilt the loop to 42 × 84 inches. According to the formula for small loop radiation resistance[6] this doubling of area increases the radiation resistance by a factor of four.

Testing the antenna with an antenna analyzer and recalculating the efficiency by the same method as before gave me an overall efficiency of 20%.

I again used the WSPR network to confirm the efficiency, but with less rigor than before. I was also able to make a number of contacts on 20 meters using the JT65 HF mode.

It's a judgment call. I felt that the 42 × 84-inch footprint of this loop was as big as my neighbors would ignore. My next step was to add a second turn to the existing loop with the hope of a practical 80% efficiency. The formula for loop radiation resistance predicts that this doubling in area would further quadruple the radiation resistance.

I used Teflon spacers to separate the turns 8 inches center to center and kept the tuning capacitor and L match from the original W6NBC design.

Nothing worked as expected. No matter how much I worked with the tuning capacitor, the matching network or turn to turn spacing I could not get the antenna to resonate. My observations suggested that the loop antenna impedance had become capacitive at 14.1 MHz.

Thinking that using a closer spacing between turns would return the loop to an

[1]Notes appear on page 6.

inductive reactance, I closed the spacing to 4¾ inches between turns. A chart in one of my references[7] cautions that too close a spacing increases the antennas skin resistance. A good rule of thumb is a minimum spacing of a 3X the wire diameter. My 4¾ inch spacing comfortably exceeds this limit.

Unable to resonate the loop with a series capacitor, I then tried several schemes to tune the antenna with an antenna tuner. I installed a tuner at the antenna to reduce the VSWR on my feed line to 1:1. I experimented with various balun arrangements. Nothing gave me satisfactory results.

Then, in frustration, I borrowed a Vector Network Analyzer from another lab. After a quick sweep the answer was obvious – the loop had a natural resonance at 21 MHz and below that frequency was capacitive at 14 MHz, not the inductance I had expected. After a little trial and error I resonated the antenna with a 5 µH inductor.

At 14.1 MHz resonance the antenna impedance measures 16 Ω. On a Q meter the 5 µH inductor measures a resistance of about 1 Ω. A 3:1 VSWR presents no difficulty for a transmission line and antenna tuner.

The operation of the antenna on 20 meters was now entirely satisfactory as evidenced by several sequential contacts with European stations using JT65, and as with WSPR, the S/N ratio was consistent with 50% efficiency.

The antenna then tested well on 15 meters with the tuning inductor replaced by a short circuit, and also on 30 and 17 meters relying on an external tuner near the antenna.

Surprises

This has been a long journey with a number of surprises.

Figure 2 – The 42 × 42 inch loop under test in the author's basement.

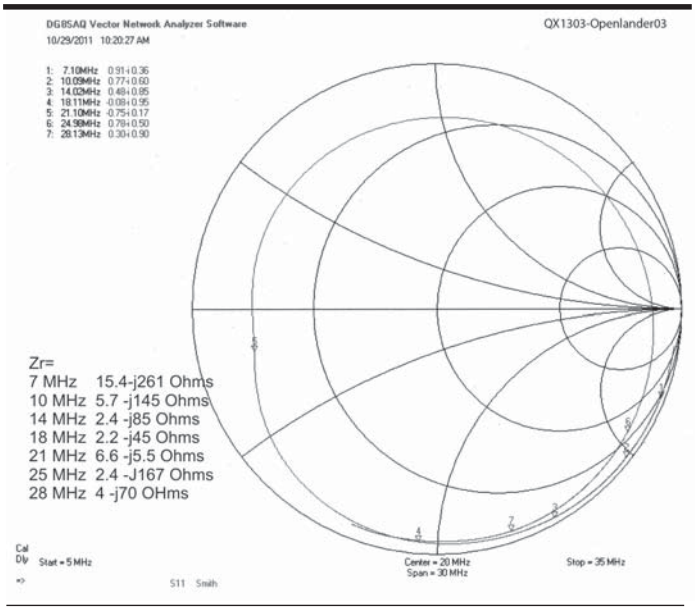

Figure 3 – A frequency sweep of the two turn loop from 5 to 35 MHz.

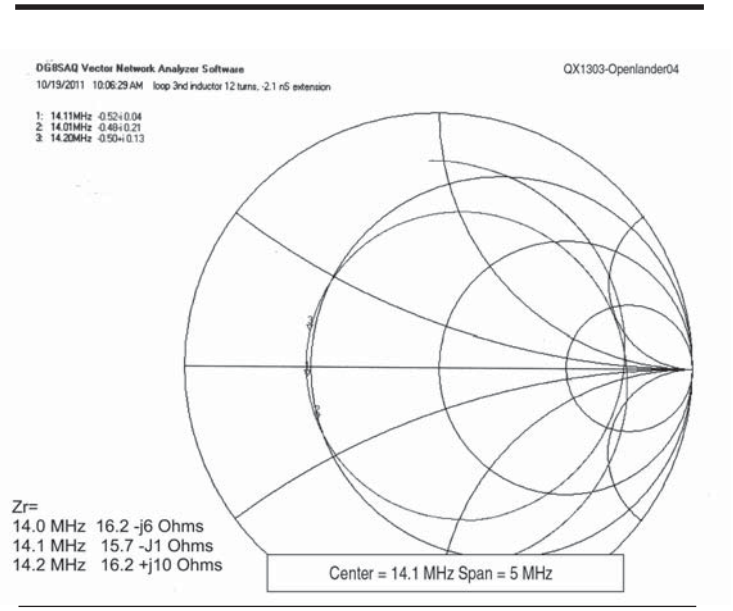

Figure 4 – Antenna Impedance at 14 MHz after added Inductance.

The first surprise was that basing the efficiency of small loop antennas on the ratio of skin resistance to radiation resistance proved grossly in error. I invested endless hours in an effort to find sources of loss by looking for hot spots. After eliminating the insulator and the balun I was stumped.

Antenna engineers recognize that, because there are practical limits on the Q of the antenna's impedance matching network, very low radiation resistances will be accompanied by high matching network losses. The degree of loss is a surprise.

The second surprise is that adding a second turn to the antenna changed the antenna's impedance behavior from an inductive loop to a dipole below resonance. It was no longer a magnetic loop, but something like a 15 meter dipole. The small loop formulas no longer applied.

My thinking here is that the inter-turn capacitance exceeds the loop inductance. My evidence for this is that when I reduced the inter-turn spacing from 6 inches to 4¾ inches the self resonant frequency dropped from about 25 MHz, as measured with an antenna analyzer, to 21 MHz as measured with the VNA.

The third surprise is the increase in radiation impedance when the antenna was resonated. It is possible that there are errors in my VNA measurements. I tend to distrust impedances near the edge of the Smith chart. The angles are usually accurate but the magnitudes can be in error. This might explain some of the difference in radiation resistance between the matched and unmatched loops.

A fourth surprise is that a careful search of the antenna literature found very little information about rectangular, asymmetrically fed, multi-turn transmitting loops.

Improvements?

There is considerable room for future improvement and experimentation. If one is willing to accept the moderate VSWRs, around 3 or 4 to 1, band switching is very simple – little more than a matter of changing matching coils. A simple relay board would do this.

If one is limited to low power operation it is possible to tolerate even higher VSWR on the feed line. My initial experiments suggest that with the 20-meter tuning inductor replaced by a short circuit, and an antenna tuner in the shack, efficient operation is practical on all the bands from 30 through 10 meters.

Finally there are many still unanswered questions. The bottom of the antenna is only about four feet off the ground. What happens if the antenna elevation is raised?

What about adding more turns. If two turns is efficient on 20 meters, would 4 turns be similarly efficient on 40 meters?

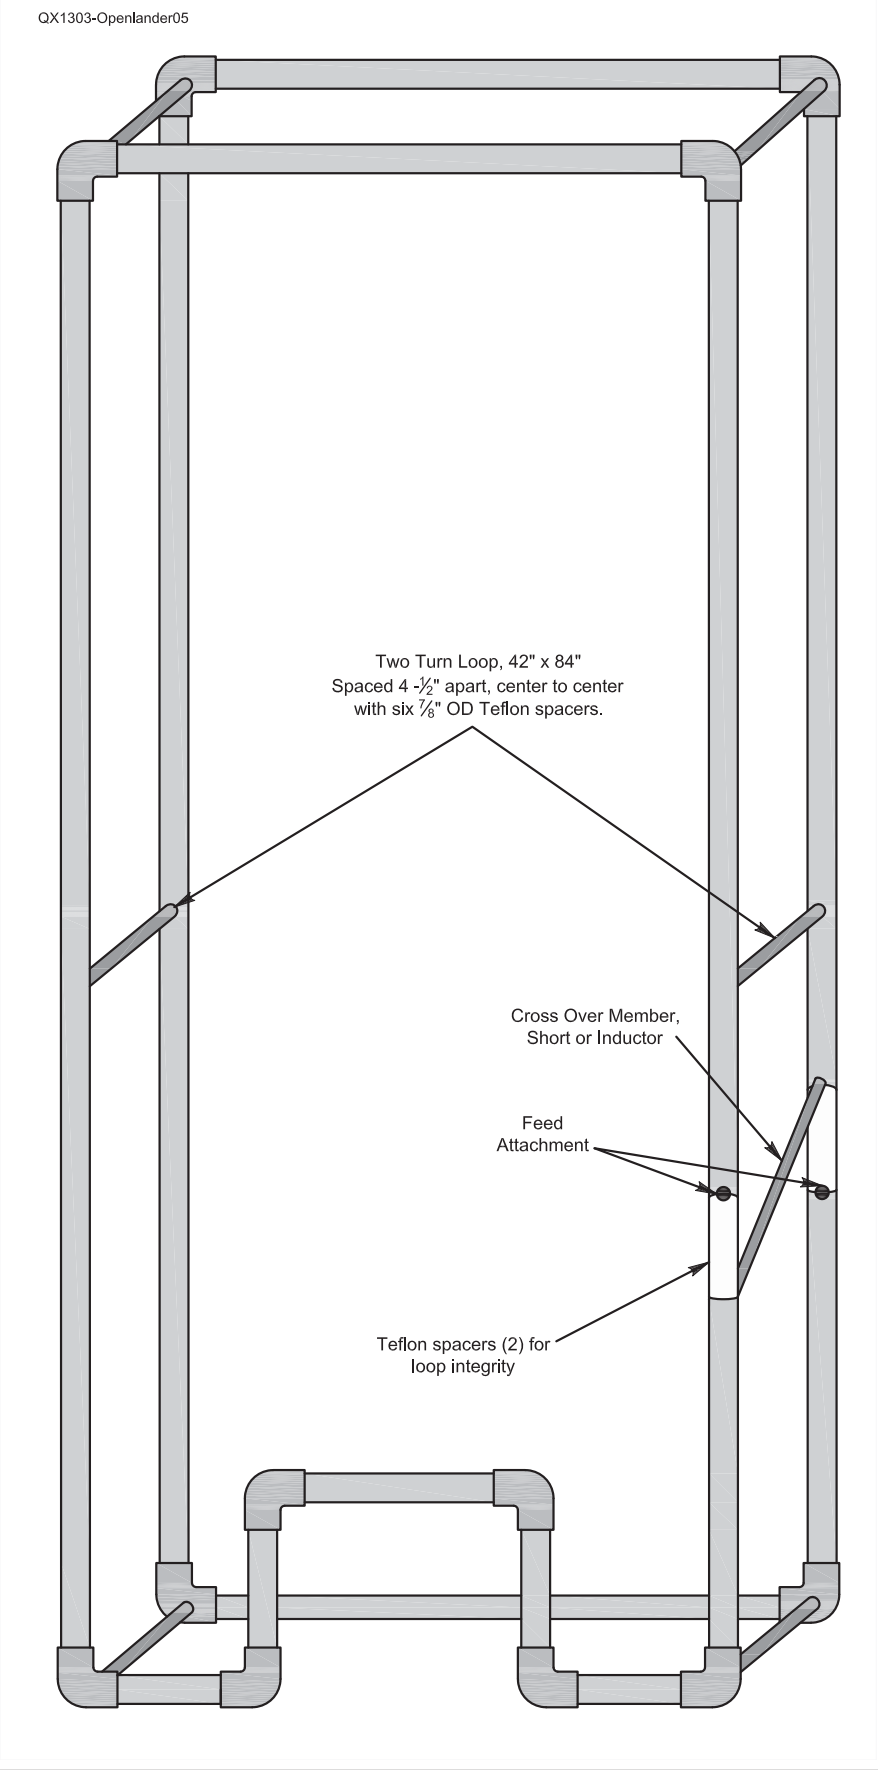

Figure 5 – A diagram of the 42 × 84-inch version of the loop.

Figure 6 – The two turn loop mounted on the author's deck.

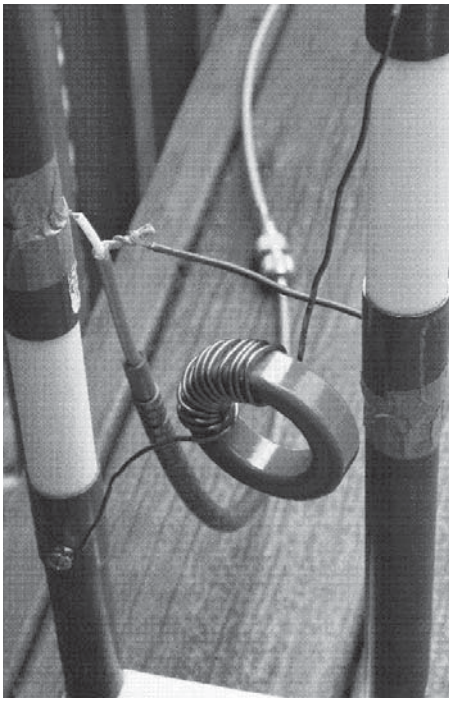

Figure 7 – A balun feed and 5 µH inductor tunes the loop for 20 meters. Replacing the balun with a metal strap resonates the loop at 21 MHz.

Notes

[1]Portune, John, W6NBC, "Compact 40 Meter HF Loop for Your Recreational Vehicle," *QST*, March 2007, pp 41-43.

[2]Portune, John, W6NBC, "Compact HF Loops for Your RV," **www.w6nbc.com/hfrvloop.html**, January 3, 2009.

[3a]Newman, Edward et al, "Two Methods for the Measurements of Antenna Efficiency," *IEEE Transactions on Antennas and Propagation*, July 1975, pp 457-461.

[3b]Belrose, John S., "Performance Analysis of Small Tuned Transmitting Loop Antennas Evaluated by Experiment and Simulation," *IEEE Antennas and Propagation Magazine*, June 2007, pp 128- 132,

[4]Air Core Inductor Inductance Calculator, **www.daycounter.com/Calculators/AirCoreInductorCalculator.phtml**

[5]Reciprocity: The idea here is that the path conditions between two stations are identical in both directions and, as a consequence, signal strengths measured at both sites will be in linear proportion to their transmitter power. In actuality, because of the volatility of skip conditions this path symmetry is rarely a true case at any individual moment in time. However, if measurements are averaged over a large enough number of samples, the rule holds. That is, if 50 W to the antenna yields an average S/N ratio report consistent with 2.5 W radiated power, the antenna efficiency is 5%.

[6]Reference formula for radiation resistance from Krauss, John D., *Antennas*, McGraw Hill Book Company, New York, 1950, pp 167. Radiation resistance is given by the formula: R = 31171 × Number of Turns^2 × (Area^2/Wavelength^4)

At 14.1 MHz

42 × 42 inches = 0.20 Ω

42 × 84 inches = 0.80 Ω

2 turns × 42 × 84 inches = 3.2 Ω

[7]Balanis, Constantine A, *Antenna Theory, Analysis and Design*, 2nd ed., John Wiley and Sons, New York, 1997, pp 211.

ARRL member Wayne Openlander, Amateur Extra class W9NZB, holds five antenna and two instrument patents. He recently retired as Vice President of Resonance Instrument, Inc, a custom manufacturer of microwave instrumentation. Previously he was Vice President of Engineering at Antenex, Inc.

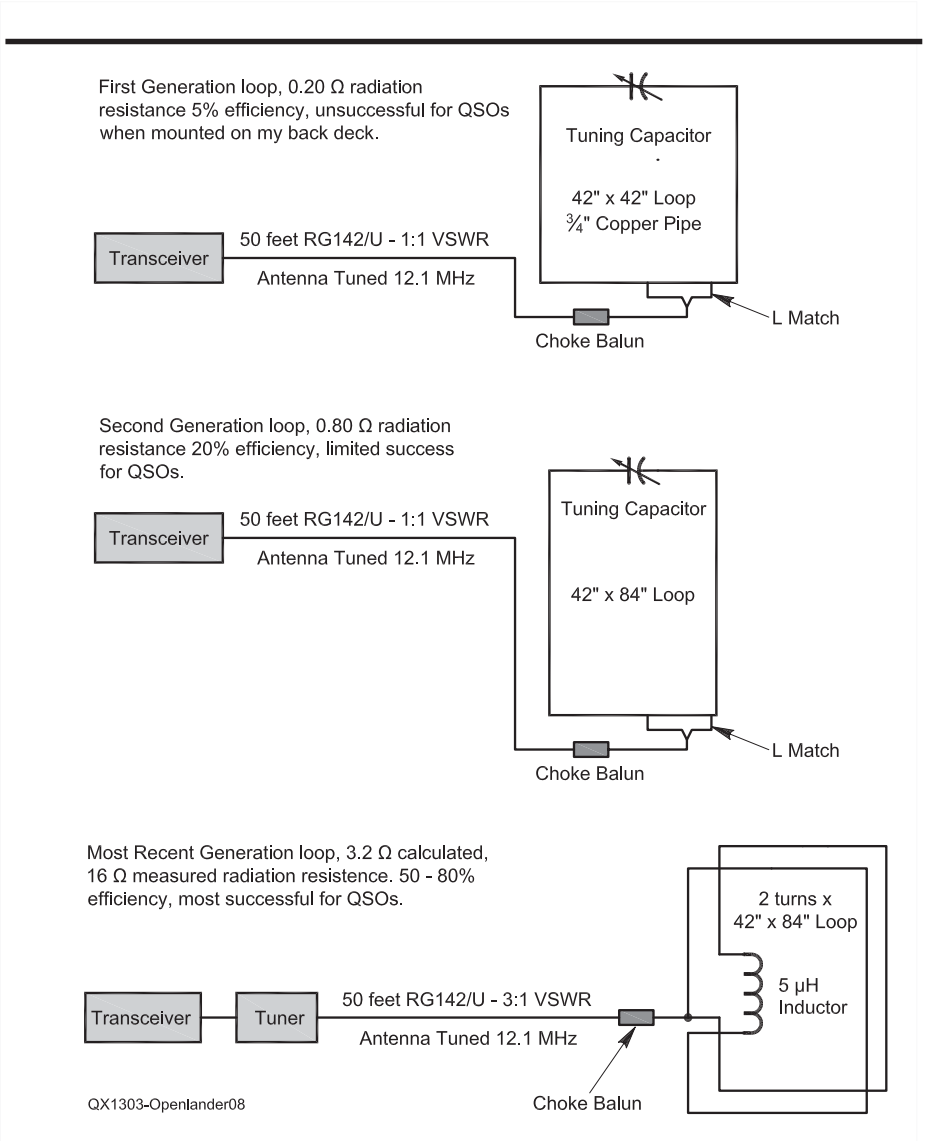

Figure 8 – System block diagrams at various stages of development.

Dr Sam Green, W0PCE

10951 Pem Rd, St Louis, MO 63146, **w0pce@arrl.net**

Fun with Voltage-to-Frequency Converters

These converters have many applications. Consider using one in a project.

A voltage-to-frequency (VFC) converter is a type of voltage controlled oscillator. It is often only useful at lower frequencies and may have a narrow pulse output instead of a nice sine wave or square wave. I used them in many applications over the years. I will describe several to give you an idea of how you may use them to benefit your own projects.

Many of these examples render signal measurements into audible signals that the user hears. This allows the user to perform other tasks that may require visual attention such as making precision adjustments while listening to the relative results of those adjustments. This method offers resolution better than that available from the close examination of an analog meter because the human ear is very sensitive to small changes in frequency or pitch.

Remote Photocurrent Measurement

The application was to measure the photocurrent of a photodiode at a remote receiver in a free-space laser communication link circa 1978. I steered the laser telescope at the transmitter telescope and used the photocurrent measurement to optimize pointing. A trans-impedance amplifier turned the photocurrent from the receiver into a voltage to control the frequency of VFC that in turn drove a loudspeaker transducer from an old acoustic modem in a rubber cup (remember those?) that attached to the microphone of a telephone handset. I controlled the frequency range to vary between 300 Hz and 3000 Hz to accommodate the telephone system. A telephone link carried the audio from the receiver back to me at the laser transmitter so I could optimize link performance by maximizing the audio frequency and thereby the received optical signal level. It was a simple and effective solution to my problem.

The ICL8038[1] is the first chip I used as a VFC. Intersil called it a Precision Waveform Generator and Voltage Controlled Oscillator. The ICL8038 worked up to 300 kHz and had the decided advantage of simultaneous sinusoidal, square and triangle outputs. That made it more of a function generator chip than a VFC, but it's the first one I used. It is obsolete now and frequency control by a voltage was not straightforward, so I recommend against using it and provide no schematic diagram.

Better choices today are the National LM331[2] and the Analog Devices ADVFC32.[3]

[1]Notes appear on page 10.

These are both much simpler to interface and very linear. I prefer the LM331 because it is in an 8 pin DIP package and works with a single supply voltage from 4-40 Vdc while the ADVFC is in a 14 pin DIP and calls for symmetrical ±15 Vdc supplies. Analog Devices also makes the AD7741, which is in an 8 pin DIP and takes a single +5 Vdc supply, but I have no experience with it.

All of these are super linear, at least to the degree that you can read an input voltage by displaying the output frequency on a frequency counter if you scale the input voltage range and output frequency range accordingly. Figure 1 is an example from the LM331 data sheet that shows the setup and suitable component values.

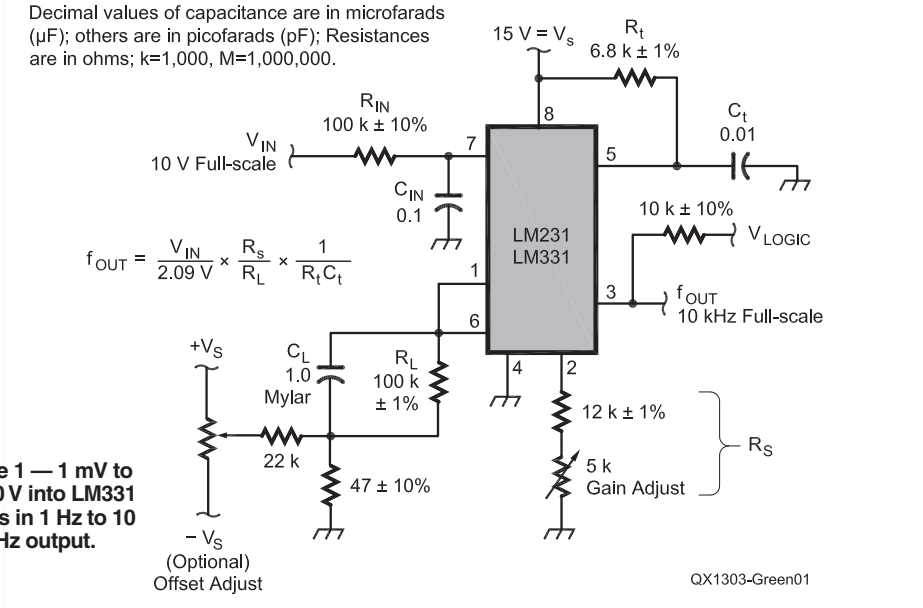

Figure 1 — 1 mV to 10.000 V into LM331 results in 1 Hz to 10 kHz output.

Audio Link

The ADVFC32 and LM331 data sheets show they are also able to serve as frequency-to-voltage converters. Figure 2 shows a fiber optic link I built in which an audio voltage modulates the frequency of a first ADVFC32 that in turn modulates light pulses into an optical fiber. A second ADVFC32 accepts pulses from the optical receiver to regenerate the original audio waveform. Think of this as an exercise in optical isolation. The input limiter helps maintain FM transmitter deviation in a desired range.

The BERT Trick

Both the LM331 and the ADVFC32 output narrow pulses. There is little audio energy in low duty cycle pulses, so the waveform is unsuitable for listening. There isn't much to hear. I had the same problem with bit-error-rate testers. These instruments monitor the performance of serial digital communication links and output a narrow one bit wide pulse each time an error occurs. At 1 Gbps that pulse lasts only 1 ns. Most people lock their eyes on the error counter display to determine error rate or when errors occur. It's pretty hard to make useful adjustments while you monitor the display. Instead, I chose to listen to the errors. Simply feeding the error pulses to a toggle flip-flop yields a high duty cycle waveform with lots of audible energy[4] as in Figure 3. The same trick works with these VFCs. A toggle flip-flop divides the frequency by two and converts the pulse train to a 50% duty cycle square wave.

Alternately, a Don Lancaster trick divides the frequency by 10 and synthesizes a decent sounding sine wave.[5] All audible solutions below use one of these frequency divider tricks to convert a train of narrow pulses into audible waveforms.

Another Photocurrent Measurement

The next project was another photocurrent measurement. This one utilized the sensitivity of the human ear to sense very small changes in pitch. I sent a fixed amount of unmodulated light into an optical fiber and used this instrument to indicate how much light came out the other end. The relative measurement was so sensitive that bending the fiber caused a noticeable decrease in the audible frequency due to the light lost in the bend. I intended this as a way to monitor fiber bending and strain when fiber is pulled during installation, where it often breaks. This method indicated a dramatic decrease in audible frequency before the fiber would break and so give the installer a chance to back off and try something else.

I patented this method.[6] My company sold it to another party who subsequently paid the maintenance fee that will keep this patent in force until 2015.

Radio Frequency Signal Strength Measurement

I designed and built a field strength meter that indicates higher power as a higher audio frequency pitch. The circuit is simply an Analog Devices AD8307[7] log-detector with 500 MHz response driving an LM331 VFC and toggle flip-flop. The output voltage of the log detector represents the logarithm of the RF signal input strength, compressing a wide dynamic range input into a couple volts of output swing. Applying that voltage to control the frequency of an oscillator in the audio range provides a simple and effective way to communicate the signal value to your ear. I constrained the audio frequency range to between 300 – 3000 Hz to match telephone

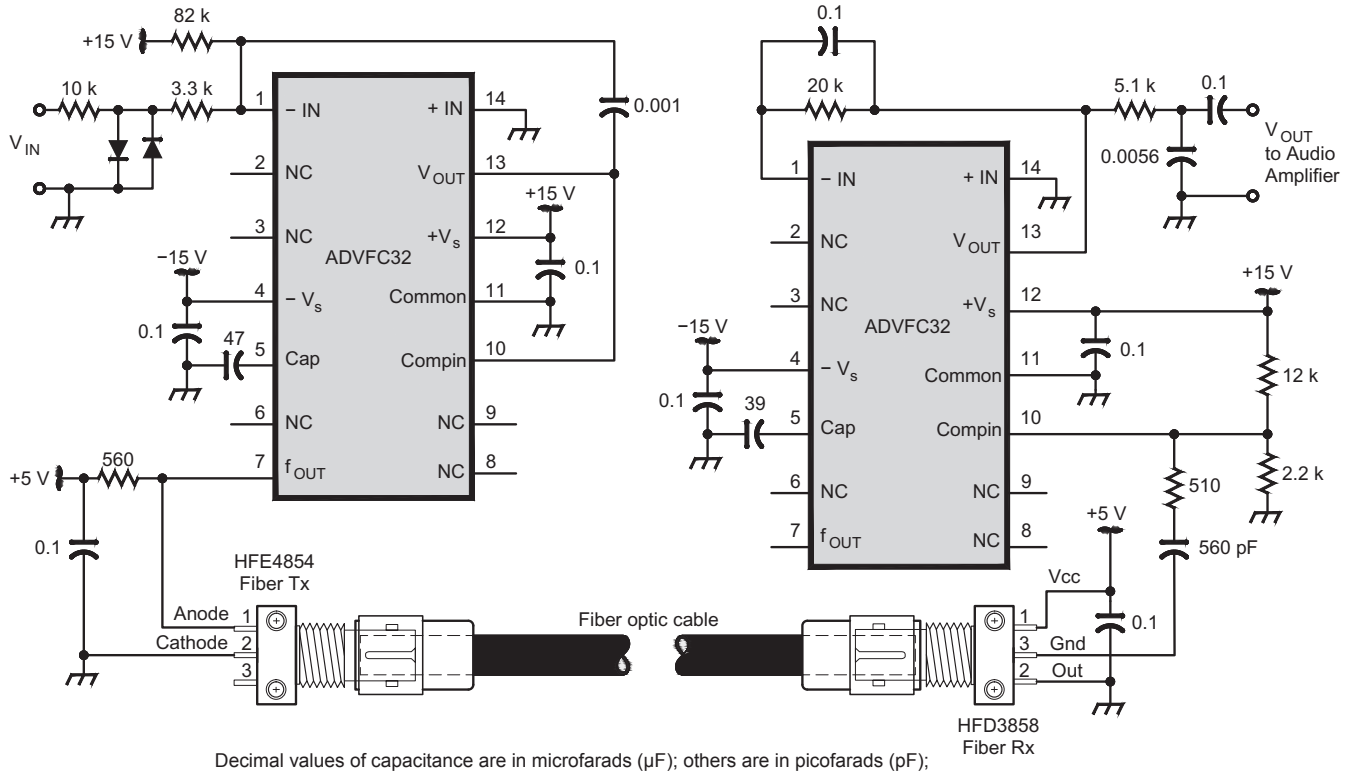

Decimal values of capacitance are in microfarads (μF); others are in picofarads (pF);
Resistances are in ohms; k=1,000, M=1,000,000.

QX1303-Green02

Figure 2 — The fiber optic link uses ADVFC32s to transmit and receive audio.

line and FM transceiver capabilities.

The human ear easily differentiates very small frequency changes, allowing very high resolution and high sensitivity to small changes. Sensitivity to amplitude changes is very much poorer, as the human ear is able to resolve variations of about 1 decibel which corresponds to about –20% or +25%.

I meant this to be an instrument for the amateur radio community to make remote measurements of relative field strength. I thought it would prove useful to measure radiated signal level at some distance from the transmitting source, perhaps for antenna comparisons or to see if an antenna tuner adjustment is real and not a phantom dip in VSWR. In the shack, you would read the meter of an ordinary instrument, but remote operation requires a way to return the measurement data to the user while making adjustments that affect the remote reading. Conversion of the field strength reading to a varying audio frequency signal allows return of the measurement data via a phone line or an audio modulated FM radio link. Rotate an antenna and listen to measure the effect in the far field.

I walked around my workplace with the prototype and listened to different sources of RF. When I heard the data exchange as the RF reader interrogated my employee badge, I got the idea that this could become a cell phone and wireless detector. That might prove useful to sense such devices operating where they shouldn't. I built up another prototype based on an AD8313 log detector with 2500 MHz response to sense cell phones and wireless. With it I could tell which laptop had its wireless enabled when I walked up and down the aisle on a commercial airline flight with headphones connecting to the unit in my pocket.

Because airlines want RF emitters turned off, and because I worked for an airplane company that offers incentives for intellectual property, I patented this method as well.[8] Again, you can look at the details in the patent online.

This is what I built into the two Altoids tins in the article "Hints & Kinks: Altoids Times Two" that appeared in *QST* and the *ARRL Letter*.[9] Figure 4 is an excerpt from the patent. The decision circuit and light emitting diode indicator never existed and were an "attempt at completeness" by a patent attorney whom I will be glad never to work with again.

Company attorneys discouraged me from publishing actual schematic diagrams, but this is fairly simple to understand from the discussion. My hope is someone will notice this and license it for manufacture and sale. I wouldn't benefit, but I am grateful for the benefits already received during my years of employment.

For extending the AD8307 logarithmic detector to work at low frequencies, refer to my *QEX* article "Fully Automated DDS Sweep Generator Measurement System".[10]

Digital Multimeter Accessory

This last project is a fairly recent attempt to overcome the worst problem of digital multimeters (DMM), the lack of a decent analog indication. Digital meters offer precision, resolution and accuracy that old moving needle analog meters never could. Analog meters did excel at indicating relative increases and decreases in magnitude.

Some DMMs add a crude bar graph display to indicate relative magnitude, but these are worthless except for large variations. No one would use such a display for any fine adjustment.

The use of a VFC in this application offers resolution that exceeds the resolution available from the moving needle of an analog meter. This is an obvious improvement, yet no vendor offers such a capability. Adding this to a DMM gives precise analog information that enables a user to perform adjustments that maximize or minimize any measurable parameter without even having to look at the DMM!

I set out to build an accessory for a DMM to demonstrate the concept. The nice thing about a DMM is that it performs all the scaling by converting all possible input signals, whether current, voltage, resistance, or any of a number of other parameters, to the range of the internal digital meter. This was ±200 mV for my demonstration unit. I added a gain of 10 before the VFC input to accommodate a 200 mV input range.

At this point, there is no advancement to the state of the art, and anyone can reproduce it.

Then I added three improvements that I thought would be desirable in a DMM.

1. I added an absolute value circuit, so that the VFC receives only positive going signals from the input.

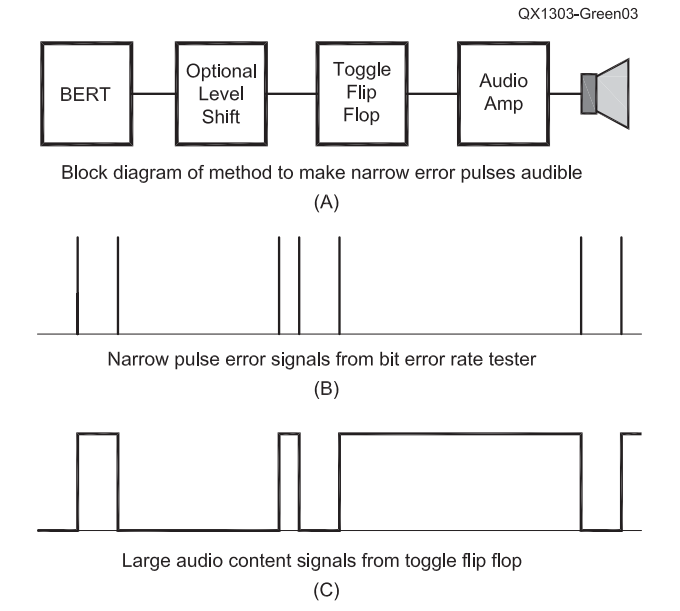

Figure 3 — A method to make narrow pulses audible.

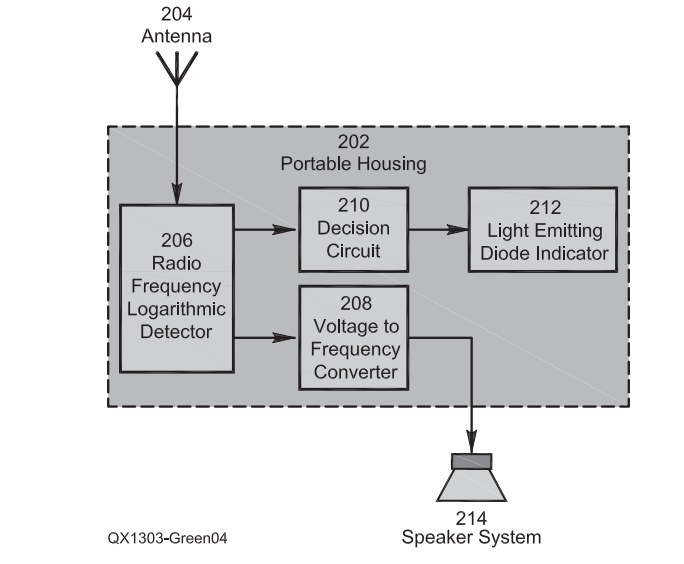

Figure 4 — A method and apparatus for detecting radio frequency signals.

2. I sensed when the input is negative and used this signal to change the way the audio sounds when the input goes negative to be able to tell the difference when listening. I tried switching between different waveforms such as sine waves and square waves and then settled on tremolo. When the input signal goes negative, a low frequency oscillator interrupts the voltage indicating waveform. This was very effective and much simpler to implement.

3. At the suggestion of my co-inventor, I added a squelch circuit to turn off the output when it wasn't changing. This turned out to be difficult to implement, but it makes for much less annoying operation.

We applied for patent protection on these key features. Figure 5 from the pending patent application[11] provides a general description.

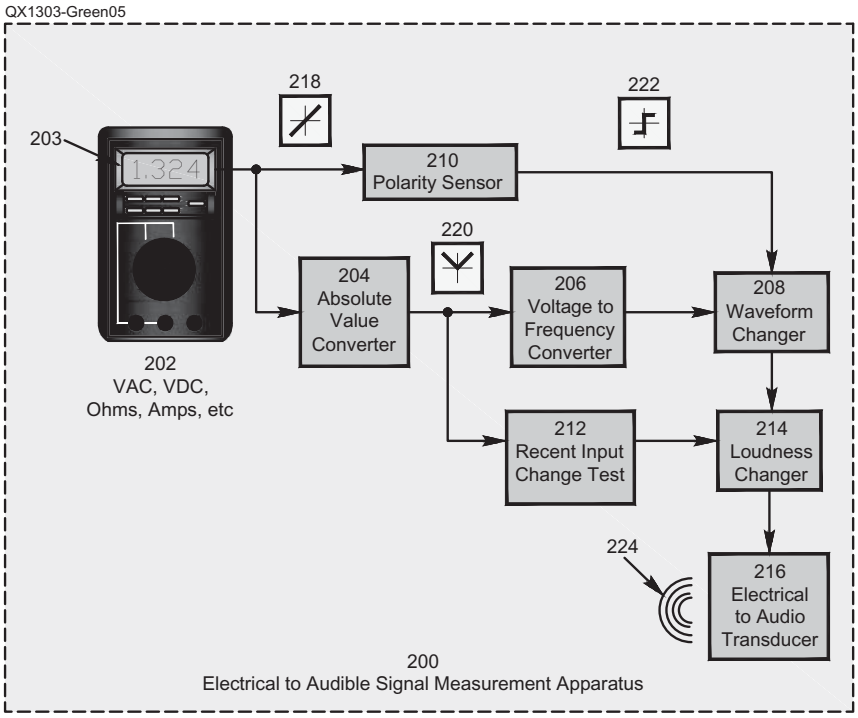

Figure 5 — An audio frequency device for audible eyes-off measurements.

Online Patent Information

Note that you need a TIFF viewer plug-in for your Internet browser to view the actual images of a patent or patent application on the USPTO website. The USPTO recommends several free TIFF viewers.[12] Of these, I have the best results with AlternaTIFF.[13] Without a TIFF viewer you can only read the text of a patent or patent application.

The message buried in this article is that you can search and read patents online. The USPTO site and the TIFF viewer are free. A huge amount of information resides on the site.

Eighteen months after the USPTO receives a patent application, they publish it on their website, whether or not they ever grant patent protection, and usually while the application is still pending. Did you know that? That's the deal you make to get patent protection. You explain and make obvious how your invention works, the USPTO makes it public and then they may or may not award patent protection. The adjective "patent" means readily visible, intelligible, or obvious. At that point, the information becomes public. If the USPTO agrees that the invention is sufficiently unique, they grant patent protection and you or your company gain exclusive rights for a number of years, while you pay increasingly larger fees to renew it.

So when you research a field of technology, consider looking among applications as well as patents. You can search on the names of known inventors, companies and areas of technology and look into what your competitors are doing. See the next great thing.

The USPTO also has a site called PAIR for Patent Application Information Retrieval that tells the licensing status of each patent and application and whether patent protection is expired for non-payment of fees.[14] PAIR even shows the correspondence between the patent examiner and the applicant's patent attorneys. It is all public information.

In Conclusion

Consider the utility of VFCs in your projects. You can send analog signals over a digital link, and you can listen to signals as they change. Enjoy!

Dr Sam Green, WØPCE, is a retired aerospace engineer. Sam lives in Saint Louis, Missouri. He holds degrees in Electronic Engineering from Northwestern University and the University of Illinois at Urbana. Sam specialized in free space and fiber optical data communications and photonics. Sam became KN9KEQ and K9KEQ in 1957, while a high school freshman in Skokie, Illinois, where he was a Skokie Six Meter Indian. Sam held a Technician class license for 36 years before finally upgrading to Amateur Extra Class in 1993. He is a member of ARRL, a member of the Boeing Employees Amateur Radio Society (BEARS), a member of the Saint Louis QRP Society, and breakfasts with the Saint Louis Area Microwave Society. Sam is a Registered Professional Engineer in Missouri and a life senior member of IEEE. Sam holds sixteen patents, with two more patent applications pending. Contact Sam at **w0pce@arrl.net**.

Notes

[1]**www.intersil.com/data/fn/fn2864.pdf**
[2]**www.national.com/ds/LM/LM231.pdf**
[3]**www.analog.com/static/imported-files/data_sheets/ADVFC32.pdf**
[4]"Bit error rate test accessory makes errors audible", Samuel Green, *Rev. Sci. Instrum.* 72(12) 4472 (01 Dec 2001).
[5]*TV Typewriter Cookbook*, Don Lancaster, ISBN 0-672-21313-3, pp 180-181, (Jan 1976).
[6]**http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=%2Fnetahtml%2FPTO%2Fsrchnum.htm&r=1&f=G&l=50&s1=5729335.PN.&OS=PN/5729335&RS=PN/5729335**
[7]**www.analog.com/en/specialty-amplifiers/log-ampsdetectors/products/index.html**
[8]**http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=%2Fnetahtml%2FPTO%2Fsrchnum.htm&r=1&f=G&l=50&s1=7898395.PN.&OS=PN/7898395&RS=PN/7898395**
[9]**http://www.arrl.org/arrlletter?issue=2009-10-15**
[10]Fully Automated DDS Sweep Generator Measurement System, *QEX*, Nov/Dec 2008, pp13-22.
[11]**http://appft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PG01&p=1&u=%2Fnetahtml%2FPTO%2Fsrchnum.html&r=1&f=G&l=50&s1=%2220100079131%22.PGNR.&OS=DN/20100079131&RS=DN/20100079131**
[12]**www.uspto.gov/faq/plugins/tiff.jsp**
[13]**www.alternatiff.com/**
[14]**http://portal.uspto.gov/external/portal/pair**

QEX

John C. Roos, K6IQL

23061 Old Kansas City Rd, Spring Hill, KS 66083, **johncroos@aol.com**

A Precision DDS for the Frequency Measurement Tests

K6IQL describes how to create a DDS with laboratory quality performance.

The Frequency Measurement Test (FMT) is held several times a year and it provides an opportunity to test our skills at determining the actual frequency transmitted by each of several test stations. These events are great fun with friendly, but stiff, competition to determine which operator has best mastered this technically challenging activity. I have participated in several of these events and had a lot of fun. Acceptable results were achieved using surplus laboratory quality synthesizers as RF frequency references. However, not all operators have access to lab-quality synthesizers.

After some personal pontification on various Internet forums pointing out that the "resolution" limitation inherent in Direct Digital Synthesizers (DDS) may compromise their use for Frequency Measurement Tests (FMT), the mental light bulb went on. I realized that perhaps I was dead wrong; no surprise there, it has happened before. Perhaps a DDS could provide an acceptable alternative to more costly equipment?

So, this project started as a simple experimental breadboard to determine if some of the difficulties inherent in using a DDS as part of a FMT measurement set up could be overcome. Like many such projects it grew in both scope and complexity, resulting in the design presented here. By the time I finished with the enhancements I had a useful PC-controlled laboratory source capable of generating highly accurate frequencies from 10 Hz to 40 MHz. The final result (see Figure 1) is a bit different from many of the DDS and Phase Locked Loop (PLL) synthesized frequency generator projects intended for amateur use. Most are aimed at applications such as a receiver Local Oscillator (LO) or transmitter Variable Frequency Oscillator (VFO). A number of clever and innovative circuits have been developed in recent years.[1] Usually these designs incorporate an onboard processor, controls and displays for the interface between the operator and the synthesizer. The control interface typically limits the frequency "setability" to 1 to 10 Hz, which is more than adequate for their intended use.

Due to the DDS "resolution" issue and lack of readout resolution of less than 1 Hz, the indicated frequency may not be the actual output frequency. That is one limitation that must be overcome for the FMT application. Additionally, the reference oscillators, while often very good, are not adequate for precision frequency measurements.

In approaching this project I wanted the synthesizer to do three things:

1. Operate from a very good internal clock, or an external GPS, Rubidium (Rb), or metrology lab quality quartz 10 MHz frequency standard.

2. Accept serial data control signals from *any* PC or processor and associated program capable of sending and receiving 5 bytes at 9600 bits per second. That allows the PC software designer to freely define the functionality and user interfaces.

Figure 1 — The Direct Digital Synthesizer is packaged in extruded aluminum housing. Front panel BNC connectors are provided for the 10 Hz to 100 kHz and the 100 kHz to 40 MHz output. LED lamps indicate lock status of the reference PLL and the presence of an external 10 MHz reference input.

[1]Notes appear on page 34.

The DDS would be implemented completely in hardware. It would be capable of executing any function supported by the DDS chip subject only to the limitations of the serial data rate. By hosting the controlling software in a PC, programs could be written in any high-level language to implement the desired test functions. Examples include frequency stepping, sweeping, or sequencing through a defined set of test frequencies.

3. Accurately report, to millihertz (mHz) accuracy, the actual output frequency generated by the DDS, which as we shall see, is not necessarily the frequency, the operator requests.

Before getting into the details of the actual design, it is useful to briefly review the Frequency Measurement Tests and some methods of accurately measuring the frequency of a HF signal. Then I will, somewhat whimsically, discuss the math behind the DDS resolution issues and the calculations required both to program a DDS and to determine the actual output frequency.

The Frequency Measurement Tests

Frequency Measurement Tests are conducted several times during the year. In April and November the ARRL publicizes and coordinates the tests.[2] In addition, an FMT interest group, FMT-nuts, is hosted on Yahoo. This group provided the test signals used in the recent ARRL FMTs. It also conducts a number of additional FMTs throughout the year. Connie Marshall, K5CM, coordinates the activity. Signals have been provided by K5CM, W6OQI, W8KSE, WA6ZTY and others.

For additional information visit the ARRL FMT website, and Connie's website, **K5CM.com.**[3] Here you'll find a wealth of information on past test results, future tests, and links to other sites describing test methods and techniques. The FMT-nuts group is very active with daily e-mail exchanges on techniques and equipment. There is a link to join the group at **K5CM.com**. If these topics are of interest, there is a lot of information on frequency standards, measurement techniques and other technical issues.

The test stations provide highly accurate signals, generally traceable to the National Institute of Standards and Technology (NIST) via GPS or WWVB. Test signals are transmitted near previously announced frequencies and are accurate to at least 0.001 Hz or 1 millihertz (mHz). This degree of accuracy is required because participating stations have gotten quite good at measuring frequency. For example, in an FMT conducted on 11 September 2011, 12 stations posted an average frequency error between 1.15 and 7.75 parts in 10^8. In the ARRL FMT conducted in April 2011, five stations posted average errors between 2.15 and 2.86 parts in 10^8. An error of 1 part in 10^8 is an error of only 70 mHz at 7 MHz!

From this the question arises, "How much measurement accuracy is needed?" The averages quoted above are one measure of performance and useful to determine overall performance. A detailed look at the results for each test frequency will often show results of 1 part in 10^9 or even better. Due to Doppler shifting of the frequency by movement of the ionosphere, a given station may have an error of only a few mHz on one test frequency and yet do poorly on another. So a good overall average may still be achieved provided there are enough good measurements to compensate for a bad one. Since it has been shown a number of times that measurement errors of less than 10 mHz are seen frequently, a system error (exclusive of ionospheric effects) of about +/- 1 mHz is in the right ballpark.

While FMTs have probably been conducted on most of the bands at one time or another, the HF tests are typically on 160, 80, 40 and occasionally 20 meters. Taking 40 meters as the typical worst case; the accuracy required of the station is simply:

$$\text{Required Accuracy} = 7\ \text{MHz} / 0.001\ \text{Hz}$$
$$= 7{,}000{,}000 / 0.001 = 7\ \text{parts in}\ 10^9$$

While this is a formidable number, the recent availability of surplus GPS-disciplined oscillators and rubidium frequency standards

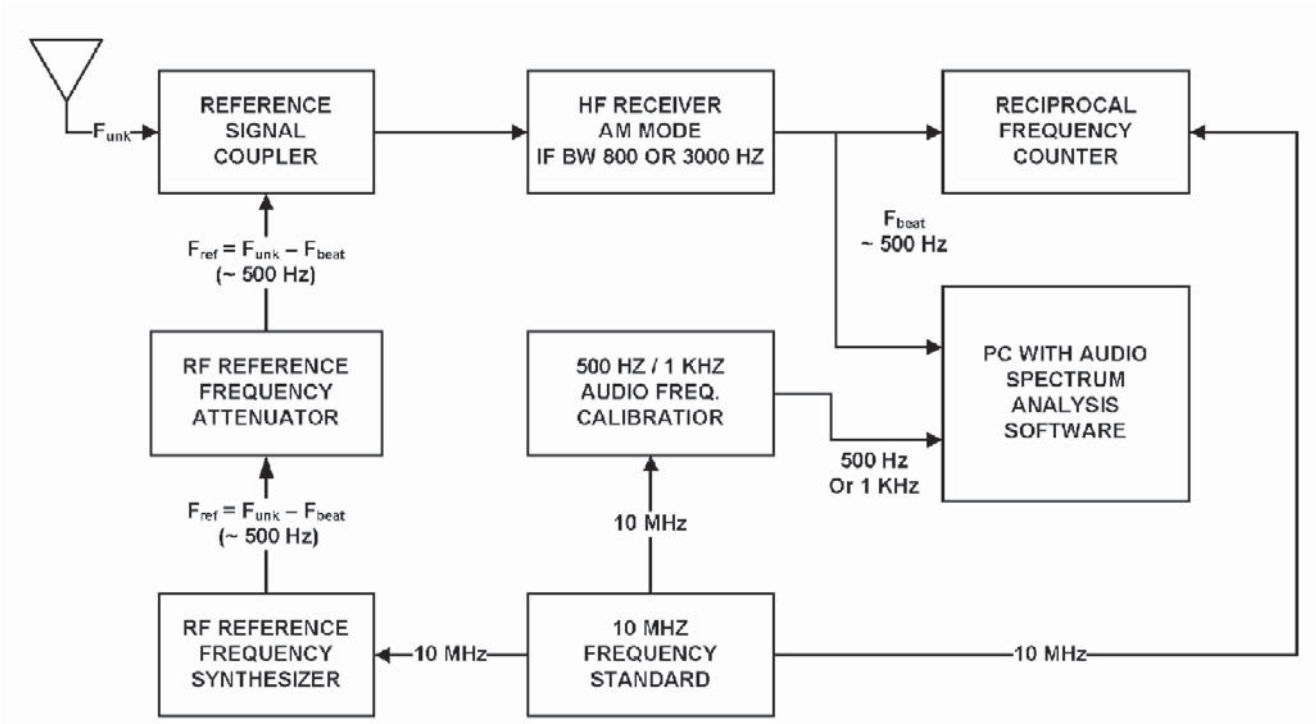

Figure 2 — The Audio Beat Method of measuring the frequency of a HF signal. All measurements are referenced to a 10 MHz Rb, GPS, or quartz frequency standard. Using this approach the stability of the receiver does not affect the measured result.

at low cost provide access to the accuracy required. Given a good frequency standard, the challenge then becomes to implement an accurate measurement system.

FMT Measurement Methods

A review of the comments posted on the ARRL FMT web page[4] will reveal a wide variety of techniques from a simple transceiver with a WWV-calibrated Temperature Compensated Crystal Oscillator (TCXO), to use of custom-designed receivers and software. Some operators, including myself, use some version of the "audio beat method" shown in Figure 2.

In this scheme, which is almost as old as radio, we have an HF receiver that receives the unknown frequency. The receiver is operated in the AM mode (the Beat Frequency Oscillator (BFO) is definitely off). An RF reference generator, generally a synthesized signal source, is coupled via an attenuator into the input of the HF receiver. Accuracy of the RF reference is determined by a local frequency standard. The frequency of the RF reference is adjusted to produce an audio beat note in the receiver. A frequency counter or PC running audio spectrum analysis software is used to determine the frequency of the beat note. Now, if the RF reference is set about 500 Hz *below* the unknown frequency, we can define:

F_{unk} = Unknown frequency from the test station

F_{ref} = Frequency of the RF Reference Synthesizer

$F_{beat} = F_{unk} - F_{ref}$

By determining F_{ref} and F_{beat} the unknown frequency is simply

$$F_{unk} = F_{ref} + F_{beat} \quad \text{[Eq 1]}$$

An often underappreciated feature of this method is that the accuracy is determined *solely* by the accuracy of the RF reference and the beat frequency measurement. Stability of the receiver LO does not appear as part of the frequency calculation. As long as the receiver is stable enough to maintain the unknown and RF reference signals within the IF pass band, and has sufficient selectivity to keep out interference, it has no effect whatever upon the measurement. However, to avoid having the receiver Automatic Gain Control (AGC) suppress the desired signal, the attenuator must be adjusted so the injected reference frequency is at a level equal to or slightly below that of the incoming signal.

Nifty as this method may be, measuring the beat frequency to the required precision within the two minute FMT key-down measurement period is a challenge. In Figure 2, two methods are indicated.

It is possible to use a frequency counter, as shown, but not of the conventional type. In a conventional counter, the resolution is determined by the gate time of the counter. A one-second gate time yields a +/-1 Hz resolution. This is hardly adequate for our purposes. Even if the gate time were extended to the entire key down period, 120 seconds, the resolution is only +/- 0.083 Hz. This is not too practical due to loss of counts during fades, false triggers and other problems. A "reciprocal" counter should be used. Reciprocal Counters[5] count a very high-speed internal clock that is gated by the input signal and counted. It then computes the frequency. At audio frequencies a reciprocal counter can produce very high-resolution measurements with short gate times. These devices easily achieve 1 mHz accuracy at 500-1000 Hz with only a 1 second gate time. Multiple measurements may be collected during the two minute key down period. These measurements may be averaged to obtain a better solution.

However, as explained in the referenced papers, all counters suffer triggering errors under poor Signal to Noise Ratio (SNR) conditions. A better solution is to use a PC that employs an audio spectrum analysis program in conjunction with the PC sound card. The sound card acts as a very high dynamic range A/D converter. One very capable program used for FMTs is *Spectrum Lab,*[6] authored by Wolfgang Buscher, DL4YHC. It is available on the Internet as a free download.

Spectrum analysis software employs the Fast Fourier Transform (FFT) technique to realize the equivalent of a contiguous group of very narrow bandwidth frequency bins (filters). This has many advantages:

• The bin bandwidth can be very narrow, 1 Hz or less, thus greatly improving the processing signal to noise ratio. A 90-dB dynamic range is common and signals you cannot hear are readily visible on the spectrum display.

• Some programs, including *Spectrum Lab*, can interpolate the frequency measurement across bins resulting in truly outstanding measurement resolution – 1 mHz or better is easily demonstrated.

• Averaging can be used, further improving both SNR and accuracy. It is possible to get 100 or more individual measurements in a 120 second key-down period.

• The measurements may be output to a text file and then to a spreadsheet for later processing, averaging, application of statistical functions or making graphs of the data.

• Some programs provide a "waterfall" as well as a conventional spectrum display. The waterfall will show real time shifts in the frequency of the unknown due to ionospheric effects. This can be very useful in determining the unknown frequency.

• The audio beat tone may be recorded in real time and played back later using multiple passes to better optimize the processing.

A potential disadvantage is that while great precision is possible, it is the crystal in the processor that controls the sound card sample rate that determines the ultimate accuracy.

Considering the accuracy goal to be 1 mHz, this at first seems a problem, but it really is not. Just how accurate is a 1 mHz measurement at 500 to 1000 Hz? Using 1000 Hz as about the highest beat frequency one would want to employ we have:

Required Accuracy = 1000 Hz / 0.001 Hz
= 1 X 10^6 or 1 part per million.

As very low cost crystals are used for computer clocks, the required initial calibration accuracy is not inherent. However, some spectrum analysis software, including *Spectrum Lab*, have routines to calibrate out the clock frequency error. Inputting a known accurate frequency from an audio frequency calibrator such as shown in Figure 2 does this. The operator simply tells the routine what frequency is displayed and the true frequency. From this the routine computes and stores the appropriate correction.

How well does this work? Actually, surprisingly well. If the computer is warmed up for an hour or so before calibration, the issue is reduced to the thermal drift of the crystal, which is not nearly as serious as the initial accuracy error.

Using a 1000 Hz tone derived from a rubidium frequency standard I did some tests. The processor was old Compaq EVO desktop machine purchased surplus for about 100 bucks. After a 1 hour warm-up the clock was calibrated. After calibration the indicated frequency was recorded every 10 minutes. It was determined to be within +/- 0.0005 Hz of 1 kHz for 6 hours at which point I stopped the test. Repeated checks over a period of months have shown the calibration to be stored and maintained. Ensuring the computer is calibrated and used after it has reached a stable temperature is the only precaution required.

With this overview of a typical FMT measurement technique I'll now get into the innards of typical DDS chips, their limitations and how to use one to achieve accurate results.

DDS Operation and Programming

Without doubt the advent of the DDS in integrated circuit form has eased the life of RF engineers. I have used them in two commercial designs: a VHF/UHF ILS transmitter

synthesizer, and as the LO for a VHF/UHF ILS monitor receiver. In both instances the performance was comparable to earlier PLL designs and with considerably less complexity and headache.

However, with any technology there is no free lunch. Some of the issues that must be considered when employing a DDS include:

• Spurious Outputs – DDSs produce spurious outputs. There are many articles[7] in which the spurious issues and means to mitigate spurs are treated, and since our focus is frequency measurement, please refer to the sources listed in the Notes.

• Finite Frequency Resolution – In DDS speak, frequency resolution is a more polite term for frequency steps that have a predictable but varying relation to the requested frequency. As I will show, the frequency you ask for is almost never the frequency the DDS produces. Close, but not exact.

The frequency resolution issue is common to all of the simple DDS chips and at the heart of the FMT accuracy problem. To illustrate this I will define a fictional generic DDS, the GDDS 200-32, a product of Mythical Semiconductors Corp. (MSC). A block diagram and specification sheet for the MSC GDDS 200-32 is shown in Figure 3.

The first four specifications listed in the figure essentially define the overall performance of almost any simple DDS:

• Maximum clock frequency. This is obvious, but as we shall see using a lower clock frequency improves the tuning resolution.

• Width of the tuning word. This is fixed by the DDS design. The width of the tuning word in conjunction with the actual clock frequency determines the worst case frequency error.

• Digital to Analog converter resolution. This determines the achievable spurious level. More D/A bits imply lower spur levels. The spur level will be approximately 20 x Log (2^N), when N is the number of D/A bits.

• Frequency resolution. This is the minimum step size for a given clock frequency. It is also an indication of the worst case frequency error to be expected when the DDS is operated at a specified (normally maximum) clock frequency.

Programming and determining the actual output frequency of a DDS comes down to two versions of the same simple equation:

$$F_{out} = (T_{wrd} \times F_{clk}) / 2^N \qquad [Eq\ 2]$$

And solving for T_{wrd} we have

$$T_{wrd} = (2^N \times F_{out}) / F_{clk} \qquad [Eq\ 3]$$

Where:

F_{out} = The Output Frequency from the DDS

T_{wrd} = The Tuning Word Input to the DDS from the controlling processor

F_{clk} = The DDS clock frequency derived from a frequency standard

N = Number of Bits in the Tuning Word

All frequencies are in the same units, Hz, kHz, or MHz. T_{wrd} is the decimal equivalent of the binary word to be loaded into the DDS chip, and therein lies our error problem.

At first glance the process seems simple; use Equation 3 to solve for T_{wrd} and then load that value into Equation 2 for the DDS. Out pops the desired F_{out}. Well, not quite, because in most instances the result of the calculation for T_{wrd} will be a floating-point number. That is, it will have a fractional part. While Eq 2 will accept a floating-point number, in real life the physical DDS chip will not. The T_{wrd} input to the DDS hardware chip must first be converted to an N-bit binary number, which is (unless otherwise defined) an integer, with no fractional part.

To illustrate this I will refine Equations 2 and 3 and redefine some of the variables to help keep things sorted out. Now we have:

$$F_{pred} = (T_{wint} \times F_{clk}) / 2^N \qquad [Eq\ 4]$$

$$T_{wfp} = (2^N \times F_{req}) / F_{clk} \qquad [Eq\ 5]$$

Where:

F_{pred} = The *Predicted* Output Frequency from the DDS chip as given by Eq 4

T_{wint} = The Tuning Word Input to the DDS Eq 4, a N-bit *integer* number

F_{req} = The *Requested* Output Frequency from the DDS, the input to Eq 5

T_{wfp} = The computed tuning word from Eq 5, a *floating point* number

F_{clk} = The DDS clock frequency derived from a frequency standard

N = Number of Bits in the Tuning Word

Now for an example using our MSC GDDS 200-32 DDS chip. We have the following:

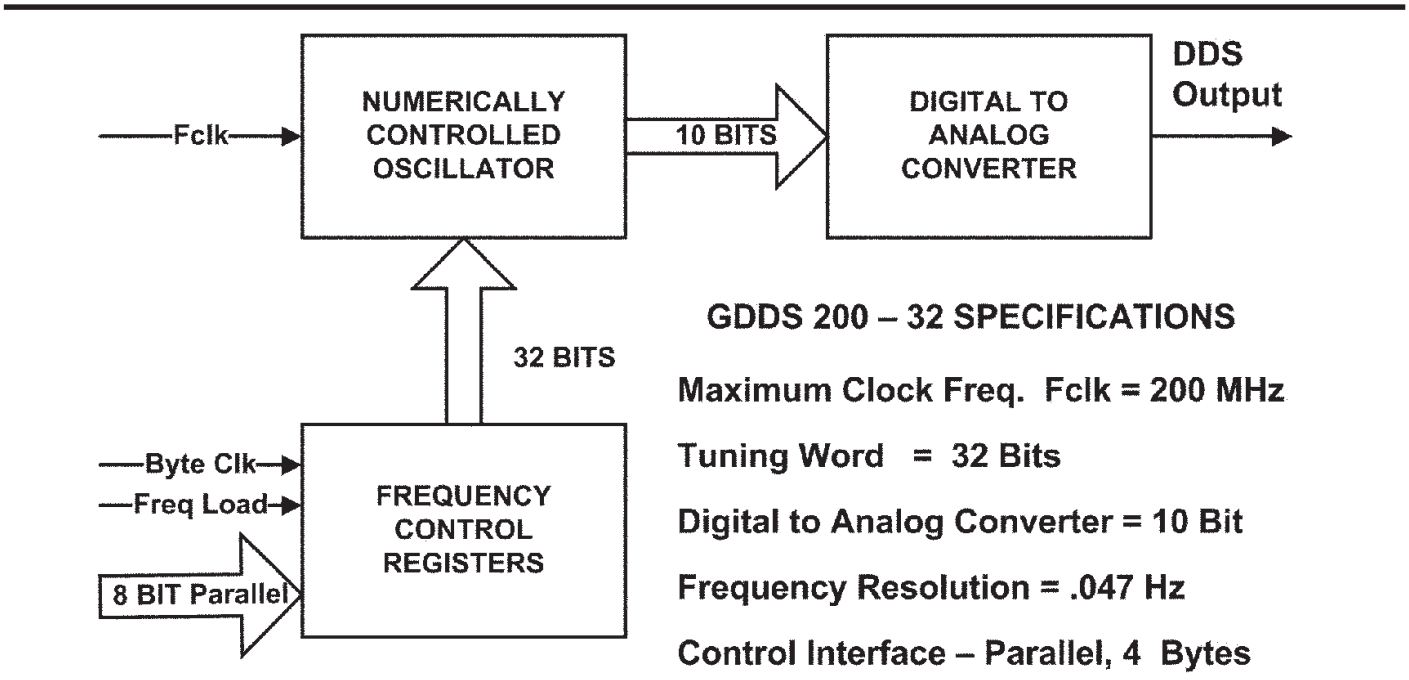

Figure 3 — Data sheet for the MSC GDDS 200-32 DDS IC.

N = 32 bits

F_{clk} = 100 MHz = 100,000,000.0 Hz

And suppose we want a 1.000,005.0 MHz output frequency, so

F_{req} = 1,000,005.0 Hz (i.e. 5 Hz above 1 MHz)

Using Equation 5 we compute

$T_{wfp} = (2^{32} \text{ X } 1{,}000{,}005.0) / 100{,}000{,}000.0$
$= 42, 949, 887.708$

Now since the tuning word input to the DDS must be an integer, we drop the .708 on the right of the decimal point, leaving; T_{wint} = 42, 949, 887.

Next this value is plugged into Equation 4 to predict the actual output frequency from our MSC GDDS 200-32 DDS chip.

$F_{pred} = (42, 949, 887 \text{ X } 100{,}000{,}000.0) / 2^{32}$
= 1, 000, 004. 9835 Hz

Subtraction of the Requested Frequency from the Predicted Frequency gives the DDS error:

Error = $F_{pred} - F_{req}$ = 1, 000, 004. 9835 – 1,000,005.000 = –0.0165 Hz error

So what happened? Our data sheet, Figure 3, specified the resolution of the GDDS 200-32 as 0.047 Hz and *the calculated error is only* 0.0165 *Hz*.

Two factors influence the DDS error:

1. The clock frequency. If the clock frequency is halved, the "resolution per step" number is also halved. Further the specifications for resolution are always given at the highest acceptable clock frequency for the particular chip — in this case 200 MHz.

If we take Equation 4 and make T_{wint} =1, we get Equation 6

$$F_{act} = (1 \text{ X } F_{clk}) / 2^N = F_{clk} / 2^N \text{ when } T_{wint} = 1$$

The specified resolution is found from simply dividing…

$$\text{Resolution Hz} = F_{clk} / 2^N = 200{,}000{,}000 / 2^{32} = 0.047 \text{ Hz} \quad \text{[Eq 6]}$$

So we see that as T_{wint} = 1,2, 3, etc; then F_{act} = 0.047, 0.094, 0.141 and so on throughout the entire frequency range of the DDS. So if the clock frequency is halved to 100 MHz then the step size is also halved to 0.0235, 0.047, 0.0705 Hz. etc.

2. The "resolution" does not imply a constant error; rather it is the worst case frequency error that will be encountered for a given clock frequency and number of bits in the control word. In fact, the actual error changes continually as a function of the requested frequency. This is due to rounding off of T_{wfp} to obtain the integer value for T_{wint}.

I call this variable error the "Round Off Error". Because of the round off, *the actual frequency produced will always be less than the frequency requested*, but the error will not exceed the "resolution" as determined by Equation 6 above. The closer the part of the floating point tuning word to the right of the decimal approaches 1 (and the next higher integer value) the greater the round off tuning error. If the output from Equation 5 is a perfectly even number (nnnn.0) then there is no round off and there will be no frequency error. Outputs with no frequency error occur many times across the tuning range of a DDS.

As the requested frequency is incremented, this round off error is repeated in a cyclical manner. This may be demonstrated analytically. By using a spreadsheet to repeat the above calculation starting at a base requested frequency and then incrementing it in small steps like 1 Hz we see what happens.

Figure 4 is a graph plotted from an *Excel*

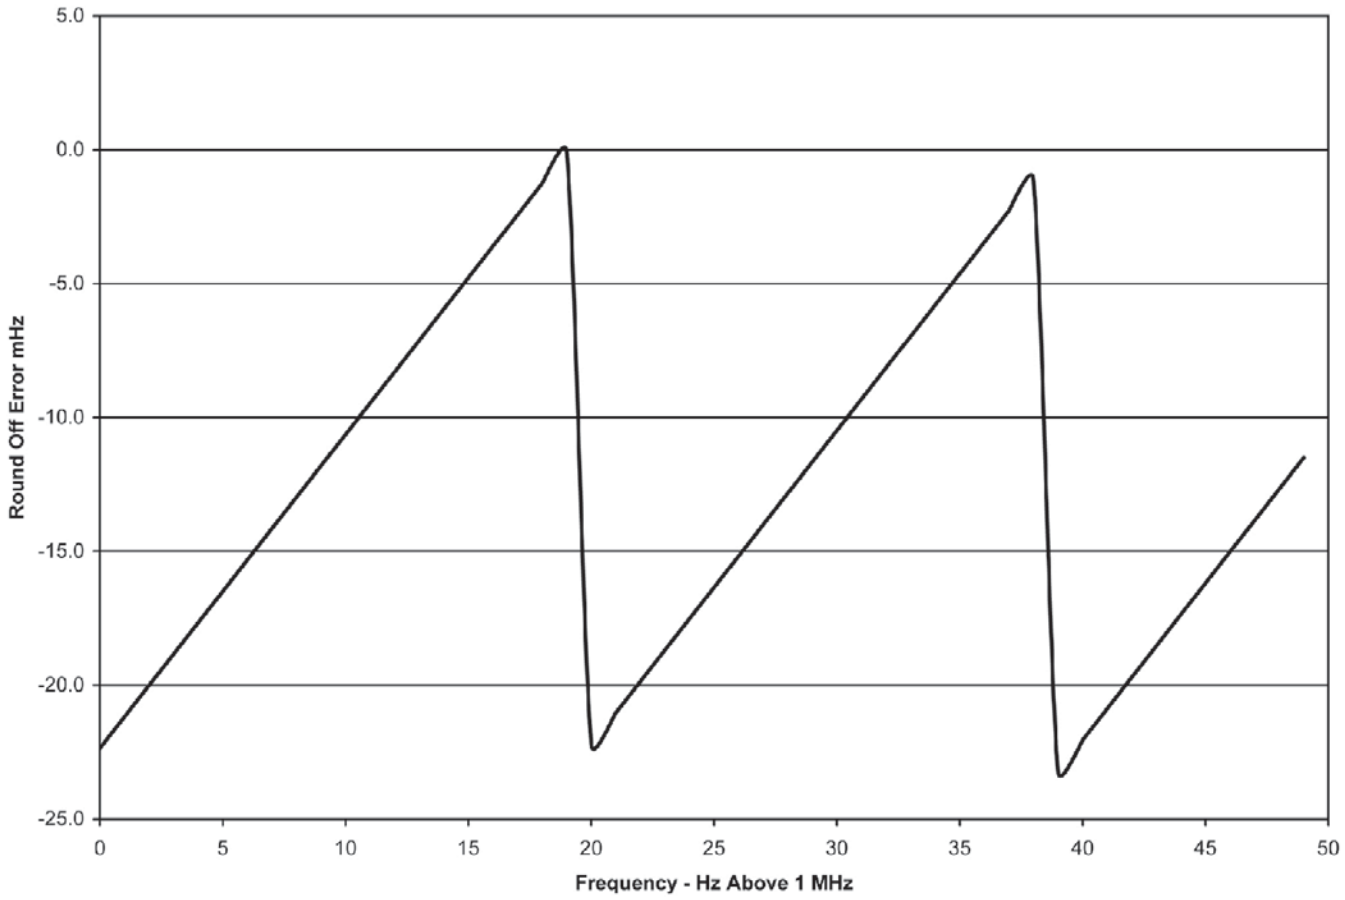

Figure 4 — The predicted DDS output frequency is compared with requested frequencies in 1 Hz steps from 1,000,000 Hz to 1,000,050 Hz. The error varies cyclically from 0 to about –23 mHz. Note the error is always below the requested frequency.

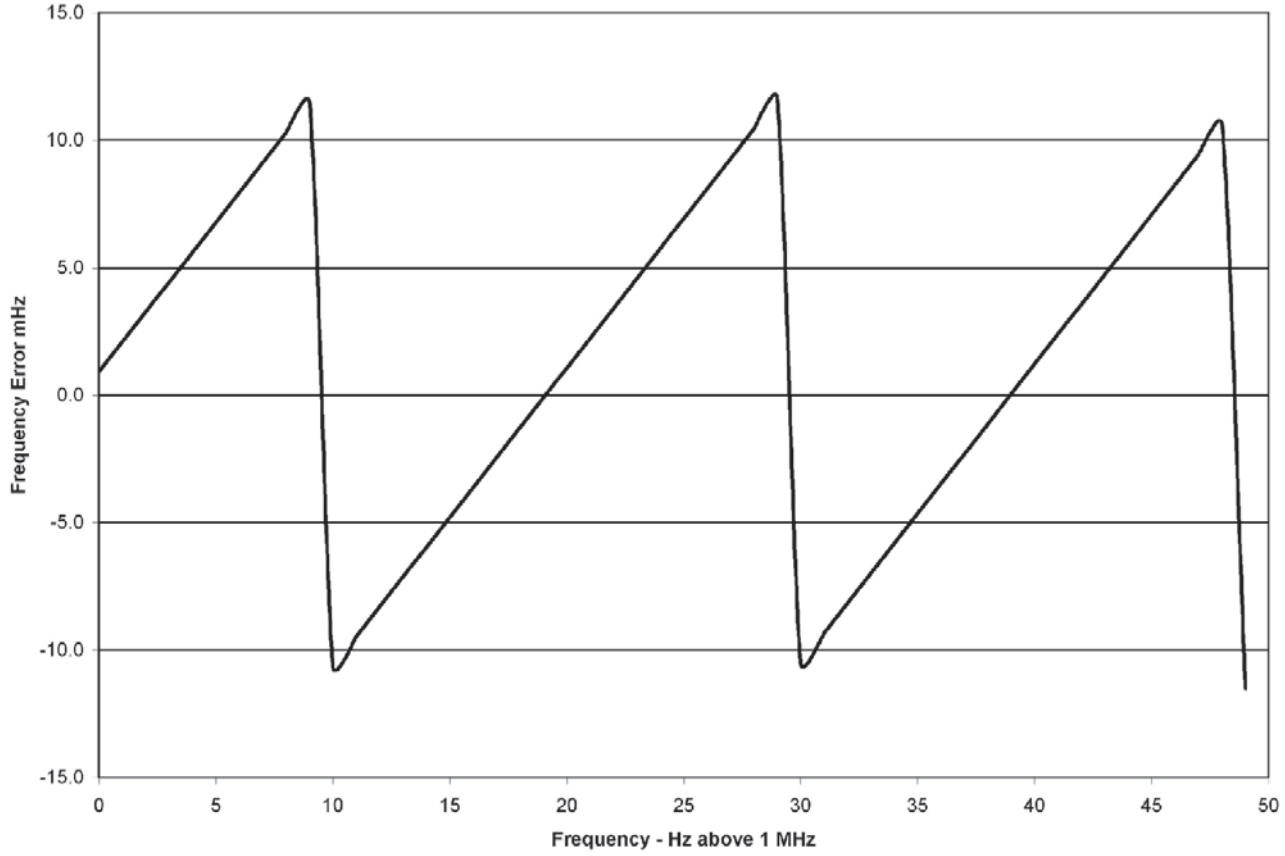

Figure 5 — In this plot the frequency error is made symmetrical about the requested frequencies by adding 0.5 to the floating point tuning word prior to converting it to an integer. The error is now about +/– 11.5 mHz instead of +0 to –23 mHz.

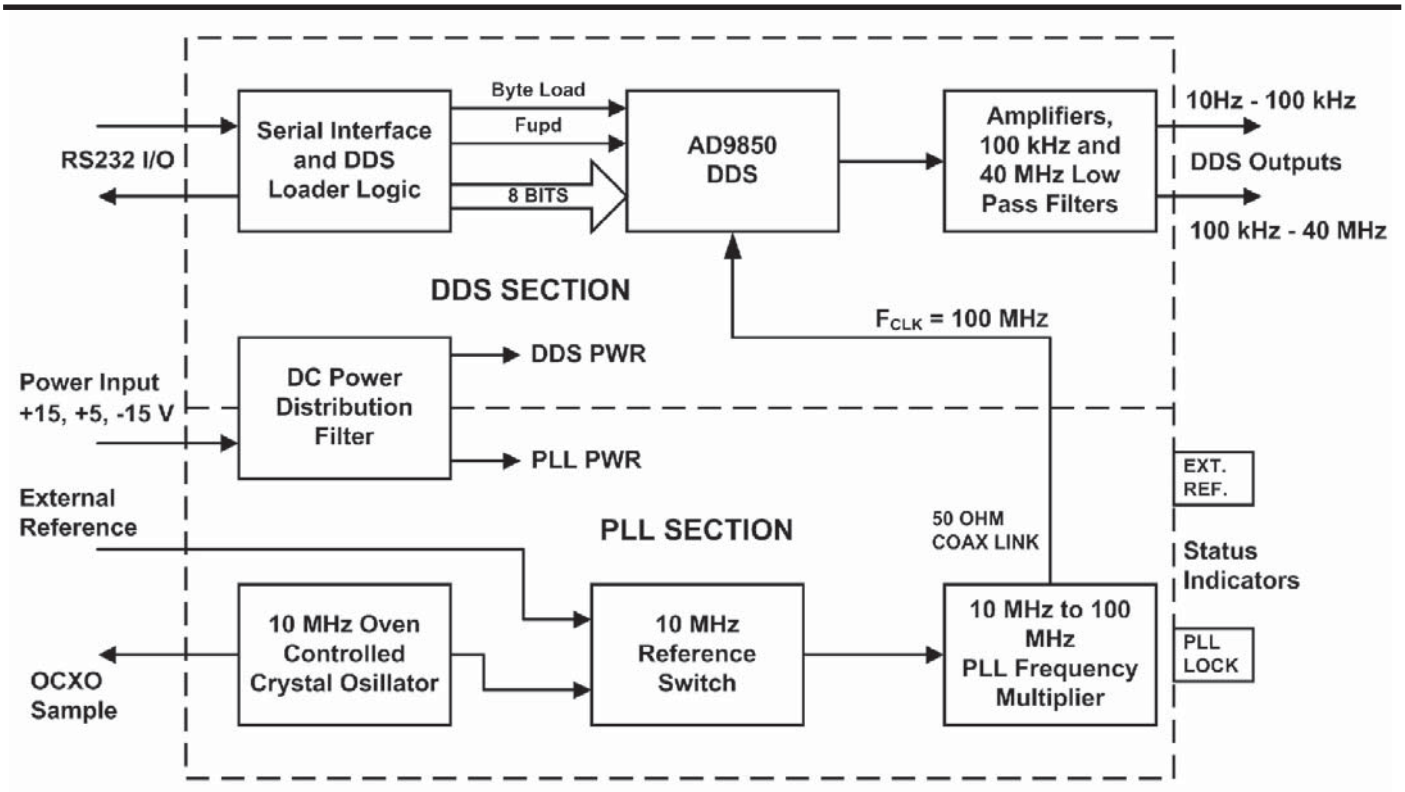

Figure 6 — The FMT DDS is divided into two sections: the DDS circuits and the 10 to 100 MHz clock frequency multiplier. DC power filtering and a coaxial connection from the 100 MHz clock multiplier to the DDS increase the isolation between sections.

spreadsheet. It shows how the frequency error changes, using a 100 MHz clock. The requested frequency was incremented in 1 Hz steps from 1,000,000 Hz to 1,000,050 Hz. Note how the DDS output frequency is always below the requested frequency due to rounding off T_{wfp} to obtain T_{wint}.

If desired, the actual DDS output frequencies can be compensated to vary symmetrically about the requested frequencies by a simple mathematical trick. Referring again to Equation 5 we can force the round off process to round T_{wfp} up to the next higher integer value if the fractional part of T_{wfp} is more than 0.5 simply by adding 0.5 to T_{wfp}.

$$T_{wfp} = (2^N \text{ X } F_{req}) \ / F_{clk}$$

Equation 5 is replaced by Equation 7 in which 0.5 is added *prior* to converting T_{wfp} from a floating point to an integer number.

$$T_{wfp} = [(2^N \text{ X } F_{req}) \ / F_{clk}] + 0.5 \qquad \text{[Eq 7]}$$

Figure 5, shows that if the compensated expression 7 is used and the previous spreadsheet calculation is repeated and plotted, we see that the round off frequency error is now symmetrical about the requested frequency. Now the peak error is halved, and can be specified as being +/- 11 mHz as opposed to + 0 / - 22 mHz. The DDS does not work any differently but the worst case peak error has been halved. This is useful in applications where a simple +/- tolerance on the frequency uncertainty, rather than the exact frequency suffices.

After working through the above analysis, my next step was to implement a DDS in hardware. The objective was simply to prove that F_{pred}, as predicted by the math would be identical to the DDS *actual* output frequency (F_{act}) to an accuracy of better the 0.001 Hz. Recalling Equation 1…

$$F_{unk} = F_{ref} + F_{beat}$$

There is no requirement that F_{ref} is a nice integer frequency like the one I obtained from my lab generators, all we really need for an FMT is to know the frequency accurately. If the DDS control software can provide this, then a DDS can be used to implement a very capable FMT facility.

FMT DDS Hardware Specification

To frame the discussion of my FMT DDS design, first look at Table 1. It lists the specifications achieved after the dust had settled. For this table the following terms are defined as:

Requested Frequency — The operator input to the DDS control software

Predicted Output Frequency — The DDS output frequency *predicted* by the software using Equations 4 and 7; including the Round Off Error

Maximum Round Off Error — Maximum expected difference between Requested and Predicted Frequency

Actual Output Frequency — Actual frequency produced by the DDS hardware and confirmed by measurement.

DDS Output Frequency Error — Difference between the Predicted Output Frequency and an Actual Measured DDS Output Frequency. Ideally this is less than 0.001 Hz.

FMT DDS Block Diagram

Figure 6 is a top-level block diagram of my FMT DDS showing the overall configuration and the external interfaces. The design is divided into two distinct sets of circuitry, the DDS Section, and the 10X PLL Multiplier Section. While both are on the same board the two circuits are not physically mingled and the power is filtered to keep signals from one section from entering the other. The 100 MHz DDS clock, generated by the clock frequency multiplier is routed directly to the DDS clock input via a matched 50-Ω coax cable. I did this to make the PLL output and the DDS input accessible for testing and some additional experiments. The only other connection between the sections is the filtered dc power.

Referring again to the block diagram, the DDS section has three major elements; the serial interface and DDS loader logic, the Analog Devices AD9850 DDS chip and the output amplifiers and low pass filters.

The serial interface and DDS loader logic sequentially receives each of five bytes from the PC. As each serial byte is received it is converted to parallel format and routed to the DDS as an 8 bit parallel byte. It is also converted back to serial and sent back to the PC to check the integrity of the serial link. This proved valuable in de-bugging the design. Additionally, the PC software awaits the reply for each byte before sending the next one. This provides a simple form of software flow control.

Table 1
FMT DDS Performance Specifications

Output Frequency Ranges
Audio: 10 Hz to 100 kHz
RF: 100 kHz to 40 MHz

Output Levels
Audio: Open Circuit 0.7 V_{rms}, 0 to + 2 dB, 10 Hz to 100 kHz
RF: 50 Ω load 12.7 dBm @ 100 kHz to 7.5 dBm @ 40 MHz

Round Off Frequency Error (Maximum): +/– 11.6 mHz

Measured Output Frequency Error relative to Predicted Output Frequency: + 0.1 to – 0.2 mHz

Required Accuracy of External Reference to Achieve 0.001 Hz Accuracy at 40 MHz: 4 parts in 10^{10}

External Reference Input Frequency: 10 MHz

External Reference Input Level Range: –10 to +16 dBm (50 Ω)

Internal Reference: Bliley NVG47AE 10 MHz OCXO

Internal Reference Stability (0 to 50° C.): Less than +/– 30 parts in 10^9

Internal Reference Aging Rate per Year: Less than +/– 50 parts in 10^9

Internal to External Reference Switchover: Automatic

DDS Clock Frequency: 100 MHz

DDS Tuning Word: 32 Bits

PC Control Interface: RS-232 Bi-directional, 9600 Bits/Second, Five 8 bit bytes, Start, Stop, No Parity

Supported DDS Control Functions: Phase — MS Byte, Frequency 4 Bytes (32 bits). See AD9850 Data Sheet for details.

dc Power Requirements: +15 V @ 168 mA, –15V @ 13 mA, 5V @ 823 mA

While savvy DDS aficionados will be quick to point out that the AD9850 is not the latest or the best DDS chip around; it was perfect for my application. As indicated on the block diagram, the control protocol is extremely straightforward. The AD9850 supports both a serial and a parallel protocol. I used the parallel protocol that is explained in detail on page 10 of the AD9850 data sheet.[8] It is summarized as follows. The control word has a total of 40 bits. The first 8 bits are used for phase control and test functions, while the remaining 32 comprise the 32-bit frequency control. To update the frequency the entire 40 bits must be used. While a serial input to the DDS is convenient, it is not especially compatible with the serial output from a PC COM port because the 40 bits must be input to the DDS in one long string without any start or stop bits. It is easier to send the serial data from the PC COM port to the DDS as 5 normal bytes, convert it to parallel at the DDS, and then load it into the DDS as 5 individual bytes using the parallel load format. Note the bytes are sent in reverse numerical order. Byte 4, the most significant byte is sent first. It is the phase control byte. Then bytes 3, 2, 1 and 0 are sent and comprise the 32-bit frequency control integer.

Because the DDS hardware requires the phase control byte to be sent first, all 5 bytes must be sent, though the AD9850 only uses four to achieve 32-bit frequency control. The fifth byte can be used to implement RF phase control. This is not used in the FMT application, though the DDS hardware described here will support it. For simple frequency control byte 4 (the first sent) is simply sent as all "0."

Upon receiving the first byte, the F_{upd} line is set low by the loader logic. The first byte is presented to the DDS and a rising edge on the Word Load line clocks it into the DDS. This process continues for four subsequent bytes. The loader logic counts the bytes and after the fifth byte has been loaded, the F_{upd} is set high which updates the DDS to the new frequency. For me the appeal of this scheme is the simplicity. The entire control logic was implemented with a few Complementary Metal Oxide Semiconductor (CMOS) chips and an on-board micro-controller was not needed.

From the DDS output the synthesized frequency is routed to a pair of low pass filters and output amplifiers. Two bands are provided having separate outputs. The AF band extends from below 10 Hz to 100 kHz. An active low pass filter rolls off output above 100 kHz, primarily because the output amplifier will not perform too well above this frequency. The RF output is active from below 100 kHz to 40 MHz. Above 40 MHz an anti-aliasing filter rolls off the response. This filter prevents the 60 MHz alias that would show up at the output, when the DDS is set for 40 MHz, from appearing on the output. The AF and RF output levels are internally adjustable so they can be set to known levels. External step attenuators are then used to reduce the signals as required by the application.

The 10X PLL frequency multiplier, which is shown in the lower half of Figure 6, generates a 100 MHz clock for the DDS. This PLL, referenced to a 10 MHz frequency standard, locks a 100 MHz Voltage Controlled Oscillator (VCO) to provide a low spurious 100 MHz clock. The 10 MHz reference is derived either from an on-board Oven Controlled Crystal Oscillator (OCXO) or from an external reference depending upon the required frequency accuracy. Selection of the reference source is automatic. Application of an external reference is sensed and the OCXO is shut down and the reference switch routes the external reference to the PLL. A sample of the OCXO frequency is routed to the rear panel enabling the oscillator to be calibrated by an externally accessible pot.

Front panel external interfaces, Figure 1, include the AF and RF output BNC connectors and two status indicating LEDs. One status indicator is illuminated when the 10X multiplier PLL is correctly in lock. The other illuminates when an external 10 MHz reference is present and of sufficient amplitude to enable switching from the internal to the external reference.

Rear panel interfaces, Figure 7, include BNC connectors for the 10 MHz external reference input, and the OCXO frequency sample output. A hole provides access to the adjustment pot for calibration of the OCXO frequency. Also located on the rear panel are the DB-9 connector for the serial control data, and the dc power connector.

DDS Section Schematic

Full size copies of the schematics and other pertinent data will be posted on the *QEX* website at **www.arrl.org/qexfiles** in a zip file named *ROOS.zip*. Just look for the March/April 2013 issue in the Table.

I will also provide full (B) size PDFs directly to anyone requesting copies by e-mail. It is probably easier to follow the discussion with a full size schematic at hand.

Figure 8 is the schematic of the DDS section. The logic timing diagram, Figure 9, is a mixed signal logic and analog scope[9] screen capture showing events as the DDS is programmed to step from 200 to 400 Hz. I used a step from 200 to 400 Hz so the frequency step would be displayable on the same time scale as the logic events. The logic analyzer probe display lines are labeled D0 to D15 for identification

Referring to Figure 8, power and optionally the control data are input to the board at TB1. TB1 is a plug-on terminal block that mates with a header on the Printed Wiring Board (PWB). This is convenient in the test environment. A separate DB9 connector is

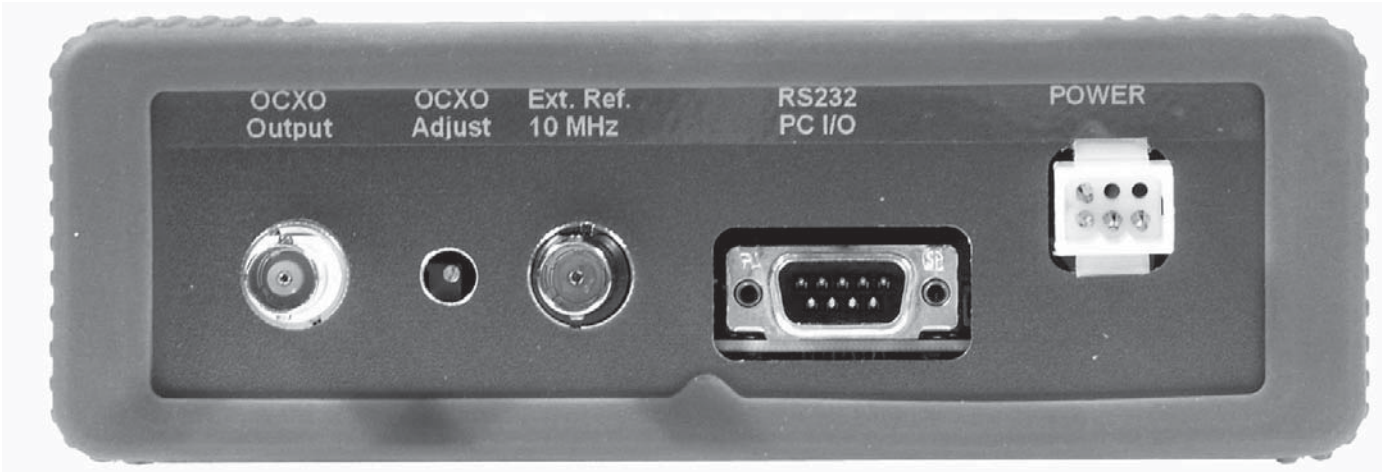

Figure 7 — Located on the rear panel of the FMT DDS are the External Reference Input, the OCXO Sample output, the OCXO adjustment, and the PC and power connectors.

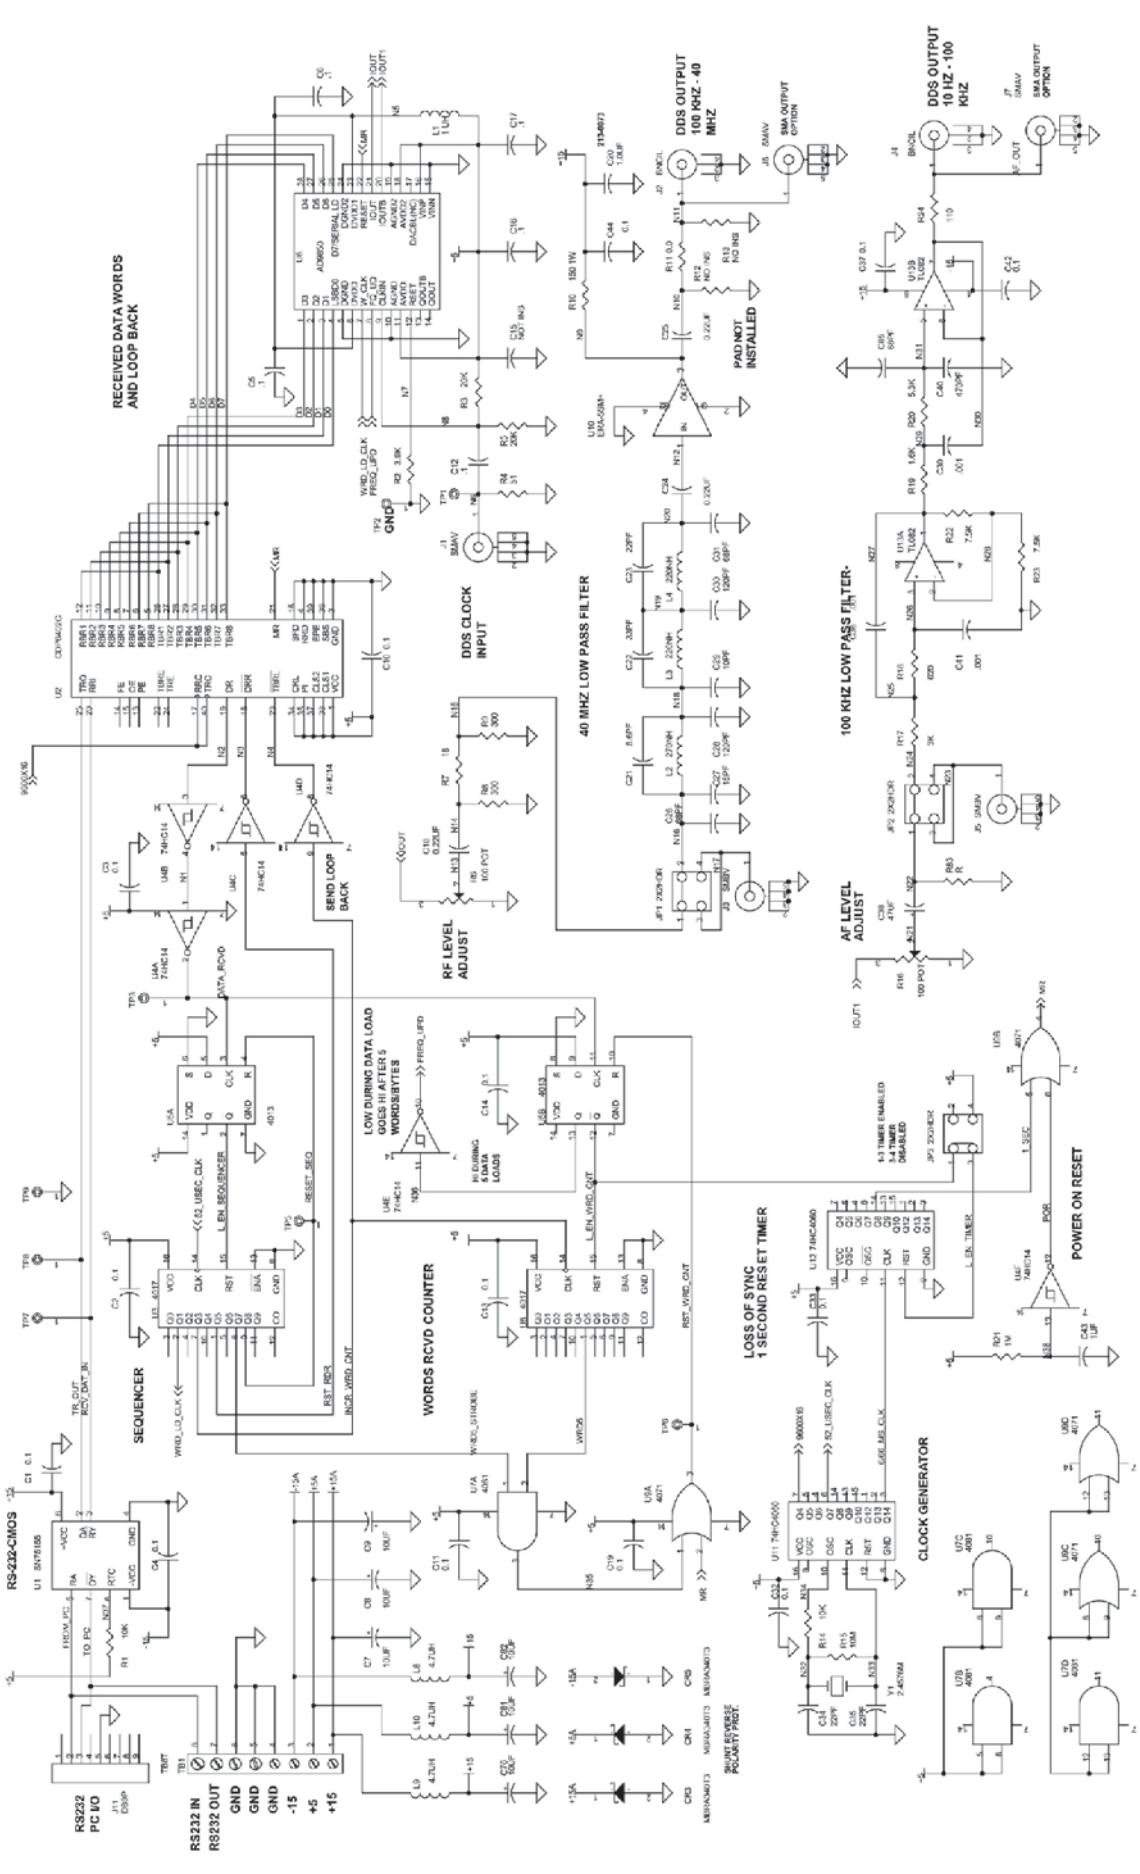

Figure 8 — Schematic of the DDS section of the synthesizer. Please see *Roos.zip* on the *QEX* website at www.arrl.org/qexfiles for a larger version of this drawing.

provided for the PC I/O. When installed in the cabinet, a short jumper cable connects TB1 to the rear panel Molex power connector.

Key to operation of the logic are the two 4017 CMOS counters, U3 and U8. The 4017 counter is handy when a simple set of events must take place in sequence because the chip has 10 internally decoded outputs. These may be connected to provide strobes as required without a lot of decoding logic. The scheme is effective and avoids race conditions, provided super high speed is not needed. U3 is the "sequence" counter. It steps the logic through its paces as each byte is received. U8 is the "words received counter." U8 keeps track of the number of bytes received and asserts the "frequency update" after receipt of the 5th byte. (Though 8 bits are commonly a Byte, the Analog Devices D data sheet uses the term Word and I have followed that convention where it applies).

RS-232 level frequency data from the PC is converted by U1 to TTL/CMOS logic levels and sent on to U2, a CDP6402C Universal Asynchronous Receiver Transmitter (UART).[10] The UART converts each byte from serial to parallel and the parallel data is routed to U6, the AD9850 DDS. The data in parallel format is also looped back to the transmit side of the UART for transmission back to the PC. The data rate is fixed at 9600 bits per second.

Clocks for the data communications and the sequencer are provided by Y-1, a 2.4576 MHz crystal, and U11. The sequencer uses a 52-microsecond clock, and while a faster clock would speed up the logic it would have little effect on the DDS update rate. Figure 9 shows that most of the update time is consumed by the PC

Referring to the schematic and Figure 9, D15 shows the first serial byte (all zeros) being received. The UART outputs the pulse seen at D14, TP3 Data Received. This pulse clocks U5A, which clears the reset at U3-15 to assert D13, Sequencer Enabled Low. At the same time it also clocks U5-B which clears the reset from U8-14 to enable U8 to count the received words. The Q output from U5-B is inverted by U4-E and routed to the DDS asserting the "Frequency update" line low. This is shown as D11, U4-10 on Figure 9. The DDS can now accept new data.

Next U3 begins counting up from 0 to 1 asserting U3-2, Word Load Clock high, indicated as D12, in Figure 9. This signal is routed to the DDS U6-7 and clocks the first 8-bit word into the AD9850.

U3 continues counting up to 3 where the decoded output on U3-7, Increment Word Count, performs two actions. It is routed to the word counter U8 clock input at U8-14 to increment the received word from 0 to 1. It is also sent to U4, where after inversion it is applied to Pin 23 of the UART. See trace D10, U3-7 Increment Word Count / Send Reply in Figure 9.

The high to low transition at Pin 23 loads

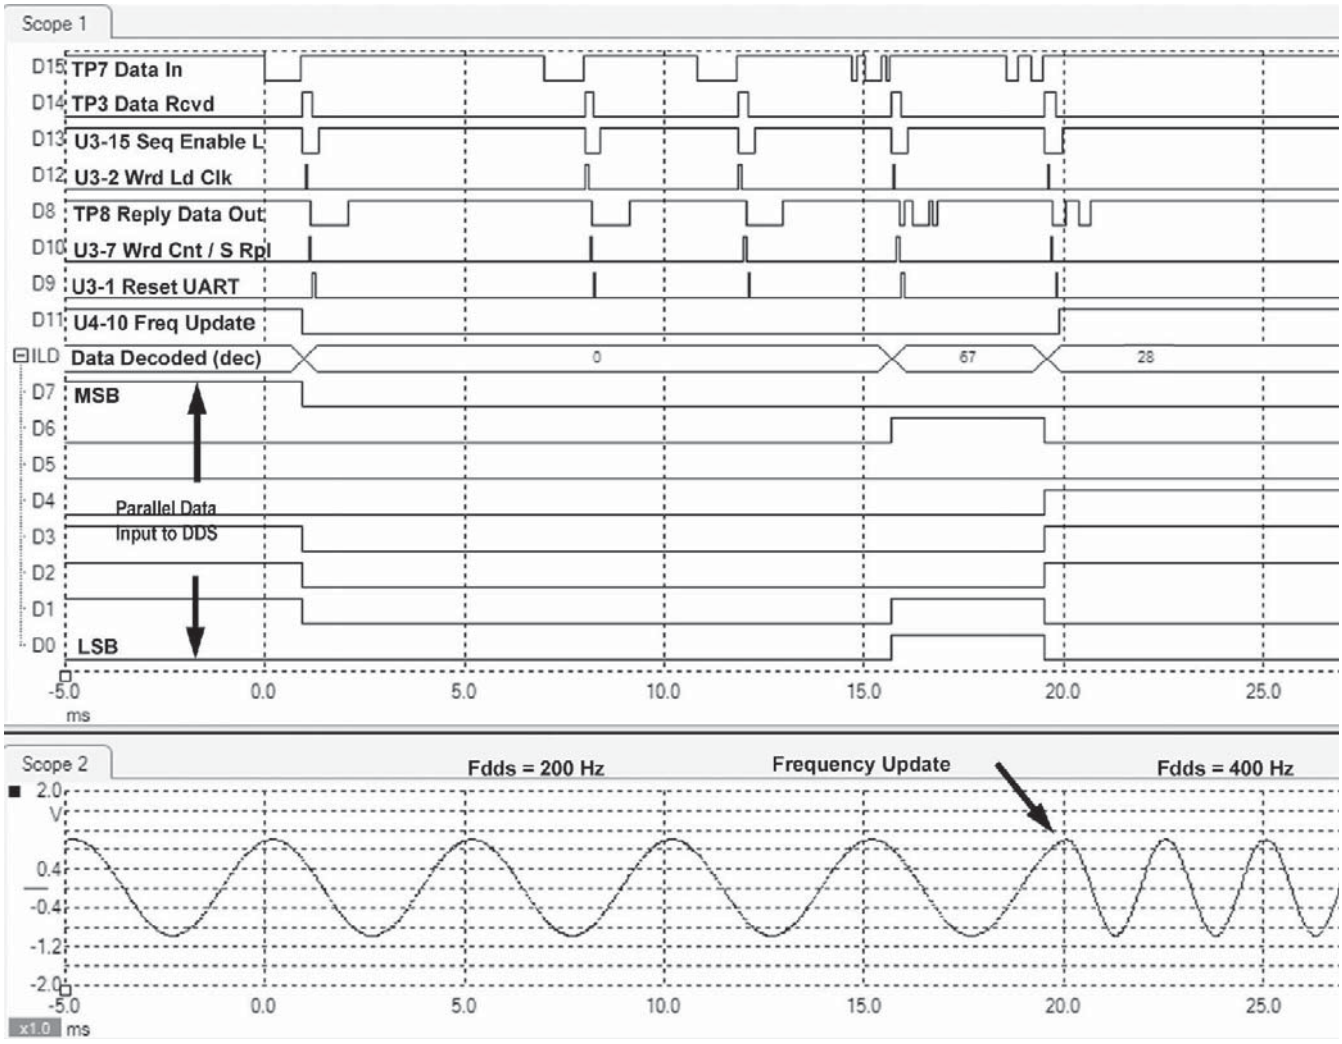

Figure 9 — Mixed signal 'scope capture of the DDS Frequency Update Process as the frequency is switched from 200 to 400 Hz.

the copy of the received data word at the UART receiver output back into the UART transmitter buffer. When Pin 23 is returned low as the sequence counter increments from 3 to 4, the UART sends the copy of the data back to the PC. This is trace D8, TP8 Reply Data Out, in Figure 9.

Sequencer U3 continues counting up from 3 to 5. When the count reaches 5 output from U3-1 is routed via the inverter U4C to Pin 18 of the UART, U2. This pulse resets the UART receiver data buffer and readies it to receive the next serial byte. See trace D9, U3-1 Reset UART on Figure 9.

Next sequencer U3 counts up from 5 to 7 and outputs a strobe that is applied to one input of the AND gate U7A-1. For the first four received bytes nothing happens because the Word Counter U8 has not yet counted up to 5.

The sequencer continues counting up to 8. At a count of 8, U3-9 goes high and outputs a pulse that is looped back to the reset input of U5 at U5-4. Resetting U5 in turn causes Pin 2 of U5 to go high which holds a fixed reset on Pin 15 of U3, halting the sequencer and resetting its count to 0. At this point the logic halts until the PC sends another byte. However the words received counter, U8 is not reset and holds its count at 1.

Upon receipt of the next 3 bytes, the above process is repeated, with the words received counter U8 being incremented each time. When the 5th and final byte is received, the clock input of U8 causes it to be incremented by the "3" count output of U3 and the "5" count output of U8 goes high. This is applied and held at Pin 2 of the AND gate U7A. As the sequencer, U3 is incremented, a count of "7" is reached. The "7" output from U3 is applied to Pin 1 of the AND gate U7A. With both inputs high U7A outputs a high that passes through OR gate U9A and is routed to Pin 10, the reset input of U5B. This resets U5B, which in turn resets the word counter U8.

Remembering, the Q output of U5B was inverted by U4E to generate the Frequency Update waveform for the DDS. This signal shown as D11, U4-10 Frequency Update, on Figure 9, has been held low from the beginning of receipt of the first byte from the PC until returned high when the word count reached 5. When Frequency Update returns high, the DDS switches from the previously programmed frequency to the new frequency.

The lower portion of Figure 9 is an oscilloscope capture of the output of the DDS before and after the frequency update. The frequency switches virtually instantaneously upon the rising edge of the frequency update waveform.

Also of interest is the Data Decoded (dec) trace. It shows the numerical value of each binary byte decoded by the UART as a decimal number. The presentation may be in binary, octal, decimal or hex. I used decimal for compatibility with the decimal presentation used by my control program. In the example shown the first 3 bytes sent are 0, the 4th is 67, and the 5th is 28 with the most significant byte being set first. If these numbers are converted back to their binary values in the tuning word and Equation 2 is solved for the output frequency, the result is

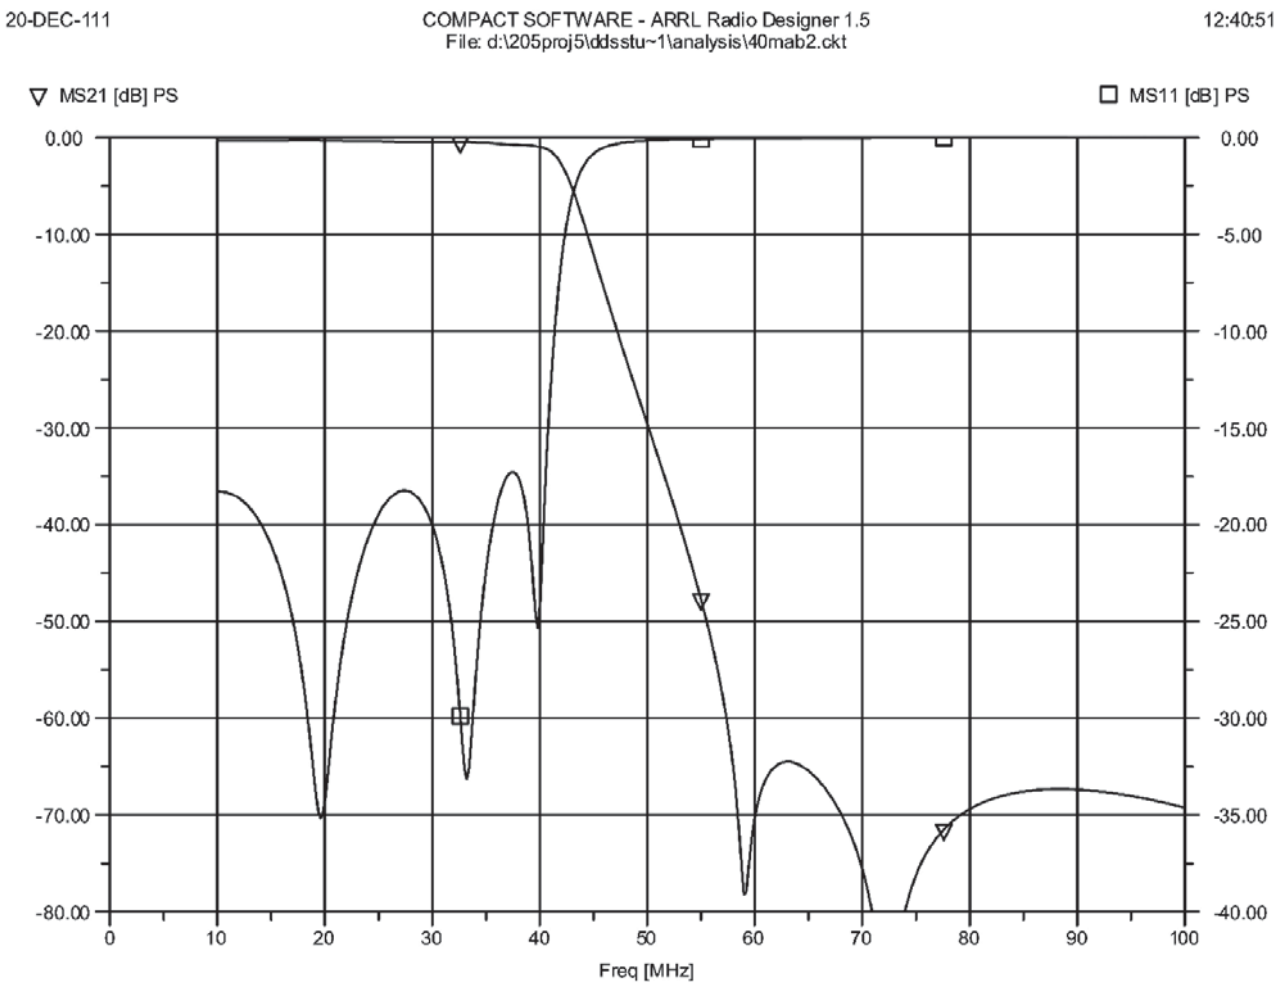

Figure 10 — Simulated response of the 40 MHz low pass filter with actual component values.

Figure 11 — 10 to 100 MHz Phase Locked Loop frequency multiplier.

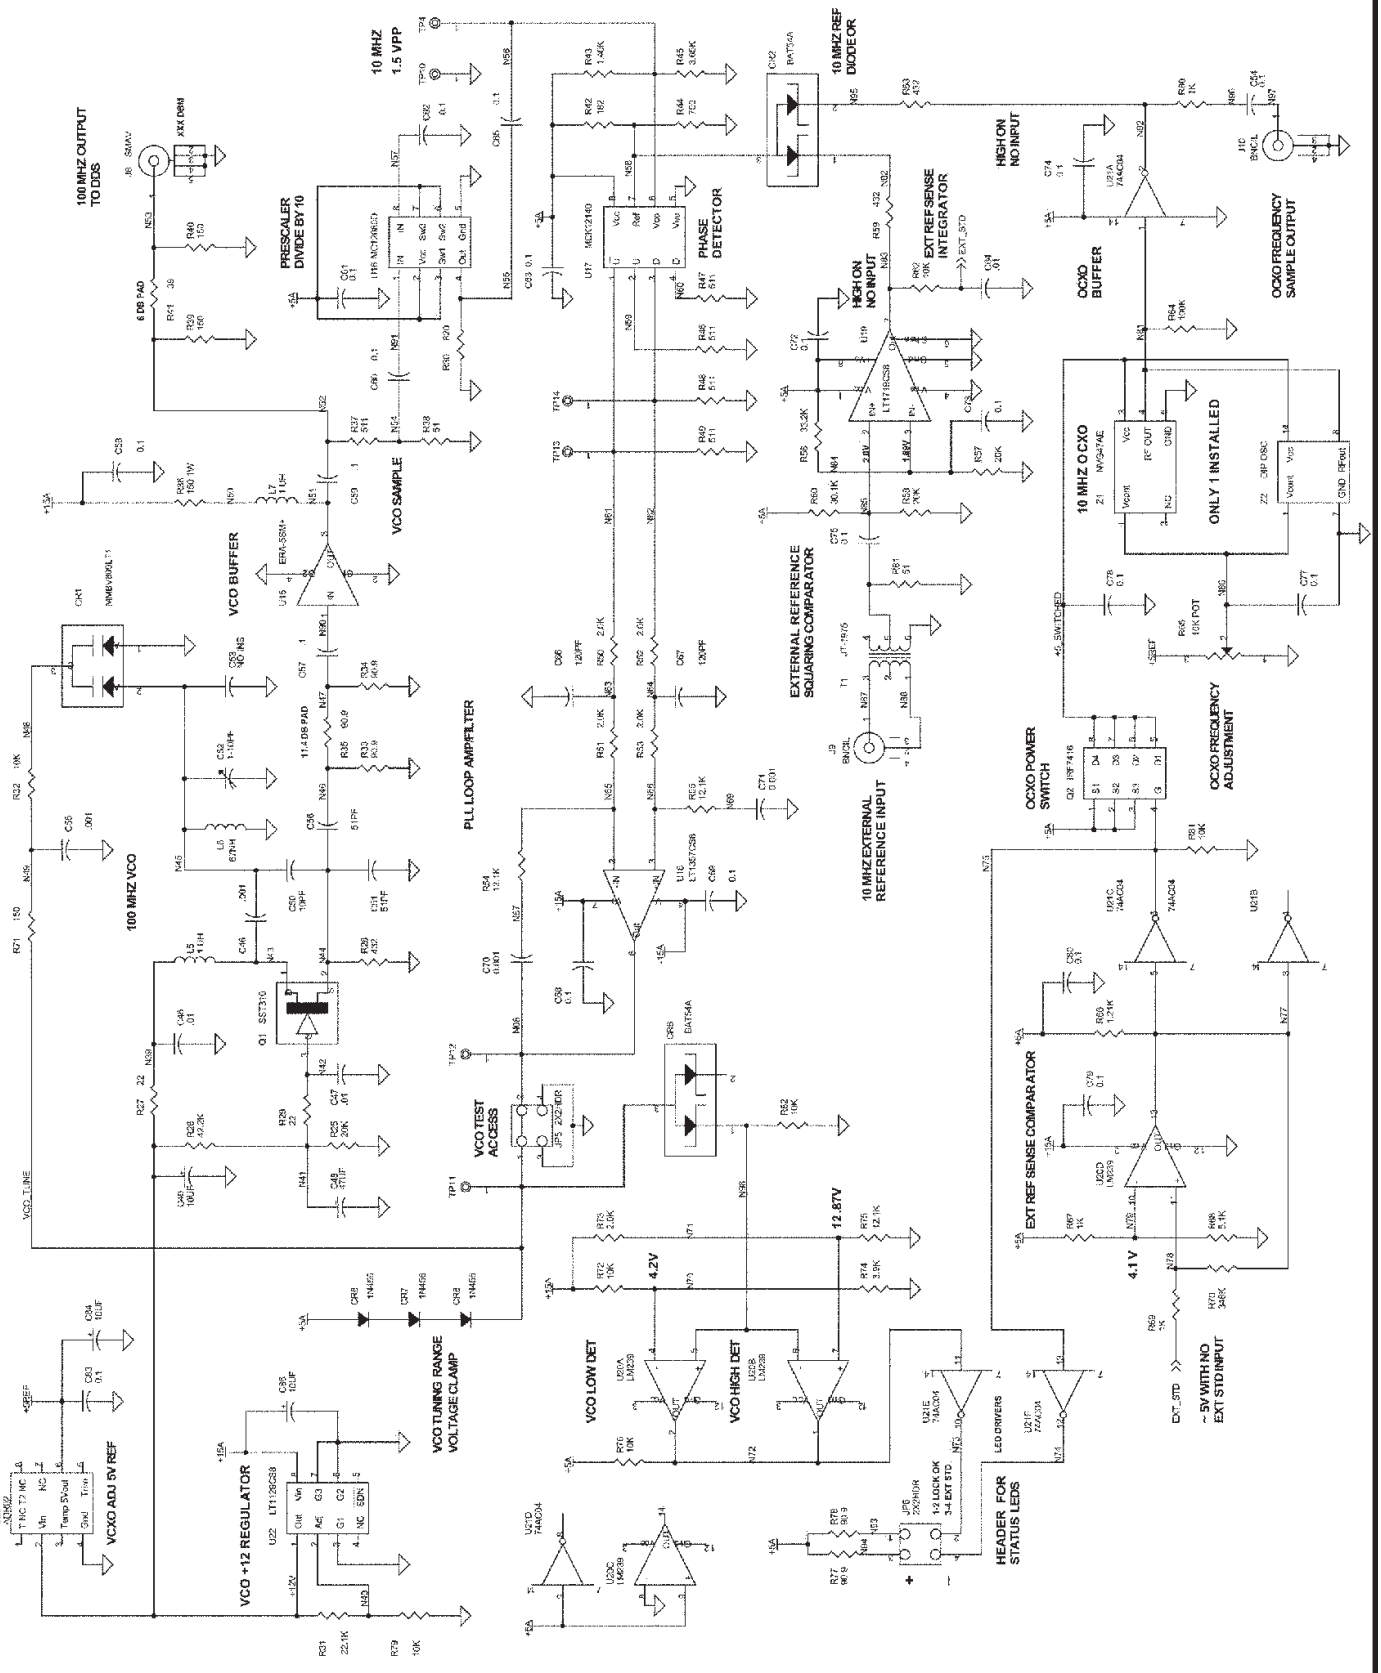

400.003 Hz, as expected.

Other digital circuits in the DDS section include the power on reset, and the loss of synchronization timer. The power on reset uses a standard RC timer and a Schmidt trigger, U4F. The power on reset pulse lasts about 1 second, which was required to ensure a solid reset with one bench supply that came up slowly.

If for some reason a total of 5 bytes are not received the logic can hang up with an incorrect word count. The design has no overall frame synchronization so the "Loss of Synch" timer serves that purpose. U12 counts 6.66 millisecond clocks from U11 until 256 are counted. It then outputs a pulse to the master reset OR gate U9B that applies the required resets. U12 is held reset by U5B except during the interval when word counting is enabled. So it is reset at the end of each frequency update and starts its timing interval from zero at the beginning of the next frequency update. To aid in logic trouble shooting a header JP-3 is provided. It may be jumpered to inhibit U12 so an automatic reset cannot occur while making measurements.

The 100 MHz DDS clock is input to the DDS at J1 an SMA connector. TP1 and TP2 enable measurement of the peak to peak clock voltage. The clock is a sine wave and is ac coupled to Pin 9 of the DDS. For sine wave operation the clock input is biased at 1/2 the supply voltage by R3 and R5. When driven with +10.3 dBm the clock amplitude is about 2.3 V_{pp}. The AD9850 has two current outputs, I_{out} and I_{out1}. They are intended for use with a transformer to convert from balanced to single ended configuration. However they may also be used individually. Due to the wide frequency range I required, no transformer is suitable. So, I used each output individually.

The signal from I_{out} is routed to the 40 MHz low pass filter and amplifier via an RF level adjustment, R6, a dc de-coupling capacitor, and a 3-dB pad. The filter is an elliptic function low pass with 3 loss poles. Spurious alias signals move down in frequency as the DDS programmed frequency moves up.[11] I set the practical upper frequency limit at 40 MHz. With a 100 MHz clock, when the output is 40 MHz, the alias is at 60 MHz. The low pass filter is designed to achieve at least 60 dB attenuation at 60 MHz and above. Figure 10 is a simulated filter response, (from ARRL *Radio Designer*) using the standard component values selected.

Following the filter the signal is sent to the output amplifier U10, a MiniCircuits ERA-5SM+.[12] A monolithic microwave integrated circuit (MMIC), this chip provides a minimum gain of 19 dB and a minimum power output of 16.5 dBm at 1 dB gain compression. From the amplifier the signal is routed to the output connector. The schematic shows provision for an output pad that is not installed. The output is taken from J2, a BNC connector. In the PWB design, provision was also made to use an alternative SMA connector, shown as J6 on the schematic. J6 is not installed.

The audio frequency output is configured similarly to the RF channel. Output from I_{out1} of the DDS is routed to the level set pot, R16, and then via a dc blocking capacitor, C38, to the 100 kHz low pass filter. This is an active filter using ICs U13A and U13B and has a voltage gain of 2. The AF output frequency response extends from below 10 Hz to 100 kHz. Above 100 kHz the response is rolled off. From the filter the output from U13B is routed via a 100 Ω resistor to J4, a BNC connector. The 100 Ω resistor ensures stability with capacitive loads.

As in the case of the RF output, the PWB has provision for alternative use of an SMA output connector. This is J7 on the schematic. It is not installed.

The circuit has numerous test points. JP1 in conjunction with the SMB connector, J3 implement a low VSWR connection for either measuring the DDS output or testing the response of the low pass filter and output amplifier. JP2 and J5 enable a similar test capability for the AF channel. Positioning of the jumper routes signals from the DDS to the SMB connectors or alternatively routes input signals to the filters for test purposes.

A somewhat rudimentary form of reverse polarity protection is provided using shunt connected 3-A Schottky diodes CR3, CR4 and CR5. These will short out any reverse polarity power inputs without compromising the supply voltage accuracy or regulation. My lab gadgets, of which this is one, are powered from purpose-built current-limited supplies. In event of a reverse polarity connection, a cheap three-terminal regulator may die, but the expensive electronics on the board will not!

PLL Section Schematic

The PLL section, Figure 11, generates a stable 100-MHz clock for the DDS. A 100-MHz VCO is locked to either of two 10 MHz references. For precise work an external reference may be supplied. If no external reference is connected, then the internal OCXO is used.

A very conventional PLL design is used. The VCO, Q1, an SST 310,[13] is a surface mount version of the popular U310 junction FET. Doubly regulated and hopefully low

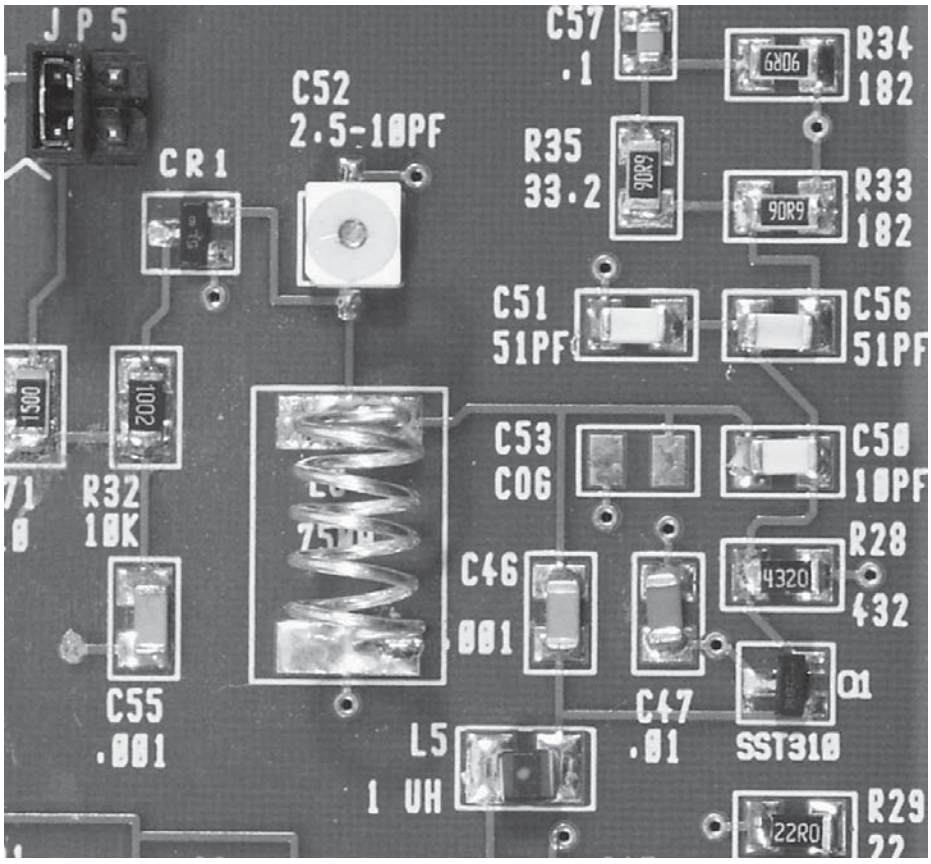

Figure 12 — The VCO Q1, with the inductor, L6, trimmer capacitor, C52, and tuning varactor CR1. L6 is 67 nH though 75 nH was screened on the PWB.

noise, +12 V power for the VCO is supplied by U22, a low noise programmable voltage regulator that operates from the +15 supply. Inductor L6, in conjunction with C50, C51, C52, C56 and varactor CR2, determines the frequency of the VCO. Trimmer, C52, is a Johanson 2322-2G "Cera-Trim", 2.5-10 pF capacitor having a Q of 1000 at 100 MHz. It is used to center the VCO control voltage at 6 to 7 V with the PLL locked. L6 is a hand wound air-core 67 nH inductor. This coil, Figure 12 has 3 to 4 times the Q of a surface mount part and was used to improve the oscillator noise performance. It is simply soldered to two surface mount pads spaced 0.400 inch apart on the PWB. The tuning range of the VCO is 93 to 107 MHz with a control voltage range of 2.5 to 13 V. Use of this relatively wide range L-C VCO instead of a 100-MHz VCXO in this single frequency application may seem a bit odd. Actually, it turned out that I could not easily obtain a suitable VCXO at a good price, so I used this design. It had worked well in other applications and the PLL would clean up any noise issues.

The VCO output is isolated by an 11 dB pad, amplified by U15, and routed to the output connector J8, via a 6 dB pad. The pads were adjusted to obtain a power output of 10.3 dBm at J8. When connected to the DDS clock input, J1, a clock voltage of 2.3 V_{pp} is applied to the DDS.

A 11X voltage divider, R37 and R38, provides a sample of the VCO frequency to the divide-by-10 prescaler, U16, an MC12080D. This is an emitter-coupled part and is operated from the +5 V supply. Since the signals are continuous, and not logic, capacitive coupling may be used on the prescaler input, and also on the output to the phase detector U17, a MCK12140. The 10 MHz output from the prescaler at TP4 is about 1.5 V_{pp} and is routed to the VCO input of the phase detector, U17. A 10 MHz reference signal from the selected frequency standard is applied to the reference input of the phase detector from a diode OR connection formed by CR2.

Reliable phase locking is assured by use of the MCK12140, a phase-frequency detector. When the loop is not locked, U17 acts as a frequency discriminator and outputs an error voltage that drives the VCO toward phase lock. R46, R47, R48, and R49 are the output load resistors for U17. The outputs are a series of pulses that vary in width as a function of the phase error. It is necessary to keep these very high-speed pulses out of the loop amplifier, U18. R50, C66, and R52 with C67 filter them, before they are routed to U18.

The loop filter uses a very high speed, low noise, operational amplifier, a Linear Technology LT1357CS8.[14] This device has a 25 MHz gain-bandwidth product and good phase margin, making it ideal for use in fast phase lock loops. Due to my use of a relatively wide bandwidth, and therefore-noisy VCO, a fast loop is required to clean it up. The loop natural frequency, damping factor, and bandwidth are set by C70, and C71 in conjunction with R50, R51, R54, and R52, R53, and R55. Since our focus is on the DDS accuracy, I am avoiding the temptation to branch off into the details of the PLL design. Table 2 summarizes the parameters of the loop filter. These were verified by measurement on the final hardware.

For those interested in PLL design and test techniques there are no better references than the books by Floyd M. Gardner.[15]

From the output of the loop amplifier, U18, the VCO control voltage is routed via the test point jumper JP5 to the frequency control varactor CR1. JP5 is included for test convenience. It enables the loop to be opened for measurement of the VCO gain constant and other tests. The 150- Ω resistor, R71, is inserted in the VCO control line to prevent U18 from seeing a stiff capacitive load that could cause oscillations. While U18 supposedly can operate into capacitive loads, adding the resistor is good design practice so long as excessive phase lag is not introduced.

Diodes CR6, CR7, and CR8 limit the VCO control voltage from going below about 3 V. At voltages below 2.5 V the VCO power drops as CR1 starts to conduct. Reduction in power caused the prescaler, U16, to miss-count and the loop would not always lock. That's a major oops! With the diode clamp, lock up is fast and reliable.

Diode CR8 conducts a sample of the VCO control voltage to U20A and U20B, LM239 comparators that comprise the "PLL Lock" Detector. This detector causes the front panel lock indicator to be illuminated as long as the VCO control voltage is in an acceptable range, 4.2 to 12.8 V. Note: CR-8 was left over from a previous design iteration to prevent exceeding the allowable negative

Table 2
100 MHz Phase Lock Loop Design Parameters

Parameter	*Specified*	*Measured*
Loop Natural Frequency — Radians/Second	120,000	125,600
Loop Damping Factor	0.7	0.72
Frequency for Loop Response – 3 dB — Radians/Second	245,000	439,000

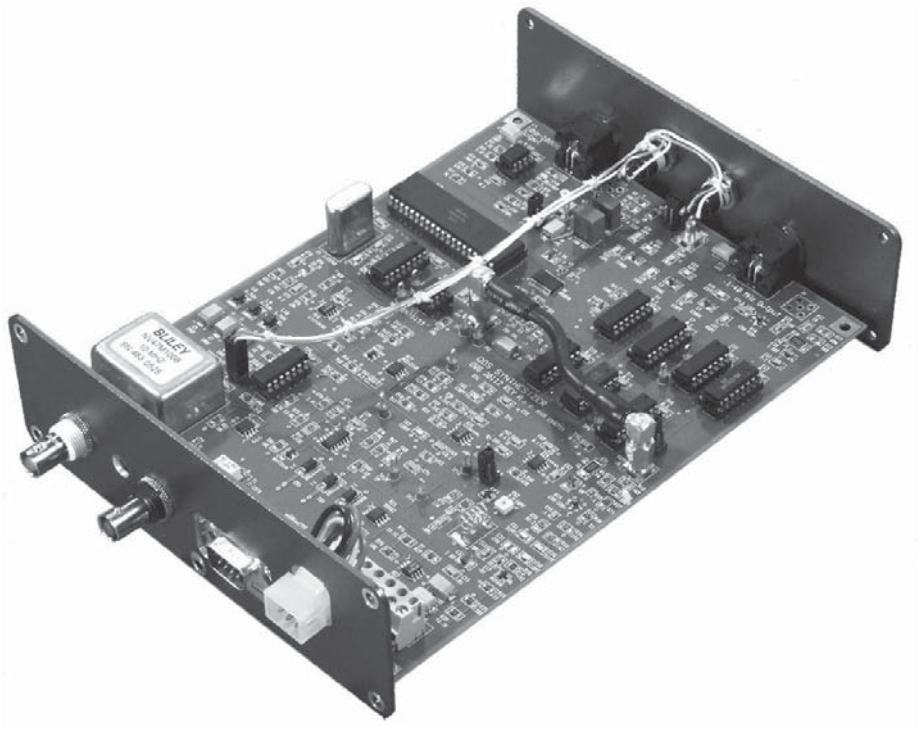

Figure 13 — The Printed Wiring Board with front and rear panels. To install the board into the enclosure the rear panel is removed and the power terminal block unplugged to separate the board from the rear panel.

input range of U20A and U20B. With the addition of CR6, CR7, and CR8, it is no longer required.

A wired OR connection routes the outputs of U20A and U20B to the LED driver U21E, a 74AC04 inverter. The front panel LED indicators are connected to the PWB using JP6, a 4 terminal header. The indicators I used have internal resistors and were intended to operate directly from 3.6 V. So, R77 and R78 were included to drop the +5 supply down to 3.6 V at 15 mA.

If an external 10-MHz reference is used, the internal Bliley OCXO[16] must be shut down and the reference input of the phase detector disconnected from the internal reference and connected to the external reference. After considering toggle switches, and 50-Ω PIN diode switches, I came up with the automatic scheme used here.

Comparator U19, a Linear Technology LT1719CS8, squares up the sine wave input from the external reference and also translates the resulting square wave to logic levels acceptable to the phase detector, U17.

However, it also forms a part of the reference switch function. The inputs of U19 are intentionally biased so that absent an external reference input, the output of U19 at Pin 7 always assumes a logic high. It goes to the +5 rail and stays there. The signal at Pin 5 is integrated by the RC filter R62 and C64 and also goes to 5 V. This signal, labeled "ext_std", on the schematic is routed to the input of the LM239 comparator, U20D, causing the output of U20D at Pin 13, to go high. This signal is sped up and inverted by U21C, a 74HC04 inverter, and then applied to gate of Q2. Q2, an International Rectifier IRF7416 P-Channel power FET, is switched on and supplies +5 V power to the 10 MHz OCXO. The IRF7416 is a remarkable little part. It is packaged in a SOIC8 surface mount package, but can handle several Amperes of current with a very low ON resistance. Current drawn by the OCXO when the oven is cold is less than700 mA. So, Q2 is operated well within its ratings.

Two oscillators are shown on the schematic, though only one can be installed. I did this so that the PWB layout would accommodate installation of either the Bliley OCXO I used or one of the more commonly available DIP-14 packaged oscillators. (The DIP oscillator footprint is hidden under the larger OCXO).

A pot, R65, powered from U23, a very stable Analog Devices ADR02 voltage reference, provides a variable voltage for adjustment of the OCXO frequency.

From the OCXO the 10 MHz signal is inverted by U21A and routed through R63, and the diode OR, CR2, to the reference input of the phase detector, U17. This diode OR enables a seamless and inexpensive way to switch between the internal OCXO and an external reference. The reference input to the phase detector, U17 Pin 7, is biased at the high ECL logic level by the voltage divider R42 and R44. Current pulled toward ground through either diode of CR2 switches the reference input to a low logic level at a 10-MHz rate. For this to work as an OR the unused input to CR2 at either Pin 1 or Pin 2 must go to +5 when that input is inactive. As discussed above the output of U19 is at +5 without an external reference input.

If an external reference is present at J9 then U19 switches at a 10 MHz rate. This causes the voltage across the integrator capacitor C64 to drop toward ground. When the duty factor is sufficient to drop the voltage to below 4.1 V, U20D switches and its output goes high. The external reference indicator is illuminated and Q2 is turned off. This removes the supply voltage from the OCXO. To ensure the OCXO output goes to ground it is pulled down by R64. This forces a low on the input of the inverter, U21A, forcing its output to Pin 2 of CR2 to +5 V. This results in the input to the phase detector to come only from U19 via CR2 Pin 1. By this means the appropriate reference source

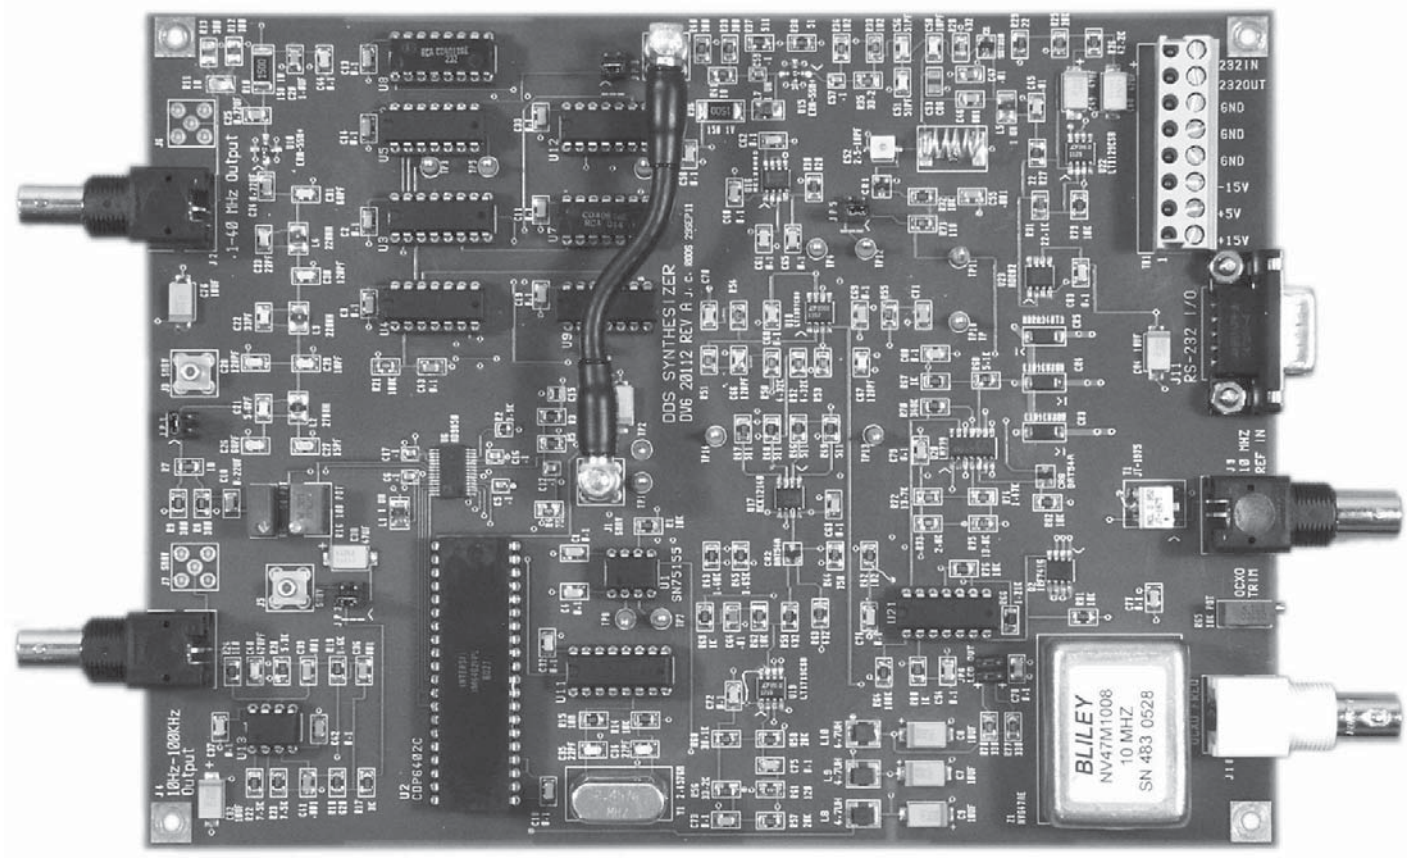

Figure 14 — Printed Wiring Board layout. The DDS circuits are located to the left of the SMA jumper cable with the 100 MHz Phase Locked Loop and internal OCXO to the right.

is automatically selected. Usable external reference input levels range from –10 to +16 dBm. Levels greatly above +16 dBm may damage the inputs of U19.

J9, the external reference input is isolated from the PWB ground by transformer T1. Breaking a potential ground loop avoids possible low-level hum modulation of the reference PLL. A sample of the OCXO frequency is routed to J10 so the oscillator may be calibrated. The level at J10 is approximately – 6 dBm with a 50-Ω load.

Construction

Figure 1 shows the DDS packaged in an extruded aluminum instrument case. A product of Box Engineering Ltd, these cases are offered in a variety of sizes and a selection of anodized aluminum colors, complete with front and rear panels and hardware. Matching color elastomer bezels are offered for each case to make a nicely finished package. I used Box Engineering part number B4-220BL[17]. It measures roughly 6.5 by 8.7 inches and the extrusion features internal longitudinal mounting slots to support a PWB. Figure 13 is an overall view of the DDS PWB mounted between the enclosure front and rear panels. It shows the assembly without the enclosure.

To remove the board from the cabinet, the rear panel is removed from the box by removing the nuts from the two BNC connectors and five panel screws. Then the power terminal block is unplugged from the header on the board. The Molex power connector, power leads and terminal block remain with the rear panel. After removing the five front panel retaining screws, the board can then be removed or inserted from the front into the correct mounting slots in the interior of the extruded case. After reinstalling the board, the plug-on power terminal block is installed and the rear panel is reattached with the mounting hardware. The elastomer bezels simply stretch into place after the panels are installed.

To fit the box correctly the printed wiring board measures 6.3 x 8.5 inches. A four-layer board was used. Layer 2 is a solid ground plane that is vital to suppressing crosstalk between circuits, ensuring low spurious, and providing short ground returns for the VHF signals. Sierra Circuits[18] fabricated the board for me using their "no-touch" prototype process.

Figure 14 shows the layout of the printed wiring board. While a single board is used, the DDS and PLL circuits are physically and electrically separated. The AD9850 DDS, associated digital circuits, and the output filters and amplifiers are all grouped to the left of the SMA cable that connects the PLL to the DDS clock input. The DDS is located just above the 40 Pin UART and is dwarfed by it. The DDS and UART were oriented to enable a direct connection of the 8 parallel data lines from one to the other. At the lower left of the board, the AF output filter, amplifier, and the output connector J4 are located. To the left of the AD9850 are the two level set pots. From the RF level set pot, the 40 MHz low pass filter is routed vertically to the RF output amplifier, U10. Surface mount inductors and stable COG ceramic capacitors were used to build the filter. Generous spacing between the inductors avoided any undesired coupling. From the RF amplifier, the signal is routed to J2, the RF output connector.

To avoid undesired coupling between sensitive circuits, the clock crystal Y1 and the logic clock oscillator-divider U11, are located at the bottom center of the PWB. It places this potential spurious source far from the AF and RF output amplifiers. Similarly the 100-MHz VCO is at the top right of the board, as far possible from Y1 and U11. This layout and the use of RF filter chokes, L8, L9, and L10, to isolate the dc supply for the PLL from the DDS logic and U11, resulted in no detectable spurs attributable to the 2.4576 MHz clock or the clock divider.

Inductor L6, the air core coil, is located just below the VCO oscillator circuit at the upper edge of the board. The VCO buffer amplifier, U15 and associated components are above the coil and deployed along the upper edge of the board, with the 100-MHz clock output connector, J6, centered on the upper edge of the PWB. The short SMA jumper cable connects J6, to the DDS clock input, J1.

To minimize the length of the 100-MHz VCO sample run to the prescaler, U16 is placed near the VCO buffer amplifier output. After division by 10, the 10 MHz prescaler output is routed, microstrip fashion, on layer 3, just below the ground plane to the input of the phase detector, U17 Pin 6. Placing this signal on the first layer below the ground plane reduced its ability to radiate 10 MHz spurious. Similarly, the lines from the on board OCXO, and the external reference input J9 are located at the lower right, away from the VCO, and routed on layer 3. As a result no 10 MHz spurious is detectable on the synthesizer output.

Artwork for the PWB was prepared using the now discontinued Ivex PWB board layout software. It did not have a canned foot print for the AD9850s tiny 28 pin Shrink Small Outline Package, 28SSOP. However, like many PWB design programs, a module footprint compiler is included, so I designed the footprint from the 28SSOP dimensions. I did make the pads a bit longer than the data sheet suggested to ensure a point to contact with the soldering iron.

While I have worked with 2 and 3 lead surface mount (SMT) parts for many years, this was my first experience with trying to install such a fine pitch multi-lead component myself. In spite of my worries, it was quite easy. Sierra Circuits left a nice layer of solder on the pads. Installation was simple; the pads and the leads on the AD9850 were treated with flux. Then, holding the chip in place, the solder on one corner pad was reflowed to secure that lead. This was repeated for the three remaining corner leads, and

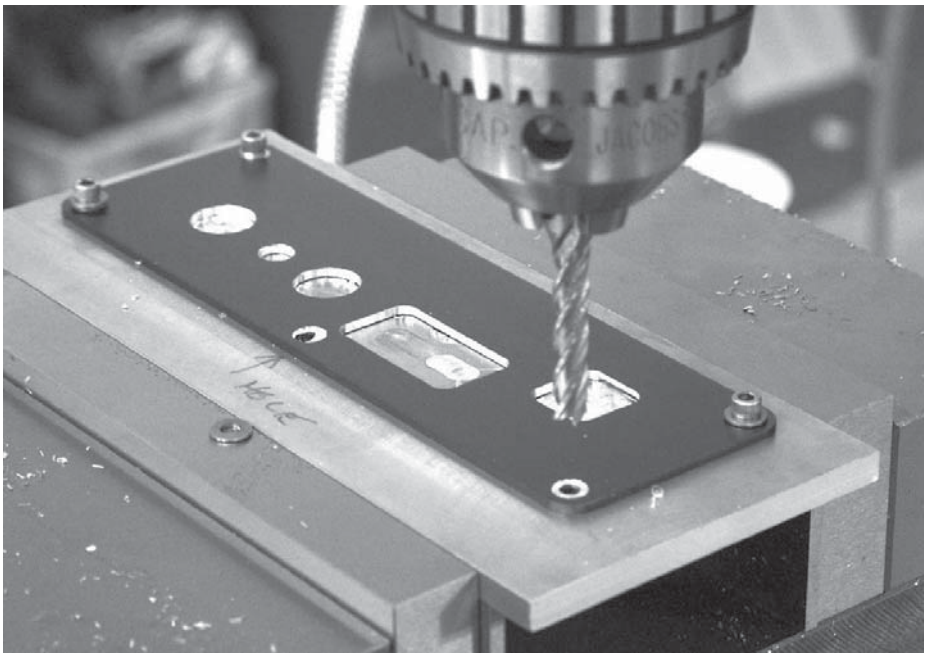

Figure 15 — The rear panel mounted on the tooling plate. The anodized finish is nonconductive so each mounting screw hole was spot faced to ensure a good RF ground.

then the rest were re-flowed. A Weller iron with a .010 diameter tip was used. The whole job took about 10 minutes. No solder was added. For those who have not yet jumped into the surface mount arena, it is far easier to work with these parts, particularly the two leaded resistors, capacitors, and diodes, than their leaded counterparts. This is especially so when removing parts from a board. Two small irons will lift a two-leaded component in about 3 seconds and there are no holes that have to be cleaned out before the new component can be installed.

Use of the Box Engineering instrument cabinet entailed a few mechanical challenges because dimensions had to be coordinated between the board, the mounting slot locations and width in the extrusion and the front and rear panel holes. Fortunately, Box Engineering provides excellent mechanical drawings for their product with dimensions for every physical feature. Using this data the PWB was made 0.020 inches narrower than the depth of the mounting slots in the extrusion, and 0.050 shorter than the long dimension of the extrusion to allow for tolerances. Using a couple of spacer nuts on the connectors at one end, this worked out perfectly. Another issue was placing the holes in the front and rear panels relative to the connectors and other components on the board. This had to be exact or some filing was inevitable. The panels are located relative to the extrusion by five pre-drilled and tapped mounting holes.

Using a spreadsheet to do the math, all of the holes in the panels were dimensioned in X-Y coordinates relative to a single point; the existing upper left panel mounting hole. One of my favorite ham radio tools is a small milling machine with X-Y digital readout. Using this, a tooling plate was made having the same hole pattern as the supplied panels. The holes were tapped for 6-32 machine screws to secure the panels to the tooling plate. This arrangement supported the thin panels and made it possible to install and remove a panel without re-zeroing the X-Y readout. Figure 15 shows the setup on the mill. After machining the entire assembly went together without need of the file.

Software

As explained above, my FMT DDS is an all hardware device controlled by software on an external processor, in my case a PC. The hardware does not limit the capability of the DDS. It will do anything the AD9850 is capable of subject to the limitation of the 9600 bit/second serial data rate, and the 40 MHz maximum output frequency.

To program the DDS, the software must accept a *requested* output frequency, perform the math described above, and then send the 40 bit control word to the DDS as five 8 bit serial bytes. After each byte the software should await the return byte echoed from the DDS prior to sending the next byte. As explained on the AD9850 data sheet the data is sent most significant byte first, i.e. byte 4, 3, 2, 1, and 0. Unless phase control is required byte 4 is always 0.

For the FMT application, time is of the essence, so a simple and fast operator interface is required. Here is how my control software functions.

Upon startup, operating instructions are displayed and then the operator is queried for a base frequency and a frequency increment. The DDS goes to the base frequency. I used the compensated formula Equation 7 to make the round off errors symmetrical about the requested frequencies. With the 100 MHz clock, the actual DDS frequency is never more than +/- 11.6 MHz from the requested frequency.

The requested frequency, predicted frequency, and the round off error in mHz are displayed. After the initial entry all functions are single keystrokes, using the function keys on the PC. Here are the available controls:

F1 — Gracefully exits the program.

F2 — Queries for a new base frequency and frequency increment. Used to move to the next FMT base frequency.

F4 — Writes the time, requested frequency, predicted DDS frequency and round off error to disk file C:\freqdata\ddsfreq.txt. F4 creates the directory and file if they do not exist.

F5 — Same as F4 but prints the data on LPT1.

F7 — Steps the requested frequency down by the frequency increment.

F8 — Steps the requested frequency down by 1/10 the frequency increment.

F9 — Steps the requested frequency up by 1/10 the frequency increment.

F10 — Steps the requested frequency up by the frequency increment.

After each frequency step the requested frequency, predicted frequency, and the round off error are displayed. This data and the time may be stored or printed whenever desired using F4 or F5. Table 3 is a copy of a frequency data file stored by the software using the F4 key. The file is opened for "append" so new data is simply added to the end of a file if it exists. Thus, it is not necessary to rename a new file for each test frequency. The time tag identifies the DDS frequency associated with each FMT test frequency.

When using the DDS for a FMT this scheme makes it easy to quickly adjust the RF reference frequency to obtain a desired audio beat frequency from the receiver. Typically I make the frequency increment 100 Hz, so the fine steps are 10 Hz. That way I can hit my desired 500 Hz beat tone within 10 Hz.

Not being a software guy, the program is written in the very useful but now defunct Microsoft *Quick Basic 4.5*. Not surprisingly, the compiled executable has run correctly under *MS-DOS* on every PC I have tested. This included 7 different processors, from an ancient 20 MHz Compaq Desk Pro to Dell and Compaq desktops and laptops of various vintages. There are a couple of ways to use the software. One is to put it on a floppy disk or thumb drive with *MS-DOS* and boot up and run under *MS-DOS*. It also has run well under *Windows XP*. *Windows 7* just seems to ignore it.

The software source code and the executable will be put up on the *QEX* files website at **www.arrl.org/qexfiles** as part of *ROOS.ZIP* mentioned earlier. It will also be available from me by e-mail.

I will also include a couple of useful utilities. One simply emulates the DDS by

Table 3
DDS Frequency Data File with Time, Requested Frequency, Predicted DDS Output Frequency and the Round Off Error

TIME: 21:05:32		
Req Frequency	*DDS Frequency*	*Error mHz*
3,500,900.0000	3,500,899.9985	–1.5289
TIME: 21:07:39		
Req Frequency	*DDS Frequency*	*Error mHz*
3,500,000.0000	3,499,999.9916	–8.3819
TIME: 21:07:44		
Req Frequency	*DDS Frequency*	*Error mHz*
3,501,000.0000	3,500,999.9992	–0.7674
TIME: 21:11:04		
Req Frequency	*DDS Frequency*	*Error mHz*
4,000,200.0000	4,000,200.0052	5.2482

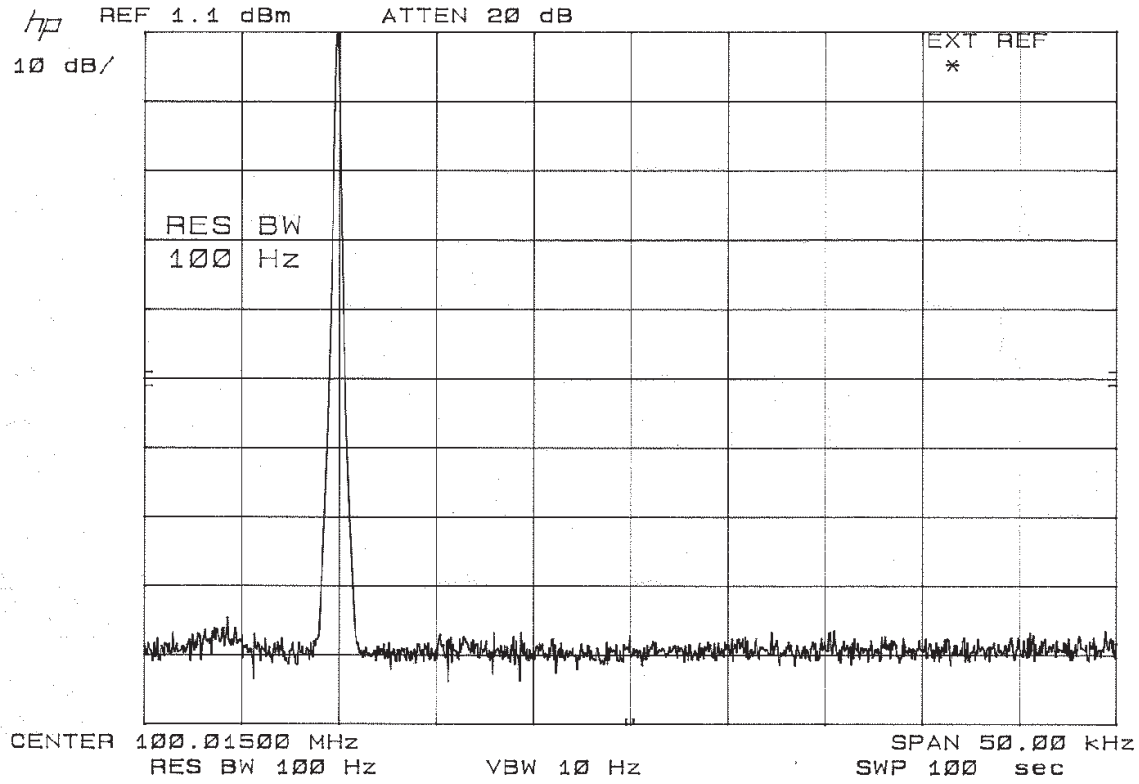

Figure 16 — Output spectrum of the 100 MHz PLL clock multiplier. When corrected for resolution bandwidth the phase noise is about –107 dBc/Hz.

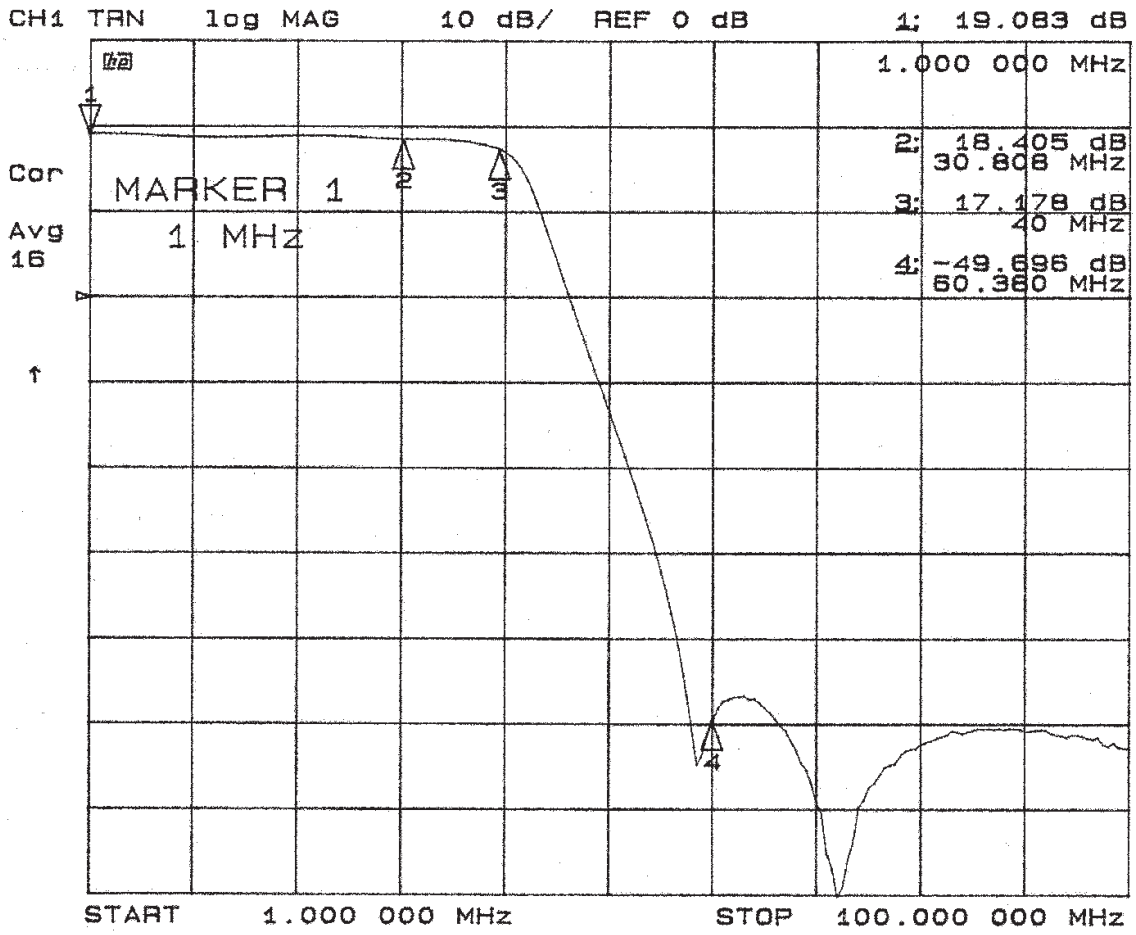

Figure 17 — Measured frequency response of the 40 MHz low pass filter and output amplifier. The filter attenuation is 68.7 dB at 60 MHz.

echoing back the control bytes. It also displays the sent and received decimal value of each byte. Hosted on a second PC, I used it to develop the software before the DDS hardware was completed. Another utility will accept typed in decimal value of bytes 3, 2, 1 and 0, and compute the binary values and the predicted DDS frequency (assuming a 100 MHz clock). This was very useful in finding an error in the control software. The core code in the program was written more than 20 years ago to control an early Standard Telecommunications DDS in a lab environment. Any skilled programmer can improve on it!

Test and Evaluation

I assembled the board complete, except for the rather costly AD9850. Testing started by getting the 100 MHz PLL going. This required first measuring the gain constant of the VCO so the PLL loop filter could be designed. Measurements with a spectrum analyzer revealed that the selected loop natural frequency was too low to completely strip the phase noise of the VCO near the carrier. The natural frequency was increased to 120,000 radians per second and that did the trick. Figure 16 shows the final spectrum plot using a 100 Hz resolution bandwidth. The loop hang-up during acquisition, mentioned previously, was found and solved by adding the diode clamp. The pads in the VCO buffer chain were adjusted to obtain a 100-MHz output of 10.3 dBm at J6.

Next the analog signal paths from the DDS to the RF and AF outputs were checked. I did have to remove the original 3 dB pad at the RF output, J2, to achieve a nominal 10-dBm output level. The performance of the 40 MHz low pass filter and the ERA-5SW+ MMIC amplifier was measured with a network analyzer. Figure 17 is a plot of the frequency response from 1 to 100 MHz. Including the loss of the 3 dB pad, and the filter, the gain is 19 dB at 1 MHz and decreases to 17.1 dB at 40 MHz. The 40 MHz low-pass filter performance compared favorably with the computer model, Figure 10. The gain at 60 MHz was reduced to –49.7 dB which when subtracted from the 19 dB gain at 1 MHz, indicated a filter attenuation of 68.7 dB.

Next the DDS control sequencer logic was tested and no difficulties found, other than the slow power on reset discovered later. With that working it was time to install the DDS chip. Initially things looked good until the frequency counter was connected. There were frequency errors, but only at some frequencies, so it appeared we had a stuck bit, but there were no wiring errors or open leads to the DDS chip. Using the software utility that calculated the predicted DDS frequency from the decimal value of the 4 frequency control bytes, sent to and echoed, from the hardware, it was quickly discovered that my frequency control program *always* set the LSB of each byte to zero. A typo that changed a variable name in the code was corrected. Then things worked as expected.

The variation in the AF and RF output levels as a function of frequency was measured with the results shown in Figures 18 and 19. The AF level was extremely flat up to about 90 kHz where a slight peak occurred due to imprecise capacitors in the active low pass filter. In Figure 19 the RF output varies from 12.7 dBm at 100 kHz to 7.5 dBm at 40 MHz. Output at 30 MHz is 9.8 dBm. As explained by the AD9850 data sheet and on the Analog Devices web site, a DDS output is not flat with frequency and falls off with increasing frequency. The slight roll off in the 40 MHz Filter and Output Amplifier, Figure 17, also contributes to the reduction in output with increasing frequency.

Spurious output frequencies are always an issue with a DDS, particularly if the intended application is a transmitter VFO or receiver LO. My applications: a precision RF reference for the FMTs, and general purpose lab generator are more forgiving. For this the performance of the AD9850 with its 10-bit DAC is adequate. However I did make a lot

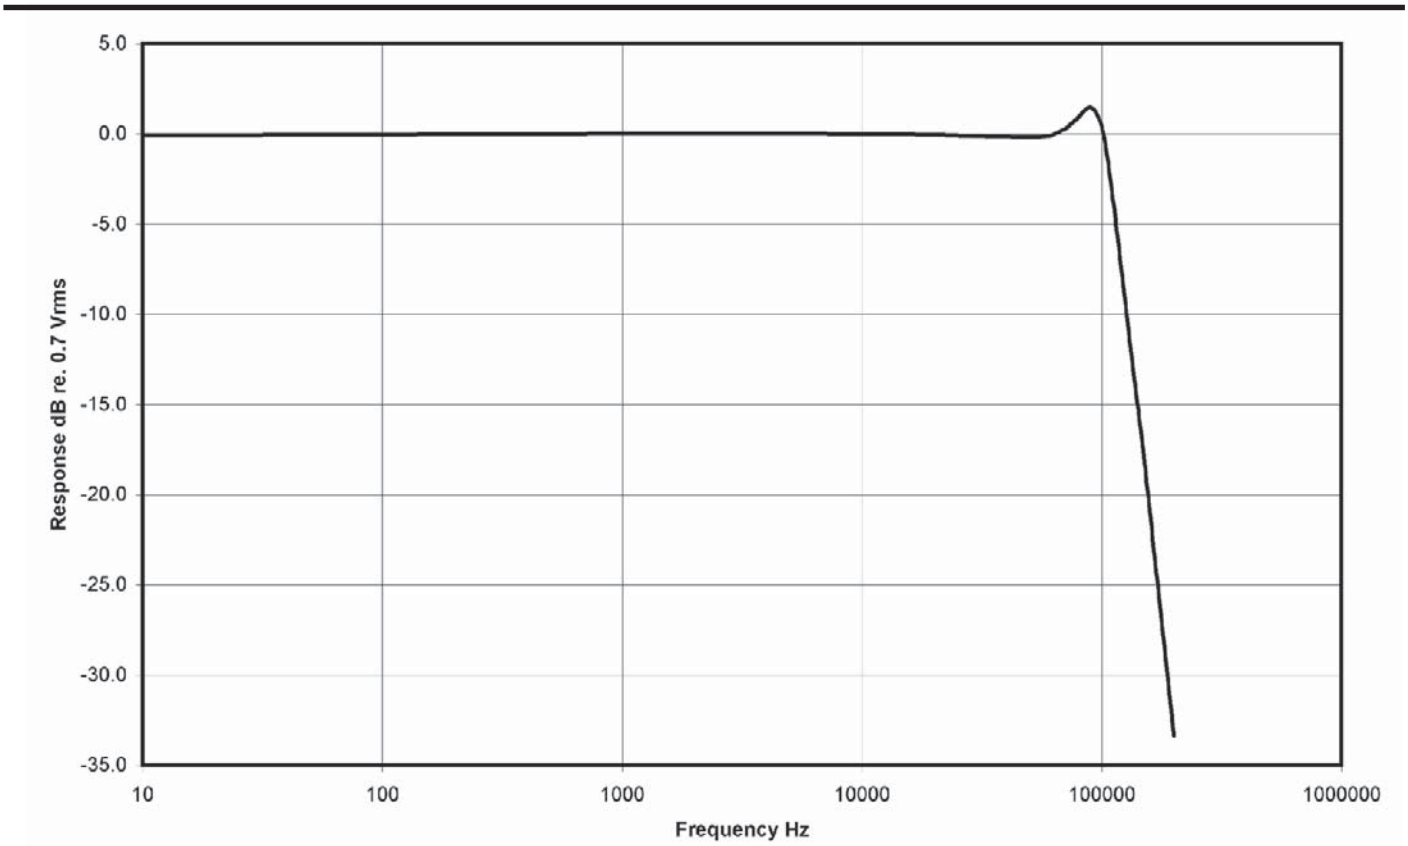

Figure 18 — Audio frequency output from 10 Hz to 100 kHz. The plot is referenced to a level of 0.7 V_{rms} at 1 kHz.

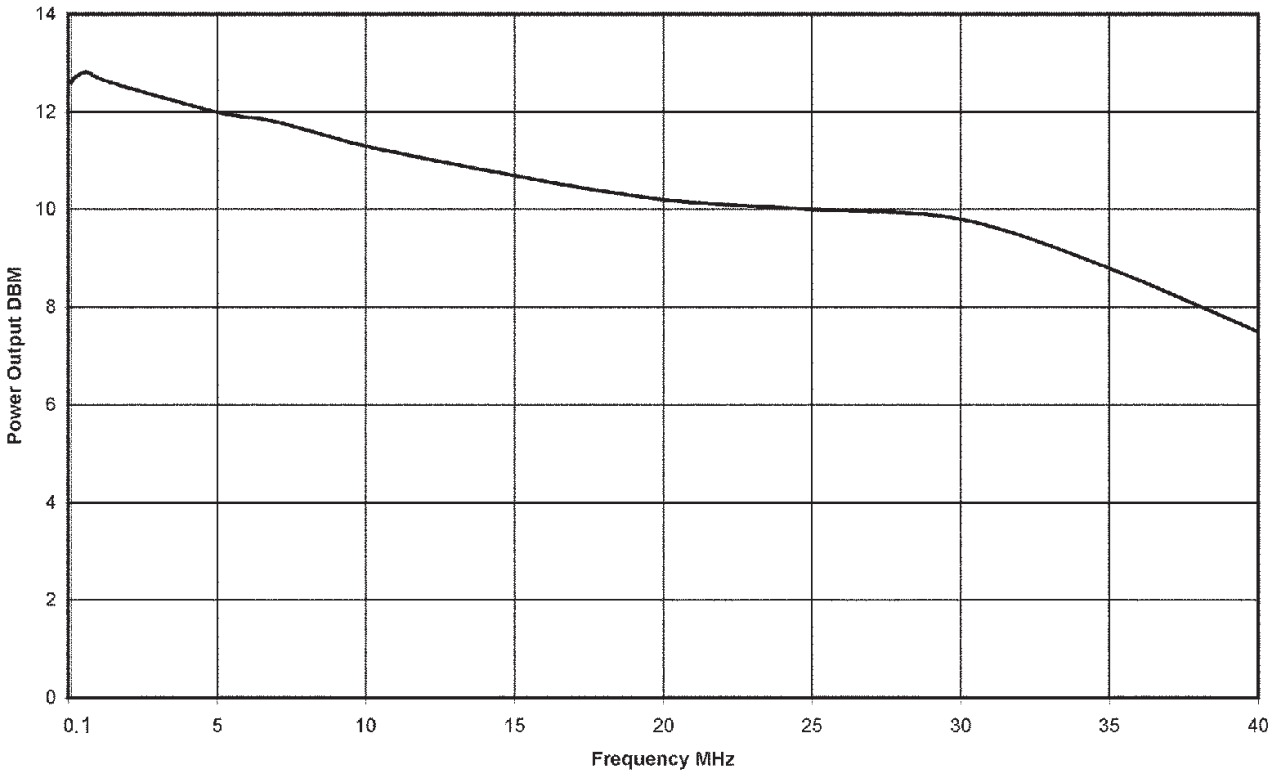

Figure 19 — The DDS RF power output from 100 kHz to 40 MHz. Power decreases as the frequency increases due in part to the increased loss of the 40 MHz filter and amplifier circuit, but the majority of the variation is due to the normal decrease in DDS output with frequency.

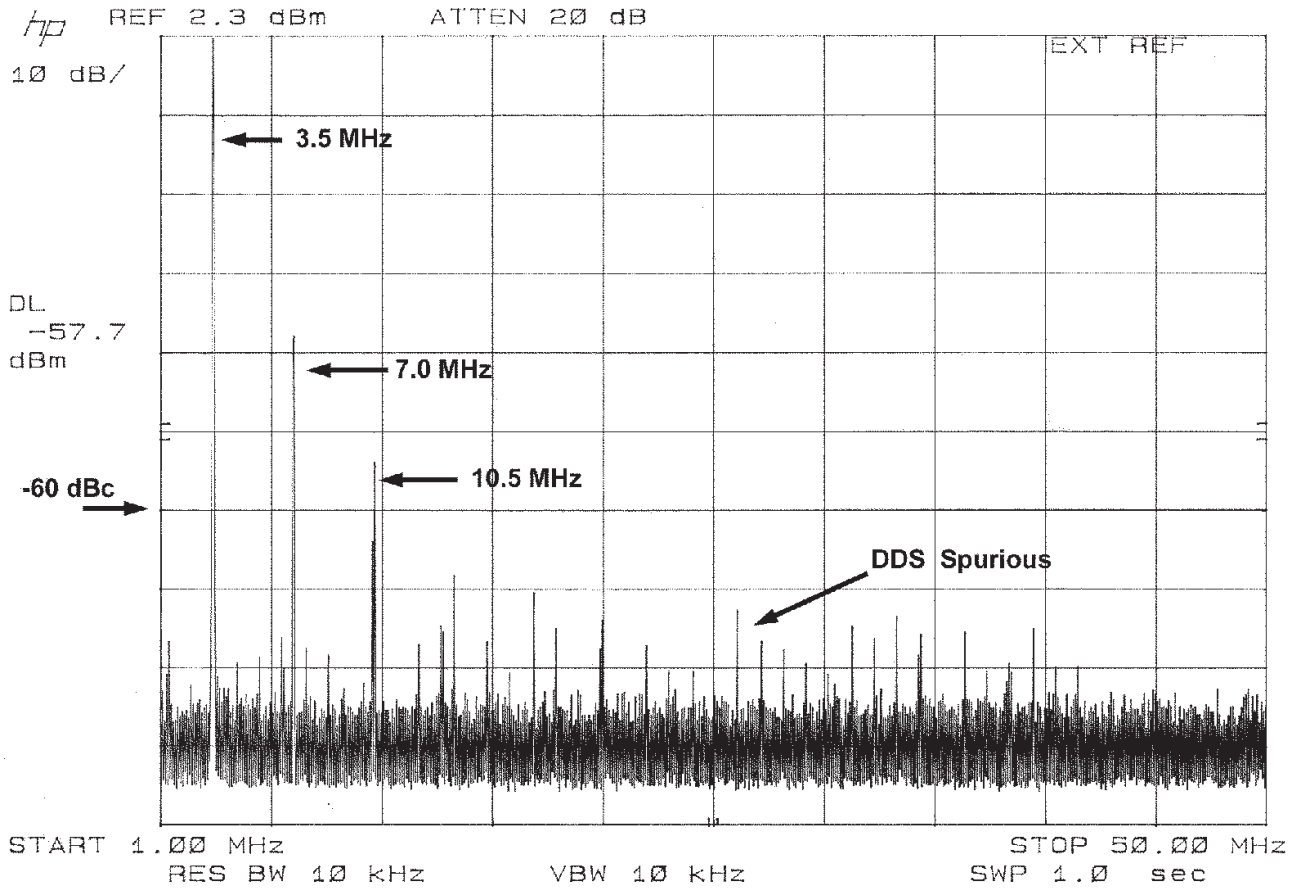

Figure 20 — DDS RF spectrum with the output set to 3.5 MHz. In addition to the expected 2nd and 3rd harmonics, numerous DDS spurs are present at levels below – 60 dBc.

of measurements with a spectrum analyzer to evaluate the spur characteristics. Stepping the frequency in 1 MHz and smaller steps from 100 kHz to 40 MHz revealed a variety of spurs. However for frequencies up to 25 MHz the non-harmonic spurs were generally at –60 dBc or lower. I have included three of the many spectrum analyzer plots to give an idea of the spur performance.

Figure 20 is typical. Here the DDS frequency is set at 3.50 MHz and the spectrum analyzer span is from 1 to 50 MHz. As noted on the figure, the fundamental, 2nd, 3rd, and higher harmonics are present. The spectrum analyzer display line (DL) is set at –57.7 dBm or –60 dBc. Numerous DDS spurs can be seen poking up above the noise level, but most are below –70 dBc. This type of spur comes and goes as the DDS is tuned from one frequency to another.

Figure 21 shows the phase noise from the 3.5 MHz carrier out to 40 kHz above the carrier. The display line is set at –87.3 dBm, or at –90 dBc, and just above the noise. After correcting for bandwidth, the phase noise level is approximately –110 dBc/Hz. However, those spurs shown above are still lurking just beyond view.

In Figure 22, spurs due to aliases are shown. These become more evident as the output frequency is increased above 25 MHz. In Figure 22, the DDS carrier is set to 30.0 MHz. In this plot each X-axis division is 5 MHz. Spurs greater than –70 dBc are evident at 10, 20 and 40 MHz, with the 10 MHz spur at –52 dBc. They are caused by harmonics of the output mixing with the 100-MHz clock. So we have 100 – (3 X 30) = 10, 100 – (4 x 30) = –20, and 100 – (2 x 30) = 40.

Table 4
DDS Frequency Accuracy Test Results
(The actual DDS output frequency agreed with the frequency predicted by the software to within +/- 0.2 mHz worst case.)

Test Frequencies (Hz)	*Max DDS Freq Error vs Requested Freq*	*Max DDS Freq Error vs Predicted Freq*
Band 1		
1000-1010 Hz –1 Hz Steps	+11.1 MHz & –11.0 MHz	+0.0 mHz & –0.1 mHz
Band 2		
1,600,000 to 1,600,010 –1 Hz Steps	+10.9 MHz & –11.4 MHz	+0.1 mHz & –0.2 mHz
Band 3		
29,550,000 to 29,550, 055 – 5 Hz Steps	+ 6.9 MHz & –10.9 MHz	+ 0.1 & - 0.1 mHz

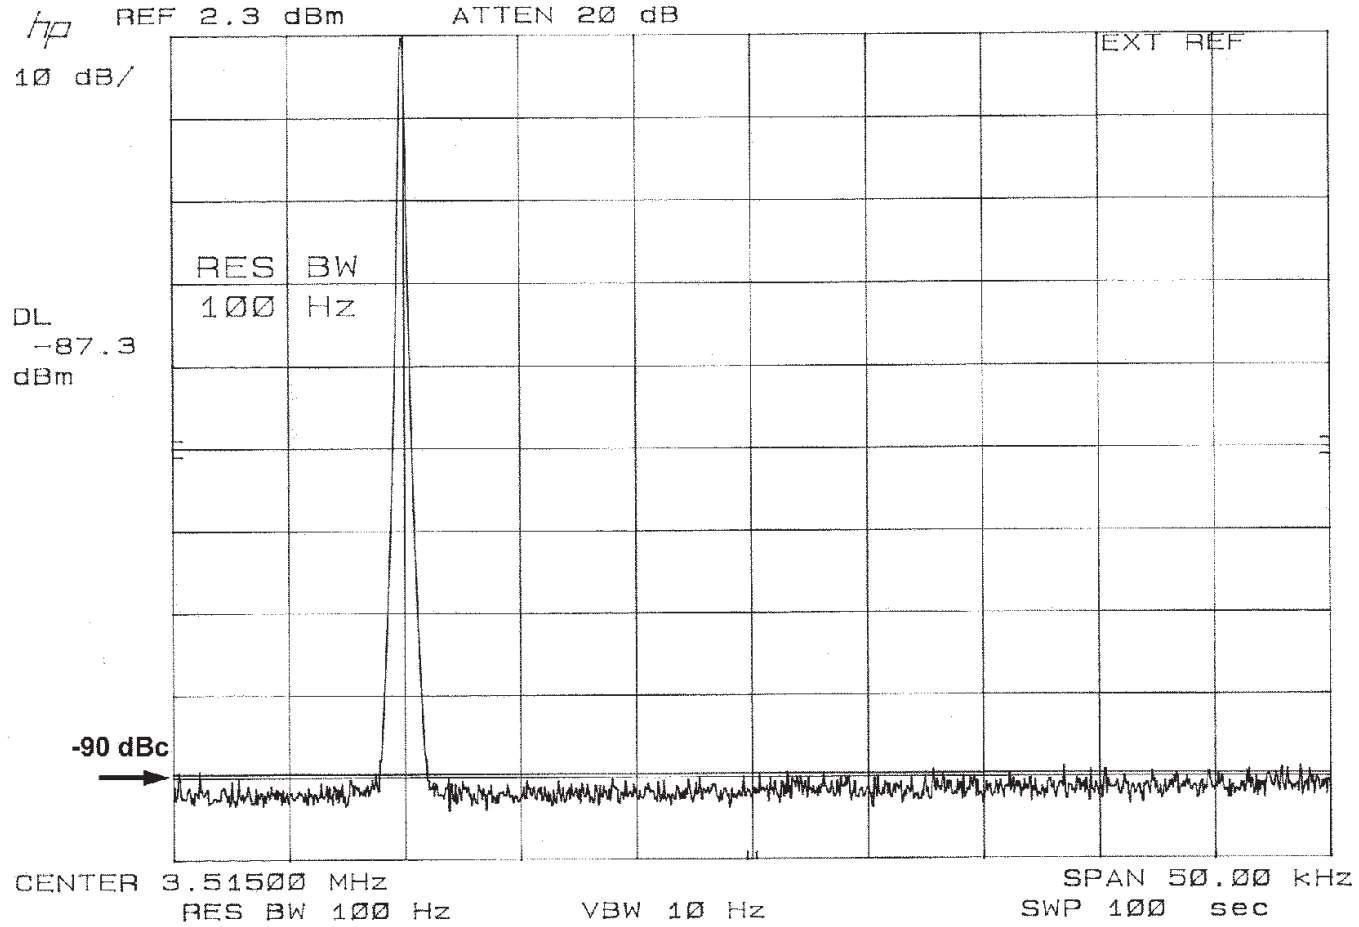

Figure 21 — A spectrum plot showing the DDS phase noise from 0 to 40 kHz from the carrier. The indicated phase noise level at –90 dBc in a 100 Hz bandwidth is equivalent to approximately –110 dBc/Hz

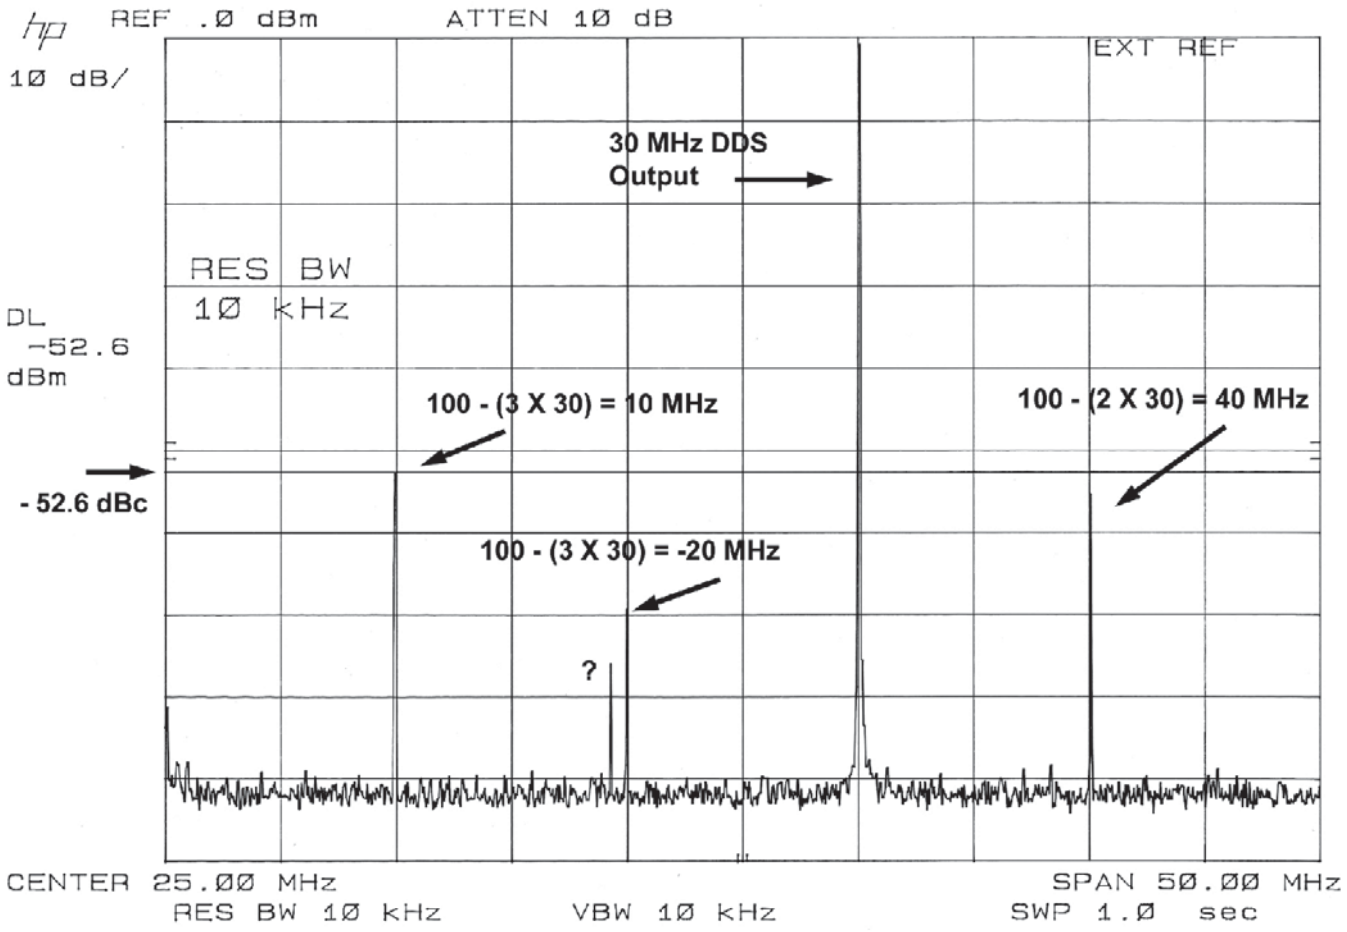

Figure 22 — This spectrum plot shows typical DDS spurs due to aliasing. These show up at output frequencies above 25 MHz. The source of the spur marked "?" is unknown.

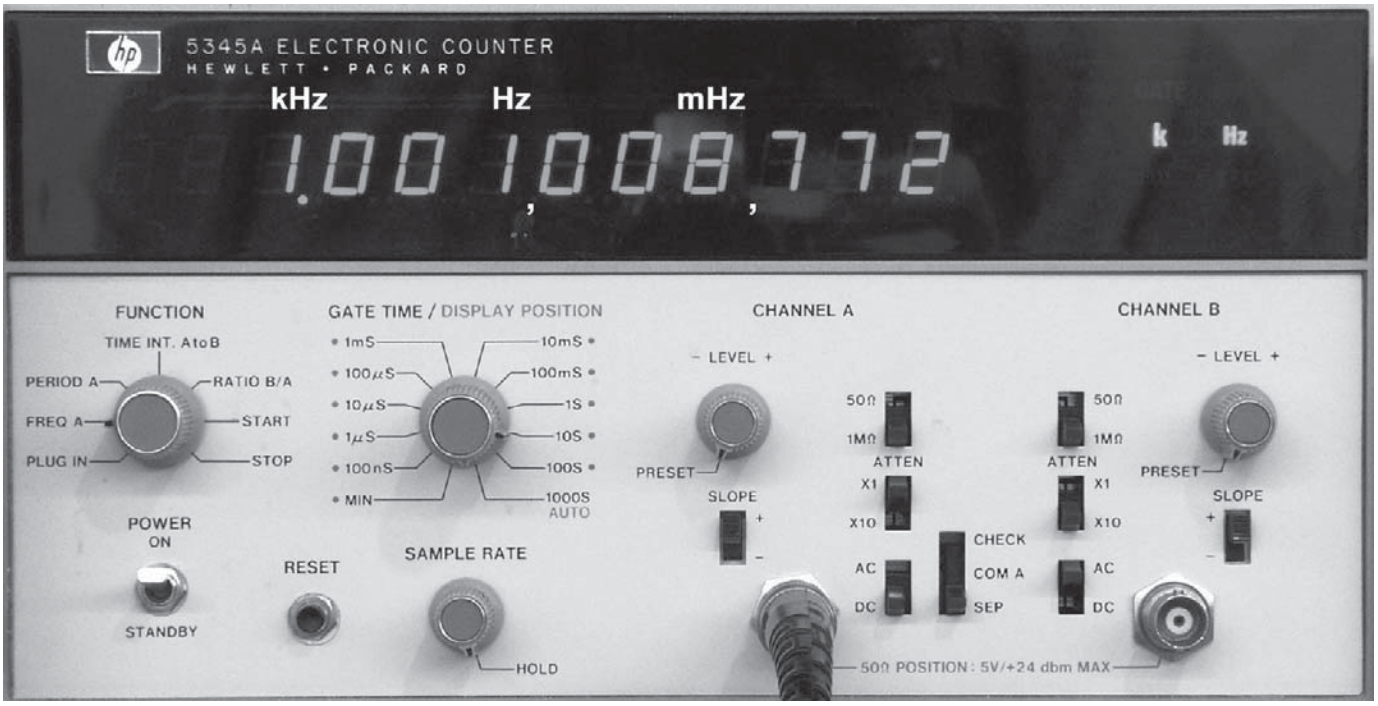

Figure 23 — The HP 5345 counter displaying the DDS output with a requested frequency of 1.001 kHz. The DDS control software predicted the actual output frequency to be 1.001 0088 kHz, just as indicated by the counter when rounded off to the nearest 0.1 mHz.

Frequency Accuracy Measurements

My objective in undertaking this project was to show by actual measurement that a DDS with appropriate control software will reliably display the true output frequency to an accuracy of 1 mHz or better. For frequencies in the audio range, this was fairly straightforward. An HP5345A reciprocal counter was used. This counter has a specified resolution of 9 digits for a 1-second gate time.[19] If the frequency is 1000 Hz, we can have 4 digits to the left of the decimal and 5 to the right, This makes the least significant digit = 0.00001 Hz. HP specifies the uncertainty of the 5th and least significant digit as +/- 2 counts. I could ignore this, since my interest was to the nearest mHz and the software displays the frequency to only 0.1 mHz.

All of this assumes the use of superbly accurate clocks to reduce the clock error that would become significant at the upper end of the frequency range. However, I cheated a bit and used the *same* Rb clock to reference both the test gear and the DDS. With this arrangement any errors due to different clock frequencies are eliminated and the results reflect only the inaccuracies of the DDS synthesizer.

Figure 23 shows the 5345 measurement of the DDS when set for a nominal 1001.0000 Hz. The DDS control program predicted the actual output to be 1 001.0088 Hz and that is what the counter shows when the measurement is rounded off to the nearest 0.1 mHz.

Measurement of the RF output frequency accurately at HF necessitates preserving the counter resolution to the right of the decimal point. In the HF measurement setup, Figure 24, a variation on Figure 2 was employed. Here a double balanced mixer, a Fluke 6070A frequency synthesizer, and a low pass filter are used to translate the unknown DDS frequency down to about 1 kHz. Knowing the Fluke synthesizer frequency and adding it to the audio frequency measured by the counter results in an accurate measure of the DDS frequency to far better than 0.001 Hz.

While many tests were made, the results in all cases were identical, so I have limited the data presented here to three representative bands, audio at 1 kHz, and RF at 1.60 and 29.550 MHz. I set the counter gate time at 10 seconds. Longer gate times would provide more resolution, but greatly lengthen the test exercise. At 10 seconds the resolution was more than adequate.

In each band the DDS requested frequency was stepped in increments of 0.1, 1, and 5 Hz respectively. There were 11 steps in the 1 kHz and 1.6 MHz bands and 12 steps in the 29 MHz band. Table 4 shows the results. The raw data spreadsheet with calculations will be included in *ROOS.ZIP*.

This data and more informal observations made during the development indicate that to at *least* the mHz level, my FMT DDS produces predictable output frequencies the accuracy of which is limited only by the accuracy of the clock.

So What Did I learn?

A few things come to mind:

• Most importantly, if the DDS clock frequency and the integer value of the actual DDS tuning word are known, the output frequency can be predicted with such precision that a Rb or GPS based frequency standard is needed to fully utilize the achievable accuracy.

• With an Rb or GPS frequency standard, and a DDS with appropriate software, one can assemble a capability for the Frequency Measurement Tests equivalent to one comprised of expensive laboratory test equipment. It is certainly more than good enough to compete with the most capable FMT participants.

• Recalling Equation 7 by adding 0.5 to the floating point tuning word value before converting it to an integer, the peak DDS round off error is halved and made symmetrical about the requested frequencies.

• The most exotic DDS chip is not required. Outstanding frequency accuracy was obtained with the humble AD9850 device.

• Controlling the DDS with an old PC using an easily programmed high level lan-

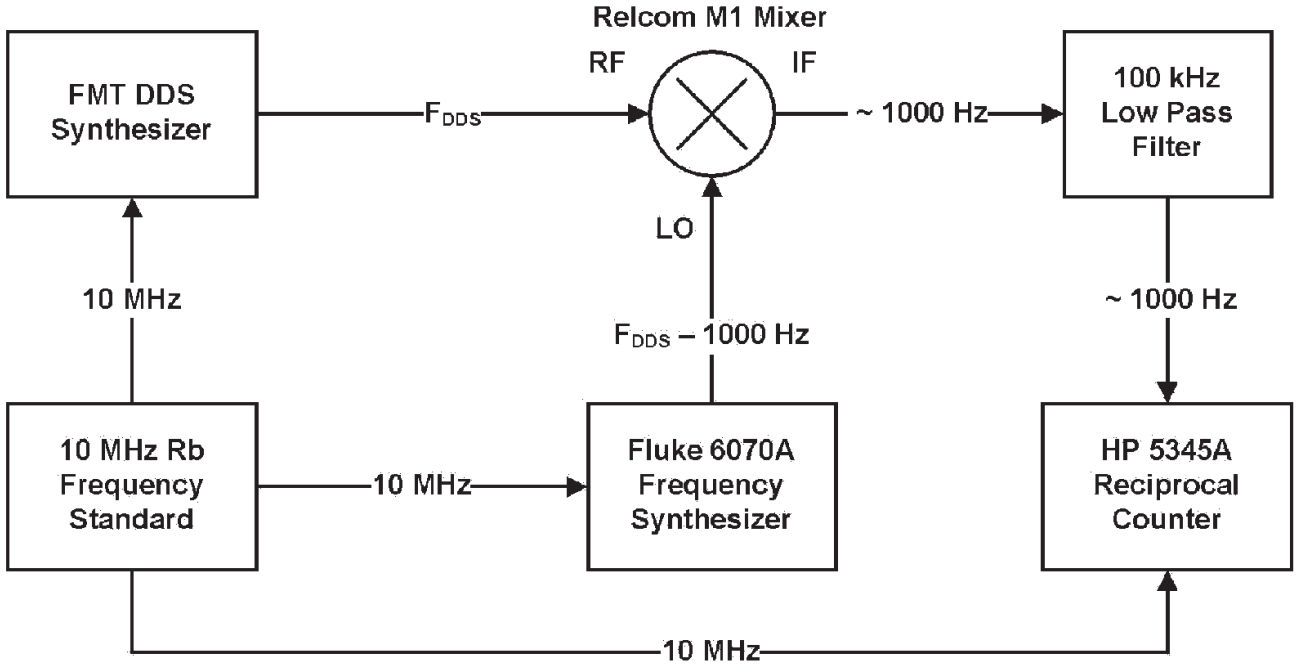

Figure 24 – Test setup for measuring the DDS RF output frequency to better than 1 mHz accuracy. Heterodyning the DDS frequency down to approximately 1 kHz increases resolution. Referencing all clocks to the same 10 MHz frequency standard eliminates clock errors.

guage facilitates complex test capabilities, perhaps automated, in roles other than the FMT application. A PC display and keyboard enable more complex and easier to use human interfaces, than 2 line LCDs and an optical encoder.

In Retrospect

This project started out as a one-off test vehicle to look at the frequency accuracy issue. As such, sections of the system block diagram and actual circuits were lifted from previous projects. While I am completely pleased with the final outcome, I have to say that it is not particularly elegant or inexpensive! Here are some improvements worth considering:

• Consider any simplification that reduces the board area. Good multilayer boards go for about 3 dollars per square inch in prototype quantities. The PWB is the single most costly component in this design.

• Eliminate the 2.4576 MHz crystal. The 10 MHz references will serve if divided by 4. Because the data stream has breaks between bytes, the minor clock skew will not matter.

• Implement the PC to DDS control interface with a micro-controller, thus eliminating most of the logic chips.

• Use only surface mount parts where possible, again to reduce the board area.

• Consider using the AD9851 with its internal X6 clock multiplier and doubling the 10 MHz reference to 20 MHz to drive it. That results in a DDS clock at 120 MHz and gets rid of most of my PLL components.

All in all, it was an interesting project that has proven useful in my lab. I would be pleased to hear from anyone with questions or comments regarding the design or ideas for design improvements.

While I have not intended this to be a construction project article, there is sufficient data available for others to reproduce the circuit. However, the existing PWB design does require a few cuts and jumpers to operate correctly and, as I pointed out, the design is not particularly cost effective. For anyone wishing to reproduce the design as is, I can provide the PWB artwork (Gerber Files) and the PWB rework instructions directly by e-mail.

Also please visit **www.K5CM.com** and join the FMT fun!

Notes

[1]Recent DDS frequency synthesizer designs: Matteo Campanella, IZ2EEQ, "A DDS Based QRSS (and CW) Beacon", *QEX*, September/October 2007, Page 29. James D. Hagerty, WA1FFL, "An Advanced Direct-Digital VFO," *QEX*, May/June 2008, Page 19.

[2]H. Ward Silver, NØAX, "Frequency Measuring Test — November 2011", *QST*, November 2011, Page 77.

[3]**www.k5cm.com**

[4]ARRL FMT Web Page — **www.b4h.net/fmt/fmtentry.php**

[5]An excellent series of application notes on electronic counters, frequency standards, and related subjects may be found at **www.agilent.com**. In "Electronic Test and Measurement" search for "AN-200." This paper "Fundamentals of Electronic Counters" explains the resolution advantages of the reciprocal counter starting on page 22. The same search will also find AN-200-1, AN-200-2, AN-3.200-3, and AN-200-4 all of which are relevant to frequency measurement.

[6]*Spectrum Lab* may be downloaded from DL4YHF at his homepage, **www.qsl.net/dl4yhf**, or at a backup site, **www.mydarc.de/dl4yhf**. These links and other related links may also be found using Goggle.

[7]See Analog Devices tutorial at **www.analog.com/static/imported-files/tutorials/450968421DDS_Tutorial_rev12-2-99.pdf.** Additionally, a Goggle search will turn up hundreds of articles, application notes, and various forums on the topic of DDS spurs and how to mitigate them. Some of the solutions are clever and worth a read.

[8]The AD9850 DDS data sheet may be downloaded from **www.analog.com**. Search for AD9850.

[9]Screen captures were made using a Pico Technology 2205 MSO (Mixed Signal Oscilloscope). For further information look at **www.picotech.com**. Scope Kit PP798 has the MSO scope and the analog and logic probes.

[10]The 6402UART is an obsolescent but very useful chip. It was used here because I have a lot of them and it is perfect for this application. Intersil made it as part number IM6402/IM6403. Harris Semiconductor also made it as the HD-6402. Jameco (**www.jameco.com**) still supplied these parts as of December 2011. Other parts such as the Texas Instruments TL16C450 might be suitable with some circuit changes.

[11]See Figure 5, page 9, of the AD9860 data sheet.

[12]Data on the MiniCircuits parts is found at **www.minicircuits.com**. MiniCircuits accepts online orders. Minimums are based on parts cost. Transformers, mixers and similar parts are often available in single quantities. The ERA-5SM+ has a 10 piece minimum order requirement.

[13]Originally a Siliconix \ Vishay part, it may be obsolescent from them. However, it is second sourced by ON Semiconductor as the SST310/MMBFU310LT1G. Various others make it as well. In this design I used the ON part.

[14]The data sheet for the LT1357CS8 and other Linear Technology parts used in this design may be obtained from **www.linear.com**. Not all of the parts are common distributor stock. Linear accepts online orders with a two piece minimum.

[15]Gardner, Floyd M., *Phaselock Techniques*, 1st Edition, John Wiley & Sons, New York, 1966. Gardner, Floyd M. *Phaselock Techniques*, 3rd Edition, John Wiley & Sons, New York, 2005, ISBN-13 978-0-471-43063-6.

[16]The Oven Controlled Crystal Oscillator is Bliley Part Number NV47AE1008. Several were obtained as surplus. According to the data sheet the temperature stability and aging rate for the lowest grade parts in the series are +/–30 ppb for 0 to 50° C and an aging rate of less than +/–50 ppb/year. The oven requires about 700 mA cold and less than 200 mA when hot. It is still shown as a current part at **www.bliley.com**.

[17]Box Engineering Ltd. enclosures are stocked by Allied Electronics. The Box Engineering Part Number is B4-220BL. The present Allied stock number is 278-0235. The part number for the protective rubber end caps (bezel) is EC4-BL and the Allied stock number is 278-0698. The mechanical drawing may be downloaded from the Allied website **www.alliedelec.com.**

[18]Sierra Circuits may be found at **www.sierracircuits.com** or **www.protoexpress.com** Their no-touch process is fully automated and not expensive for a four day manufacturing turn around. A 4-layer board requires 8 files: An Excellon drill file, a Gerber file for each layer, a top silk screen Gerber file, and a top and bottom solder mask Gerber file.

[19]*Operating and Service Manual, Electronic Counter HP 5345A*, p 1-3, Table 1-3 Specifications. Hewlett Packard Company, Palo Alto, CA. 1974.

John Roos, K6IQL, was first licensed in 1955 and obtained his BSEE from California State Polytechnic University in 1964 where he specialized in RF and microwave Engineering. Until recently he has worked in the RF area in various positions from Junior Engineer to Engineering VP while designing a variety of transmitters, receivers, and other RF devices for the Electronic Defense and Aviation Navaids industries. Projects have included a spread spectrum data link and time distribution system operating above 20 GHz; 500 MHz to 18 GHz high-probability-of-intercept electronic warfare receivers; telemetry transmitters and receivers, and more recently, Instrument Landing System (ILS) transmitters. He has two patents in the areas of spread spectrum communications and electronic warfare receiver calibration. Amateur Radio interests have always included precision frequency measurement as well as building VHF low noise receivers and high power transmitters. During the Vietnam conflict he operated an all-homebrew AM 2-meter repeater for Navy MARS on 5000-foot. Mt Wilson in Southern California. He has published articles in QEX, 73 Magazine *and others. He is a member of the IEEE, ARRL, NRA and a former member of the Association of Old Crows. Presently retired, he lives in Spring Hill, Kansas with his wife, Barbara, and is enjoying tinkering with assorted ham radio and machine shop projects.*

QEX

By Luiz Amaral, PY1LL/AC2BR

PÇA Augusto Ruschi, 31/102, Recreio, Rio De Janeiro 22790-045, RJ, Brazil, **py1ll@yahoo.com**

An Output Filter for My 500 kHz CW Transmitter

When an off-the-shelf solution isn't available, the answer may be found in your own creativity.

I was proud to put the finishing touches on my homebrewed 500 kHz CW transmitter, but since the signal was generated as a square wave it needed a filter at the output to meet spectral purity requirements. Having no commercial filters, and no way to easily buy them, I decided to craft my own filter using homemade 2.55 and 7.53 nF capacitors as described in my article "Some Homemade Capacitors" in the January/February 2012 *QEX*. My filter design would pair these capacitors with similarly homemade inductors. Assembling the capacitors was relatively easy, but creating the inductors presented a special challenge.

Due to the high inductances needed at this low frequency, it was necessary to build the inductors around good quality ferrite cores. Odd as it might seem, I found the ideal cores in the yokes of cathode ray tubes. With the decline of CRT televisions throughout the world, these cores are relatively easy to acquire in the surplus market. Not only are they inexpensive, they perform quite well at MF frequencies with very low loss.[1]

Figure 1 — A cathode ray tube yoke core. The steel clips hold the two halves together.

Inductor Construction

My filter design required two 14.93 µH inductors and one 19.66 µH inductor. Due to the nature of what was available at the time, I ended up using somewhat dissimilar cores.[2] Figure 1 shows one yoke core separated into two ferrite halves. The steel clips hold the halves together.

In Figure 2 you'll see the drawing of a typical yoke core with the dimensions indicated for External Top Diameter (ETD), Internal Top Diameter (ITD), Height (H), External Bottom Diameter (EBD) and Base Thickness (BT). The cores I used for the two 14.93 µH inductors had the following dimensions: ETD = 60 mm; ITD = 47 mm; H = 43 mm; EBD = 91 mm and BT = 4.5 mm. For the 19.66 µH inductor I used a core with these dimensions: ETD = 60.5 mm; ITD = 47.5 mm; H = 43 mm; EBD = 89 mm and BT = 5 mm.

The cores were wound with 2.5 mm diameter insulated copper wire. To make sure I achieved the desired inductance values, I repeatedly measured the cores with an LRC SmartTweezer meter.

The relationship between the inductance with a given core and the number of turns on it is called its *AL* value. Of the three cores I had available, two had AL values of approximately 414 and the third had an AL value of 400. The AL values were determined for each core by winding three

[1]Notes appear on page 38.

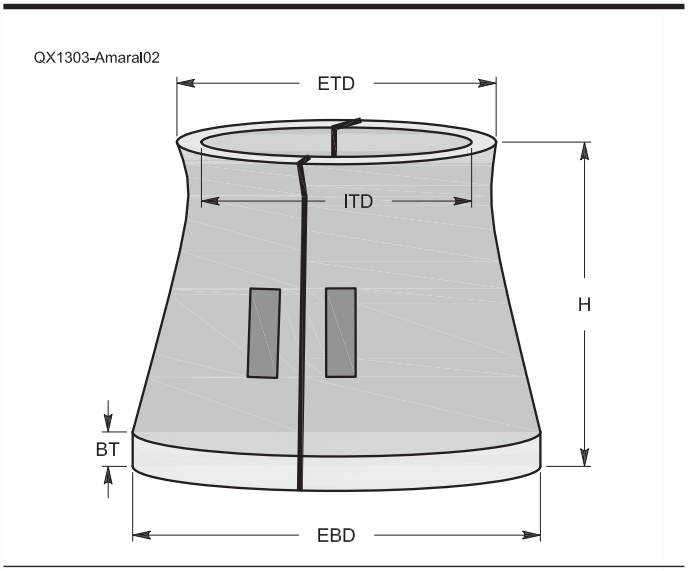

Figure 2 — A typical yoke core labeled for External Top Diameter (ETD), Internal Top Diameter (ITD), Height (H), External Bottom Diameter (EBD) and Base Thickness (BT). See text.

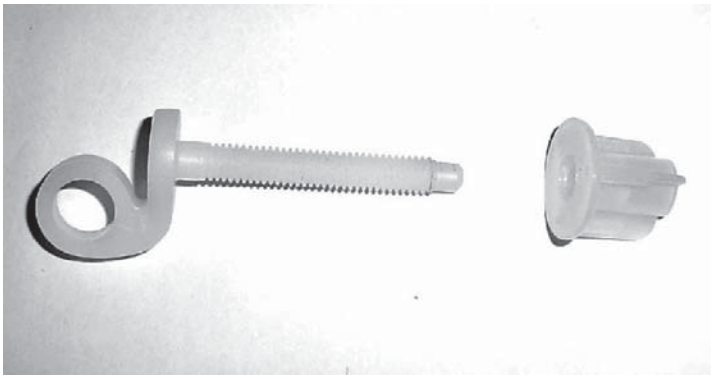

Figure 3 — The original bolt had a ring on its head, as you can see here.

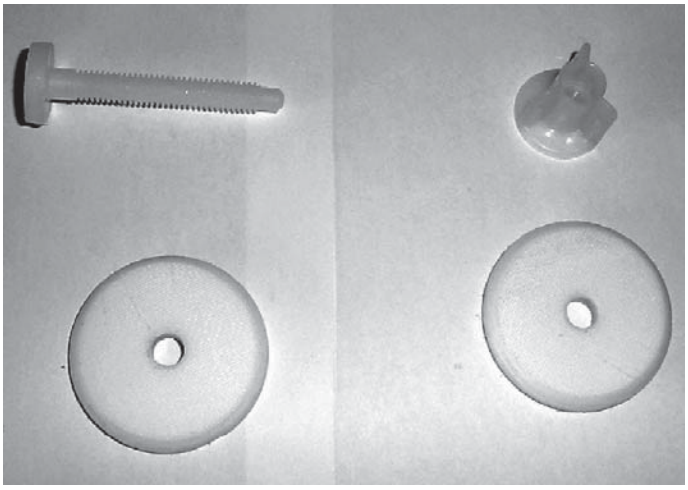

Figure 4 — Here the bolt ring is now removed. The plastic washers are shown as well.

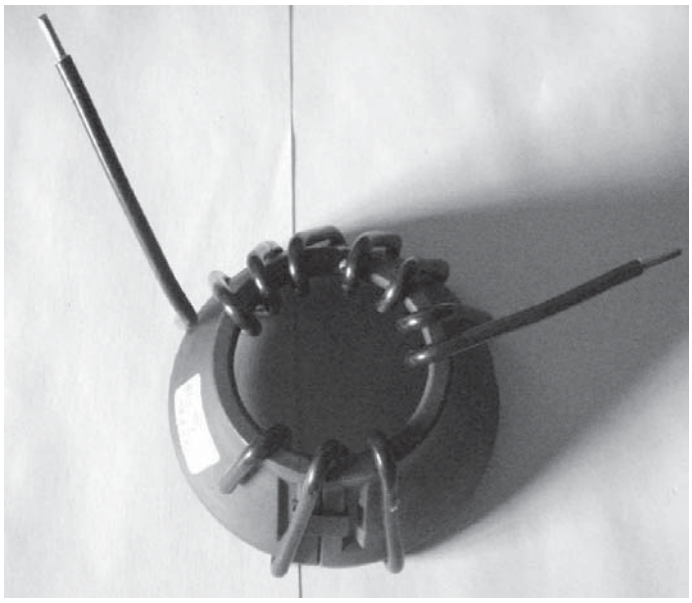

Figure 5 — Notice that the secondary winding is not connected to anything. It serves strictly as a means to balance the assembly and make centering easier.

Figure 6 — A side view of the finished inductor showing the plastic mounting hardware.

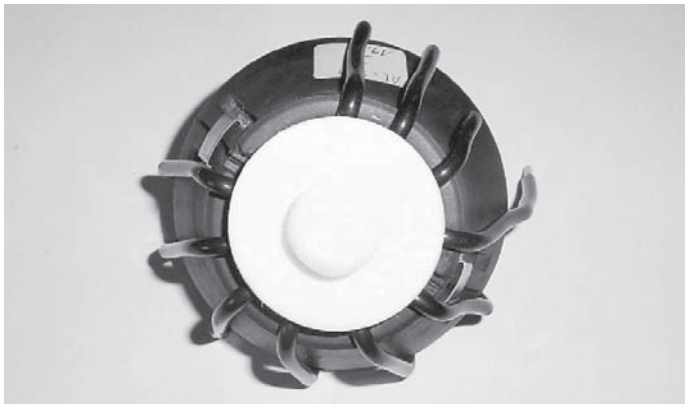

Figure 7 — A view of the top of one of the inductors, looking down.

turns and using the following formula:

$AL = L / n^2$

…where L is the inductance measured in nH and n the number of turns. AL values are given in nH/turn^2, as this appears to be the current convention.

Because of the differing AL values, I had to make some adjustments when winding the cores. For the two 14.9 μH inductors I had to wind six turns without gaps on the AL 414 core and seven turns *with* gaps (adjusted with Teflon tape) on the AL 400 core. The 19.6 μH core was made with seven turns on an AL 414 core without gaps. The gaps are necessary because the required numbers of turns are integers, or at least half-integers and fine adjustments are not feasible.

Putting it All Together

I mounted the inductors in an aluminum enclosure using plastic discs, nuts and bolts.[3] The original bolts had rings on their heads as shown in Figure 3. These rings were cut off to create the bolts shown in Figure 4. Note the angled bezels at the rims of the plastic washers. These bezels are handy for making adjustments to get the inductors in place.

To correctly mount the inductors I used open circuit (unconnected) *secondary* windings strictly as a means to balance the assembly and make centering easier (using the washer bezel), as shown in Figure 5. In the Figure 5 image you'll notice that the top turns comprise the actual inductor while the bottom turns are the open circuit windings used only for centering purposes, with the external and internal washers as in Figures 6 and 7. (Figure 8 shows the three inductors wound, but still without the extra centering turns.)

The aluminum enclosure appeared to have no effect on the inductance values, according to my measurements. There was no mutual coupling between the inductors and the filter response curve was exactly as I had anticipated.

In Figure 9 you can see the enclosure before the components were installed. There are eight metal **L**s for mounting the filter capacitors and plastic bolts for inductor mounting. You'll notice the input and output PL-259 connectors with a copper

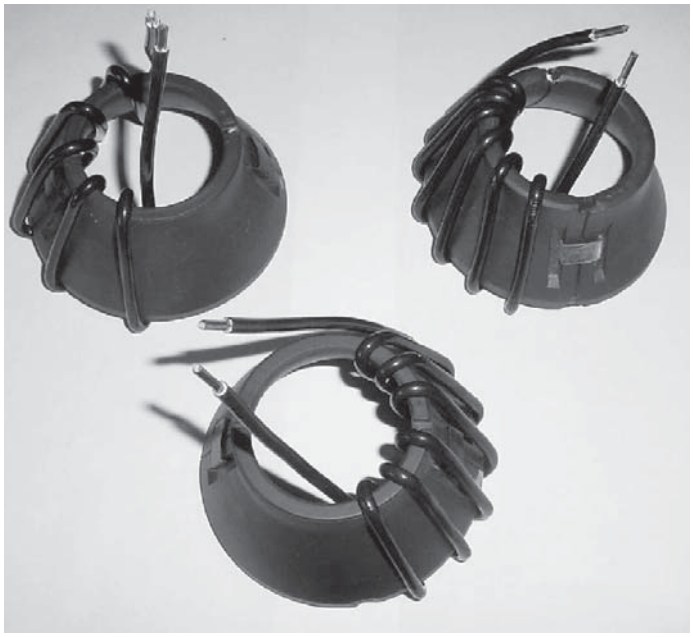

Figure 8 — The three inductors shown with their primary windings before the centering turns were added.

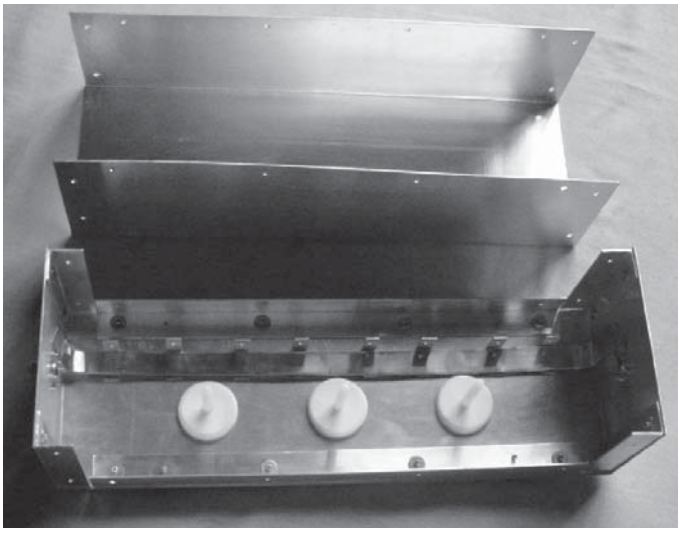

Figure 9 — The aluminum enclosure before the components were installed.

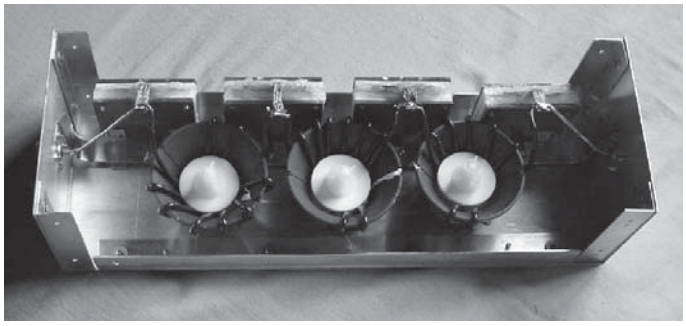

Figure 10 — The enclosure with all components installed. Note the row of homebrew capacitors along the top.

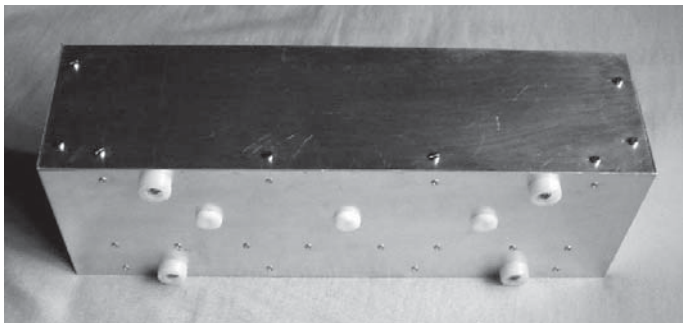

Figure 11 — The underside of the filter.

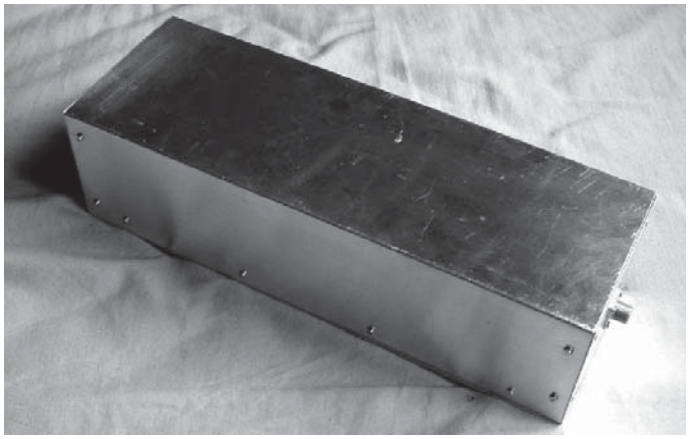

Figure 12 — The completed filter in its normal position.

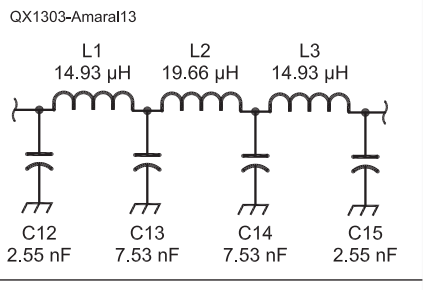

Figure 13 — Schematic diagram of the homebrew output filter.

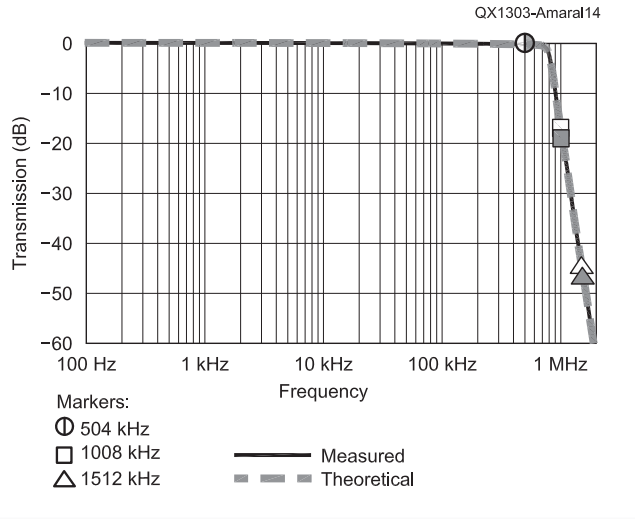

Figure 14 — The predicted response is almost identical to the actual measured response.

conductivity. Figure 10 shows the box open, but with all components mounted.

The closed enclosure is shown in Figure 11 with a bottom view to illustrate the plastic feet. The finished filter appears in its normal orientation in Figure 12.

Results

The filter diagram is shown in Figure 13 and its predicted and actual responses are depicted in Figure 14. I used *ELSIE* software to calculate the LC filter performance. In Figure 14 you can see theoretical result beside the measured result – they are nearly identical. I took my measurements using a MINIPA function generator, a Hitachi oscilloscope and a 50 Ω load connected to the filter output.

The actual measured values are: point 1, 0 dB; point 2, –17.9 dB; point 3, –46 dB. It is interesting to note that even at 1 MHz the attenuation was 17 dB.

It is remarkable what can be done with just a collection of "junk" components, especially when you have a custom application that requires an affordable custom solution.

Luiz Amaral is a retired nuclear physicist and university professor. Born in Rio de Janeiro, Brazil, he is presently residing in New York City. A ham since 1958, Luiz enjoys all modes including ATV, packet radio and satellite communication and has designed and built several of his own receivers and transmitters. On HF, Luiz is fond of 160, 80 and 40 meters and he is currently exploring the 630-meter band. He is a widower with three children and five grandchildren.

Notes

[1]For example, see the article here: **http://py2wm.qsl.br/balun/Balun_with_free_ferrite.pdf**

[2]It is difficult to get surplus yoke cores that are reasonably similar in shape and size. My friend, Celso, PY2TS, got these cores for me.

[3]The plastic bolt is one of those used for mounting toilet covers.

Ray Mack, W5IFS

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

SDR Simplified

A Look at Noise Reduction and Adaptive Filters (Any sufficiently advanced technology is indistinguishable from magic.—Arthur C. Clark*)*

Statistics and the Nature of Noise

I will give you a little exposure to statistics (in case you haven't used it before) because noise concepts heavily involve statistics. I recommend Chapter 2 of the book *The Scientist and Engineer's Guide to Digital Signal Processing* by Steven Smith for a good starting point or refresher on statistics as well as a good DSP book.[1]

In general, we mean Gaussian white noise when we talk about noise in radio systems. Figure 1 shows the Gaussian distribution and the spectrum (which has a value of 1 from $-\infty$ to ∞). The Gaussian distribution is the bell curve we probably all experienced when our grades were "curved" by our teachers. The sample size in that case was no more than probably 30 samples. In electronics, we have a continuous system with an infinite number of samples of voltage in the system we are measuring. The Gaussian distribution is a very close approximation to what we see in the real world for 99.9997 percent of noise voltages. The equation allows for an extremely low probability of a negative infinite and positive infinite voltage which, of course, will never happen in a real system (at least hopefully not in our lifetimes). Perhaps those occurred during the Big Bang. The 99.9997 percent probability corresponds to ±4.5 standard deviations and is the number used in the incorrectly named 6 Sigma manufacturing goal.

Another important characteristic of white noise is that it is uncorrelated. This is impor-

[1]Notes appear on page 46.

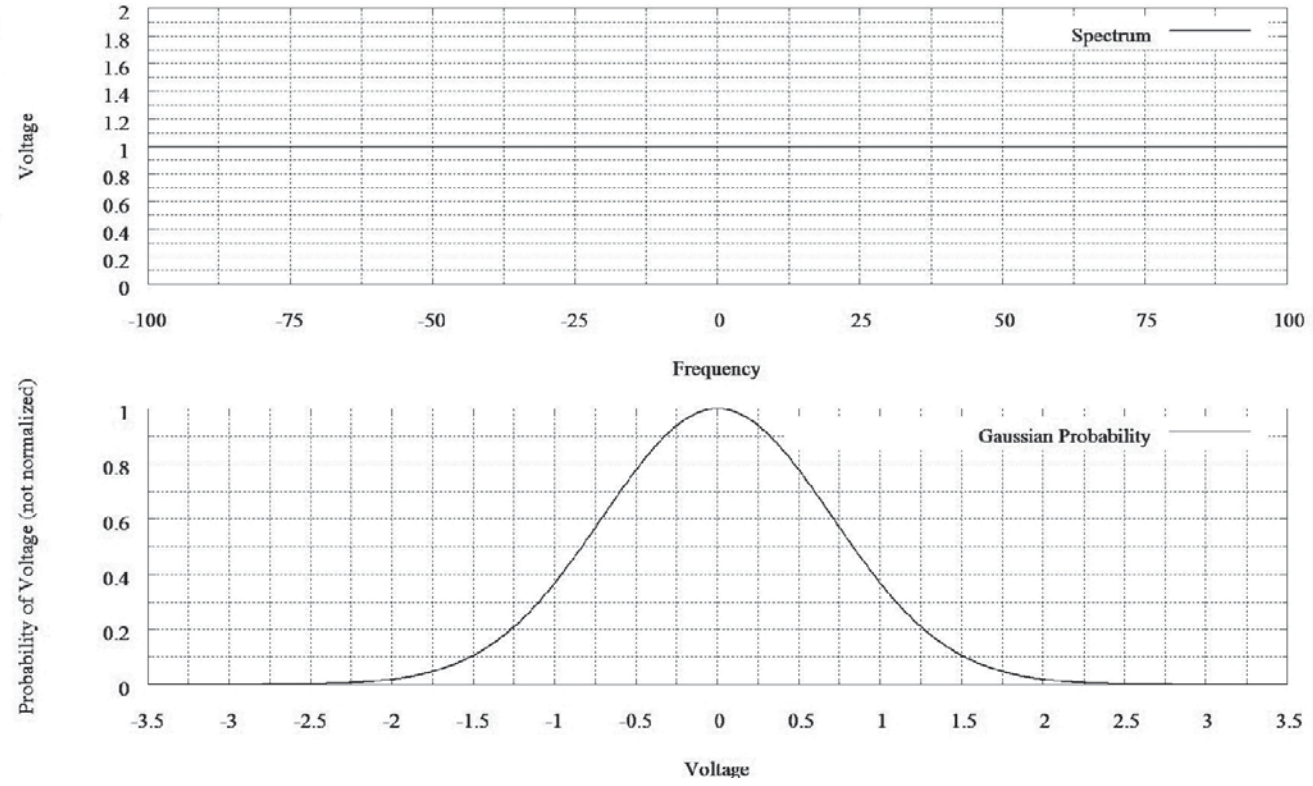

Figure 1 – (Bottom) A plot of the Gaussian probability distribution. (Top) The spectrum that corresponds to Gaussian white noise from -100MHz to 100 MHz.

tant if we implement an algorithm that acts to correlate an input signal. When I first started studying noise, I had trouble understanding the probability shape and the frequency spectrum. The Gaussian shape describes the ***probability*** function of any particular voltage at any particular time. However, noise is a totally random process, so it is reasonable for the voltage to change very quickly or very slowly as well as change from a large negative value to a large positive value. These describe the voltage as a ***time*** function. Remember that very quick changes in time translate into very broad frequency spectrum. The dirac delta function has an infinite frequency spectrum with a constant value. Random noise actually contains signals that come very close to random instances of very small dirac delta functions.

The Moving Average Filter

Smith observed that the moving average filter is frequently the first choice by engineers looking to reduce noise because of its simplicity. The moving average filter is optimal when one wants to reduce noise in time encoded signals. It is also one of the worst filters for frequency encoded signals. The moving average filter excels at reducing noise in signals where we need to keep the sharpest step response. Two radio examples that can benefit from the moving average filter are receiver AGC and transmitter ALC. Receiver AGC is an excellent example where we are interested in detecting the edge of a step in signal level. Figure 2 shows an example of a 100 mV square wave such as an AGC input with 25 mV_{P-P} white noise (top plot) and the results of passing it through an eight element moving average filter (bottom plot) or four element filter (middle plot). The rise times of the edges are lengthened but the edges remain very sharp. It is interesting how much noise even a four element filter will remove. A moving average filter is actually a special case of FIR filter where each tap has the value 1/N with N being the number of taps.

The moving average filter is especially easy to implement in *C* code because all of the values have equal weight. That means that we do not care about the order we add each value. Listing one shows two different implementations in *C*. The first implements a circular buffer of arbitrary size that replaces the oldest sample with the newest sample and then computes the average. The second example illustrates improved efficiency of execution if the circular buffer size is a power of two.

The Matched Filter

The matched filter is another optimal time domain signal noise reduction technique. If you know the exact wave shape of the desired signal ahead of time, you can use the cross correlation function to determine when that signal occurs. A matched filter looks very similar to an FIR filter in implementation. An FIR filter works by performing a convolution of the input time samples with the impulse response of the desired frequency response. A cross correlation filter also passes the time samples through the filter, but the values in the filter are an exact replica of the desired signal. Figure 3 shows the damped sinusoid expected waveform that we used for an implantable defibrillator using magnetic pulses with pulse position modulation for encoding. The top graph shows the sample values used for the matched filter calculation. Figure 4 shows the expected signal with noise and the output of the correlation machine. The signal plus noise contains random noise (not really white noise) and two pulses. One pulse is a positive version of our signal that begins at sample 23 and the second is a negative version that begins at sample 57. The noise is roughly 25 mV_{P-P} and the signals are also 25 mV_{P-P}. You really have to use a lot

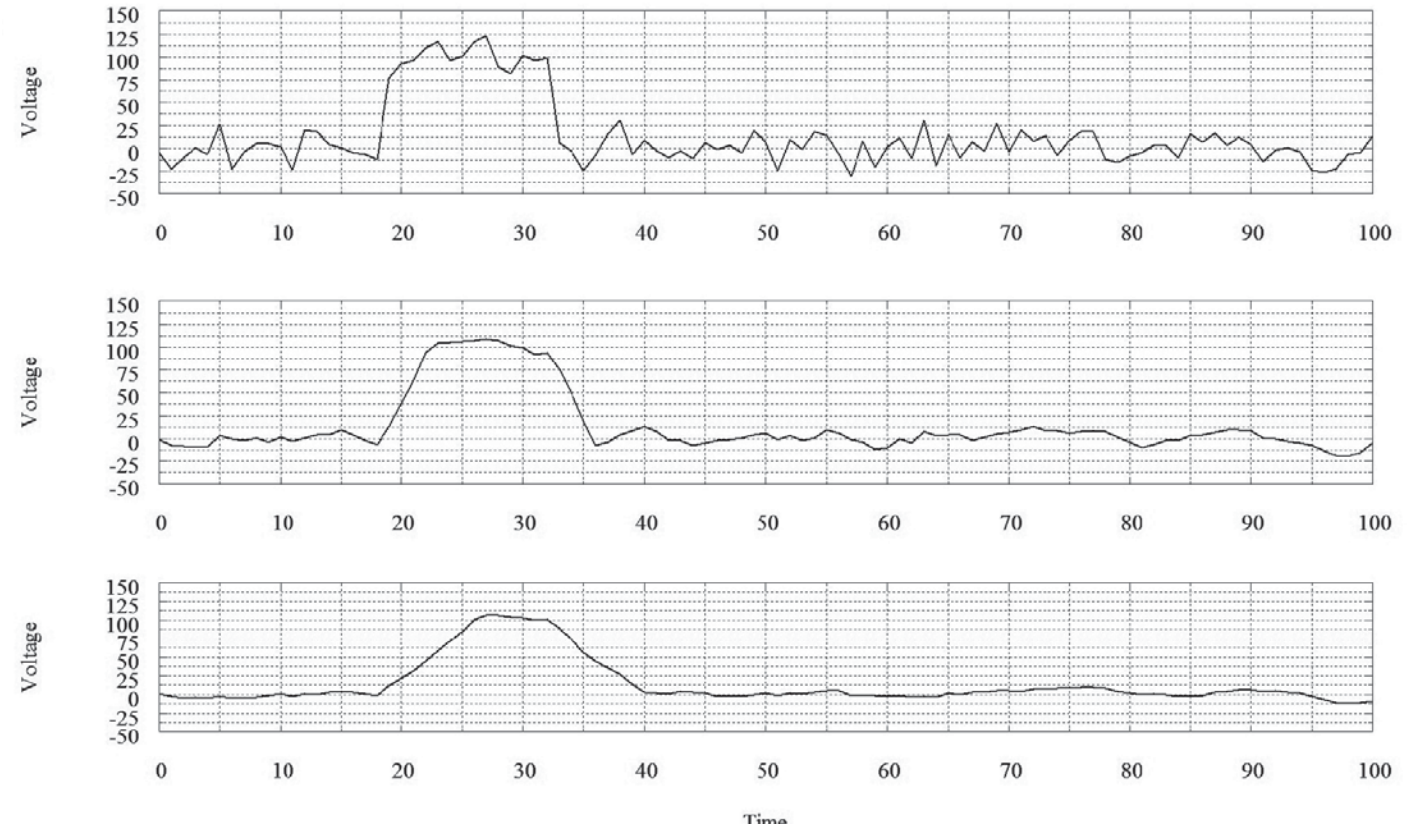

Figure 2 – (Top) A plot of an example AGC square pulse contained in random noise. (Bottom) The output of an 8 tap moving average filter. (Middle) The output from a four tap moving average filter.

Listing 1

Two Moving Average Filters

```
#define NUMBER_OF_SAMPLES  11@@// an arbitrary number that fits
                               // the number of samples we want
                               // to average
@
int@sample_store[NUMBER_OF_SAMPLES][ @
int next_index;

int arbitrary_moving_average(int new value)
{
unsigned int i;
int          accumulator;

    // store the new sample in the ring buffer
    sample_store[next_index] = new_value;
    // set up the pointer to the next entry in the ring buffer
    next_index++;
    if (next_index ==  NUMBER_OF_SAMPLES)
    {
      next_index = 0;
    }
    // we do not care about the sample order
    // because of the commutative property of addition
    accumulator = 0;
    for (i=0; i < NUMBER_OF_SAMPLES; i++)
    {
      accumulator += sample_store[i];
    }
    // the function result is the new moving average value
    return (accumulator / NUMBER_OF_SAMPLES);
}

#define ARRAY_SIZE   8       // This value must be a power of two
                              // for the logic to work
#define MASK         0x7     // the bit pattern to mask for the array
                              // size
#define BITS_IN_SIZE 3       // the number of bits that correspond to
                              // the exponent of the array size
                              // 8 == 2**3
int binary_moving_average(int new value)
{
unsigned int i;
int          accumulator;

    // store the new sample in the ring buffer
    sample_store[next_index] = new_value;
    // set up the pointer to the next entry in the ring buffer
    // we save several instructions because the bit masking
    // replaces a compare and several load instructions
    next_index++;
    next_index &= MASK;
    // we do not care about the sample order
    // because of the commutative property of addition
    accumulator = 0;
    for (i=0; i < ARRAY_SIZE; i++)
    {
      accumulator += sample_store[i];
    }
    // the function result is the new moving average value
    // a shift operation takes an order of magnitude fewer
    // CPU cycles compared to an arbitrary divide
    return (accumulator >> BITS_IN_SIZE);
}
```

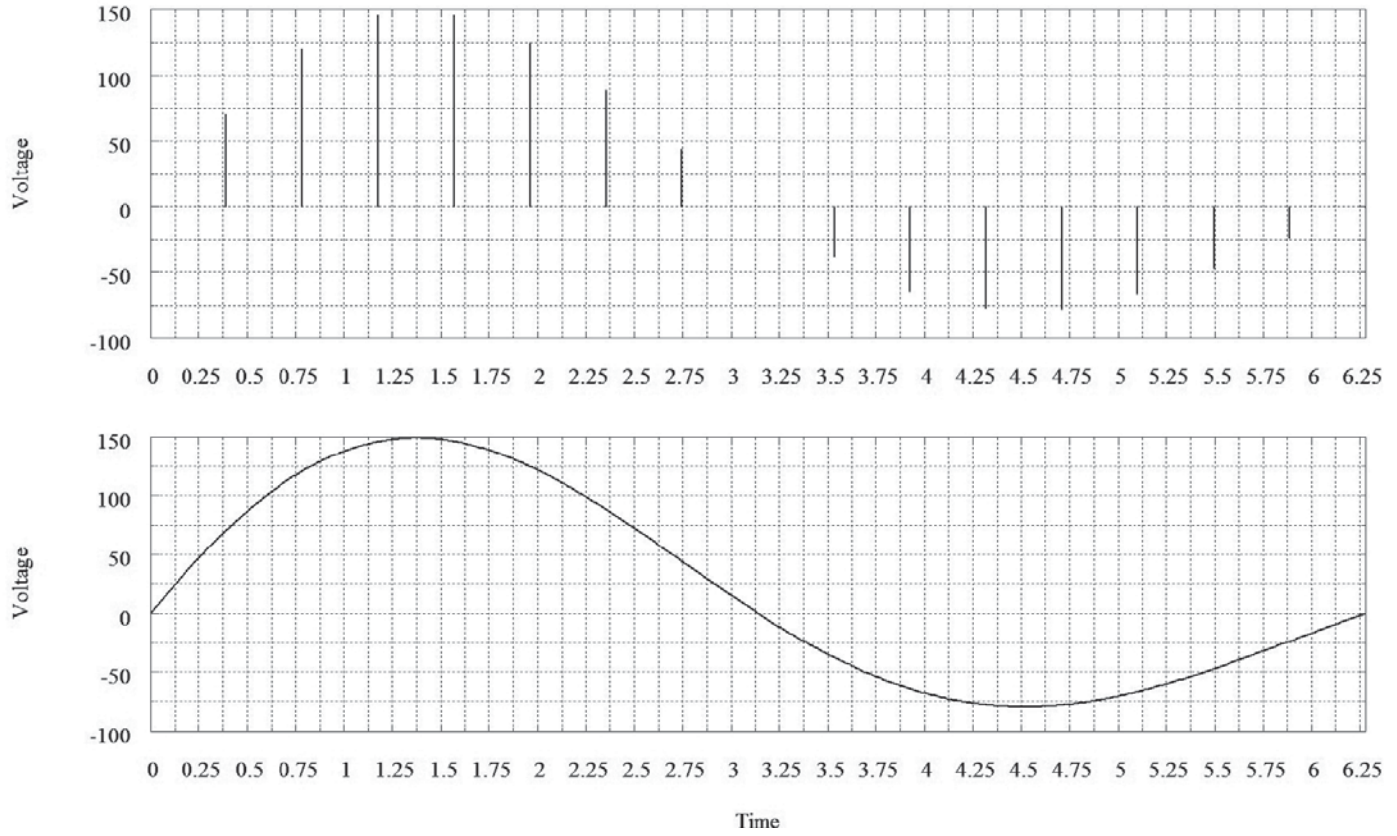

Figure 3 – (Bottom) The waveform used for our matched filter example. (Top) The sampled version of the waveform used in the matched filter.

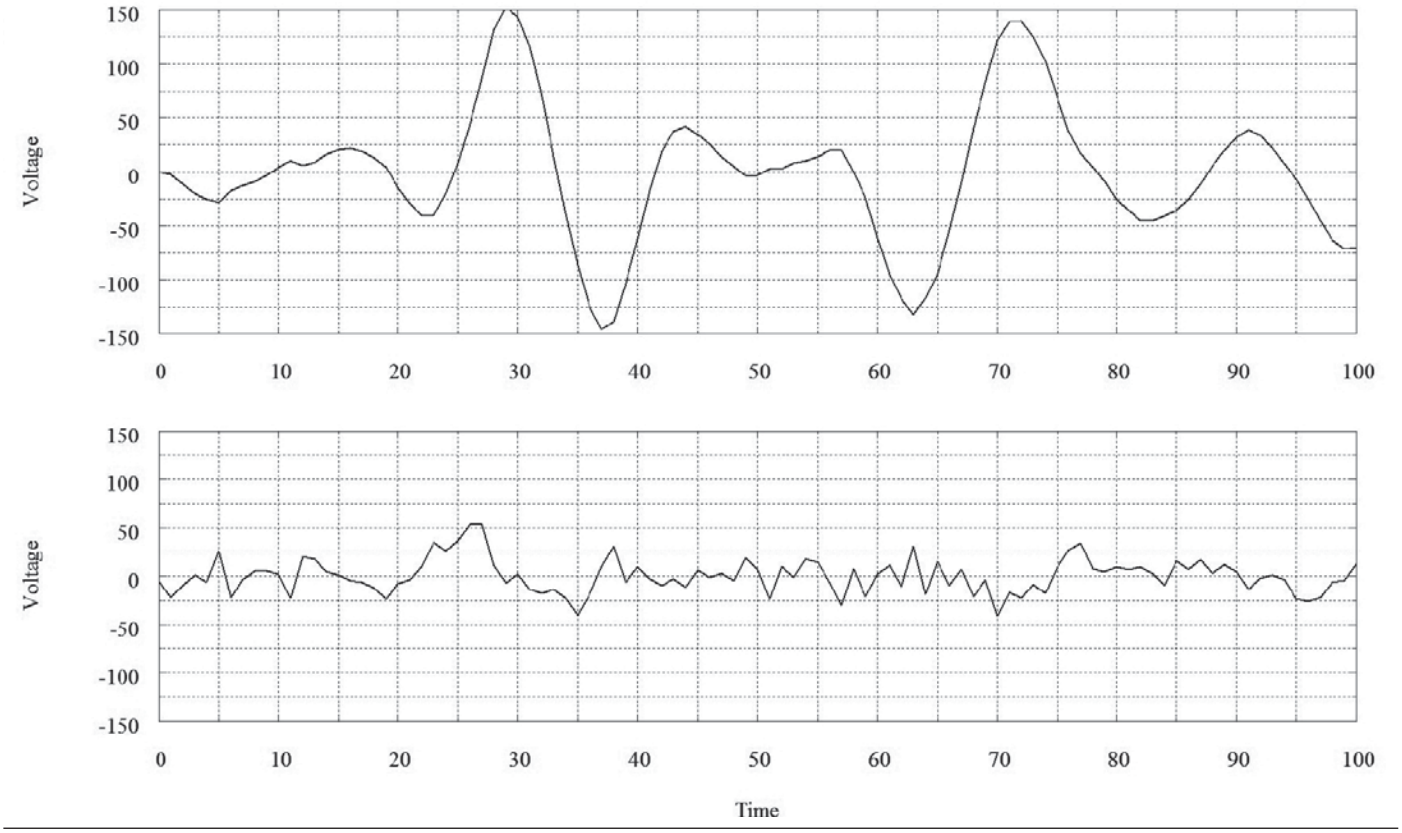

Figure 4 – (Bottom) Two matched filter pulses buried in noise. The first pulse is a positive version and the second pulse is an inverted version. (Top) The output of the matched filter showing two real pulses and a third pulse that is actually noise.

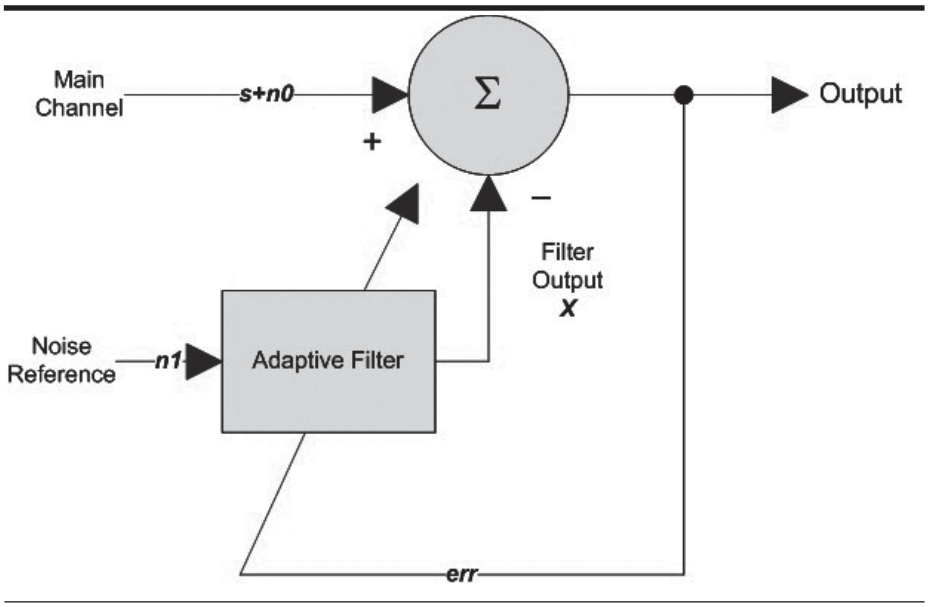

Figure 5 – A block diagram of an adaptive noise cancelling system..

Listing 2

An LMS Example

```
#define NUMBER_OF_TAPS      64
int   LMS_weights[NUMBER_OF_TAPS];
int   sample_data[NUMBER_OF_TAPS];

int   newest_sample;

void LMS_weight_adjuster(int new_sample, int error, int gain)
{
int   i, j, k;

   sample_data[newest_sample] = new_sample;
   // our pointer to input data in its circular buffer
   i = newest_sample;
   // our pointer into the weight array
   j = 0;
   // going backwards through an array is
   // an excellent example of where a do- while
   // construct works best
   do
   {
      LMS_weights[j] = LMS_weights[j] + (2 * error * gain * sample_data[i]);
      j++;
      i--;
      if (i < 0)
      {
         i = NUMBER_OF_TAPS - 1;
      }
   } while (i != newest_sample);
   // point to the next position in the data circular buffer
   newest_sample++;
   if (newest_sample == NUMBER_OF_TAPS)
   {
      newest_sample = 0;
   }
}
```

of imagination to see the two pulses among the noise. We get two output pulses that look like a single cycle of a sine wave. This is a consequence of our signal being a damped sine. Other signal shapes would have other shapes. We identify our two pulses because there is a positive peak followed by a roughly equal negative peak or a negative peak followed by a positive peak. The peaks of the output will always be 8 samples apart for this particular waveform. There is noise in the output since even noise will have some resemblance to the desired signal. In this particular sequence, the noise very much looks like our signal beginning around sample 83, but the signal level is not enough to declare a pulse if we require the output to be above 50. The beauty is that the correlation greatly enhances the detection of the signal compared to the noise. The code for this issue contains an *Excel* spreadsheet where I have implemented the moving average filter and the matched filter.[2]

The Adaptive Transversal Filter

An FIR filter is a transversal filter because the signal traverses the filter from input to output in a serial fashion. An FIR filter is a fixed function, however. There are many algorithms that can be used to modify the filter coefficients in order to accomplish varying filter performance as the signal environment varies; it adapts to the changing environment. We normally encounter two such scenarios in amateur radio. The first is adaptive notching of heterodynes or CW signals in the pass band of an SSB signal. The other is reduction of white noise on an SSB or CW signal. In the second case, the desired signal actually occupies small numbers of bins in the received spectrum and the signals are highly correlated where the noise is totally uncorrelated to the desired signal. My first DSP system to implement these functions

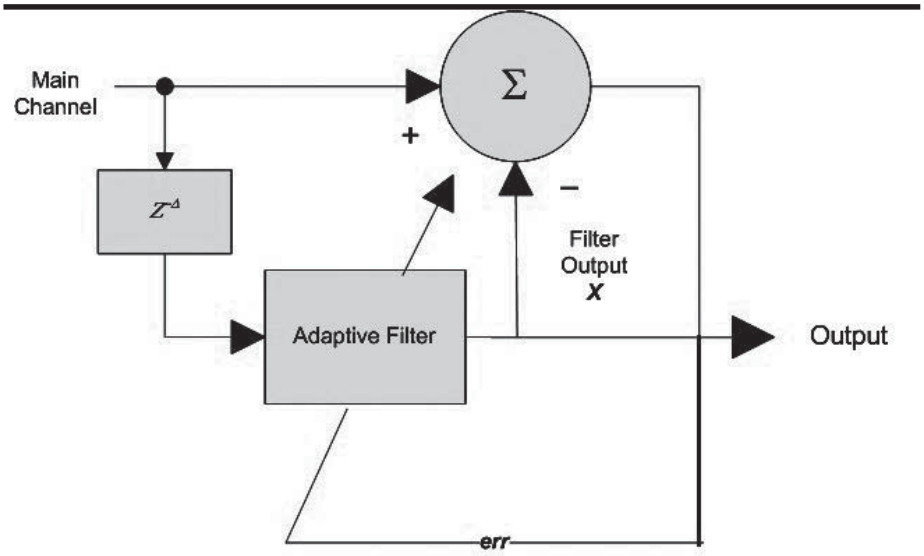

Figure 6 – A modified adaptive system using the input signal for both the main channel and the noise reference channel.

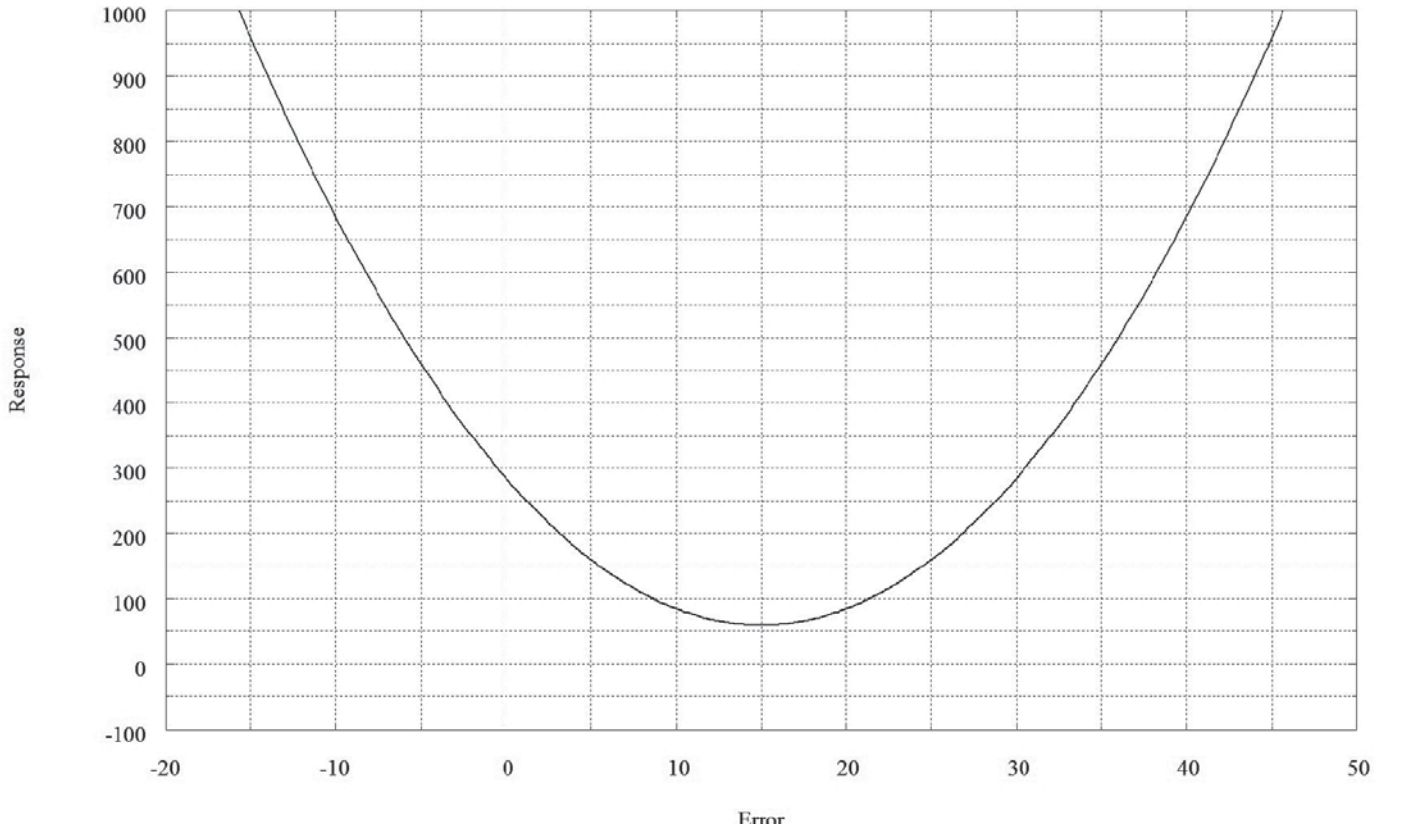

Figure 7 – A plot of a Quadratic response function. The LMS function has a response that corresponds to this type of response.

was the RadioShack DSP-40. Timewave and others make newer systems that are even better at adaptive notching and white noise reduction. Both functions are very easy to implement at baseband because the noise is confined to our audio bandwidth and we can sample fast enough to have narrow frequency bins.

Figure 5 shows one method for adaptive noise reduction. We have a signal plus noise input from the main antenna plus an auxiliary noise antenna feeding the system. The noise at the main and auxiliary channels are related, but may vary over time so that a fixed filter implementation will not be optimum. The noise channel is applied to an adaptive filter that creates a second signal that is subtracted from the main channel. If the adaptive filter is perfect, the control signal will exactly match the noise contained in the main channel and the output will be the desired signal with no noise. This is rarely (if ever) possible, but our goal is to reduce as much as possible the noise in the output signal.

This is where the math gets messy again. We start by assuming that s, n_0, n_1, and x are statistically stationary, meaning the mean and standard deviation do not change with time, and that they have a zero mean value. We also assume that s is uncorrelated with n_0 and n_1, but n_0 and n_1 are correlated. These are pretty safe assumptions in the real world. It is way beyond our scope to describe why, but the expected value of the product of two signals that are uncorrelated with zero mean is equal to zero. Likewise, the expected value of the square of any signal is the square of its value (auto-correlation). The output of our system is:

$$err = s + n_0 - x$$

Squaring (which gives us power):

$$err^2 = s^2 + (n_0 - x)^2 + 2s(n_0 - x)$$

Now we take the expected value of all of the elements:

$$E[err^2] = E[s^2] + E[(n_0 - x)^2] + 2E[s(n_0 - x)]$$
$$E[err^2] = E[s^2] + E[(n_0 - x)^2]$$

The last term falls out because the two signals are not correlated. (Yes, it seems like magic to me too.) If we adjust our filter to minimize the power in the error signal (which is also our output signal) we will maximize the signal power in relation to the noise power. We do not accidentally reduce the signal, s, instead of the noise (which would also reduce the total power) because the signal, x, is derived only from the noise that is present in both channels.

A voice signal is actually composed of a fundamental frequency plus harmonics with some small sidebands around each of the fundamental plus harmonics. If we could create a filter that was a series of band pass filters that passed only the fundamental and harmonics while attenuating all of the other frequencies in the input spectrum, we could get rid of all of the noise power at those other frequencies. We would still have the noise power that falls in the same bins as our voice signal, but the potential improvement in signal to noise is huge. This is exactly what an adaptive noise canceller does. It constantly tracks the changing characteristics of the voice signal (which, for the purposes of tracking, change rather slowly) and implements just the number of narrow band pass filters to pass the voice and eliminate the noise.

What do we do if we do not have a viable reference noise channel? Figure 6 shows a modification of the system of Figure 5. Our input consists of a correlated signal and an uncorrelated broadband component (white noise in our receiver systems). We place a delay between the input channel and the reference channel sufficient to make the signal in the input uncorrelated from the reference. The information in the reference channel remains correlated with the main channel. The result is that we attempt to minimize the correlated signal power in the error signal. The output is taken from the output of the adaptive filter rather than from the adder.

The separation of the correlated signal such as a constant sine wave from an uncorrelated signal can also be used as an adaptive notch filter. Again, by placing a significant delay between the reference and the input channels, we can take a signal such as voice that becomes less correlated as the time difference increases and create a noise signal which is basically uncorrelated. The sine wave interference has almost perfect correlation even after significant delay so the adaptive noise canceller will create an adaptive notch at the frequency.

The LMS (Least Mean Squared) Algorithm

Figure 7 shows a plot of a quadratic equation (one having squared terms). We created a function with the same general shape when we squared the error signal in Figure 5. The square of the error signal will always fall on the line and our goal is to adjust the filter values so that we find the bottom of the curve. The equation for a real system is actually a three dimensional surface, but we will simplify it to be a simple curve in one plane. All adaptive algorithms look at where we were with the last error estimate and where we are with this estimate. The goal is to always have the difference move in a negative direction on the curve. The Least Mean Squared (LMS) algorithm is constrained to work with a transversal filter, so it is not a general algorithm for all systems. The LMS algorithm is probably the most used adaptive technique for electronic DSP implementations because DSP processors are designed to implement transversal filters. Figure 8 shows an implementation of an adaptive linear combiner. It is just a normal FIR

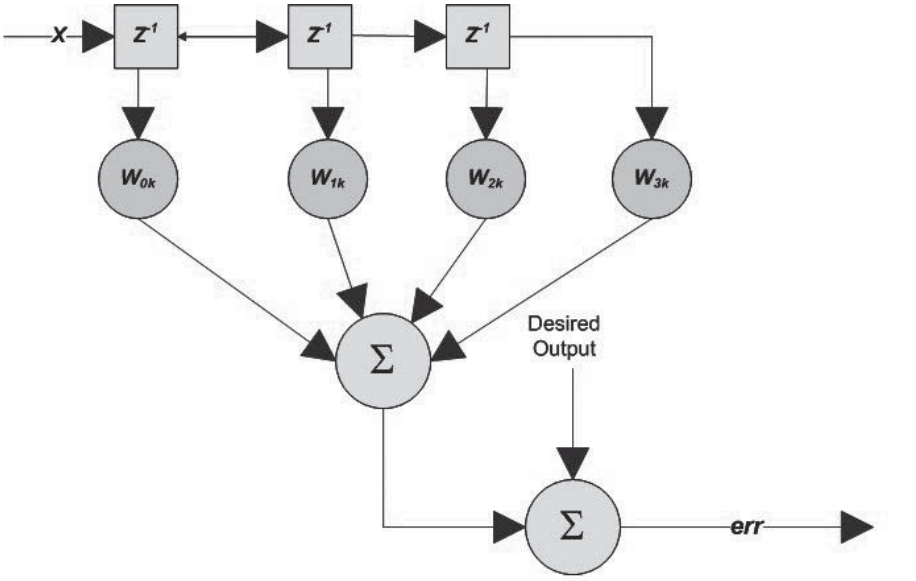

Figure 8 – A block diagram of the LMS algorithm implementation.

filter but instead of each tap value being a filter coefficient, it is an adaptive weight value. The math guys call this a weight vector. The math gets really messy with vectors and matrices and such, but as with a lot of DSP, we can ignore the math and go right to the implementation. We need just a little math to describe the LMS equation, though. We call the set of weights a vector where $\boldsymbol{W}_k$ is just $[W_{0k},W_{1k},W_{2k},W_{3k},\ldots]$ from Figure 8. Likewise, the input signal vector X_k is just $[X_{0k},X_{1k},X_{2k},X_{3k},\ldots]$. In an FIR filter, the coefficients are static so $\boldsymbol{W}_{k+1}$ (the values at z^{-1}) would be exactly identical to $\boldsymbol{W}_k$. That is not the case in an adaptive filter. The adaptive process first calculates the output of the filter and then replaces all of the weight values every time we add a new data sample to the filter. In essence we have two filter processes going on in parallel. Fortunately, the math is actually quite simple:

$$W_{k+1} = W_k + 2\mu\varepsilon_k X_k$$

where μ is just a constant gain value and ε_k the single data value that is the output of the error calculation. The software simply walks down the present set of weight values and data values and does two multiply operations and one addition per element to create the new set of weights. Listing 2 shows how simple the process is in *C*. The listing is a representative expression of the algorithm. An actual implementation is included in the source code for this issue, but can also be downloaded as part of version 1.8.0 of *PowerSDR*.[3] The file *lmadf.c* contains all of the code used in *PowerSDR*.

Reader Feedback

Peter Traneus Anderson wrote the first article about digital down conversion for *QEX* in 1996.[4] He reminded me that one must be careful when looking at the commercial digital down converter chips that are designed for broadband digital applications. Digital may work just fine at 100 dB spurious free dynamic range (SFDR), but narrow band applications frequently need 120 dB to 160 dB SFDR. Commercial applications have improved in the past 17 years, though, and some of the commercial chips should do quite well in narrow band applications.

Gary Heckman, KC7FHP, has been following our work for some time and did some work on his own to understand the Hilbert Transform. He implemented the algorithm in *MS-BASIC/QBASIC* and had good results. I have included his e-mail to me (which includes the *BASIC* source code) in the ZIP file for this issue.

More Reading

I recommend two books that I used as reference for this installment. I have already mentioned the book by Smith.[1] It is very well written and uses a conversational style. He also tries to keep the ugly math to a minimum. I also used *Adaptive Signal Processing* by Bernard Widrow and Samuel Stearns.[5] This book is full of really nasty math, so it will not help you much unless you can suffer through sophomore level engineering math. It is interesting that Widrow has a couple of equations named after him and that a lot of the earliest work in adaptive DSP signal processing only goes back to the 1960s. This is a very new area of study and products like the DSP-40 were available very shortly after the concepts were developed in academia and Bell Labs.

Notes

[1]Steven Smith, *The Scientist and Engineer's Guide to Digital Signal Processing*, Prentice-Hall, 1997.
[2]**www.arrl.org/qexfiles**. Look for *SDR Simplified 3-2013.zip*.
[3]**support.flexradio.com/Downloads.aspx?fr=1**
[4]"A Better and Simpler A/D for the DDC-Based Receiver", by Peter Traneus Anderson, KC1HR, *QEX*, August 1996, pages 21 to 24.
[5]Bernard Widrow and Samuel Stearns, *Adaptive Signal Processing*, Prentice-Hall,

Exciting Antenna Projects and Design!

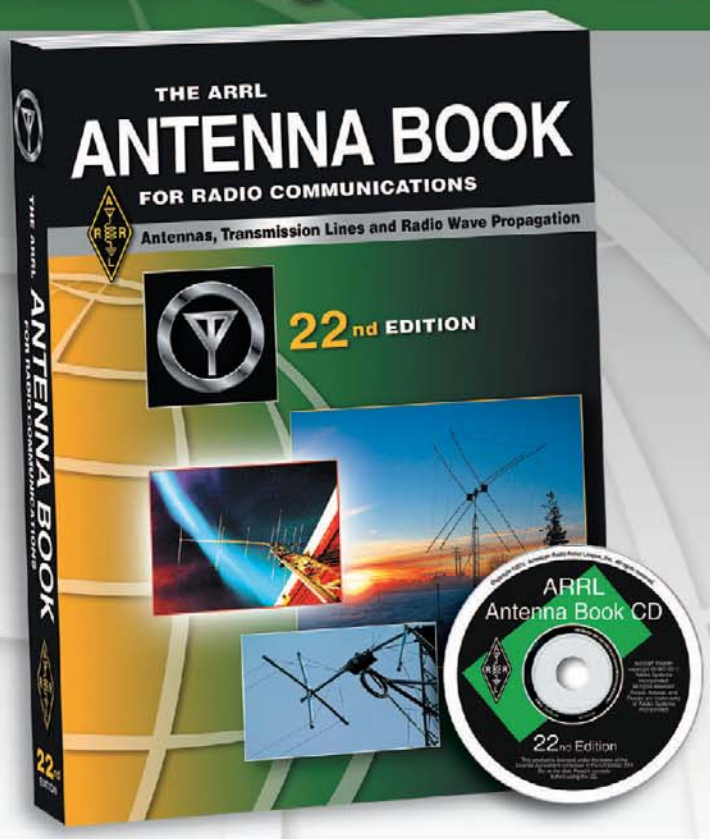

CD-ROM Inside! Includes all of the fully searchable text and illustrations, plus utility programs and supplemental content from expert contributors.

System Requirements: A fully searchable PDF version of this book is included for Windows® 7, Windows Vista®, or Windows® XP, as well as Macintosh® systems, using Adobe® Acrobat® Reader® software. The Acrobat Reader is a free download at **www.adobe.com**. PDF files are Linux readable. Utility programs are Windows® compatible, only. Some utilities have additional limitations and may not be compatible with 64-bit operating systems.

Table of Contents:

The ARRL Antenna Book

22nd Edition

Antennas, Transmission Lines and Radio Wave Propagation

Build one antenna, and you'll quickly find yourself planning the next. ***The ARRL Antenna Book for Radio Communications*** is devoted entirely to that pioneering spirit. In this one book, you will find all of the information you need for complete antenna systems—from planning, to design and construction. It describes antennas for nearly any frequency range and operating application: from the HF low-bands through VHF, UHF and microwave; fixed station, portable, mobile, maritime, satellite and more!

This twenty-second edition includes a complete reorganization, new content, and exciting new antenna projects:

- Every chapter updated or rewritten, including **limited-space antennas, portable and mobile antennas, receiving antennas, building towers, and antenna system troubleshooting**
- New designs, including a **C-pole ground-independent HF antenna, a 40 meter Moxon beam, and improved half-element designs for Yagis**
- Updated **EZNEC-ARRL** version 5.0 antenna modeling software
- Expanded content on CD-ROM, including and an extensive collection of antenna models
- New **Author and Project Indexes**

…and more!

The ARRL Antenna Book 22nd Edition
Softcover Book with CD-ROM.
ARRL Order No. 6948
Only $49.95*

*Shipping and handling charges apply. Sales Tax is required for all orders shipped to CT, VA, and Canada. Prices and product availability are subject to change without notice.

Order Online www.arrl.org/shop
or Call Toll-Free 1-888-277-5289 (US)

ARRL *The national association for AMATEUR RADIO®*
SHOP DIRECT *or call for a dealer near you.*
ONLINE WWW.ARRL.ORG/SHOP
ORDER TOLL-FREE 888-277-5289 (US)

QEX 3/2013

Quicksilver Radio

See us at Hamfests
From Maine to Florida
Visit our Huge Display at Dayton

Discrete Component Analyzer

Identifies and measures transistors, MOSFETs, J-FETs, diodes, LEDs and more.
Pocket sized and battery powered.
Visit our Web Site for more details.

LCR Analyzer

Identifies and measures inductors, capacitors, and resistors. Optional tweezers for SMD components
More info and downloadable manuals on our Web Site.

Quicksilver Radio Products

Sign up on our Web Site for your free newsletter.
Ham Radio news, articles, & special discounts.

www.qsradio.com