QEX

$7

November/December 2017

www.arrl.org

A Forum for Communications Experimenters

Issue No. 305

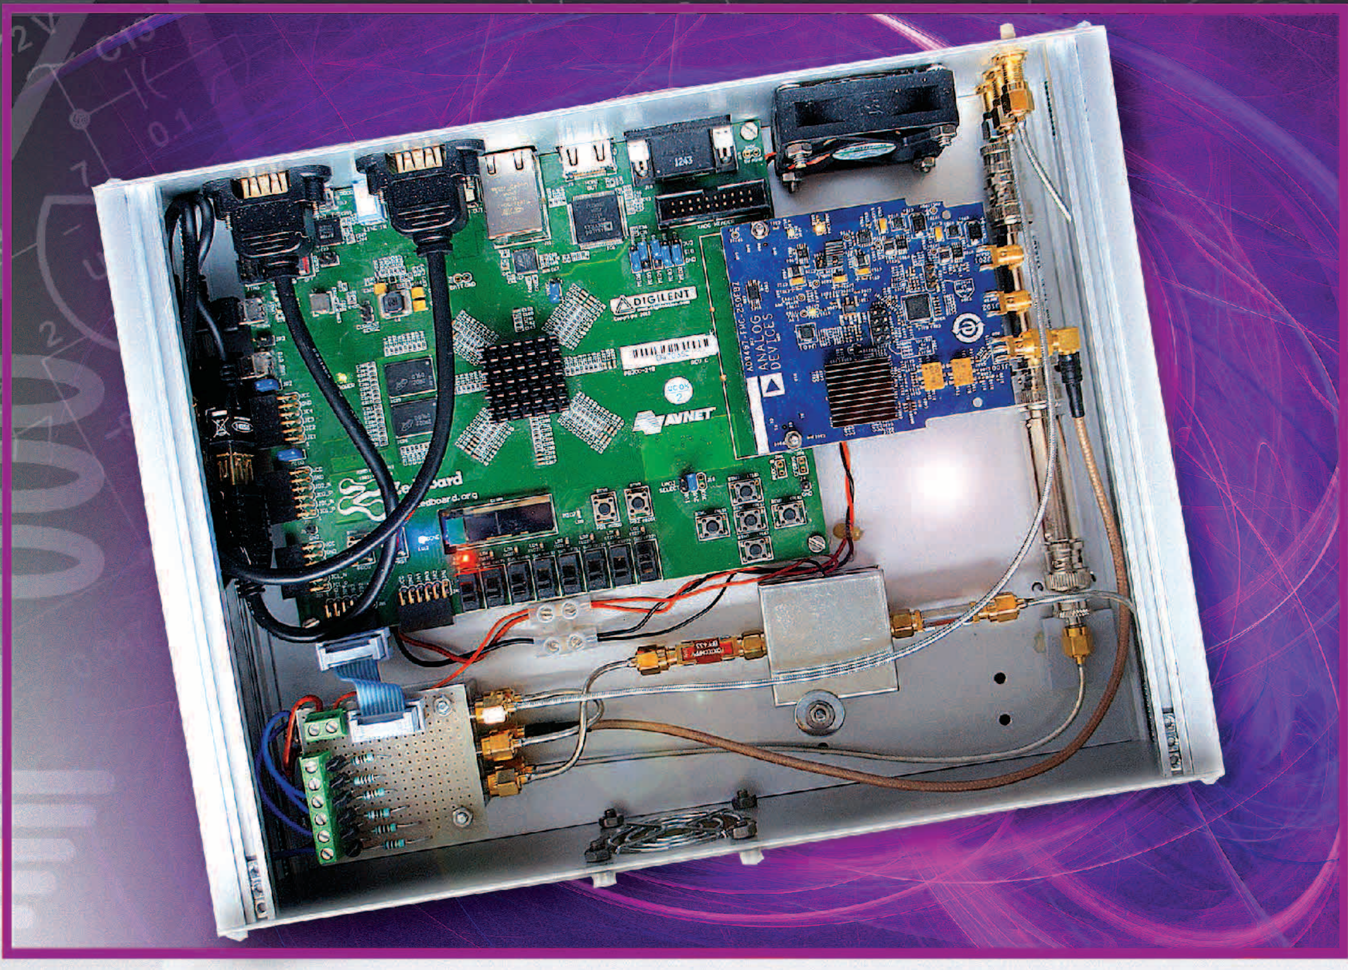

DC9ST describes a direct sampling software defined receiver that captures a seamless view from dc to VHF, and includes undersampling reception of the 70 centimeter band.

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

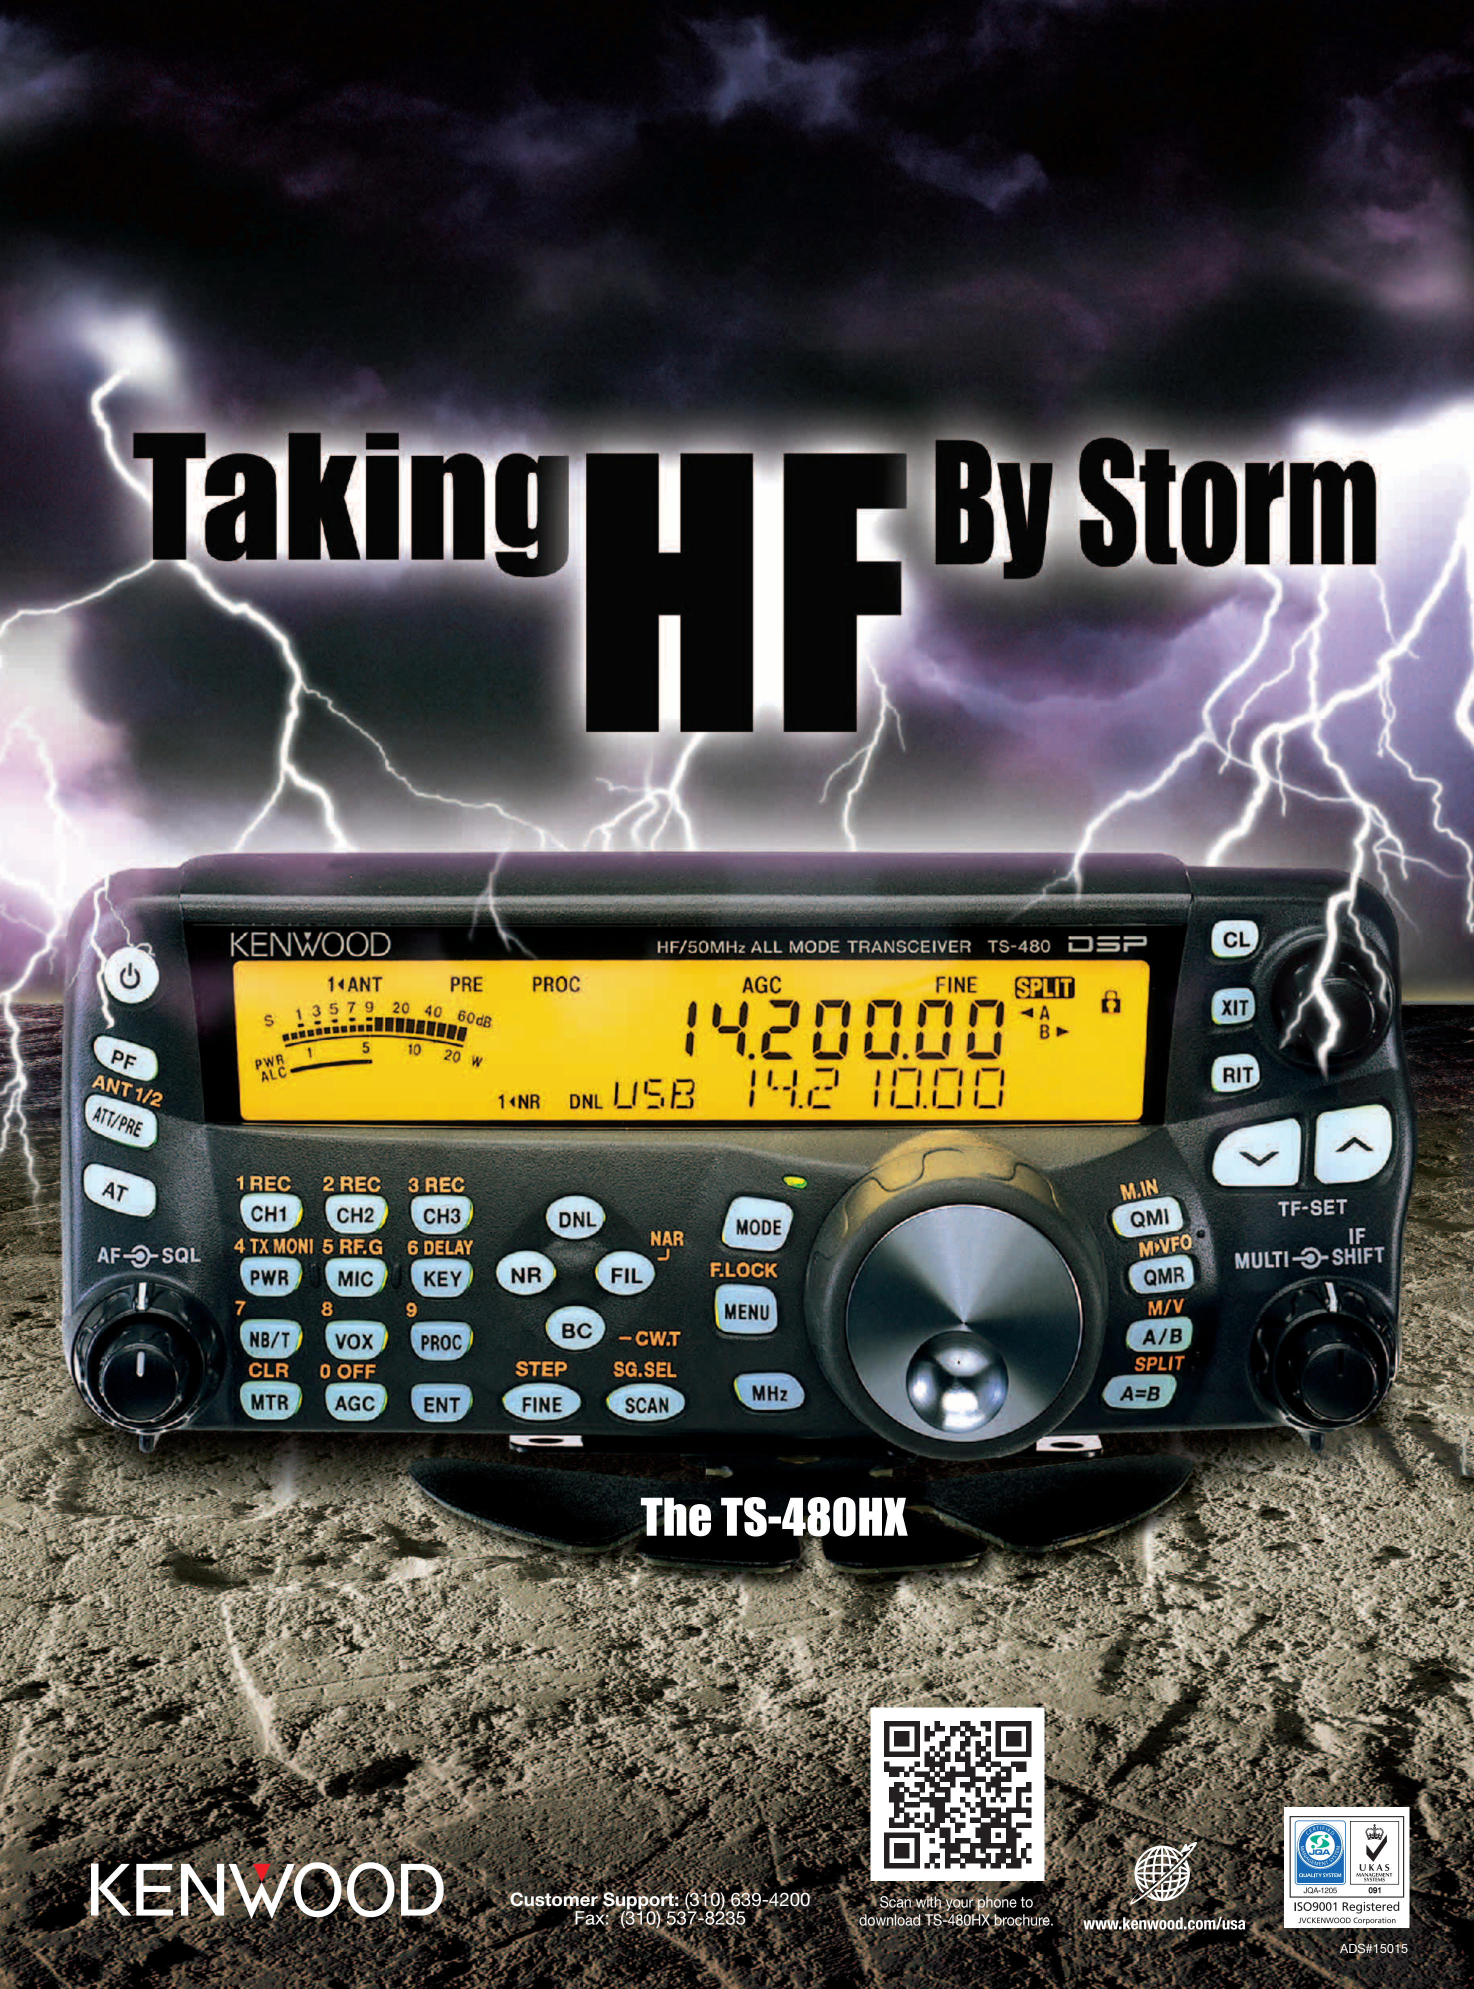
Taking HF By Storm
KENWOOD
HF/50MHz ALL MODE TRANSCEIVER TS-480
The TS-480HX
KENWOOD
Customer Support: (310) 639-4200
Fax: (310) 537-8235
Scan with your phone to download TS-480HX brochure.
www.kenwood.com/usa
ISO9001 Registered
JVCKENWOOD Corporation
ADS#15015

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 304

Publisher
American Radio Relay League

Kazimierz "Kai" Siwiak, KE4PT
Editor

Lori Weinberg, KB1EIB
Assistant Editor

Zack Lau, W1VT
Ray Mack, W5IFS
Contributing Editors

Production Department

Steve Ford, WB8IMY
Publications Manager

Michelle Bloom, WB1ENT
Production Supervisor

Sue Fagan, KB1OKW
Graphic Design Supervisor

David Pingree, N1NAS
Senior Technical Illustrator

Brian Washing
Technical Illustrator

Advertising Information Contact:

Janet L. Rocco, W1JLR
Business Services
860-594-0203 – Direct
800-243-7768 – ARRL
860-594-4285 – Fax

Circulation Department

Cathy Stepina, *QEX Circulation*

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: $29;

US by First Class Mail: $40;

International and Canada by Airmail: $35

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

Stefan Scholl, DC9ST, describes a direct sampling Panoradio software defined receiver that captures up to 100 MHz bandwidth at once and provides an extra-wide panorama view seamlessly from dc to VHF. The Panoradio has a second band pass filtered front end to enable undersampling reception of the 70 centimeter band.

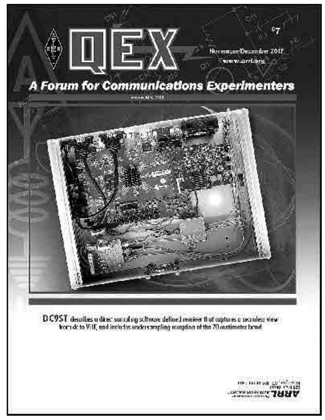

In This Issue

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: Rick Roderick, K5UR
PO Box 1463, Little Rock, AR 72203

Chief Executive Officer: Tom Gallagher, NY2RF

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**.

Kazimierz "Kai" Siwiak, KE4PT

Perspectives

Time and Leaps in Technology

Earth experienced another equinox in September marking the onset of spring south of the equator and fall north of the equator. With it comes the semi-annual assault on our body clocks known as local summer time or local winter time. When I phone my daughter Dai, KE4QXL, in New Zealand from Florida, USA, I must take into account the local summer/winter times and their mini-jetlag effects on our lives. The local time-change boundaries don't transition on the same dates across political boundaries, so her local time is either 16, 17, or 18 hours ahead of my local time depending on the date of the year — additionally, the transition dates vary from year to year. It's confusing and would be a potential mess for logging ham contacts. Ah, but we hams and the airlines have that problem licked! We use Coordinated Universal Time (UTC). Even the GPS satellites apply relativistic corrections to stay in sync with UTC. Everyone logs the same time UTC everywhere in the world, and everything works out, right?

Not necessarily! Ham technology has leapt to the limits of time accuracy where even one second matters. Some Amateur Radio modes — like several digital modes in *WSJT-X* — rely on fractional-second coordination and accuracy of UTC in our local computers. Since 1972, UTC has implemented 27 leap seconds to align Earth time with the clockwork of the Universe. The latest adjustment was on 2016 December 31 when the **www.time.gov** clock read 23:59:60. These occasional leap seconds have been branded a potential hazard to navigation, but rest assured, your computer clock is safe this year; no such micro-jetlag moment is planned for the end of 2017.

In This Issue

Our *QEX* authors touch upon a wide variety of Amateur Radio topics. These are at the top of the queue.

James L. Tonne, W4ENE, discusses asymmetrical audio waveforms, the problems this creates, and ways of minimizing the asymmetry.

Raymond F. Gurney, KDØFYF, measures the signal strength of digital mode transmissions and proposes that this might be a useful tool for acquiring propagation data.

Peter DeNeef, AE7PD, shows that unusual antennas near ground require careful treatment to assess RF exposure safety compliance.

Joseph M. Haas, KEØFF, shows how and why relays can be operated significantly below their nominal rated voltage.

Stefan Scholl, DC9ST, describes a software defined radio that captures up to 100 MHz bandwidth, as well as direct sampling of the 70-cm band.

John Flood, K4DLX, measures the characteristic impedance of transmission lines using an antenna analyzer.

Keep the full-length *QEX* articles flowing in, but if a full length article is not your aspiration, share a brief **Technical Note** that is perhaps several hundred words long plus a figure or two. Expand on another author's work and add to the Amateur Radio *institutional memory* with your technical observation. Let us know that your submission is intended as a **Note**.

QEX is edited by Kazimierz "Kai" Siwiak, KE4PT, (**ksiwiak@arrl.org**) and is published bimonthly. *QEX* is a forum for the free exchange of ideas among communications experimenters. The content is driven by you, the reader and prospective author. The subscription rate (6 issues per year in the United States is $29. First Class delivery in the US is available at an annual rate of $40. For international subscribers, including those in Canada and Mexico, *QEX* can be delivered by airmail for $35 annually. Subscribe today at **www.arrl.org/qex**.

Would you like to write for *QEX*? We pay $50 per published page for articles and Technical Notes. Get more information and an Author Guide at **www.arrl.org/qex-author-guide**. If you prefer postal mail, send a business-size self-addressed, stamped (US postage) envelope to: *QEX* Author Guide, c/o Maty Weinberg, ARRL, 225 Main St, Newington, CT 06111.

73,

Kazimierz "Kai" Siwiak, KE4PT

Peter DeNeef, AE7PD

HamRadioAndVision@gmail.com

RF Exposure Safety Compliance for Three Multi-band Inverted L Antennas

Unusual antennas near ground require careful treatment to assess RF exposure safety compliance.

In *RF Exposure and You*[1] Ed Hare, W1RFI, observes that "End-fed wires worked against earth ground almost always result in more exposure in the shack or nearby rooms than would an antenna located farther away." You can find safety compliance distances for many antennas in the General Purpose Tables in Section 4 of Supplement B to FCC OET Bulletin 65.[2] The General Purpose Tables are based on a formula that depends on frequency, power, and antenna gain. [The formula is given in Section 2 of Supplement B to FCC OET Bulletin 65 along with an EPA-recommended power factor of 2.56 that is used to account for ground reflections.— *Ed*.]. The Paul Evans,VP9KF, online calculator[3] uses this same formula. ARRL Lab comparisons with numerical electromagnetic code (NEC) simulations showed that predictions from the general purpose formula gives conservative compliance distances for a number of resonant antennas — dipoles, ground planes, and Yagis.

In 2007, Kai Siwiak, KE4PT, reported a similar comparison for a non-resonant end-fed inverted L antenna.[4] The KE4PT attic-mounted antenna comprises a pair of 48 ft long parallel horizontal wires separated 1 m apart and shorted at each end, with a 14 ft vertical section fed at the bottom against a ground rod connected to a 21 ft high vertical mast. Predictions from the general purpose formula were useful at 7 MHz and above, but they are too low for the lower frequency bands. Strong RF magnetic fields near the wires result in compliance distances greater than even the conservative tables predict.

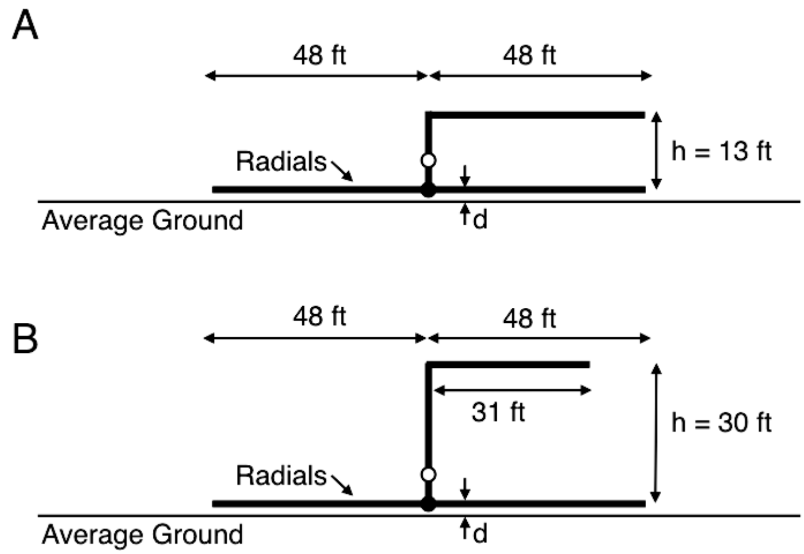

Figure 1 — NEC models for the 13 ft high (A), and 30 ft high (B) inverted L antennas. The height *d* of the radials depends on the wavelength.

I calculated FCC compliance distances using NEC models for three end-fed inverted L antennas, and compared them with predictions from the OET Bulletin tables. My NEC models[5,6] are available on the *QEX* files web page, **www.arrl.org/qexfiles**.

Figure 1 (A) shows my NEC model of a 13-ft high inverted L. It comprises a 61 ft length of #14 AWG copper wire with a 48 ft horizontal section, and is similar in overall size to the KE4PT antenna. To be conservative the model includes a very good radial ground system that is described in the NEC model section below. To avoid problems with wire-to-earth connections in NEC2 the buried radials are represented by slightly elevated wires.

Figure 2 shows FCC compliance distances for controlled areas at 100 W average power. Distances are measured from the closest part of the antenna to the point of exposure. Figure 3 shows distances for uncontrolled/public areas. The diagonal lines

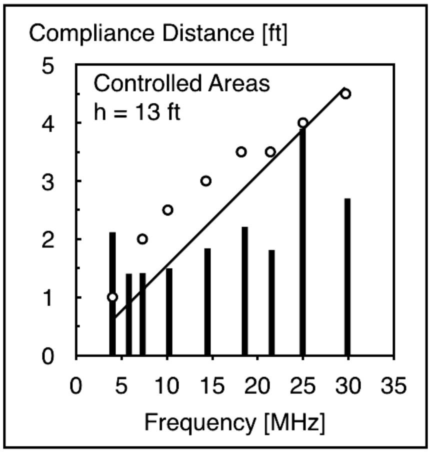

Figure 2 — The bars show the controlled area compliance distances for the 13 ft high inverted L at 100 W. The straight line represents the General Purpose Formula, and open circles are from the ARRL NEC Tables.

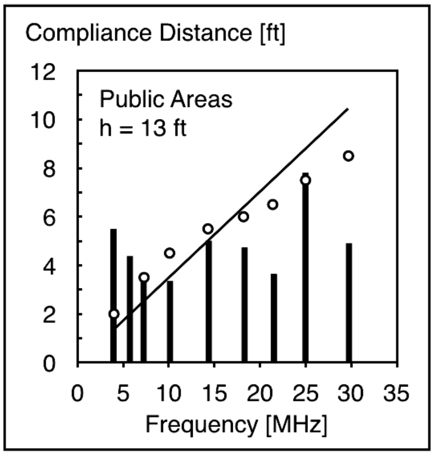

Figure 3 — The bars show the uncontrolled/ public area compliance distances for the 13 ft high inverted L at 100 W. The straight line represents the General Purpose Formula, and open circles are from the ARRL NEC Tables.

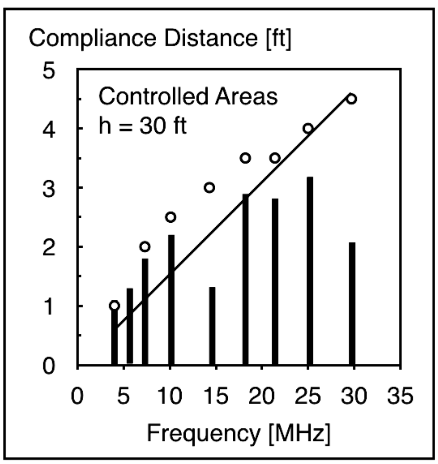

Figure 4 — The bars show the controlled area compliance distances for the 30 ft high inverted L at 100 W. The straight line represents the General Purpose Formula, and open circles are from the NEC Tables.

are predictions from the general purpose formula when the gain is 0 dBi. They are too low in the 40 to 80 meter bands.

The open circles are from the ARRL NEC Tables (from Note 1) for resonant quarter-wave ground-mounted verticals. W1RFI notes that the compliance distances from this NEC model can be used for end-fed random wires longer than a quarter-wave. Accordingly, the ARRL NEC Table predictions above the quarter-wave resonance at 4.09 MHz are close to or above the simulation results. At 25.1 MHz the two public-area distances are within 0.5 ft, a difference that is not significant when the predictions are used conservatively. There is no ARRL NEC Table for the 60 meter band.

At 25.1 MHz the feed point impedance is nearly 2,000 Ω — beyond the range of a typical tuner — because the wire length is close to three half-wavelengths.

A Higher Wire

Figure 1 (B) shows my NEC model for a 30 ft high inverted L with the same wire length and radial system as in Figure 1 (A). Figures 4 and 5 show the 100 W compliance distances for controlled and for public areas respectively. The general purpose formula predictions are again too low in the lower frequency bands, and the ARRL NEC Table predictions are close to or above the simulation results.

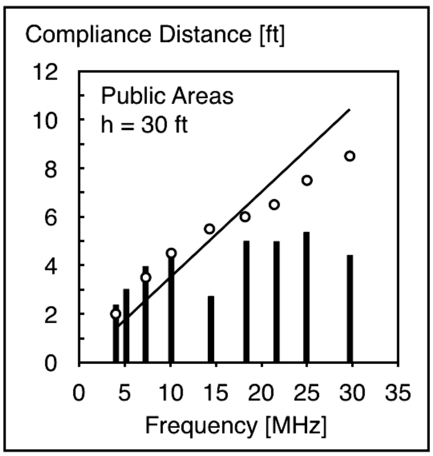

Figure 5 — The bars show the uncontrolled/ public area compliance distances for the 30 ft high inverted L at 100 W. The straight line represents the General Purpose Formula, and open circles are from the ARRL NEC Tables.

A Lower Wire

Figure 6 shows my NEC model for an inverted L mounted along the top of a 6 ft high fence. A grounded counterpoise runs parallel to the wire at a height of 1 ft. Figures 7 and 8 show that the 100 W compliance distances for this antenna are longer than for the previous examples for both uncontrolled and public areas respectively. For instance, at 14.3 MHz the distance in Controlled areas is 3.9 ft, compared with 1.8 ft for the 13 ft high wire and 1.3 ft for the 30 ft high wire. Both types of predictions fail for many of the bands in both controlled and uncontrolled/ public areas.

The fields near this antenna are more intense because the wire is coupled more closely to the ground system. Figure 9 shows contour plots from the simulation, showing two magnetic field strengths at 14.3 MHz and 100 W. H_C = 0.34 A/m and H_U = 0.15 A/m are the maximum permitted exposure in controlled areas and uncontrolled/public areas respectively. Arrows indicate the horizontal compliance distances, 3.9 ft and 8.3 ft. The cross-section view in Figure 9 is at a peak of the standing wave current, 12 ft. from the feed-point end of the wire. At a peak of the standing wave voltage — 17 ft farther along the wire — the contours of the electric field are similar to Figure 9. The compliance distances for electric field exposure are 3.6 ft and 7.5 ft, so the magnetic field determines the compliance in this case. Figure 10 shows 14.3 MHz compliance distances versus average power. Decreasing the power from 100 W to 50 W reduces the distances by about 30%.

NEC Models

I used NEC2 models to compute electric and magnetic field strengths, and compared

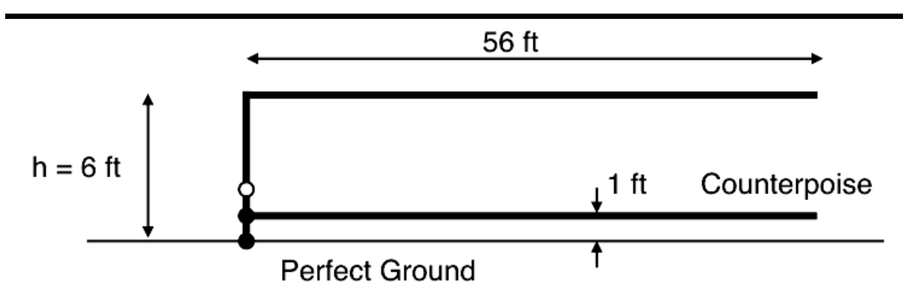

Figure 6 — NEC model for the 6 ft high inverted L antenna.

them to FCC maximum permitted exposures (Note 2). Each reported compliance distance is the longer of two calculated distances, one for the electric field and another for the magnetic field.

All wires are #14 AWG copper. To be conservative, each model includes a very good ground system. NEC2-based programs do not calculate wire-to-earth connections accurately, so buried radials are represented by wires that are slightly elevated — between 0.001 and 0.01 wavelength above earth ground. For example, at 14 MHz the elevation *d* is 0.5 ft. The ground systems for the 13 ft and 30 ft high inverted L antennas are thirty-two 48 ft radials over an NEC-average ground. The segment lengths are smaller near the central junction, gradually increasing at greater distances from the junction to reduce computation time.

The counterpoise for the 6 ft high inverted L is connected to earth ground. To avoid problems with the wire-to-earth connection I used the NEC perfect ground in the model. KE4PT compared his NEC2 simulations (including wire-to-ground connections) versus spot calculations with an NEC3 engine, which can model buried wires and ground rods. The relative field strengths away from the ground post connection were typically within 10%. Ground-level field values were similar only when he used a perfect ground in the NEC2 model.

Incidental Radiation

W1RFI lists end-fed wires with one end in the shack as a problem that can cause excessive incidental radiation. Incidental RF fields are from sources not included in the NEC simulation, such as cables, equipment, and longer wires to an RF ground system. They can create unexpected hot spots.

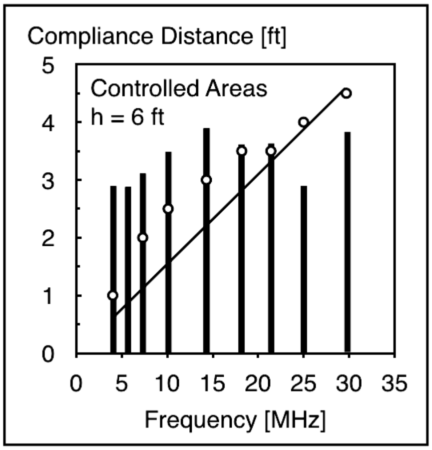

Figure 7 — The bars show the controlled area compliance distances for the 6 ft high inverted L at 100 W. The straight line represents the General Purpose Formula, and open circles are from the ARRL NEC Tables.

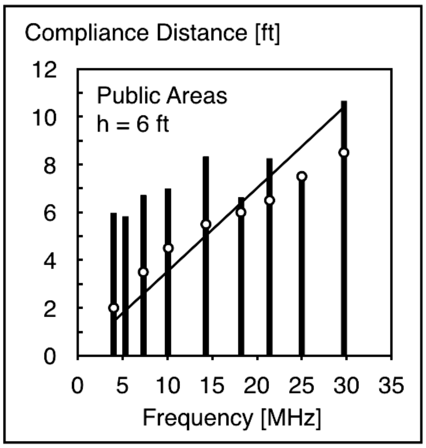

Figure 8 — The bars show the uncontrolled/public area compliance distances for the 6 ft high inverted L at 100 W. The straight line represents the General Purpose Formula, and open circles are from the ARRL NEC Tables.

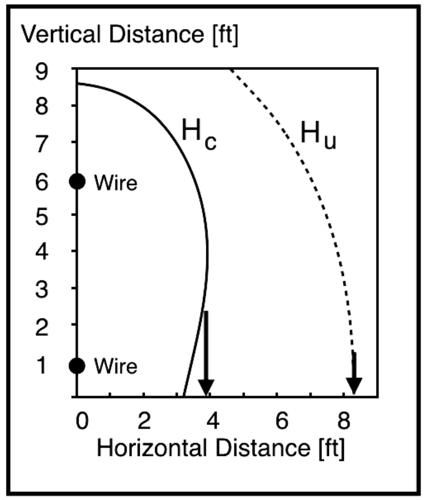

Figure 9 — Magnetic field strength contours for the 6 ft high inverted L at 14 MHz and 100 W. H_C and H_U are maximum permitted exposures in controlled and uncontrolled areas respectively. The arrows indicate compliance distances.

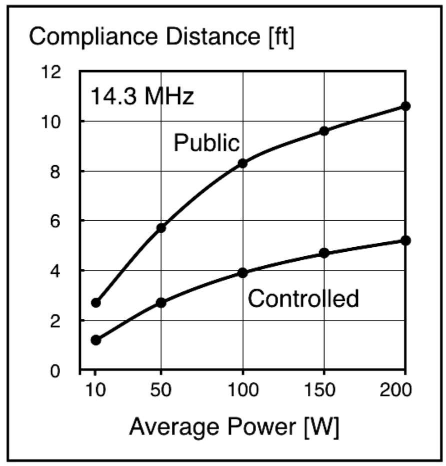

Figure 10 — Compliance distances versus average power for the 6 ft high inverted L antenna.

Conclusions

If you use a fixed or portable end-fed antenna, a remote tuner is more than just a good idea for convenience. It enables you to locate the antenna a safe distance away. These compliance distances are important if the wire is located close to people in controlled or public areas. Neither the general purpose formula nor the ARRL NEC Tables work for all bands in any of my three examples. For the two classic inverted L antennas the predictions from the ARRL NEC Tables are useful when the wire is longer than a quarter-wavelength.

For the lowest inverted L, ground coupling increases the calculated compliance distances to the point that ARRL NEC Table predictions are often too low. In the analysis of his end-fed antenna, KE4PT concludes that unusual antennas require careful evaluation, especially if a ground or ground post is part of the system.

Peter DeNeef, AE7PD, received his first license as KF7FPX in 2009. He has written articles about international RF safety guidelines (in connection with the International Commission on Non-Ionizing Radiation Protection (ICNIRP)). More articles can be found on his web site for vision-impaired hams, ***www.HamRadioAndVision.com****, which receives more than 29,000 visitors a year.*

Notes

[1]E. Hare, W1RFI, *RF Exposure and You*, (**www.arrl.org** currently out of print), use OET Bulletin 65b (1997).

[2]Evaluating Compliance with FCC Guidelines For Human Exposure to Radiofrequency Electromagnetic Fields, OET Bulletin 65b (1997), **https://transition.fcc.gov/Bureaus/Engineering_Technology/Documents/bulletins/oet65/oet65b.pdf**.

[3]P. Evans, VP9KF, Amateur Radio RF Safety Calculator, **hintlink.com/power_density.php**.

[4]K. Siwiak, KE4PT, "An All-Band Attic Antenna," *QST*, Oct 2007, pp. 33-37.

[5]Several versions of *EZNEC* antenna modeling software are available from developer Roy Lewallen, W7EL, at **www.eznec.com**.

[6]Arie Voors, 4nec2 NEC based antenna modeler and optimizer, **www.qsl.net/4nec2/**.

Raymond F. Gurney, KDØFYF

1220 Button Rock, Longmont, CO 80501; **gurneyr@hotmail.com**

Observations of Sinusoidal Variations in Signal Strength of Digital Mode Transmissions and their Causes

Measurement of the signal strength of digital mode transmissions may be a useful tool for acquiring propagation data.

Hundreds of hours of strip chart recordings of digital mode transmissions have produced unexpected results. Recordings of signal strength levels often produce sinusoidal waveforms in the 20 – 800 mHz (millihertz) range. These sinusoidal waveforms can be used to extrapolate the state of the ionosphere through which the transmissions passed. A strip chart recording of a nearby digital mode transmission (for example JT65 or JT6) is flat, as expected. We expect recordings at great distances would be similar except for path loss and interactions produced by a birefringent media, see Figures 1 and 2.

'Normal' means, as in Figure 2, that a test result taken at far distances that resembles, more or less, test results at close distances before variables in the path add ambiguity. The data shows sinusoidal patterns appearing more often than expected, see Figure 3. Faster sinusoidal waves often appear to be modulated by slower, longer wavelength waves, as seen in Figure 4. Frequencies generally range from 20 to 800 mHz with most in the range 100 to 500 mHz.

General Findings and Observations

Observations and experimentation reveal the following.

(1) Sinusoidal waveforms are not present in close-by propagation.

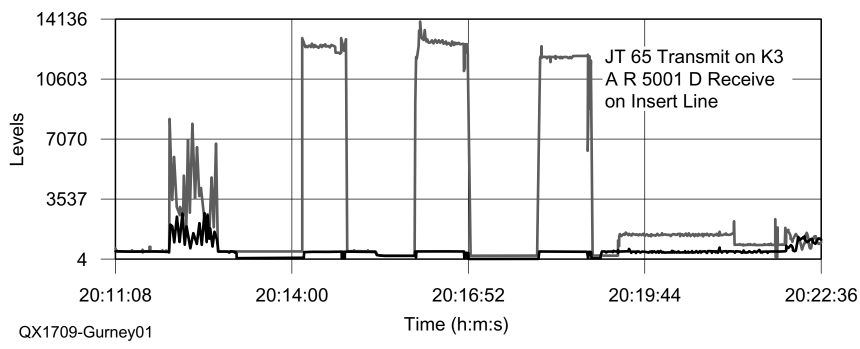

Figure 1 — Strip chart of a JT65 signal measured at close distances.

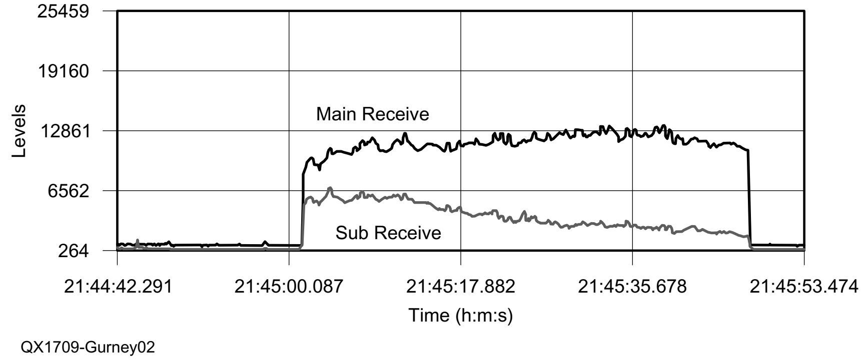

Figure 2 — Strip chart of a JT65 signal measured at the far distances.

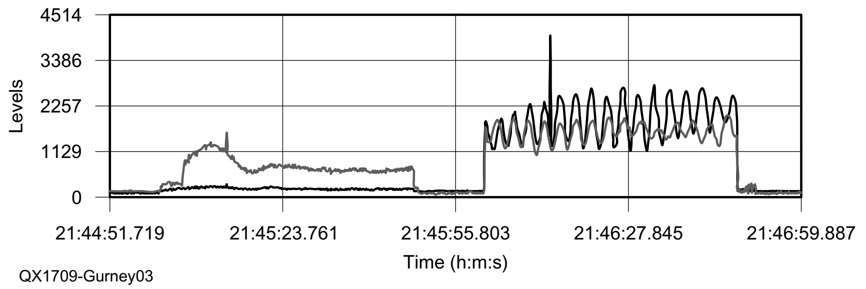

Figure 3 — Strip chart of two JT65 signals. The left hand trace represents the expected 'normal'. The right hand trace exhibits a sinusoidal variation with a frequency of approximately 352 mHz.

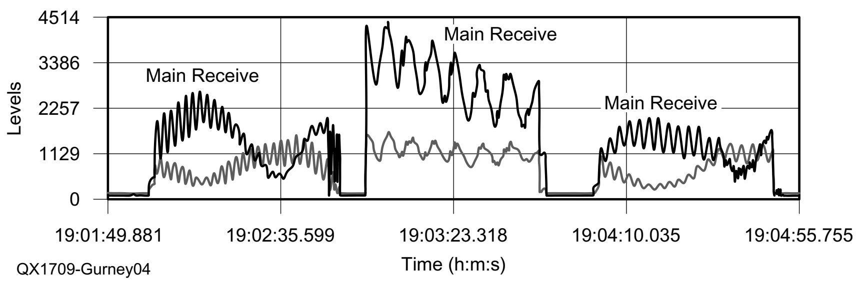

Figure 4 — Strip chart of several JT65 signals with a 492 mHz wave modulated by a 31 mHz wave, followed by a 146 mHz wave.

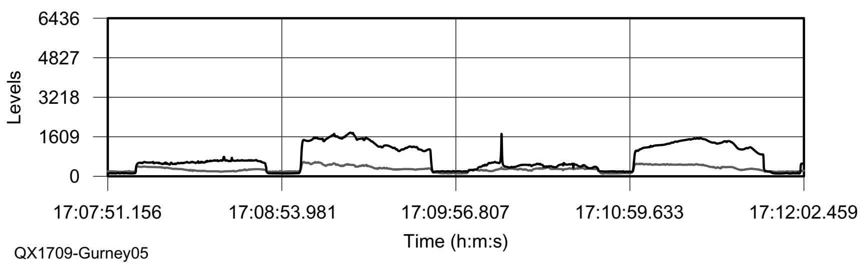

Figure 5 — Strip chart of several JT65 signals with no sinusoidal variations, that is, the expected normal looking trace.

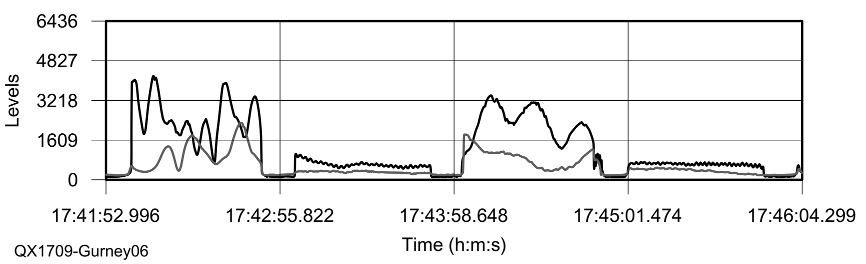

Figure 6 — Strip chart of several JT65 signals showing sinusoidal variations in both weak and strong signals.

(2) Sinusoidal waveforms are not always present in distant propagation, see Figure 5.

(3) Sinusoidal waveforms can appear in weak and strong signals, as in Figure 6.

(4) The sinusoidal waveforms change with antenna orientation. When compared to a fixed antenna, the phase relationship of the sinusoidal wave forms changes with antenna orientation. The charts of Figure 7 show a JT65 signal originating at Tewksbury, MA. On the upper chart the bottom trace is the main receiver and the top trace is the sub receiver. On the upper half of the figure, the main receiver dipole antenna is orientated north/south. On the lower half the main receive antenna is rotated to east/west. The sub receiver antenna is a fixed traveling wave antenna orientated east/west with the termination on the east end.

(5) Switching sense on a circularly polarized antenna, say, from left-hand circular polarization (LHCP) to right-hand circular polarization (RHCP), produces a nearly 3 dB change in level indicating a predominance of handedness, but with no change in pattern, see Figure 8.

(6) Sinusoidal waveforms from a transmitter at a fixed location change over time. Sinusoidal frequencies on JT65 transmissions from Gloversville, NY over a thirty minute period are shown in Table 1.

(7) Transmissions on similar azimuths and relative distance have a similar sinusoidal pattern. In Figure 9 a station in central Ohio starts a JT65 transmission one second past the UTC minute (lower trace). A second station (upper trace) with higher power starts a JT65 transmission about 12 seconds late. The second station is located in Toronto, Canada. There is difference in the two 157 mHz patterns.

(8) Sinusoidal waveforms have similar characteristics across bands and digital mode types, indicating they are not related to frequency but to some other variable.

(9) Frequency of observed sinusoidal waveforms generally fall in the range of 20 – 800 mHz, with the *mean* between 100 – 500 mHz. The patterns can be reproduced by mixing closely spaced tones, or through amplitude modulation using a mHz tone.

(10) The sinusoidal waveforms appear in north/south, east/west transequatorial propagation paths.

Doppler Shift and ULF Waves

The mechanisms involved in the

Table 1.
Time and Frequency over a 32 minute period.

Time, minutes	0	3	15	17	19	22	24	26	28	30	32	34
Frequency, mHz	570	176	761	916	508	352	229	218	124	44	37	63

propagation of high frequency radio waves have been studied since the inception of radio. Doppler variations in HF waves were noticed by Harang.[1] Alfvén waves and other perturbations have been studied for many years by many researchers. The study of the interaction between ULF waves and the ionosphere is on-going. In general, anything that causes a time dependent refractive index can cause Doppler shifts of HF signals.[2] In addition, experimentation has shown that ULF waves can be produced in the ionosphere using transmitters with slightly different frequencies.[3] More recently, HAARP and other arrays have produced similar results. In this section we will show the relationship of path and refractive index to Doppler shift. Then look at the Sutcliff-Poole[4] model that defines three mechanisms in terms of Doppler velocity that contribute to Doppler shift in HF signals due to ULF waves. Finally, we will demonstrate that the observed sinusoidal waveforms are related to the amount of Doppler shift and that some patterns are related to perturbations typically caused by ULF waves.

Over the past several decades the Amateur Radio community has devised numerous digital protocols. A few have enjoyed widespread acceptance especially those used for small signal work. For the rest of this paper we limit our discussion to JT65. The JT65 mode is perfect for observation. The frame starts one second after the UTC minute and each frame is 46.8 seconds long. Modern band plans have all the users in one spot and depending on band and conditions, transmissions are more or less continuous for extended periods of time. JT65 signals are phase continuous and of constant amplitude and there are no key clicks.[5]

Like many of the digital modes, JT65 transmits an audio waveform using an upper SSB mode. That USB signal can be represented by:

$$\begin{aligned} f_s(t) &= \mathrm{Re}\left[Z(t)e^{j\omega_C t}\right] \\ &= f(t)\cos(\omega_C t) - \hat{f}(t)\sin(\omega_C t) \end{aligned} \quad (1)$$

The modulated signal is recovered by demodulation or detection. The pre-filter term of the detected signal is:

$$f(t)\cos(\omega_C t) - \hat{f}(t)\sin(\omega_C t)\cos(\omega_C t) \quad (2)$$

This is where errors can arise. If there is a shift in phase in the base band term or the carrier term, for example $(\omega_C t + \theta)$, a distortion term $\hat{f}(t)\sin(\theta)$ appears, and the output signal can be distorted. It can be shown that frequency error results in all frequency components translated by $\Delta\omega$. The phase error results in all components being phase shifted by θ.[6]

For HF terrestrial communication the path often passes through anisotropic media. The Y_T and Y_L terms in the Appleton-Hartree formula[7] remind us that the refractive index

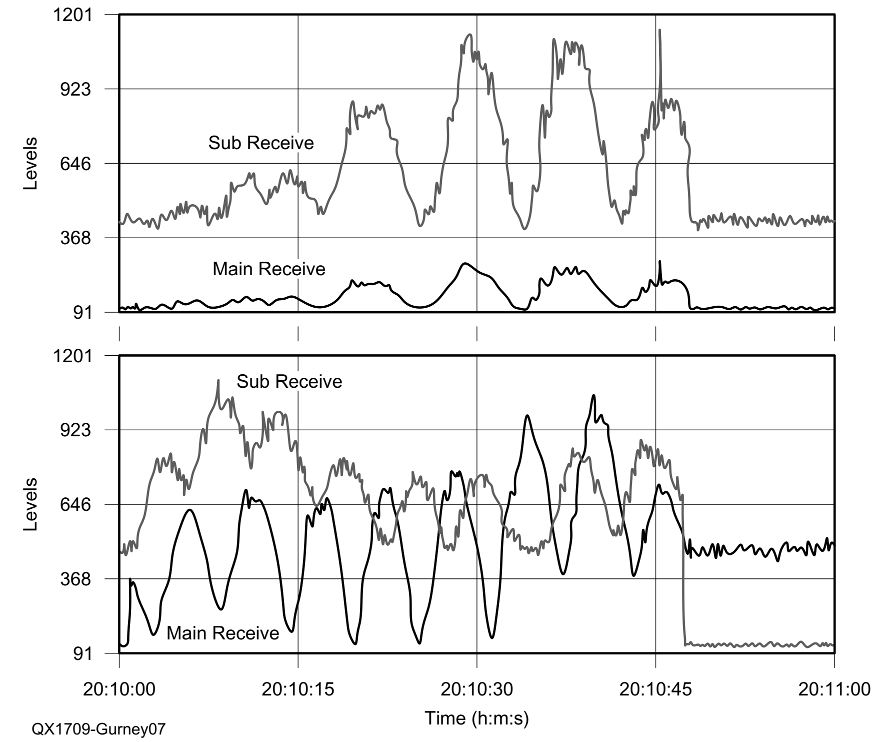

Figure 7 — Two strip charts demonstrating a phase shift with a change in dipole orientation.

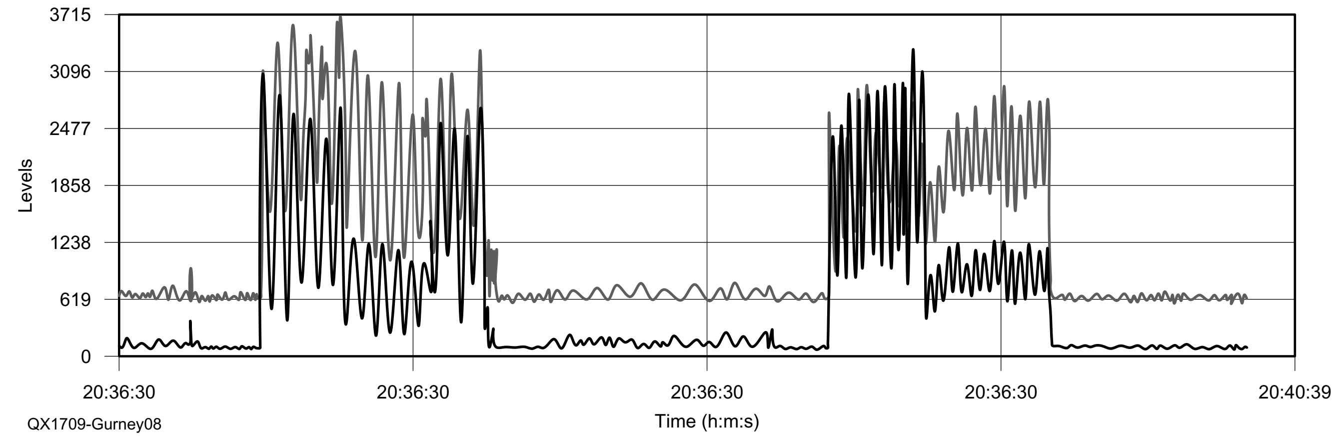

Figure 8 — Strip chart of JT65 signals received using a circular polarized antenna showing a near 3 dB loss when the sense is changed. Notice that the frequency of the sinusoidal pattern does not change.

depends on the transverse and longitudinal components of Earth's magnetic field. This tells us that a wave front that travels for time t will not be a circle. The transit time to travel between two points is defined by

$$T_p = \frac{1}{c}\int_s n\cos(\alpha)ds = \frac{P}{c} \qquad (3)$$

The phase path is

$$P = \int_s n\cos(\alpha)ds \qquad (4)$$

For the present discussion of JT65 we can ignore group path. If the phase path varies, say from variations in path length, the receive frequency will not be equal to the transmit frequency.

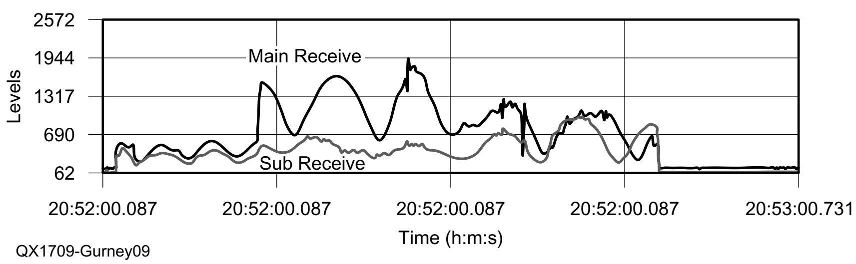

Figure 9 — Strip chart of JT65 signals the same azimuth and relative distance producing the same sinusoidal pattern.

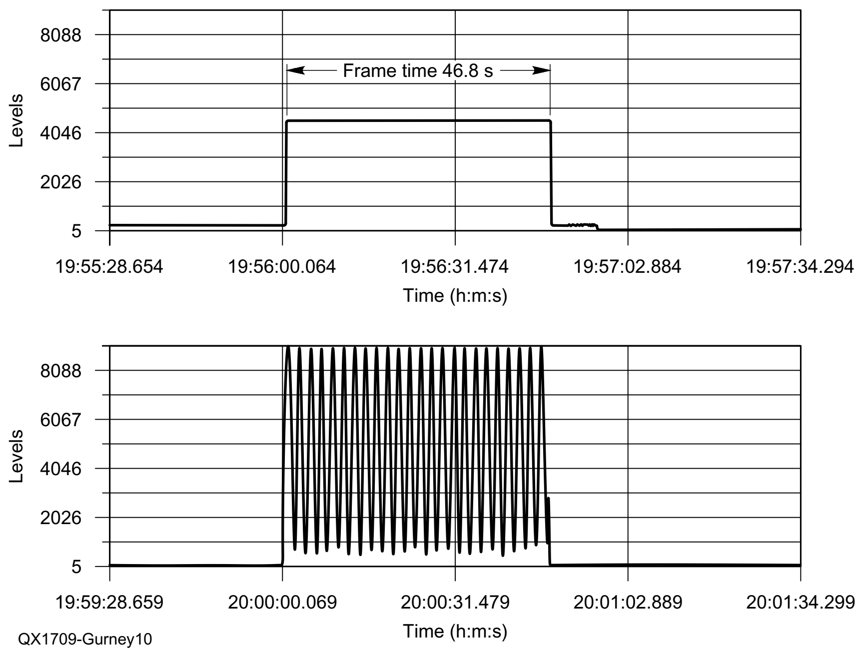

Figure 10 — Bench test demonstrating frame time and the mixing of tones with a Δf of 0.5 Hz.

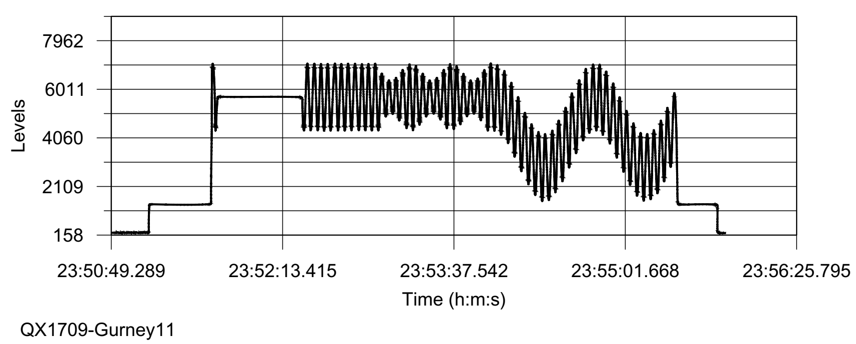

Figure 11 — Bench test of an AM modulated tone with various parameters.

$$\Delta f = -\frac{1}{\lambda_0}\frac{dP}{dt} = -\frac{f_r}{c}\frac{dP}{dt} \qquad (5)$$

where Δf is the Doppler shift in frequency, c is the speed of light and P is the phase path. This relationship is well known to those involved with vertical incidence sounders.

Bennett in his enlightening 1964 short paper[8] shows us that,

$$\frac{dP}{dt} = \int \frac{\partial \mu}{\partial t}\cos(\alpha)ds \qquad (6)$$

so that Δf can be restated

$$\Delta f = -\frac{f_r}{c}\int \frac{\partial \mu}{\partial t}\cos(\alpha)ds \qquad (7)$$

where c is the speed of light, f_r is the signal frequency, μ is the refractive index as defined by Appleton-Hartree and s is the propagation path. A more complete model of Doppler shift in HF signals due to ULF waves was developed in 1989 by Sutcliff and Poole.[9] They demonstrate, with more granularity, how sensitive the refractive indices are to changes in magnetic field and electron density. Δf is described as an effective Doppler velocity V^* with the following relationship,

$$\Delta f = 2 f_r \frac{V^*}{c} \qquad (8)$$

where $V^* = V_1 + V_2 + V_3$, c is the speed of light, f_r is the radio frequency, V_1 is the magnetic mechanism, V_2 is the advection mechanism, V_3 is the compression mechanism.

The presence of ULF wave signatures was detected by Sutcliff and Poole (Note 4). A 2001 paper by Sinha[10] reported periodicities in the 20 – 25 second range (418 – 523 mHz) with simultaneous ground magnetic data. A more recent study[11] investigated the relationship between ULF energy and HF signals using the Sutcliff-Poole model. Two sets of results were presented. The first set produced Doppler shifts in the ~ <50 – 650 mHz range with the mean between ~ <50 and 420 mHz. The second set of results produced Doppler shifts in the ~ <50 to more than 900 mHz with the mean between <50 and 420 mHz. These results are consistent with the observations presented in this paper.

If the signal arriving at the receiver has components that are Doppler shifted and components that are not shifted (or shifted a different amount), the signals will mix in the detector resulting in the observed sinusoidal waveforms. For JT65, the frequency of the sinusoid gives the value of Δf directly,

$$\Delta f = \frac{R_{2\pi}}{F_t} \qquad (9)$$

where $R_{2\pi}$ is the number of 2π rotations in the observed frame, and F_t is the frame time; see Figure 10.

The chart in Figure 11 is a bench test. Two tones, 1275 Hz at -15 dB and 1275.5 Hz at -22 dB are mixed. As the chart advances, the -22 dB tone is AM modulated by a 0.05 Hz tone. Lastly, the AM modulation is moved to the -15 dB tone. 1275 Hz is used because it is near the center of the pass band of some popular JT65 codec software.

Figures 12 and 13 are JT65 transmissions with the SSB dial frequency of 21.076 MHz. There is a correlation between them and the bench test chart. In fact, most observed sinusoidal waveforms can be reproduced with careful choice of Δf and level.

The causal mechanism that describes the observed sinusoidal waveforms can be described as Doppler shift. A practical communication is not a 'ray' but an aperture, which illuminates an area of the ionosphere. The HF signal, everywhere, breaks up into its component *X* and *O* waves and remixes in various ratios or, at times, separates into different paths. The birefringent media produces various Doppler shifts that appear as a sinusoid on the strip chart. At various times the signals are modulated by ULF waves through the mechanisms described earlier. These often have a longer wavelength and produce a characteristic signature. Careful examination of the charts show these long wavelength waves are spread across several JT65 frames, see Figure 14. The idea that more than one periodicity might be present was briefly mentioned by Rishbeth and Garriott in 1964.[12]

Conclusion

Investigation of ULF waves often involve rockets, satellites, sounders and other expensive data sources. This paper has demonstrated an observational method of determining Doppler shifts over HF signal paths, and the ability to observe perturbations in the ionosphere along those paths with the use of low cost data. The next step is to develop software to map the data to geographic area and to combine it with magnetic data such as SuperMag. With a distributed network — receivers as a sensor array — vast amounts of ionospheric data can be collected at low cost.

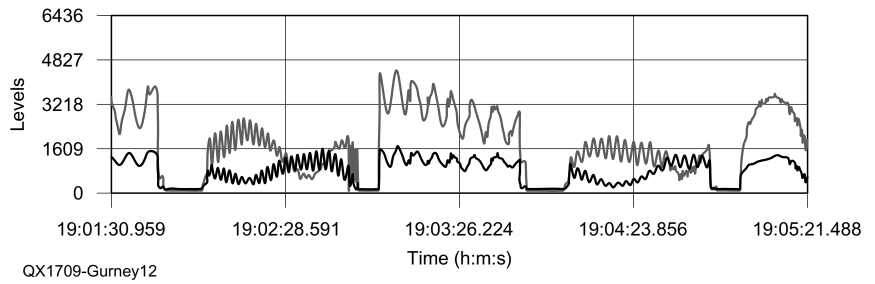

Figure 12 — Strip charts of several JT65 signals that show a correlation with the forms in Figure 11.

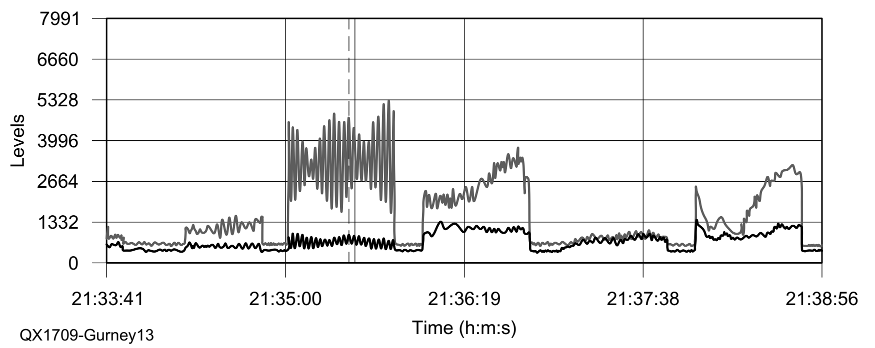

Figure 13 — Strip charts of several more JT65 signals that show a correlation with the forms in Figure 11.

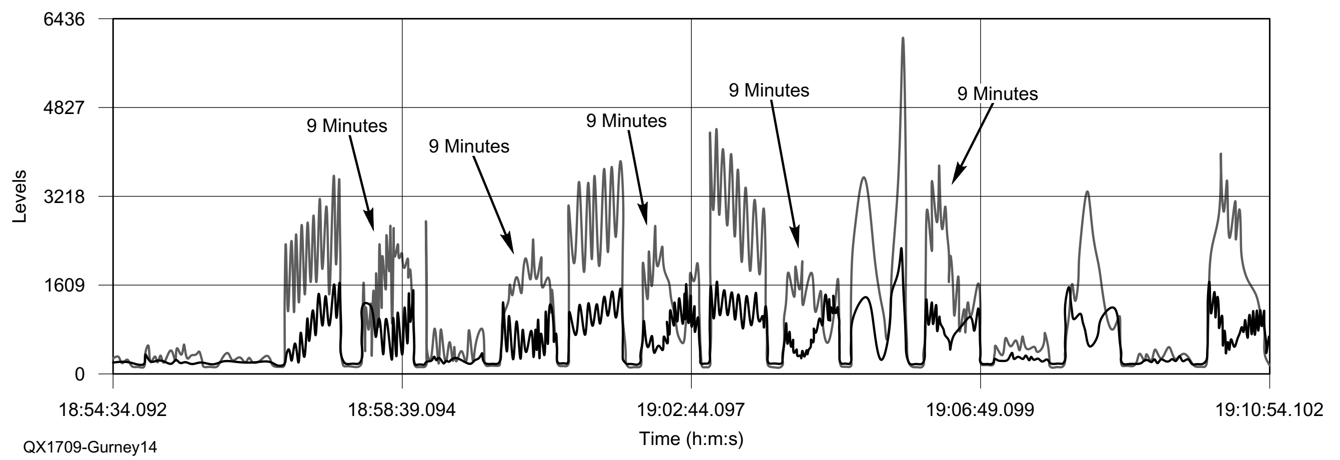

Figure 14 — Strip chart of several JT65 signals that demonstrate the continuity of longer wavelength forms over several minutes.

Raymond F. Gurney, KDØFYF, joined the Marine Corp in 1966. On returning to civilian life, he earned the AA and BS degrees in Business and Management. He also attended a one-semester class in Amateur Radio at a community college that led to his Amateur Radio license, and a subsequent life-long interested in radio. Raymond retired in 1994 from a career in the telecommunications and electronic industries. He has since pursued an interest in radio.

Notes

[1]L. Harang, "Pulsations in an Ionized Region at Height of 650 – 800 km During the Appearance of Giant Pulsations in the Geomagnetic Records", Terr. Mag. Atmos. Electr., 44, 17, 1939.

[2]A.T. Waterman, J. W. Strobehm, "Reflection of radio waves from undulating tropospheric layers", Journal of Research of the National Bureau of Standards, Vol. 67D, No. 6, Nov. – Dec. 1963.

[3]S. Ganguly, "Experimental observation of ultra-low-frequency waves generated in the ionosphere", Nature 320, 511. – 513,10 Apr. 1986, doi:10.1038/320511b0.

[4]P. R. Sutcliffe and A. W. V. Poole, "Ionospheric Doppler and electron velocities in the presence of ULF waves", J. Geophys. Res., 94, 13 505–13 514, 1989.

[5]J. H. Taylor, K1JT, Version 1.7.0 WJST-X User Guide, **www.physics.princeton.edu/pulsar/k1jt/wsjtx-doc/wsjtx-main-1.7.0.html**.

[6]A.J. Wilkinson, EEE3086F Signals and Systems II 2014, 508, p1, Apr. 14, 2014.

[7]A .Bourdillon, presentation notes, "Ionospheric radio wave propagation Part 1", L'Aquila, 7-12 April, 2008, **ftp://geosp-server.aquila.infn.it/isss/Geom&Ion08/bourdillon.pdf**.

[8]A. Bennett, "The calculation of Doppler shifts due to changing ionosphere", J. Atmos. Terr. Phys., 29, 887 – 891, 1967.

[9]P. R. Sutcliffe and A. W. V. Poole, "The relationship between ULF geomagnetic pulsations and ionospheric Doppler oscillations: Model predictions", Planet. Space Sci., 38, 1581 – 1589, 1990.

[10]A. K. Sinha, B. M. Pathan, R. Rajaram, and D. R. K. Rao, "Low frequency modulation of transionospheric radio wave amplitude at low-latitudes: possible role of field line oscillations", Annales Geophysicae (2002) 20: 69 – 80 European Geophysical Society 2002.

[11]C. L. Waters, T. K. Yeoman, M. D. Sciffer, P. Ponomarenko, and D. M. Wright, "Modulation of radio frequency signals by ULF waves", Ann. Geophys., 25, 1113 – 1124, 2007, **www.ann-geophys.net/25/1113/2007/ European Geosciences Union 2007**.

[12]H. Rishbeth and O. K. Garriott, "Relationship between simultaneous geomagnetic and ionospheric oscillations", Radio Sci., 68D, 339 – 343, 1964.

Joseph M. Haas, KEØFF

4301 Yellowstone Dr., Taylor, TX 76574: **joeh@rollanet.org**

Using Direct-Current Relays at Lower Coil Voltages

Here's how and why relays can be operated significantly below their nominal rated voltage.

Relays are common electronic components found in many devices. Solid-state switches and circuits have supplanted these devices in many ways, but they remain viable as a simple and cost-effective way to switch power and signals, RF in particular, with a relatively low-power control signal. However, a common problem for the hobbyist or experimenter is that one might locate the perfect relay for an application, only to find that the relay coil voltage is something like 24 V dc or 28 V dc with only 12 V dc power available. While others have explored the use of these higher voltage dc coils, the news of their successful application often falls on deaf ears. I was once one of the "deaf ears".

The Lowly Coil

The central element of a traditional relay is the coil. It is simply an electromagnet, and its sole purpose is to create a magnetic field of sufficient strength to attract a mechanical lever. This lever, in turn, moves electrical contacts causing a change to the electrical connections. One must be aware of the effect of the collapsing magnetic field when the relay is de-activated. However, for dc relays, this is typically addressed by a reverse-biased diode to limit the voltage-spike that can result when the source of coil current is removed. For much of my career, I simply did not give relays much more attention than that.

My cursory treatment omits many of the important subtleties of coils and relays in general, but for most applications, this level of design scrutiny is sufficient. The shape and construction of the coil is the same as with any inductor intended for ac circuits.

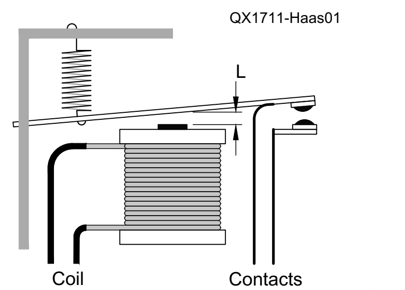

Figure 1 — A simplified relay system.

These coils have inductance, resistance, and capacitance, just as with any other inductor. The resistance tends to dominate the behavior of the device, so inductance and capacitance can usually be ignored by the designer.

It is very easy to forget that relay coils are current-mode devices since most relays have a specified operating voltage for the coil. The operating voltage divided by the coil resistance is the current responsible for the operation of the relay. Specifying the coil voltage is tantamount to specifying its current. Still, forgetting that the current is the real star can cause one to miss some of the core concepts at work here.

Anatomy of a Relay

Figure 1 shows drawing of a simple relay. The coil current produces a magnetic field along the coil axis. If strong enough, this field will attract the nearby ferromagnetic lever, which will move towards the coil. As the lever moves closer to the coil armature (the metal core of the coil), the force on the lever increases. This force is roughly an exponential function of the air gap, *L*, which is the distance between the coil armature and the nearest point of the lever. Once the lever contacts the armature, the force required to dislodge the lever is typically orders of magnitude greater than the force that initially started moving the lever from its rest position. This concept is key to the method employed in this article.

Some Definitions

Pull-in voltage (or current) V_{pi} or I_{pi} is the minimum voltage or current that will cause the relay to engage from its rest position. When presented in a datasheet, this value generally includes allowances for variances in the relay construction and environment.

Switching delay t_d is the time that it takes for the electrical contacts to close in the engaged position once the relay coil is activated. Note that this is not the time it takes for the lever to reach the full extent of its movement.

Pull-in delay, t_{pi} is the time that it takes for the relay lever to contact the coil armature.

Pull-out delay t_{po} is the time that it takes for the relay lever to reach the open position after the relay supply is removed. This is greatly impacted by the method(s) applied to the coil to address the back-EMF that results when the supply voltage is removed.

Contact bounce time t_b is the time duration that the contacts bounce against each other one or more times for each activation event.

Holding voltage (or current) V_h (or I_h) is the minimum voltage (or current) that will hold the relay engaged once it has been pulled in and settled.

The pull-out delay t_{po} can be of concern for systems that require a fast transition between states. There isn't much that can be done to decrease the pull-in delay t_{pi} since this is a function of applied current, coil inductance, and the mass of the lever assembly, among other effects. However, t_{po} can be controlled to some extent. In general, anything that increases the coil voltage during the field collapse duration will reduce the time it takes for the field to fully collapse. The diode that is commonly used to control the coil voltage during field collapse will result in the longest possible t_{po}. Reducing t_{po} can be accomplished by using a Zener diode in series with a switching diode, oriented so that the Zener diode is reverse-biased when the switching diode is forward biased. The main concern is limiting the reverse coil voltage to below the point of failure for the coil-drive electronics.

While coil current is the ultimate quantity of interest, we generally reach this current via an applied voltage. Voltage control is easier than current control. Thus, the calculations discussed herein will deal in voltages and coil resistance. However, the resistance of the relay coil varies with temperature, so the *voltage – resistance* relationship is valid only at temperatures where the resistance value is known. Since copper resistance increases with temperature, coil current decreases with temperature, which may need to be considered if a datasheet is not available, or if operating outside of the specified temperature range

Dynamic Relay Current

It may be a surprise that a typical relay holding current is rarely more than half the minimum pull-in current. In fact, it is often less than half the minimum pull-in value. This is the effect we will exploit. By varying the relay current such that at least the minimum operating current is applied for enough time to engage the relay, followed by a reduction of this current to something greater than the holding current, the relay can be reliably operated with less than full rated current (or voltage). Here's how it works.

One popular way to accomplish a dynamic current drive is to carefully switch an appropriately sized capacitor in series with the relay coil.[1] By pre-charging this capacitor before switching, we can induce a voltage across the relay coil that is greater than the pull-in voltage. If the capacitor can hold sufficient charge — determined by the value of its capacitance, the voltage of the pre-charge, and the load current — the relay will engage. Once the capacitor is discharged, we switch in the available dc voltage supply, which must be greater than the hold-in voltage, to keep the relay engaged.

This can be accomplished with just four components in addition to the relay and its associated clamping diode. We need a DPDT switch, a capacitor, a diode, and a resistor as seen in Figure 2. In the OFF state, the supply voltage is presented to the capacitor through the current limiting resistor *R*. The value for *R* should be large enough to prevent the current rating of the switch contacts from being exceeded, but small enough that the capacitor charges quickly. I generally set *R* to 100 Ω as a starting point.

Once the switch is moved to the ON position, the charged capacitor appears in series with the relay coil such that its voltage adds to the supply voltage to provide nearly double the supply voltage across the relay coil. The capacitor then begins to discharge, and will provide current to the relay coil until its voltage dips below the forward voltage drop of diode *D*1. This diode does not clamp the coil transients, a separate diode across the relay coil is needed. From this point on, the voltage presented to the relay coil is the supply voltage minus the forward voltage drop of *D*1.

The Capacitor Value

The capacitor value must be large enough so that most of its charge is still retained after discharging for the duration of pull-in-delay. This can be calculated effectively using the equation for the discharge of a capacitor, the coil resistance, and the holding voltage of the relay. Figure 3 shows a plot of capacitor voltage versus time for a capacitor discharging through a resistance *R*. The time constant *RC* is represented by τ.

You can estimate *RC* from Figure 3. Subtract the supply voltage V_{dd} from the value of V_{pi}, then divide by V_{dd}. This ratio sets the location of interest on the *y*-axis. Read the *x*-axis value where the exponential curve intersects this value to determine the number of time-constants that are involved. Set this equal to t_{pi} and solving for *C*.

Consider a relay with a V_{pi} of 18 V, a coil resistance of 3 kΩ, a t_{pi} of 15 ms, and an available minimum V_{dd} of 10 V, then

$$\frac{V_{pi} - V_{dd}}{V_{dd}} = 0.8$$

The *x*-axis value extracted from the chart is about 0.24 *RC*. Setting this equal to t_{pi} and solving for *C*, we get 0.24 $RC = t_{pi}$, so

$$C = \frac{0.015}{0.24 \cdot 3000} = 21\ \mu\text{F}.$$

As can be seen in Figure 3, the capacitor is essentially discharged after 5 time constants. The formula is,

$$V_c(t) = V_{dd} e^{t/RC}\ ;\ t \geq 0.$$

As long as $[V_c(t) + V_{dd}] > V_{pi}$ for a duration of at least t_{pi}, the capacitor is of sufficient value to pull in the relay. $V_c(t)$ is the capacitor voltage with respect to time, and V_{dd} is the available supply voltage. Solving for *C* gives,

$$C = \frac{-t_{pi}}{R \ln\left(V_c(t)/V_{dd}\right)} \quad \text{F}$$

where $V_c(t)$ is set to $V_{pi} - V_{dd}(min)$ to find a minimum value for *C*. For example,

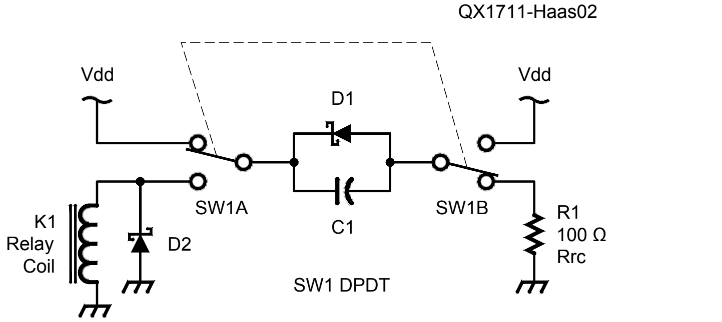

Figure 2 — A simple step-up relay drive circuit. The switch is shown in the "off" position.

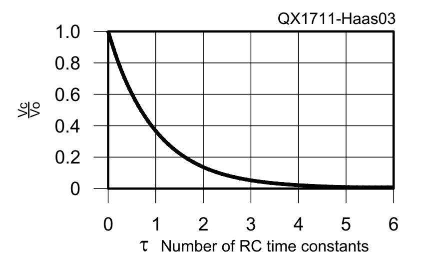

Figure 3 — This voltage versus time-constant plot for a capacitor may be used to graphically determine a value for *C*.

considering a relay with a V_{pi} = 18 V, a coil resistance R = 3 kΩ, t_{pi} = 15 ms, and $V_{dd}(min)$ = 10 V, the minimum capacitance needed is,

$$C = \frac{-0.015}{3,000\ln\left((18-10)/10\right)} = 22.4\ \mu F$$

which agrees well with the graphical estimate.

The relay coil exhibits inductance and parasitic capacitance. The parasitic shunt capacitance is often negligible — in the hundreds of pF at most. The inductance is usually significant — as much as several millihenry. The effect of this inductance is to resist the change in coil-current flow. When activating the relay, this will delay the buildup of the magnetic field and is reflected in the pull-in time t_{pi}. Thus, if the t_{pi} is specified or measured, the inductance need not be considered as this effect is already taken into account.

More capacitance is certainly better, at least that seems to be the popular belief. However, one must consider the charge time of the capacitor for the application at hand. In the previous example, A 20,000 μF capacitor would certainly be large enough, but its size and charge time — about 10 s — might make a noticeable impact to its application. Doubling or tripling the calculated value might be appropriate, but going beyond that is only warranted if one can tolerate or mitigate the charge-time implications.

Lacking a Data Sheet

What if there is no data sheet available for a given relay? This can be a common problem since many desirable relays can be found in surplus equipment. Manufacturer's data may be scarce or difficult to find. Fortunately, it is relatively easy to measure the required parameters.

The coil resistance can be measured with a DMM. If a clamping diode is attached connect the DMM such that the diode is reverse biased by the meter. If several of the same model relays are available, an average can be calculated. This is good practice to make sure that the relay-under-test isn't defective.

Copper resistance varies by temperature, and can be calculated for a given temperature T in °C,

$$R(T) = R(T_0)\left(1 + \alpha\Delta T\right)$$

where T_0 is the reference temperature (usually 25 °C), $\Delta T = (T - T_0)$ and α is the temperature coefficient (per °C). For copper α= 0.004 per °C near 25 °C.

If the maximum operating temperature is known, this can be used to calculate the worst-case coil resistance. Most coils are wound with copper wire. A moderate temperature of 50 °C is realistic for open air mobile operation. For resistance measurements made at 25 °C, this results in a 20% increase of coil resistance at 75 °C. Calculate ΔT for the maximum expected temperature rise and this result will be used to scale readings to reach a safe over-temperature-range result. This will decrease the coil current for a given applied voltage and needs to be considered to ensure proper operation at the worst-case temperature.

Pull-in delay t_{pi} is more difficult to measure directly, but it can be estimated by measuring t_d, the time to contact closure. Doubling t_d should be a reasonable estimate for most relay geometries. This assumes that the actuation lever moves about half of its travel before activating the relay contacts. In truth, the relay lever continuously accelerates as it moves toward the armature, so this approach features a built-in "fudge-factor". If the relay mechanism is visible, this can be easily verified. For mechanisms that are not visible, a 3x value for t_d should provide more than enough margin to cover the uncertainty of the measurement.

A function generator that can achieve sub-hertz frequencies is a useful tool for this measurement. An NPN or MOSFET driver, 2-channel oscilloscope, and a variable power supply are also needed, see Figure 4. The variable supply needs to support the normal coil voltage, or must achieve at least half the rated coil voltage and allow for another offset supply to be added to achieve the desired coil voltage. Connect the transistor base — a 1 kΩ series resistor is needed for an NPN transistor — or gate to the function generator, the transistor emitter or source to ground, and the collector or drain to one side of the coil. A reverse biased diode across the coil is also needed to protect the drive circuits. Connect the other coil terminal to the power supply and set the voltage to the nominal coil voltage. The common terminal of the relay contacts is connected to ground. Connect one channel of the oscilloscope to the function generator output, and the other to the normally open terminal of the relay contacts. The normally open connection also needs a pull up resistor, 10 kΩ is a reasonable value, to either the coil supply voltage, or some other convenient power source that shares the same ground as the relay common terminal.

Operate the function generator at about 1 Hz, with a square-wave output that varies from ground to a positive voltage level sufficient to saturate the switching transistor. Observe the two waveforms on the oscilloscope and adjust the timing and triggering to produce several horizontal divisions of separation between the rising edge of the function generator, and the falling edge of the relay contact signal. Measure the time delay between these edges. Observe this time delay over multiple edges and select the largest value. Some relays will exhibit a noticeable difference in the timing

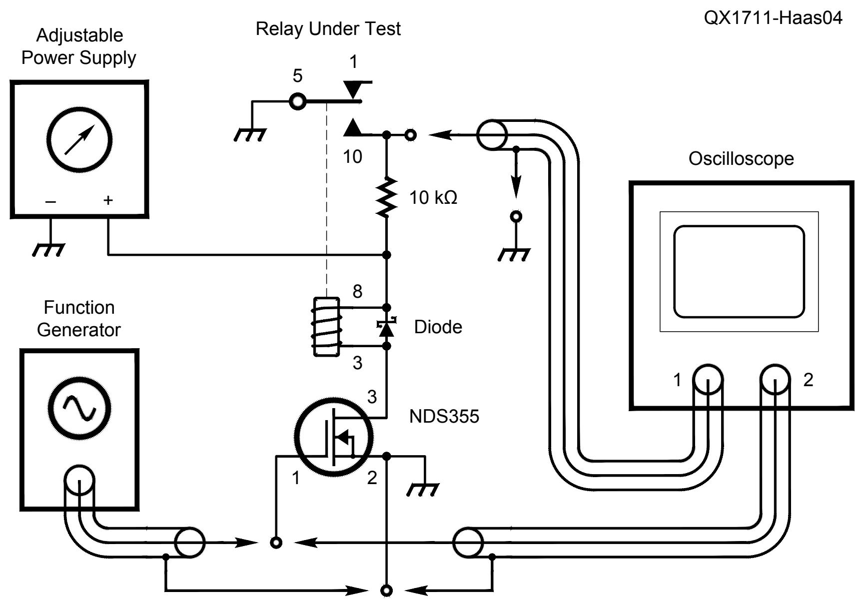

Figure 4 — Test setup for measuring critical relay parameters.

of these edges. This value is the nominal t_d of the relay. If bounce is observed, use the first edge to establish the timing value. The bounce delay t_b may be of interest to the end-use application, so make a note of it for later reference.

The same setup used above is also used to measure V_{pi}. Reduce the coil supply voltage slowly until the relay no longer engages. Increase the voltage until the relay begins to operate normally and record this as V_{pi1}. Then, repeat the t_d measurement. Note the values for voltage and timing. This value for t_d will generally be larger than the one determined at nominal voltage and it is this value that should be used in calculations. The pull-in voltage is,

$$V_{pi} = V_{pi1}\left(1 + \alpha \Delta T\right)$$

The undisturbed V_h (no mechanical or thermal changes) can be measured by forcing the relay to engage at nominal voltage and slowly lowering the coil voltage until the relay contacts release. This measurement does not use the function generator switch; simply ground the relay coil connection that was connected to the function generator driver. This value is generally not of much use except to get an idea of the ideal lower limit for V_h. A practical value for V_h requires the relay to be subjected to vibration and temperature extremes that represent real exposure conditions. Unfortunately, the equipment to produce these stimuli are generally out of reach of most experimenters. Controlled shocks can be induced to try to release the relay lever, but these may easily miss a real-world target without careful planning and execution.

Generally, if the undisturbed V_h value is less than 75% of the rated coil voltage value, the relay should reliably hold in the engaged position using the techniques described here. Shock mounting the relay can also be a risk mitigation to further ensure that the relay not disengage improperly.

Adding Digital Control

A simple way to implement this control methodology is to use a DPDT switch to execute a dynamic relay current control circuit, but this simplicity comes at a cost. First, the circuit is not directly transferrable to digital control unless a second relay with a coil voltage that is compatible with the available V_{dd} is used. Second, the switch version must be physically cycled to cause relay activation. If power is applied with the switch in the ON position, the relay will not activate since the capacitor will be discharged, with no way to charge it until the switch is manually returned to the

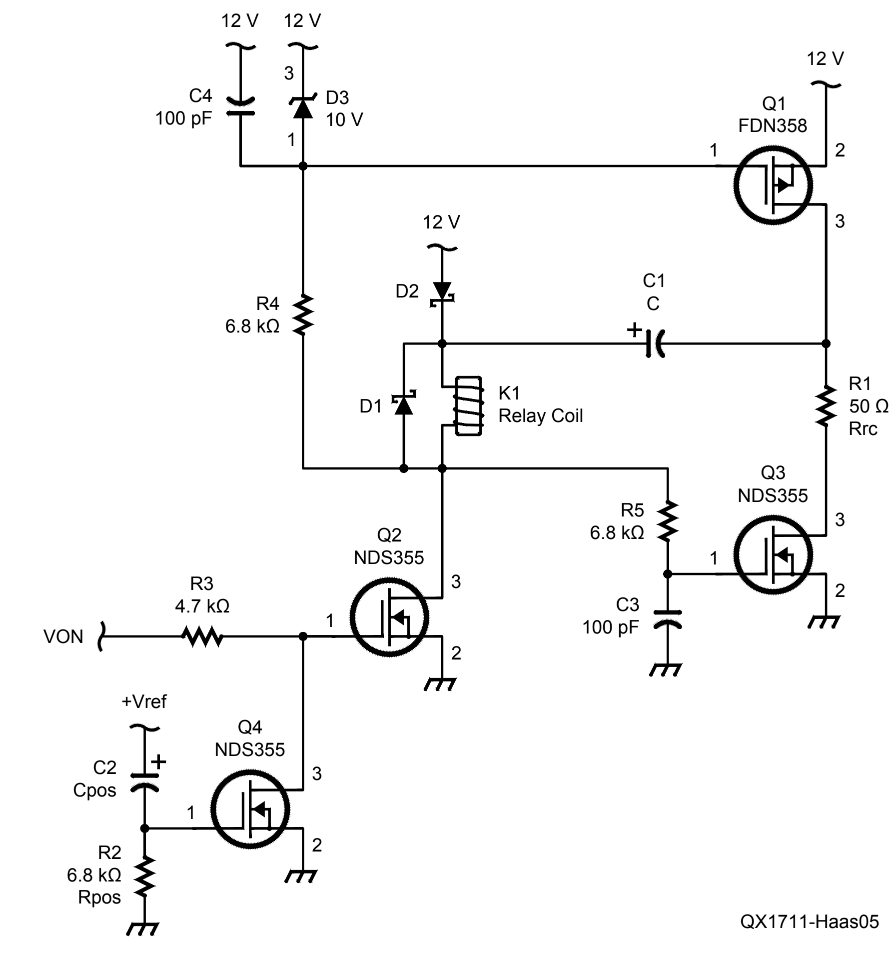

Figure 5 — Digitally controlled relay drive circuit. The $+V_{ref}$ is a regulated voltage reference. An unregulated reference may be used, but the highest likely V_{ref} voltage must be used to calculate the value of C_{pos}.

de-activated position long enough to charge the capacitor.

Another relay can be used in place of the DPDT switch, but it just seems wrong to use a relay to activate another relay. However, a solid-state switching circuit can also accomplish the required task with just a few added components. This is implemented in Figure 5.

This circuit has several features. Some of these features can be eliminated for some applications, while others are central to the basic function. The transistors $Q1$ and $Q2$ form the core of the relay switch. When $Q1$ and $Q2$ are off ($V_{ON} < 0.5$ V) no current flows in the coil. $D2$ is forward biased which allows $V_{dd} - V_{D2}$ to appear at the bottom of the relay coil which turns $Q3$ on. MOSFETs are preferred since they conduct only a small leakage current which greatly minimizes the off current that flows through the relay coil. This allows C to charge via R_{rc}. A minimum charge time of $5R_{rc}C$ is sufficient to bring C reasonably close to full charge.

V_{ON} drives $Q2$ and when this voltage is above about 3 V, $Q2$ is guaranteed to turn on. This grounds the coil, turns off $Q3$ and turns on $Q1$. With $Q1$ on, the voltage at the top of the relay coil becomes the sum of V_{dd} and the voltage across the capacitor, which is essentially V_{dd} - $V_{D2}(fwd)$. Current then flows through the coil and C begins to discharge. As C discharges, V_c drops as does the voltage across the coil. When V_c drops to about $V_{D2}(fwd)$, $D2$ begins to conduct and the voltage across the coil settles to $V_{dd} - V_{D2}(fwd)$. When V_{ON} is reduced below 0.5V, the coil turns off, and the process repeats.

$Q4$ represents a power-on-set circuit. When power is first applied, C_{pos} begins charging which briefly turns on $Q4$. This forces the circuit to keep the relay in the OFF mode while C charges. For most small FETs satisfying the inequality,

$$5R_{rc}C \le R_{pos}C_{pos}\ \ln\left(1 - \frac{V_{gs}(th)}{V_{ref}}\right)$$

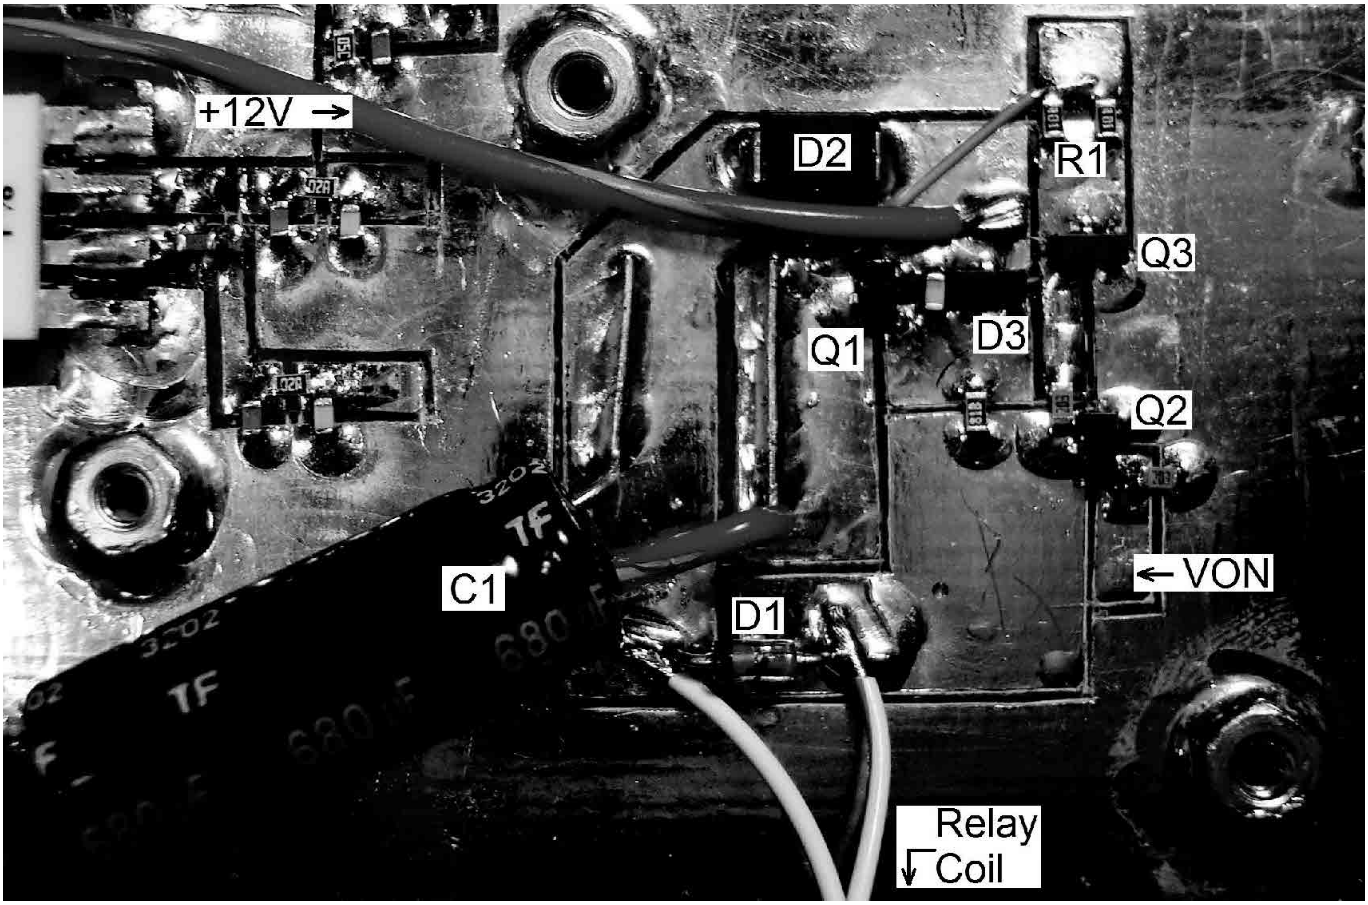

Figure 6 — Photo of an implementation of the circuit of Figure 5. The power-on set circuit was not use in this implementation. The area surrounding the circuit is ground.

will ensure that the circuit will reliably activate the relay if V_{ON} is activated before (or at the same time as) V_{dd} is applied and C has a chance to charge. $V_{gs}(th)$ is the threshold voltage of the MOSFET. Of course, this feature can be eliminated if V_{ON} will never be activated before power is applied to the relay circuit.

Figure 6 shows an implementation of this circuit using predominantly SMD components. It was implemented to drive a DowKey RF relay that requires 24 V dc, where V_{pi} = 18 V, T_{pi} = 0.025 s, R_{coil} = 250 Ω; Applying Equation (1), C = 448 μF rounded up to 680 μF. The core of the circuit is relatively compact, even with generous routing rules. The capacitor $C1$ takes up much of the circuit volume.

Conclusion

This treatise illustrates that operating a relay at lower than specified coil voltage has a sound basis in physics, and is suitable for some applications. This method reduces the power consumed by the relay coil by roughly a factor of 4 when V_{dd} is roughly half the rated coil voltage, it can have a noticeable impact when operating from battery power with high ON duty cycles. I've used this technique on several occasions, and I'm glad to have finally taken the time to examine the lowly relay in greater detail.

Joseph M. Haas, KEØFF, holds an Amateur Extra class license, and is an electrical engineer with a BSEE from the University of Missouri, Rolla, now the Missouri school of Science and Technology. He works in the MIL-AERO industry. He has worked in avionics design for 15 years, oil-field instrumentation design for 3 years, and semiconductor process equipment design for 7 years. Joseph started tinkering with electronics at the age of 6, and was first licensed as KAØGPZ just over 40 years ago. Most of his designs, both professional and hobby, are mixed-signal in nature and revolve around a microcontroller performing as a control system and/or data converter involving signals from dc up to a few megahertz. He has also designed a few RF circuits in the 2.5 GHz range and has experimented at 10 GHz. Joseph maintains a projects page at ***www.rollanet.org/~joeh/projects/****.*

Notes

[1]Didier A. Juges, KO4BB, "How to Operate 24 V relays on 12 V", **www.ko4bb.com/ham_radio/Projects/24V_Relays/**, (accessed Sep 11, 2017).

Stefan Scholl, DC9ST

Malvenweg 3, 89160 Dornstadt, Germany: **dc9st@panoradio-sdr.de**

The Panoradio: A Modern Software Defined Radio with Direct Sampling

This SDR captures up to 100 MHz bandwidth at once and provides an extra-wide panorama view seamlessly from dc to VHF, as well as direct sampling of the 70-cm band.

Direct sampling receivers are considered as the ultimate concept of software defined radio (SDR) as they digitize the analog signal right after the antenna with a minimum of analog processing. This is a desirable approach, since digital signal processing has many advantages over analog signal processing.

In recent years several developments of direct sampling receivers occurred in Amateur Radio, such the “All-Digital Transceiver” from James C. Ahlstrom, N2ADR[1], the well-known WebSDR, from Pieter-Tjerk de Boer, PA3FWM[2] , the HPSDR project[3], SDR version of the Red Pitaya[4] of Pavel Demin, and many more including numerous commercial products. The sampling rates of many of these SDRs are usually in the range of 60 to 130 MHz and cover frequencies in the short wave band from 0 to 30 MHz or even up to 65 MHz. An overview with a selection of direct sampling receivers in Amateur Radio is shown in Table 1.

The direct sampling Panoradio SDR

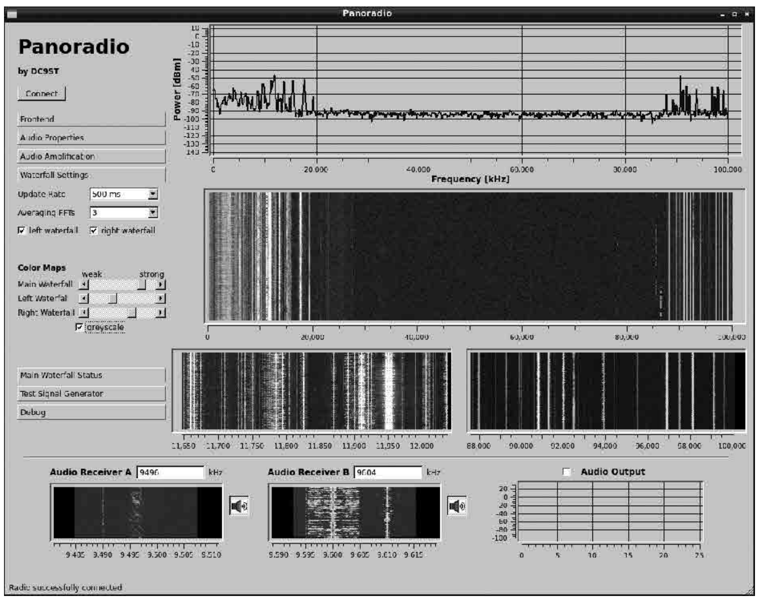

Figure 1 — Panoradio GUI shows the main waterfall in the full span from 0 to 100 MHz. Short wave signals occupy only a fraction at the left of the complete span. Signals above 88 MHz are FM broadcast stations and the sporadic emissions below 88 MHz belong to emergency services. The small left waterfall shows the 15 m broadcast band. The small right waterfall zooms on the VHF broadcasting signals.

Table 1.
Some examples of direct sampling receivers in Amateur Radio.

Direct Sampling SDR	*AD Converter*	*Sampling Rate (Msps)*	*Frequency Range (MHz)*
N2ADR All-Digital TRX (note 1)	ADS 5500 (14 Bit)	123	0 – 61
WebSDR (note 2)	LTC 2216 (16 Bit)	78	0 – 30
HPSDR (note 3)	LTC 2208 (16 Bit)	130	0 – 65
Red Pitaya SDR (note 4)	LTC 2145 (14 Bit)	125	0 – 50
Panoradio SDR	AD 9467 (16 Bit)	250	0 – 100, 425 – 445

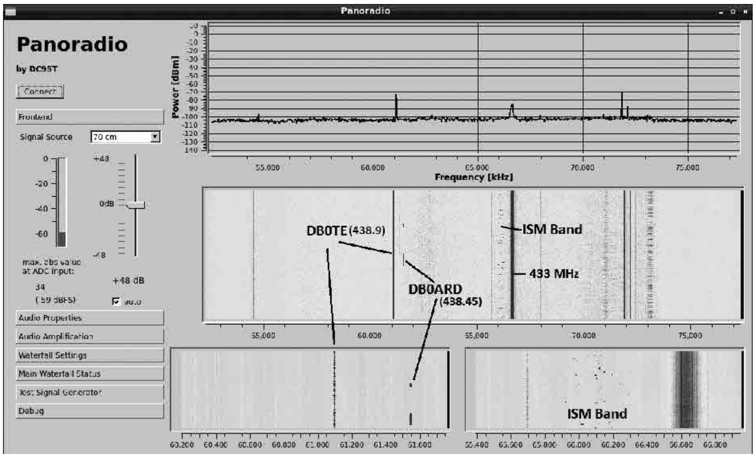

Figure 2 — Experimental reception of 70-cm band signals with undersampling shows various signals, such as the 433 MHz ISM band and two 70-cm band ham repeaters. Scales show the aliasing frequencies.

Figure 3 — Stand-alone operation with mouse and monitor connected directly to the SDR.

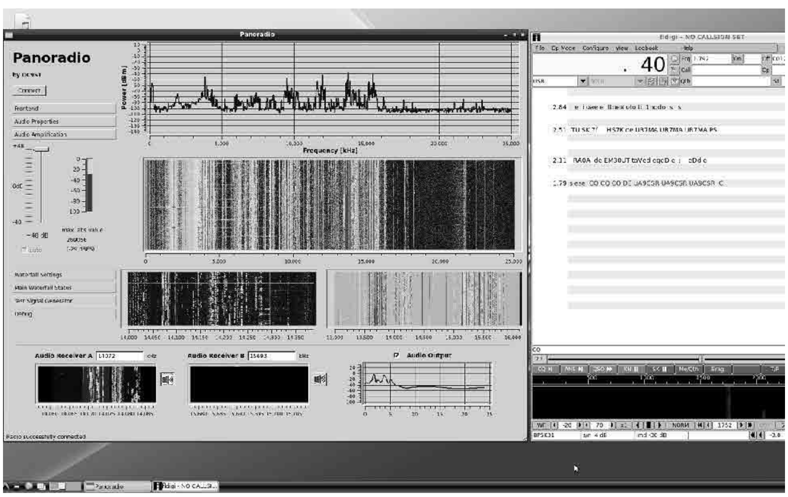

Figure 4 — Demodulation of PSK31 signals with Fldigi running directly on the SDR hardware.

receiver presented here digitizes the signal with a very fast 250 Msps sampling rate with 16 bits. It also has an experimental mode to receive UHF signals in the 70-cm band (425-445 MHz) using undersampling. The receiver has a unique panorama function. It captures and displays aliasing-free signals from 0 to 100 MHz in three zoomable waterfall displays for spectrum monitoring. These panoramic displays can independently zoom in to a bandwidth of 6 kHz and thus deliver a minimum resolution down to 7 Hz. They allow the observation of different bands in the complete frequency range from medium wave to VHF simultaneously. For example it is easy to watch signals in the 80 m, 10 m and 4 m band at the same time (Figure 1). A screenshot of the GUI during experimental reception of the 70-cm band can be seen in Figure 2.

In addition to the panoramic displays, two independent audio receivers tunable from 0 to 100 MHz are available. Each audio receiver covers a bandwidth of 22 kHz (with optional filters for 6 or 2.4 kHz), and displays the signal spectrum in additional small waterfalls. The two audio signals can be multiplexed or mixed prior to SSB demodulation using a Weaver demodulator.

The SDR is a standalone embedded system, where mouse, monitor and keyboard can be attached directly for full operation as shown in Figure 3. The hardware is built around the Zedboard, a modern embedded hardware platform. The Zedboard contains the Xilinx Zynq device, that combines a FPGA for fast data processing with a dual ARM A9 processor core running a Linux operating system. With the operating system running on the Zedboard, the software for controlling the radio and demodulation software, such as Fldigi[5], runs directly on the SDR (Figure 4). The inside view of the Panoradio is depicted in Figure 5.

The Panoradio is intended as a demo, that shows what is possible with today's technology in AD conversion and signal processing hardware. It is an open source project, so the design files can be accessed from the project website: **panoradio-sdr.de**.

The Panoradio consists of three parts: an analog front end, an AD converter and digital signal processing, which this paper presents in detail in the following sections.

Basics of Direct Sampling

First we briefly introduce and comment on direct sampling.

Direct Sampling Receivers

In a SDR receiver the AD converter marks the boundary between analog and digital signal processing (DSP). Historically the AD converter moved from the end of the reception path (where bandwidth is low) further towards the front (where bandwidth is higher). Doing so, more and more analog circuits are replaced by DSP. This is desirable, because DSP is superior to analog processing. The ultimate concept is "direct sampling", sometimes called "all-digital receiver", or "digital radio". Then the antenna is more or less directly connected to the AD converter with no analog down conversion. Usually an analog (anti-aliasing) filter and sometimes an analog preamplifier are included.

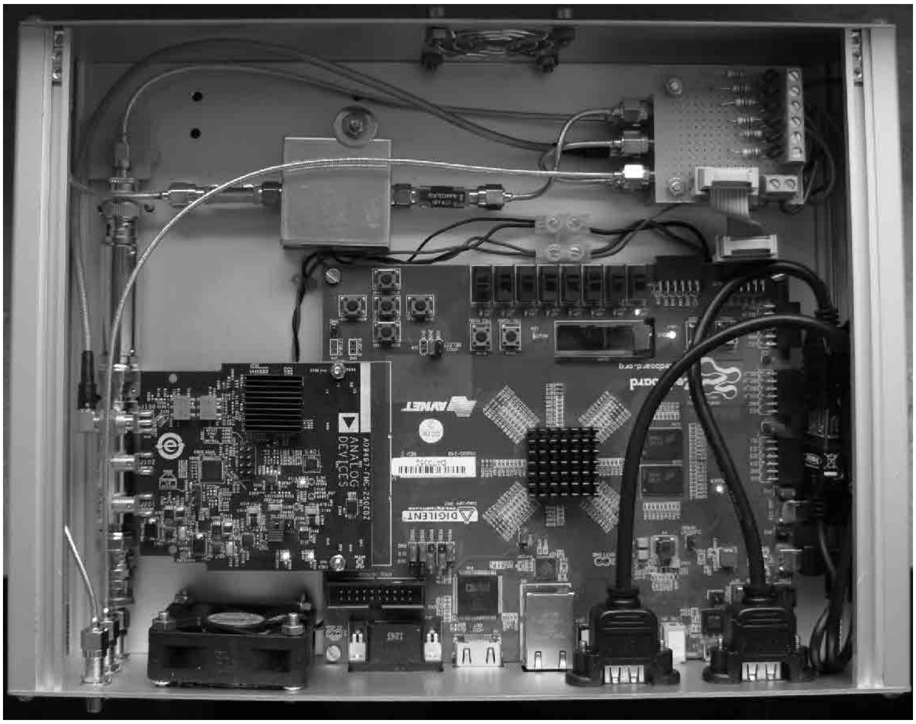

Figure 5 — The inside view shows the Zedboard (large center PCB) and the ADC board (small left-side PCB) with the anti-aliasing filter in the lengthy metal case on the left. The relays for switching the front end and their driver circuit are located In the upper right corner. The square metal box contains the 70-cm band front end.

Advantages of DSP

The extensive use of digital signal processing has many advantages over analog circuits: Analog processing is often limited by the laws of physics, that can hardly be overcome. Digital processing is limited only by circuit complexity — a better performance (sensitivity, dynamic range, spurs, agility, etc.) is achieved by more complex calculations and larger bit widths. Since semiconductor technology has continuously advanced following Moore's law, very complex systems can be built today and it is possible to achieve extraordinary accuracy and performance for digital signals with comparatively little effort.

The advantages of digital signal processing are apparent at different stages of the receiver:

— Digital FIR filters can be built with virtually any filter response. Low insertion loss, steep transition and high stop-band attenuation are only a matter of calculation resources. In many cases even medium-sized FPGAs can handle complex high-performance filters extremely well.

— Digital mixers and amplifiers are implemented as multipliers that do not introduce any spurs, such as harmonics and unwanted IMD and do not have any gain imperfection or other parasitic behavior.

— Digital oscillators based on direct digital synthesis (DDS) achieve extremely high spectral purity with low hardware complexity, and have virtually no spurs or harmonics. The frequency can be changed instantaneously without phase discontinuity.

— In general, effects of aging, impedance mismatch and all sorts of EMC issues (grounding, coupling between HF components, noisy power supply, etc.), that can bug analog designers, are completely eliminated in digital signal processing. The problem of additional noise that is introduced by most analog circuitry and measured as noise figure, is also widely eliminated in the digital domain. Noise is present only as quantization noise, caused by discrete sampling in amplitude. However, this quantization noise can be made arbitrarily low using larger bit widths.

— A sometimes overlooked property of digital radios is that they can be copied easily, since the radio usually consists of software or other design files, such as programming code, VHDL, FPGA design files or a GNU radio project on a PC. Before trying modifications, the designer can "backup" the radio and retrieve it, if the changes do not work.

Direct sampling receivers, indeed, have the potential to build high-performance receivers, since DSP can be implemented with almost arbitrary accuracy. Once the signal has been digitized, it is not prone to further degradation. The only bottleneck that may degrade the signal is the AD converter. It is important to carefully select the AD converter, since its properties define the overall performance of the radio in terms of linearity, sensitivity and bandwidth.

Challenges of Direct Sampling

When implementing a direct sampling receiver, the designer faces some special challenges regarding clock jitter and data rates.

First, the phase noise or clock jitter of the sampling clock is critical for high-speed SDRs. Even a small clock jitter may introduce additional noise in the ADC, which limits the SNR of the ADC. SNR is:

$$\mathrm{SNR} = 20\log\left(1/(2\pi f_a t_{jitter})\right) \qquad (1)$$

where t_{jitter} is the clock jitter in seconds and f_a is the maximum analog input frequency.

The SNR is not only dependent on the amount of clock jitter, but also increases with the frequency of the analog signal. Since direct sampling receivers often deal with higher analog frequencies than other SDRs, it is important to consider the noise introduced by clock jitter.

Second, the AD converter outputs the samples for further DSP at a very high data rate. The sampling theorem relates sampling rate with bandwidth: The sampling rate must be twice the highest frequency of the analog signal. The larger the processed analog bandwidth, the higher data rates are produced. In case of the Panoradio data rates up to 16-bit × 250 MHz = 4 Gbit/s occur at the input, and with I and Q, up to 2 × 22 bit × 250 MHz = 11 Gbit/s internally. This requires powerful hardware platforms like FPGAs because the computational load is too high for PCs or other processors.

Analog Front End and AD Converter

The Panoradio analog front end is mainly responsible for conditioning the analog signal coming from the antenna to prepare it for AD conversion, which includes proper anti-aliasing filtering.

Analog Front End

The Panoradio offers three analog input paths and analog front ends for different frequency ranges.

— A 90 MHz low-pass filter for a full panoramic view of HF and VHF.

— A 425-445 MHz band-pass filter plus a pre-amplifier for 70-cm band direct sampling reception.

— A direct path without filtering for experiments and flexibility.

AD converters in a direct sampling receiver require some analog front end to cope with aliasing due to sampling. Every analog input signal above $f_s/2$ appears in the band between 0 and $f_s/2$.

Wideband Reception: 0-100 MHz

To avoid aliasing in the Panoradio, with its $f_s = 250$ Msps, a low-pass filter must suppress all frequencies above 125 MHz before the AD converter. It is crucial to achieve a high stop-band attenuation, otherwise strong signals from TV transmitters or mobile phone services might interfere with signals in the 0 to 125 MHz range. Since most analog filters have a quite smooth transition from passband to stopband, the corner frequency of an anti-aliasing filter is usually significantly less than $f_s/2$, that is, less than 125 MHz.

In the Panoradio, two 90 MHz low-pass filters consisting of two cascaded Crystek CPFL 0090 filters are used. These filters provide around 55 dB stopband attenuation each, or 110 dB in total (Figure 6). With this high suppression, aliasing signals are completely eliminated before reaching the AD converter.

Undersampling for 70-cm band reception

The Panoradio has a second front end to enable undersampling of the 70-cm band. With undersampling the analog signal can have a frequency above $f_s/2 = 125$ MHz, which clearly violates the sampling theorem. So, aliasing occurs and these signals appear in the digital domain in the band below 125 MHz, according to the formula,

$$f_{alias} = \left| nf_s - f_{input} \right|$$

where f_{input} is the analog signal to be sampled and n is an integer, such that $f_{alias} < f_s/2$.

This behavior, which is not normally desired, provides an opportunity to translate signals above 125 MHz into the reception band, just as an analog mixer would do. The 70-cm band from 430 to 440 MHz, then translates to between 70 and 60 MHz, with an inverted spectrum.

Since in AD conversion many frequencies have the same alias frequency, as can be seen from the formula, analog signals potentially interfere with each other in the digital domain below 125 MHz. For example not only the band 430 – 440 MHz, but also the bands 180 – 190, 310 – 320, 560 – 570 MHz and so on, appears in 70 to 60 MHz. Therefore it is mandatory to use an analog band-pass filter before AD conversion that passes only signals from the 70-cm band, ensuring that no interference occurs.

The task of the Panoradio 70-cm band front end is twofold. First, it provides suppression of signals out of the 70-cm band. Second, it amplifies the input signal in order to improve sensitivity. The front end consists of a 433 MHz band-pass filter (Foxtech BPF 433) with 50 dB stop-band attenuation, followed by a low noise amplifier LNA4ALL[6] (+23 dB) and a second band-pass filter, shown in Figure 7. The first filter prevents the amplifier from overloading from unwanted out-of-band signals. The second filter sets the ultimate anti-aliasing and prevents broad band noise from the LNA to reach the AD converter, achieving a total amplification of +20 dB and more than 80 dB stopband attenuation (Figure 8).

High-Speed AD conversion

The AD converter is the heart of every direct-sampling receiver. Its properties are crucial for the overall receiver performance. Important ADC parameters are SNR, SFDR (spur free dynamic range), analog bandwidth and full-scale voltage. It should be noted that the number of bits is of minor importance, because the least significant bits usually carry only noise. The ENOB (effective number of bits) considers this issue and describes the number of bits that are "noise-free".

The Panoradio uses the AD9467-250, which is a state-of-the-art 16 bit, 250 Msps

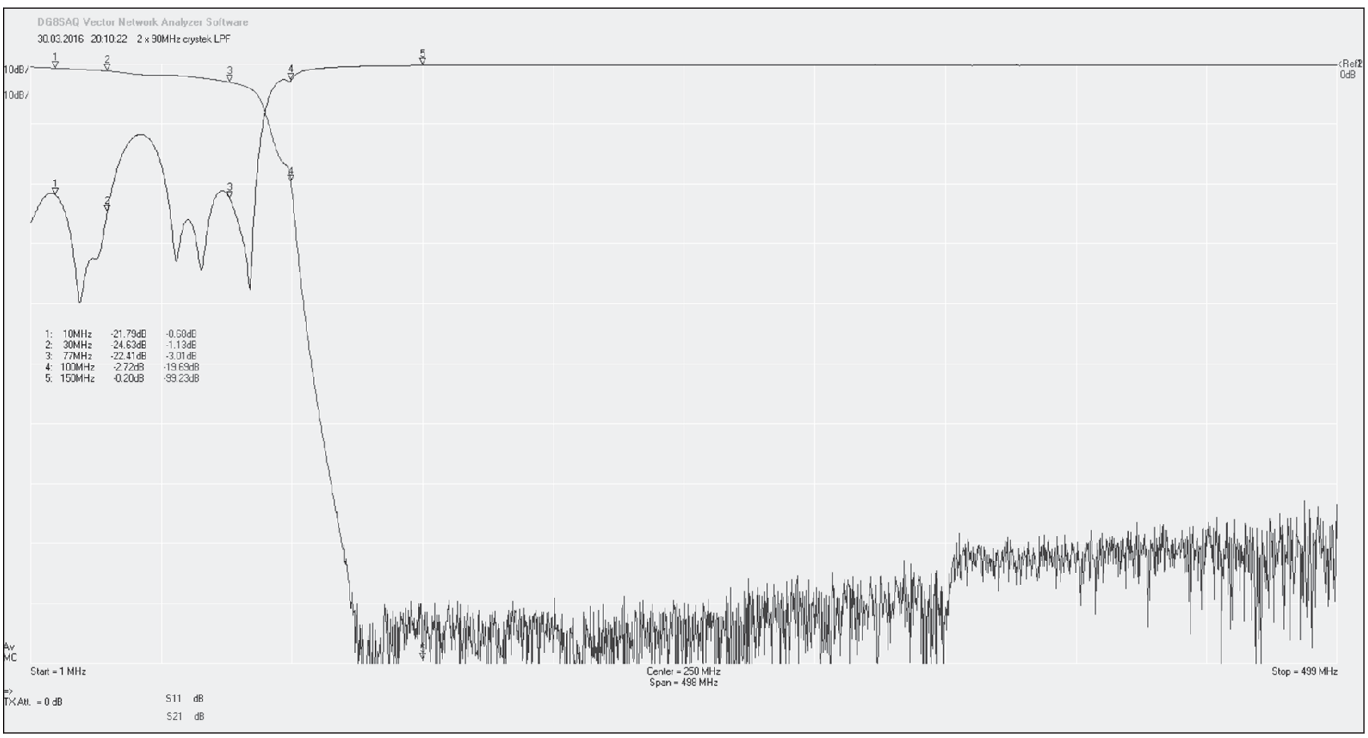

Figure 6 — 100 MHz front end with two cascaded analog low-pass filters. The dynamic range is limited by the VNWA 2 used for measurement.

AD converter with excellent properties and a wide analog bandwidth[7]. It is characterized by,

— SNR: 76 dBFS / 12.3 ENOB
— SFDR: 100 dBFS
— Analog bandwidth: 0 – 900 MHz
— Full scale voltage: 2.5 V p-p.

The Panoradio uses the AD9467 evaluation board from Analog Devices with a FMC connector, that perfectly fits into the Zedboard FMC socket. Minor soldering modifications were necessary on the evaluation board to enable the on-board clock source, and to route the analog signals correctly to the AD converter chip (circumventing the optionally available preamp and clock conditioner implemented on the board) and to provide impedance matching.

Impedance matching from a 50 Ω single-ended system to the ADC differential input with its 530 Ω in parallel with 3.5 pF consists of mainly a resistive matching plus some additional components for filtering of high-frequency noise. Resistive impedance matching has the advantage of providing a wideband match supporting frequencies up to many hundreds of megahertz. Two 1:1 transformers convert the signal from single-ended to differential. Using two transformers mitigates phase and amplitude imbalances that otherwise would cause spurs[8].

The full scale power of the sole AD converter at 2.5 V p-p for the 530 Ω input results in

$$P_{FS} = \frac{V_{p-p}^2}{8R} = \frac{2.5^2}{8 \cdot 530} = 0.00147 \text{ W}$$

which equals 1.7 dBm.

With the SNR of 76 dB and the thermal noise for a 125 MHz bandwidth of -93 dBm the ADC noise figure is around 19 dB.

However, taking the analog matching circuit of the ADC into account, the full scale voltage rises due to the resistive impedance matching network. Since a resistive match is not lossless, it introduces a loss of approximately 9 dB. Therefore the full-scale power of the complete circuit rises to $P_{FS} \approx$+11 dBm. Note that this decreases the sensitivity by 9 dB, resulting in a final noise figure of approximately 28 dB.

The full scale power of the AD converter circuitry of +11 dBm is quite large for the application of a direct sampling SDR. Noise figure is usually not a major problem for operation at low frequencies in the HF band. For higher frequencies it is expected to affect the receiver sensitivity. To cope with this shortcoming, that virtually all direct sampling receivers exhibit, a preamp is used to get an improved practical dynamic range and sensitivity.

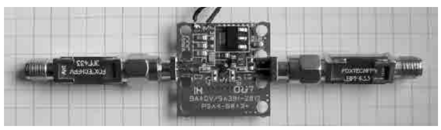

Figure 7 — 70-cm band front end with two 433 MHz band-pass filters and the LNA4ALL preamp.

For the Panoradio wideband input (0 – 100 MHz) I decided to not use a preamp to keep distortions low. An active preamplifier always introduces distortion, a problem that increases with the number of strong signals an ADC is exposed to. Due to the high bandwidth and large number of signals in the Panoradio, using no preamp excludes possible distortions. Transformers may also introduce distortions and considerable insertion loss at higher frequencies due to parasitic behavior. This effect gets worse for larger turn ratios, that is, higher voltage gain[8]. For ratios like 1:4 or 1:9 the maximum usable signal frequency is limited, which would heavily impair performance for 70-cm band reception. However, if high sensitivity is required, it is easy to attach an external preamp (preferably with band limitation filter). So the decision against a preamp or transformer in the Panoradio was in favor of wideband operation with as little non-linear distortion as possible.

Later FFT processing exhibits large processing gains such that the sensitivity for monitoring signals in a waterfall is much better than the specified ADC sensitivity (up to 73 dB for full zoom).

Low clock jitter is important to preserve the ADC SNR. The evaluation board features a high-performance Vectron VCC 6 clock generator that has a very low jitter of 133 *f*. Equation (1) shows the maximum achievable

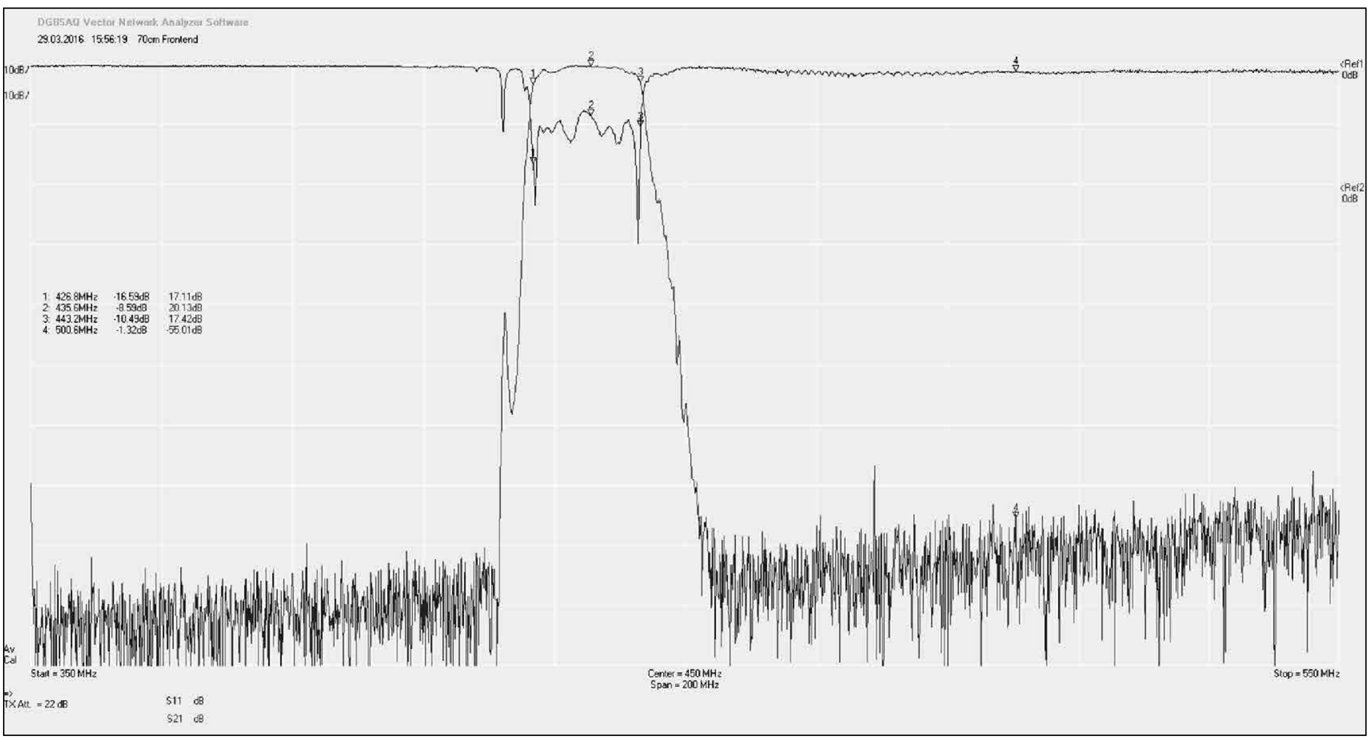

Figure 8 — 70 cm band front end measurements. The dynamic range limited by the VNWA2 used for measurement.

SNR due to clock jitter. For input frequencies from 0 - 100 MHz (the Panoradio main frequency range) the theoretical SNR is always more than 80 dB, which is clearly above the ADC SNR of 76 dB. So the clock jitter of the Vectron VCC 6 does not affect the noise performance of the receiver.

For input frequencies above 240 MHz, the impact of the clock jitter cannot be neglected. For 430 MHz (70-cm band), the ADC SNR and sensitivity is reduced by approximately 7 dB. Therefore, in contrast to the wideband input, the 70-cm band front end uses a preamp to make up for this additional loss by increasing the signal level and thus lowering the noise figure from an excessive 34 dB to moderate 13 dB. Non-linear distortion introduced by the preamp is not an issue, because it is a narrowband front end. The band-pass filters reject out-of-band signals before they reach the preamp. This also shows that sensitivity is a general issue for undersampling that needs to be countered by large amplification.

This completes the analog front end and the ADC discussion. From this point on the SDR does all processing with DSP.

The Zedboard as Digital Processing Platform

SDR algorithms run on a hardware platform, either as software on a processor such as a PC or microprocessor, or as a digital circuit in a microchip, like an FPGA. These different hardware platforms have diverse processing speed, flexibility and ease of programming. Direct sampling receivers usually deal with large bandwidths and high data rates, and therefore require powerful processing platforms. Before describing the Panoradio DSP implementation details, the available hardware platforms and requirements are further analyzed.

Processor vs. FPGA

A processor is basically a digital circuit that executes software serially and is therefore inherently slow and not suitable for large bandwidth processing in high-end SDRs. Exceptions are multi-core processors, like GPUs in graphics cards, that exhibit parallelism to some extent.

The advantage of the processor is its flexibility, since the tasks it fulfils can be changed easily by executing different software. Also, it is comparatively easy to program software. It is possible to run an operating system with the advantage of easy access to user interfaces like a monitor, keyboard, mouse, and providing access to hard disks for storage or for an internet connection.

In SDR applications, processors are very good for:

— demodulation of different modes

— providing a graphical user interface (GUI)

— controlling the receiver functionalities

— recording and storing signals

— providing network and internet access.

FPGAs are fundamentally different from processors. An FPGA is a microchip that contains a large number of configurable circuit elements, such as adders, multipliers, multiplexers and registers, that can be used to create complex digital circuits. For every algorithm, that is to be executed on the FPGA, a specialized digital circuit must be designed and loaded onto the FPGA.

The big advantage of FPGAs is that they can process highly parallel data, and can easily deal with data rates of several hundred Msps. The achievable speed-up over processors is often in the order 100 to1,000 times, even though FPGA clocks are often much slower than processor clocks. Since the circuit in an FPGA is custom designed, the bit widths of data, memories and arithmetic units can be application specific, which is very efficient. If, for example, a bit width of 6 bits is considered sufficient for the data, the processing units are laid out to use only 6 bits, whereas in a processor the bit width is restricted to a predefined number such as 24, 32 or 64 bits.

The drawback of FPGAs is that the hardware implementation (circuit design) of algorithms requires good knowledge of hardware description languages, digital circuit design and experience with the design tools. Also the implementation of interfaces to the outside (monitor, keyboard, network connection) is difficult.

In SDR applications, FPGAs are very good for high-speed, repeatedly occurring calculations like, mixing and down conversion; filtering; high-speed FFT; and parallel demodulation of a large number of signals.

In summary, processors and FPGAs each have advantages and disadvantages. In a SDR a wide range of different tasks must be performed with different speeds. That is difficult to achieve with only a processor or only an FPGA. So the ideal hardware for SDR is a combination of both — and indeed, high-end SDRs already follow this configuration. So also does the Panoradio, which uses the "Zynq" device from Xilinx, that combines a processor and an FPGA on a single chip.

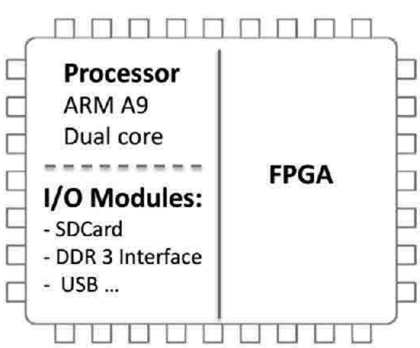

Figure 9 — The Zynq SoC features a processor and an FPGA as well as I/O interfaces on a single chip.

The Zynq

The Zynq chip is the result of a recent development in microelectronics to combine processors and FPGAs for acceleration in a system on chip. It contains a dual ARM Core A9, an FPGA and various interfaces (such as DDR3 memory interface, USB, SD Card) as shown in Figure 9. There are several thousand possible internal connections between the processor and the FPGA, which are programmable and allow fast transmission of data. This is a big advantage over a separate combination of FPGA and processor on a PCB, where the number of possible connections is small and can not easily be changed after PCB design.

The interconnect between processor and FPGA is a memory mapped interface. That means that a part of the processor memory is virtually connected to a 32-bit register in the FPGA through a AXI bus (Figure 10). The processor can write data to these special memory addresses (using the "mmap" command or a Linux driver), that is then transferred to the FPGA registers. Or it can read the registers in the FPGA by reading data from the memory. Direct memory access (DMA) is available for transferring large amounts of data.

Programming the Zynq involves two tasks; programming the processor, and designing the digital circuit for the FPGA. These two tasks can be done nearly independently by different tools.

The processor can be programmed either bare metal (programming it directly in *C*) or using an operating system. I strongly recommend using an operating system like Linux, which hides many details of the processor and provides easy access to user interfaces as well as to standard software. Then any programming language can be used to develop software for the Zynq.

The hardware design for the FPGA is done with Xilinx software "Vivado" using VHDL or Verilog and predefined IP blocks. For the Panoradio, the different receiver blocks (like FIR filters and DDS oscillators) are implemented using VHDL plus Xilinx IP cores. These blocks are connected together in a top level block-design schematics by using the graphical editor provided in Vivado.

The Zedboard

The Zedboard is a low-cost and open-source evaluation board and is ideal for experimenting with the Zynq device (**www.zedboard.com**). The Zedboard contains many peripherals that unleash the power of the Zynq, such as power supplies, 512 MB DDR3 Memory, Gigabit Ethernet, USB connectors, HDMI interface, SD card slot and an audio interface (sound codec ADAU1761). These peripherals make the Zedboard a fully stand-alone system. Additionally, the Zedboard offers a FPGA Mezzanine Card (FMC) connector, which is a standardized socket for attaching peripheral boards. I used this socket to attach the ADC board.

Receiver DSP Implementation

This section focuses on what is actually implemented in the Zynq FPGA and processor. SDR signal processing includes many different steps: digital down conversion, filtering, demodulation, FFT, control, interfaces and GUI. High-speed data processing is handled in the FPGA part of the Zynq, whereas low data rate tasks are handled by the processor. The overall implementation is shown in Figure 11.

After AD conversion the FPGA receives the samples via the ADC data interface module. It performs some ADC tests, takes care of line delays and converts the differential double-data-rate signals to single-ended data with single data rate, that can be processed in the FPGA. Optionally, a test signal generator can feed well-defined waveforms into the circuit for testing. Three digital down conversion (DDC) blocks perform the actual reception. One DDC implements functionality for the three zoomable waterfall displays. Two DDCs are dedicated to the two audio receivers. The following audio post processing block can further reduce bandwidth and contains a Weaver SSB demodulator. Interfaces for audio and HDMI to the processor are also implemented in the FPGA.

The ARM A9 processor takes care of

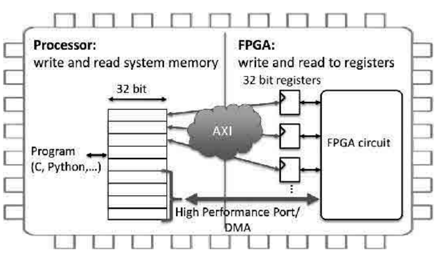

Figure 10 — Data transfer between the processor and the FPGA via memory mapped AXI bus or DMA.

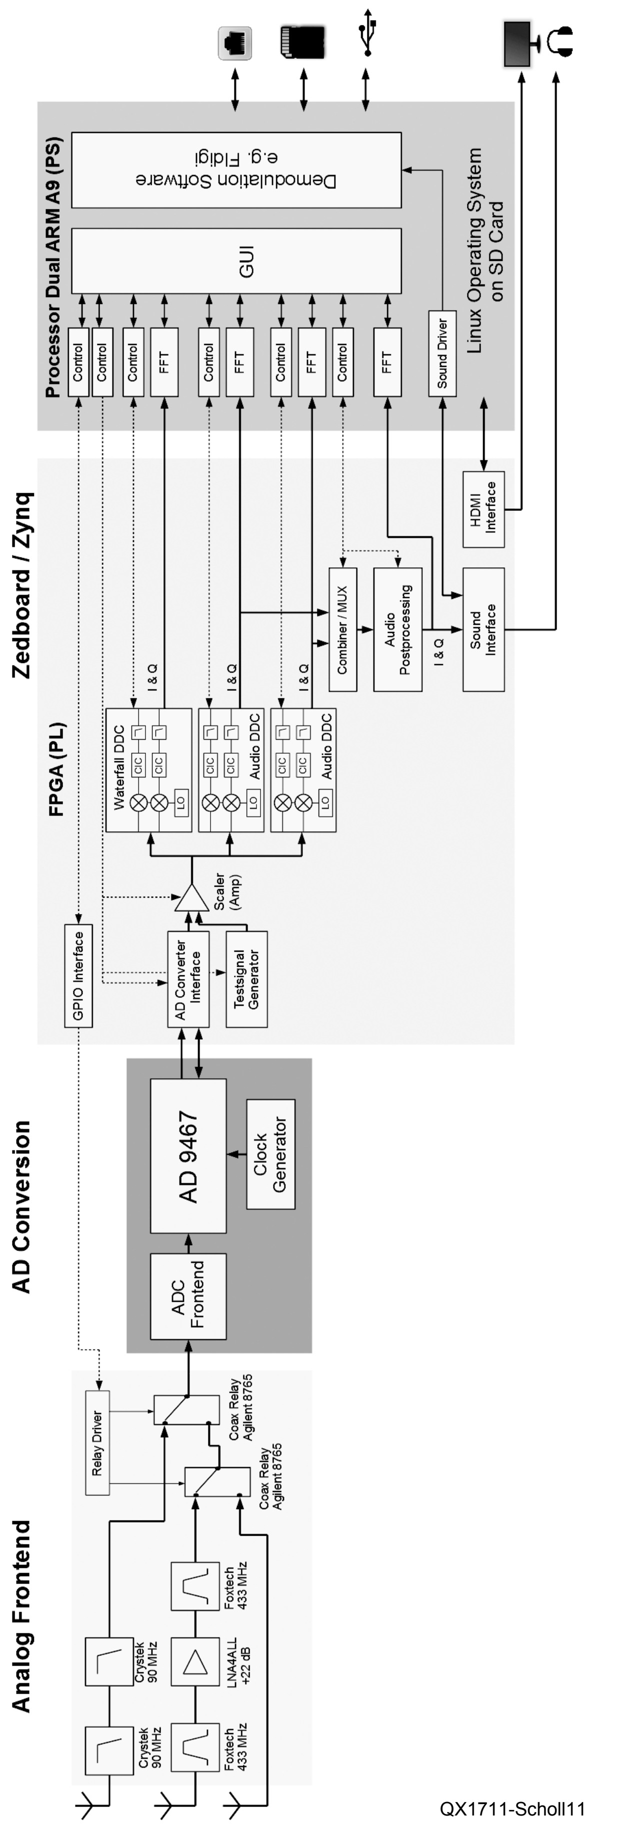

Figure 11 — Block diagram of the complete receiver showing the three main parts: analog front end, AD converter, and the Zedboard for digital processing.

all slow-speed processing, such as running a GUI to control the radio, calculating FFT, displaying received signals and executing demodulation software (Figure 4). Moreover it runs the OS for standalone operation.

DDC for Zoomable FFT

The three Panoradio waterfall displays can independently zoom on signals from the full span of 100 MHz to a span of 6.1 kHz as is depicted in Figure 12. This task is accomplished by the implementation of a zoom FFT with flexible bandwidth. This is a DDC, that down converts and filters the selected portion of the spectrum followed by an FFT. For this purpose, two challenges have to be solved: how to implement the zoom function efficiently, and how to implement three waterfalls efficiently.

DDC and Zoom Functionality

Overview: The DDC for the zoomable FFT is shown in Figure 13. It follows the typical design of a DDC with a digital local oscillator (LO), a complex mixing stage with two multipliers and CIC and FIR filters in the I and Q paths. The filters reduce the bandwidth to the desired span of the current waterfall setting. Finally, the I and Q samples are stored in a memory and sent to the processor for FFT.

Local Oscillator (LO): The implemented LO generates sine and cosine waveforms using direct digital synthesis (DDS). In DDS a sine is generated using a large look-up table. It stores samples of a sine waveform, which is then samples and interpolates at the desired frequency. A second table or a second address pointer generates the corresponding cosine. DDS LOs can achieve high accuracy and agility. For FPGA implementation I used the predefined Xilinx DDS IP block[9].

Filtering: CIC filters are low-complexity circuits that combine filtering with decimation. Decimation reduces the sampling rate from 250 Msps to a fraction determined by the decimation factor between 4 and 4,096 in order to relax the computational load for the FPGA in the subsequent stages. The following sharp FIR filters with 140 coefficients determine the final bandwidth and an additional decimation by a factor 8 sets the final sampling rate. The filter design has been done with the Matlab "Filter Design and Analysis Tool". The calculated filter parameters and coefficients were imported into the Xilinx IP blocks for CIC and FIR filters in Vivado.[10]

Zoom functionality: The required bandwidth of the DDC needs to be very flexible between 100 MHz and 6 kHz to implement the waterfall zoom function. For the cases of 100 MHz (1× zoom) and 50 MHz (2× zoom) bandwidth, the input samples are directly routed to the memory for a real FFT. For higher zoom levels, greater than 4× (span < 25 MHz), the signal passes through the DDC. The flexibility in output bandwidths from 25 MHz (4× zoom) to 6.1 kHz (16,384× zoom) is achieved only by a programmable CIC decimation rate of 4 to 4096 (16× and higher) or bypassing the CIC filters (4× and 8× zoom). By changing only the CIC decimation rate by a factor of 2, the output bandwidth also changes by a factor of 2. Thus the FIR filter remains fixed, which greatly simplifies the DDC. Figure 14 shows the resulting filter responses for different zoom levels.

Performance: The DDC is designed for excellent performance. The DDS provides LO signals with 21 bits each and 110 dB SFDR. The stop-band attenuation of CIC and FIR filters is typically greater than 105 dB (see Figure 14). The 140 coefficients for FIR filtering have a large bit width of 24 bits, such that the influence of coefficient

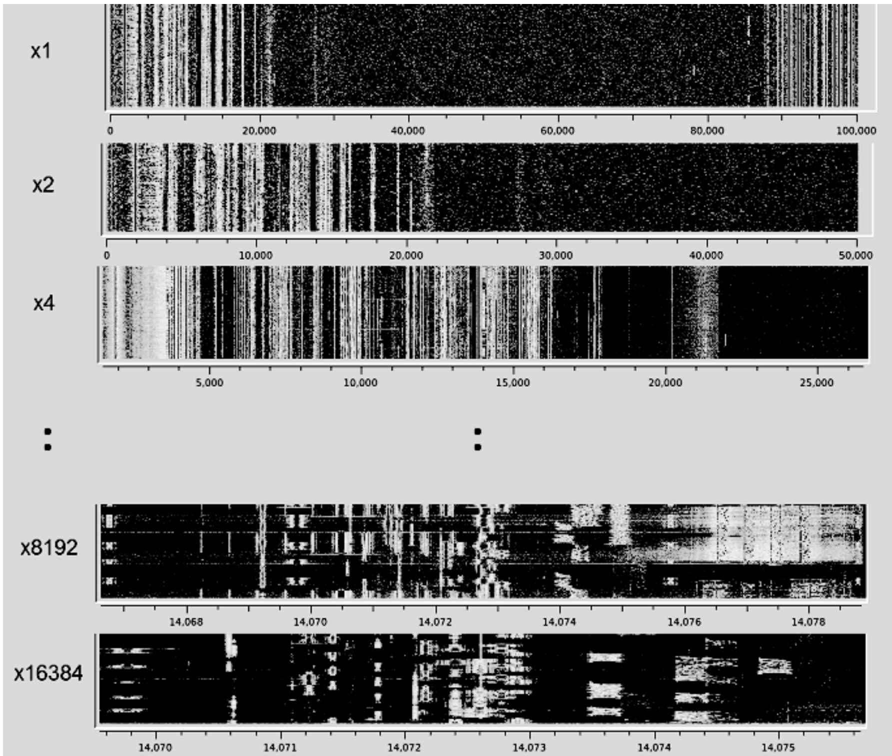

Figure 12 — Zoom function of the waterfall, here from full span to digital signals in the 20 m amateur band.

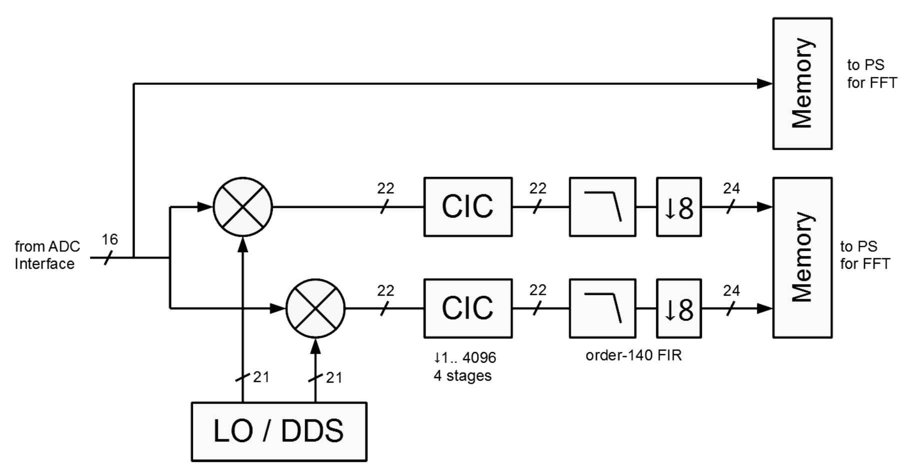

Figure 13 — DDC for zoomable FFT spectrum display.

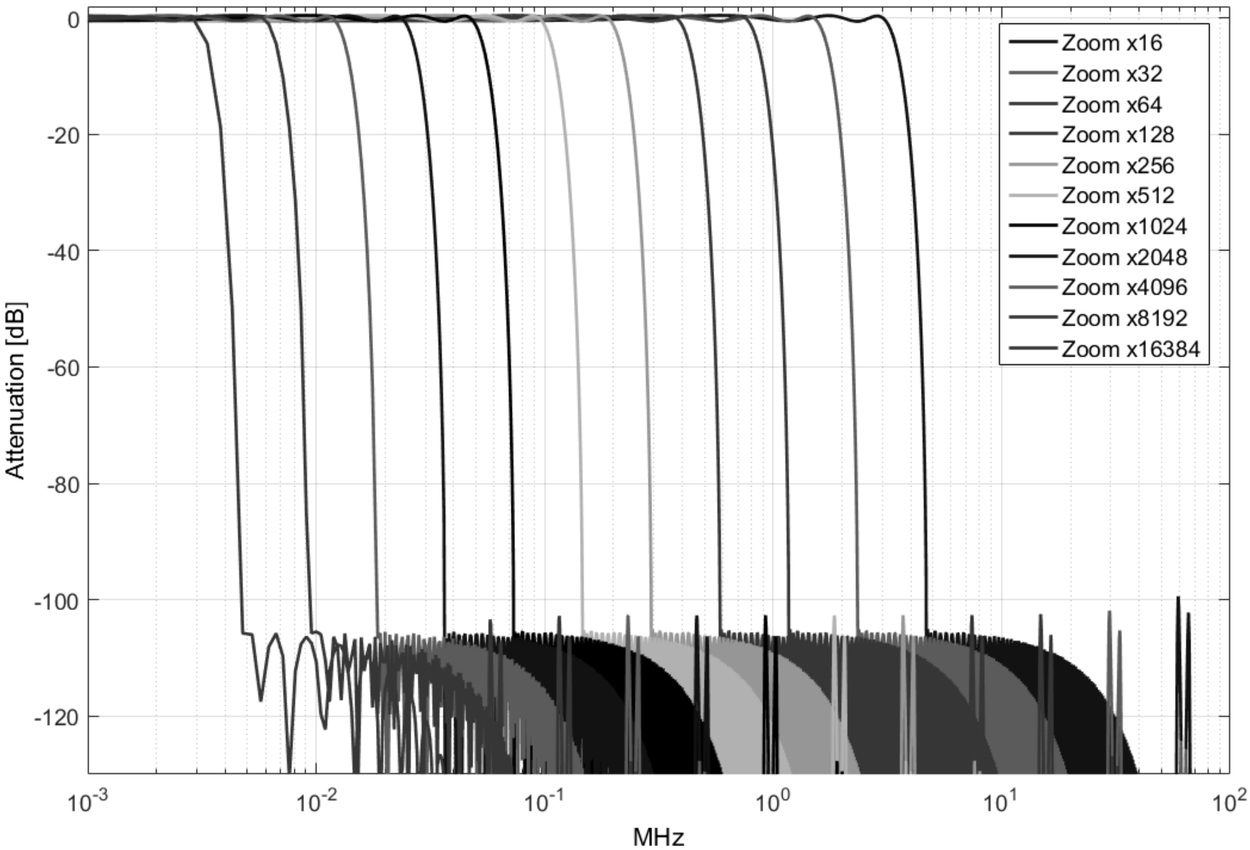

Figure 14 — Response of the DDC for the zoomable FFT for different zoom levels from x16 to x16384 achieved by only changing the decimation in the CIC from 4 to 4096.

quantization is negligible. For the overall DDC even worst case, for example full scale, out-of-band signals appear only with very small amplitude with less than 105 dBFS. In more realistic cases the DDC will not introduce any spurious signals. The bit width of I and Q signals is large enough to keep the influence of quantization errors very low (the introduced quantization noise is less than few tenths of a decibel). Note, that the bit width of I and Q channels increase with smaller bandwidth because noise power reduces, and therefore SNR increases.

One DDC – Three Waterfalls

For the three waterfall displays, three DDCs would be required. However, it is more efficient to share one single DDC between the three waterfalls. Then the DDC quickly switches continuously between the parameters of the different waterfall displays, captures data and calculates an FFT before switching to the next. The parameters that need to be switched are the center frequencies (LO) and CIC decimation rate. The LO frequency can change instantaneously since the DDS is very agile. The filters however require some time to recover from transients, which is a potential problem especially for high decimation values. However, worst case simulations of transient behavior showed, that the system recovers very quickly and this is not an issue if the switching time stays above several tens of milliseconds.

FFT Processing

For calculating the FFT, the DDC captures 8,192 I and Q samples in a memory (4,096 samples for 100 and 50 MHz bandwidth). Then the samples are weighted with a Hann window to reduce FFT leakage. A complex FFT is carried out resulting in 8,192 frequency bins. Since the monitor resolution is too low to display a 8,192 point spectrum, the spectrum is scaled down to 1,024 points by summing eight neighboring bins. This approach has a big advantage. It largely increases the resolution and sensitivity for detecting signals that are otherwise masked by the "leakage bell" of strong signals, see Figure 15. Improvements of some ten dB might occur in the vicinity of strong signals.

Optionally, the Panoradio can do several data captures and FFT operations consecutively to display the average spectrum. This does not reduce the noise floor, but reduces fluctuations in the spectrum plot.

Audio Reception and Demodulation – Audio DDC

The audio DDC follows the basic principle of the waterfall DDC with the exception, that the CIC has a fixed decimation rate of 2,048, see Figure 16. A FIR filter reduces the bandwidth to 11 kHz, so the total bandwidth for I and Q is 22 kHz, which fits well to a 48 kHz rate sound card. The audio I and Q outputs are passed to the audio post processing block for further bandwidth reduction and SSB demodulation. In addition a memory is included to capture samples for a 256 point FFT for the audio receiver waterfall displays. The additional decimation-by-two before FFT memory storage reduces the sample rate and thus increases FFT resolution. This decimation is not applied to the I and Q outputs, so as not to violate the sampling theorem at the subsequent SSB demodulator output. The audio DDC suppresses out-of-band signals by more than 80 dB, which provides very good attenuation of unwanted signals for practical operation.

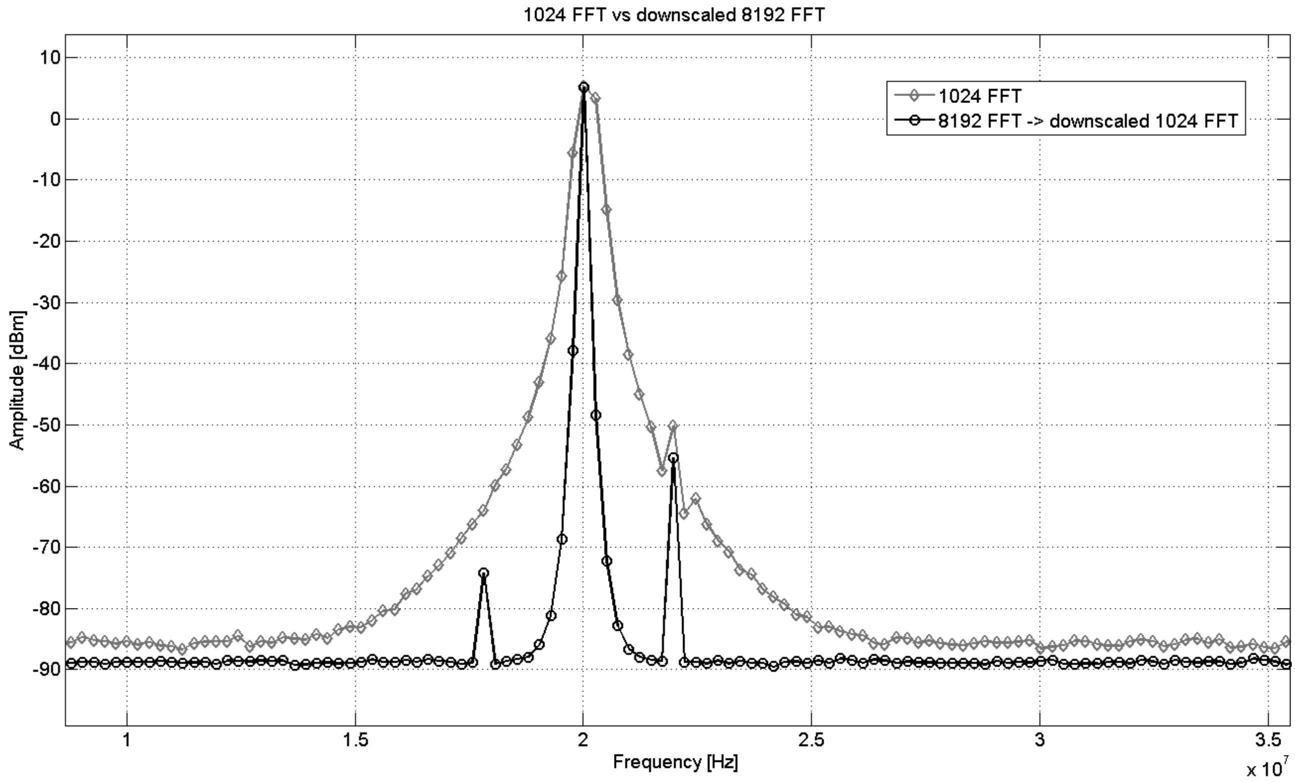

Figure 15 — Simulation and comparison of a native 1024-point FFT and a 1024-point FFT that originates from a downscaled 8192-point FFT.

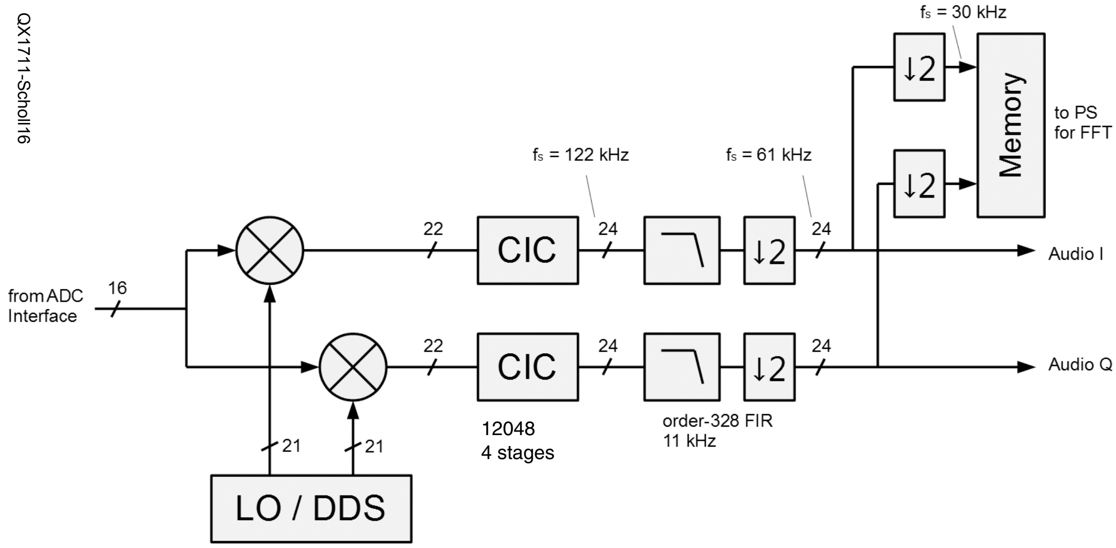

Figure 16 — One of the two audio DDCs implemented for audio reception.

The Panoradio features two of the above described audio DDCs. The I and Q output of the two DDCs are either multiplexed or summed. Summing allows for audio reception of two different bands simultaneously, that is, to receive and demodulate PSK31 signals from two bands at the same time.

Audio Post Processing and SSB Demodulation:

The I and Q samples from the audio DDC are fed to the audio post processing for volume setting, bandwidth reduction, SSB demodulation and sample rate conversion as shown in Figure 17. The first step is a clock domain crossing reducing the FPGA clock frequency from 250 MHz to 100 MHz, which greatly relaxes timing constraints for FPGA synthesis and place and route. This is possible because the signal bandwidth at this point is just several tens of kHz wide and the sampling rate is just 61 kHz. The clock domain crossing is followed by a software controlled variable gain amplifier for audio volume setting. Additional optional FIR filters can reduce the ultimate audio bandwidth from 22 kHz to 6 or 2.4 kHz. The signals can then be passed as I and Q values to the sound card directly or with prior SSB demodulation. A Weaver demodulator[11] is implemented for optional SSB demodulation. I consider the Weaver demodulator as one of the best options for SSB demodulation in DSP, since the required mixers, oscillators and low-pass filters are of low complexity and can be implemented with virtually perfect phase and amplitude balancing.

The sample rate after filtering (or optional SSB demodulation) is 30 kHz or 60 kHz depending on the selected filters. In any case this does not match with the 48 kHz sample rate of the Zedboard audio codec. Therefore an additional asynchronous sample rate converter interpolates the input signal at 30 or 60 kHz and filters the output to retrieve the signal at the desired 48 kHz.

FPGA Design

The high-speed digital processing parts of the Panoradio are implemented on the FPGA. The Zedboard Zynq 7020 is a mid-sized Zynq SoCs with 53,200 look-up tables and 106,400 flip-flops, 560 KB SRAM and 220 DSP slices. The FPGA for the Panoradio design is less than 50% occupied (46% look-up tables, 34% flip-flops, 39% memories, 45% DSP slices / multipliers). This quite low utilization would allow for an even more complex DSP. Main contributors to the FPGA utilization are the waterfall DDC, which operates at high sample frequencies and requires high parallelism (especially DSP slices for multiplications for FIR filtering), and the interconnect circuitry between processor and FPGA.

Control Software & Linux OS

The dual core ARM A9 processor runs the Linux operating system Linaro (based on Ubuntu for embedded devices). The OS and file system are stored on an SD card. With the Linux OS the Panoradio offers comparable possibilities as a PC (although the ARM A9 does not have as much computing power). The SDR can run any general purpose software for Linux. For this project, I compiled Fldigi for the ARM A9. Other useful software, like file browser, editors, audio players or the program "Scilab" for scientific calculations can run on the processor. I used QtCreator as the programming IDE. It offers the possibility to program directly on the Zynq instead of applying a cross compiler. Furthermore, the OS provides Ethernet, USB and HDMI interfaces for attaching network, monitor and input devices. During system start-up, a bootloader programs the FPGA with the bit stream containing the FPGA implementation.

The task of the software for controlling the Panoradio is mainly to provide a GUI and to control the radio functionality. This control includes setting the analog front end path, initializing the ADC and audio interface, setting the DDC frequencies and filter properties, determining amplifications for ADC and audio interface, and reading I and Q data from the DDCs for the FFT. Furthermore, the

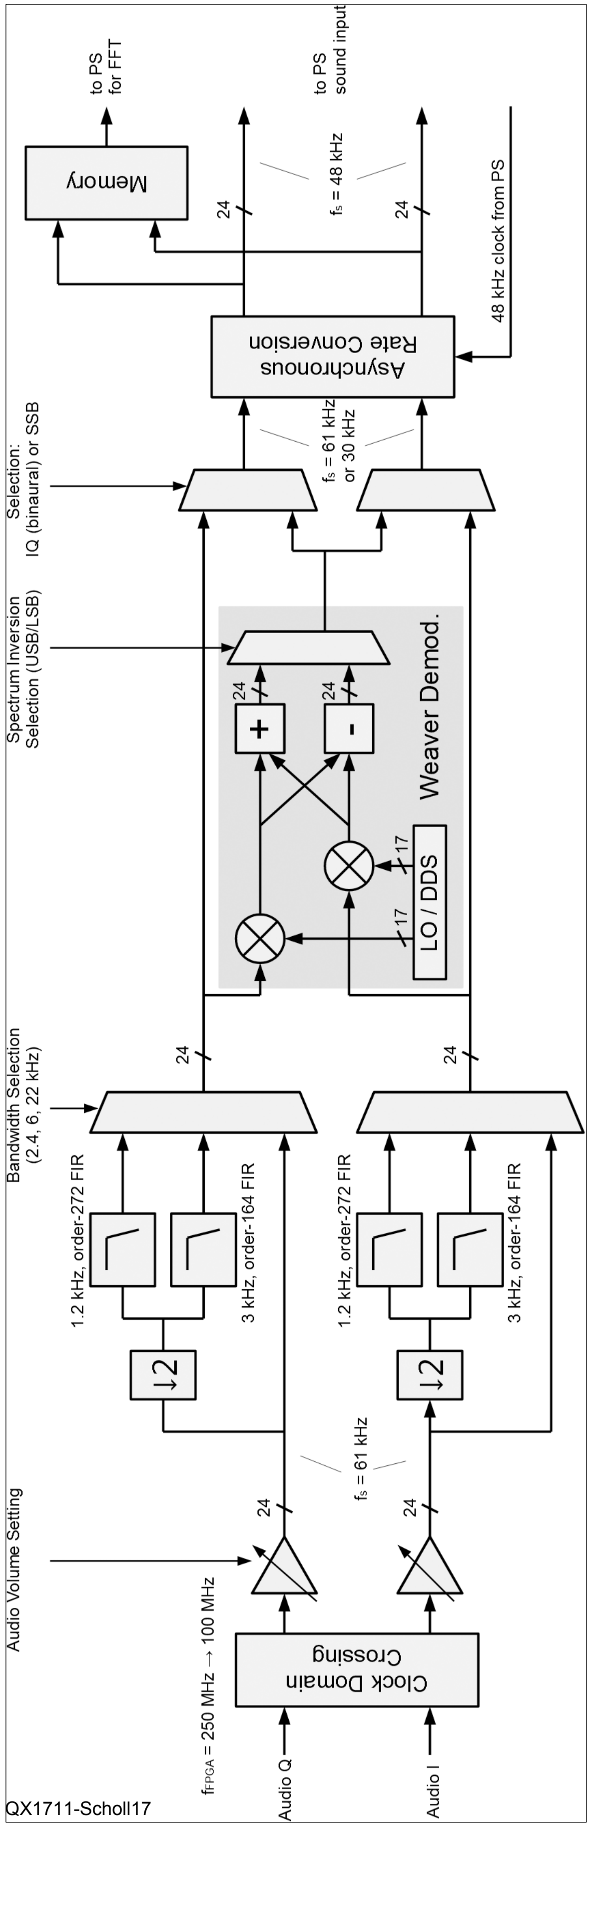

Figure 17 — Audio post processing: volume setting, final bandwidth selection, SSB Weaver demodulation, and sample rate conversion.

FFT calculations (including windowing, averaging, decimation, etc.) for the waterfall displays are done in software with the processor.

The control software is written in C++ and Qt. Qt provides a powerful framework for implementing a GUI and enhances C++ in many different useful ways. Qwt[12], a library for scientific applications, that provides various kinds of plots, is used to implement the waterfall displays.

The software GUI has the following features.

— Selection of the analog front end (100 MHz, 70-cm band or direct).

— AD converter input level monitoring.

— AGC control after the ADC.

— FFT averaging, update, time, and color scheme.

— Different modes for audio reception (bandwidth, SSB, spectrum inversion).

— Monitoring of audio signal level and setting for audio amplifier.

— Selection of a single audio receiver or the sum of both.

— Control of digital test signals.

— Monitoring of internal receiver states for debugging (signal and its spectrum at the different stages, setting of filters, LO, etc.).

For FFT calculation the software implements a wrapper for the FFTW library[13] a C library with fast FFT algorithms. The wrapper adds functions for applying a window, FFT averaging, converting to power scale in dBm and provides C++ style access to the FFTW library.

I chose C++ as the programming language, because the execution speed of the software is critical. Interestingly, the bottleneck is not the FFTs or the communication with the IP cores, but the drawing routines for the waterfall plots. The Zynq does not have any graphics acceleration core, which could speed up the drawing process. However, Xilinx has realized this bottleneck and included graphics acceleration in the Zynq successors; the UltraScale MPSoc, and the Zynq UltraScale+.

Conclusion

The Panoradio was designed to demonstrate the potential of today's analog and digital devices for SDR. Modern AD converters can digitize signals both fast and accurately. Recent digital FPGA devices can process the high data rates of direct sampling receivers and provide the possibility to run software and even a Linux OS on a single chip. The combination led to a stand-alone SDR covering frequencies from dc to VHF as well as 70-cm UHF band with direct sampling. The Panoradio is an open-source project available on **www.panoradio-sdr.de** in order to encourage Amateur Radio enthusiasts to try, learn and further develop direct sampling SDR technology.

Stefan Scholl, DC9ST, was originally licensed in 2006, as a German Class A license. He received the Diploma in electrical engineering in 2010 from University of Kaiserslautern, Germany, where he worked for 6 years as a researcher in the area of forward error correction for communication systems. He has nearly completed his PhD. Currently he is an engineer for radar technology at Hensoldt (formerly Airbus Defence and Space). Stefan's interests include homebrew equipment and hardware and algorithms for software defined radio.

Notes

[1]James C. Ahlstrom, N2ADR, "An All-Digital Transceiver for HF", *QEX* Nov./Dec. 2010.

[2]Pieter-Tjerk de Boer, PA3FWM, WebSDR: **https://wwwhome.ewi.utwente.nl/~ptdeboer/ham/sdr/**.

[3]High Performance Software Defined Radio, **www.openhpsdr.org**.

[4]Pavel Demin, Red Pitaya SDR, **pavel-demin.github.io/red-pitaya-notes/sdr-transceiver/**.

[5] Fldigi by W1HKJ, **https://sourceforge.net/projects/fldigi/files/**.

[6]Adam Alicajic, 9A4QV, **lna4all.blogspot.de/**.

[7]Analog Devices, AD 9467 Datasheet, Rev D ed.

[8]Rob Reeder (Analog Devices), "Transformer-Coupled Front-End for Wideband A/D Converters", Analog Dialogue, Vol. 39, Apr. 2005.

[9]Xilinx, DDS Compiler v6.0 LogiCORE IP Product Guide, PG141 ed., Jun. 2015.

[10]Xilinx, FIR Compiler v7.2 LogiCORE IP Product Guide, PG149 ed., Jun. 2015.

[11]Donald K. Weaver, "A Third Method of Generation and Detection of Single-Sideband Signals", Proc. of JRE, Nov. 1956.

[12]U. Rathmann. (2015, August) Qwt - Qt Widgets for Technical Applications. **qwt.sourceforge.net/**.

[13]M. Frigo and S. Johnson, "FFTW," MIT, Tech. Rep. 3.3.3, Nov., 2012.

James L. Tonne, W4ENE

118 Garnet Lane, Oak Ridge, TN 37830; **tonne@comcast.net**

All-Pass Networks in a Speech Chain

W4ENE discusses asymmetrical audio waveforms, the problems this creates, and ways of minimizing the asymmetry.

It is a well-known fact that a speech waveform viewed on an oscilloscope is quite commonly lopsided. That is, one side of the waveform, say the top side, has a greater peak amplitude than the bottom side. The degree of this asymmetry is highly dependent on the voice of the individual. Here, we will explore this phenomenon and outline some problems that may arise from it. We will also discuss methods for minimizing the problems.

Possible Problems – Entirely Linear Systems

In a high quality public address system where the speaker's voice is simply augmented by a power amplifier and loudspeaker, there would be no problem if the volume levels are such that the system is entirely linear. The lopsided waveform would pass with its asymmetry unnoticed.

If such a lopsided waveform were to be used to modulate an AM radio transmitter, and if the modulation index is adjusted to a relatively low level, then again such a waveform would offer no problem. However, it has been found best for ordinary AM systems to set the polarity of modulation so that the peaks with the greater amplitude are modulated upward. This minimizes distortion found in the typical AM signal envelope. However, in a purely linear system, asymmetric waveforms are not in themselves a problem.

AGC Loops with Fast Attack Times

There was at one time — circa 1950s, 1960s — a competition among various broadcast equipment manufacturers to see who could develop the fastest-acting AGC system to control the modulation in a broadcast transmitter. These devices were commonly called volume limiters or volume limiting amplifiers. They generally reacted to an overload situation within a millisecond. Following the overload they would restore the gain to normal over a period of perhaps a few seconds. They were without question better than a human controlling the modulation levels.

It was interesting to view the output of these devices on an oscilloscope. There was no visible clipping or other artifact added to the waveform, just a (usually) well-controlled modulation level. The winner in the war-of-speed used a system that had a zero attack time. It used a delay line to delay the audio signal while the gain-controlling voltage was being generated. But in every one of these units there was an annoying tendency for the device to respond to signals that were not the same as that to which the human ear responded. Rephrased, they were controlling modulation, not volume. Maximizing volume was becoming an issue at the time. So, while the outputs of these fast-acting devices looked nice on an oscilloscope, they didn't accomplish the broadcaster's needs.

If a lopsided waveform were to be applied to one of these units, the peak with the greatest magnitude, whether positive or negative, would cause the generation of AGC voltage. If those peaks could be made equal in amplitude then less AGC voltage would be generated and modulation would increase. This must be done, however, in a manner that does not increase the peak-to-peak value.

Transient Clipping

Research at CBS Laboratories showed that if the AGC loop could have a reaction time (attack time) of a few milliseconds, and a recovery time (release time) of perhaps 200 ms, such an AGC system would best

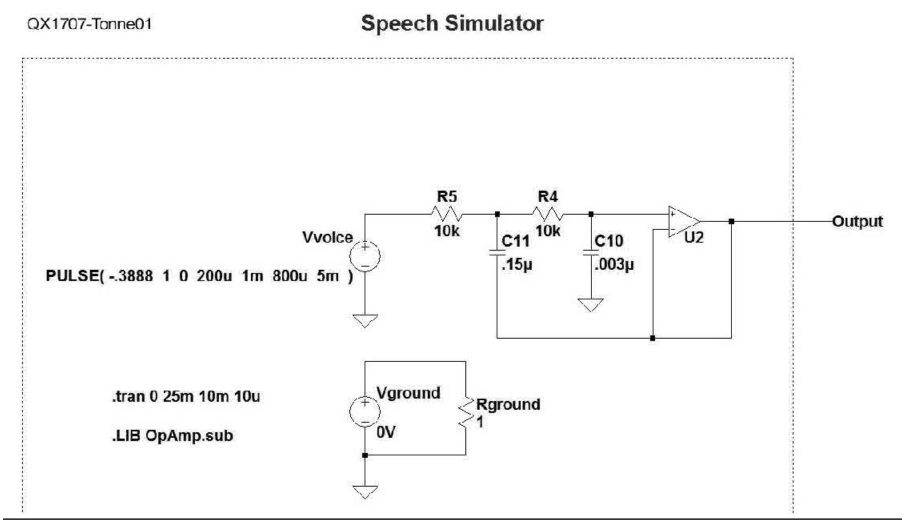

Figure 1 — LTspice rendering of the speech simulator.

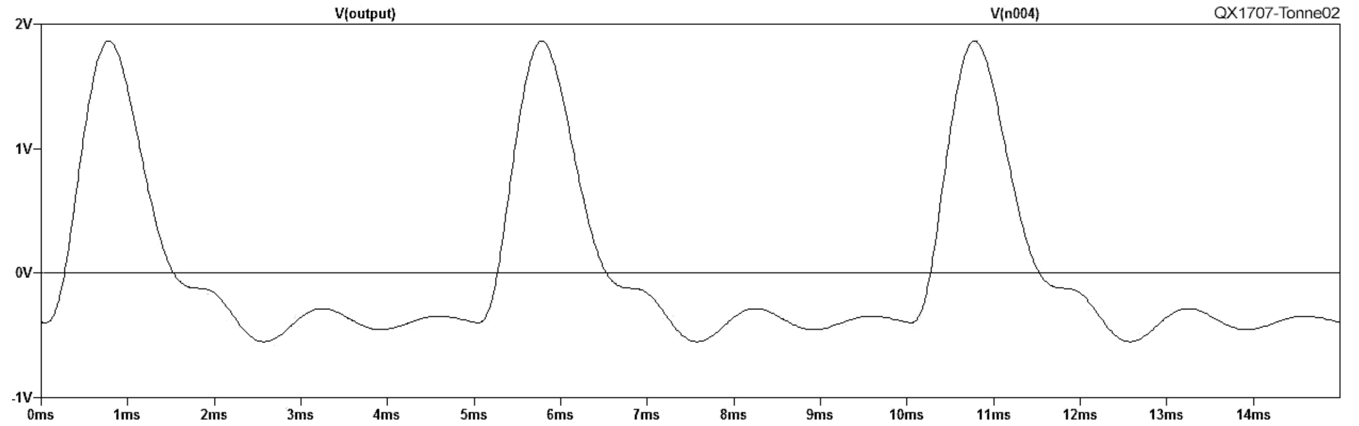

Figure 2 — Waveform from the speech simulator.

match the response of the human ear. Such an AGC system could control loudness and it would match perception by the ear, allowing the volume to be maximized. This is what broadcasters were looking for.

There was a drawback to such a scheme. The relatively long attack time required that such an AGC system must be followed by a clipper to catch the transients that escaped the AGC system. One of the first commercial units to use these techniques was the *CBS Volumax* system. Such a long attack-time system was quite a departure from the conventional wisdom of the time.

In a system with such a long attack time, an asymmetrical applied waveform causes less AGC voltage generation than in a system with a short or fast attack time. However, the signal from the AGC system must then be applied to a clipper. Clipping one audio peak more than the other results in a dc, or at least a sub-audible syllabic, signal component from the clipper. In an AM transmitter this appears as a form of amplitude "carrier shift". In an FM transmitter this appears as a center frequency shift, and interferes with the FM transmitter Automatic Frequency Control (AFC) system. It is certainly disconcerting to watch an analog frequency meter on an FM broadcast transmitter kick violently when such a unit is used to control the modulation with an applied program containing asymmetric waveform components involving a clipper.

To minimize this problem, the waveform should be processed in a manner such that prior to clipping, the peaks are rendered symmetrical, preferably without increasing the peak-to-peak value. Next we look at ways to handle the problem.

A Standardized Waveform

Let us generate a standard waveform that can be reproduced easily and will allow us to compare various approaches to processing. The proposed signal has an approximate 10 dB (3:1 voltage ratio) positive to negative amplitude ratio, but it has no dc component. The areas under the curve for the positive and for the negative portions of the waveform are equal. A circuit to generate such a speech waveform is shown in the LTspice model rendered in Figure 1. The signal generated by that circuit has a fundamental frequency of 200 Hz and is a believable replica of a steady speech signal, perhaps the sound "ohhhh". The generated signal is shown here in Figure 2.

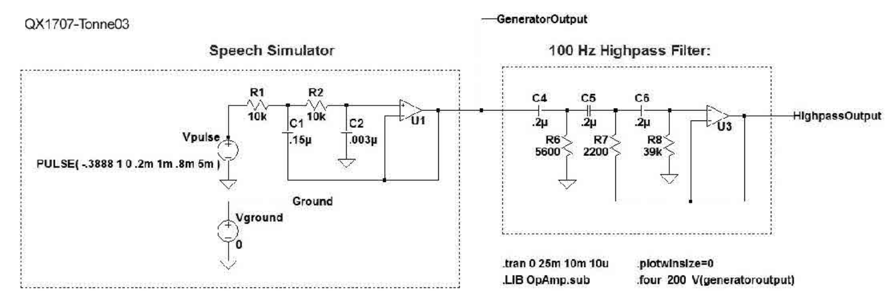

Figure 3 — LTspice rendering of the speech simulator of Figure 1 followed by a 100 Hz high-pass filter.

There is about 10 dB of asymmetry in this waveform and there is no dc component. The areas above and below the centerline are precisely equal. If this waveform were to modulate a transmitter directly, the upward peaks would require 10 dB more power than the downward peaks. If the transmitter could not handle this degree of asymmetry, then the modulation level would have to be reduced until the positive peaks were in a linear region and the negative peaks would be reduced in amplitude. This is certainly an inefficient use of transmitter capability.

A Solution– Use a High-pass Filter

One way to make the waveform symmetrical "top to bottom" would be to apply it to a high-pass filter. Indeed, this may be a part of the speech-processing chain already. Such a high-pass filter would do double duty. It would remove those components that serve no purpose if transmitted, and in fact might cause mischief. The high-pass filter might also make the positive and negative peaks more nearly equal. If the speech signal were applied to a high-pass filter and then to a clipper, there would be similar amounts of clipping applied to the positive and negative peaks. There would be a reduction in axis shift due to any sub-audible components generated by asymmetric clipping. The ear normally tolerates clipping of both modulating waveform peaks better than clipping just one side of the waveform.

In Figure 3 we see our standard speech waveform generator of Figure 1 connected to a 100 Hz high-pass filter. Figure 4 shows the output of the filter (lagging waveform) compared with the output of the generator (leading waveform). The horizontal line depicts the zero voltage level. Observe that the peak-to-peak voltage value has actually increased. This is not our objective. The use of a high-pass filter is not helpful in this respect.

A Better Solution: Use an All-pass Network

Another way to process the speech signal is to pass it through a network that has a flat frequency response but rearranges the relative phases of the signal frequency components to make it less asymmetric. Because such a network passes with all frequency amplitudes equally, it is called an all-pass network.

Let us look at an active version — using op-amps, resistors and capacitors — of such an all-pass filter. At various frequencies the phase through the circuit shifts. If a complicated waveform, made up of a fundamental and various harmonics, is applied to the circuit, the harmonics have their phases altered relative to the fundamental. This can be accomplished with a lumped-element circuit — using inductors and capacitors — or it can be accomplished at much lower expense using active circuitry using op-amps, resistors and capacitors as shown in the Figure 5.

This network will have a flat frequency response if R1 and R2 have the same value. R3 and C1 can be interchanged. The delay performance will be identical, just the phase will be inverted.

The Kahn SymmetraPeak Circuit

Using an all-pass network to make an audio signal more symmetrical is certainly not new. A commercial product to provide this function was called the SymmetraPeak and was marketed by Kahn Communications about 1959. It was a lumped-element device for the simple reason that at the time op-amp circuitry was not available. The schematic of the SymmetraPeak, rendered in LTspice, is shown in Figure 6. Our quite asymmetric test waveform has become relatively symmetrical, as seen in Figure 7.

Active Circuitry

The SymmetraPeak was bulky and expensive. When op-amp circuits were developed that accomplished the same thing it faded away. An example of an op-amp equivalent (Figure 8) was designed by Gary Blau,W3AM. Our standard signal as it exits from that active all-pass is shown in Figure 9. It is even more symmetrical than from the SymmetraPeak.

The signal peak-to-peak amplitude is not changed. The magnitude of the higher-amplitude peak has been reduced while at the same time the magnitude of the lower-amplitude peak has been increased. There is no axis shift, and no sub-audible components have been added. The areas under the curve above and below the zero axis are equal. Clipping of such a waveform would cause a minimum of "mischief" compared with clipping of the original asymmetric waveform. The phase shift of this network goes from near zero degrees at very low audio through near 1440 degrees at extremely high audio frequencies.

RMS-sensing AGC and Clipping

If the modulation level in the transmitter is controlled by a peak-sensing audio AGC unit (a "limiter"), that AGC system will respond to the peak with the highest instantaneous magnitude. But if the modulation level is controlled by an rms-sensing AGC unit, asymmetry does not enter into the picture at all with respect to the AGC portion of the speech processing. But an rms-sensing AGC unit must be followed by a clipper to catch those waveform excursions that escape the AGC unit. Be advised that those excursions will be of significance. But if the clipper operates on one side of the waveform more than the other, a dc or sub-audible component will be developed by the clipper. This will normally cause trouble in the modulator proper. We have a situation wherein the modulator must be direct-coupled to properly handle the signal.

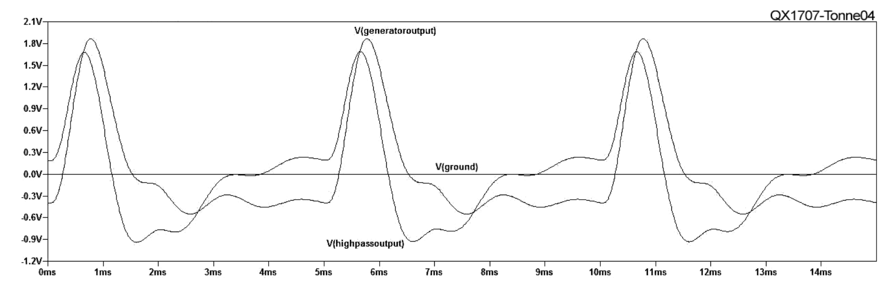

Figure 4 — Output of the filter (lagging waveform) compared with the output of the generator (leading waveform). The horizontal line depicts the zero voltage level.

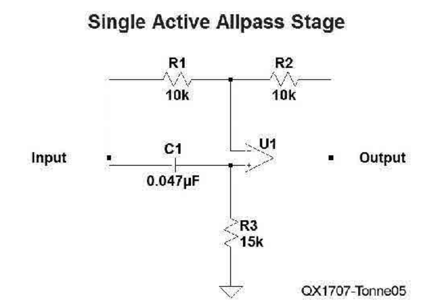

Figure 5 — A single stage active all-pass network rendered in LTspice.

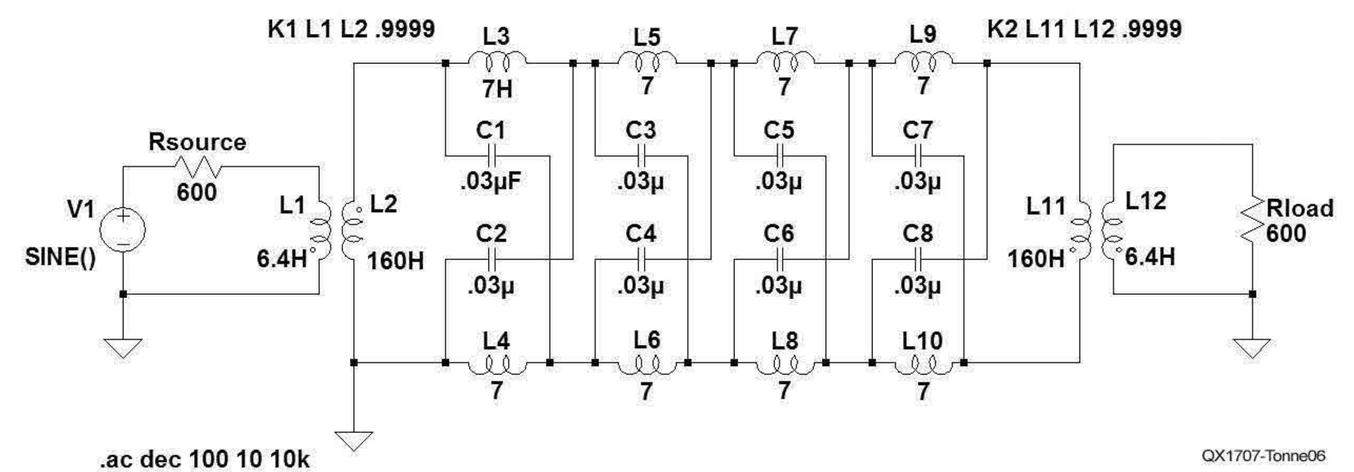

Figure 6 — The Kahn SymmetraPeak circuit rendered in LTspice.

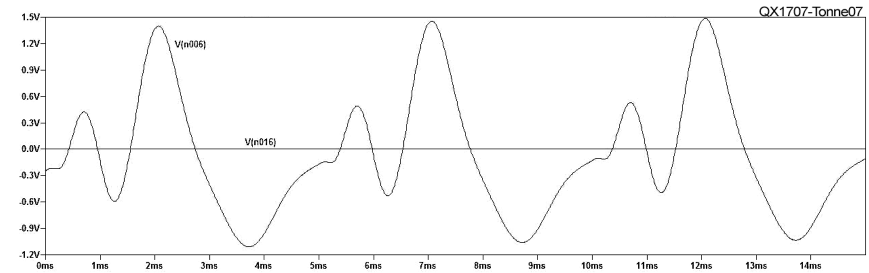

Figure 7 — Output of the SymmetraPeak. The trace is V(n006), and the straight line is V(n016).

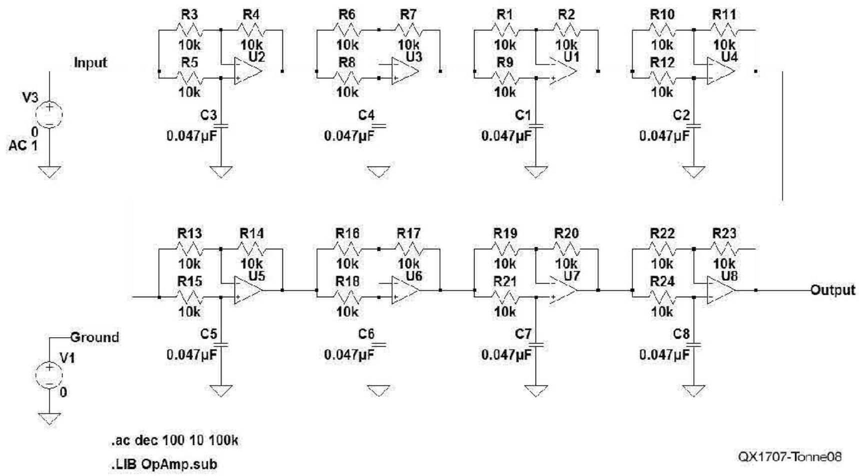

Figure 8 — An active all-pass network rendered in LTspice.

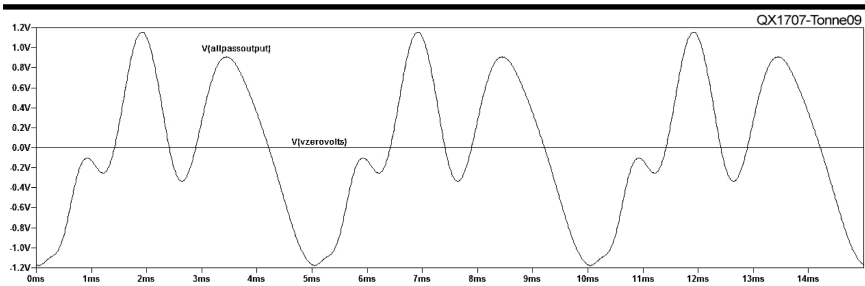

Figure 9 — Output of the active all-pass network. The trace is V(allpassoutput), and the straight line is V(vzerovolts).

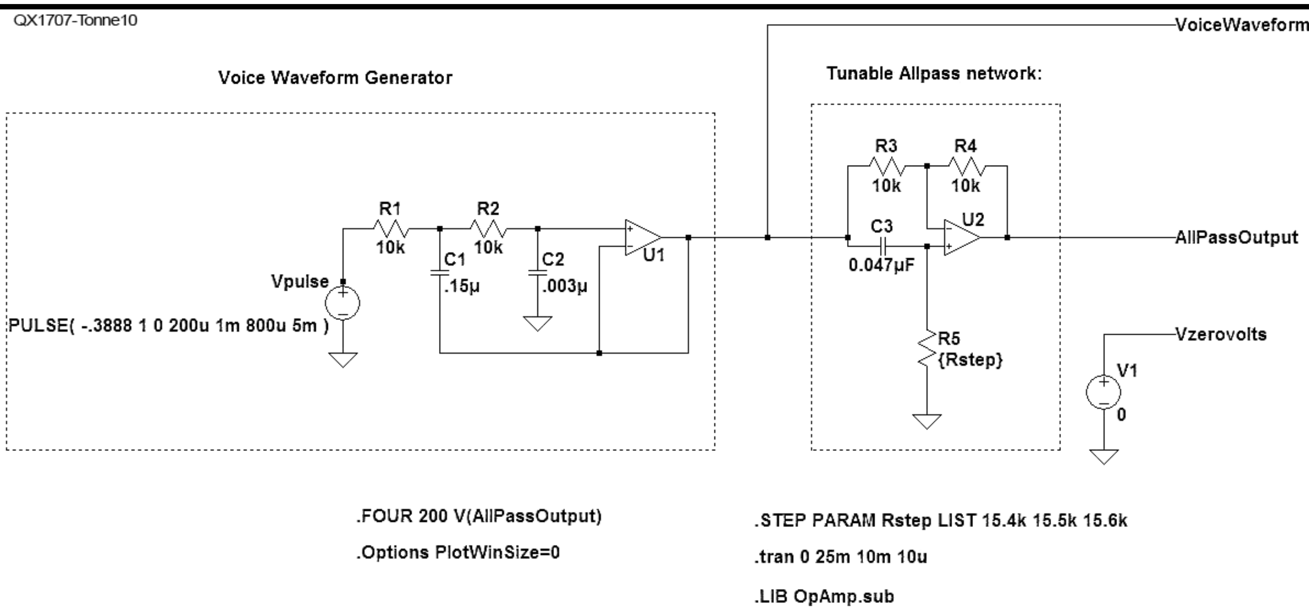

Figure 10 — Schematic of a tunable (with Rstep) all-pass network driven by the waveform generator.

Here again the all-pass can come to the rescue. Insert the all-pass between the rms-sensing AGC block and the clipper. By adding the all-pass network at this point, the clipper will clip symmetrically and no sub-audible components will be involved.

A Tunable All-pass Network

A tunable single-stage all-pass network has been devised that may very well be suitable in most cases. The schematic is shown Figure 10 as it underwent development in LTspice. By using an oscilloscope, R5 (“Rstep”) can be adjusted for best waveform symmetry on the individual’s voice. However, the peak-to-peak amplitude will not be as nicely controlled as with the more complex network. This simple circuit does give a significant amount of performance as seen in the output waveforms of Figure 11.

This all-pass network is quite comparable with the more complex circuits for this speech waveform. It is the customization that allows the high degree of performance seen with this simple circuit. The caveat is that various speech waveforms may require readjustment of the potentiometer. This should not be a problem in the usual one-user application.

A Possible Failure Mode

The use of an all-pass network is not a cure-all. If the applied waveform is symmetrical top-to-bottom in the first place — it has no even-order harmonic content — then the all-pass can actually increase the signal peak-to-peak amplitude. In Figure 12 we see an example of an applied waveform with top-to-bottom symmetry. The signal consists of a 200 Hz fundamental and an equal amplitude 600 Hz third harmonic. Now let us apply that signal to a Blau all-pass circuit. The resulting output signal is shown in Figure 13.

The signal still has top-to-bottom symmetry but the peak amplitude has actually increased. This can be seen by comparing of this waveform with the waveform of Figure 12. This illustrates an interesting point. If the waveform to be corrected has no even-order components (unlikely in practice) and so is symmetrical top-to-bottom in the first place, then the using an all-pass network might not be beneficial, and in fact will be harmful. This aspect of the all-pass is usually glossed over.

A Suggestion

An all-pass block should be placed at an appropriate point in a speech processor used in a radio transmitter. It should be placed ahead of a clipper if any preceding AGC circuit is slow-acting or especially if it is mis-sensing.

The all-pass network is inexpensive if constructed using op-amps and associated components, and generally ensures lower distortion by virtue of less clipping or at least symmetrical clipping. Symmetrical clipping always causes less “mischief” than does asymmetrical clipping. Placement of the all-pass block should always be prior to the point where clipping occurs or where it might occur.

The amateur radio fraternity has been relatively slow to pick up on the idea of using an all-pass network, although Gary Blau, W3AM, see **www.w3am.com/8poleapf.html**, also writes on this subject.

James L. Tonne, W4ENE, holds the Amateur Extra class license, and was first licensed in 1951. His current Amateur Radio interests are largely focused on speech processing and filter design. He has written several articles for QST *and* QEX *and was a major contributor to the RF and Filters chapter in the ARRL Handbook. He is the author of the Tonne Software package on the CD accompanying the ARRL Handbook and included as part of the downloadable package available on the ARRL web page.*

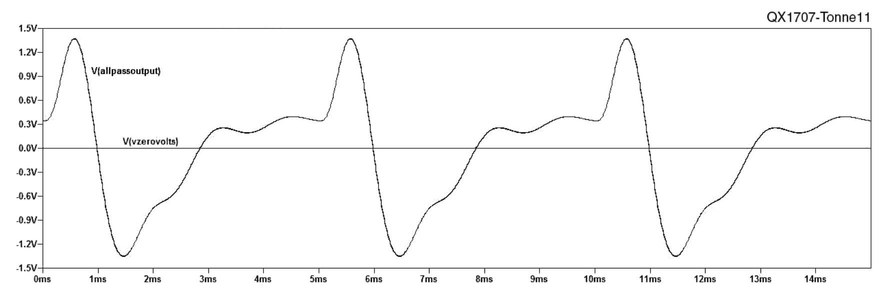

Figure 11 — Output of the simple tunable all-pass network. The trace is V(allpassoutput), and the straight line is V(vzerovolts).

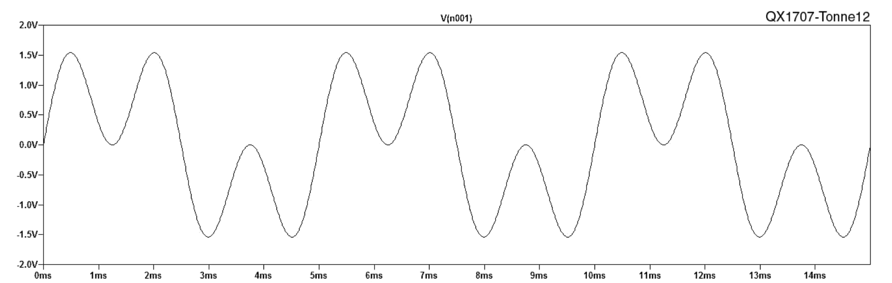

Figure 12 — A waveform V(n001) with top-to-bottom symmetry

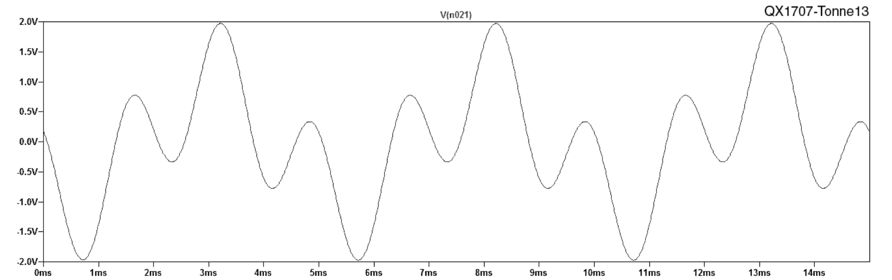

Figure 13 — Illustrating an all-pass failure, the result of applying the waveform of Figure 12 to a Blau all-pass circuit.

Technical Notes

Measuring Characteristic Impedance of Coax Cable in the Shack

Here is a simple method for measuring characteristic impedance, Z_0, of coax cable using equipment readily available in the shack. The technique presented here is based on a well-known property of transmission lines (John D. Kraus, *Electromagnetics*, 1953, pp. 433ff.). If one applies a short circuit at the far end of a transmission line that is exactly one-eight electrical wavelength long, the input impedance of the line will be exactly jZ_0 for a lossless line. Losses of short pieces of coax have negligible effect on the results.

An antenna analyzer with real and imaginary readout, ($R + jX$), is needed to see the jZ_0 term. I used the MFJ 269 (MFJ) and Rig Expert AA-230 (Rig Ex.) antenna analyzers for these measurements, and obtained good results.

Procedure

- Prepare the sample coax, or transmission line, to be measured by shorting the far end.
- Connect the ($R + jX$) meter to the near end of the transmission line.
- Tune the meter to find the half-wavelength null frequency $F1$ where R and X are approximately zero.
- Accurately divide $F1$ by 4 to get $F2$.
- Tune the ($R + jX$) meter to exactly $F2$.
- The display will read ($R + jX$) of the transmission line, where $X = Z_0$.

The results shown in Table 1 agree within an ohm or two of published data for the coax cables tested. I also measured the characteristic impedance of a twisted pair transmission line, seen in Figure 1.

Practical considerations for such tests

The quality of the terminations, both the cable short and connection to the ($R + jX$) meter have a significant effect on the results. Using standard coax connectors at both ends of the coax is recommended, but usable results can be obtained with carefully wired short circuit connections. A high quality short should be used (Figure 2).

Meter accuracy is a first order factor — the results are no better than the meter

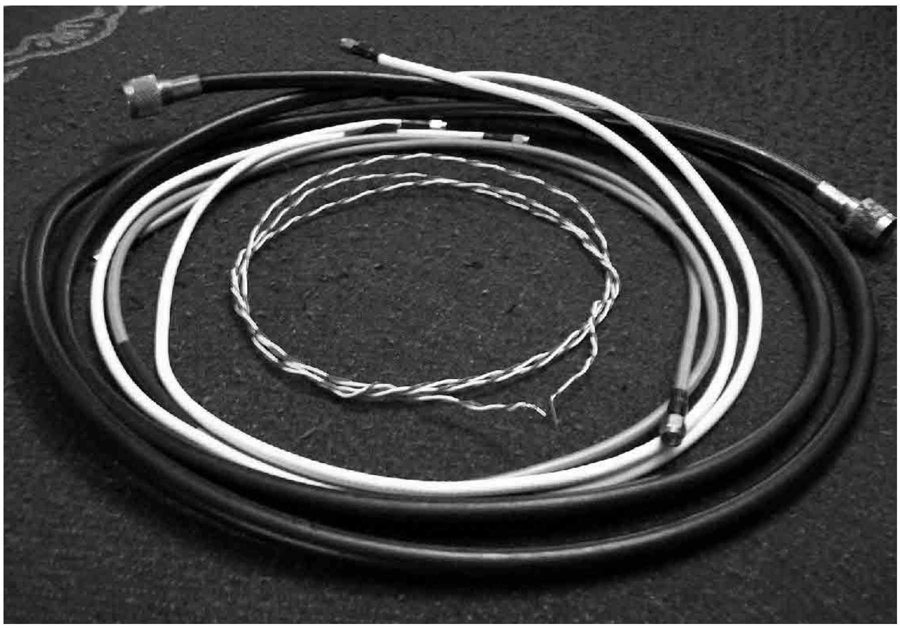

Figure 1 — Tested coaxial cables, including a twisted-pair transmission line.

Table 1.
Test results with a variety of coaxial cable samples.

Transmission line type	*Length, in.*	*Meter type*	*F1, MHz*	*"Short", Ω, at F1*	*F2, MHz*	*Z_0 impedance reading, Ω*
RG58C	70	MFJ	53.80	(1+j0)	13.45	0+j53
RG58C	70	Rig Ex.	53.90	(2+j0)	13.48	3+j54
RG59B	72	MFJ	52.80	(8+j0)	13.20	10+j73
RG59B	72	Rig Ex.	52.92	(10-j2)	13.23	15+j76
RG62A	49	MFJ	98.80	(2+j0)	24.70	0+j99
RG62A	49	Rig Ex.	98.95	(1-j0)	24.74	4+j100
RG142B	51	MFJ	79.30	(1+j0)	19.83	0+j49
RG142B	51	Rig Ex.	79.40	(2-j1)	19.85	2+j50
LMR400	96	MFJ	51.30	(1+j0)	12.83	0+j51
LMR400	96	Rig Ex.	51.43	(1+j2)	12.84	2+j52
2 by #22 AWG twisted pair	60	MFJ	63.40	(5+j0)	15.85	3+j111
2 by #22 AWG twisted pair	60	Rig Ex.	63.45	(4-j1)	15.86	6+j110

Figure 2 — Various short-circuit terminations.

accuracy. Fortunately, most coax cables have characteristic impedance near 50 Ω where the $(R + jX)$ meters perform best. The frequency range of the $(R + jX)$ meter places a limitation on the cable length that can be evaluated. A maximum length is determined by the lower frequency limit of the meter for $F2$, and a minimum length is determined by the high frequency limit of the meter for $F1$. Since meter accuracy is usually best near the low end of its range, best results will be obtained with longer test cables. — *Best regards, John Flood, K4DLX*; **k4dlx@bellsouth.net**.

Send your short *QEX* Technical Note to the Editor, via e-mail to qex@arrl.org. We reserve the right to edit your letter for clarity, and to fit in the available page space. "*QEX* Technical Note" may also appear in other ARRL media. The publishers of *QEX* assume no responsibilities for statements made by correspondents.

UNITED STATES POSTAL SERVICE®

Statement of Ownership, Management, and Circulation (All Periodicals Publications Except Requester Publications)

1. Publication Title	2. Publication Number	3. Filing Date
QEX	0 8 8 6 - 8 0 9 3	October 1, 2017
4. Issue Frequency	5. Number of Issues Published Annually	6. Annual Subscription Price
Bi-Monthly: Jan/Mar/May/July/Sept/Nov	6	$29.00

7. Complete Mailing Address of Known Office of Publication *(Not printer) (Street, city, county, state, and ZIP+4®)*
225 Main Street, Newington, Hartford County, CT 06111-1494
Contact Person: Yvette Vinci
Telephone *(Include area code)*: 860.594.0257

8. Complete Mailing Address of Headquarters or General Business Office of Publisher *(Not printer)*
225 Main Street, Newington, CT 06111-1494

9. Full Names and Complete Mailing Addresses of Publisher, Editor, and Managing Editor *(Do not leave blank)*
Publisher *(Name and complete mailing address)*
American Radio Relay League, Inc., 225 Main Street, Newington, CT 06111-1494
Editor *(Name and complete mailing address)*
Kazimierz "Kai" Siwiak, 225 Main Street, Newington, CT 06111-1494
Managing Editor *(Name and complete mailing address)*
Kazimierz "Kai" Siwiak, 225 Main Street, Newington, CT 06111-1494

10. Owner *(Do not leave blank. If the publication is owned by a corporation, give the name and address of the corporation immediately followed by the names and addresses of all stockholders owning or holding 1 percent or more of the total amount of stock. If not owned by a corporation, give the names and addresses of the individual owners. If owned by a partnership or other unincorporated firm, give its name and address as well as those of each individual owner. If the publication is published by a nonprofit organization, give its name and address.)*

Full Name	Complete Mailing Address
American Radio Relay League, Inc	225 Main Street, Newington, CT 06111-1494

11. Known Bondholders, Mortgagees, and Other Security Holders Owning or Holding 1 Percent or More of Total Amount of Bonds, Mortgages, or Other Securities. If none, check box → ☒ None

Full Name	Complete Mailing Address

12. Tax Status *(For completion by nonprofit organizations authorized to mail at nonprofit rates) (Check one)*
The purpose, function, and nonprofit status of this organization and the exempt status for federal income tax purposes:
☒ Has Not Changed During Preceding 12 Months
☐ Has Changed During Preceding 12 Months *(Publisher must submit explanation of change with this statement)*

PS Form **3526,** July 2014 *[Page 1 of 4 (see instructions page 4)]* PSN: 7530-01-000-9931 **PRIVACY NOTICE:** See our privacy policy on *www.usps.com.*

13. Publication Title			14. Issue Date for Circulation Data Below	
QEX			Sept/Oct 16-July/Aug 17	Sept/Oct 2017
15. Extent and Nature of Circulation			**Average No. Copies Each Issue During Preceding 12 Months**	**No. Copies of Single Issue Published Nearest to Filing Date**
a. Total Number of Copies *(Net press run)*			6800	6700
b. Paid Circulation *(By Mail and Outside the Mail)*	(1)	Mailed Outside-County Paid Subscriptions Stated on PS Form 3541 (Include paid distribution above nominal rate, advertiser's proof copies, and exchange copies)	4328	4260
	(2)	Mailed In-County Paid Subscriptions Stated on PS Form 3541 *(Include paid distribution above nominal rate, advertiser's proof copies, and exchange copies)*	0	0
	(3)	Paid Distribution Outside the Mails Including Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Paid Distribution Outside USPS®	1110	1062
	(4)	Paid Distribution by Other Classes of Mail Through the USPS (e.g., First-Class Mail®)	549	472
c. Total Paid Distribution *[Sum of 15b (1), (2), (3), and (4)]*			5986	5794
d. Free or Nominal Rate Distribution *(By Mail and Outside the Mail)*	(1)	Free or Nominal Rate Outside-County Copies included on PS Form 3541	57	53
	(2)	Free or Nominal Rate In-County Copies Included on PS Form 3541	0	0
	(3)	Free or Nominal Rate Copies Mailed at Other Classes Through the USPS (e.g., First-Class Mail)	83	70
	(4)	Free or Nominal Rate Distribution Outside the Mail *(Carriers or other means)*	82	82
e. Total Free or Nominal Rate Distribution *(Sum of 15d (1), (2), (3) and (4))*			222	205
f. Total Distribution *(Sum of 15c and 15e)*			6209	5999
g. Copies not Distributed *(See Instructions to Publishers #4 (page #3))*			591	701
h. Total *(Sum of 15f and g)*			6800	6700
i. Percent Paid *(15c divided by 15f times 100)*			96.42%	96.58%

* If you are claiming electronic copies, go to line 16 on page 3. If you are not claiming electronic copies, skip to line 17 on page 3.

UNITED STATES POSTAL SERVICE®

Statement of Ownership, Management, and Circulation (All Periodicals Publications Except Requester Publications)

16. Electronic Copy Circulation	**Average No. Copies Each Issue During Preceding 12 Months**	**No. Copies of Single Issue Published Nearest to Filing Date**
a. Paid Electronic Copies		
b. Total Paid Print Copies (Line 15c) + Paid Electronic Copies (Line 16a)		
c. Total Print Distribution (Line 15f) + Paid Electronic Copies (Line 16a)		
d. Percent Paid (Both Print & Electronic Copies) (16b divided by 16c × 100)		

☒ **I certify that 50% of all my distributed copies (electronic and print) are paid above a nominal price.**

17. Publication of Statement of Ownership

☒ If the publication is a general publication, publication of this statement is required. Will be printed in the **November/December 2017** issue of this publication.

☐ Publication not required.

18. Signature and Title of Editor, Publisher, Business Manager, or Owner	Date
[signature]	09/27/17

I certify that all information furnished on this form is true and complete. I understand that anyone who furnishes false or misleading information on this form or who omits material or information requested on the form may be subject to criminal sanctions (including fines and imprisonment) and/or civil sanctions (including civil penalties).

 PRIVACY NOTICE: See our privacy policy on *www.usps.com.*