QEX

$1.75

ARRL Experimenter's Exchange

December 1994

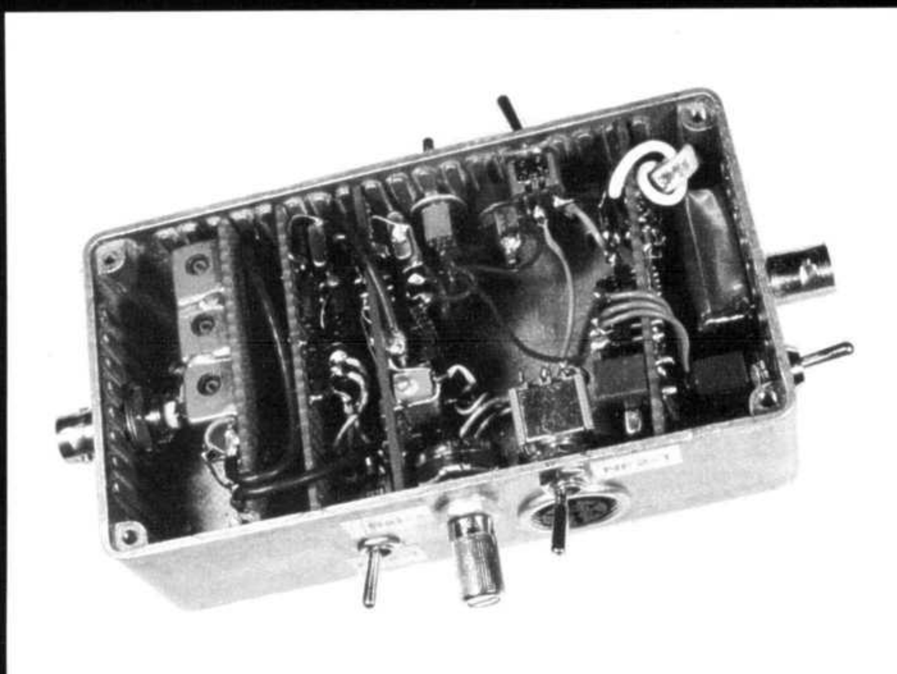

A Multitone Test-Signal Generator

QEX: The ARRL
Experimenter's Exchange
American Radio Relay League
225 Main Street
Newington, CT USA 06111

QEX (ISSN: 0886-8093 USPS 011-424) is published monthly by the American Radio Relay League, Newington, CT USA.

Second-class postage paid at Hartford, Connecticut and additional mailing offices.

David Sumner, K1ZZ
Publisher

Jon Bloom, KE3Z
Editor

Lori Weinberg
Assistant Editor

Harold Price, NK6K
Zack Lau, KH6CP
Contributing Editors

Production Department

Mark J. Wilson, AA2Z
Publications Manager

Michelle Bloom, WB1ENT
Production Supervisor

Sue Fagan
Graphic Design Supervisor

Joe Costa
Technical Illustrator

Joe Shea
Production Assistant

Advertising Information Contact:

Brad Thomas, KC1EX, *Advertising Manager*
American Radio Relay League
203-667-2494 direct
203-666-1541 ARRL
203-665-7531 fax

Circulation Department

Debra Jahnke, *Manager*
Kathy Fay, N1GZO, *Deputy Manager*
Cathy Stepina, *QEX Circulation*

Offices

225 Main St, Newington, CT 06111-1494 USA
Telephone: 203-666-1541
Telex: 650215-5052 MCI
FAX: 203-665-7531 (24 hour direct line)
Electronic Mail: MCIMAILID: 215-5052
Internet:qex@arrl.org

Subscription rate for 12 issues:

In the US: ARRL Member $12, nonmember $24;

US, Canada and Mexico by First Class Mail: ARRL Member $25, nonmember $37;

Elsewhere by Surface Mail (4-8 week delivery): ARRL Member $20, nonmember $32;

Elsewhere by Airmail: ARRL Member $48, nonmember $60.

QEX subscription orders, changes of address, and reports of missing or damaged copies may be marked: *QEX* Circulation. Postmaster: Form 3579 requested. Send change of address to: American Radio Relay League, 225 Main St, Newington, CT 06111-1494.

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

About the Cover

DL7IY's handy test-signal generator can be used to test receivers and transmitters.

ISSUE NO. 154

Features

Columns

December 1994 QEX Advertising Index

THE AMERICAN RADIO RELAY LEAGUE

The American Radio Relay League, Inc, is a noncommercial association of radio amateurs, organized for the promotion of interests 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 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 two 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 at 225 Main Street, Newington, CT 06111 USA.

Telephone: 203-666-1541
Telex: 650215-5052 MCI
MCIMAIL (electronic mail system) ID: 215-5052
FAX: 203-665-7531 (24-hour direct line)

Officers

President: GEORGE S. WILSON III, W4OYI
1649 Griffith Ave, Owensboro, KY 42301

Executive Vice President: DAVID SUMNER, K1ZZ

Purpose of *QEX*:

1) provide a medium for the exchange of ideas and information between Amateur Radio experimenters
2) document advanced technical work in the Amateur Radio field
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 correspondence for publication in *QEX* should be marked: Editor, *QEX*.

Both theoretical and practical technical articles are welcomed. Manuscripts should be typed and doubled spaced. Please use the standard ARRL abbreviations found in recent editions of *The ARRL Handbook*. Photos should be glossy, black and white positive prints of good definition and contrast, and should be the same size or larger than the size that is to appear in *QEX*.

Any opinions expressed in *QEX* are those of the authors, not necessarily those of the editor or the League. While we attempt to ensure that all articles are technically valid, authors are expected to defend their own material. Products mentioned in the text are included for your information; no endorsement is implied. The information is believed to be correct, but readers are cautioned to verify availability of the product before sending money to the vendor.

Empirically Speaking

Overseas *QEX* for Less

For some time we've been hearing from our friends outside North America that, while they enjoy reading *QEX*, they distinctly *don't* enjoy paying $48 or $60 per year for the privilege of doing so. Those rates, for members and nonmembers, respectively, aren't intended to gouge our offshore readers. Rather, they reflect the high cost of getting *QEX*—or any magazine—mailed for timely delivery from the US to other continents via air mail.

What many have told us, particularly in Europe, is that they would happily opt for slower delivery if we could bring the subscription price down substantially. At the same time, there are still those readers who want their *QEX* right away, despite the cost. So, we've decided to let the reader make the decision. Beginning immediately, new and renewal subscribers to *QEX* from outside North America can opt to have their copy delivered by surface mail, at an annual cost of $20 for ARRL members and $32 for nonmembers. This is a substantial savings over the airmail rates. But for those in a hurry, the airmail rates—and the consequent fast delivery—are still available.

Of course, those who do opt for the surface rates have to understand that surface mail is *slow*. If you subscribe to *QEX* at the lower rate, don't be surprised if your copy takes 1 to 2 months to reach you. You can have it fast or you can have it cheap, but we've not found a way to give it to you fast *and* cheap.

More on HF Simulators

Last month in this space we noted the activity on the TAPR HF-SIG mailing list, which is discussing the possibility of developing a DSP-based HF channel simulator to use to test modem designs and implementations. We're happy to say that the work is progressing. Those participating are rapidly finding the relevant published material, including CCIR Recommendations, papers in the professional literature and books that address the subject of how to implement channel simulation.

While it's still a bit early—no code has been written yet, as far as we know—it is becoming clear that the problem is a tractable one for the current generation of low-cost DSP hardware available to hams. We have little doubt that the next few months will see an HF channel simulator implemented and running.

One of the lessons we are all relearning from this effort is that putting multiple minds on a project can make it easier, and that effective digital communications—via the Internet, in this case—makes the meeting of minds easier to achieve. Now if only we could carry out these discussions via an *amateur* digital network. Sigh.

This Month in *QEX*

It's one thing to build a home-brew receiver or transmitter. It's another to know how it stacks up performance-wise. "The MTG1 Multitone Test Generator," by Detlef Rohde, DL7IY, provides audio and 80-meter two-tone signals and an HF comb generator, all of which can help you test your latest creation.

How well do these modern HF digital modes really work? A partial answer is provided this month by Peter Reynolds, KE4BAD, who reports on some "HF Channel Simulator Tests of Clover."

Been "Hearing Strange UHF Signals Lately?" If so, maybe your 70-cm preamp can't handle the high levels of commercial signals in nearby bands. John Reed, W6IOJ, provides an integrated preamp and filter design that may help.

There is much confusion among amateurs as to where the reflected power on a transmission line ends up. The answer is: in the load. But why? This month, your editor, Jon Bloom, KE3Z, takes his shot at explaining "Where Does the Power Go?"

In this issue, Ladimer S. Nagurney, WA3EEC, contributes a review of the book: *Digital Signal Processing in Communication Systems*, by Marvin E. Frerking. Does the book tell you how to generate and process modulated signals with DSP? Read the review!

This month's "Proceedings" column lists the papers in the *Proceedings of the AMSAT-NA 12th Space Symposium and AMSAT Annual Meeting*.

Finally, in his "Digital Communications" column this month, Harold Price, NK6K, reports the interesting reactions of several readers to his recent gloom-and-doom columns about packet.—*KE3Z, email:jbloom@arrl.org (Internet)*

The MTG1 Multitone Test Generator

A clean multitone signal generator is needed to test home-brew equipment. Here's one that's easy to build.

by Detlef Rohde, DL7IY

Described here is a test generator that is useful for checking the performance of short-wave transmitters and receivers. It consists of the TTG1, an audio frequency two-tone generator that was presented in the German Amateur Radio magazine *cq-DL*.[1] This device is useful for linearity measurements of SSB transmitters and amplifiers.

The MTG1 adds an RF two-tone signal that is necessary for testing of a receiver's large-signal capabilities. Also included in the same box is a device called a "needle-pulse generator" that delivers a comb of constant amplitude marks (–73 to –67 dBm) spaced 50 or 100 kHz apart. This part of the MTG1 is useful for receiver testing and other measurements. With

[1]Notes appear on page 6.

Titiseestrasse 12
D13469 Berlin Germany

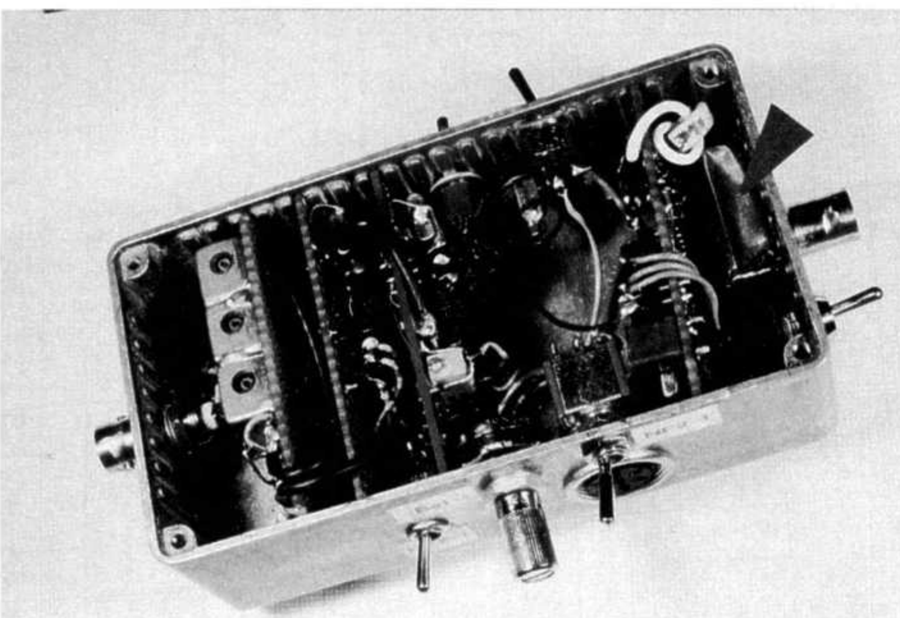

modern components, a very compact assembly of the unit is possible as shown in the photo. Fig 1 shows the circuit diagram.

Mode of Operation

The MTG1 contains three independent crystal-stabilized oscillators from which the output signals are derived and sent to three separate jacks.

The audio frequencies, F1 and F2, are derived from a reference oscillator operating at 14.24 MHz. As in the TTG1, low-pass filtering of two square-wave signals provides a two-tone sine-wave signal with very low distortion. A very simple CMOS oscillator chip with binary dividers (4060) generates a clock signal at 222.5 kHz for the switched-capacitor filter, a MAX293 from Maxim. The 3-dB corner frequency of the low-pass is set by this clock signal to 2.225 kHz. Square-wave frequencies of 1.738 kHz (F2) and 13.9 kHz are available on pin 2 and pin 13 of the oscillator chip, respectively.

To get two non-harmonic output frequencies which fit into the window of an SSB-filter, the 13.9-kHz wave is

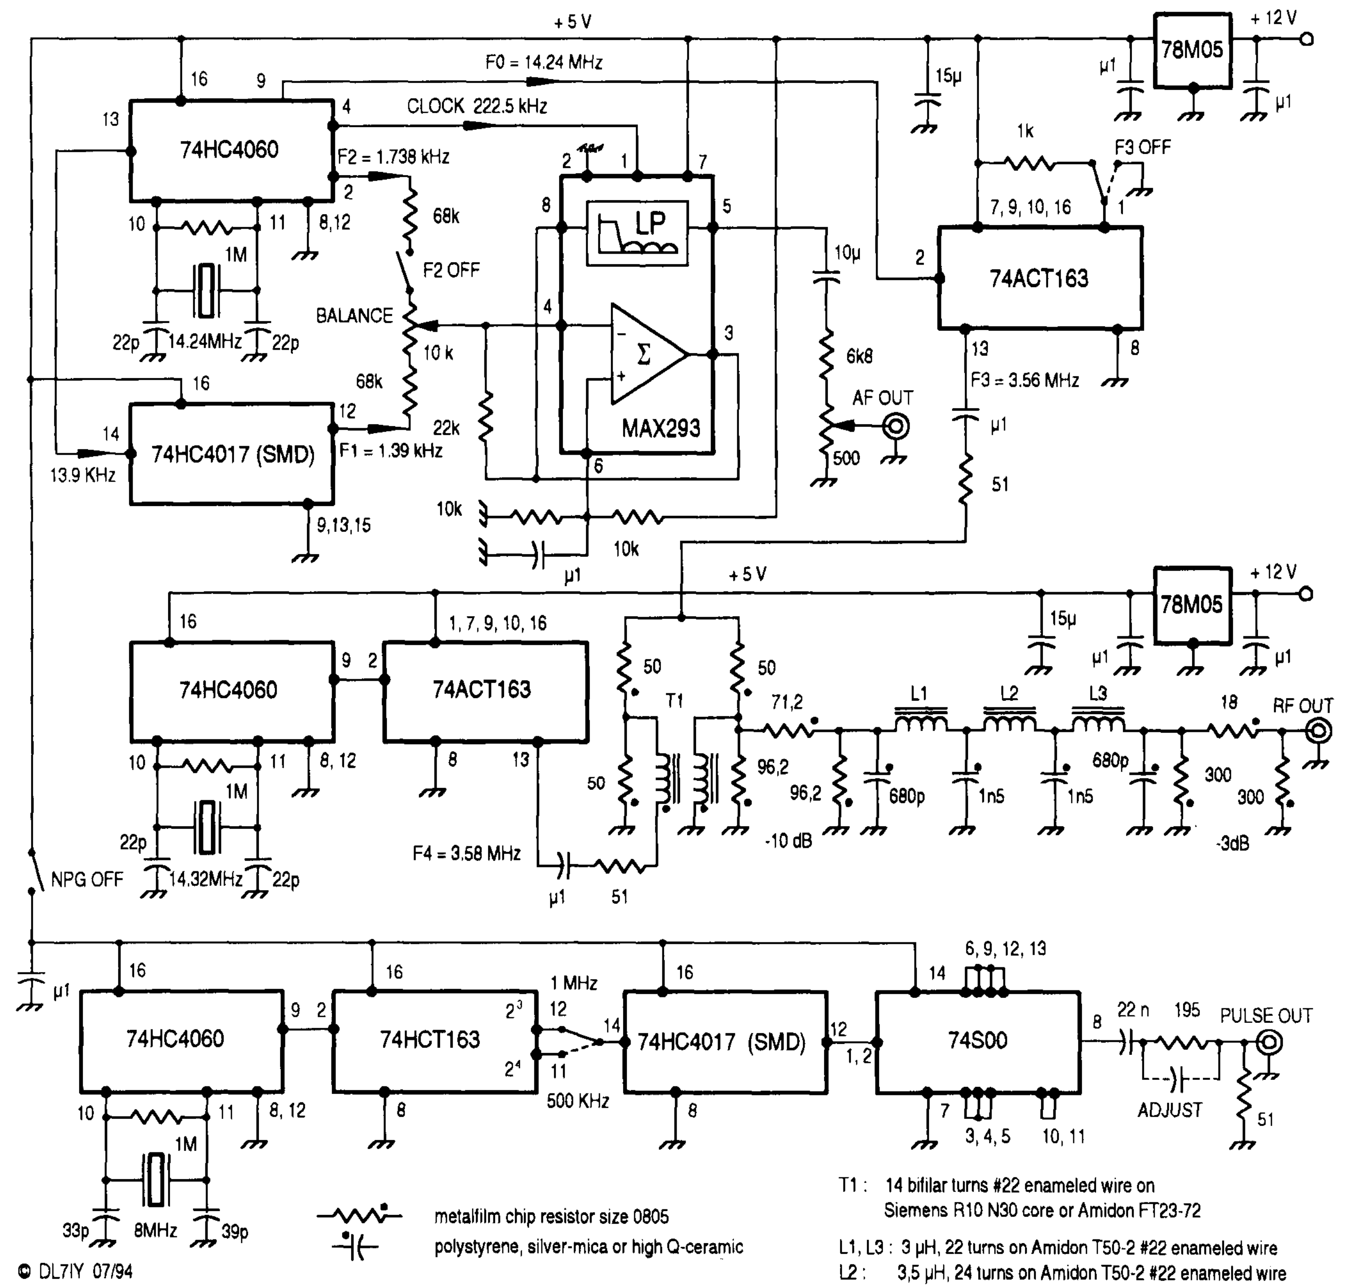

Fig 1

divided by 10 for an output frequency of 1.39 kHz (F1). Normally, SSB filters have bandwidths between 2 and 3 kHz, so modulation products of the upper harmonics of the basic frequencies, F1 and F2, are suppressed by the SSB filter and will not occur at the final stage of the transmitter if there is good sideband filter performance and good linearity in the signal path. For one-tone measurements, frequency F2 may be switched off. An RF two-tone signal is generated as follows. The reference frequency of 14240 kHz, available on pin 9 of the oscillator chip, is selected and divided by 4 using a very fast binary divider (74ACT163). A TTL-level square wave at 3560 kHz (F3) appears at pin 13 of the 74ACT163. Another reference oscillator delivers the second RF frequency of 3580 kHz (F4) in the same manner. This oscillator operates at 14320 kHz.

A passive combiner bridge is used to add the two square-wave signals.[2] Series 50-Ω–resistors minimize the possibility of internal intermodulation in the output stages of the binary dividers. A 10-dB attenuator is inserted into the signal path at the output of the combiner so that the symmetry of the bridge, which is carefully adjusted by use of selected metal-film chip resistors, is less sensitive to load variations.

After the signals are combined, a 7-element low-pass filter with a Chebyshev characteristic is used to suppress the upper harmonics of the signals in order to get a sinusoidal output.

The elements of this filter should not be overdriven by the signals on its input (–3 dBm each). This especially refers to the coils used, as their cores might be saturated causing the generation of harmonics and IMD products. Iron-powder toroids (Amidon T50-2) are a good choice, although fine tuning of the inductance is not possible with toroids. Alignment for best harmonic suppression performance must then be done by correcting the final values of the capacitors used. If no spectrum or network analyzer is available, the station receiver can be used to find the levels of the harmonics. Be careful! The receiver might be overdriven by the strong –7-dBm input signals causing spurious frequencies not coming from the generator. Use a step attenuator in front of the receiver's input and increase the attenuation until no more reduction of harmonics is achieved.[3] (This device is also necessary for most of the measurements possible with the MTG1, such as measuring the third-order intercept point, IP3, of a receiver.) After finding an appropriate input power level, tune the filter for best harmonic suppression. At the filter output, a fixed attenuator of 3 dB reduces the influence of receivers that don't present a 50-Ω impedance. The output level of –7 dBm of each carrier is sufficient to cause IMD products in amateur receivers. A spectrum analysis of the generator output shows no detectable spurious signals within the dynamic range of the analyzer (Fig 2). A time-domain analysis of the two-tone signal is shown in Fig 3 as it would appear on an oscilloscope.

A high-level receiver front end with IP3=36 dBm, developed by Michael Martin, DJ7VY, was used to determine the generator's own IP3, which proved to be 26.5 dBm. This is adequate for most amateur-radio gear but not for commercial equipment, where higher IMD dynamic ranges of >90-dB are necessary for measuring the state-of-the-art performance of modern (high-priced) professional receivers. To make such measurements, generators of 100-W output can be attenuated and combined to achieve a higher IMD dynamic range.

The marker comb generator in the MTG1 is similar to those published by other authors.[4] Another reference oscillator, operating at 8-MHz, delivers selectable 500-kHz or 1-MHz square-wave signals after binary division by 16 (74HCT163). A decade divider (74HC4017) provides an output signal of 50 or 100 kHz. To form sharp needle pulses, a classic pulse shaper is realized using a fast TTL NAND gate (74S00). The short propagation delay leads to output pulses of about 4 ns or less. This is not an ideal needle pulse that has an infinite spectrum of harmonics in the frequency domain, but it's sufficient for measurements in the short-wave range. The output signal is adjusted to a level of –67 dBm for 100 kHz marks by use of a fixed attenuator. The 50-kHz marks are 6 dB below this level at –73 dBm. The accuracy of the measurements made with the needle-pulse generator

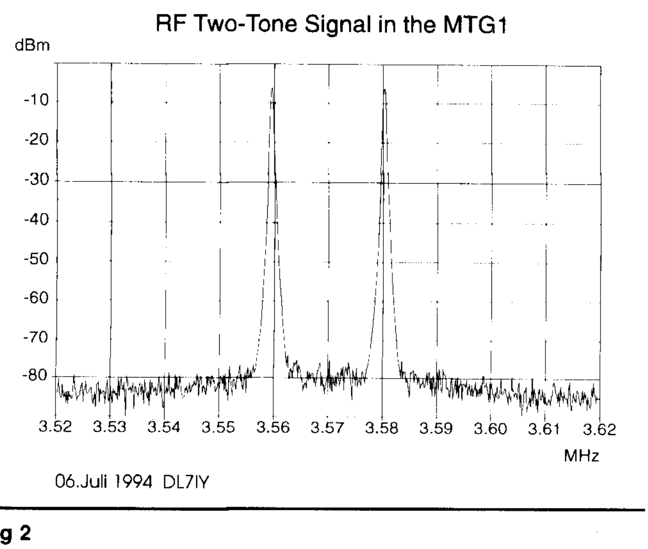

Fig 2

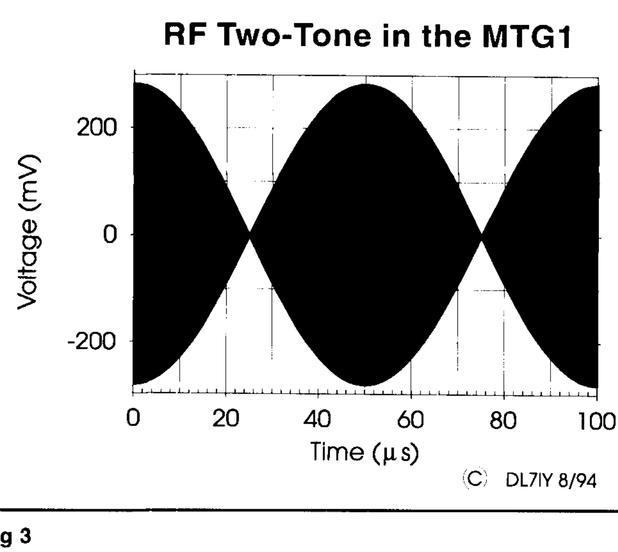

Fig 3

depends mainly on the correct level of the marks. There may be a difference of about 0.5 dB between odd and even harmonics, so the uncertainty of measurements will be around that much.

Measurements with the MTG1

The 350-Hz spacing of two-tone audio signals from the MTG1 is good for oscilloscope monitoring of the transmitted signals. You can easily see the effects of overdriving or of improperly adjusted operating points in your so-called linear amplifiers. Following the signal with an oscilloscope, you should see a display like that of Fig 3 in the IF or final-frequency range when all stages are properly adjusted. It is easy to detect a wrong operating stage when you find a difference from this waveform. The output level of the two-tone signal from the MTG1 should be adjusted to the same level as the one from a microphone.

The low phase noise of the MTG1 signals allows testing of the phase-locked oscillators in a receiver. By watching the various beat frequencies of the needle-pulse generator, for example, I detected a badly locked VCO in my ICOM IC-735. It was easy to hear that some markers sounded noisy or jittered. The use of two-tone tests for testing receiver performance is well known and described in the *ARRL Handbook*. The very clean RF signals of the MTG1 with their high spectral purity also show the bad performance of some amateur receivers that are equipped with noisy PLL synthesizers. A couple of years ago I measured an FT-757GX from Yaesu. You could see that the S-meter did not go down to the S9 mark between the two 20-kHz spaced carriers. It might be impossible to detect weak signals when operating between strong signals as a consequence of using such equipment.

The frequency responses of receivers or amplifiers may be determined by measuring the amplitudes of the needle-pulse generator spectral lines at the output of the device with a level meter. Your receiver's S-meter may be used to see if your receiver is less sensitive in any range. The marker generator is also useful for noise-floor measurements. It may also be used as a broad-band signal source in the same manner as a noise bridge. In connection with a frequency-selective detector (the station receiver), RF bridge and step-attenuator, impedance-matching measurements are possible. It is much more comfortable to watch a weak signal on a defined frequency than to listen for differences in a noise signal in order to measure the high return loss of a well-matched device.

Hints for Construction

The MTG1 was built in a small aluminum box with slots in the side walls to serve as holders for the printed-circuit boards. The PC boards are double-sided universal boards with a ground plane on one side. The components used are, in many cases, SMD-type to minimize the size of the MTG1. In practical measurements, the three parts of the generator are used separately, so it is not necessary to shield between them. The active section may be selected by switching the supply voltages. Good shielding is necessary at the output of the needle-pulse generator (see arrow in the photo). Make sure that this part is switched off when performing measurements with AF or RF two-tone generators.

The combiner and output filter are mounted together with the input and output attenuators on a separate board. When using tunable coils, coupling between them must be avoided by the use of good metal shielding, which is not necessary when using toroids.

Conclusion

All the circuits in the MTG1 are variations of circuits already published in amateur-radio literature.[5] This solution is easy to build compared to other published designs because there are fewer steps in alignment and because of the use of modern integrated circuits that are cheap and easy to find. A diminutive radio shack may become a serious laboratory when the operator owns such measuring devices. They are of special importance when there is no possibility of access to such high-tech devices as spectrum or network analyzers. But this equipment is fine for calibrating your own measuring aids, which can then be of much help to the home-brewing amateur who wants to measure the performance of his self-constructed equipment. The MTG1 may help you decide whether or not you're making a good choice when buying used or new equipment for your station.

The author is member of the technical staff of the Heinrich-Hertz-Institut Berlin, known world-wide as a research facility, especially in the field of fiber-optic communication. He is grateful for being able to use the infrastructure of this professionally equipped working place to finish this work.

Notes

[1]Rohde, Detlef, DL7IY, "Zweitontestgenerator TTG1," *cq-DL* October 1994.

[2]*The ARRL Handbook for the Radio Amateur*, 1994, p 2-45.

[3]*The ARRL Handbook for the Radio Amateur*, 1994, beginning with p 25-37.

[4]Waxweiler, Richard, DJ7VD, "Eichmarkengeber mit definierter Markenamplitude...," *cq-DL*, August 1978, pp 348-349.

[5]Waxweiler, R., "Hochfrequenz-Zweitongenerator", *cq-DL*, September 1980, pp 412-414.

HF Channel Simulator Tests of Clover

A look at the performance of some Clover modulation formats.

by Peter Reynolds, KE4BAD

Introduction

Clover is a new waveform for data transmission for HF ("waveform," as I use it here, comprises the modulation and coding techniques used to transmit data). It was developed by Ray Petit, W7GHM, specifically to overcome the limitations of the modes of HF data communication currently widely used by amateurs. The limitations of RTTY, AMTOR (including its derivatives SITOR and Pactor), and AX.25-based packet radio for data transmission over HF are well known.[1] Clover combines sophisticated modulation techniques with error-correction-and-detection coding and error-control protocols. The result is a waveform that efficiently uses the radio spectrum to provide error-free throughput in a narrow bandwidth over a wide range of HF channel conditions. In this article I will describe a brief test of the HAL PCI-4000 PC-board implementation of Clover. The testing was conducted at SHAPE Technical Centre using an HF channel simulator.

Clover has evolved from the original Clover-I waveform into Clover-II. I will not discuss Clover-I further and will use Clover to refer to the current Clover-II waveform implemented in the HAL PCI-4000 product. More details on Clover-I can be found in Note 2.

Many HF radio users (both inside and outside the amateur community) still use 75-baud FSK radio teletype (RATT) or some minor variation for HF data communication. Clover came to my attention because of a program underway within the North Atlantic Treaty Organization (NATO) to upgrade the FSK currently used in RATT systems to a more modern waveform. The upgrade program is to be implemented incrementally; the first step will be to simply adopt a new modem without immediately changing the existing procedures and other equipment.

A new modem for this application must consider the bandwidth of the available HF channels, as well as compatibility with existing equipment. Clover is designed to operate in a –50-dB bandwidth of 500 Hz . Much of the maritime mobile HF spectrum has been reallocated into 500-Hz channels by a recent World Radiocommunication Conference, or WRC. We decided to evaluate the performance of Clover because, as a commercially available product offering data transmission on HF in a widely available bandwidth, it offered a potential low-cost solution for maritime mobile HF data communication. Even if not adopted directly for use, Clover offers a good data point for comparison for other HF data modems that use more bandwidth, such as those defined in

[1]Notes appear on page 12.

Middelweg 9, Wassenaar 2242BK
The Netherlands

MIL-STD-188-110A and STANAG 4285.[3,4] Although our users are military, it should be noted that the 500-Hz bandwidth applies to all new maritime mobile HF users, many of which are commercial.

The Clover system offers a significant improvement over the current waveforms for HF data communication used by maritime mobile users and over the systems currently in wide use in the amateur community. In the crowded amateur HF bands, as in other parts of the HF spectrum, any more efficient use of the spectrum deserves close attention.

Clover Description

The Clover modem described in this article is implemented on an ISA-bus PC card, the HAL PCI-4000. Software (for DOS) is provided to control the modem and to provide a data interface via the PC bus. Data I/O is over the PC bus only; an external data I/O port is not provided. The card has connectors for a radio interface (audio I/O, PTT, SEL-CAL and tuning indicator), and appropriate cables are provided. A detailed manual is also provided, which covers setup and operation as well as details of the system. The following sections give a brief overview of the Clover system including modulation, coding, and ARQ. For a more detailed description, see Notes 1 and 2.

Modulation

The Clover output is an audio signal intended to drive an SSB transmitter. Clover uses four subcarriers at 2062.5 Hz, 2187.5 Hz, 2312.5 Hz and 2437.5 Hz. Each subcarrier is constantly on-off keyed at 31.25 baud, timed such that the pulses on successive frequencies reach their peak amplitude at 8-ms intervals. The data is encoded by the difference in phase and/or amplitude of successive pulses on a single subcarrier. Changes in the phase and/or amplitude of the subcarrier occur only when the instantaneous amplitude of that subcarrier is zero due to the on-off keying. This technique avoids the broad spectrum usually associated with phase modulation. To distinguish between classical FSK and PSK modulation, in which a single audio subcarrier is shifted in frequency or phase, and the Clover technique, in which changes in frequency or phase occur only in the zero-amplitude intervals of the constant on-off keying, the Clover modulation techniques are called frequency shift modulation (FSM) and phase shift modulation (PSM). Clover uses eight different forms of modulation, varying in throughput and robustness. Each will be described here briefly. A general description and discussion of different types of modulation may be found in Notes 5 and 6.

The most robust Clover modulation (also having the lowest throughput) is binary PSM (BPSM) with four-fold diversity. Each subcarrier is modulated with BPSK, meaning that the phase of a pulse can have two values, one indicating binary "one" and the other "zero." The result is 31.25 bit/s on each subcarrier. The data is repeated on each of the four subcarriers, giving four-fold diversity and a total throughput of 31.25 bit/s. The second modulation is binary "FSM" with two-fold diversity, in which each of the two pairs of subcarriers are used to transmit data in a fashion similar to an FSK tone pair. Because there are effectively two 31.25-baud FSK signals carrying redundant information, the throughput is also 31.25 bit/s (the same as for the four-fold diversity BPSM).[7] The third modulation is BPSM, in which each subcarrier carries different data, for a net throughput of 4×31.25=125 bit/s. Fourth is QPSM, in which each symbol can be one of four phases; each phase represents a pair of bits. The throughput is therefore 4×31.25×2=250 bit/s. Next is 8-ary PSM; here each symbol can be one of eight phases, each representing three bits. The throughput is then 4×31.25×3=375 bit/s. The next two modulations have the same throughput of 500 bit/s. Each also has sixteen symbols, each symbol carrying four bits. One uses sixteen different phases; the other uses eight different phases combined with two different amplitudes. Finally, there is a modulation using sixteen different phases and four different amplitudes, for a total of 64 different symbols, each carrying six bits. The total throughput of this mode is then 4×31.25×6=750 bit/s (perhaps giving new meaning to the expression "making hay while the sun shines!").

Error-Correction Coding

The above throughput figures do not include any error-correction coding. Error-correction coding involves the systematic addition of redundancy to the transmitted data, which allows some number of errors to be corrected at the receiving end of the link. This technique is often called "forward error correcting" or FEC coding. Clover includes a form of FEC coding called *Reed-Solomon* coding. A discussion of the principles and advantages of Reed-Solomon coding is beyond the scope of this article but can be found in Note 8. Suffice it to say that Reed-Solomon, or RS, coding introduces redundant bits into blocks of data in a complex way, which allows a certain number of errors in received data to be detected and/or corrected. As more redundant bits are added, the throughput that the user sees falls, but the number of errors that can be detected or corrected increases. Clover includes RS codes that use 10%, 25% or 40% of the available bits for error correction. The coding can also be turned off, in which

Table 1. Clover Data Rates

	4× BPSM	*2×"FSM"*	*BPSM*	*QPSM*	*8PSM*	*16PSM*	*8P/2ASM*	*16P/4ASM*
Raw bit/s	31.25	31.25	125	250	375	500	500	750
Code rate								
100%	31.25	31.25	125	250	375	500	500	750
90%	28.13	28.13	112.5	225	337.5	450	450	675
75%	23.44	23.44	93.75	187.5	281.25	375	375	562.5
60%	18.75	18.75	75	150	225	300	300	450

case no bits are added at all. Looking at it another way, the user sees a channel that provides between 60% and 100% of the throughput figures calculated above. The ratio of the coded bit rate to the uncoded bit rate is called the "code rate." These factors are summarized in Table 1.

It should be noted that the most robust modes are those toward the bottom and left side of the table; the channel must be quite good to support modes to the right and top of the table. In our tests, the code rate was fixed at 60%; therefore only the data rates in the bottom row of Table 1 were available. The data rates are the lowest available but can correct the largest number of errors. This attempts to compensate for the fact that errors cannot be retransmitted. Fortunately, this code rate also provides user data rates of 75, 150 and 300 bit/s, which are compatible with the maritime systems currently in use.

ARQ Mode

The Clover system includes, and is generally intended to be used with, an automatic repeat request, or ARQ, protocol. An ARQ protocol divides the data to be transmitted into blocks. A form of error detection, usually a checksum or a cyclic redundancy check (CRC), is computed and added to each block of data. The checksum is simpler to implement, but the CRC provides better error detection. The data blocks (including the checksum or CRC) are then transmitted. When received, the error detection code is recalculated based on the received data. If there are errors in the received data, the error detection code calculated at the receiving end will not match that sent with the data. The receiving ARQ protocol then knows that an error has occurred and requests that the block in error be retransmitted.

The above is a fairly simple description of an ARQ protocol; quite a lot of technical detail is omitted. More detailed descriptions may be found in Note 9.

The Clover system uses the feedback inherent in an ARQ protocol to adapt to changing channel conditions. If a significant number of data blocks are received with errors, the system shifts to a more robust modulation and coding scheme (down and to the left in Table 1). If no errors are detected, the system shifts to a faster throughput scheme (up and to the right in Table 1). This adaptivity makes modulations such as 16-phase, 4-amplitude PSM (occasionally) usable on HF. While the channel is good enough, such a modulation can provide high throughput in the bandwidth. When the channel deteriorates, a more robust scheme maintains a reasonable throughput.

The Clover system combines ARQ with the RS error-correction code. Only errors that are not correctable by the RS code will cause an ARQ retransmission.

Because of our application, the ARQ mode of the Clover equipment was not used in our tests. This is because the problem for which we are defining a solution requires a modem that can be used as a direct replacement for existing modems in existing systems. These existing systems are not compatible with the variable delay introduced by ARQ protocols.

Channel Simulator Tests

Test Objectives

The immediate goal of the tests that we conducted was to determine the performance in terms of bit error rate (BER) as a function of E_b/N_o achievable in a 500-Hz bandwidth from a commercially available product. This information was needed to support the selection or development of an HF modem to replace existing 75-baud FSK modems in maritime communication systems.

The other candidates for this application were the MIL-STD-188-110A and STANAG 4285 waveforms, described in Notes 3 and 4, respectively. These are sophisticated and expensive HF modems that use PSK modulation with a 3-kHz bandwidth: coding and DSP techniques are used to provide user data rates from 75 to 2400 bit/s over a wide range of HF channel conditions. The waveforms will be described briefly later.

Test Equipment and Procedures

The tests were carried out using the Cossor 1250 HF channel simulator. This is a baseband channel simulator, which means that the simulation of multipath and fading is carried out using the audio output of the modems. A baseband simulator is significantly simpler and less expensive than an RF simulator and provides valid results for our purposes. The test procedures are also simpler because radios are not needed. The Cossor 1250 simulator digitizes the audio from the transmitting modem and operates on the digitized signal to produce fading and multipath effects. The Cossor 1250 can simulate one channel with a single groundwave path and up to four skywave paths or two channels simultaneously, each with one groundwave and two skywave paths. Any of the paths, including the groundwave, can be turned off if desired. The characteristics of all the skywave paths are completely independent. The output of the simulator is an audio signal that is applied to the audio input of the receiving modem.

For this test, we focused on a Rayleigh fading channel with two equal amplitude paths and 1-ms relative delay, and an (independent) 2-Hz fade rate on each path. We consider this to be a "typical HF channel" and would like to be able to operate on such a channel. The first characteristic, which an effective waveform demonstrates, is an error rate that decreases fairly quickly with increasing SNR. If a waveform demonstrates, an error rate that asymptotically approaches some value with increasing SNR, the waveform has an irreducible error rate on that channel. This type of behavior is not desirable because it means that it is impossible to lower the error rate below the asymptotic value on that channel.

When the Clover performance on this channel exhibited an irreducible error rate, a different (better) channel

Table 2. HF Simulator Channels Attempted with Clover Modem

Test #	*Doppler Spread Path 1*	*Doppler Spread Path 2*	*Relative Power*	*Relative Delay (ms)*	*Test Description*
1	0	NA	Path 2 off	NA	AWGN
2	0.2 Hz	NA	Path 2 off	NA	Slow flat fading
3	2.0 Hz	NA	Path 2 off	NA	Flat fading
4	0.2 Hz	0.2 Hz	Equal	1.0	2 path Rayleigh
5	2.0 Hz	2.0 Hz	Equal	1.0	2 path Rayleigh

was selected from the list in Table 2, which is a subset of a standard set of test conditions from Note 10. The additive white gaussian noise (AWGN) channel (test 1 in Table 2) is not representative of an HF channel. It was not used in these tests, but is included because of the channel identification numbering. Tests 2 and 3 are considered to represent good HF channel conditions. Test 4 is close to the CCIR good and moderate channels. A number of test channels more severe than the CCIR poor channel are defined in Note 10 to more thoroughly evaluate the performance of modems over a range of operating conditions. We recognize that the use of nonCCIR channels will make it more difficult to compare our results with results collected elsewhere. However, the additional information provided from testing on these channels justifies the use of nonstandard channels.

Two personal computers with PCI-4000 boards installed were the data terminals for this test. The PCI-4000 software was configured for "FEC" (that is, nonARQ) mode. As has been mentioned, the PCI-4000 implementation of the Clover modem accepts data for transmission only over the PC bus. Further, only ASCII file transfer has been implemented in the Clover control software (HAL has stated that binary file transfer will be supported in future versions). This prevented the use of bit error rate (BER) test sets for direct measurement of BER. For this test, a "quick brown fox" ASCII data file with numbered lines was constructed for ease of analysis. The data file was transmitted from one PC, via the PCI-4000, over the HF channel simulator, to the other PCI-4000 board, and finally logged on the destination PC. After the test was ended, the received data file was reduced on another PC to obtain a character error rate (CER). Data in error is indicated in the received traffic by a block of underscore characters; when the data correction capability of the RS code is exceeded, the entire block of data is lost and is replaced by underscore characters. The error counting algorithm, however, did not count underscore characters; rather, it remembered the last line number read and used that information to determine the number of character errors. In this way we avoided the possibility of undercounting errors due to a loss of synchronization in low SNR conditions. Because the test was conducted with some time limitations, data points were collected at fairly high error rates (that is, low signal-to-noise ratios).

Results

The results are shown in Figs 1 and 2. As discussed above, performance is shown in terms of character error rate (CER) rather than BER because of limitations of the current Clover control software. In addition to the Clover performance curves, two additional waveforms are shown. One is the STANAG 4285 at 300 bit/s with long interleaving (described below). The other is the standard military 75-baud, 850-Hz shift FSK. In order to allow presentation of data for modems operating at different data rates and bandwidths, results are plotted against E_b/N_o rather than SNR. E_b/N_o is related to SNR by the following formula:

$$\frac{E_b}{N_o}(\mathrm{dB}) + \frac{R_b}{B}(\mathrm{dB}) = SNR(\mathrm{dB})$$

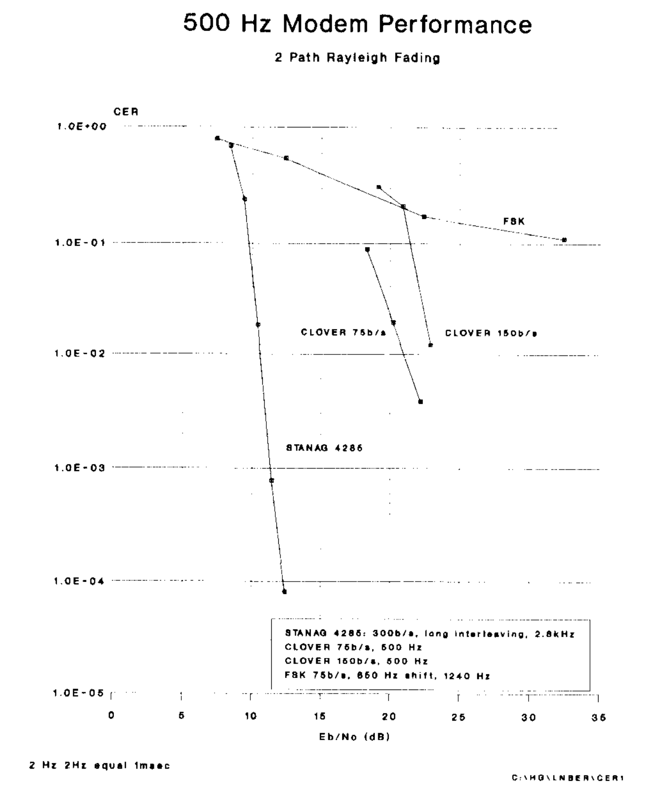

Fig 1

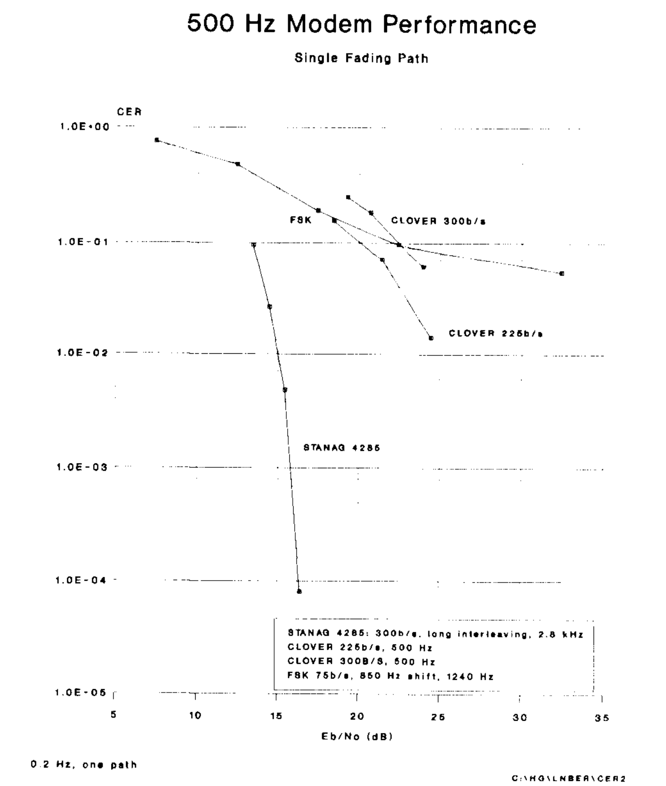

Fig 2

where R_b is the user bit rate and B is the bandwidth of the SNR. To convert from E_b/N_o to SNR, use the formula above with the R_b/B values from Table 3.

On Fig 1, only points for Clover at 75 and 150 bit/s are plotted. This is because Clover, at higher data rates, demonstrated an irreducible error rate on this channel. At 225 and 300 bit/s, we found that the poorest channel on which Clover could produce the desired "waterfall curve" was the single path channel with a 0.2-Hz fade rate. This is one of the best channels on which we conduct tests. It appears that the 2-Hz fade rate (approximately one fade every half second) has an unfortunate effect on the Clover waveform. Table 4 reviews the results of the Clover waveform tests on the different simulated HF channels in terms of whether the CER curve shows the desired "waterfall" type curve or an irreducible error rate.

Discussion

Figs 1 and 2 clearly show the superiority of the Clover waveform to FSK. In fact, on every channel in Table 2 except the AWGN channel, 75-baud FSK exhibits the same type of irreducible error rate shown in Figs 1 and 2. Very high SNR at the receiver is required to get the CER for uncoded 75-baud FSK below 10%. Clover, on the other hand, provides quite good performance relative to FSK (about 1% CER) at 75 and 150 bit/s with E_b/N_o around 23 dB.

The significant difference in performance between the 150 and 225-bit/s data rates (see Table 4) is consistent with modulation theory and with our experience with other modems. Referring to the bottom row of Table 1, the 75 and 150 bit/s data rates use BPSM and QPSM respectively, while the 225 and 300 bit/s data rates use 8PSM and 16PSM. Modulation theory predicts that BPSK and QPSK will have similar performance characteristics; 8 and 16 PSK are much more vulnerable to errors than are BPSK and QPSK. This significant change in performance predicted by theory has been observed in our extensive experience with the 3-kHz PSK modems discussed below, which also use BPSK, QPSK, and 8PSK.

The other candidate HF waveforms under consideration for the NATO upgrade were the MIL-STD-188-110A and STANAG 4285 waveforms, described in Notes 3 and 4. These are sophisticated and expensive HF modems that use a 3-kHz PSK modulation with coding and DSP techniques to provide user data rates from 75 to 2400 bit/s over a wide range of HF channel conditions. They represent the current state of the art in commercially available HF waveform performance. The MIL-STD and STANAG waveforms are similar in some respects but are not identical, nor are they interoperable. Notes 3 and 4 are intended to describe the waveforms in such a way that any manufacturer can produce a modem that will interoperate with a modem built by any other manufacturer. Because the descriptions in the references are somewhat inaccessible, they will be described here briefly.

In contrast to Clover, the above waveforms use 2400-baud PSK modulated onto a single subcarrier tone for all user data rates. User data rates from 75 to 2400 bit/s are supported (actually higher data rates are supported but no coding is provided so they are not practically usable). The modulation varies from binary PSK (BPSK) at the lower data rates to 8PSK at the higher data rates. At all data rates, fully half of the bits sent over the channel (channel bit rate = baud rate × m, where m = 1 for BPSK, m = 2 for QPSK, and m = 3 for 8PSK) are reserved for channel equalization. Equalization is a powerful DSP technique that compensates for multi-path

Table 3. E_b/N_o to SNR Conversion Factors

Modem	*R_b (bit/s)*	*B (Hz)*	*R_b/B (dB)*
Clover	75	500	−8.2
Clover	150	500	−5.2
Clover	225	500	−3.5
Clover	300	500	−2.2
STANAG 4285	300	3000	−10.0
FSK	75	300(see note)	−6.0

Note to Table 3: Although the military FSK signal uses an 850-Hz shift and a nominal occupied bandwidth of 1240 Hz, much of this bandwidth is unused. The usual FSK demodulator implementation will not see a noise bandwidth of 1240 Hz; the actual noise bandwidth depends on the demodulator implementation, but 300 Hz is a reasonable value.

Table 4. Summary of Simulated HF Channels Attempted

Test #	*75 bit/s*	*150 bit/s*	*225 bit/s*	*300 bit/s*
2	waterfall	waterfall	waterfall	waterfall
3	waterfall	waterfall	irreducible	irreducible
4	waterfall	waterfall	irreducible	irreducible
5	waterfall	waterfall	irreducible	irreducible

Note: Test number is from Table 2.

Table 5. STANAG 4285 Modulation and Coding

User data rate bit/s	*Modulation*	*Code rate*	*Coding*
2400	8 PSK	2/3	punctured rate 1/2
1200	4 PSK	1/2	rate 1/2
600	BPSK	1/2	rate 1/2
300	BPSK	1/4	rate 1/2 repeated 2×
150	BPSK	1/8	rate 1/2 repeated 4×
75	BPSK	1/16	rate 1/2 repeated 8×

and other channel impairments. Modems that implement these waveforms actually perform better on channels with two equal amplitude paths with up to about 6-ms relative delay and independent fading than they do on single path channels. This can be seen by comparing the STANAG 4285 curves in Figs 1 and 2. Half of the channel bits are available for user data and encoding. As an example, the range of modulation and coding used in STANAG 4285 are shown in Table 5. In all cases, the basic coding technique is a rate 1/2 convolutional code with interleaving.

Comparing the performance of a 500-Hz waveform like Clover to a 3-kHz waveform like STANAG 4285 is only reasonable if one keeps in mind the fundamental differences in cost and bandwidth. Given the bandwidth used and the cost of the Clover waveform, one cannot expect performance comparable to the 3-kHz waveforms. However, it is interesting and instructive to compare the two when these differences are kept in mind.

Conclusion

Both Clover and the STANAG 4285 performance demonstrate the capabilities of PSK-based modems (multi-tone PSK in one case, single-tone PSK in the latter) on HF. Clover offers the amateur community an efficient waveform for data transfer with performance far exceeding that of FSK. However, it remains to be seen if amateurs will pay the price in sufficient numbers.

We have seen the price of military HF modems fall from around $25,000 in 1985 to a range of $6,000 to $10,000 today. This reduction was caused by both advances in technology and competitive pressures. A similar price reduction should occur over time with the Clover equipment.

While this article focuses on the Clover waveform, it should be noted that there are other systems with similar characteristics available. For example, the ARTOR modem (not to be confused with AMTOR), by Ascom Radiocom Ltd in Switzerland, offers similar characteristics (ARQ, FEC, 400-Hz bandwidth). Tests of the ARTOR modem are planned. Also, G-TOR has recently been described and discussed in the amateur community.[11] G-TOR uses 100, 200 or 300-baud FSK, ARQ and FEC. It remains to be seen how Clover will fare in competition with systems like G-TOR, which has clear cost advantages; however, the results here clearly indicate the performance advantage of Clover over any system based on FSK.

About the Author

Peter D. Reynolds, KE4BAD, is a Senior Scientist at SHAPE Technical Centre, The Hague, The Netherlands. His work there is with digital radio communication systems in the VLF, LF, HF and UHF bands. Prior to coming to STC, he was with the MITRE Corp, McLean, Virginia, in a similar capacity. He holds an MSEE from George Washington University and a BS in Physics from Rensselaer Polytechnic Institute.

Notes

[1] Henry, B., and Petit R., "HF Radio Data Communication: CW to Clover," *Communications Quarterly*, Spring 1992.

[2] Henry, B., and Petit R., "Clover: Fast Data on HF Radio," *CQ Magazine*, May 1992.

[3] *Military Standard: Interoperability and Performance Standards for Data Modems (MIL-STD-188-110A)*, September 1991.

[4] *Standardization Agreement: Characteristics of 1200/2400/3600 Bits per Second Single Tone Modulators for HF Radio Links (STANAG 4285)*, February 1989.

[5] Wickwire, K., "The Status and Future of High Frequency Digital Communication: Part II," *QEX*, July 1992.

[6] Shanmugan, K. S., *Digital and Analog Communication Systems*, Wiley and Sons, New York, 1979.

[7] *Baud* indicates the symbol rate in symbols per second. For binary modulation techniques, in which each symbol represents one bit, the bit rate is equal to the baud rate. For nonbinary modulation, the bit rate will be greater than the baud rate.

[8] Michelson, A. M. and Levesque, A. H., *Error-Control Techniques for Digital Communication*, Wiley and Sons, New York, 1985.

[9] Schwartz, M., *Telecommunication Networks*, Addison Wesley, Reading, Massachusetts, 1987.

[10] *Standardization Agreement: Minimum Technical Equipment Standards for Naval HF Shore-to-Ship Broadcast Systems (STANAG 4481)*, 1994

[11] Prescott, G., et al, "G-TOR: A Hybrid ARQ Protocol for Narrow Bandwidth HF Data Communication," *QEX*, May 1994

Hearing Strange UHF Signals Lately?

This simple 70-cm filter/preamp is a positive step for eliminating spurious responses.

by John Reed, W6IOJ

UHF receiver designs in past years often did not seriously consider the possibility of interference from neighboring signals. The minimal activity at UHF made the likelihood of this problem relatively low. But recently there has been an explosion of commercial UHF transmitters, mainly hand-held and mobile transceivers for cellular phones and other community-service activities. Monitoring my 70-cm antenna with a broadband 0.1 to 1-GHz receiver shows a typical continuous signal level of –32 dBm, or about S9 + 40 dB (add 10 dB when monitoring with a 2-m vertical dipole). It is my understanding that this level—or more—is not unusual. It is unlikely that your receiver will reject all of these signals. This article describes an integrated filter/preamp assembly that has been specifically designed to discriminate against out-of-band signals. It has a 6-dB, 10-MHz bandwidth with a continued skirt fall-off to more than 60 db. The low-noise preamp has 20 dB of gain. It is a simple assembly that can be built from parts that are easy to procure.

770 La Buena Tierra
Santa Barbara, CA 93111

Spurious Response

A good receiver will have a spurious response rejection of at least 90 dB. But even this will allow some interference when receiving weak signals. A typical minimum signal is –123 dBm (10-kHz BW, 1-dB NF and a 10-dB threshold). Relating this to the 90-dB spurious response rejection capability shows that interference is possible from spurious signals at a –33-dBm level. Although this is a relatively strong signal, many repeaters cover their operating area with a signal at that level when received on a dipole antenna. Also, a nearby transceiver, cellular or otherwise, can cause a similar signal. This need not be a serious problem if the receiver is well designed. The possible interference will likely be restricted to the image frequency or possibly to spurious signals related to the local oscillator second harmonic. However, many of us are using older or home-brew equipment having spurious response rejection ranging from 40 to 60 dB rather than 90 dB. This poor response is most likely due to front-end filtering where the skirt tails off at the lower levels. This, coupled with additional design compromises such as a poorly filtered frequency-multiplier-type local oscillator, can be a source of many interfering spurious responses.

Another source of unwanted signals is intermodulation products from large signals exceeding the receiver's linearity capability. The offenders are usually second- and third-order products and blocking. Although much of the UHF activity uses low-power transmitters, there still is a possibility of large-signal interference when a receiver has a relatively wide-band input circuit followed by low-level circuitry that blocks with only a fraction of a milliwatt of input signal.

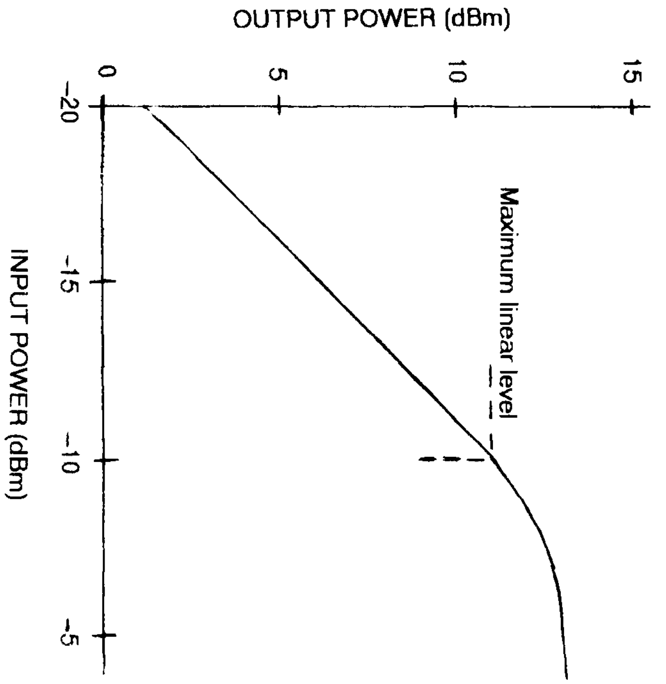

Fig 1—Output power versus input power at 435 MHz.

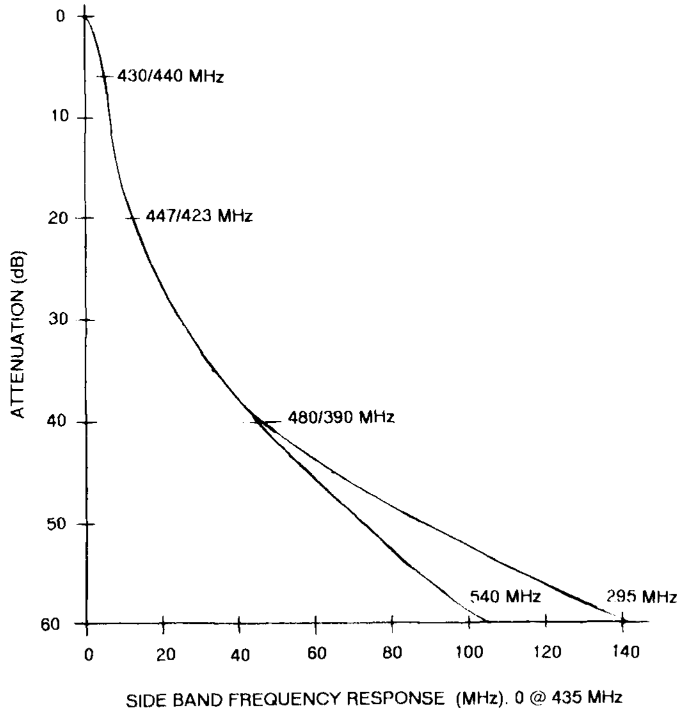

Fig 2—Frequency response of the filter/preamp.

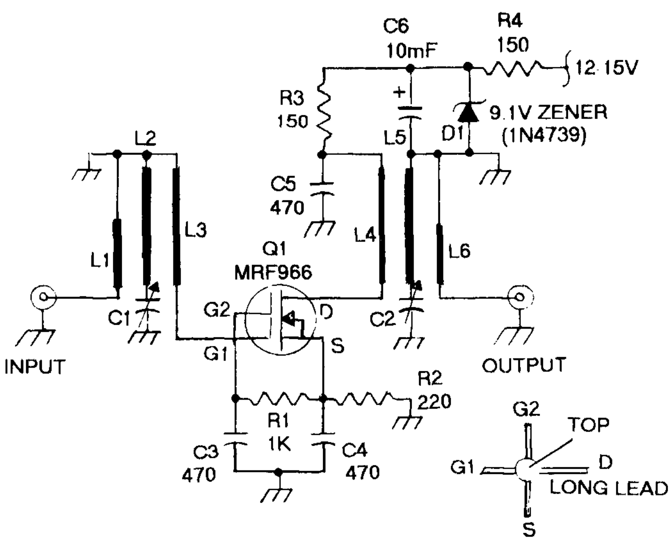

Fig 3—Schematic of the filter/preamp. The filters are parallel #14 wires, see Fig 4 for details. C1 and C2 are 1.6- to 5-pF Sprague Goodman plastic dielectric capacitors (part # GYA5R000). The 470-pF capacitors are 50-V disc ceramic.

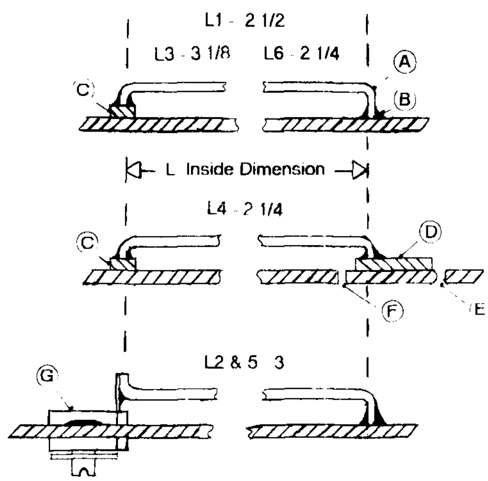

Fig 4—Details of the filter. Each filter consists of three parallel #14 wires mounted ⅛ inch from the common base PC board and spaced from each other by ⅛ inch. (A) Each wire end is peened to form a small flat surface as a soldering aid. (B) Ends are tinned before final soldering. (C) Connection to the component PC board, ⅛-inch stripline. (D) Glue-down PC board mounting pad, ⅜×⅛ inch, cemented to the base PC board. (E) 1/16-inch hole for the R3 feedthrough lead. It is insulated from the base PC board foil by beveling the hole with a large drill. (F) #60 drill hole for the C5 common lead. It is soldered to both top and bottom foils. (G) C1, C2, FILTRIM trimmers.

Filter/Preamp Performance

The filter/preamp active device is a MOSFET having double-tuned filters on the input and output. Fig 1 shows the measured response of the preamp versus signal level, indicating linear response up to a maximum input of −10 dBm (S9 +60 dB) and an output of 11 dBm. This input level, while providing a low-noise preamplifier gain of 21 dB, will still permit a dynamic range exceeding 100 dB in a typical application such as the −123-dBm system noted earlier with a receiver that has a blocking level of 1 dBm.

The frequency response curves shown in Fig 2 were measured using a variable frequency source, a combination of attenuators and a wide-band 38-dB amplifier. The curves show that the upper and lower sidebands are symmetrical to −40 dB. Skirt divergence after that makes it appear that the upper sideband response is more effective than the lower sideband. But the data beyond −60 dB indicates significant filter-multimoding effects in the upper sideband data. The strongest multimode peak is less than −60 dB. Applying this filter/preamp to a receiver having a 10-MHz IF will increase the image rejection by 27 dB. For one having a 30-MHz IF the improvement will be 45 dB. Attenuation of cellular phone signals will be improved by over 60 dB.

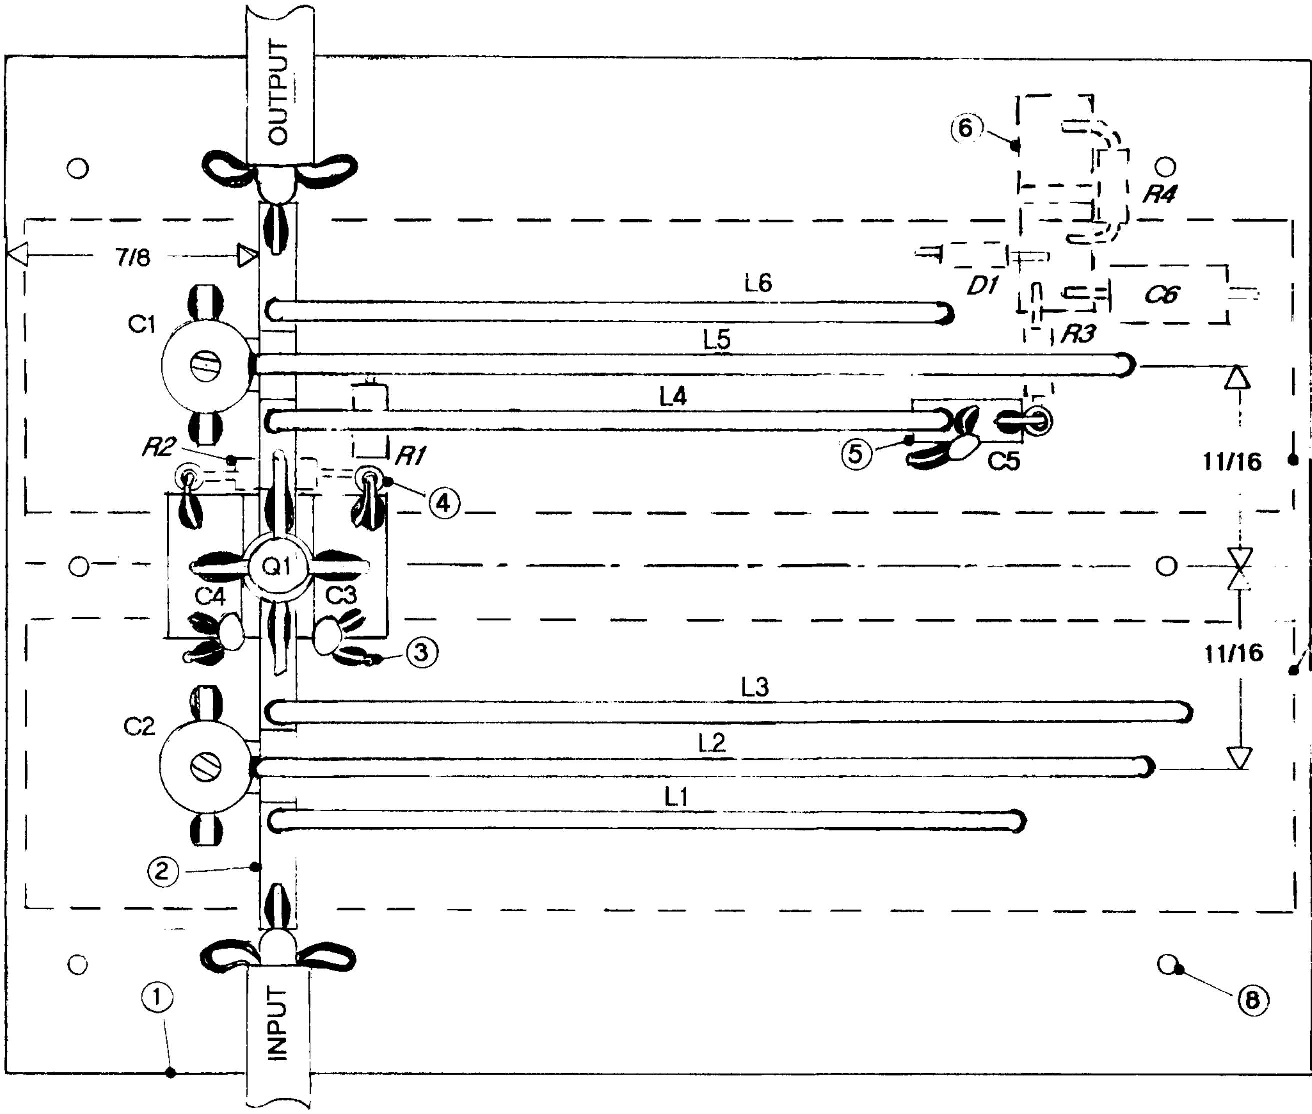

Fig 5—Detailed layout of the filter/preamp. Components mounted on the underside are shown with dotted lines and noted in italics. The heavy lines indicate soldering areas. (1) The base PC board is 1½×3½×¹⁄₁₆-inch with foil on both sides (Radio Shack 276-1499). (2) The component PC board (Fig 6) and the other two glue-down pads are fastened to the base PC board with a clear cement (I use Elmer's Clear Household Cement). (3) Note the near-zero lead lengths of C3, C4 and C5. The common leads of these capacitors are formed into small holes and soldered on both the top and bottom foils. (4) R1, R2 and R3 feedthrough resistor leads are through ¹⁄₁₆-inch holes. Contact with the surface foils is avoided by surface reaming with a large drill. (5) Glue-down PC board mounting strip, ⅛×⅜ inch. (6) glue-down PC board mounting strip, ¼×¾ inch. (7) Open end U-shields, ½×1×4½ inches with ¼-inch mounting flange. Mounting flanges overlap in the center. (8) Shield mounting holes for 2-56 screws.

The effectiveness of the input filter in avoiding the generation of nonlinear distortion products in the preamp was measured by noting a 2% change in the MOSFET source current as a function of inputs at different frequencies. It indicates that the linearity threshold gradually increases from the on-frequency −10-dBm value (a level consistent with that shown in Fig 1) to +10 dBm (S9 +80 dB) at 485 and 385 MHz.

The MRF966 MOSFET has a typical NF of 0.8 dB at 70 cm. The question is, how much is lost in the input filter as compared to that of the standard NF test circuit using a slug-tuner input? I have no calibrated NF source so was unable to make this measurement. However, I would be disappointed if calibrated NF testing showed a degradation of more than 0.5 dB. My tests with an uncalibrated noise source indicate that the filter/preamp outperforms a second MRF966 preamp using a conventional input circuit (wire inductance and piston-trimmer capacitor). Tuning of the filter/preamp input circuit is slightly different between optimizing it for NF or gain. Peaked for best NF, the BW is about 12 rather than 10 MHz.

Circuit

The schematic (Fig 3) shows the 12- to 15-V dc input voltage clamped at 9.1 V. This is to prevent accidental overvoltage to the MOSFET (10 V max). The 9.1 V together with the 150-Ω drain series resistor and the 220-Ω source bias resistor results in the MOSFET operating at close to 5 V at 10 mA, the recommended values for optimum NF performance.

The #14 wire-line inductors of the filters connected to the MOSFET (L3 and L4) are sized close to resonance and closely coupled to parallel wire lines tuned with capacitor trimmers (L2 and L5). In this manner, adjustments have been minimized to two for

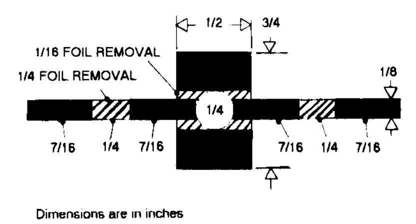

Fig 6—Detail of the component PC board. Material is 1/16-inch double-sided PC board (Radio Shack 276-1499). It can be easily made with a hacksaw and file, together with a clamped straight-edge along the 1/8-inch dimension for ensuring stripline width accuracy. Note: Foil pattern is required on both sides (see text).

the complete filter/preamp assembly. Both filters are inductively coupled to the input/output cables. There may be a question of how well the untuned resonant lines can be duplicated. The resonance mainly depends upon the line length, but there are other influencing factors such as spacing of the wire from the printed-circuit board, spacing between the parallel lines and effects from shield tolerances. I have made three assemblies using the indicated dimensions and all have shown almost identical performance.

The filter construction details are shown in Fig 4. The parallel wire lines are easy to assemble by temporarily inserting a flat 1/8-inch-thick plastic or metal piece between the filter wire and the PC board. This permits maintaining the wire position with your fingers while soldering. Solder L2 or L5, the center wire, to the capacitor first, then assemble the adjacent lines to the indicated dimensions (1/8-inch spacing above the PC board). I try to keep assembly tolerances to within about 1/32 inch. As indicated in Fig 5, the filters are shielded with simple open-ended aluminum U-shields. To be stable the filter/preamp requires a shield on either the input or the output, and it needs both to have the specified overall filter performance.

The key components are mounted on a simple glue-down PC board, detailed in Fig 6. The primary advantage of using this technique rather than the conventional etched board is that it permits an unbroken ground plane, a particularly important consideration in this application because the top foil is the filter ground reference, and to have this common reference for the rest of the RF circuitry is highly desirable. The glue-line insulates the bottom foil from the base PC board, in effect making it a floating capacitor that results in undesirable feedback if left as a reactive elongated section. This effect is avoided by duplicating the top-side foil pattern on the bottom side.

The layout detailed in Fig 5 illustrates the required near-zero lead-length mounting of the 470-pF capacitors, and also the desirable short feedthrough leads of the 220-Ω, 1-kΩ and 150-Ω resistors. Mount the MOSFET last. *A word of caution: the MRF966 has no protective diodes.* I have found that the safest working method is to completely isolate the assembly (no electrical connections) except for a wire connected between the base PC board and the soldering iron. Using this procedure, I have reused a number of MOSFETS from circuit to circuit without damage.

With only two trimmer adjustments, the assembly is easy to peak to any given frequency in the 70-cm band. It is stable regardless of input/output loading.

Parts Suppliers

MRF966—RF Parts Co, 435 South Pacific St, San Marcos, CA 92069, tel: 800-737-2787.

FILMTRIMSs (C1 & C2)—Digi-Key Corp, PO Box 67, 701 Brooks Ave S, Thief River Falls, MN 56702-0677, tel: 800-862-5432.

Where Does the Power Go?

The age-old question of what happens to reflected power still bedevils amateurs—but it shouldn't.

by Jon Bloom, KE3Z

The title of this article comes from a discussion thread (a set of messages) on the Usenet rec.radio.amateur.homebrew news group. The power in question is reflected power, and the question that started the discussion is: where does the reflected power on a transmission line go when it gets back to the transmitter? This is a question that has provoked a lot of discussion over the years. There is general agreement that the reflected power does not get dissipated in the transmitter, but the explanations of what does happen to it range from the absurd to the incomprehensible. In this article, I'll try to explain what really does happen where the transmitter meets the feedline and do so in a way that, hopefully, will clear up some of the many misconceptions about reflected power and transmitter matching.

225 Main St.
Newington, CT 06111
email: jbloom@arrl.org (Internet)

Transmission Lines

I think the first problem amateurs run into when considering questions like these is a feeling that transmission lines are somehow "magic," or "special," and need to be treated as a component in their own right, like an inductor or capacitor. In fact, the way we analyze transmission lines reflects (ooof! pun!) a basic fact: in circuit analysis we make the simplifying assumption that the time it takes a signal to propagate is negligible, but we can't do that directly when a transmission line is involved. The circuit of Fig 1A, for example, is a simple series circuit. We assume that the current will be the same in all parts of the circuit at any instant. For a circuit that is physically small, this is a reasonable assumption. The current travels at a significant fraction of the speed of light, after all. But now make that circuit physically larger by stringing long wires from one component to the other, as in Fig 1B, and the situation has changed. Now, even at nearly the speed of light, by the time a signal has propagated through the wires to the resistor the voltage source is outputting a different voltage. In a circuit such as this, we cannot make the

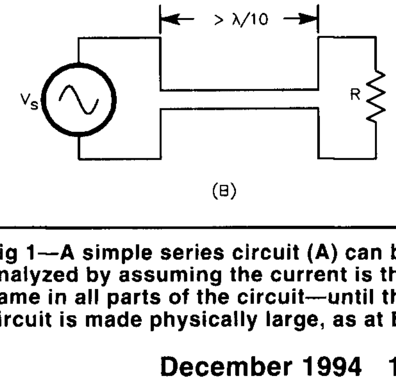

Fig 1—A simple series circuit (A) can be analyzed by assuming the current is the same in all parts of the circuit—until the circuit is made physically large, as at B.

simplifying assumption that the current is the same everywhere in the circuit; it isn't. Analyzing a circuit like this would be quite difficult if we didn't have some (relatively) easy way of treating the delay of those long wires as a part of the circuit. And that's what we do by treating them as a transmission line.

Except for the effect of propagation delay, there is nothing happening in a circuit containing a transmission line that doesn't happen in a "regular" circuit. An intuitive example of this is the use of a short length of coax. Say we connect together two parts of a circuit operating at 3.5 MHz through a 3-inch piece of coax. We know from experience that the added coax makes a negligible difference to the circuit operation. That's because the delay through 3 inches of coax is so small compared to the period of a 3.5-MHz waveform.

Transmission-line theory eases the analysis of a circuit like that of Fig 1B by letting us treat the transmission line as an impedance transformer. (If the transmission line is lossy, it's a little harder to figure out the impedance transformation, but we still can do it.) That lets us use the tools of circuit analysis—Ohm's and Kirchoff's laws, for example—instead of having to deal directly with the time delay imposed by the size of the circuit. In treating the transmission line as an impedance transformer, all we have done is reduce the effect of time delay to a more tractable form; we haven't changed the basic fact that it is the time delay that causes the "transmission line" effect.

Transmitters are Sources

The biggest problem in these discussions tends to be that amateurs view the transmitter impedance as just a simple impedance. With the transmitter end of the feed line terminated by the transmitter's impedance, we struggle to devise a way by which all of the reflected power gets re-reflected, to add to the transmitter-supplied forward power. We do that because we know from measurements using directional watt meters that the power delivered to the load is the forward power minus the reflected power. And we also know that if the feed line is lossless all of the power delivered by the transmitter reaches the load, regardless of the SWR—and thus the reflected power—on the feed line. How else to explain that except to say that all of the reflected power gets re-reflected at the transmitter end of the line?

How else, indeed, if the transmitter presents just a passive impedance to the returning signal. But that's not the case. The transmitter comprises not only an impedance, but a *source* as well. We must include the source in our analysis if we expect to find out what happens to the reflected power. If we don't, we can't find a way to get the reflected power back up the line to the antenna, where it belongs. Why? Because in order to re-reflect all of the power, the transmitter impedance—if that's all the transmitter consisted of—would have to be either zero or infinity. And that's unreasonable.

The transmitter comprises not only an impedance, but a source *as well*

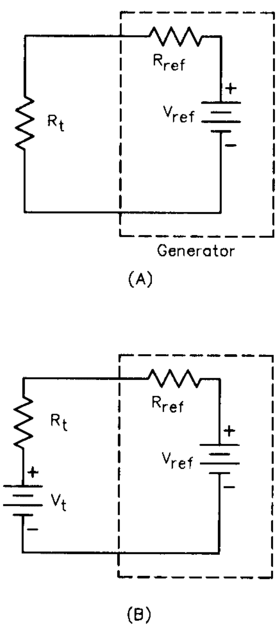

Fig 2—Analyzing the power dissipated in R_t in these two circuits shows that we must include the effect of any sources present in the circuit to get a valid result.

A simple example using a dc circuit will serve to show why we can't neglect the sources when determining power dissipation in a system. Fig 2A shows a circuit with a generator, consisting of an ideal voltage source and a series resistance, and a load resistor. (An ideal voltage source has a 0-Ω impedance and a constant output voltage, no matter what is connected to it.) I've labeled the components in this circuit as though the source is a reflected signal and the load is a transmitter output stage. In this case, the amount of power dissipated in the load resistor, R_t, depends on the values of the resistors and the voltage source. If the voltage source is producing voltage, power will be dissipated in R_t. Now consider Fig 2B. Here, I've added a voltage source to the "transmitter." If this source is producing a voltage other than zero, the power dissipated by the resistors will be something different from what it was in Fig 2A. For example, if the voltages of the two sources were equal, no power would be dissipated in the system at all! We're not looking to find the values of the powers in this example, only to show that the power dissipation cannot be determined unless we take all of the sources into account.

Modeling the System

Fig 3 shows the electrical model of our transmitter-and-load system. The transmitter is a generator, consisting of a voltage source and a series resistance, and the load is some impedance. The first question we should ask ourselves is: Is this model valid? The answer is a qualified yes. Using this model we can determine the voltages, currents and powers appearing at the transmitter output, along the transmission line and at the load. We cannot, however, determine the power dissipated in the transmitter. That's

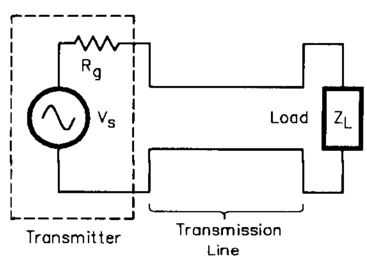

Fig 3—The model we will use for a transmitter, transmission line and load.

because the way we have chosen to model the transmitter doesn't describe the internal circuitry of the transmitter. All our source/resistor model really tells us is what the transmitter "looks like" to whatever is connected to it—the feed line and load in this case. (Thévenin's theorem tells us that we can model the effect of the transmitter on the external load in this way.) But this doesn't tell us what's happening *inside* the transmitter. That may seem to be a problem, since we want to show that the reflected power doesn't get dissipated in the transmitter. But if, using this model, we can show that the reflected signal ends up combining with the generator-supplied signal to produce the forward power, we've accomplished what we set out to do. We don't need to worry about what's happening inside the transmitter because the reflected power never gets there. It's also worth noting that the transmitter output impedance, R_g, may not be purely resistive in reality. But we can afford to neglect the possibility of reactance for this discussion because any reactance that is present will be offset by an opposite reactance in a real system.

For the transmitter to deliver its design power, it must see the proper impedance as a load. The proper load impedance will depend on the characteristics of the transmitter. It may depend on the power the transmitter can safely deliver, on the effect of load impedance on the transmitter's linearity or on some other parameter. I've chosen for an example a load resistance that is two times the transmitter's resistance.

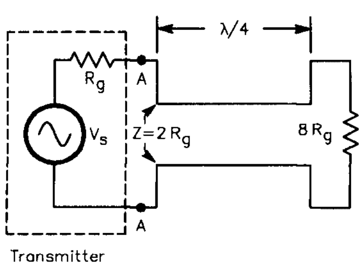

Fig 4—An example system. The quarter-wave line acts to transform the impedance of the load to the design load impedance the transmitter wants to see. We can analyze this system to show that the reflected power combines with the forward power. It does *not* get dissipated in the transmitter.

We can find the sum RMS voltage at any point on the line—including the end-points—by adding the forward and reflected voltage at that point

A Simple Example

Fig 4 is the example system we will analyze. It consists of a generator (transmitter) with an impedance of R_g, connected to a load of $8R_g$ through a lossless quarter-wave transmission line. The design load of the transmitter is $2R_g$, so the task of the transmission line section is to transform the $8R_g$ load to $2R_g$. Using the well-known formula for a quarter-wave transformer:

$$Z_0 = \sqrt{Z_1 Z_2} \qquad \text{Eq 1}$$

we can calculate the needed impedance of the transmission line:

$$Z_0 = \sqrt{(8R_g)(2R_g)} = 4R_g \qquad \text{Eq 2}$$

Now we need to figure out the voltages in the system. We will do so using the principle of superposition, which states that if multiple voltages appear at a point in the system, the resulting voltage is the sum of the voltages. In this case, the voltages of interest are the forward and reflected voltages. The RMS amplitude of the forward voltage is the same at all points along the line. So is the RMS amplitude of the reflected voltage. We can find the sum RMS voltage at any point on the line—including the endpoints—by adding the forward and reflected voltage at that point, but we need to take the relative phase between the two voltages into account. (If we did this for all the points on the line and plotted the result, we would see a *standing wave.*)

We'll begin at the load by calculating the reflection coefficient:

$$\rho = \frac{R_L - Z_0}{R_L + Z_0} = \frac{8R_g - 4R_g}{8R_g + 4R_g} = \frac{1}{3} \qquad \text{Eq 3}$$

This tells us that the voltage reflected from the load is ⅓ of the incident (forward) voltage. This reflected voltage is in phase at the load since the load resistance is greater than the line impedance, so we can find the voltage across the load by summing the forward and reflected voltages:

$$V_L = V_F + V_R = V_F + \frac{1}{3}V_F = \frac{4}{3}V_F \qquad \text{Eq 4}$$

At the source end of the line, point A, the reflected voltage is 180° out of phase with the forward voltage since the reflected voltage had to traverse the quarter-wave line going and coming. Thus the voltage here is the forward voltage *minus* the reflected voltage:

$$V_A = V_F - V_R = V_F - \frac{1}{3}V_F = \frac{2}{3}V_F \qquad \text{Eq 5}$$

Since the impedance looking into the line at point A is $2R_g$, we can compute the voltage applied to this impedance from the generator by simple voltage division between the generator impedance and the impedance at point A:

$$V_A = \frac{2R_g}{R_g + 2R_g}V_s = \frac{2}{3}V_s \qquad \text{Eq 6}$$

Equating the expressions for V_A in Eqs 5 and 6 shows that:

$$V_F = V_S \qquad \text{Eq 7}$$

which also means, using Eqs 3 and 4:

$$V_R = \frac{1}{3}V_S \qquad \text{Eq 8}$$

and

$$V_L = \frac{4}{3}V_S \qquad \text{Eq 9}$$

Now that we know the voltages throughout the system, we can calculate the power delivered to point A by the generator and the power delivered

No *reflected power is dissipated in the generator for the simple reason that there is none left—it's all delivered to the load* to the load. If our analysis is correct, these two values should be the same:

$$P_A = \frac{V_A^{\ 2}}{2R_g} = \frac{\left(\frac{2}{3}V_s\right)^2}{2R_g} = \frac{\frac{2}{9}V_s^{\ 2}}{R_g} \qquad \text{Eq 10}$$

$$P_L = \frac{V_L^{\ 2}}{8R_g} = \frac{\left(\frac{4}{3}V_s\right)^2}{8R_g} = \frac{\frac{2}{9}V_s^{\ 2}}{R_g} \qquad \text{Eq 11}$$

What this analysis shows is that *no* reflected power is dissipated in the generator for the simple reason that there is none left—it's all delivered to the load. In this example, the quarter-wave transmission line matches the load to the desired load impedance of the generator. It does *not* match the load to the generator's internal impedance. We could choose other examples—this one was chosen because the calculations are simple—and get a similar result. If the load impedance included reactance, the matching system would have to include a reactive component to ensure that the impedance presented to the generator at point A was wholly resistive and equal to $2R_g$, since that's the load the generator was designed for.

It's interesting to consider what happens if the transmission line length is reduced to nearly zero, where the delay through the line produces a negligible phase shift. I said earlier that it is only the delay that makes a transmission line "special." If that's true, we should be able to repeat the above calculations for a very short line and get a consistent result. With a negligible line length, we still have a reflection coefficient of ⅓. Now, however, the reflected wave arrives back at point A without the 180° phase shift that occurred because of the quarter-wave line. So, instead of subtracting the reflected voltage in Eq 5 we have to add it. That makes it:

$$V_A = V_F + V_R = V_F + \frac{1}{3}V_F = \frac{4}{3}V_F \qquad \text{Eq 12}$$

which is the same voltage as is present at the load. Clearly, the very short length of line is having no effect on the circuit operation—it might as well not be there, which is what we expect. Of course, this also means that no impedance transformation is taking place, and the generator will see a load of $8R_g$ instead of $2R_g$. That's not the way we want the system to operate, which is why we needed the matching section in the first place, but it does show that our results are consistent.

> *We should not expect that an actual conjugate match will be present in the system during operation*

Reflected Power

What we have seen so far is that the forward and reflected waves combine to produce the voltage at the load and at point A. What may still be causing you some confusion is the idea of reflected "power." That's because we don't think of power as having a polarity, so we can't imagine that the reflected power does anything but zip into the transmitter, despite the fact that there is power coming *out* of the transmitter at the same time. But remember that power is just voltage times current. We are far better off analyzing the system voltages and currents until we've determined the voltages across and currents through the components of the system. Then we can calculate power with certainty.

The Conjugate Match

One of the explanations that has been advanced for the total re-reflection of the reflected wave is that of the conjugate match. Unfortunately, this explanation is unsatisfying at best. A conjugate match occurs at the connection of two impedances that have equal resistances but reactances that are equal in magnitude and opposite in sign (one capacitive, one inductive). Many amateurs take this to mean that the load is matched to the conjugate of the transmitter's impedance, but that isn't so. We can easily see this in our example of Fig 4. Here, the transmitter impedance is R_g and the load presented to it is $2R_g$. The resistances are not the same, hence there is not a conjugate match. That also ensures that there cannot be a conjugate match anywhere else in the system, which we can confirm by calculating the impedance seen looking into the transmission line from the load end. Here, the generator impedance is transformed by the quarter-wave line. From Eq 1, we get:

$$Z_1 = \frac{Z_0^{\ 2}}{Z_2} = \frac{\left(4R_g\right)^2}{R_g} = 16R_g \qquad \text{Eq 13}$$

Since the load is $8R_g$, there is no conjugate match present here, either.

Where the conjugate match does come into play is in designing a matching system that will transform the impedance of the load to the *design load impedance* of the generator. If the load is complex, we design a matching system that would result in a conjugate match *if the impedance at the generator end was the design load impedance*. The reciprocity of a matching network means that this approach will ensure that the load seen by the generator is the load it was designed to see. The generator will operate properly, and all of the power supplied by the generator will reach the load. The conjugate match is a useful design tool, but we should not expect that an actual conjugate match will be present in the system during operation.

Conclusion

What I have attempted to show in this article is that we don't need to resort to any magic to find out what happens to the reflected power on a transmission line. Once we analyze the entire system, considering the voltages and impedances present, it turns out that the reflected power exists as an independent entity only on the transmission line itself. In the transmitter and the load there is only one power to consider, which depends on the voltage present and the impedance the voltage appears across.

Digital Communications

by Harold E. Price, NK6K

Gloom III

I had more response to the last two gloomy columns than any previous effort. Some of the responses should be of interest to the group as a whole, so I'm including them in this (hopefully last) gloomy column.

The first is a response to Lyle Johnson's letter from the October column, where Lyle says that UNIX isn't a suitable platform for the average packet shack. The writer is Bdale Garbee, N3EUA (bdale@gag.com). Bdale has been active in packet for many years as a guru in his local area, and nationally through TAPR, AMSAT and as a contributor to NET/NOS. Bdale:

> I just got done reading your column in the October *QEX*. Take the fact that I'm currently spending most of my spare time on the Phase 3D GPS receiver instead of on packet as a first-order indication of my level of agreement with your overall deep blue funk, but...
>
> First, I was intrigued by Lyle's assertions, in part because in my current position as manager of a group maintaining computer-based design tools for an R&D lab, I'm seeing very similar trends with the vendors I talk to. With increasing frequency, vendors are talking about *Windows NT*. However, none of the significant packages we use have actually been ported to NT yet, so I think it's too soon to say whether NT will really work out to be tool nirvana at some point in time.
>
> However, despite these trends, I'm really bothered by Lyle's statement "I contend that UNIX is completely impractical for the average ham, and maybe even a fraction of the techies." I think his point is in general well taken, that hardware vendors aren't well equipped to sell hardware to non-Windows users, and that commercial CAE tools are unlikely to be ported to BSD or Linux in our lifetimes. However, it all depends on why the average ham is buying a computer, and whether you believe the average ham will have one, or more than one, computer within his reach.
>
> In every case that I'm aware of where one of the free BSD variants or Linux has been put up by a ham involved in packet, that individual has become wildly excited about how cool the concept of a shell is, and how many networking things are either already there out of the box, or are really easy to add. In each case, that ham has "another computer," frequently referred to as "my wife's machine," which runs *Windows* to support the handful of "appliance" applications that are needed around the house, like *Word*, *Excel* and *Quicken*, or their equivalents. In my case, the "DOS box" also hosts schematic, PCB and programmable logic tools. But there's this distinction being made between the machine which is "just an appliance" and the machine on which hobby computing is being done. Sort of like the difference between grabbing the cellular

In every case where one of the free UNIX variants has been put up by a ham involved in packet, that individual has become wildly excited—N3EUA

5949 Pudding Stone Lane
Bethel Park, PA 15102
email: nk6k@amsat.org (Internet)

I believe that mostly-text HTML is in fact a very efficient way to provide information access to folks deploying portable stations in emergency conditions—N3EUA

phone or the mic to the 2-m mobile rig in your car.

I have not suggested to anyone seeking my advice on such matters in the last few years that they should run one of the free BSD or Linux variants to the exclusion of all other computing in their home. However, either operating system with AX.25 support in the kernel is such a powerful, flexible, robust and easy to use environment for doing cool things with Amateur Radio, that I think it would be wrong for even the "average ham" to ignore it.

Maybe I'm just insufficiently capable of thinking "like an appliance operator," but I'd like to be more optimistic than that.

Second, I was sort of surprised at how excited I found myself becoming recently when Phil Karn, KA9Q, announced publicly that he was porting his NOS package over to the DJGPP DOS-hosted DOS-targeting version of the GNU C compiler. The more I've thought about it, the more cool this is. The DJ compiler targets full 32-bit mode, so that a 386 of some kind will be the minimum platform, but the combination of it being a free compiler, and supporting large memory spaces rationally, means that for the first time since 1989, I'm actually thinking from time to time about things it might be cool to add to NOS when he's done.

The reason I mention this in the same breath as the thoughts above, is that for amateur packet radio, a mixture of BSD/Linux machines being used as application servers and home stations for power users, and NOS running on even the cheapest of the 386SX-based portable/notebook PCs, has the potential of being a really feature-capable environment for applications development.

Lyle may scoff at the notion of running Mosaic over a 1200-baud packet link, but I believe that mostly-text HTML is in fact a very efficient way to provide information access to folks deploying portable stations in emergency conditions, and so forth. It can get pretty close to ASCII for density of information content per unit of transmission time, and adds hyperlink capabilities, which I suspect could revolutionize some ARES and/or RACES operations, for example. I fully expect to see WWW usage by these kinds of folks over at least our 9600-baud UHF repeater in the next year. There won't be many pictures, but that's irrelevant.

The Web is another one of those areas where we can figure out how to use the technology effectively for amateur radio purposes, or ignore it and find ourselves one step closer to irrelevancy a year or two down the road.

Who knows where we might or might not invest energies in the next year or two, but Phil's move to the DJ compiler suite is likely to cause an interesting ripple of new applications and IP infrastructure improvements to appear.

My mom says "All the good games need a 486, a sound card and a CD-ROM now." My mom *for crying out loud*

Two Chickens in Every Pot

Bdale, as did Lyle before him, has touched on the haves and have-nots problem again. In the 1980s, Amateur Radio was split into two parts, those with computers and those without. Bdale, and others, now point to the next split, those with one computer and those with more.

Bdale, Lyle, Phil, myself and others are "in the business" and have several computers. In the early and mid 1980s, packet's formative years, a home computer was still an unusual item, one that your spouse had to be talked into, or sneaked in under the guise of "ham gear." Now, computers are an appliance. My dad (N3ECV) and mom, as parents of a boomer and not "in the business," little suspected that they'd ever own a computer, and even I wouldn't have guessed when packet started that they'd have upgraded three times, from an 8088 to a 286 to a 486DX2. My mom says "All the good games need a 486, a sound card and a CD-ROM now." My *mom* for crying out loud. Who would have thought?

Anyway, this brings up the point that maybe we should stop worrying about running our wide-area LANs (only a slight oxymoron) on the household computing appliance. After all, no one complained about not being able to use the toaster, electric typewriter or stereo system on two meters. We never tried to pay the gas bill on a piece of ham gear before, why start now?

The only problem is that the older PCs that are typically replaced (8088) aren't suitable for UNIX, *Windows NT* or even *Windows*, which is why they are being replaced. Still, the concept is sound—don't use your "home" computer to do ham radio. Use your "ham" computer. The successful bidder for our 2400-MHz spectrum could give each digital ham a free 486 and still make a profit. This is called foreshadowing—more about 2400 MHz later.

Neither Bdale nor I, nor Lyle, think that the majority of hams are in the two-or-more computer category. Still, if the network is limited to current capabilities and an 8088 base, we suffer three problems:

1) Our network isn't relevant in the larger world, making it less likely to meet the technical advancement goals in Part 97, and less likely to help us retain spectrum (more foreshadowing!).

2) The Haves, who are doing the work, are going off to more interesting pastures.

3) Many of today's Have-nots are going to be Haves, in the 18-month computer industry hardware capability quanta, and the network won't be ready for them. This has already happened once or twice; see the Packet History section.

The decision to implement AX.25 was not based on what would be best for the network in 1994

Buyer's Guide?

Bdale and I both agree that it is time to publish "The Ham's Guide to Linux—how to buy a Linux system or convert your current system." As usual, there are three components to networking; software, hardware and education. On the subject of information, Steve Stroh, N8GNJ, writes:

> Is Amateur Radio getting eclipsed by the commercial wireless activity? Yes. Is this a bad thing? Probably, but it doesn't have to be, and depends on how you look at it. Amateur Radio is an incubator. It's really rare that an idea can evolve in an amateur radio environment and sustain itself completely within Amateur Radio. Clover is a good example. I'm pretty sure that it's being used commercially now, to great success. Amateur sales of Clover, from what I've read, are decent, but not great. I think there's money to be made—maybe not the obscene amounts that there is in the commercial marketplace, but there is enough activity there to justify some effort.
>
> What can change this? Organization is a big item. Every fall we seem to see lots of enthusiasm gushing forth from students anxious to do a thesis on some aspect of Amateur Radio. They seem to expend more than half their energy in running down the sources of information. There's a *lot* of it out there. But it's buried—pretty deep, in a lot of cases. I think Bdale Garbee demonstrated recently that WWW is *the* way to go. Once the information is posted, it's *out* there, probably forever, to anyone who is sufficiently motivated to look for it (in contrast to some poor grad student who was given a reference to an article published in *Packet Radio Magazine* that no library has ever heard of). In a perfect world, those who publish any kind of amateur radio information would be public-spirited enough to allow its copyright to lapse after 2 years or so, thereby permitting any information to be scanned, updated and posted to a WWW site for anyone to browse. Then we wouldn't spend so much precious time and momentum reinventing things and rediscovering information.

Packet History

Someone on the TCP-GROUP mail list recently said:

"The choice of AX.25 is just astonishing! Who thought that adapting a point-to-point LAPB protocol for use in a multipoint, multiaccess environment is a good idea? Where's the advancement of the state-of-the-art? As far as I can tell, AX.25 came to be due to blind devotion to the ISO god, and not due to any sound protocol design."

Lest you youngsters forget, here is my take on how we got to the odd place we find ourselves. Two groups, closely followed by others, implemented AX.25 on devices that were, or became widespread and spawned countless clones: Hank Magnusky, KA6M, and his group in San Francisco and TAPR in Tucson and Los Angeles. I was on the conference call in 1982 between a group in Los Angeles (Dave Henderson, KD4NL; Wally Linstruth, WA6JPR; Skip Hansen, WB6YMH, and me) and Hank, where it was decided to implement AX.25 on the TAPR and VADCG TNCs. The ironic part, for ISO conspiracy theorists, is that Hank was a staunch supporter of RFC-based protocols.

The decision was not based, as you would hope, on what would be best for the network in 1994, when multi-MIPS, 500-MB hard drives, and 8-MB memory are considered "entry level." It was not based on 100,000 users, graphical interfaces and internetworking.

It was based on using the AO-10 spacecraft and a small number of point-to-point links. It was based on taking a step up from a protocol that had a 4-bit address field. It was based on a PAD for "dumb terminals." It was based on the protocol that the people at that time were offering to implement.

You've heard, in this column and elsewhere, how AX.25 was invented at an AMSAT meeting to develop a protocol for its satellite. You haven't heard, and I'd almost forgotten, what else was involved.

AX.25 may have remained a paper protocol, as many do, if it were not for the fact that KD4NL and NK6K had just written a LAPB protocol, Dave for a 6809 and me for a Z-80. We wrote it for a point-to-point interface so that we could link our Field Day site at 8000 feet in the mountains above Los Angeles down to my shack in Redondo Beach on two meters. We wanted to do computerized logging and duping, and we didn't want to risk my 8-inch floppy drives on my S100 ($500 for two, 160 kB each) in the dust on the hilltop. This was regular LAPB, used for the purpose for which it was created.

TAPR was soon to come out with its TNC—more than one hundred were presold—but the original software group hadn't come through. TAPR had to have a protocol, quickly. As I was on the hook to 16 locals as "TAPR Beta Coordinator" for the area, and since the LAPB Dave and I had written was

There has been advancement in the state of the protocol art, but the protocol has little to do with our current problems—no advancement in our 1940's FM/FSK radios close to the AX.25 that had been proposed for use on AO-10, TAPR picked AX.25 to implement, and Dave and I (and Margaret Morrision, KV7D, in Tucson) ported LAPB to the TAPR TNC.

Hank volunteered to do an AX.25 for the VADCG boards so his user community in San Francisco would be able to talk to the new boxes.

Unfortunately, perhaps, for the 1994 network, TAPR was in the right place at the right time, and thousands upon thousands of TNCs were sold in a few years, all with the protocol derived from a point-to-point satellite service and a Field Day logging system, in ROM.

At the time, the state-of-the-art was 5-bit Baudot at less than 110 baud, with no addressing and no error correction. AX.25 was a huge advancement. It looks nonoptimal now, of course.

There has been advancement in the state of the protocol art (the use of TCP/IP with incremental work to make it more relevant to slow shared links), but the protocol we use has little to do with our current problems—no advancement in our 1940's FM/FSK modulation techniques, no low-cost off-the-shelf high-speed data radios, and no organized nationwide network-building entity.

The goal at the time (1984) was to allow people to experiment with digital radio, to give them a common building block (the TAPR TNC), and to see what happened. What was supposed to happen was better radios, better protocols and a network. What we got was linked RLI BBS systems. That's not to say that RLI forwarding, for the time, was bad. It gave the impression of a network where none existed, however.

I think Phil wrote NET not long after it became possible to do so, meaning the size of the computing resource available to the average ham was not able to support NET in the years preceding its development.

"What about stupid things like AX.25 putting the call in *every* packet!?"

You had to be there at the time. The problems with the existing protocol were thought to be:

1) limited address space, and

2) the requirement for a central authority to distribute address.

The call signs were a solution to the problem to be solved: No one wanted a big brother to assign addresses (one big brother [the FCC] was enough), and no one wanted to transmit a CW ID once every 10 minutes. An ID in each packet made the FCC happy and isn't all that much overhead. Now, of course, we have many little brothers handing out IP addresses—some doing a good job, some not—and we have no requirement to ID in CW because of the legacy of calls in packets.

When people speak of the shame of admitting to coworkers that they're using 1200-baud AX.25, the shame is in the 1200 baud, not the AX.25. Would they laugh any less if you proudly said 1200-baud IP? They'd think you were touched for even bothering.

To some extent, amateur packet radio turned out to be more of a sociological experiment than a technical experiment. Can a large number of poorly linked groups (how ironic that this is a communications hobby), with no central leadership, no defined goal and no funding, come up with a viable network, or even the tools to build one?

The jury is still out.

Tom proposed that three 100-kHz-wide channels, and several narrow channels, be allocated to packet. There was no response from the digital community

Peaked?

I received a letter from Tom O'Hara, W6ORG. Tom is well known for his ATV and spectrum management work. He sent a copy of a letter circulated among the members of SCRRBA, the southern California frequency management group, discussing the new definition of repeater. It seems that Part 97.3(35) defines repeater as a device that "simultaneously" retransmits the transmission of another amateur station on a different channel or channels. As 97.201(b) and 97.205(b) prohibit only auxiliary and repeater stations from the 431- to 433-MHz weak-signal sub band, this could allow properly planned and coordinated simplex packet links, well away from weak-signal work centered on 432.0.

To gauge interest, Tom proposed that three 100-kHz-wide channels, and several narrow channels, be allocated to packet. Though he expected some excitement, there was no response from the digital community. Tom says:

> I read your depressed article in the August *QEX* today. Has packet in general passed its peak interest period? Enclosed is a letter I circulated around the SCRRBA Technical Committee, including Jim Fortney, K6IYK, who represents the digital interests. I thought all the new channels would be received by the digital community with great excitement, but nothing. I remember when you where here we had quite a time trying to find some space for the one high-speed slot at 439.0 MHz and it was quickly taken over by the 1200-baud people. I have not had any applications for high-speed or back-bone digital links on 23 cm for a few years now.

This is very distressing. Our future is passing before our eyes, and it's use 'em or loose 'em. It may already be too late for 2400 MHz.

Doom—2400-MHz Prospects

As foreshadowed, here is some bad news. If you thought the 220-MHz grab was big, this could be bigger. Report No. DC-2666, ET Docket No. 94-32, re: Notice of Proposed Rulemaking FCC 94-272 on October 20, 1994, "Commission Proposes Allocation of Spectrum Transferred From Federal Government to Private Sector," says in part:

> In compliance with the provisions of Title VI of the Omnibus Budget Reconciliation Act of 1993, the Department of Commerce released a report on February 10, 1994, which made preliminary identification of 200 megahertz of spectrum for reallocation from Federal Government to private sector use, including 50 megahertz at 2390-2400 MHz, 2402-2417 MHz, and 4660-4685 MHz that is immediately available. The Reconciliation Act requires the Commission to adopt rules by February 10, 1995, to allocate the spectrum.

The Commission also believes that most of the services to be provided in this spectrum would likely meet the statutory criteria for auctions. Therefore, the Commission proposed to make licenses for this spectrum available through competitive bidding to the extent possible and practicable.

"Auctions" means big money is in play. Find out what this is all about, and write in to the FCC. This will no doubt be in all the usual ham information sources by the time you read this, so view this as a reminder to do your part. The full text of the FCC NPRM can be obtained in file FCC94272.TXT from ARRL's BBS (203-666-0578), via anonymous FTP from oak.oakland .edu or from the ARRL's email INFO server (info@arrl.org—send a message to that address containing the single line: send fcc94272.txt). The comment deadline is December 19, so time is of the essence.

If you thought the 220-MHz grab was big, the loss of 2400 MHz could be bigger

Low(er) Cost 900 MHz?

In the only good news I could find, here is a note from Matt Kastigar, NØXEU, on a project he's working on to reclaim previous-generation cellular phones:

I have converted (cut-up, etc) a Hitachi (early) floor/trunk mount cellular phone; and with the help of Jay Underdown, WØOGS, put the duplexer at 900 MHz. Anything that was a computer (versus "radio") was surgically removed. Then I deciphered the bit-patterns for the PLL and (via common 1N914's) "forced" the LO to where the output is 902.500000 MHz (LO is shared transmit and receive). I fixed the output at (approx) 1 watt. Yes, it works, sort of, and needs the computer section reactivated for all the 90's-type features—scanning, memories, etc, plus the control of output power (enough to get a response from the cell site and no more).

It's all there, and with care and attention to details like the output filters, mixer peak, etc, cheap radios are here. This was not a serious effort; it started out because someone gave me a BC-454 (Command) receiver and I started thinking (here's the dangerous part) that there must be something out there that's cheap and usable by the amateur community like those surplus rigs. A trip to a local surplus vendor supplied the radio ($5.00 with handset and antenna—a take-out from a rental car company).

Since the older radios use "full size" components, this is the logical first step in the experimental stage.

Please

Please send me some good news. Three depressing columns is enough.

Book Review

Digital Signal Processing in Communication Systems, ISBN: 0-442-01616-6. By Marvin E. Frerking. 1994 Van Nostrand Reinhold, New York. Hardbound, 576 pp.

Reviewed by Ladimer S. Nagurney, WA3EEC, Associate Professor of Electrical Engineering, University of Hartford, Connecticut

Digital signal processing is becoming pervasive in communications systems, in both professional and amateur equipment. But in most DSP books, the topics discussed are the FFT and FIR/IIR filters, with communication topics such as modulation and demodulation discussed only briefly, if at all. And practical topics such as implementation are almost never discussed.

Marvin Frerking has changed this emphasis with the book *Digital Signal Processing in Communication Systems*. Finally, a book has been written that directly relates DSP to its communication applications. For completeness the book includes a chapter on the fundamentals (FFT and z-transforms) and one on digital filters, but even these chapters emphasize communication applications.

Chapter 3 discusses A/D conversion and spends a large amount of time on quantizing noise, intermodulation distortion, aperture jitter, and sampling of band-pass signals. These directly affect the signal-to-noise ratio of the system.

The heart of the book is Chapters 6 to 8, entitled "Digital Algorithms for Communication Systems," "Digital Receiver/Exciter Design," and "Data Transmission," respectively, that comprise 277 of the book's 576 pages.

The "Digital Algorithms for Communication Systems" chapter (Chapter 6) begins with a discussion of DSP techniques used to generate a sine wave, paying special attention to practical issues such as frequency resolution limitations due to finite-length arithmetic. The author then discusses amplitude modulation, including the approximations necessary for DSP processors that do not have a square-root function. The chapter continues with FM and SSB and concludes with a section on digital AGC and squelch circuitry.

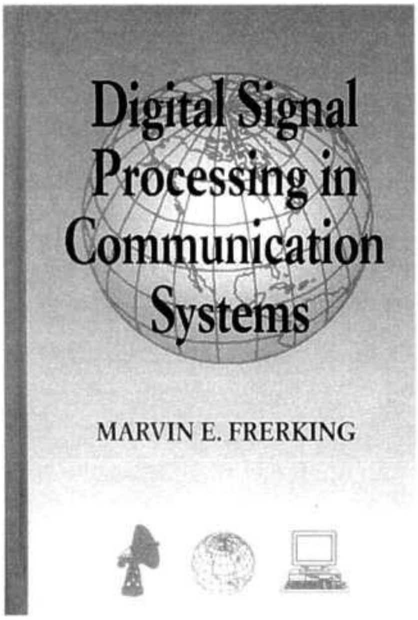

Chapter 7, "Digital Receiver/Exciter Design," begins with the design of a receiver, thoroughly discussing gain and noise distribution and the tradeoffs between analog and digital approaches for each stage. Frerking then discusses alternative architectures such as narrowband receivers using high-speed A/D converters, harmonic-sampling receivers and direct-sampling receivers. Transmitter design is then discussed based on the assumption that the baseband (modulating) signal will be digitized, processed and finally converted to an analog signal at the carrier frequency. He concludes the chapter with discussions of DSP for ALC and using DSP to linearize a nonlinear power amplifier.

Data transmission is included in Chapter 8 beginning with a discussion of matched filters. FSK and PSK are then discussed including all variants, such as m-PSK and MSK. The chapter concludes with a discussion of equalizers and a short note on HF channel models and their simulation using DSP.

Chapter 9 reviews the major topics of speech coding including Linear Predictive Coding and the implementation of various algorithms.

The concluding chapter is "DSP Hardware," which is a potpourri of topics such as DSP processor design, data flow structures and various bus designs.

The book has 7 appendices that provide rigorous derivations of results presented in the book. The reference section of the book lists 90 references that represent a good cross section of communication and DSP literature.

Having taught both communication engineering and DSP at the undergraduate level for a decade, I found the book a pleasure to read. Most topics were presented at the right level, with a good balance between theory and practice. A radio amateur with some knowledge of linear systems, communication engineering and elementary DSP should find this book accessible. For those without the latter two prerequisites, having Couch's *Digital and Analog Communication Systems* and Proakis and Manolakis's *Digital Signal Processing* nearby would be useful.

I would expect that technically inclined radio amateurs would benefit from this text.

The author, Marvin E. Frerking is WØEQC.

Feedback

In "Testing and Calculating Intermodulation Distortion in Receivers," by Ulrich L. Rohde, KA2WEU, in the July, 1994, *QEX*, several measurements of receiver intercept points are reported under the "Example Measurements" heading. The equations that show the calculations may be a source of some confusion. In those equations the author used a number that is the difference in attenuator settings between the reference level measurement (P_{IM_n}) and the level of one of the two signals used for the two-tone measurement (P_A). Using the difference in the attenuator readings is valid only if the attenuator is set to 0 dB for the two-tone measurement, as was done by the author. The following derivation makes this clear:

$$IP_n = \frac{nP_A - P_{IM_n}}{n-1}$$

$$P_A = -20 \text{ dBm} - A_2$$

$$P_{IM_n} = -20 \text{ dBm} - A_1$$

Here, P_A is the level of one tone applied to the receiver input during two-tone testing, and P_{IM_n} is the reference level measured using the single-frequency signal. Each of these levels is the level into the attenuator, –20 dBm, minus the attenuator setting in dB: A_2 for the two-tone signal and A_1 for the reference level. Substituting to form a single equation gives:

$$IP_n = \frac{n(-20 \text{ dBm} - A_2) - (-20 \text{ dBm} - A_1)}{n-1}$$

$$= \frac{(-20 \text{ dBm})(n-1) + A_1 - nA_2}{n-1}$$

Note that the two-tone attenuator setting, A_2, appears in the equation n times. Thus, the difference in attenuator settings can't be used directly unless A_2 is 0 dB.

The most error-free procedure is to calculate the actual receiver input power levels and plug them into the first equation above. Trying to make a measurement such that the difference in attenuator settings can be used directly can lead to confusion and errors.—*KE3Z*

Proceedings

Available from both the ARRL and AMSAT (PO Box 27, Washington, DC 20044; tel: 302-589-6062; fax: 301-608-3410; or Internet: martha@amsat.org).

Proceedings of the AMSAT-NA 12th Space Symposium and AMSAT Annual Meeting. ISBN: 0-87259-487-4; ARRL Order Number: 4874; $12.00 plus s/h.

Welcome to the 1994 AMSAT-NA Space Symposium & Annual Meeting, Bill Tynan, W3XO

Phase 3D, A New Era for Amateur Satellites, The Phase 3D Design Team

Phase 3D Critical Component Radiation Testing, Paul A. Barrow, VE6ATS

Contingency ACS Configurations for the AMSAT Phase 3D Spacecraft, Walter K. Daniel, KE3HP

Telecommunications Satellites from the World's Garage; The Story of the Amateur Radio Satellites, Keith Baker, KB1SR, and Dick Jansson, WD4FAB

Launch Opportunities Beyond Phase 3D, Philip Chien, KC4YER

The Re-Entry of OSCAR-13, James Miller, G3RUH

1993-94 Report on DOVE Recovery Activities, Robert J. Diersing, N5AHD

A SAREX Case Study: Getting Teachers Interested in Amateur Radio and Space Education, Joan Freeman, KD4SRD, and Philip Chien, KC4YER

UoSAT-3: Lessons Learned from Three Years of Serving the Development Community, Eric Rosenberg, WD3Q

A Long-Term Examination of AO-16 and UO-22 Activity Logs, Robert J. Diersing, N5AHD

UNAMSAT-A: An Operations Guide, David Liberman, XE1TU

LUSAT-1 CW Beacon Transmitter, Gustavo Carpignano, LW2DTZ

VOXSAT Review, Gustavo Carpignano, LW2DTZ

The Stanford SQUIRT Micro Satellite Program, Christopher A. Kitts and Richard A. Lu

Investigating the Integrated Control of Payloads with Amateur Satellites, Christopher A. Kitts

Use of Star Cameras for Attitude Determination of Amateur Radio Satellites, Walter K. Daniel, KE3HP

Orbital Analysis by Sleight of Hand, Dr. H. Paul Shuch, N6TX

Dish Feeds for Mode S Reception, Ed Krome, KA9LNV

MMICs at S-Band, Doug Varney, WA1UVP

A Major Field Day Satellite Operation, New Stearns, AA7A

Arctic HF Satellite Radio Propagation Studies, John Branegan, GM4IHJ

1994 Index

January 1994 to December 1994 QEX *(Issues 143 through 154)*

Proceedings

RF (Lau)

Specialized Communications Techniques

Upcoming Technical Conferences