QEX

May/June 1998

$4

Forum for Communications Experimenters
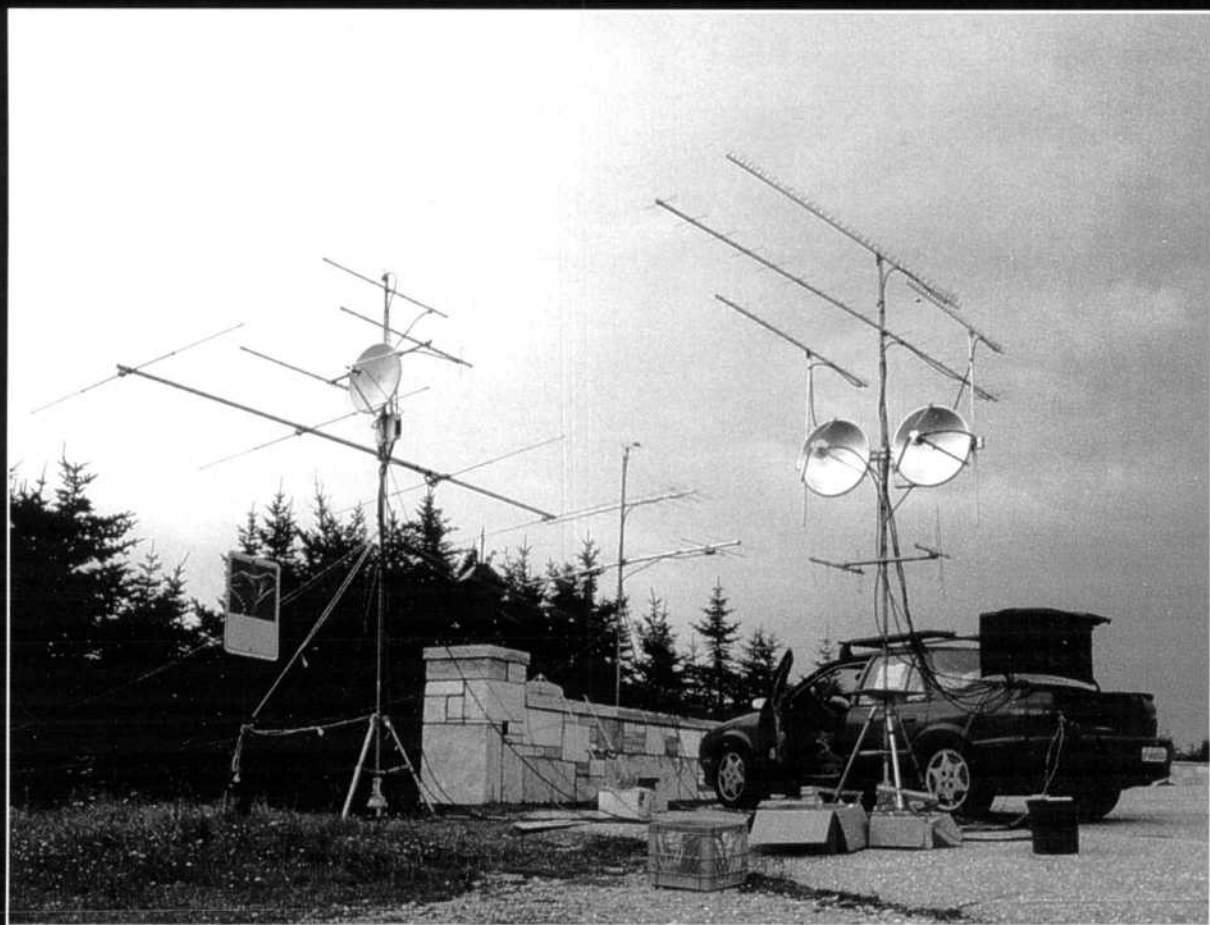
VHF to Microwave!—Portable Contest
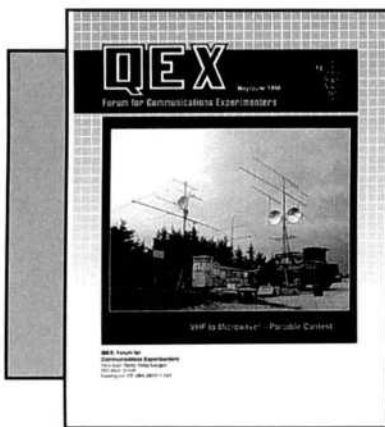
QEX

David Sumner, K1ZZ
Publisher

Rudy Severns, N6LF
Editor

Robert Schetgen, KU7G
Managing Editor

Lori Weinberg
Assistant Editor

Zack Lau, W1VT
Contributing Editor

Production Department
Mark J. Wilson, K1RO
Publications Manager

Michelle Bloom, WB1ENT
Production Supervisor

Sue Fagan
Graphic Design Supervisor

David Pingree, N1NAS
Technical Illustrator

Joe Shea
Production Assistant

Advertising Information Contact:
Brad Thomas, KC1EX, *Advertising Manager*
American Radio Relay League
860-594-0207 direct
860-594-0200 ARRL
860-594-0259 fax

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

Offices
225 Main St, Newington, CT 06111-1494 USA
Telephone: 860-594-0200
Telex: 650215-5052 MCI
Fax: 860-594-0259 (24 hour direct line)
Electronic Mail: MCIMAILID: 215-5052
Internet:qex@arrl.org

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

About the Cover

Zack's portable contest station is impressive. Learn the hows and whys in his RF column.

Features

Columns

May/June 1998 QEX Advertising Index

Empirically Speaking

Antennas are a perennial hot topic for amateurs because they represent one of the few areas in our hobby where homebrew, or at least home-adapt, is still common. Over the past few years, software for applying the many versions of the Numerical Electromagnetic Code (*NEC*) has become widely available at prices affordable to amateurs. These programs are now widely used, and they have truly energized antenna discussion in amateur circles. In this issue, we have two articles on *NEC* programs. Nevertheless, *NEC* and *MININEC* are not the only games in town. As we have become more sophisticated in the use of antenna modeling software, we have also become more aware of its shortcomings and limitations.

The whole subject of antenna modeling has generated great interest in commercial applications as well as amateur work. To address this interest, the Applied Computational Electromagnetics Society (ACES) has been formed to provide a forum for electromagnetic modeling. The society was organized at the Naval Postgraduate School in Monterey, California, but it has a worldwide membership. ACES publishes both a journal and a newsletter several times a year. Much of the material is quite advanced, but there is often material of interest to amateurs, particularly in the newsletter. In addition, ACES has been a source of software code, especially of a developmental character, which may be of interest to advanced amateurs. I have been a member of ACES for several years and found it a useful resource. You can obtain more information on ACES from: Richard Adler, 833 Dyer Rd Room 437, ECE Dept/code ECAB, Naval Postgraduate School, Monterey, CA 93943-5121.

In addition to ACES, the IEEE Antenna and Propagation Society and the IEEE Broadcast Technology Society publish newsletters and transactions. These publications and the older IRE transactions can yield a rich trove of antenna ideas. (For example, Al Christman, KB8I, has published several articles about elevated ground systems for low-frequency vertical antennas in the Broadcast Technology transactions. Al's articles are interesting to amateurs and quite easy to follow.) Readers' letters are especially helpful; they are usually much shorter and simpler than the articles. You need not join the IEEE to get these publications; they are widely available in public libraries, especially at colleges and Universities.

In This *QEX*

We have two articles on antenna modeling using the *NEC* and *MININEC* software. They are particularly interesting in that they show that the latest version of *NEC* (4.1), while very capable, is not superior in all respects. You really must choose the software to fit your problem and pocketbook. Logan and Rockway are the originators of *MININEC*. Their article describes the latest version of *MININEC*, and it's application. *MININEC* lives! L. B. Cebik, W4RNL, gives a very careful overview of the limitations of *NEC* 4.1, with comparisons to results for other programs, as well as some sound modeling advice. This article may save you a bundle if you have been considering purchasing *NEC4*, which has a hefty user fee, in addition to paying for any *front-end* software. This program is just the thing for some users, but not everybody.

Many of you have given us a lot of great feedback about Doug Smith's DSP article in the last issue. Part 2 continues the journey with a trip through a DSP-IF transceiver. It's full of more great revelations. Check it out.

As we build more-sophisticated receivers, the subject of phase noise (it's measurement and reduction) often arises. We have a pair of articles on this subject: Bruce Pontius, N0ADL, shows how to measure phase noise with low-cost equipment. William Cross, KA0JAD, describes his Flexible Frequency Generator (synthesizer). Along the way, he explains how phase noise is generated and how to minimize it.

Yes, that's a Zack Lau photo on the cover—again! We show no prejudice; Zack simply has a gift for great color photos. Come on folks; let's see some good color shots for the cover! Look in the RF column to learn how Zack developed his formidable and versatile VHF to microwave contest station.—*73, Rudy Severns, N6LF,* **rseverns@ arrl.org**

NEC-4.1: Limitations of Importance to Hams

Antenna modeling seems deceptively easy.
Come tour the pitfalls of the latest software.

By L. B. Cebik, W4RNL

lthough most hams will not use NEC-4 because of its cost, I asked L. B. to write this article because QEX readers are more likely to consider purchasing NEC-4 software than any other Amateur Radio publication audience. After reading what L. B. has to say, you may very well decide not to get NEC-4 for your applications—even if price is not a limiting factor. This article shows how important it is to select a modeling program with regard to how its strengths and weaknesses relate to your application. After reading this article, I will use NEC-4 only when it is clearly the better program for the problem at hand. Of particular interest to me is the modeling of sailboat antennas over salt water. That problem requires the modeling of multiple wires of

*1434 High Mesa Dr
Knoxville, TN 37938-4443
e-mail cebik@utk.edu*

radically different diameters, connected at common points with very small angles. For that application, the new MININEC Pro (with its new algorithms) is much superior to other software packages. In addition to the NEC-4 discussion, there are a number of comments relating to other commonly used antenna-modeling programs that are well worth reading. Overall, this is a very important cautionary tale. If you are into antenna modeling, I think you will find this article very interesting, maybe even a bit disturbing.—Rudy Severns, N6LF, QEX Editor

Although *NEC-4* (current version 4.1) has appeared to be a large jump from *NEC-2*, it is simply another step in the evolution of method-of-moments antenna modeling pro-

[1]Notes appear on page 16.

grams.[1] Three factors make *NEC-4* seem like so large a leap forward. First, it resolves the problem *NEC-2* has with stepped-diameter elements, so common in HF Yagi construction. Second, it adds the capability of handling buried radial systems (just when they are going out of style in favor of elevated radial systems). Third, it's not in the public domain, it's proprietary and requires a license, in addition to the purchase or development of interface software.[2]

Numerous ham users of *NEC-4* appear to have overlooked that, like all of its predecessors, *NEC-4* has some limitations that users must heed if they are to successfully model and analyze antennas with this program. Some of these limitations involve common rules of modeling that are covered

extensively in software manuals and elsewhere. These I shall bypass. Some of the limitations are published or unpublished recommendations that many ham modelers seem to overlook. These I shall review briefly. Finally, some of the limitations are relatively unpublished, and these I shall look at in more detail.

First, a couple of conventions: *NEC* handles angles of radiation as zenith angles, increasing from 0° at the zenith, directly overhead to 90° at the horizon. Hams are more accustomed to elevation angles, which count from 0° at the horizon to 90° overhead. I shall keep to the amateur convention wherever matters of up-down angles arise. Amateur and commercial interests are also more at home speaking of antenna element *diameter*; whereas wire *radius* is native to *NEC* calculations. Again, I shall adopt the convention more familiar to hams. Readers of the basic *NEC-4* manual must, of course, translate wherever appropriate.[3]

Some Commonly Abused, Known Limitations of *NEC-4*

Method of moments modeling requires the use of straight wires to form a close approximation of the antenna geometry. For longer linear wire lengths, wires may be segmented within a given wire specification. The recommended maximum segment length is 0.1 λ, with a more conservative limit of 0.05 λ recommended for critical regions of the antenna. Although the absolute limits permit as few as five segments for a 0.5 λ dipole, most modelers use 9 or 11 (adhering to the need for an odd number of segments for a single-source center-feed of the antenna). With 11 segments, each segment for a 3.5 MHz center-fed wire would be about 12.5 feet long. Quite good results emerge from this segmentation for simple antennas.

However, unconscious error occurs when modelers simply shift the frequency of the antenna to other amateur bands without changing the segmentation of the wire. At 28 MHz, the shortest segment should be no more than 3.5 feet, or more conservatively about 1.7 feet. If one is exploring the use of a 3.5 MHz dipole as an "all-band" doublet, the antenna should be resegmented for each frequency band or segmented sufficiently for all frequencies to be explored. *NEC-4* has no specific limit for segment shortness of thin-wire antennas.

NEC limitations on the minimum segment length occur in relationship to the wire radius or diameter. In general, the maximum wire diameter is limited so that π times the diameter, divided by the wavelength should be much smaller than one. HF antennas are highly unlikely to approach anywhere near this limit. However, highly segmented antenna wires with large diameters may approach the length-to-diameter limitations of the
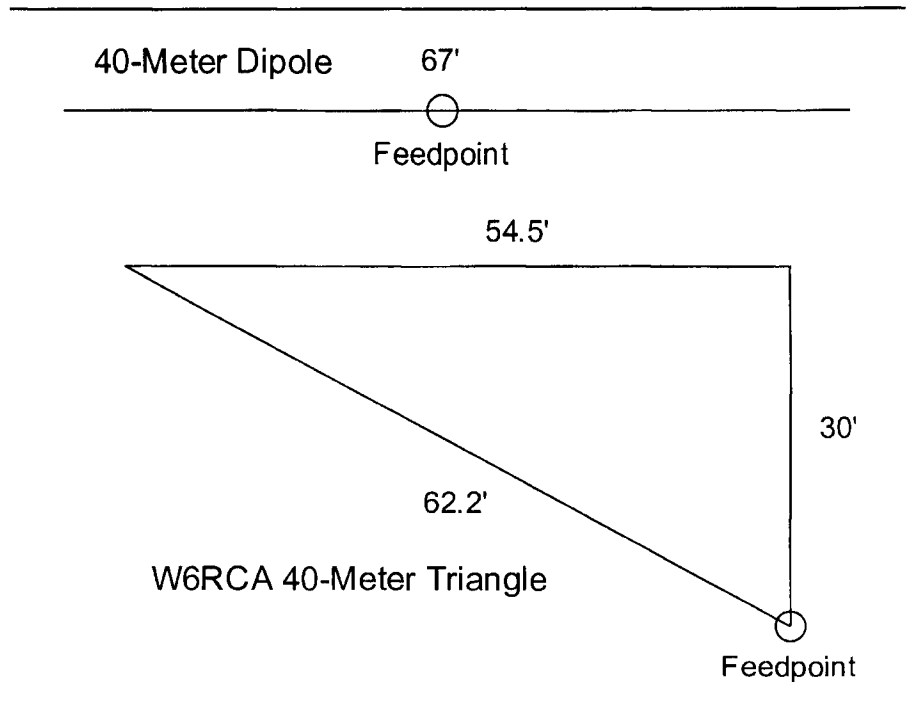

Figure 1—Two antennas for comparative convergence tests.
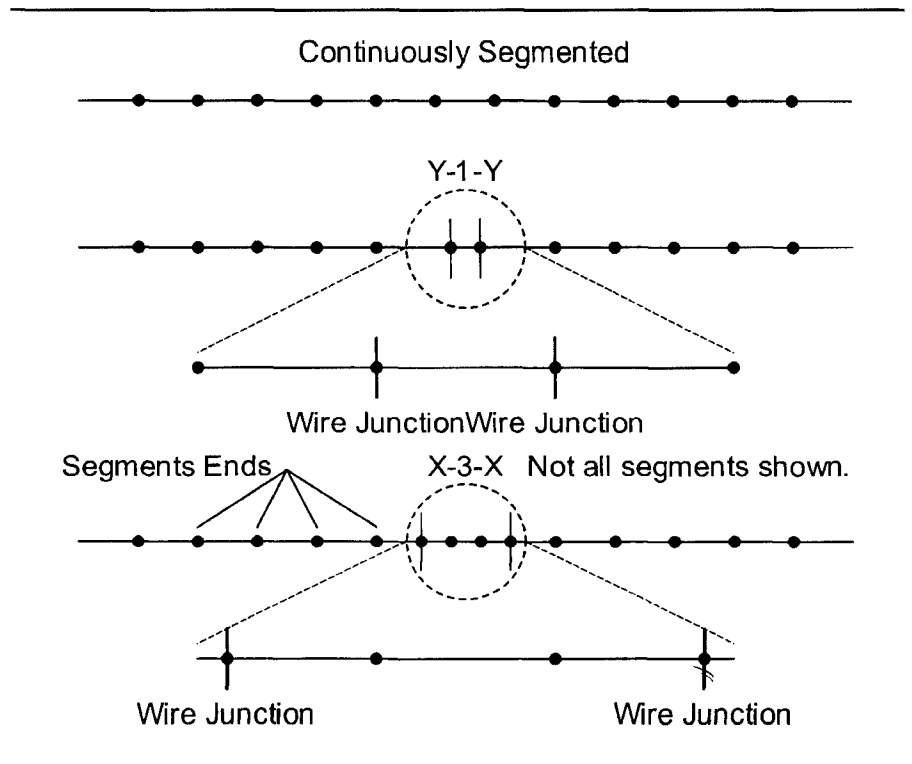
Figure 2—Three methods of segmenting a dipole.

program. The known error point occurs when the segment length is less than 25% of the wire diameter, but programmers strongly urge ratios of length to diameter many times larger, at least 1:1. Although this limitation is seemingly easy to avoid, it is often crossed when modelers begin increasing wire diameters in order to explore various affects on antenna length, resonant frequency, current phasing, etc.[4]

The exact number of segments needed in a half wavelength can be determined by *convergence* testing. Although this common test is rarely mentioned in modeling literature anymore, it remains a fundamental test of result reliability for antenna models in any of the *NEC* and *MININEC* programs. Convergence testing consists of increasing the number of segments per unit of length equally throughout the antenna structure and observing changes in the output data for parameters significant to the modeling exercise. Often, gain and feed-point impedance are used as leading indicators. If these figures change seriously, convergence has not been achieved. If they do not change seriously, we can consider the model to be converged. Exactly what counts as a serious change is subject to the nature of the test and the complexity of the antenna being modeled. Moreover, every change in segmentation will produce mathematically detectable changes.

Linear, simple wire antennas rarely require more than the minimum recommended number of segments per half wavelength to achieve conver-

gence in free space. However, even seemingly simple antennas over ground may require denser segmentation. For example, an asymmetrical triangle that is vertically oriented may require considerably more segments to achieve convergence. Figure 1 and Table 1 illustrate the difference by showing some *NEC-4* output numbers for a standard dipole and a 1 λ triangular loop. Note that the simple dipole shows almost no change, except for minor changes in the feed-point impedance, after 11 segments per half wavelength. However, not until the triangle uses about 73 segments or about 36 per half wavelength, does the elevation angle of maximum radiation (or take-off angle, TO) stabilize, along with the feed-point impedance. The gain continues to vary even at 104 total segments.

One problem related to segmenta-

tion concerns the length of adjoining segments. Traditional cautions suggest keeping the ratio at 2:1 or less. However, the ratio required may be even less, especially in the region of the antenna feed-point. Consider a thin-wire center-fed dipole of 0.01 inch-diameter aluminum wire. One standard way to segment the dipole is to use a single wire with equal segments along its length. (We may call this continuous segmentation.) A second way to segment the antenna is to use three wires. The center, or feed-point, wire might consist of a single segment, which we shall arbitrarily set to 0.2 feet for a 20 meter operating frequency (a rate of about 170 segments per half wavelength). The outer wires may be segmented at the same rate or a different rate. We may designate this the Y-1-Y arrangement, as shown in Figure 2.

Table 1—Convergence Testing for Two Different Antennas

Total Segments	Gain (dBi)	TO Angle (°)	Feed-point Impedance (R ±jX Ω)
40-meter λ / 2 Dipole (See Fig. 1)			
5	5.82	49	102.90 − j38.83
11	5.86	49	92.38 + j6.47
15	5.87	49	91.08 + j8.59
19	5.87	49	90.51 + j9.18
23	5.87	49	90.21 + j9.42
40-meter 1 λ Triangular Loop (See Fig. 1)			
21	3.75	38	67.62 − j7.07
32	2.57	27	53.91 − j4.72
42	2.28	24	49.54 − j1.34
52	2.18	22	47.54 + j1.19
63	2.13	22	46.40 + j2.94
73	2.11	21	45.78 + j4.86
83	2.10	21	45.48 + j5.71
94	2.09	21	45.26 + j6.28
104	2.08	21	45.14 + j6.79

Figure 3—Segment length versus dipole gain for different methods of segmentation.

Figure 4—Segment length versus dipole impedance for different methods of segmentation.

Figures 3 and 4 graph the outputs of *NEC-4* with respect to gain and to feed-point impedance for both the continuous and the Y-1-Y configurations. So long as the total segmentation is such that adj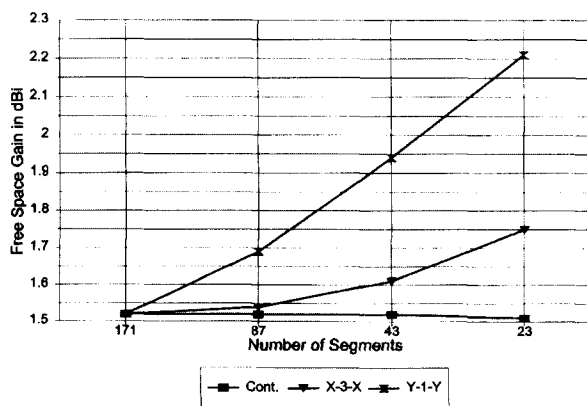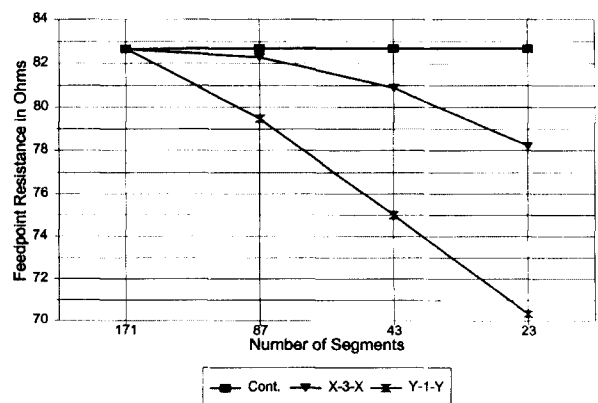oining segments are of equal length, results converge. However, even at an adjacent segment ratio of 2:1 (87 segments), reported gain for the Y-1-Y configuration is erroneously higher and feed-point impedance of the resonated antenna is erroneously lower.

An often-heard recommendation requires that the segments immediately adjacent to the feed, or source, segment are equal in length to that segment. The middle line of the graphs, designated X-3-X, traces this configuration in Figure 2. The source and adjacent segments are all 0.2 feet long at 14 MHz, while the end wires are variously segmented so as to increase the ratio of segment lengths in the outer wires to segment lengths of the center wire. Even in this configuration, a 2:1 ratio (number of segments = 87) yields a clearly detectable set of deviations from the continuously segmented model. Segment-length equalization is perhaps considerably more significant than many modelers suppose. An alternative is to taper segment lengths along each outer wire so that the segments adjacent to the center wire are nearly equal in length to the segments in the center wire. Existing software either supports the segment-length tapering (GC) input card or provides for segment tapering externally to *NEC-4* calculations. A further alternative, recommended where multiple di-

verging dipoles meet, is to use a three-segment wire of no less than 0.02 λ, with due caution paid to the length of adjoining segments.[5]

Segment-length equalization should not be divorced from the idea of selecting segment length as a function of the operating wavelength. When so segmented, wires in a more complex geometry adhere to the recommendation that the segments and their dividing points parallel each other to the degree permitted by the material structure of the antenna. Although this recommendation is made specifically for closely spaced wires, adherence to a general segmentation scheme ensures adherence to it as well.[6]

There are numerous other modeling cautions enumerated either in *NEC-4* documentation or in user's manuals for commercial implementations of the program. Among the areas in which modelers need to use caution are wire junctions and "near junctions," minimum angles of angular wire junctions and minimum loop-antenna element sizes. However, the items discussed above represent perhaps the most numerous problematic practices that I have encountered in looking at several hundred ham-generated models in all versions of *NEC*.

A Lesser-Known Limitation: Stepped-Diameter Difficulties

NEC-4 implemented changes in the method of handling currents and boundary conditions to overcome known inaccuracies that occurred in *NEC-3* and *NEC-2* with respect to antenna elements having stepped di-

ameters. For many cases, these measures are completely effective, and the results of direct *NEC-4* modeling of standard Yagi designs with "taper schedules" have been very accurate.[7]

However, anomalies begin to appear when modeling, in *NEC-4*, Yagi designs developed by K6STI for the program *YA*.[8] When modeling in *MININEC*, no problems of convergence were noted with reasonable numbers of segments per half-wavelength. However, when modeled with *NEC-4*, convergence only occurred with very large numbers of segments per half-wavelength.

A feature of the K6STI designs is the use of a large-diameter, short-length center section for each element to simulate boom-mounting plates. This technique has proven sound, relative to real-world antenna construction. However, only after reducing the diameters of these element centers is *NEC-4* able to achieve convergence with a reasonable number of segments per half-wavelength.[9] The initial conclusions reached from these investigations are:

1. *NEC-4* has limits in dealing with stepped-diameter elements, especially where the step in diameters between adjacent element segments is large and the large step occurs in the region of maximum element current.

Achieving convergence under these circumstances may require quite large models relative to the number of antenna elements involved. These models grow larger for every diameter step involved in the structure of the element.

2. Inadequate segmentation in

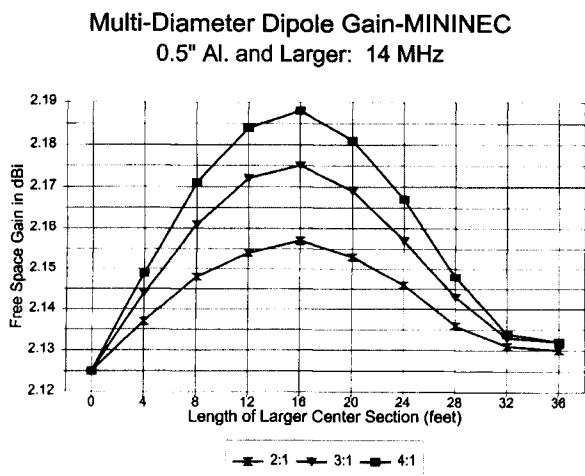
Multi-Diameter Dipole Gain-MININEC
0.5" Al. and Larger: 14 MHz

Figure 5—Multi-diameter dipole gain—*MININEC*.

Multi-Diameter Dipole Length-MININEC
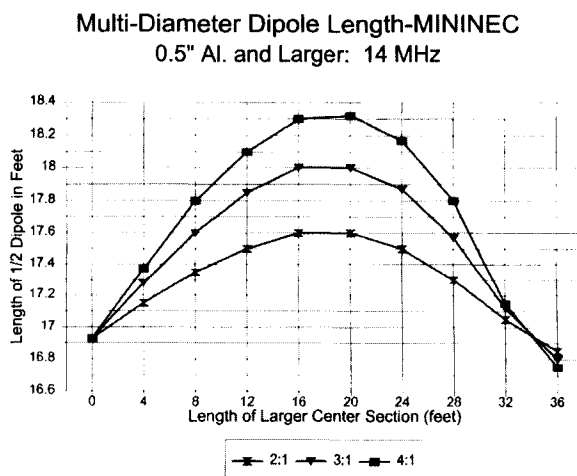**0.5" Al. and Larger: 14 MHz**

Figure 6—Multi-diameter dipole length—*MININEC*.

stepped-diameter elements in *NEC-4* may result in unrealistically high values of forward gain and low values of feedpoint impedance. Adequacy of segmentation includes the number of lements, equalization of segment lengths within the limits of the element structure and alignment of segments among the elements.

To explore these anomalies more systematically, I set up a small test systematically modeling dipoles with larger center sections. The test frequency was 14 MHz. The initial material was 0.5 inch-diameter aluminum. In increments of two feet each side of center, I increased the length of a larger diameter center section in progressive steps for total lengths of 4, 8, 12 feet, etc, up to and including the total antenna length.

Each antenna was then resonated to less than 1 Ω reactance. Thus, the length of each model differs. In the graphs that follow, the far right entry labeled "36" is a placeholder for the actual total length of the antenna at the increased diameter. That column, alone, violates the linear progression of the other enlarged center sections.

Each model was tested using ratios of 2:1, 3:1 and 4:1 relative to the original dipole diameter of 0.5 inch, for center-section diameters of 1, 1.5 and 2 inches. All antennas were modeled in free space.

The models were first run in *MININEC* 3.13 via *ELNEC 3.* Segmentation was set at 34 segments overall, using one segment per foot of enlarged center section, with the remaining segments split between the smaller-diameter end sections. This yielded segment lengths that are well within all *MININEC* boundaries for accurate results, as well as very reasonably close in length between the antenna wires.

Figure 5 shows the progression of gain figures for the *MININEC* runs. Interestingly, the gain of the models peak when the enlarged center section is just under half the total length of the antenna. As the larger center section is further lengthened, *MININEC* shows a decrease in gain. The curves for the three ratios are nicely congruent.

The length of the resonant antenna also changes with the length of the larger-diameter center section, as shown in Figure 6. Overall antenna length actually peaks with center-section lengths slightly longer than those for maximum gain. *MININEC* models do not reach a final shortened length associated with fatter elements until the entire antenna is at the larger diameter.

The same antennas were run with *NEC-4*, initially with EZNEC Pro and later with a beta version of *GNEC.* Segmentation was virtually the same as with the *MININEC* models with a single additional segment in the center section to permit the required midsegment feed-point. As with the *MININEC* models, each antenna was resonated to less than ±1 Ω of reactance for models using 2:1, 3:1 and 4:1 ratios of the center segment to the end sections.

The pattern of gain produced by the *NEC-4* models, shown in Figure 7, is quite unlike that yielded by *MININEC.*

Maximum gain occurs with the shortest possible larger-diameter center section and progressively decreases as the center section is lengthened. The curves for *NEC-4* are less smooth than for *MININEC* because the former, in the commercial versions noted, yields gain figures to two decimal places, while the latter yields figures to three decimal places. Hence, *NEC-4* rounding yields somewhat stair-step curves. Within those limits, the curves for the three different diameter ratios are congruent.

Equally congruent are the overall antenna-length curves, as demonstrated by Figure 8. Notice that in all graphics in this series, antenna length is shown in terms of lengths each side of the feed-point.

Interestingly, despite the vastly different gain curves, the overall antenna length curves for *MININEC* and *NEC* are exceedingly comparable. *MININEC* yields slightly longer resonant lengths for each modeled case, a phenomenon long noted (and corrected for in some commercial versions of *MININEC*). Nonetheless, *MININEC* and *NEC-4* show the longest resonant length at just about the same length of larger-diameter center section, as shown in Figure 9. This graph uses the 4:1 ratio curves because they produce the sharpest length peaks and would be most sensitive to significant differences in the peaks for each modeling system, with these models.

The question that remains is "Which of the two gain curves is the more reliable?" The *MININEC* gain curves with a 2:1 diameter ratio of center
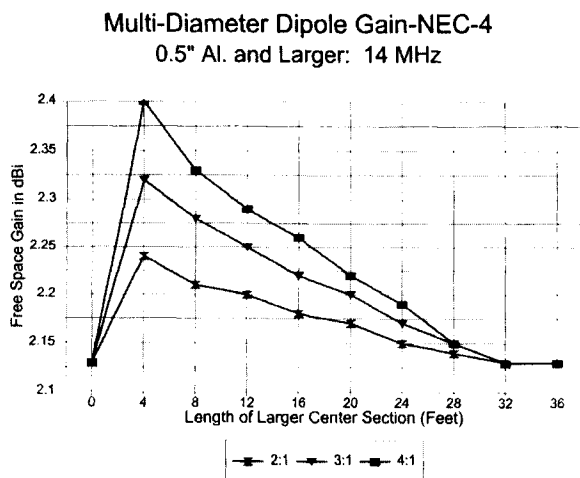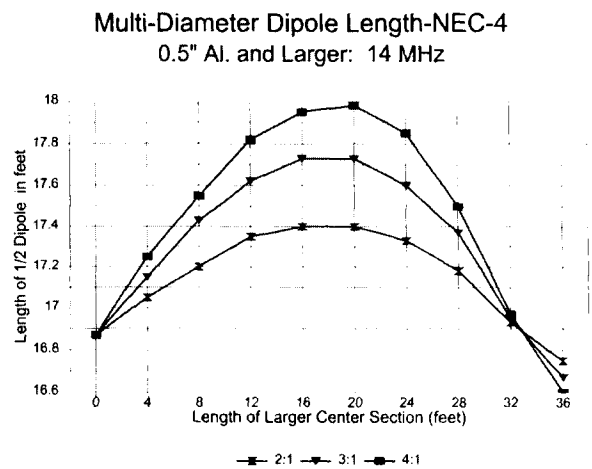

Figure 7—Multi-diameter dipole gain—*NEC-4.*

Figure 8—Multi-diameter dipole length—*NEC-4.*

section to end sections were rerun with twice as many segments. Gain figures are convergent, within a maximum divergence of 0.003 dB. In the comparable *NEC-4* models, segment numbers were doubled (and one segment removed from the wire containing the feed-point to retain an odd number of segments). The results show an order of magnitude less convergence with short center sections of 4 and 8 feet. Convergence within two decimal places occurred only when the center section reached 24 feet long. This result is consistent with the difficulty of convergence testing Yagi models employing short, large-diameter center sections for each element.

Moreover, the *NEC-4* curve totally envelopes the *MININEC* curve (as shown in Figure 10), for diameter ratios of 2:1. Within the limits of *NEC-4* rounding, nowhere does the *MININEC* curve exceed the *NEC-4* curve in gain value. In contrast, for shorter center sections, the reported gain of the *NEC-4* model is significantly higher.

The anomalous results yielded by *NEC-4* with shorter, large-diameter center sections of multidiameter dipoles gains importance as these figures accumulate in multielement antenna arrays. Since these gains are usually cumulative in most antennas designed for maximum or close to maximum gain, the reported gain may be significantly higher than reality.

The convergence test and the fact that the *NEC-4* curve envelopes the

MININEC curve are strong indicators that *NEC-4* may be simply inaccurate when antenna elements consist of multidiameter sections, such that the center section is short and significantly larger in diameter than succeeding sections of the element. The degree to which such an inaccuracy becomes operationally significant to an antenna design depends on many variables of both design and engineering goals and cannot be independently estimated. Unless *NEC-4*'s results can be independently confirmed as accurate (with *MININEC*'s curves consequently invalidated), the phenomenon will be, at least, disconcerting to antenna modelers.

Attributing an anomaly to *NEC-4* in this case does not, itself, certify the accuracy of the *MININEC* result. Even if correct, the increase of gain may be of more mathematical interest than operational significance in many cases. At a diameter ratio of 4:1, the maximum gain is only about 0.06 dB relative to a dipole of the thinner size. However, where such gain increases accumulate on multielement antennas, the resulting gain figures may yield unrealizable expectations of real antenna performance. Those who model antennas should note the disparity between the two modeling programs wherever it emerges.

A Second Lesser-Known Limitation: Closely Spaced Wires

With respect to closely spaced par-

allel wires, *NEC-4* documentation notes that minimum separation is limited by the thin-wire approximation and that the actual error versus separation had not been well determined at the time of the manual release. The manual sets a "reasonable" limit on separation between wire axes of two to three times the largest diameter.[10]

Numerous Amateur Radio antennas contain quite closely spaced structural elements. Among these structures are open-sleeve coupled elements, Tee and gamma-match rods and folded dipoles. Therefore, it seemed reasonable to explore more methodically the possibility of a further systematic error. Apparently anomalous results occurred when antenna wires in *NEC-4* were placed in close proximity, despite following recommended guidelines for aligning the segments of the wires to the degree possible.

Therefore, I performed a simple modeling test. I modeled a 1 inch diameter aluminum dipole for 14 MHz. Then I created three different models of resonant 21 MHz dipoles, all aluminum, but having diameters of 0.5, 1.0 and 1.5 inches.

In separate tests, I placed each of these 21 MHz dipoles in proximity to the 14 MHz dipole at distances of 2 through 12 inches, in 2 inch increments. The 2 inch spacing was deemed the least permissible that would prevent the wire surfaces from touching. This spacing falls below the recommended *NEC-4* guidelines of wire
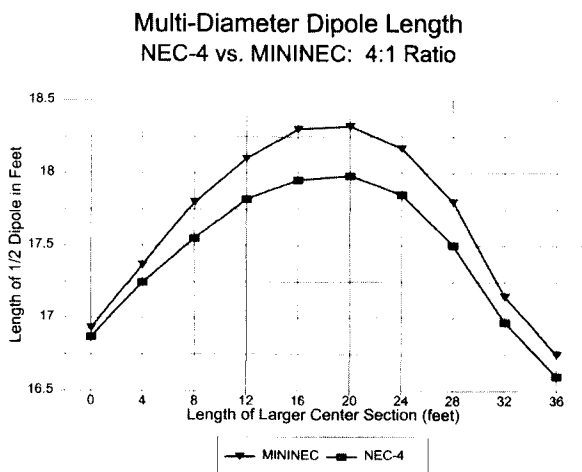
Figure 9—Multi-diameter dipole length—*NEC-4* versus *MININEC*.
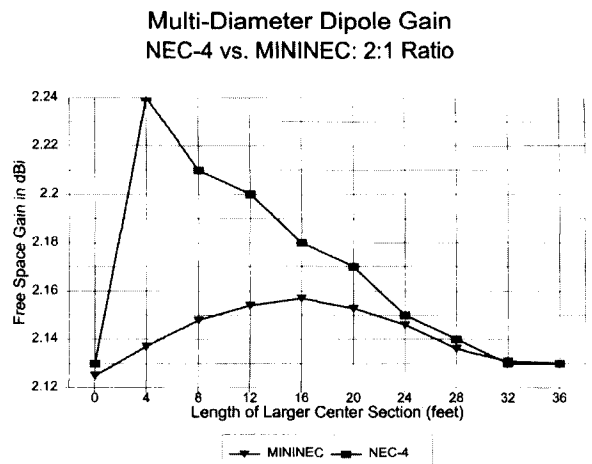
Figure 10—Multi-diameter dipole gain—*NEC-4* versus *MININEC*.

separation. For each test, the 14 MHz fed dipole was adjusted to resonance and the gain figure was recorded. As in past tests, resonance is defined as a feedpoint reactance of less than 1 Ω. Since the test was designed to reveal the effect of the shorter wire on the longer, the original lengths of the 21 MHz antennas were preserved for each size throughout the test runs.

The tests were performed on both *MININEC* 3.13 within *ELNEC* 3.0 and on *NEC-4* within *EZNEC Pro*. The lengths of the 14 MHz and the 21 MHz elements differed between programs from 0.5 inch for the smallest diameter element to about 1 inch for the largest.

The 14 MHz antenna used 34 segments in *MININEC* and 35 segments in *NEC-4*. The 21 MHz element was assigned 22 segments in *MININEC* and 23 in *NEC-4*. This segmentation aligned the segments quite reasonably. Since this segmentation already exceeds common practice in linear antenna design, convergence testing was not systematically undertaken, although the same performance curves appear with both fewer and more segments per half-wavelength.

All gain figures were recorded as free space gain in dBi. With respect to closely spaced elements, there are two gain figures of note: the in-plane gain and the out-of-plane gain. The former is the maximum gain of the dipole and extra wire in the plane that contains them both. The latter is the gain in a plane that is perpendicular to the wire axes and passing through the wire midpoints. Out-of-plane gain will ordinarily be less than in-plane gain, although in-plane gain will show a front-to-back ratio that is approximately double the difference of the two gain figures.

The tests were first run with *MININEC*. Figure 11 shows the in-plane gain for each diameter of extra wire as the distance between elements is decreased from 12 to 2 inches. Most noticeable in the graph is the flattening of the curves as the spacing reaches 2 inches, despite a reasonably linear progression to that point.

Some of the reason for the flattening appears in Figure 12, which records the out-of-plane gains for the wire pairs over the same range of spacings. As the distance reaches 2 inches,

Close-Spaced Wires: MININEC
In-Plane Gain vs. Wire Spacing
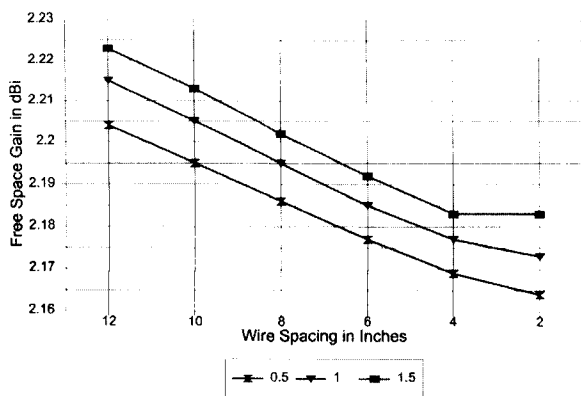
Figure 11—In-plane gain of close-spaced wires—*MININEC*.

Close-Spaced wires: MININEC
Out-of-Plane Gain vs. Wire Spacing
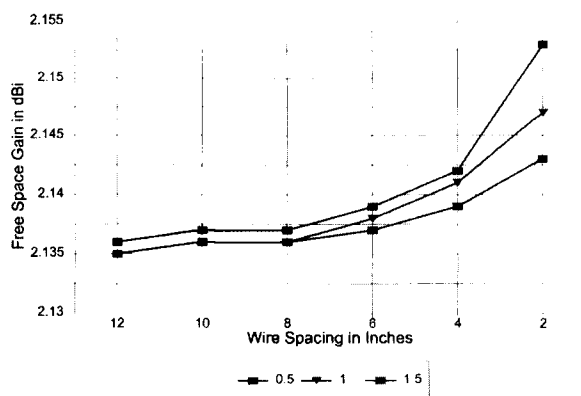
Figure 12—Out-of-plane gain of close-spaced wires—*MININEC*.

Close-Spaced Wires: NEC-4
In-Plane Gain vs. Wire Spacing
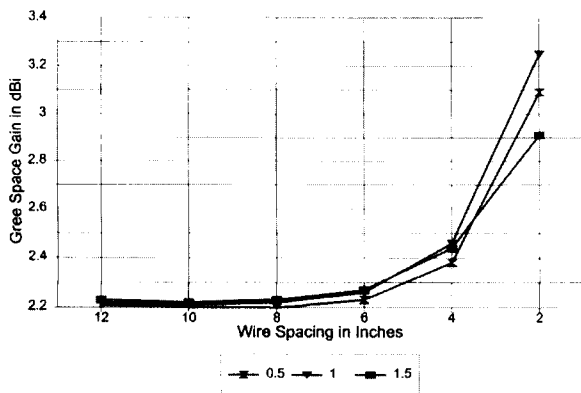
Figure 13—In-plane gain of close-spaced wires—*NEC-4*.

Close-Spaced Wires: NEC-4
Out-of-Plane Gain vs. Wire Spacing
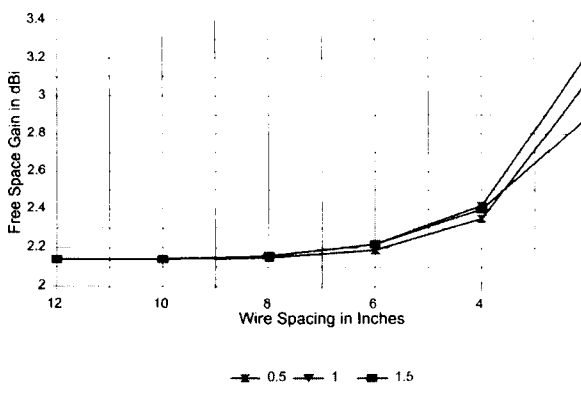
Figure 14—Out-of-plane gain of close-spaced wires—*NEC-4*.

the gain begins a steep increase. To all appearances, the two wires begin at this close proximity to act as a single fat wire. Whatever the true accuracy of the *MININEC* results, they do at least accord with normal expectations for closely spaced wires.

The *MININEC* results acquire a greater degree of confidence when one examines the range of variation. The total in-plane gain variation is less than 0.06 dB, while the out-of-plane gain variation is less than 0.02 dB.

The gain figures encountered with *NEC-4* show a quite different pattern. For example, the in-plane figures, which appear in Figure 13, show an overall increase through the same range of distances separating the two wires. The range of variation between 12 and 6 inch spacing is about 0.04 dB, but over the entire span of separations, the range increases to more than 1 dB.

Equally notable is the fact that *NEC-4* shows the highest gain when the two wires have the same diameter. How exact this equality is cannot be determined by this test, since the secondary wire diameters are widely separated.

A similar curve accompanies the figures for out-of-plane gain in Figure 14. There is a larger spread of gains in the 12 to 6 inch range (0.06 dB), but the overall gain increase with closing separation is greater than 1 dB.

At least for the case at hand, which uses a secondary wire about 66% as long as the fed wire, there appears to be a critical distance at which anomalous results begin to emerge. Figures 15 and 16 compare the in-plane and out-of-plane gains for *MININEC* and *NEC-4* when the elements have the same diameter.

Careful examination of the graphs shows that the curves overlap for spacings of 12 and 10 inches. However, between the 10 and 8 inch marks, the curves begin to diverge ever more radically.

It is also interesting to contrast *MININEC* and *NEC* with respect to the required length of the 14 MHz element for resonance with the 21 MHz wire in close proximity. As
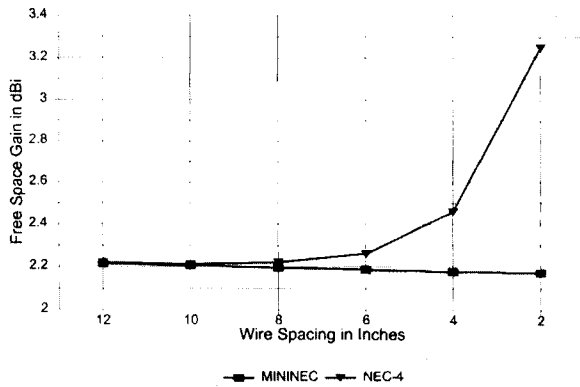
Figure 15—In-plane gain of close-spaced wires—*MININEC* versus *NEC-4*.
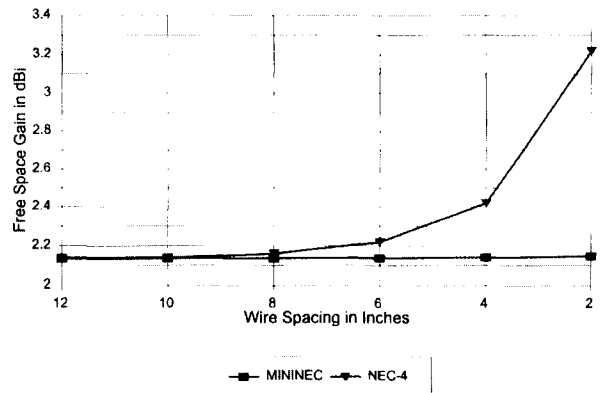
Figure 16—Out-of-plane gain of close-spaced wires—*MININEC* versus *NEC-4*.

Figure 17—14 MHz wire length—*MININEC* versus *NEC-4*.

Figure 18—14 MHz resistance—*MININEC* versus *NEC-4*.

Figure 17 shows, once we allow for the slight length variation of the two systems, the *MININEC* curve is much steeper than the *NEC-4* curve at the closest spacings. However, the curve is in fact smoother than the *NEC-4* curve, as the break in the *NEC-4* curve at 6 inches is real and not a function of rounding.

Moreover, there are also interesting differences in the feed-point-resistance curves. As noted, the 14 MHz elements were resonated to less than 1 Ω reactance. In Figure 18, the *MININEC* feed-resistance curve shows a small dip at the 8 inch spacing and then a smooth progression upward. In contrast, the *NEC-4* curve shows a rapid progression downward past the 6 inch point. These same phenomena occurred with scaled VHF antenna models.

These tests are only the beginning of a systematic exploration of the differential in gain and other figures from the *MININEC* and *NEC-4* modeling systems. Here are a few further developments of these tests.

Scaling: Scaling all the dimensions of the situation by a factor of 10 upward (including frequency, lengths and wire diameters) produces curves that tightly fit those produced so far. This applies to both the *MININEC* and *NEC-4* curves. The resonant 140 MHz (0.1 inch diameter aluminum wire spaced from 1.2 to 0.2 inches) offsets from the resonant 14 MHz curves (1.0 inch diameter aluminum wire spaced from 12 to 2 inches) are so slight that they are negligible.

Wire Length Ratios: The 20-meter

antenna with a 15-meter wire forms a length ratio of 3:2. A series of *NEC-4* runs compared this scenario to wires with a 2:1 ratio and with a 4:3 ratio, focusing on wires of the same diameter (1 inch aluminum). The results appear in Figure 19.

Interestingly, the departure of the gain from a typical *MININEC* curve is greatest when the fed wire is about 50% longer than the closely spaced unfed wire, at least when the wires have the same diameter. Once more, the widely separated selection of test ratios does not lend precision to this conclusion. Since wires having a 1:1 diameter ratio appear to have a greater departure from typical *MININEC* curves than other wire diameter ratios, it appears that (by chance) my initial tests have fallen into at least the ball park of greatest deviation.

Wire Diameters and Spacing: Just as wire-length ratios may be isolated for specific investigation, so too may be the relationship between wire diameters and spacing. A series of models were undertaken in both *MININEC* and *NEC-4* using a constant 3:2 wire-length ratio between the driven wire and its closely spaced undriven companion. Both wires were assigned the same diameter and checked at the standard 2 inch spacing increments. Wire size was varied through aluminum wire diameters of 1.0, 0.5, 0.1, 0.05 and 0.01 inch. A limitation of this test is that, for any spacing, the surface-to-surface distances vary as the wire size is changed.

Figure 20 correlates, for *MININEC*

runs, the wire size and gain for 14 MHz dipoles with companion wires about 2/3 as long. Only with the largest wire diameter (1.0 inch) does the figure reveal a graphically detectable departure from otherwise flat results. (Note: The very low gain of the 0.01 inch diameter wire is due largely to the losses associated with using aluminum wire.)

Figure 21 shows the same runs (with wires adjusted in length for resonance) with *NEC-4*. Even with the thinnest wire size used, a graphically evident departure from a relatively constant gain is shown between 8 and 6 inch separations. For fatter wires, the increase in reported gain appears as early as between 10 and 8 inch separations. Although the erroneous gain increase report may not be operationally significant in many instances, the trends are at odds with the minimum separation recommendations in *NEC-4* documentation.

Frequency and Spacing: The effects of frequency on gain and feed resistance reported by *NEC-4*, when the wire size is held constant and spacing is varied, are also interesting. Models were constructed using 1 inch aluminum wire, with a 3:2 ratio of driven wire to companion wire lengths for 3.5, 7.0, 14.0, 21.0 and 28.0 MHz. Wire spacing was varied in 2 inch increments from 12 to 2 inches as with preceding models.

Figure 22 shows the reported antenna gain for the models in *NEC-4*. The 80-meter rise is steepest because the spacing represents a smaller fraction of a wavelength. Figure 23 presents the reported feed-point resistance
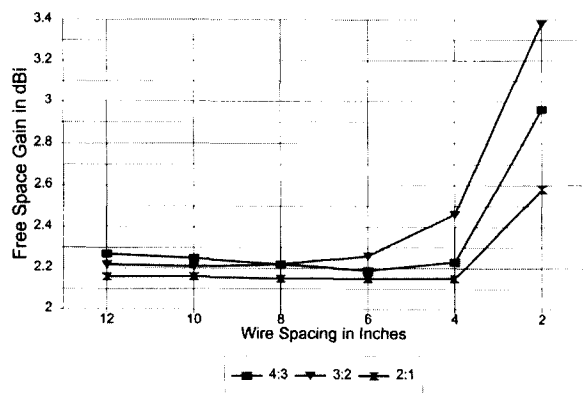
Figure 19—14 MHz gain versus wire spacing for different length ratios.
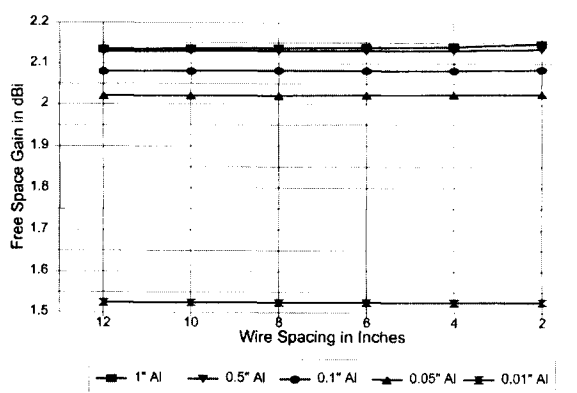
Figure 20—Close-spaced gain for wire sizes from 0.01 to 1.0 inch diameters—*MININEC.*

for each of the resonated models for each frequency and spacing. For all but 28 MHz, the downward curve of reported feed-point resistance is roughly proportional to the rise in reported antenna gain.

The seeming exception is 28 MHz, where the impedance appears to rise continuously. In actuality, there is a knee to each curve. For 3.5 and 7.0 MHz, the knee lies off the graph to the left. For 14 MHz, the knee that shows a higher feed-point resistance than adjoining values is at the 8 inch spacing marker, while for 21 MHz, the knee is visible in the graph at 4 inch spacing of the wires. The knee peak for 28 MHz is approximately at the 2 inch spacing marker, giving the illusion of further rise. However, with closer spacings, the feedpoint resistance may begin to decline. Since a 2 inch spacing is already beyond even the limits recognized by *NEC-4* documentation, the graph was not carried further than the values shown. However, it appears clear that close-spacing phenomena bear a relationship to the fraction of a wavelength that a given spacing represents.

Close Spacing and Multielement Arrays: The effects of a close-spaced wire on a dipole model are only indicators, not predictors of the effect of a close-spaced wire on a parasitic beam model. Indeed, the disparity of gain and other performance figures between *MININEC* and *NEC-4* might well be either profound or quite trivial.

I created a three-element Yagi model to check the potential for divergent readouts. An extra wire was placed ahead of the driven element at the spacing indicated in the tables. The results appear in Table 2.

The table holds some surprises. First, the *NEC-4* gain values diverge more radically than the *MININEC* numbers, especially for the 2:1 ratio of driven element to extra wire. Second, for the 3:2 and 4:3 ratios, *MININEC* and *NEC-4* gain numbers diverge in opposite directions. Nonetheless, the *NEC-4* figures are still farther from the "no-wire" baseline than those of *MININEC*. Third, unlike the simple dipole examples, the gain of some models may decrease in the presence of the extra wire.

Before we draw any conclusions, let me reveal that the above table is erroneous. It is based on a defective model that is nevertheless all too common in amateur modeling practice. All four elements, the three 20-meter elements plus the added wire, whatever its length, were assigned 10 segments in *MININEC* and 11 segments in *NEC-4*. The segments do not align, which is especially important in *NEC*, but significant in *MININEC* with some models. Moreover, neither model converges well with models having twice as many elements. These problems cast doubt on the reliability of the results.

Table 2—Closely Spaced Wires in *MININEC* and *NEC-4* Models of a 3-Element Yagi*

| Space (inches) | MININEC | | | NEC-4 | | |
	Gain (dBi)	FB (dB)	Z (Ω)	Gain (dBi)	FB (dB)	Z (Ω)
No wire	8.03	24.97	26.9 + j2.1	8.08	27.94	26.6 + j4.8
2:1 Ratio						
4	8.02	24.86	26.9 + j1.1	8.68	27.81	23.4 + j3.4
7	8.03	24.87	27.0 + j1.3	8.26	27.84	25.7 + j3.9
10	8.03	24.88	27.0 + j1.4	8.14	27.85	26.4 + j4.1
13	8.03	24.88	27.0 + j1.5	8.10	27.86	26.6 + j4.2
16	8.03	24.89	27.0 + j1.5	8.09	27.88	26.7 + j4.2
3:2 Ratio						
4	7.88	24.77	28.0 + j0.0	9.35	27.68	20.3 + j1.8
7	7.96	24.79	27.6 + j0.4	8.61	27.76	24.1 + j2.6
10	8.00	24.82	27.5 + j0.6	8.33	27.81	25.6 + j3.0
13	8.01	24.84	27.4 + j0.7	8.21	27.86	26.4 + j3.2
16	8.02	24.87	27.3 + j0.7	8.15	27.91	26.7 + j3.3
4:3 Ratio						
4	8.20	24.74	26.2 − j1.0	6.81	27.93	37.7 + j3.3
7	8.12	24.78	27.0 − j0.5	7.34	27.91	33.1 + j3.6
10	8.08	24.82	27.4 − j0.3	7.67	27.93	30.6 + j3.4
13	8.06	24.87	27.5 − j0.2	7.85	27.98	29.2 + j3.1
16	8.05	24.91	27.6 − j0.1	7.94	28.04	28.5 + j3.0

*This table is based on erroneously constructed models. See text for an explanation and see Table 3 for corrected figures.
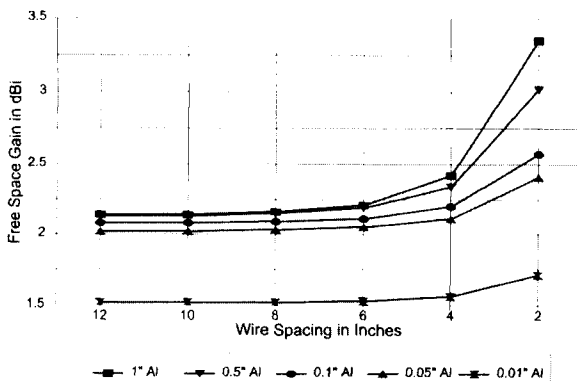
Figure 21—Close-spaced gain for wire sizes from 0.01 to 1.0 inch diameters—*NEC-4*.
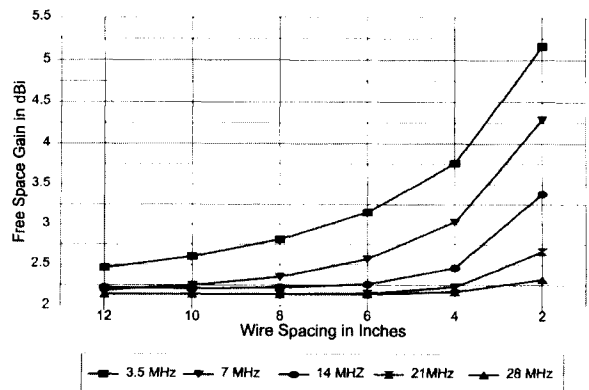
Figure 22—Gain versus frequency and spacing.

So I reset the models, using 34 segments per 20-meter element (35 in *NEC*) and a proportionate number for the shorter extra elements. Convergence with models having twice as many segments was excellent, with a gain difference of about 0.01 dB. Running these models resulted in a change of extra-element spacing, to begin at closest with 3 rather than 4 inches. Despite the closer spacing, interesting results developed, as shown in Table 3.

A comparison of the tables shows two very significant facts. First, when models are developed carefully, rather than casually, any tendencies for a program to deliver potentially erroneous results is lessened. All figures for each length of extra element are far more tightly grouped.

Second, despite the tighter grouping, the same types of curves develop as with the casual model. To two decimal places, *MININEC* results are totally stable, although the third decimal place shows the mathematical progressions that appeared in the earlier model. Likewise, *NEC-4* progressions show increasing gain with closer spacing for the two shorter lengths of extra elements and reduced gain with closer spacing for the longest extra element. In general, the instability with the figures occurs when spacings are closer than 6 to 9 inches at 14 MHz. Maximum deviations from the norm run from 0.2 dB to 0.5 dB for the example used. While less than with the casual model, the amounts of deviation from the normal can be significant, especially when compared to the extremely stable *MININEC* figures.

Conclusion: Each of these directions of research will require many more runs of wire combinations at many different frequencies before precise conclusions and systematic formulations of the *MININEC/NEC-4* differentials can be drawn. Nonetheless, the simple tests performed here are sufficient to suggest strongly that users of *NEC-4* model closely spaced wires with great caution. If *NEC-4* proves to be the anomalous case, then it may not be possible to routinely model closely spaced antenna structures with any presumption of accuracy with respect to resulting gain figures.

As with all such tests, the appearance of results is not a sufficient validation of a modeling system. Nonetheless, it appears safe to note that closely spaced wires modeled in *NEC-4* should always be approached with more caution than confidence.
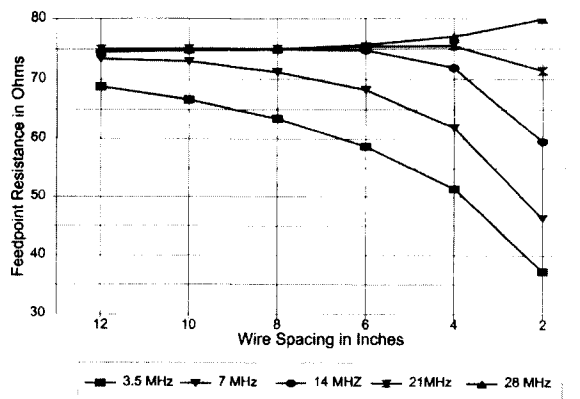
Figure 23—Feed-point resistance versus frequency and spacing.

Table 3 —Closely Spaced Wires in *MININEC* and *NEC – 4* Models of a 3-Element Yagi (Corrected Models)

Space (inches)	*MININEC* Gain (dBi)	FB (dB)	Z (Ω)	*NEC – 4* Gain (dBi)	FB (dB)	Z (Ω)
No wire	8.01	26.66	28.4 + j5.4	8.12	26.86	26.3 + j6.6
2:1 Ratio						
3	8.01	26.64	28.9 + j4.1	8.40	26.72	25.0 + j5.2
6	8.01	26.65	28.7 + j4.5	8.15	26.74	26.4 + j5.8
9	8.01	26.66	28.7 + j4.7	8.12	26.75	26.5 + j6.0
12	8.01	26.67	28.6 + j4.8	8.12	26.77	26.5 + j6.0
15	8.01	24.68	28.6 + j4.8	8.12	26.78	26.5 + j5.2
3:2 Ratio						
3	8.00	26.52	29.3 + j3.7	8.33	26.62	26.0 + j4.1
6	8.00	26.57	29.2 + j3.7	8.14	26.66	27.0 + j4.9
9	8.00	26.61	29.1 + j3.6	8.12	26.69	27.0 + j5.2
12	8.00	26.65	29.1 + j3.3	8.12	26.73	26.9 + j5.3
15	8.00	26.70	29.1 + j2.6	8.12	26.76	26.9 + j5.3
4:3 Ratio						
3	8.01	26.73	30.0 + j1.6	7.40	26.69	33.3 + j3.9
6	8.01	26.80	29.8 + j2.6	8.03	26.66	28.3 + j4.4
9	8.01	26.87	29.7 + j2.9	8.10	26.70	27.7 + j4.5
12	8.01	26.95	29.6 + j3.1	8.11	26.76	27.5 + j4.6
15	8.01	27.02	29.6 + j3.1	8.12	26.82	27.4 + j4.7

Table 4—Single Quad Loops of a Single-Wire Diameter

Antenna	Output	*MININEC*	*NEC-2*	*NEC-4*
0.0808" wire	Gain		3.26	3.26
9.146'; 31 segs	Feed Z		126.9 + j0.02	126.9 – j0.13
0.0808" wire	Gain	3.25	3.26	3.26
9.146'; 61 segs	Feed Z	126.3 – j7.93	127.0 – j0.27	127.0 – j0.69
0.0808" wire	Gain	3.26		
9.146'; tapered	Feed Z	126.0 – j3.38		
0.5" wire	Gain		3.37	3.37
9.364'; 31 segs	Feed Z		129.7 + j0.21	129.7 – j0.10
0.5" wire	Gain	3.36	3.36	3.26
9.364'; 61 segs	Feed Z	129.5 – j4.07	130.0 + j0.49	129.8 – j0.41
0.5" wire	Gain	3.36		
9.364'; tapered	Feed Z	129.1 – j0.42		

Two Practical Limitations of *NEC-4* and a Validation Test

The importance of respecting the limitations of *NEC-4* appears in many potential applications to amateur HF (and VHF) antennas, even of relatively standard design. I shall illustrate with just two examples, the latter of which will also comprise a validation of *MININEC* as the more accurate program in these regards.

Wires of Different Diameters Joined at Sharp Angles: A problem in the *NEC-2* calculation engine is the unreliability of results when wires of unequal diameter join at right or acute angles. Although *NEC-4* improves upon this situation, its results are not wholly reliable.

Single Quad Loops of a Single Wire Diameter: The foundation for testing the reliability of *NEC-4* outputs when wires of different diameter join at right angles is the single quad loop. The test employed loop materials of 0.5 inch and 0.0808 inch diameters. When only a single diameter wire is used, all programs perform credibly, as long as models adhere to the antenna geometry criteria of the specific program. All loops were modeled at 28.5 MHz, with copper wire in free space. All loops are square. Dimensions and segmentation are given for one side of the loop. The tapered-segment *MININEC* model employs the internal values of the *ELNEC* program.

Table 4 provides the results of the modeling. The initial models were created in *NEC-4* and tested on *NEC-2* and *MININEC*. A tapered-segment-length model was created in *MININEC* for comparison with the equal-segment models. Convergence of the two *MININEC* models is good for practical purposes, although a slight numerical difference shows.

For the *NEC* models, there is no significant numerical, let alone practical, difference between *NEC-2* and *NEC-4* models. Moreover, and especially significant for this test, there is no significant difference between the values achieved at 31 segments per side and 61 segments per side. Practical convergence of results is achieved at much lower levels of segmentation.

Single Quad Loops with Different Wire Diameters: To test the ability of the programs to handle wires of different diameter joining at right angles, I modeled a single square quad loop. The top and bottom wires were 0.5 inches in diameter, while the vertical wires were 0.0808 inches in diameter.

This might be a model of a portable quad loop using tubing for the horizontal members and wire for the vertical pieces, thus allowing the assembly to collapse for transportation.

The initial model was constructed in *MININEC* and then tested in *NEC-2* and *NEC-4*, using 61 segments per side. The *MININEC* model required 10.15 foot side lengths to approach modeled resonance. Table 5 shows the modeling results.

The *MININEC* model converges well with its tapered alternative model. However, the *NEC* models diverge in values. Although the gain values are plausible, the feed-point impedance values indicate a condition far from resonance. The divergence from *MININEC* is worse for *NEC-2* than for *NEC-4*, suggesting that the *NEC-2* figures are least reliable. Since there is no simple theoretical calculation with which to compare the overall results, one cannot claim that *MININEC* qualifies as a standard against which to measure the other programs. However, given *MININEC*'s ability to handle wires of different diameters in other contexts and the general trend of *NEC-4* results under such conditions to be closer than *NEC-2* results to the *MININEC* figures, it seems likely that *MININEC* may yield outputs that are closest to reality among the three.

The *NEC-4* model with 61 segments per side can be brought closer to resonance by shortening each side to 9.94 feet. This figure might seem equally reliable with the *MININEC* lengths of 10.15 feet per side, except for one significant factor: The *MININEC* figures achieve convergence, while the *NEC* figures do not, especially with respect to feedpoint impedance. I ran the revised *NEC-4* model through various segmentations ranging from 21 to 121 segments per side. In Table 6, "ΔR" and "ΔX" indicate changes in the feedpoint impedance values from the preceding level of segmentation.

The values for gain are well converged, but those for feed-point impedance are not. Compare, for example, the differences among figures for 31 and 61 segments per side for the equal-diameter wire loops, using either the 0.0808 or 0.5 inch models. *NEC-4* varies by only 0.1 Ω resistance and under 0.5 Ω reactance across that spread. With the present unequal wire-diameter loop, the same difference in segmentation yields a difference of 2.8 Ω resistance and 10.991 Ω reactance, a +200% difference for each output figure.

Moreover, the progression of values shows no signs of closure within the limits of practical modeling. Although there is a trend downward in the delta numbers, where closure will occur remains unclear. Without convergence, the figures cannot be regarded as reliable.

Table 5—Single Quad Loops of Different Wire Diameters

Antenna	Output	MININEC	NEC-2	NEC-4
10.15' sides	Gain	3.61	3.57	3.60
61 segs/side	Feed Z	$137.2 - j5.71$	$175.4 + j140$	$150.3 + j44.3$
10.15' sides	Gain	3.61		
tapered segs	Feed Z	$136.7 - j2.30$		

Table 6—Revised NEC-4 Model of a Quad Loop Using Different Wire Diameters

Segments Per Side	Gain (dBi)	Feed-Point Impedance (Ω)	ΔR*	ΔX*
21	3.54	$133.4 - j13.280$		
31	3.54	$134.5 - j9.375$	1.1	3.905
41	3.53	$135.5 - j5.589$	1.0	3.786
51	3.53	$136.4 - j1.943$	0.9	3.646
61	3.53	$137.3 + j1.616$	0.9	3.559
71	3.53	$138.3 + j5.119$	1.0	3.503
81	3.53	$139.2 + j8.663$	0.9	3.544
91	3.53	$140.0 + j11.880$	0.8	3.217
101	3.52	$140.9 + j15.36$	0.9	3.48
111	3.52	$141.7 + j18.57$	0.8	3.21
121	3.52	$142.4 + j21.60$	0.7	3.03

*See text for an explanation of ΔR and ΔX.

Conclusions and Implications: Because there is no independent standard at hand against which to measure the modeled results, the *MININEC* figures for the single quad loop cannot be certified as in fact closer to reality than those yielded by *NEC-4* for antennas constructed of different-diameter wires joining at right angles. However, *MININEC*'s achievement of reasonable convergence of results and *NEC-4*'s inability to achieve converged results suggests that the *NEC-4* results are less trustworthy than those of *MININEC*. *NEC-2* figures are most divergent and least reliable of the three modeling calculation engines.

It is clear that *NEC-4* will yield lower gain numbers and higher feed-point values than *MININEC* for a loop of a given size. Otherwise expressed, *NEC-4* will call for a loop of smaller dimensions to approach resonance.

These trends also apply to other antennas using wires of different diameters joining at right and acute angles. Models of folded X-beams show lesser gain and greater feed-point values on *NEC-4* than on *MININEC*.

Given the limitation of *NEC-4* with respect to parallel wires of different diameters, it is probable that the present limitation of *NEC-4* is an extension of the same root mechanism. Therefore, it is likely that *MININEC* remains the modeling engine of choice for antennas employing angular junctions of different-diameter wires.

Folded Dipoles: As a final test of *NEC-4* limitations, let us turn to even more compact closed antenna geometry—the folded dipole. Because the characteristics of the folded dipole are so well known, it is possible to calculate in advance the feedpoint impedance of a folded dipole using antenna wires of any ratio. This will provide a test of whether the presumed greater accuracy of *MININEC* in cases of the order discussed here is, in fact, justified.

Equal Diameter Folded Dipoles: The actual test consists of modeling a folded dipole. A folded dipole, where the long parallel wires have the same diameter, effects an impedance transformation of 4:1 for any spacing within reason. Thus, the anticipated feed-point impedance should be in the region of $288\,\Omega$ (72×4). Since folded dipoles also act like fat wires and are thus shorter at resonance than single-wire dipoles, the anticipated modeled feedpoint impedance was slightly lower than the theoretical calculation. The modeled folded dipoles used 0.5 inch diameter elements spaced 0.25 foot (3 inches).

MININEC tends to chop corners and give erroneous results unless one of two procedures is followed:

One may use as many segments as the program allows to minimize the size of the corner chopped.

One may taper the segment lengths approaching the corner so that corner segments are small while the overall segment count is held to a practical minimum.

The basic *MININEC* folded dipole used 66 segments longitudinally and 2 segments at the ends. *NEC* models added one segment to each longitudinal wire to maintain parallel segmentation. Tapered *MININEC* models used the internal segmentation values of the *ELNEC* program. Since these produced eight-segment midlength wires, the *NEC* models added one segment to this section to satisfy the need for an odd number of segments for center feeding. Finally, a more highly segmented model, using 120 segments per longitudinal wire was created to equalize the segment lengths with those of the 2-segment end wires. This last model was not adjusted for resonance.

Table 7 shows the modeling results. In practical terms, all programs do a satisfactory job of modeling a simple folded dipole when both wires have the same diameter. When sufficient segments are used in *MININEC*, tapering proves less accurate, assuming that the balance of results represents a consensus close to reality.

Systematically, *NEC-4* shows slightly lower feed-point impedances for these closed models than does *NEC-2*. Nonetheless, when all wires have the same diameter and other modeling geometry guidelines are met, all modeling programs give equally usable results.

Unequal Diameter Folded Dipoles: When the wires of a folded dipole differ in diameter, they effect (relative to a single-wire dipole) a different feed-point impedance-transformation ratio than do folded dipoles with equal diameter wires. The theoretical impedance transformation ratio is given by

$$R = \left(1 + \frac{\log \dfrac{2s}{d_1}}{\log \dfrac{2s}{d_2}} \right)^2 \qquad \text{(Eq 1)}$$

Where R is the impedance transformation ratio, s is the wire spacing (center-to-center), d_1 is the diameter of the fed wire and d_2 is the diameter of the second wire, and where s, d_1 and d_2 are given in the same units.

If we use a wire 0.0808 inch in diameter (#12 AWG) for the fed wire and a wire 0.5 inch in diameter for the second wire, maintaining the 3 inch spacing, then the impedance transformation ratio will be approximately 7.47. A folded dipole of this construction would have a calculated feed-point impedance of about $533\,\Omega$. In practice, due to "fat wire" effect, we might expect a feed-point impedance slightly lower than this.

It should be noted that the impedance-transformation equation does not account for the end wires. In this

Table 7—Equal Diameter Folded Dipoles

Antenna	Output	MININEC	NEC-2	NEC-4
FD: equal seg 16.1'; 66/2×2	Gain	2.22	2.22	2.22
	Feed Z	285.7 + j0.90	285.9 + j4.10	285.8 + j3.99
FD: tapered 16.06'	Gain	2.21	2.21	2.21
	Feed Z	281.0 − j0.68	284.2 + j9.87	284.0 + j8.66
FD: equal seg 16.1'; 120/2×2	Gain	2.22	2.22	2.22
	Feed Z	285.8 − j1.80	286.0 + j2.27	285.8 + j0.51

Table 8—Unequal Diameter Folded Dipoles

Antenna	Output	MININEC	NEC-2	NEC-4
FD: equal seg 16.2'; 66/2×2	Gain	2.21	0.69	1.59
	Feed Z	530.5 + j1.47	375.2 + j25.8	462.6 + j17.4
FD: tapered 16.2'	Gain	2.21	0.37	1.22
	Feed Z	526.5 + j10.8	347.2 + j38.5	423.4 + j37.5
FD: equal seg 16.2'; 122/2×2	Gain	2.21	0.56	1.53
	Feed Z	527.6 − j2.99	364.1 + j25.1	456.0 + j15.43

test, the end wires were also 0.0808 inch in diameter.

If either version of *NEC* can handle parallel wires of differential diameters, then the results should coincide reasonably with those of *MININEC*, which takes such cases in stride. The test used models of similar construction to those for equal-diameter folded dipoles. A basic model used 66 segments per longitudinal wire and two segments per end wire; and a tapered-segment version of the antenna was created using internal tapering values. The results appear in Table 8.

The *MININEC* models clearly come very close to expectations. Since the tapered model was not adjusted to resonance, its values are lower, but the large equal-segmented model is likely more accurate.

NEC-2 models of parallel wires of different diameters (as has been well established) produce highly erroneous values. Tapering throws the values even farther off the mark. Although somewhat better, *NEC-4* values are also highly unreliable. Moreover, reducing segmentation of the *NEC-4* models produced nothing reliable. An autosegmented model at conservative minimums of 11 segments for the longitudinal wires and one segment each for the ends yielded a gain of 1.82 dBi and a feed-point impedance of 443.8 + j39.6 Ω. Further reducing segmentation to the absolute minimums of five segments per long wire and one segment per short calculated a gain of 2.64 dBi and a feed-point impedance of 371.3 + j26.07 Ω.

Conclusions and Implications: Because the behavior of a folded dipole is well-established and easily predicted, the antenna forms a very good test of the present modeling question—the adequacy of *NEC-4* to deal with parallel wires of unequal diameters. The conclusion is that *NEC-4* remains deficient in this regard, and antenna modelers are duly cautioned.

The inadequacy of *NEC-4* to model

this situation casts doubts on a number of possible modeling challenges. For example, modeling gamma and Tee matching sections as physical elements contributing to the radiation pattern as well as effecting an impedance transformation is now dubious. Direct physical modeling of phasing lines and other close-spaced structures with closed geometries or variations in wire sizes at junctions will generally not yield reliable results for HF antennas unless validated by comparison with comparable *MININEC* models. There are alternative means of modeling some structures using the network input card or transmission lines. Moreover, careful construction of substitute models may yield results that can be tested against *MININEC* models. However, direct physical modeling of such structures often pushes *NEC-4* beyond its limits of reliability.

For situations with parallel wires of unequal diameter, close-spaced wires, or unequal diameter elements with large changes in diameter between wires, *MININEC* remains the modeling program of choice, despite its other limitations.[11] These other limitations are as important to respect as those we have uncovered in *NEC-4*.

Notes

[1]For a short history of antenna modeling software, see R. P. Haviland, W4MB, "Programs for Antenna Analysis by the Method of Moments," *The ARRL Antenna Compendium* (1995), pp 69-73.

[2]*NEC-4* is a proprietary code of the Lawrence Livermore National Laboratory, University of California, from whom a user-license must be obtained. Export restrictions apply. To obtain a user-license, contact Gerald J. Burke, L-156, Lawrence Livermore National Laboratory, PO Box 5504, Livermore, CA 94550. The price of the license is $850 ($150 for an approved educational site). There are only two commercial programs covering *NEC-4* available currently. One source is Roy Lewallen, W7EL. *EZNEC Pro* has an option for *NEC-4* (*EZNEC/4*), if the purchaser has a confirmed license for *NEC-4* ($600). The second source is Nittany-Scientific's *GNEC*, which is scheduled for

appearance before you receive this article. The price is $795. Contact information for each of these sources was provided in my earlier article, "*NEC* and *MININEC* Antenna Modeling Programs: A Guide to Further Information," *QEX*, Mar/Apr 1998, pp 47-49.. All modeling for this article was initially done in *EZNEC Pro*, with confirmation models developed on a beta version of *GNEC*.

[3]The fundamental document for *NEC-4* users is Gerald J. Burke, *Numerical Electro-magnetics Code—NEC-4: Method of Moments*, UCRL-MA-109338 (LLNL), 1992. The manual appears in three parts: I. *NEC* User's Manual; II. *NEC* Program Description—Theory; and III. *NEC* Program Description—Code. References here will be confined to Part I, which we shall abbreviate *NEC-4-I* in further notes.

[4]The limitations noted so far appear in *NEC-4-I*, pp 3-4.

[5]Roy Lewallen, W7EL, *EZNEC Pro* (User's Manual, 1997), p 52.

[6]*NEC-4-I*, p 4.

[7]*NEC-4-I*, p 195.

[8]*YA*, a specialized version of *Yagi Optimizer*, by Brian Beezley, K6STI, is distributed with current editions of *The ARRL Antenna Book* (ARRL), 17h and 18th editions.

[9]Details of these initial studies appear in notes titled "*NEC-4* versus *NEC-2* with Stepped-Diameter Correction and Autosegmentation," available at my Web site, along with numerous other entries from my antenna notebooks. **http://funnelweb.utcc.utk.edu/~cebik/radio.html.**

[10]*NEC-4-I*, p 4.

[11]Although *MININEC* 3.13 is public domain and the calculating engine behind numerous implementations, commercial versions may show variations in output data due to the inclusion of different correction algorithms. For example, *AO*, by K6STI, contains a frequency correction to align *MININEC* results with those of *NEC-2*. *ELNEC*, by W7EL, contains a parallel-wire correction factor. *NEC4WIN*, to the best of my knowledge, contains no correctives. *MININEC Pro* is a new proprietary version of *MININEC* by Rockway and Logan that is said to overcome many limitations of public domain *MININEC*, but I have not yet calibrated the program against others. Sources of these programs were provided in *QEX* (Mar/Apr 1998, pp 47-49). For the exercises involving *MININEC*, *ELNEC*, with the parallel wire corrector in operation, was used throughout for ease of input file transfer to *EZNEC Pro* and from there to *GNEC*.

Wire Modeling Limitations of NEC and MININEC for Windows

Come listen to the authors of MININEC as they describe its workings and compare it with NEC-4.

By John Rockway and James Logan, N6BRF
EM Scientific, Inc

The History of *MININEC*

The original *MININEC* was written by John Rockway with a little prodding and support from Jim Logan. Over the years, the Rockway-Logan team has been responsible for the development of this code into one of the best known and most useful method-of-moments antenna modeling codes available. A number of other individuals have contributed small, but not necessarily insignificant, pieces to *MININEC*'s capability, but it has been

[1]Notes appear on page 21.

EM Scientific, Inc
2533 N. Carson St, Ste 2107
Carson City, NV 89706
e-mail **emsci@aol.com**
http://www.emsci.com

the dual efforts of the Rockway-Logan team that has made *MININEC* into the powerful antenna-design and analysis tool it is today.

Because of the similarity in names, it is often stated that *MININEC* is but a personal computer (PC) version of its big brother, *NEC*.[2] This could not be farther from the truth, however. There are significant differences between these two codes. Both codes use the method of moments to solve for currents on electrically thin wires. However, each code starts with a different version of the integral formulation for the currents and fields for wires. Then, each follows significantly different algorithms to implement the method of moments.

In 1980, when the first version of *MININEC* was written, PCs had not been on the market for very long. They were relatively expensive and very limited in capability. PCs were generally regarded as mere novelties or toys. PCs were typically limited to 16 kB of memory with an eight-bit word length. There was no *FORTRAN* for the PC, so *MININEC* was written in *BASIC*. *NEC* was (and still is) a very powerful computer code, with tens of thousands of *FORTRAN* statements, originally written for use on large mainframe computers. In those days, PCs could not support such a large program. The formulation had to be changed to allow a simpler implementation of the method-of-moments in order to produce a more compact code. It would not be possible to include many of the powerful modeling options provided by *NEC*. Following the

advice of Professor Don Wilton, at the University of Mississippi (now with the University of Houston), the first version of *MININEC* was written in 500 lines of *BASIC* that required 32 kB of memory. Nonetheless, this version proved surprisingly accurate for dipoles and monopoles.

The first public release of *MININEC* occurred in 1982.[3] The code was 550 lines of *BASIC* that would run on an APPLE II computer with 64 kB of memory. It could compute the current distribution, impedance and far-field pattern of an arbitrarily oriented set of wires in free space or over a perfectly conducting ground plane. Lumped-impedance loads were allowed at segment junctions except for segments intersecting with the ground plane. In addition, wires intersecting the ground plane were restricted to right angles. In interpreter *BASIC* (there were no *BASIC* compilers then) the problem size was limited to 10 wires and 50 currents (or 70 segments with junctions).

MININEC was an instant success. Almost immediately, a small user group developed and began to grow. In 1984, partly to meet the demand for *MININEC,* but also to share other computer algorithms, the authors teamed up with two colleagues: Peter Li and Dan Tam. They published a book that contained an improved version of *MININEC* along with some other useful algorithms.[6] *MININEC2,* as it became known, was not significantly different from its predecessor, but the limitation for wires intersecting the ground plane was removed. Wires could intersect the ground at any angle.

The power of PCs began to grow. Computers were getting faster, had more memory, and used math coprocessors. *BASIC* compilers also became available. These factors opened new vistas for *MININEC*. In 1986, the authors released *MININEC3*.[7] This code featured a new user interface that automatically determined wire connections from the user inputs for wire end coordinates. It could also read and interpret a limited *NEC* input data set. There was no way to save and edit geometry data, however. *MININEC3*

included near-fields, a Fresnel reflection-coefficient correction to the real-ground patterns and an expanded lumped-parameter loading option. *MININEC* had grown to just over 1600 lines of *BASIC*. With a math coprocessor and a *BASIC* compiler, *MININEC3* could solve antenna problems with up to 50 wires and 50 current unknowns.

The next *MININEC* effort by the authors produced the *MININEC* system in 1988.[9] This was a valiant effort by the authors to provide improved problem definition, the ability to save features and improve on-line graphics. The release of the *MININEC* system happened to coincide with the introduction of Microsoft *Windows,* which took the PC world by storm. The authors were too close to publication to backtrack and implement a *Windows* system. Nonetheless, there were many worthwhile innovations in this code. This was the first version of *MININEC* that *required* a compiler, a *BASIC* compiler. All previous versions could run in interpreter *BASIC*. The solution time and storage requirements for rota-
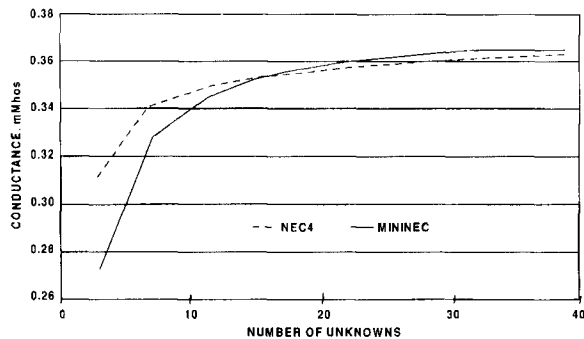

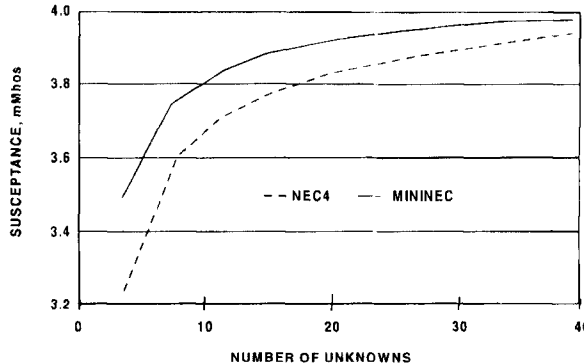Fig 1—Conductance versus unknowns for a short dipole.

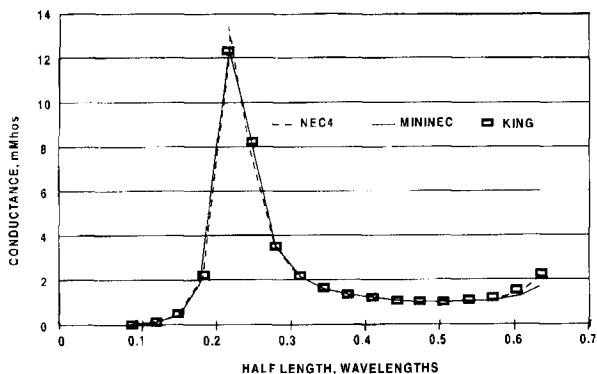Fig 2—Susceptance versus unknowns for a short dipole.

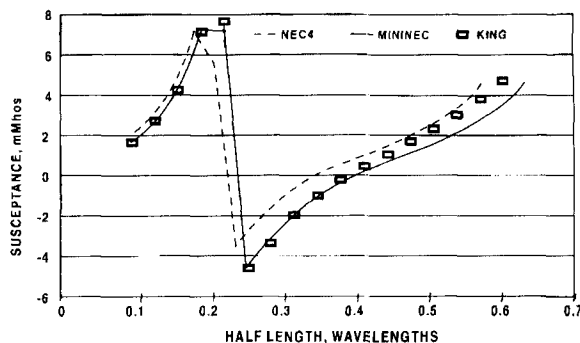Fig 3—Dipole conductance vs. frequency (29 unknowns).

Fig 4—Dipole susceptance vs. frequency (29 unknowns).

tionally symmetric antennas were greatly reduced. The transpose elimination algorithm was available as a user-selectable option to allow computation of larger problems—up to 50 wires and 90 current samples or 190 segments were permitted—without recompiling.

Many others have also attempted to improve on *MININEC*. Most notable are the innovative user interfaces and graphics displays offered by Roy Lewallen[5] in 1991 and Brian Beezley[1] in 1992.

In 1995, the authors published the first of a series of *MININEC for Windows* codes. These codes represented the development of a *new version* of *MININEC*. An improved solution of the potential-integral formulation for the currents resulted in a more-accurate solution for currents on wires. In addition, *FORTRAN* was used for the computationally intensive portions of *MININEC*. This led to an increase in speed over previous versions.

The first code was *MININEC Professional for Windows*.[10] Because it is a *Windows* application, text and graphical outputs are easily transferred to other *Windows* applications such as spreadsheets and word processors. Mouse support and printer drivers are also supplied by the *Windows* environment. The input is a node-based geometry. That is, nodes define points in space (in Cartesian, cylindrical or geographic coordinates), and wires are defined between nodes. Entries are made in tables through individualized window screens. On line, context-sensitive help is provided along with preprocessing diagnostics. *MININEC Professional* is dimensioned for 1000 wires and 2000 unknowns.

In 1996, the authors published *MININEC Broadcast Professional* for *Windows*,[11] which is similar to its predecessor, but more powerful. Additional features include an improved voltage-source model, a plane-wave-source model, automated convergence testing, design analysis post processing, array synthesis and ground-wave calculations. *MININEC Broadcast Professional* is dimensioned for 2000 wires and 4000 unknowns.

Also in 1996, the authors published *MININEC for Windows*,[12] a simplified version of *MININEC Professional* that is more suitable for first-time users and their pocketbooks. This code is dimensioned for 400 wires and 800 unknowns.

The Modeling Process

We do not attempt to present the *new MININEC for Windows* formulation and method-of-moments procedure in this paper. Interested readers should refer to the *MININEC* documentation or other suitable texts. For the uninitiated, however, we will describe a few realities of modeling.

The *MININEC* formulation defines the relationship between the currents and charges on wires and the associated electric and magnetic fields. A number of simplifying assumptions make the results more tractable for computer programming. For example, all conductors are straight cylinders with lengths much greater than their diameters (ie, the thin wire approxi-
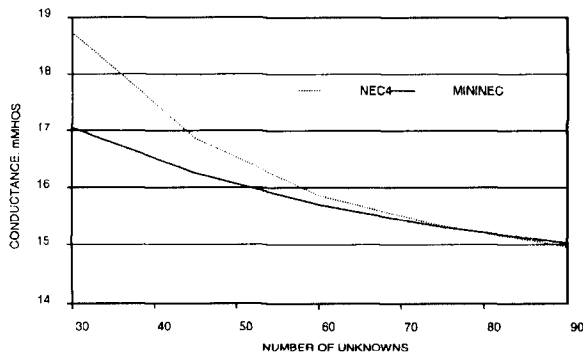

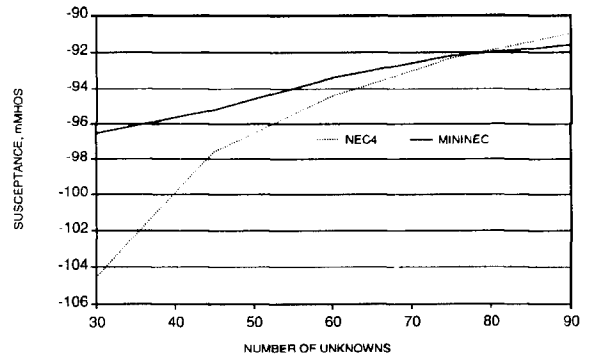Fig 5—KH = 0.2 TEE antenna conductance vs. unknowns.

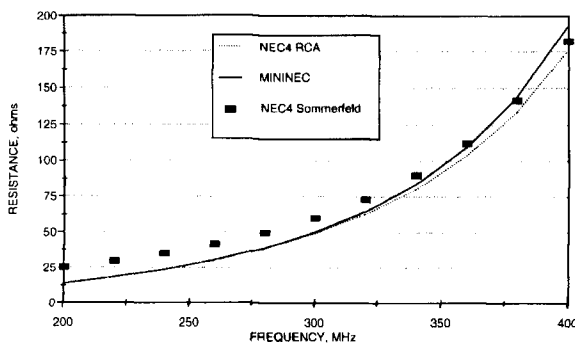Fig 6—KH = 0.2 TEE antenna susceptance vs. unknowns.

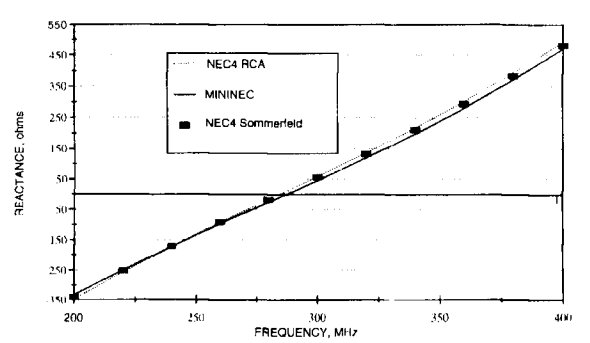Fig 7—Dipole resistance of a dipole over average ground.

Fig 8—Dipole reactance of a dipole over average ground.

mation). The method of moments is used to translate the formulation into a system of equations that can be readily solved on a computer. The method of moments introduces further constraints on the model and the solution. For example, the solution requires that each wire is divided into a number of short segments and that any sources and loads are collocated with the segment connections. The simplifying assumptions and method-of-moments constraints combined with user experience defines the validity range for the antenna model and solution.

The modeling process, highly automated by the *MININEC* user interface, is reduced to six principle steps: (1) Geometry Definition, (2) Electrical Definition, (3) Model Validation, (4) Solution Definition, (5) Model Execution, (6) Output Display and Solution Analysis. Each of these steps is discussed in the following paragraphs.

Step 1, Geometry Definition: The antenna geometry is defined by specifying a set of nodes in a three-dimensional imaginary grid or coordinate system in space. These nodes become the wire ends. The grid may be in free space or include a ground plane. Coordinate dimensions are selected (feet, inches, centimeters, meters, etc). Wires are defined by connecting nodes. A radius and number of segments are assigned to each wire. The number of segments required depends on the wire length in wavelengths. A current node will be associated with each segment on every wire. These nodes are discrete points on the wires where currents will be computed. In the new *MININEC* series of computer codes, a three-dimensional display of the geometry as well as a tabular listing is available.

Step 2, Electrical Definition: The electrical circuit parameters for the

antenna model are defined. The frequency (or set of frequencies) for the solution is specified. The number, location and strength of sources (eg, feed points) and loads are defined. Sources and loads must coincide with the location of current nodes defined in Step 1. The new *MININEC* codes provide a three-dimensional geometry display that indicates the location of sources and loads.

Step 3, Model Validation: The model is checked against the modeling constraints to determine that the geometry and electrical specifications are within the validity range. In the new *MININEC* codes, a three-dimensional geometry display is color coded to indicate warnings and errors. A diagnostic window provides a tabular display or warnings and errors.

Step 4, Solution Definition: The type of fields desired in the solution is specified. Types may be near-electric, near-magnetic or (far-field) radiation fields. The exact number and location for each to be computed must be defined.

Step 5, Model Execution: The currents are computed for the model defined by Steps 1 through 4.

Step 6, Output Display and Solution Analysis: The computed output is examined in tabular and graphical form to determine the validity of the results and eventual interpretation to the real world.

Modeling Accuracy

The accuracy of the results from using any numerical modeling code depends both on the user as well as on the code. The old adage "garbage in, garbage out" applies all too well. Given that a knowledgeable user has defined the antenna model within the modeling constraints, what kind of accuracy can be expected? How do the

latest versions in the *MININEC for Windows* series of codes compare to the latest version of *NEC*? The following examples give an indication of what to expect. The results of the *MININEC* for *Windows* series of codes are compared with version 4 of *NEC*.

Dipoles

Any evaluation of a wire-antenna modeling code begins with the evaluation of a dipole antenna. Figures 1 and 2 show a comparison of the new *MININEC* to *NEC* in a typical convergence test for a short dipole in free space. In a convergence test the accuracy is determined as a function of the number of unknowns. A convergence test provides the rationale for selection of segment density for a desired accuracy. It also demonstrates the stability of the analysis. The dipole half-length is 0.159155 meters and radius is 0.001588 meters. Figure 1 shows the conductance versus the number of unknowns and Figure 2 shows the susceptance versus the number of unknowns. R. W. P. King[4] reports an admittance (conductance + j susceptance) of $0.25 + j3.87$ mS. These figures show that as the number of segments (unknowns) increases, the admittance of both codes converges toward approximately the same asymptotic values.

Figures 3 and 4 show the admittance results of *MININEC* compared to *NEC* for a short dipole over a range of frequencies. Figures 3 and 4 compare conductance and susceptance to frequency. Also shown are values from R. W. P. King.[4] The dipole has a half length of 0.25 meters and a radius of 0.00351 meters. The length-to-radius ratio is 142. The segmentation scheme provides 29 unknowns over the frequency band. Both codes perform very well.
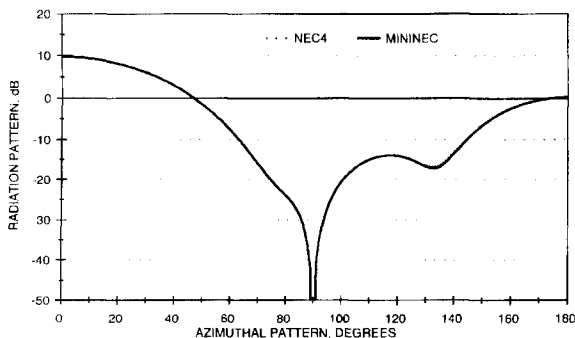

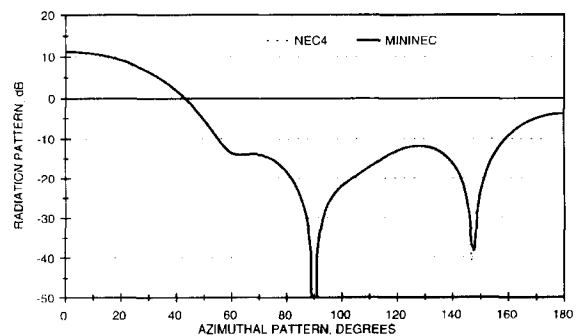Fig 9—Radiation pattern of a three element Yagi.

Fig 10—Radiation pattern of a five element Yagi.

Multiple Wire Antennas

A second step in an evaluation is calculation of a multiple-wire antenna. Figures 5 and 6 display *MININEC* and *NEC* conductance and admittance calculations for a **T** antenna described by King.[8] The specific antenna has KH = 0.2, where K = $2\pi/\lambda$, λ is the wavelength and H is the height, in wavelengths. King reports an admittance for this antenna of $29.6 - j102.6$ mS.

Real Ground

Figures 7 and 8 show the impedance computed by *MININEC* compared to the impedance computed by *NEC* for a dipole over real ground. The antenna is a 0.5 meter center-fed dipole with 0.0005 meter radius at 0.1 meters above an average ground. The dielectric constant of the ground is 15, and the conductivity is 27.8 mS. Figure 7 shows the dipole resistance versus frequency and Figure 8 the dipole reactance versus frequency. Two sets of calculations are shown for *NEC*. One set of calculations is for the solution using the Fresnel reflection-coefficient method, and the other is for the Sommerfeld solution. These results have been obtained after checking the solutions of both codes for convergence. The Fresnel reflection-coefficient method is shown a good approximation to the more general, exact Sommerfeld solution.

Radiation Pattern

Figures 9 and 10 show the Yagi radiation patterns computed by *MININEC* compared to the radiation patterns of *NEC*. Figure 9 is the pattern for a three-element Yagi, and Figure 10 is the pattern for a five-element Yagi. The Yagi dimensions in meters are shown in Table 1.

The selected examples show that *MININEC* gives comparable results to *NEC*. This is not a complete picture of the comparison of these codes, but it gives the reader a glimpse of the results to be expected. A more thorough analysis of *MININEC* is presented in References 10 and 11. (Also see L. B. Cebik's article in this issue.—*Ed.*) This analysis shows that for a wide variety of problems *MININEC* and *NEC* provide comparable results.

Table 1—Dimensions of three and five-element Yagi antennas

Yagi	3 element	5 element
Reflector length	0.482	0.482
Driven element length	0.25	0.25
1st director length	0.442	0.428
2nd director length		0.424
3rd director length		0.428
Director spacing	0.2	0.2

References

[1]Beezley, B., *The MN4 Manual*, Vista, California, 1992.

[2]Burke, G. J. and A. J. Poggio, *Numerical Electromagnetics Code (NEC) —Method of Moments*, Naval Ocean Systems Center Technical Document 116, January 1981.

[3]Julian, A. J., J. C. Logan, J. W. Rockway, MININEC: *A Mini-Numerical Electromagnetics Code*, NOSC Technical Document 516, September 1982.

[4]King, R. W. P., *Tables of Antenna Characteristics*, IFI/Plenum Data Corporation, New York, 1971.

[5]Lewallen, R., "*MININEC*: The Other Edge of the Sword," *QST*, Feb 1991.

[6]Li, S. T., J. C. Logan, J. W. Rockway, D. W. Tam, *Microcomputer Tools for Communications Engineering*, Artech House, Inc, Dedham, Massachusetts, 1984.

[7]Logan, J. C. and J. W. Rockway, *The New MININEC (Version 3): A Mini-Numerical Electromagnetic Code*, NOSC Technical Document 938, September 1986.

[8]Prasad, S. and R. W. P. King, "Experimental Study of L-, T- and Related Transmission Line Antennas," *Journal of Research of the National Bureau of Standards—D Radio Propagation*, Vol 65D, No. 5, Sept-Oct 1961.

[9]Rockway, J. W., J. C. Logan, D. W. Tam, and S. T. Li, *The MININEC System: Microcomputer Analysis of Wire Antennas*, Artech House, Inc. Dedham, Massachusetts, 1988.

[10]Rockway, J. W. and J. C. Logan, MININEC *Professional for Windows*, EM Scientific Inc, Carson City, Nevada, 1995.

[11]Rockway, J. W. and J. C. Logan, MININEC *Broadcast Professional for Windows*, EM Scientific, Carson City, Nevada, 1996.

[12]Rockway, J. W. and J. C. Logan, MININEC *for Windows*, EM Scientific, Carson City, Nevada, 1996. □□

Signals, Samples, and Stuff: A DSP Tutorial (Part 2)

As we continue our exploration of DSP techniques, let's look inside an IF-DSP transceiver.

By Doug Smith, KF6DX/7

In Part 1 of this series (*QEX*, Mar/Apr 1998, pp 3-16) we learned about fundamental DSP techniques and algorithms for use in modern transceivers. Here we'll explore an actual IF-DSP transceiver design. Many of the issues relate well to conventional analog equipment, but special emphasis is placed on the unique requirements, advantages and trade-offs in a digital radio.

While the performance of a digital transceiver may exceed that of a traditional design, the basic goals are the same. We wish to fabricate a receiver with good sensitivity and selectivity, the maximum dynamic range and minimum distortion. The transmitter must produce a low-distortion, spectrally pure signal. The frequency stability and tuning resolution should not impose undue limitations on operation.

In this age when DSP hardware

[1]Notes appear on page 37.

PO Box 4074
Sedona, AZ 86340
e-mail **dsmith@sedona.net**

capabilities are finally exploiting advances in theory discovered over the last 40 years, we can indeed set our sights quite high! As we begin, let's define the challenges facing us in receivers, so as to assess the influence of IF-DSP technology.

A Receiver: All that Gain and a Whole Lot More

Superheterodyne receivers have been around for awhile, and until DSP hardware can achieve sufficient speed and dynamic range to digitize signals straight from the antenna, we'll all continue to use them. The main advantage of a superhet is that signals are converted to a fixed IF that provides most of the gain and selectivity. To avoid spurious responses, multiple frequency conversions are common. We ought to recognize, however, that minimizing the number of conversions also diminishes the number of oscillators and, therefore, the number of possible internal signals, or "birdies."

At some stage, we'll digitize some signals and perform filtering and

other signal processing. We want this point to be as close to the antenna as possible, so we must look at the frequencies, bandwidths and dynamic ranges available in DSP components before choosing an IF. We can eliminate many traditional analog signal-processing stages if we digitize signals ahead of the point that expensive crystal or mechanical filters previously occupied. Our first trade-off is between high-speed analog-to-digital converters (ADCs) and the costly filters they would replace. This decision is driven mainly by cost, although issues of current consumption and processing power definitely come into play.

The final compromise also depends on the performance levels we expect to achieve. For example, many excellent ADCs are quite capable of digitizing signals directly from the antenna: their sampling rates are fast enough for the job, but their dynamic ranges are narrow. As we'll see below, HF receivers must handle a tremendous range of input signal levels without

flinching! So before we can make even this first decision about the receiver's conversion scheme, we must think about dynamic range: What is it, and how much do we want?

Receiver Dynamic Range (DR)

It's every receiver's job to produce a useful replica of the transmitted information and reject all other signals. In today's crowded HF bands, this is an increasingly difficult task! The desired signal might be quite weak, so we need good sensitivity and lots of gain without introducing excess circuit noise.

Sensitivity must be specified as a function of the bandwidth of interest, because we're trying to copy a narrow-bandwidth signal in the presence of noise, which exists at every frequency! In the specified bandwidth, a signal received at the antenna terminals has a certain signal-to-noise ratio (SNR). We fight to preserve this SNR throughout the receiver. Electronic circuits introduce some noise, however. The ratio of the input SNR to the output SNR of a receiver is referred to as its *noise figure* (NF), and is expressed in decibels.

Originally explained by Einstein in 1905, Brownian motion of atoms and free electrons in any conductor produces an available noise power[1,2] (in watts) of:

$$P_{noise} = kTB \qquad \text{(Eq 1)}$$

Where k is Boltzmann's constant, 1.38×10^{-23}, T is the absolute temperature in Kelvins, and B is the bandwidth, in hertz. "Plug and chug" on these numbers, and you'll find that at room temperature (293 K) and in a bandwidth of 3 kHz, this power is −139 dBm or 12.1 *attowatts!* This quantity (12.1 × 10^−18 W) represents the minimum discernible signal (MDS) in a perfect receiver using a typical voice bandwidth. Note that as the temperature decreases, the possibilities increase linearly; a receiver operating in a liquid-nitrogen bath is a real gem! Atmospheric and cosmic noise are usually much greater than this theoretical limit, however.

The best HF receivers today have NFs around 7 dB. When noise power equals signal power, the output SNR is 0 dB, and the input signal level is:

$$P_{in_{0dBSNR}} = -139 + 7 \, \text{dBm} = -132 \, \text{dBm}$$
$$\text{(Eq 2)}$$

We define this MDS level as the lower limit of the receiver's dynamic range. (In a 50 Ω resistor this power corresponds to 0.056 μV.—*Ed.*) It's not so easy to find the upper limit of the dynamic range. Because of the mani-

fold ways receivers degrade at high input levels, we'll define several dynamic ranges, one based on each of these.

Receiver Overload: Let Me Count the Ways

Normally, overload phenomena involve large, off-channel signals. Of course, it's also possible to overload on a very strong *desired* signal. For most modern receivers, this level would be so high that radio communication wouldn't be necessary; you could just shout out the window!

Large-signal performance is typically characterized[5] by measuring the following effects:
- Third-order intermodulation distortion (IMD)
- Second-order IMD
- "Blocking," or desensitization
- In-band IMD

Let's examine methods for each of these measurements along with the strengths and weaknesses of current methods.

IMD Dynamic Range and Intercept Point

To measure IMD dynamic range, we inject two off-channel signals of equal amplitude and measure the degradation of receiver performance. Degradation comes in the form of an undesired, on-channel signal produced by the mixing of the off-channel signals. We increase the off-channel signal levels until the on-channel signal power equals the noise power. This is the definition of MDS given above.

We define the IMD dynamic range to be the ratio of this off-channel signal power to the MDS power, expressed in decibels. In the ARRL method for third-order IMD, one interfering signal is placed 20 kHz from the center channel and another 40 kHz from center. The third-order intercept point (IP₃) is calculated by assuming the receiver distortion obeys a perfect cube law: For every decibel increase in interference, the third-order IMD product will increase 3 dB, and the difference will increase by 2 dB. IP₃ is extrapolated, therefore, by adding half of the third-order IMD dynamic range to the interference level obtained in the measurement above:

$$IP_3 = \left(\frac{IMD \, D \, R}{2} \right) + P_{QRM} \qquad \text{(Eq 3)}$$

This is supposed to be the level where the third-order IMD product is equal in amplitude to the interference.

Were we to actually inject interference of this level, however, we might find a *real* IP₃ much higher; receivers seldom obey perfect cube laws as they're predicted to do! This normalized procedure is a good basis for comparison, though.

In the second-order test, we inject two non-harmonically related signals and look for the undesired product at the sum or difference of the frequencies. IMD dynamic range is measured as above, and IP₂ is extrapolated by assuming the receiver obeys a perfect *square law.* For every decibel of increase in the interference, the second-order product increases 2 dB, and the difference increases by 1 dB:

$$IP_2 = (IMD \, D \, R) + P_{QRM} \qquad \text{(Eq 4)}$$

How can the receiver obey two apparently conflicting laws at the same time?! In the second-order case, we're mixing the two *fundamentals* of the interference; whereas, in the third-order case, we're mixing the fundamental of one with the internally generated second harmonic of the other.

Note that when we add two fundamental signals, the result is always greater than twice the frequency of one of the signals. For this reason, our second-order performance can be improved by using *half-octave* band-pass filters ahead of the receiver front end. Such filters—when switched or tuned as the receiver changes frequency—always attenuate one of the interfering signals, reducing the deleterious effects.

"Blocking" Dynamic Range

In this measurement, we inject a *single* off-channel source, and look for some degradation in the on-channel performance. In the ARRL method, the output power from a single on-channel input signal is monitored. The interference, 20 kHz away, is increased until the desired output power either increases or decreases by 1 dB.

A decrease is supposed to indicate that some stage or other is saturating, while an increase results in a "noise-limited" measurement. The *blocking dynamic range* (BDR) is calculated as the ratio of the interference power in the measurement above to the MDS power, expressed in decibels.

In reality, saturation seldom occurs in modern receivers before the noise takes over. This noise is the result of *reciprocal mixing,* wherein the interference mixes with the phase-noise sidebands of the LO to produce in-band noise. A state-of-the-art synthesized

LO has phase noise in a 3 kHz BW, and at a 20 kHz offset, of around 100 dB below its injection level. Were the BDR measured using the SNR instead of the average output power, we could call it desensitization or *"desense."* It would be on the order of 100 dB, and would be solely a measure of the synthesizer phase noise. This number is quite a bit lower than that usually obtained with the ARRL method.

The difference becomes evident when trying to measure an IF-DSP receiver with a *digital AGC* system. Such a system holds the peak desired output level constant, and as the SNR degrades, the average output power *decreases!* In a conventional receiver (all other things being equal) the SNR would be identical, but the output power would *increase* because of the added noise. The peak-to-average ratio of noise is high, so monitoring the average or RMS output power wouldn't indicate an increase until much more interference power were added.

To correlate the SNR method with the ARRL method, we might consider using degradation of the output SNR as our criterion, as in the EIA standard. The degradation level could be chosen to equate the new measurements to existing BDR measurements of known receivers. Let's face it, reciprocal mixing gives the most trouble these days. If a blocking measurement is still desired, we ought to use the *peak* output level, not the RMS.

In-band IMD

This is a measure of distortion produced by a receiver when the only signals present are inside the desired passband. Current ARRL methods call for a two-tone input with a frequency separation of 100 Hz. This is excellent. It's roughly the natural impulse frequency of the human voice system. The IMD product levels are examined relative to one of the tones. The AGC speed, if adjustable, is set to its fastest setting.

Digital AGC systems can cause problems here, because they are capable of *very* fast attack and decay times. If the decay time is set fast enough, clearly the two-tone will be subject to extreme distortion; it'll begin "flat-topping." It doesn't make sense to defeat the very system designed to prevent the thing being measured!

Preamplifiers and Dynamic Range

It's obvious that to achieve the best sensitivity, some gain ahead of the first mixer is required. If this gain stage has a low NF, we can improve

the sensitivity by almost the amount of the gain. This extends the receiver's dynamic range on the low end.

It's difficult to make up the difference on the high end, though. The large-signal handling will degrade by *at least* the amount of the preamplifier gain, so the dynamic range is generally reduced. Also, notice that dynamic range is just the ratio of maximum and minimum signals that can be handled, and it says almost nothing about actual large-signal handling capability! One receiver might have a greater dynamic range than another, and still have a poor IP. Its sensitivity may be excellent, but it might not be a good large-signal performer.

AGC

We've seen that HF receivers must handle very weak signals (–132 dBm) and strong signals that may approach +20 dBm, near the IP_3. Expressed this way, the dynamic range can exceed 150 dB! As we expect the output level to remain relatively constant and the distortion to stay within limits, a gain-control system is necessary. We must keep analog stages linear, so an analog AGC system is mandatory.

We intend to provide the final selectivity in our receiver using digital filters, as this eliminates the need for expensive crystal or mechanical filters. So it follows that some of the signals we digitize will be undesired —this raises a problem: The digital filters will remove the interference, but the analog AGC will still act on the total bandwidth! A strong interfering signal will reduce the analog gain, as it must, and the level of our desired signal will fall as well. This is where the digital AGC system comes in.

Digital AGC Algorithms

We decide that to keep the desired signal's output level constant, we need a system that measures the ratio of total digitized signal energy to desired signal energy.[1,4] When the interference increases, this system will compensate for the reduction in gain caused by the analog AGC. The effect will be to hold the desired signal's peak level constant. Now we must determine how we're going to measure the critical ratio, and how and when to make adjustments in the gain compensation.

Clearly, the digital gain-compensation algorithm must use two data as inputs: the ratio of total signal level to desired signal level, and the actuation or amount of analog AGC. The ratio of

the amplitudes is easily calculated by the DSP system; it need only compare the peak digitized input level with the peak output level after filtering. This isn't quite the whole solution, however, because when the desired signal decreases, the system can't tell if it was because of interference-caused analog gain reduction, or because the other station just stopped transmitting!

So we arrange to monitor the analog AGC voltage in order to find out what it is doing. It turns out we don't need to know the amount of analog gain reduction if we can adjust the digital gain fast enough. We'll examine the analog AGC to detect when the gain is decreasing rapidly and when the amplitude ratio is increasing rapidly, then quickly boost the digital gain until the desired output level is maintained. Notice that both the analog and digital AGC systems maintain a fast-attack characteristic in all situations.

In practice, this system works quite well; the digital AGC decay time can be continuously adjusted as desired. On-channel signals are digitally boosted by the amount necessary to keep the peak output constant. The main drawback is that the dynamic range of the ADC system limits the available digital gain.

ADC Limitations

We learned in Part 1 that the dynamic ranges of ADCs are limited by the bit-resolution, speed and input frequencies of the devices. Looking at current technologies, we see that 16-bit ADCs are available with 96 dB of dynamic range; the input frequency ranges of these are confined to just above audio, however. A low-frequency IF poses a problem only in image rejection because it's difficult to filter out responses close to the desired signal. The trade-off here is between the digital dynamic range at hand for gain boosting and the frequency of our last IF.

Oversampling ADCs can be advantageous in our design, because they spread quantization noise over large bandwidths, then apply digital filters of their own to eliminate most of the noise. Sampling rates and IFs can be chosen so that the ADC filter aids the selectivity of the receiver. We discover that 40 kHz is a good last IF for this reason. The Analog Devices AD7722 is chosen as our ADC.[15]

We further decide that we can apply up to 60 dB of digital gain boost using this device, since its dynamic range is 96 dB; the output SNR cannot degrade to much less than:

$$SNR_{out_{min}} \approx 96 - 60\,\text{dB} = 36\,\text{dB} \qquad (\text{Eq 5})$$

because of the ADC system. We'll adjust the analog AGC so that it provides a peak input level near the maximum input allowable for the device. Notice that exceeding the maximum input level of the ADC results in instant, catastrophic degeneration of the output signal. ADC overload is the one thing we can't tolerate. So, we allow a few decibels of *headroom* in setting the analog AGC operating point.

The First IF

Now that we've decided on a *last* IF, it's time to figure out how we get there from our RF range of up to 30 MHz. Selection of a first IF depends on the location of spurious responses, as well as availability and cost of components. Spurious and image problems are greatly reduced if the first IF is above the highest RF. *Up-conversion* has been standard in HF receivers for some time now.

Several popular IFs offer cost advantages because of commonality with other radio services.[1,2] These include, but are not limited to:

- 45 MHz—popular in cellular phones
- 70 MHz—the standard UHF and microwave IF
- 75 MHz—aviation service marker beacon

An IF above twice the highest RF is favorable because it eliminates second-harmonic spurious responses. Eg, a strong signal appearing at 22.5 MHz

generates a second harmonic at 45 MHz. An IF above 60 MHz is above that harmonic and as a high-side-injection first LO will cover less than half an octave, which simplifies its design.

We can't increase our IF arbitrarily, though; the cost of crystal filters becomes prohibitive above 75 MHz, and the loss and instability of surface-acoustic-wave (SAW) and other VHF filter technologies make them unattractive. In addition, the synthesizer would be forced to higher frequencies, which would increase phase noise and associated reciprocal-mixing problems.

If we want the receiver to cover the range below 500 kHz, we must limit the phase noise. This energy enters our IF directly because of imbalance in the first mixer as the LO approaches the IF. We decide that 75 MHz is a good choice, and now we're ready to draw a block diagram of the receiver.

The Block Diagram

See Fig 1. Beginning at the antenna, we place a bank of half-octave bandpass filters (to extend the second-order IMD dynamic range as described above). These are switched by PIN diodes or relays. Next, a preamplifier with a gain of 15 dB and noise figure of 3.6 dB is used. This amplifier can be switched out using relays, or a 20-dB fixed attenuator can be inserted. Then, we select a high-level, double-balanced mixer to translate our signals to the first IF

This first mixer is critical, because

it's likely to determine our IMD dynamic range. We can expect an IP_3 several decibels above the injection level, so we choose a mixer designed for a +17 dBm LO. We anticipate that the extra energy will warrant careful shielding, filtering and isolation to prevent birdies and LO leakage to the antenna.

While filters at the mixer ports mitigate these problems, they can have an unexpected consequence: degraded IMD performance. A mixer is an inherently nonlinear device that generates harmonics of all signals entering it, like it or not! So, although our LO might be spectrally pure, its harmonics and those of the RF and IF will be present. When a mixer port is terminated at its characteristic impedance, this harmonic energy is absorbed by the termination. A filter, however, reflects some of it back into the mixer, where it may add to IMD.

One elegant solution involves "idler" filters. These networks provide a broadband termination impedance *and* the filtering we need. Fig 2 shows the use of idler filters. The filter passing the desired band (eg, a low-pass) is designed to be singly terminated.[6] A similarly designed high-pass filter is parallel connected, and terminated in the characteristic impedance—usually 50 Ω. The mixer then sees a relatively constant broadband load. A trade-off exists between the complexity and cost of these networks and the IMD degradation we'd otherwise suffer.
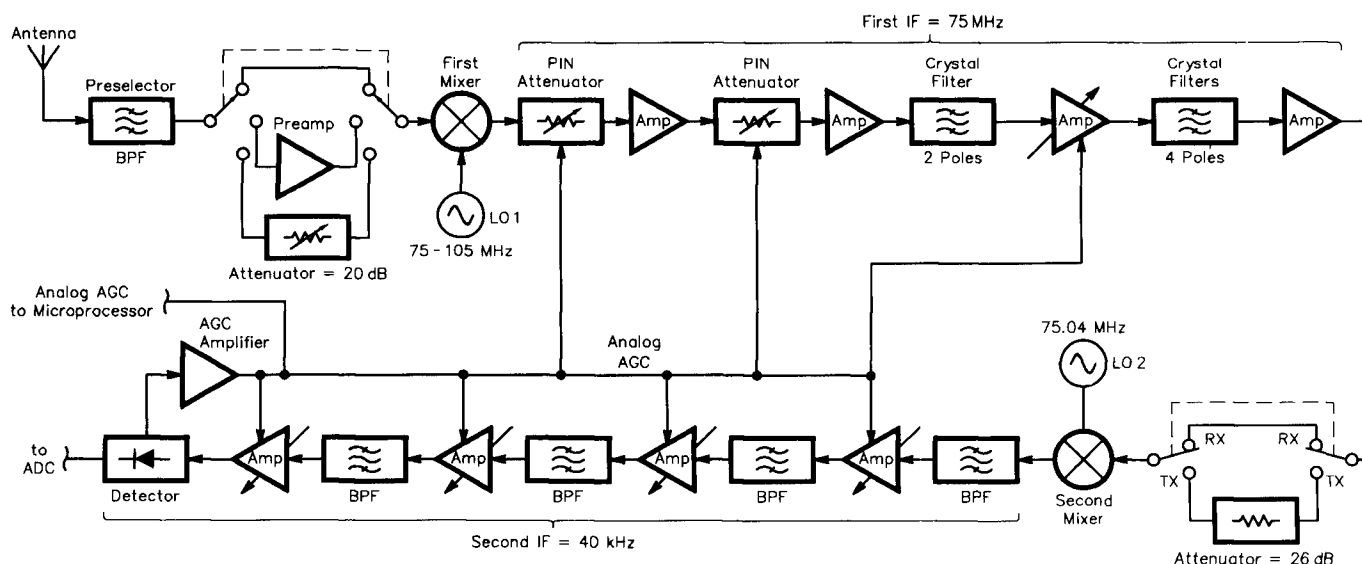

Fig 1—IF-DSP receiver block diagram.

High-Side versus Low-Side

We can elect to put the LO above or below the IF and take either the sum or difference frequency as our 75 MHz signal. High-side injection (taking the difference frequency) is attractive because the required LO range of 75 to 105 MHz is less than half an octave. With low-side injection, the LO range would be 45 to 75 MHz, nearly an octave; this would generate LO harmonics within an octave of the IF. We therefore choose high-side injection.

We'll build ourselves a first LO that covers the range in less than 1 Hz steps by combining direct digital synthesis (DDS) technology with a PLL. The second LO will have a fixed-frequency. We achieve drift cancellation by using difference mixing in both stages. I'll describe the synthesizer design in a future segment.

The First IF Strip

Following the first mixer we must have some gain to compensate for losses in the crystal filters ahead. This stage must have a low noise figure and moderate gain, yet must handle some very large signals to avoid degrading the IMD performance. It also ought to provide a good broadband termination at both its input and output—a pretty tall order so far!

Transistor circuits have been described in the literature,[1,2] which use a combination of voltage and current feedback to achieve simultaneous noise and impedance matching. In combination with FETs of moderate power rating, the performance can be impressive. Fig 3 gives an example. In this circuit, several parallel JFETs increase the current capacity while keeping cost under control. The feedback is taken from a transformer with a turns ratio designed to provide constant 50 Ω input and output impedances. The amount of feedback sets the gain at 6 dB. The measured input impedance is plotted as Fig 4 on a Smith chart normalized to 50 Ω.

The noise figure is 1.5 dB, and output IP_3 is +35 dBm! Clearly, this amplifier fills the bill. It won't affect our receiver's IP_3. Two of these stages are used before any narrow-band filters.

Gain control of transistor stages by varying the bias is impractical because of linearity problems, so PIN diode attenuators are commonly used. Fig 5 shows a typical attenuator circuit. We intersperse two of these attenuators to provide sufficient control range. This circuit must also provide constant impedance, so it uses a combination of series and shunt diodes. Bias voltage and current levels are set so that as the series diode is switched off, the shunt diodes switch on. The bias must be strong enough to avoid excessive IMD. The control range of each attenuator is 48 dB, and the IP_3 is +38 dBm.

Three two-pole monolithic crystal filters provide the "roofing" in our design, with a single MOSFET stage providing 20 dB of gain in their midst. It's advantageous to distribute component gain and loss throughout the receiver with the best noise figure and AGC effectiveness in mind. We must still maintain good SNR when gain control is applied, so concentrating control in one place is unwise.

After the roofing filters, an emitter-follower feeds the second mixer through another PIN diode attenuator. Since we intend to use the first IF in the transmit (TX) mode, this attenuator will provide precise gain reduction to restrict mixer-generated spurious signals while transmitting. It is set for zero attenuation while receiving (RX).

Fixed injection at a frequency of 75.040 MHz translates signals directly to the 40 kHz second IF. Notice that the second mixer can be a lower-level device, because AGC and filtering limit the amplitudes seen at this stage.

The Second IF Strip

Now that we're down to a frequency our ADC can handle, we have to provide enough gain so that the IF output level is near the maximum ADC input. This will allow the greatest dynamic range for digital AGC operation as described above. Getting gain at 40 kHz isn't a problem, but we also want additional analog AGC range here. Finally, we have to further bandwidth-limit the output so that *aliasing* cannot occur.

As described in Part 1, aliasing results when input bandwidth exceeds half the sampling frequency. Once incurred, nothing can be done to alleviate it! We want the sampling frequency as low as possible to minimize
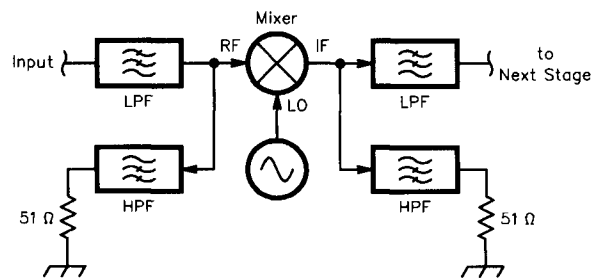

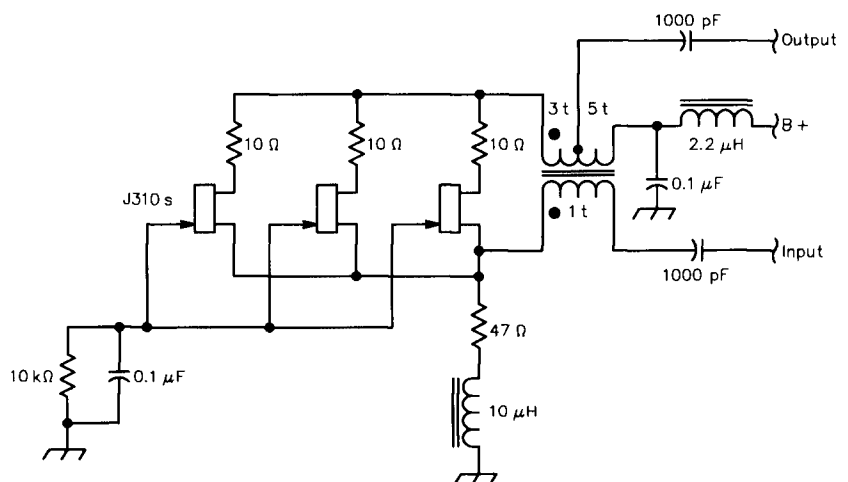**Fig 2—Double-balanced mixer employing 'idler' filters.**

Fig 3—75 MHz IF amplifier stage.

DSP "horsepower" requirements. Narrow-band FM, using bandwidth of 15 kHz, dictates the minimum passband width. This means our sampling frequency must be at least 30 kHz. To ease the design of the filters, we set the sampling frequency somewhat higher than this.

Narrow-band dual-gate MOSFET amplifier stages like that shown in Fig 6 are used in the second IF. These provide uncomplicated gain control and allow us to implement LC filtering in the 15 kHz BW. Since the ADC's maximum input level is 3 V_{P-P}, and its input impedance is high, more than 90 dB of voltage gain is required! Four stages of filtering and amplification get us what we need.

Analog AGC

The analog AGC system is implemented using the traditional detector and amplifier scheme. Gain reduction is implemented first in the later stages, then in the front-end components, so that the SNR can continue to increase with input signal levels. AGC voltage is fed to the DSP so that digital AGC can keep the final output level constant. It's fascinating to watch the second IF output level gyrating with interference while copying a weak signal on an adjacent frequency!

Summary

We've seen that in an IF-DSP transceiver, we had to start the design at the "back end" because of the limitations of available DSP components. This led us to certain decisions about the second IF, but we acknowledge that the rest of the design still resembles that of a conventional transceiver. In the next segment, we'll ex-
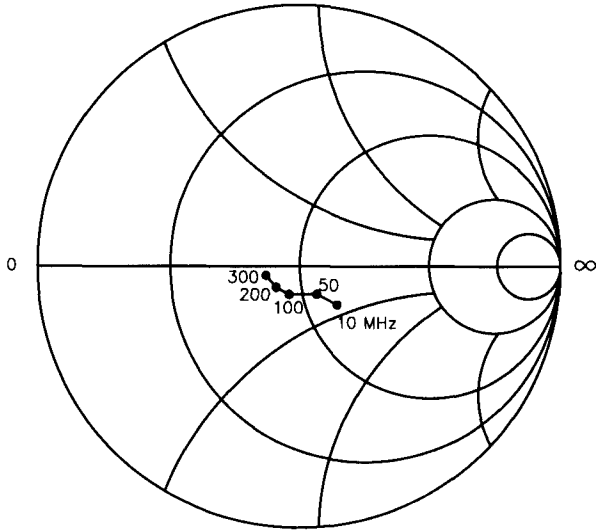

Fig 4—Input impedance of amplifier stage *versus* frequency.
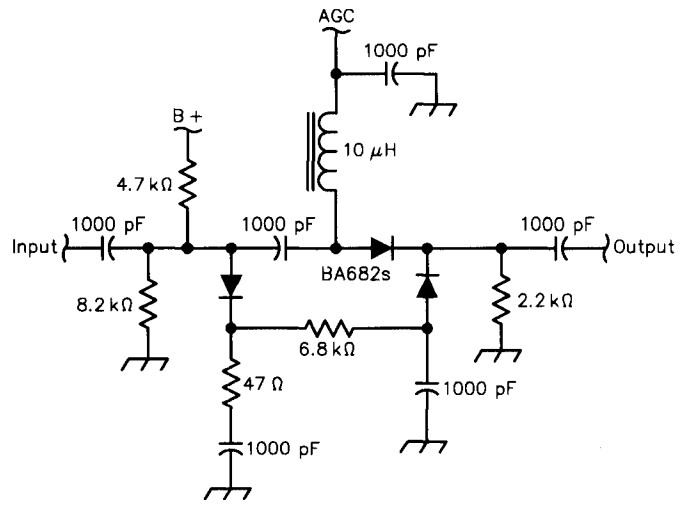

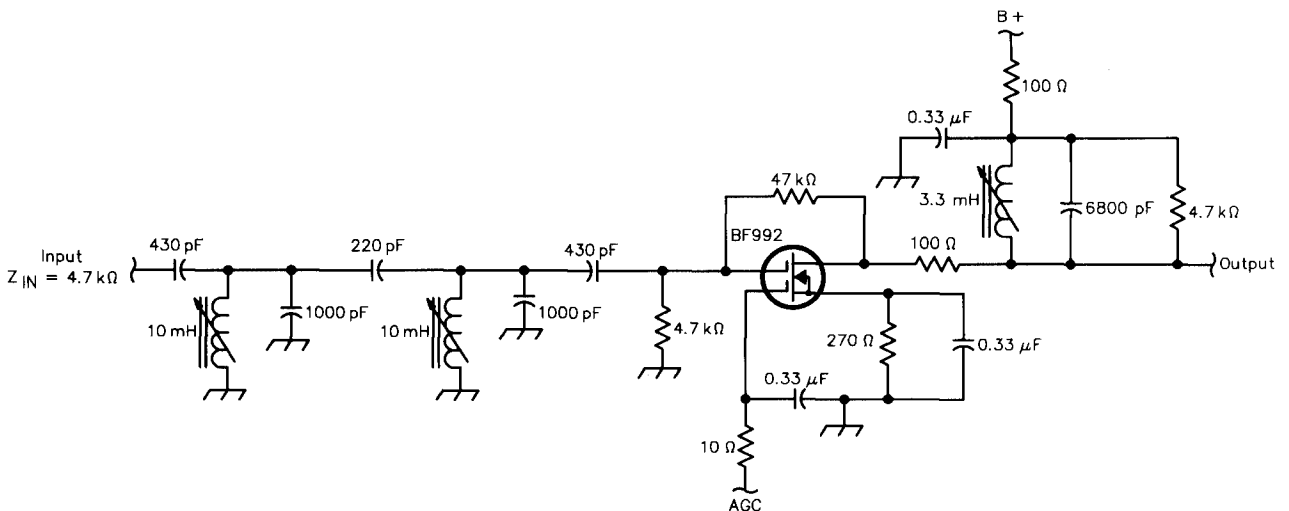
Fig 5—PIN diode attenuator circuit used in 75 MHz IF

Fig 6—40 kHz IF BPF and amplifier stage.

plore a state-of-the-art synthesizer design to see how DDS/DSP techniques help us get around the bands with the greatest of ease!

The Synthesizer: Excursions in the Frequency Domain

Synthesizers have come a long way since first becoming popular in HF transceivers in the '70s. Availability of components then lagged well behind the development of theory. Now, hardware capabilities have nearly caught up—which is the case for DSP in general—and are driving the very rapid advancement of HF equipment we are now experiencing. Although the gap has narrowed, we're still far from building an entirely digital, direct-conversion transceiver.

Paralleling breakthroughs in the microprocessor and data acquisition fields, progress in direct digital synthesis (DDS) has enabled performance levels only dreamed of a decade ago. Virtually all new designs profit from this technology.

Design Goals

In the previous segment, we defined the frequency ranges to be covered by our synthesizer design, set limits for phase noise and spectral purity and established the output levels. Now let's consider three other critical requirements: frequency stability, lock time and tuning resolution.

Amateurs are free to operate anywhere within large frequency bands, so it might seem that frequency accuracy isn't very critical. Nonetheless, prevalent narrow-bandwidth communication modes require good frequency stability, and operators have come to expect excellent stability from their rigs. It is reasonable to expect ±20 Hz stability over an anticipated temperature range of –10 to +50°C. We'll address the issues of long and short-term stability below.

We wish to attain a tuning speed that doesn't impose limitations on typical use. "Cross-band," or split-frequency operation ought to be considered. We set an upper limit of 20 ms for a frequency shift of ±600 kHz. Lock time is defined as the time required to settle within the above accuracy limits.

The smallest frequency steps should be such that they don't impede performance. A minimum step of 10 Hz used to be good enough, but now certain digital modes benefit from smaller steps. In addition, we'll discover that the narrow digital notch filter described in Part 1 of this series requires tuning within several hertz to achieve the best

null! We therefore set 1 Hz steps as our design goal. Table 1 summarizes the work for this major subsystem.

DDS Meets PLL, Object: Matrimony

DDS synthesizers achieve the fastest lock times and—when using a crystal-derived clock—the least phase noise of available methods.[4] It's also common knowledge that their spurious outputs can be excessive, especially as the output frequency approaches one half of the clock frequency. To build a DDS covering 75 to 105 MHz, directly, would be quite a feat!

Reasonably priced devices currently use clocks to about 60 MHz, which limits outputs to well under 30 MHz. Although it's possible to multiply or mix the DDS output up to the proper range, these strategies quickly become complicated and suffer from spurious problems. As a reference input to a standard PLL, however, a DDS can provide the performance we want.

LO₁ Block Diagram

In this design, we operate a VCO at 75 to 105 MHz and phase lock it to a DDS output at $F_{REF} = F_{VCO} / 100$, as

shown in Fig 7. F_{REF} is high to get fast lock times, and a DDS clock frequency much higher than F_{REF} reduces spurious content. The ratio of F_{VCO} / F_{REF} is important, because any noise and spurs in F_{REF} are multiplied by the PLL, *within the PLL bandwidth*, by the factor:

$$N = 20 \log \left(\frac{F_{VCO}}{F_{REF}} \right) = 40\,dB \qquad (Eq\ 6)$$

Our final output-spurious level is set by the DDS output purity.

DDS Spectral Purity

A DDS is a system that generates digital samples of a sine wave and converts them to an analog signal using a digital-to-analog converter (DAC; see Fig 8). Spurious outputs from such systems are caused by phase and amplitude inaccuracies in the digital and analog circuitry. First, let's look at the errors of the digital portion. These are analogous to the effects of digital word size and truncation of numerical results explained in Part 1.

In a DDS chip, a phase counter increments at each clock pulse, and the phase information is used to look up a

---

Table 1—Frequency Generation Requirements

LO₁

Frequency Range	75 to 105 MHz in 1 Hz steps
Output Level	+17 dBm ± 2 dBm
Lock Time (large step)	≤ 20 ms
Phase Noise	≤ –132 dBc/Hz @ F_c ±20 kHz

LO₂

Frequency	75.040 MHz
Output Level	+17 dBm ±2 dBm
Phase Noise	≤ –145 dBc/Hz @ F_c ±20 kHz

General

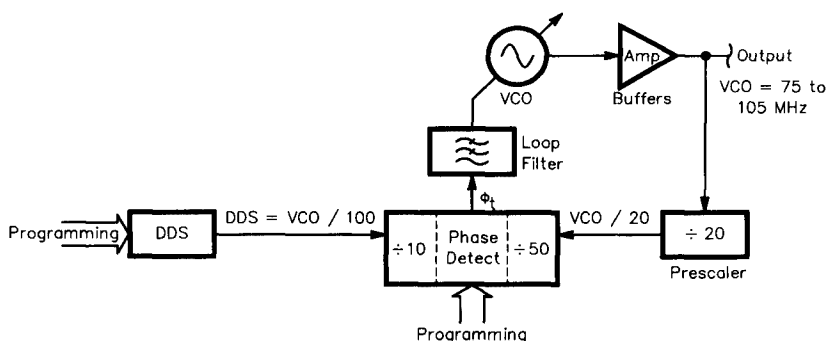|	
Frequency Stability	≤ 20 Hz over –10 to +50°C
	±40 Hz/year (aging)
Harmonics, spurious	≤ –73 dBc

Fig 7—LO₁ block diagram.

sine-wave amplitude from a table. Since the phase is represented by a binary number, with a fixed number of bits, p, errors can develop because resolution beyond p bits isn't possible—the number is *truncated*. The effect is phase-modulated (PM) spurs in the DDS output.

Further errors are related to the output resolution of the look-up table. The table values representing the amplitudes are truncated to some number of bits, a. This mechanism produces amplitude-modulated (AM) spurs in the output.

After Cercas,[7] maximum PM spurs could be:

$$P_{PM\;spurs} = -(6.02p - 5.17)\,\text{dBc} \quad \text{(Eq 7)}$$

maximum AM spur levels could be:

$$P_{AM\;spurs} = -(6.02a + 1.75)\,\text{dBc} \quad \text{(Eq 8)}$$

In the analog signal we generate, the DAC introduces more AM spurs, harmonics and IMD because of inherent nonlinearities (see Part 1). Spurs are also likely at the clock frequency and its harmonics. A ninth-order elliptical low-pass filter after the DAC removes harmonics and many of the spurs.

It turns out we can eliminate all AM spurs by squaring the DDS output at the external reference input of the PLL chip! We can do nothing about the remaining PM spurs, so we'd better keep them 40 dB below the desired output spurious level, or:

$$P_{(PM\;spurs)} \leq -73 - 40\,\text{dBc} = -113\,\text{dBc}$$

$$\text{(Eq 9)}$$

This means a bit-resolution decided by solving Eq 7 for p:

$$p \geq \frac{-113 - 5.17}{-6.02}$$

$$\geq 19.63 \quad \text{(Eq 10)}$$

The Harris HSP45106 has a 32-bit phase accumulator, 20-bit-address sine look-up table and 16-bit output resolution. Since the AM spurs will disappear, a 10-bit DAC is sufficient; we chose the Harris HI5780.[14]

The resulting DDS output feeds the reference input of the Motorola MC145159 PLL synthesizer IC, where the wave is squared and divided by 10 to establish a phase reference of 75 to 105 kHz ($F_{VCO}/1000$). This should provide very fast lock times!

Frequency Resolution

As stated above, the PLL will multiply the DDS frequency by 100. To get our 1 Hz output steps, therefore, we must tune the DDS in 10 *millihertz* steps! With a 32-bit phase accumulator, DDS step size will be:

$$df_{DDS} = \frac{f_{clk}}{2^{32}} \quad \text{(Eq 11)}$$

A clock frequency around 10 MHz exceeds the goal by a factor of 4, producing a step size of 2.3 millihertz!

The VCO

A Colpitts design was selected, with a resonant tank circuit switched in eight bands using PIN diodes and three capacitors with binary-weighted values, as shown in Fig 9. The VCO is tuned using back-to-back varactor diodes, in order to achieve maximum voltage across the tank.

The advantage of restricting the frequency range to several bands is one of decreased phase noise. According to
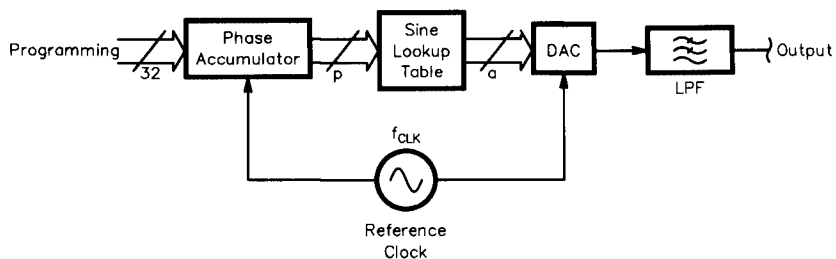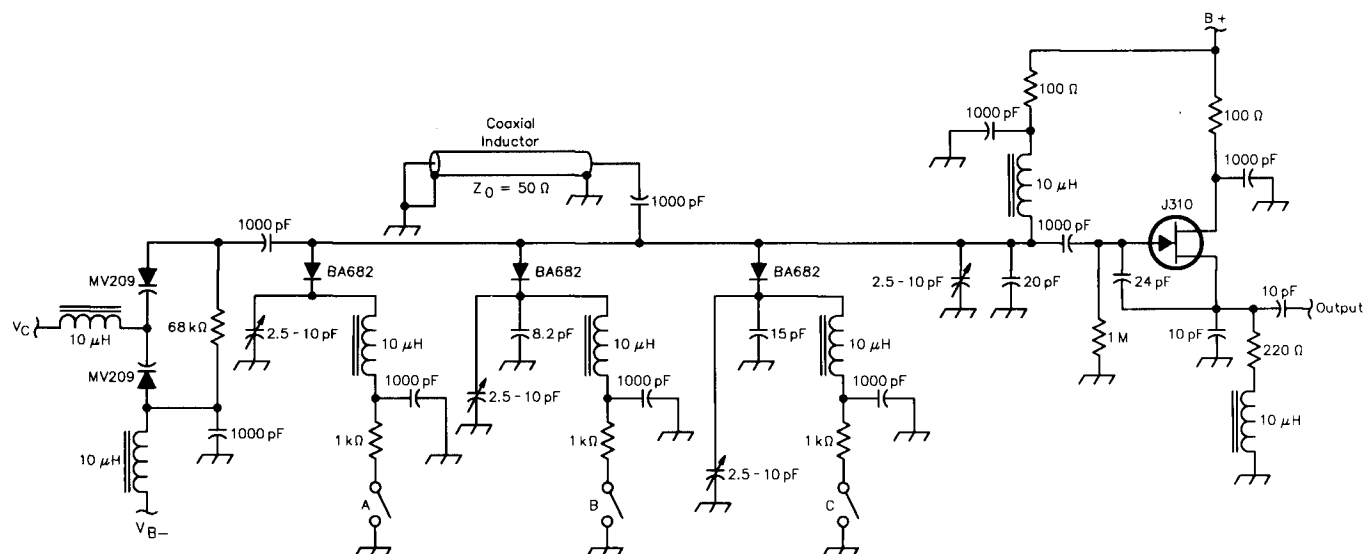

Fig 8—DDS block diagram.

Fig 9—75-105 MHz band-switched VCO for LO$_1$.

Leeson's model,[8] the phase noise of a VCO may be characterized by the equation:

$$PN_{df} = 10 \log \left\{ \left(1 + \frac{f_c^2}{\left[2Q_{load}df\right]^2} \right) \left(1 + \frac{f_f}{df} \right) \left(\frac{nkT}{2P_{avg}} \right) + \frac{2kTR_v K_\phi^2}{df^2} \right\}$$

(Eq 12)

where

PN_{df} = SSB phase-noise power relative to total power, in dBc/Hz

f_c = center frequency

df = frequency offset of noise measurement

Q_{load} = loaded Q of tank circuit

f_f = flicker frequency of active device

n = noise factor of active device

k = Boltzmann's constant, 1.38×10^{-23}

T = temperature in Kelvins

P_{avg} = average power at input to active device

R_v = equivalent noise resistance of varactors

K_ϕ = VCO sensitivity in Hz/V

Great, but what does this incredibly complex equation mean?

First, it means that the phase noise goes up about 6 dB every time we double the center frequency. Second, the phase noise goes *down* 6 dB when we double the loaded Q of the tank circuit. These effects aid each other, because it's more difficult to achieve high Qs at higher frequencies. Third, more power at the input of the oscillator device is better. Finally, the phase noise will approach a lower limit established by the other factors as VCO sensitivity is decreased.

The inductor for the resonant circuit is critical in determining Q. It can also make the VCO "microphonic," (susceptible to vibration, if not solidly mounted). We can borrow from transmission-line theory to address these issues.

A section of transmission line, shorted at the far end, and less than λ/4 long, looks like a shunt inductor[12] having a reactance of:

$$L = Z_0 \bullet \tan \beta L$$ (Eq 13)

where Z_0 is the characteristic impedance of the line, and ßL is its electrical length in radians. Eg, a 1 λ line is 2 π radians in electrical length.

When used in this way, the SWR in the line section will be quite high, making it "lossy" and degrading the Q. By keeping it short and using a carefully selected coaxial cable, however, microphonics and any chance of radiation are virtually eliminated! The shield is simply soldered to the circuit board along its entire length; it can even be coiled into a very compact size.

Output Buffering

To maintain our high resonant-circuit voltage, we want to draw as little energy from the oscillator as possible. On the other hand, we must take enough that the noise figure of the buffer amplifiers doesn't become a problem. Two buffer stages are used here. They also provide isolation; changes in load impedance caused by the presence of large receiver input signals would otherwise "pull" the VCO, inducing FM.

The PLL: Closing the Loop

Since the phase reference is at $F_{VCO}/1000$, the VCO output sample passes through an external ÷20 prescaler and a ÷50 *inside* the PLL IC before phase detection. The high comparison frequency produces fast lock times by allowing a large loop bandwidth, while making it easier to avoid "reference spurs," or phase modulation

of the VCO at the reference frequency.

The MC145159 PLL chip uses an analog sample-and-hold phase detector, which further eases loop filter design. Details of this technique can be found in the data sheet,[13] but let's note its two major advantages: It has programmable gain, and it outputs an analog level that already resembles the desired loop control voltage. We level-shift and amplify this 0 to 5 V signal until it suits the VCO's tuning range and then filter it to remove noise introduced by the amplifier and sample-and-hold circuits.

The device also outputs a coarse "frequency steering" signal, which is used to move the control voltage by large amounts when phase lock has been lost. When the phase at the reference frequency is within 2 π radians of lock, this output goes high-impedance, and the analog phase detector takes over.

LO₂: A Voltage-Controlled Crystal Oscillator (VCXO) Design

The second LO is fixed in frequency, but it needs to be phase-locked to the master reference so that some "drift cancellation" can be obtained, as explained further below. It should have very low phase noise so as not to add to the total. For these reasons, a VCXO is best.

At 75.040 MHz, an overtone "rock" must be used. A useful property of overtone oscillators is that they produce stability roughly equal to that of the fundamental mode. Because of this, fifth-overtone operation was initially considered. The tuning range must allow the VCXO to track the reference over its limits, though, and a third-overtone design is much more "pullable." The addition of a resistor across the rock aids in achieving sufficient range without lowering the Q enough to affect output noise.

Fig 10 is a modified Pierce circuit. Feedback is taken from a winding on the drain transformer and passed through the series-resonant crystal. Drain capacitors resonate the transformer at the output frequency, and transform the impedance down to 50 Ω.

A PLL using the MC145159 locks LO_2 to the reference oscillator. Two buffer amplifiers are again used.

The Frequency Reference and Drift Cancellation

As both LOs are locked to the same frequency reference, and difference mixing is used at both IFs, some degree of drift cancellation is obtained. Let's look at how this works and calculate the magnitude of the effect.

The DSP sees a 40 kHz IF signal, which has undergone two frequency translations. For a single RF input signal at f_{RF}, the mixing is expressed by the equation:

$$f_{IF} = f_{LO_2} - \left(f_{LO_1} - f_{RF} \right)$$
$$= f_{RF} - \left(f_{LO_1} - f_{LO_2} \right)$$ (Eq 14)

Now each LO varies with a change in reference frequency according to:

$$\frac{df_{LO}}{df_{REF}} = \frac{f_{LO}}{f_{REF}}$$ (Eq 15)

$$df_{LO} = f_{LO} \left(\frac{df_{REF}}{f_{REF}} \right)$$ (Eq 16)

Substituting, we can write:

$$df_{IF} = df_{RF} - \left(df_{LO_1} - df_{LO_2} \right)$$
$$= -\left(f_{LO_1} - f_{LO_2} \right) \left(\frac{df_{REF}}{f_{REF}} \right)$$ (Eq 17)

assuming the RF input isn't changing, only the reference. We see that the total error is proportional to the difference in the LO frequencies, which is approximately equal to the RF:

$$\left(f_{LO_1} - f_{LO_2}\right) \approx f_{RF} \qquad \text{(Eq 18)}$$

At low RFs, the drift cancellation is nearly perfect. As the RF increases, any error increases linearly to the maximum, at 30 MHz, of:

$$df_{IF} \approx 30 \times 10^6 \left(\frac{df_{REF}}{f_{REF}}\right) \qquad \text{(Eq 19)}$$

To get the ±20 Hz accuracy, we now know we must hold the reference to within:

$$\left|\left(df_{REF\,\text{max}}\right)\right| = \frac{20 f_{REF}}{30 \times 10^6} \text{ Hz} \qquad \text{(Eq 20)}$$

$$\frac{\left|\left(df_{REF\,\text{max}}\right)\right|}{f_{REF}} = \frac{1}{1.5 \times 10^6} \text{ Hz} \qquad \text{(Eq 21)}$$

This means keeping the reference, *whatever its frequency*, within ²/₃ parts-per-million (ppm) over the range of variables. This is difficult without a crystal oven, but DSP technology once again comes to the rescue!

Microprocessor Compensation

Considering that frequency variation versus temperature is the main factor, we arrange to place a temperature sensor that can be monitored by the DSP microprocessor near the reference oscillator. We specify the AT-cut reference crystal to a tolerance that gives us a reasonable approximation to a straight-line frequency-versus-temperature curve over the range, as shown in Fig 11. We measure the oscillator performance, and use a look-up table and digital-to-analog converter (DAC) to output a tuning voltage that precisely compensates the temperature variations.

Furthermore, the receiver can be tuned on command to an accurate external frequency standard, such as WWV, and the internal reference adjusted to match! In actual practice, this technique results in less than 0.2 ppm error. During this calibration procedure, the internal temperature is noted, and the compensation look-up table is adjusted to reflect whatever curve the oscillator happens to be following. Note that *all* the variables are inside the loop and are, therefore, canceled.

Summary

The marriage of DDS and PLL is an extremely flexible system. Although the PLL programming is normally fixed, its reference frequency can be selected to create different performance characteristics. The DDS can be similarly tuned over a wide range to suit the needs of any particular system.

We've also seen how a microprocessor-compensated crystal oscillator (MPCXO) can exceed the stability of many oven-controlled units—at a small fraction of the power. Over time, repeated calibration to a precise external standard assures the best accuracy on an adaptive basis.

Next, we'll design a transmitter that fits neatly into our plan and exploits many of the advantages obtained in the receiver and synthesizer. The unique benefits of IF-DSP technology will again be highlighted. When reducing cost while improving performance, it's hard to go wrong!

Go Ahead and Transmit!

In this final segment describing IF-DSP transceiver design, we'll define the requirements for the transmitter and look at how they merge with the structures we developed for the receiver. As before, we'll emphasize the significance of DSP-related issues—even to the exclusion of the more mundane aspects, with which most readers are already familiar.

Many receiver sections, such as mixers, filters, amplifiers and IFs, are also required in a typical transmitter design. Considerable cost savings are realized in transceivers by sharing these circuits between modes. To achieve this, we obviously must keep the same frequency conversion scheme, and do a bit of signal switching to toggle between receive (RX) and transmit (TX). First,
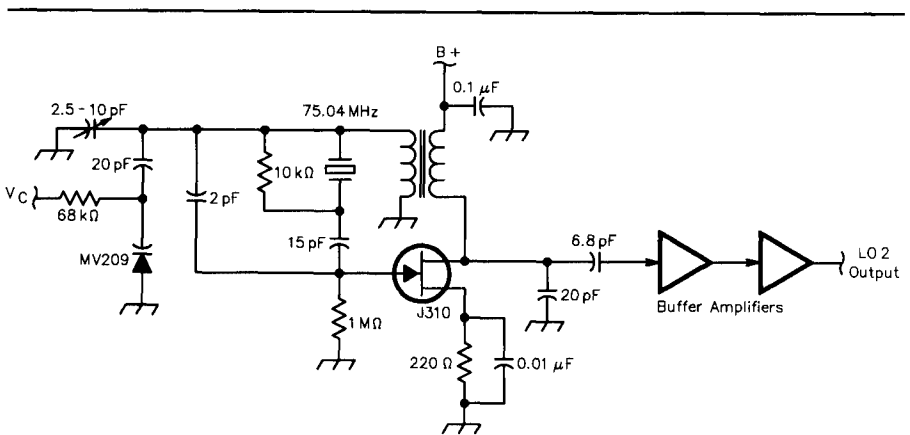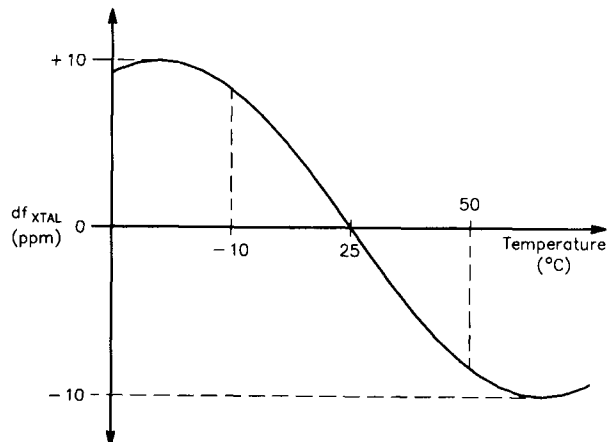

Fig 10—75.04 MHz VCXO for LO₂.

Fig 11—Reference crystal frequency *versus* temperature curve.

let's identify what we expect from the low-level TX, or exciter, function.

Modulation and Drive: The Exciter Design

The goal is to produce a signal—in one of several modulation formats—that can be amplified and transmitted over the air. For an HF transmitter, SSB, AM, FM, CW and various data modes are usually included. Our signal ought to occupy a bandwidth commensurate with good engineering practice and the quality of communication desired—and no more. It shall be, so far as possible, a faithful replica of the input or *baseband* information. We endeavor to design a system that fits simply into the existing architecture and introduces minimum distortion and noise. The receiver's block diagram is presented for review as Fig 12.

With these objectives in mind, it's clear that if we could generate a 40 kHz signal with the desired characteristics, we could translate it up to the 75 MHz IF, then to RF. We wouldn't need narrow second-IF filtering if we could ensure that the 40 kHz transmit signal were already bandwidth-limited. We *would* need the first IF's crystal filters though, to remove the image product and LO bleed-through at 80 kHz and 40 kHz away, respectively.

T/R Switching

The first IF strip can be used in the TX mode by swapping the LOs. This requires a double-pole, double-throw switch, which is implemented using PIN diodes. Good isolation between the injection signals is mandatory.

The input to the strip will be the 40 kHz transmit signal, with the output translated to RF by the second mixer. We must have switches at both ends to select input and output signals, as shown in Fig 13. Good isolation and linearity are critical in these switches. Note that we've avoided switching any low-level, 75 MHz signals in this design.

Many traditional analog signal-processing stages are unneeded, as in the receiver. The balanced modulator, sharp crystal or mechanical filters, speech processor, gain-controlled stages and carrier-null adjustments disappear! In addition, each unit will perform identically to the next, because it's all done in firmware.

Level and Gain Determination

Referring to Fig 14, as the second mixer is a low-level device, we must keep its input level low to avoid objec-

tionable spurious products. For a level-13 mixer, the resulting output level is about −28 dBm. Therefore, 48 dB of gain is necessary to achieve a +20 dBm exciter output. Once we get to about the 0 dBm level, push-pull stages are helpful in reducing second harmonic output and in obtaining sufficient levels. A final 30 dB power amplifier takes us to 100 W. It's best to shoot for at least a 3 dB gain margin, so the exciter output stages are designed to handle up to +23 dBm, or 200 mW.

At the first mix (from 40 kHz up to 75 MHz) we must have a high drive level to overcome the LO bleed through. The LO-to-IF isolation might be as low as 40 dB, so with a +17 dBm injection level, the bleed-through might be as high as:

$$P_{bleed-through} = (17 - 40)\,\text{dBm} = -23\,\text{dBm}$$
(Eq 22)

The crystal filters will attenuate this product by a further 65 dB. So, to keep it less than the design goal of −70 dBc, and considering the mixer's conversion loss is about 6 dB, we must use a drive level of at least:
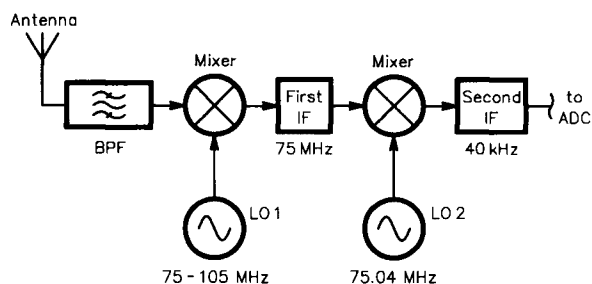

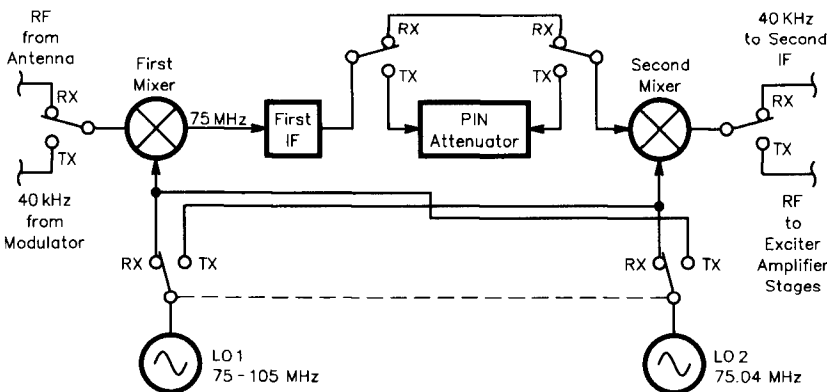Fig 12—IF-DSP receiver conversion scheme.

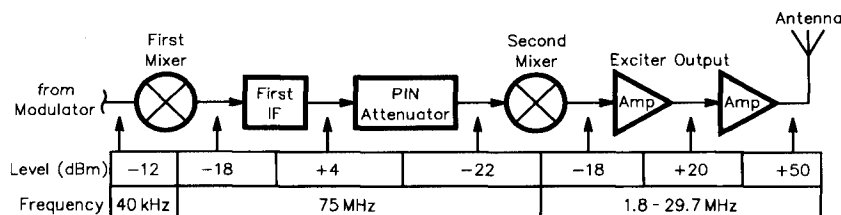Fig 13—Conversion scheme with T/R switching added.

Fig 14—Transmit level diagram.

$$P_{(40\,kHz\,drive)} = (-23 - 65 + 70 + 6)\,dBm = -12\,dBm \qquad (Eq\ 23)$$

Whereas in the receiver, the first IF strip had a gain of about 10 dB, it must now have a loss of:

$$LOSS_{IF_1} = (-12) - (-28)\,dBm = 16\,dB \qquad (Eq\ 24)$$

The PIN diode attenuator just before the second mixer must therefore have attenuation equal to:

$$ATTEN_{PIN} = 10 - (-16)\,dB = 26\,dB \qquad (Eq\ 25)$$

Receive / Transmit Gain Comparisons

Consider the total power gains in the receiver versus the transmitter. The receiver takes as little as –132 dBm from the antenna and amplifies it to around 1 W at the loudspeaker, or +30 dBm; the power gain is:

$$GAIN_{RX} = 30 - (-132)\,dBm = 162\,dB \qquad (Eq\ 26)$$

In the transmitter, a typical dynamic microphone might produce 5 mV (RMS) into 600 Ω, or:

$$P_{MIC} = \frac{\left(5\times10^{-3}\right)^2}{600} \approx -44\,dBm \qquad (Eq\ 27)$$

The gain is:

$$GAIN_{TX} = 50 - (-44)\,dBm = 94\,dB \qquad (Eq\ 28)$$

The receiver has a far more difficult task, but the transmitter is still doing yeoman's duty! Now think of the maximum path loss of:

$$LOSS_{PATH} = 50 - (-132)\,dBm = 182\,dB \qquad (Eq\ 29)$$

and ruminate on the fact that the total power gain from microphone to loudspeaker must be:

$$GAIN_{TOTAL} = 162 + 94\,dB = 256\,dB \qquad (Eq\ 30)$$

a factor of 4×10^{25}! It's a wondrously large amount of enhancement we get from our electronics!

The SSB Modulator

As described in Part 1, we intend to use the phasing method of SSB generation. We learned it's convenient to choose an output sampling rate four times that of the output frequency, because the injection to the modulator takes on values of only one or zero, simplifying matters. This will be a sampling rate of four times 40 kHz, or 160 kHz. The digital-to-analog converter (DAC) will be humming right along!
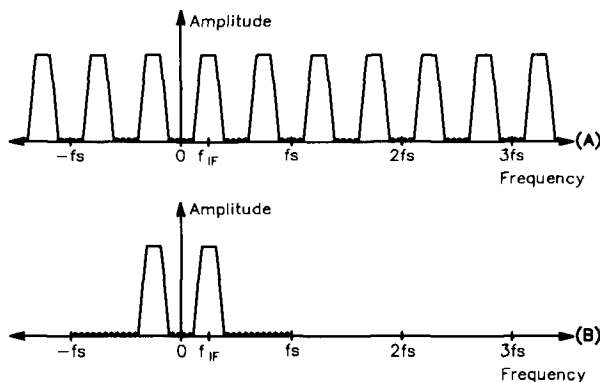
We'll bandwidth-limit the baseband information and the resulting SSB signal, but as we saw before, the sampling process will cause the output spectrum to repeat at harmonics of the sampling frequency,[3] as shown in Fig 15. Fortunately, it's easy to build a low-pass filter that will remove these *alias* products.

On inspection, we quickly see it's not practical to sample baseband signals at a rate of 160 kHz. For one thing, we must bandwidth-limit the signals to something like 3 kHz, and building a filter with a fractional bandwidth of:

$$BW_{FRAC} = \frac{3\,kHz}{160\,kHz}$$
$$\approx 0.01875 \qquad (Eq\ 31)$$

and decent attenuation characteristics would be a fantasy—even in the DSP world! Additionally, such a high sampling rate is beyond the capabilities of most DSP chips with on-board ADCs. A sampling rate equal to the IF (40 kHz) is more reasonable. This allows us to build a pair of transmit filters having the response shown in Fig 16, with 6 dB points at 180 Hz and 2.9 kHz.[9]

The filtered output is then *interpolated* (see Part 1) up to the 160 kHz rate. An *interpolation filter* removes the alias components due to the lower sampling frequency prior to application to the modulator. The result is a bandwidth-limited SSB signal, ready for translation to RF.

Other Transmitter Modes

Other modulation modes are considerably simpler than the SSB phasing method. In CW, for example, we just output a single, 40 kHz signal. Note that the DSP system makes it easy to shape the rise and fall times of the transmitted CW note. The keyer—which may also incorporate speed and weighting controls—can have adjustable dynamics, from "hard" to "soft," to suit the operator.

FM and AM modes are a little tricky, because we must limit the amplitude of the baseband information to prevent overmodulation. As we'll discover below, the automatic level control (ALC) in AM presents some interesting difficulties.

Distortion and Noise Sources

As discussed previously, numerical truncation, quantization noise and DAC nonlinearity affect the quality of any

Fig 15—(A) DAC output spectrum showing "aliases. (B) Output spectrum after low-pass filtering.
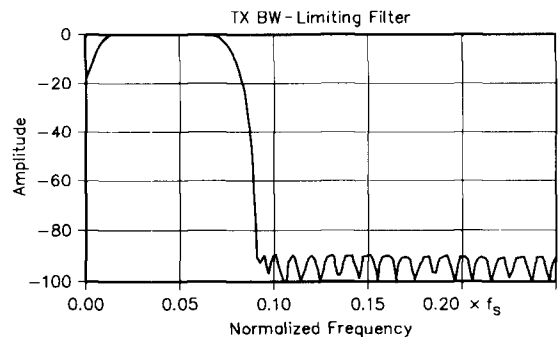

Fig 16—Response of transmit bandwidth-limiting filters.

digitally processed signal. As we revisit these topics, we'll see how they place limits on transmitter performance and why component selection is critical.

First, SSB opposite-sideband rejection and carrier suppression are of concern. Whether analog or digital, amplitude and phase inaccuracies degrade the opposite-sideband suppression in a phasing-method modulator. Maintaining 16-bit data representation throughout the system ensures that computational effects are negligible in a digital implementation.[3] The DAC performance is generally the limiting factor. The best 16-bit DACs produce amplitude and phase accuracy quite adequate for our needs, resulting in typical opposite-sideband suppression of –70 dBc. Note that this level is far lower than the intermodulation distortion (IMD) produced by the final PA.

In a digital modulator of this type, carrier rejection is established mainly by the dc offset present at baseband. This is easy to correct, since in the absence of input audio, the DSP can measure any offset. It's subtracted prior to modulation.

Secondly, noise produced by quantization effects can be significant. A 10-bit ADC has a maximum signal-to-noise ratio (SNR) of:

$$\left(\frac{3}{2}\right)2^{20} \approx 62\,\text{dB} \qquad \text{(Eq 32)}$$

This will also be the SNR of the transmitter output. This level is deemed sufficient for all applications.

Lastly, we must consider mixer performance at an RF-port frequency of 40 kHz. Certainly, the mixer must be designed to handle signals down to this range. The presence of an antialiasing filter will degrade the IMD characteristics (as described in a prior segment) unless idler networks are used.

Digital Automatic Level Control (ALC)

In this design, ALC is realized solely by controlling the amount of 40 kHz drive signal. Information about output power is obtained from a directional coupler, which measures both the forward and reflected power at the output of the low-pass filters that limit harmonic radiation. As it's our intention to place an automatic antenna-tuner unit (ATU) between this point and the antenna, a phase detector is also included.

Refer to Fig 17. As described in *The ARRL Antenna Book*,[12] this directional coupler produces output levels proportional to the square roots of both forward and reflected power, ie,

SWR information. These two signals are rectified, filtered and then fed to ADC inputs on the DSP chip. (A small bias current is passed through the detector diodes so they're more sensitive. This allows measurement of as little as 0.3 W.)

The forward output is continually compared to a predetermined threshold. When the threshold is exceeded, drive is reduced. When maximum output levels are not being reached, gain is slowly increased to a preset limit. The whole thing works much like a traditional analog ALC.

When the reflected power exceeds a certain amount, we reduce the forward power to protect the output devices. DSP makes it easy to hold the reflected power to a fixed value, such as 10 W.

ALC in AM

It's long been a problem to hold the carrier level constant in AM transmitters. Because the baseband signal may not have symmetrical positive and negative amplitudes, a suitable analog ALC system would be incredibly complex. In DSP, we can prevent *carrier shift* by using adaptive techniques.
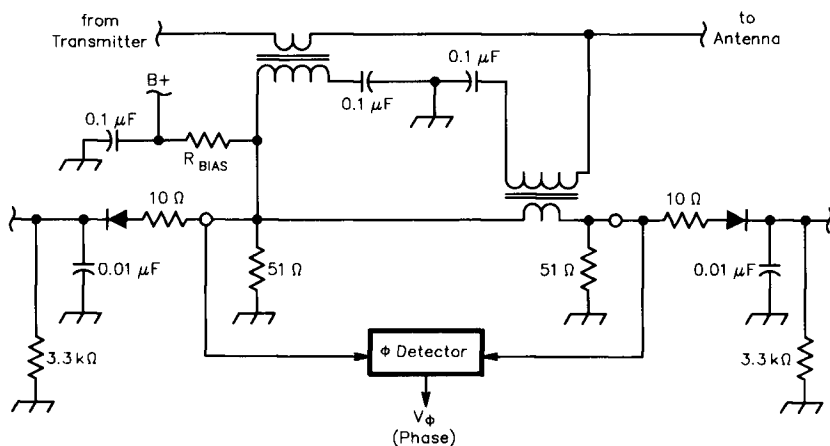
Fig 17—Directional coupler and detectors.

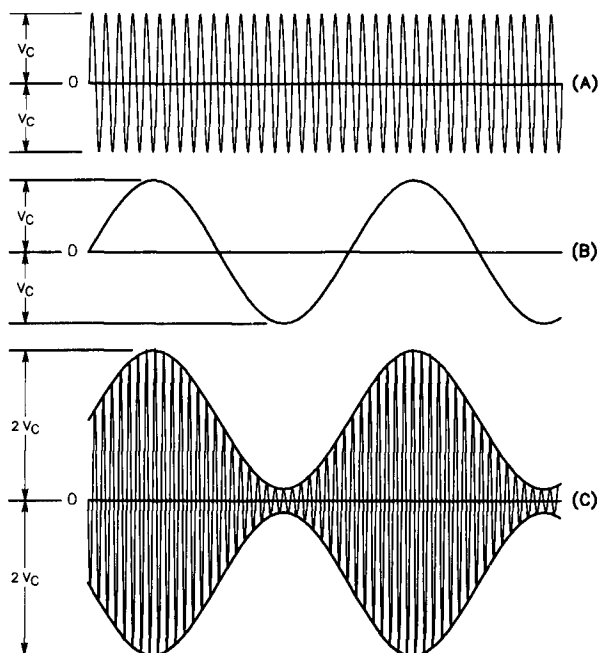Fig 18—(A) Carrier wave. (B) Baseband input waveform. (C)AM output.

First, the ratio of drive level to output power is easily computed when the transmitter is on, so we can determine the drive level required to achieve a carrier level that is 25% of the peak power setting. Second, the baseband signal must be held to a maximum peak level that equals the carrier drive level. When the carrier and peak-limited baseband levels are added, the result will be a 100%-modulated AM wave,[10] as shown in Fig 18.

This means that *two* ALC servomechanisms operate in our AM ALC: One continuously computes the drive-to-output ratio and maintains the carrier level. The second compresses the peak baseband signal to that same level. The system for this is shown in Figure 19.

Since the baseband peak detector employs a full-wave rectifier in firmware, audio inputs with asymmetrical positive and negative voltage swings can produce unexpected results. That is, either the upward or downward modulation may reach 100% before the other can do so; if the downward modulation limits baseband amplitude first, the peak envelope power cannot reach its set level without introducing a carrier shift! See Fig 20.

SWR Computation

As our detector outputs are proportional to the forward and reflected voltages, the *reflection coefficient* is just the ratio:

$$\rho = \frac{V_{REFL}}{V_{FWD}} \qquad \text{(Eq 33)}$$

and the SWR is calculated using:

$$SWR = \frac{(1+\rho)}{(1-\rho)} \qquad \text{(Eq 34)}$$

To find the actual antenna impedance, we also need to know the *phase* of the reflection coefficient. To get the relative phase of the coupler's two output signals, we'll use a digital phase detector much like those found in PLL chips.

At This Phase of the Game

To determine the phase, we can build a circuit that finds the ratio of the time, t, between the rising edges of the forward and reflected voltages to the total RF period, p. See Fig 21. In signed format, the phase (in degrees) is then:

$$\phi = 360° \left(\frac{1}{2} - \frac{t}{p} \right) \qquad \text{(Eq 35)}$$

A dual-D flip-flop can be configured to output a "1" during time t, and a "0" otherwise, with the forward and reflected voltages as the clock inputs.
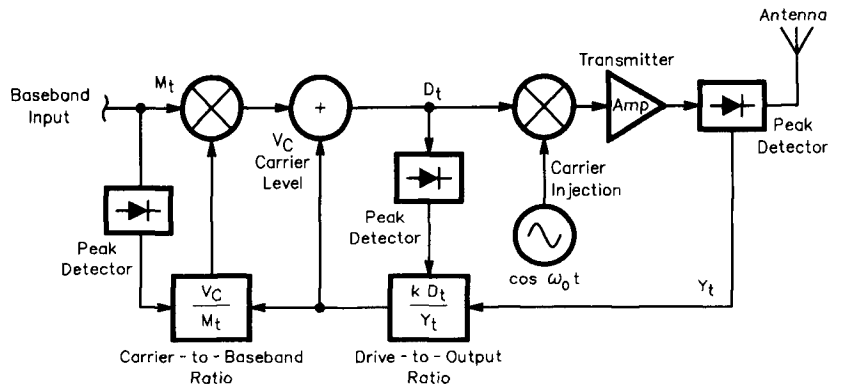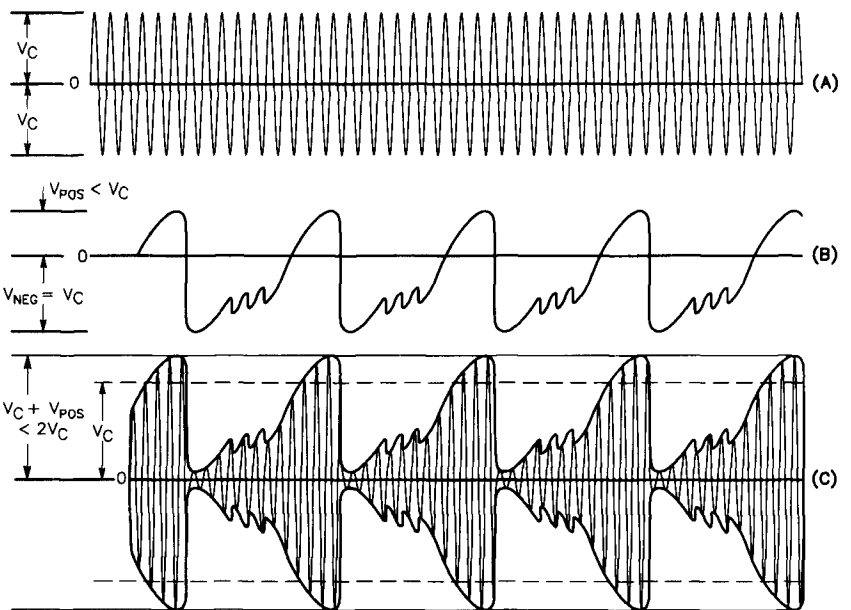
Fig 19—AM ALC block diagram.

Fig 20—(A) Carrier. (B Baseband input with asymmetrical amplitudes. (C) AM output.
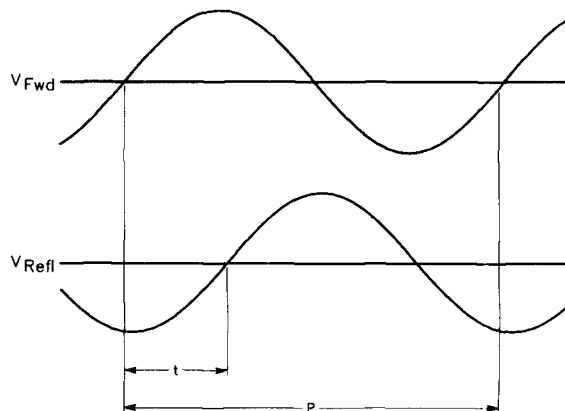
Fig 21—Phase relationship of forward and reflected voltages.

The output is then integrated with a simple RC filter, and the result is a voltage from 0 to 5 V that is directly proportional to relative phase. A circuit to do this is shown in Fig 22.

We can guarantee a constant forward voltage, but the reflected level may be quite low near 50 Ω, so an amplifier is required ahead of this clock input. Even so, the device will fail to clock properly below some value of reflection coefficient. This threshold needs to be less than a SWR of, say, 1.3:1, corresponding to a reflection coefficient of:

$$\rho = \frac{(SWR - 1)}{(SWR + 1)}$$
$$\approx 0.13$$
(Eq 36)

Almost 18 dB of gain is therefore required to match the forward voltage, and the amplifier can clip at larger input levels to avoid overdriving the chip.

Complex Impedance Transformations

Magnitude and phase information in hand, we can compute the antenna impedance directly! The transforms are easier to present if handled in two steps. The amplitude and phase angle of the reflection coefficient are in *polar form*, so the first conversion is to *Cartesian* coordinates x and y:

$$x = \rho \cos \phi$$
$$y = \rho \sin \phi$$
(Eq 37)

These numbers, which range from −1 to 1, may then be used to plot the impedance point on a *Smith Chart*[11] of unity radius.

A second step finds the normalized complex impedance $R + jX$:

$$R = \frac{1 - x^2 - y^2}{(1 - x)^2 + y^2}$$
$$X = \frac{2y}{(1 - x)^2 + y^2}$$
(Eq 38)

To convert to a 50 Ω system, we simply multiply R and X by 50. These data are going to be quite useful as we consider the ATU that must provide a *conjugate match* between our transmitter and an antenna, which may not look like a 50 Ω resistor!

ATU Configuration

In order to hold the internal ATU to a sensible cost, we must limit the range of antenna impedances to be matched. We soon discover that, outside the SWR = 3:1 circle, the voltages and currents in the matching elements rise rapidly. External ATUs must deal with a much larger range of

antennas. This places stringent demands on component Q. Physically large inductors and capacitors are employed to reduce losses.

It turns out that we can always achieve a conjugate match to the antenna using an LC network with a series inductance and a shunt capacitance. For antenna impedances with resistance greater than 50 Ω, the capacitance needs to be at the output side; those with resistance less than 50 Ω need the capacitance on the input. We arrange to switch binary-weighted inductance and capacitance values into the circuit using relays (as shown in Fig 23), so we can obtain the range of values we need. Now all we need is an al-

gorithm that drives the network toward a match under microprocessor control.

"Fuzzy-Reasoning" ATU Algorithms

Fuzzy reasoning is a process, like those in the human mind, which assesses a situation in relative terms. For example, if we see the antenna is capacitive, we know inductance must be added; if inductive, capacitance must be used. Further, if the antenna is *very* capacitive, *more* inductance must be inserted. A fuzzy-reasoning system employs *transfer functions* that describe how much adjustment to make based on detector inputs. The transfer functions can represent not only the theoretical requirements of
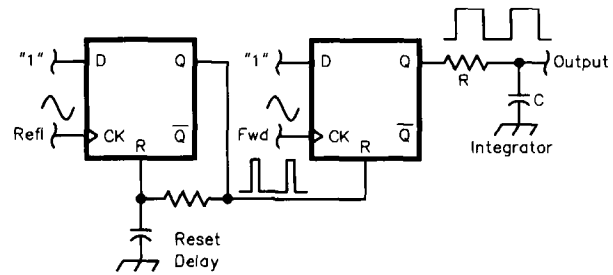

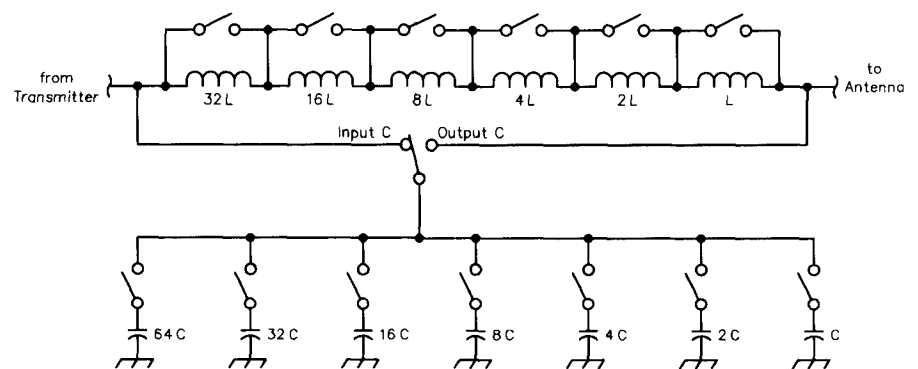
Fig 22—Digital phase detector.
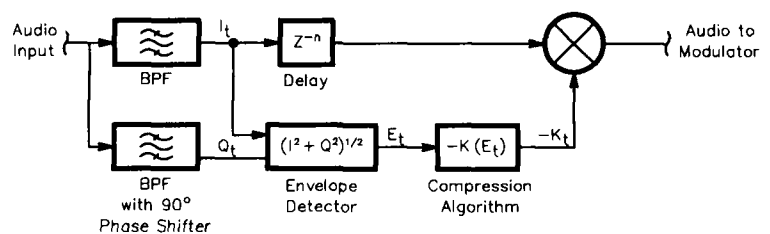

Fig 23—ATU network configuration.

Fig 24—Digital RF compressor block diagram.

the system, but can also incorporate any predictable errors from the detectors and other sources. Fuzzy reasoning tends to overcome errors in systems that can provide only roughly accurate absolute measurements, but produce good relative resolution.

In the case of the ATU, the accuracy of the phase detector degrades rapidly when SWR is less than 1.3:1, and we must rely solely on the reflection coefficient to guide us. A certain amount of "thrashing about" must be employed to find the minimum SWR. Above this level, the phase information is useful in steering toward the goal. To achieve this, we develop transfer functions that embody the matching rules and create a fuzzy-reasoning *engine* that adjusts circuit elements on a step-by-step basis until it reaches minimum SWR.

Step size must be determined by the degree of correlation between the transfer functions and the actual performance of the circuit. Tuning speed is the parameter that suffers because of inaccuracies. In actual practice, measure-and-adjust cycles of around 25 ms yield tuning times well under one second. The use of adaptive, memory-tuning techniques enhances performance.

Adaptive ATU Memory

After several tune cycles on a particular antenna, we begin to get some idea of how it performs. If a frequency is selected near one that has already been matched, a network may be estimated from the previously stored data. The tuning time is therefore greatly reduced. As additional memory points are stored, the number of steps in each tune cycle diminishes until, finally, the antenna system is wholly characterized. If enough points can be stored, tuning time shrinks to that required for switching to the correct network—and no more.

As the network data are available to the control system, the antenna impedance can be plotted across the band. This is extremely useful during initial setup of an antenna, and during continued operation.

Speech Processing: More Bang per Buck

Another distinct advantage in a digital, phasing-method SSB modulator is that the RF envelope can be calculated before the modulation is performed! This allows us to employ RF compression methods on the baseband signal prior to filtering, where they can be effective without adding to "splatter." This scheme is shown in Fig 24.

As explained in Part 1, the envelope of the SSB output is computed as:

$$E_{SSB} = \left(I^2 + Q^2\right)^{1/2} \qquad \text{(Eq 39)}$$

To avoid the time-consuming square-root calculation, we can use an approximation:[1]

$$\left\{\begin{array}{l} For : |I| > |Q|, \quad \left(I^2 + Q^2\right)^{\frac{1}{2}} \approx |I| + 0.4|Q| \\[2ex] \qquad |Q| \geq |I|, \quad \left(I^2 + Q^2\right)^{\frac{1}{2}} \approx |Q| + 0.4|I| \end{array}\right. \qquad \text{(Eq 40)}$$

This envelope amplitude is used to compress the baseband levels so that the peak-to-average ratio of the transmitted signal is reduced. That is, the average power is increased. The effect is the same as that produced by RF processing. This naturally involves the introduction of distortion, since the transmitter is no longer linear. Nevertheless, this type of distortion enhances the syllabic and formant energy in speech without introducing the "mushy" sound caused by audio clipping.

To elaborate, consider that the human voice has a peak-to-average ratio as high as 15 dB. This doesn't use a peak-limited transmitter very well, and at the 100-W PEP level, the average output power might be as little as 3 W! RF compression enhances the weaker parts of human speech such that intelligibility is improved. As shown by the studies in the reference literature, 15 dB of RF compression can produce up to 6 dB of intelligibility improvement on the receiving end. This is equivalent to quadrupling the output power!

The compressor attack and decay times can be varied to change the amount of processing introduced. As they are made faster, compression approaches the effects of RF clipping. It's widely known that this is the most effective form of speech processing.

Conclusion

We've seen that first, we must design the most linear system possible in a transmitter; then, to improve speech intelligibility, we must destroy the linearity! In Part 3 of this series, we'll examine advanced DSP techniques that further improve communication, in both the transmitter and receiver. We'll introduce adaptive signal-processing methods that wouldn't be possible without DSP technology. They can correct for many traditionally troublesome production variations. Moreover, we'll see how computer control of transceivers makes many interesting features easy to implement!

Doug Smith, KF6DX, is an electrical engineer with 18 years experience designing HF transceivers, control systems and DSP hardware and software. He joined the amateur ranks in 1982 and has been involved in pioneering work for transceiver remote-control and automatic link-establishment (ALE) systems. At Kachina Communications in central Arizona, he is currently exploring the state of the art in digital transceiver design.

Notes

[1]Sabin, W. E. and Schoenike, E. O., editors, *Single Sideband Systems and Circuits*, McGraw-Hill, New York, NY, 1987.

[2]Rohde, U. L. and Bucher, T. T. N., *Communication Receivers: Principles and Design*, McGraw-Hill, New York, NY, 1988.

[3]Oppenheim, A. V. and Schafer, R. W., *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1975.

[4]Frerking, M. E., *Digital Signal Processing in Communications Systems*, Van Nostrand-Reinhold, New York, NY, 1993.

[5]Gruber, M., *ARRL HF Product Review Test Manual*, ARRL, Newington, CT, 1992.

[6]Zverev, A. I., *Handbook of Filter Synthesis*, John Wiley & Sons, New York, NY, 1967.

[7]Cercas, F. A. B., Tomlinson, M. and Albuquerque, A. A., "Designing with Digital Frequency Synthesizers," *Proceedings of RF Expo East*, 1990.

[8]Leeson, D. B., "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, volume 54, 1966.

[9]Alkin, O., *PC-DSP*, Prentice-Hall, Englewood Cliffs, NJ, 1990.

[10]Panter, P. F., *Modulation, Noise, and Spectral Analysis*, McGraw-Hill, New York, NY, 1965.

[11]Smith, P. H., *Electronic Applications of the Smith Chart*, McGraw-Hill, New York, NY, 1969.

[12]Straw, R. D., Ed., *The ARRL Antenna Book*, 17th edition, ARRL, Newington, CT, 1994.

[13]Motorola Semiconductor, *AN969: Operation of the MC145159 PLL Frequency Synthesizer with Analog Phase Detector*, Motorola, Incorporated, Austin, TX, 1991.

[14]Harris Semiconductor, *DSP Databook 1994: DB302B*, Harris Corporation, Melbourne, FL, 1994.

[15]Analog Devices, *AD7722 Data Sheet*, Analog Devices, Inc, Norwood, MA, 1997.

Measurement of Signal-Source Phase Noise with Low-Cost Equipment

Do you want to explore phase-noise performance, but don't have a kilobuck for equipment? Learn how to make those measurements with relatively simple, inexpensive equipment.

By Bruce E. Pontius, NØADL

Methods for Measurement

There are several methods for measurement of phase noise on signal sources. The methods require various degrees of sophistication in the test equipment used. Some of the techniques are listed below in order of progressive complexity, accuracy and cost.

1. Listen to the noise on a communication receiver (as you are forced to do in a real radio-communication application where noise causes interference and lowers communication range or quality).

15802 N 50th St
Scottsdale, AZ 85254
e-mail **bepontius@aol.com**

2. Display it directly on a spectrum analyzer.

These first two methods, while simple, are limited to measurement of sources with high phase-noise levels, due to dynamic range restrictions in the receiver or analyzer caused by the measuring instrument's own phase noise.

2a. Use a narrow notch filter to attenuate the carrier, while allowing the phase noise to pass for observation on a spectrum analyzer.

3. Mix the test signal with a reference source to produce a lower frequency. Then filter the new test signal to reduce the carrier so the noise can be measured using a communication receiver or a special receiver, eg, a measuring receiver, power meter or spectrum analyzer.

4. Use phase-lock techniques to eliminate the carrier so the noise can be measured using a communication receiver, special receiver or a spectrum analyzer.

5. Measure it with a specially designed set of equipment like the HP 3048A (3048AR in 1997) Phase Noise Measurement System using reference sources like the HP 8662A Low Noise Signal Generator. This computer controlled and compensated system is the most accurate of those mentioned, but it is costly at $35,900 for the software and the pieces in the test set. With opt 001, the 8662A signal generator (for a heterodyne and phase-locking reference source) is $79,950. This combination covers frequencies up to 1280 MHz.

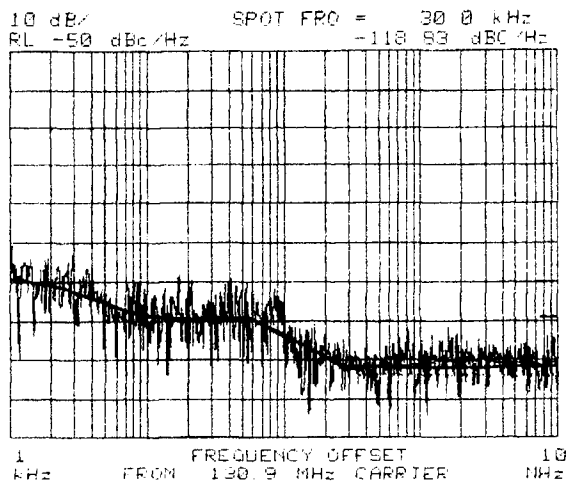We will explore the measurement of

Figure 1—Phase noise of an HP 8640B displayed on an HP 8560E spectrum analyzer with Phase Noise Utility software for controlling the display.
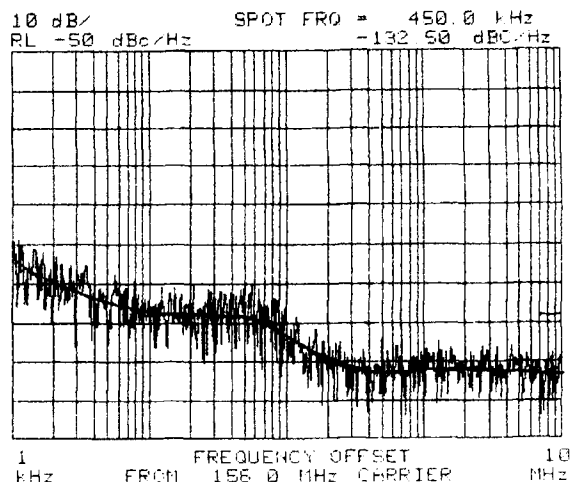

Figure 2—Phase noise of a Marconi 2024 signal generator, followed by a clean-up filter and displayed on an HP 8560E with Phase Noise Utility software for controlling the display.

phase noise with low-cost equipment and compare the results to measurements made using the HP 3048A system.

Spectrum Analyzer Method

First, how well does a high-quality spectrum analyzer do the job? An HP 8560E was used to display the output of signal generators in a direct-measurement mode. To simplify the presentation and reading of results, I used the phase-noise utility software program available for the 8560-series analyzers. The same numbers can be obtained without that software but they are not as easy to read or as nicely displayed.

An HP 8640 signal generator gave the noise display shown in Figure 1 and a Marconi 2024 signal generator caused the display in Figure 2. The 8640 noise is less than that contributed by the analyzer, so we are merely reading the phase noise of the spectrum analyzer, which is in agreement with the specifications listed in the analyzer's manual. The 8662's phase noise is even less than that of the 8640 so it was not even needed in this test. However, the Marconi generator has a higher level of noise, which can be measured with the next method I'll discuss, so it was run on this setup for another reference. A four-pole crystal filter, about 28 kHz wide, was used to clean up the Marconi's signal so that the analyzer's noise floor could be observed. Again, the analyzer noise dominates the display. The top line
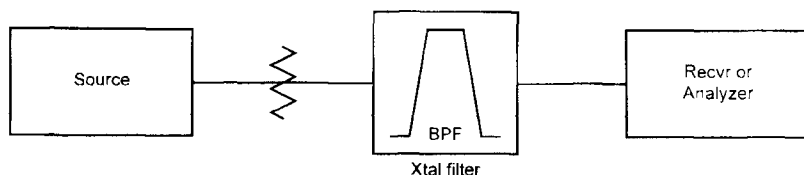

Figure 3—Filter method

represents –50 dBc/Hz and the Spot Frequency readings highlight other numbers.

So, the HP 8560E-series analyzers can be used to directly measure noise higher than about –115 dBc/Hz at offsets less than 100 kHz and about –130 dBc/Hz at greater than 400 kHz offset frequencies, as Figures 1 and 2 show. Levels below those get lost in the analyzer's own noise when the strong carrier from the source is present. Similar numbers resulted from tests with HP 8563Es, the 26.5 GHz version.

The presence of the full carrier from the source under test causes dynamic-range problems. Crystal filters were used to provide more sensitivity to enable lower noise measurements by attenuating the carrier as shown in Figure 3, resulting in the displays of Figures 4 and 5. First, the source under test is set to a frequency within the passband of the filter and a reference level is established on the analyzer. Then the source is tuned

away in frequency to the desired offset frequency, above or below the filter's passband. The carrier is attenuated, but noise in the filter's passband can be measured if the carrier attenuation is adequate.

In Figure 4, the noise power in the filter's passband, at 50 kHz offset, is shown on the 8560E's display printout as –101.2 dBm across the 1 kHz BW for a density of –131.2 dBm/Hz. The carrier power was +8 dBm going into the analyzer, so the noise is +8 + 131.2, or 139.2 dB below the carrier or –139.2 dBc/Hz. (Remember, the carrier in the picture is attenuated by the filter.)

The arrangement shown with VHF overtone crystal filters will allow accurate measurements 20 kHz or more from the carrier, down to about –143 or maybe –145 dBc/Hz. Similar filters are available at frequencies from HF through at least 200 MHz. At HF, where narrow crystal filters are readily available, the offsets could be less than 500 Hz. A communication

```
ATTEN 10dB    UAUG 100    MKR -112.6dBm
PL -50.1dBm   10dB/       130.8250MHz
```
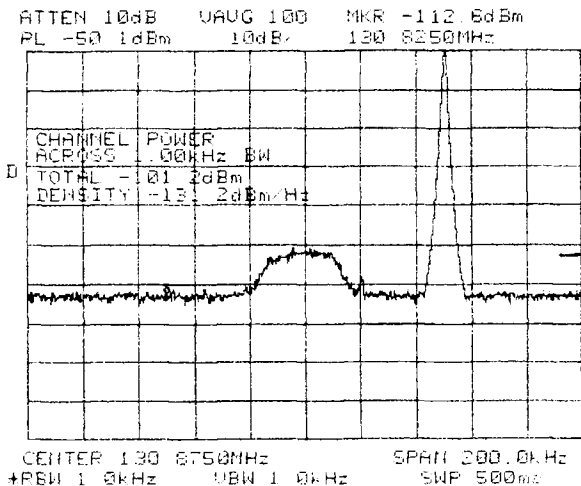
Figure 4—Marconi 2024 with a 130.875 MHz filter. Carrier offset is 50 kHz from the filter center frequency. The noise density in the center of the filter is −131.2 dBm/Hz.


```
ATTEN 10dB    UAUG 100    MKR -113.8dBm
PL -50.1dBm   10dB/       130.9257MHz
```
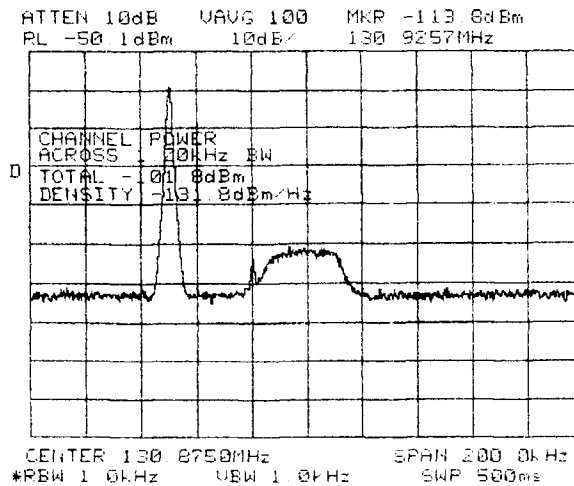
Figure 5—Marconi 2024 with the carrier on the other side of the filter.

receiver can be calibrated for noise and used to measure the noise as well. This was done in the HF range using different receivers and the numbers obtained (with a lot more labor involved) agreed closely with the analyzer numbers.

When using a receiver instead of an analyzer, an attenuator ahead of the receiver is set to give a convenient reading on the receiver's S meter, say S9, on the carrier at the filter passband frequency. Then the carrier is shifted and the attenuator is reset for the same S meter reading. The carrier power remaining must be reduced enough by the external filter to allow the receiver to work properly. Using eight-pole crystal filters usually makes the measurement possible by providing 70 dB or more of attenuation only about one filter bandwidth away on either side. With the faithful old Drake receiver I often use, the meter reading is about 3 to 4 dB higher on noise than on a CW signal when using the 500 Hz IF filter inside the receiver. The noise power is computed by algebraically adding the attenuator difference and −27 dB (for the 500 Hz filter BW) to arrive at decibels per hertz, referenced to the carrier. (A power meter can be used as well, when preceded by a filter and some gain. This type of measurement will be described later, in the discussion of heterodyne methods.) The numbers recorded using the receiver were within a decibel or so of the analyzer
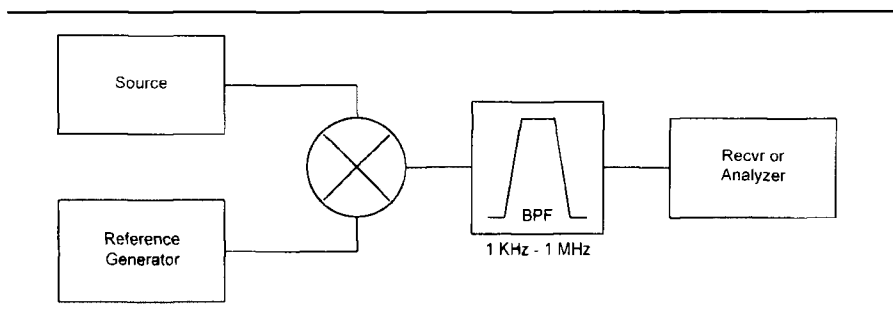
numbers, for tests in the 9 MHz range.

This method is limited by the frequency tuning range of the sources under test and the availability of narrow filters at the desired frequencies.

Heterodyne and Filter Method

Figure 6 shows a method that can provide greater sensitivity. A reference generator and mixer are used to reduce the carrier to zero hertz or a low frequency that is attenuated, by a 1 kHz to 1 MHz filter. The same analyzer or receiver can be used to measure the noise without the carrier interfering. Only one filter is required, and its frequency remains constant, but a mixer and another reference source are required. The reference source need not be an expensive signal generator, but its phase-noise level should be equal to or less than that of

Figure 6—Heterodyne and filter method.

the source under test. The reference source can simply be another source identical to the one under test, in which case the noise power measured is the uncorrelated sum of the two, or each source contributes half of the power measured. Crystal oscillators (with frequency multipliers if necessary) are low-noise reference sources. They offer enough frequency adjustment range and electronic tuning range for most applications.

This method is very good at lower frequencies or with stable sources. At VHF and above, source frequency drift can cause measurement problems because the carrier output of the mixer drifts up to higher frequencies that can overload the measurement system. A VHF or UHF VCO that is not in a synthesizer loop can quickly drift out of the measurement range of this

system. The measurements are limited by the passband of the filter. If you want to look at noise within 1 or 2 kHz of the carrier, the filter cut-off must be lowered, which requires more stability from the sources. Overall, this method is moderately simple and quite useful, and it can be used with a spectrum analyzer for displaying and measuring the noise. A receiver or power meter can also be used.

Phase Locking and Carrier-Elimination Method

Adding a little more complexity, in the form of phase-locking circuits and a preamp, can extend the measurement range and handle drifting sources.

Figure 7 shows the phase-locking and carrier-elimination method setup. We will describe this method with several examples using different equipment.

Starting with measurements at VHF, some transmitters were connected with the test setup as shown in Figure 8. I used a +23 dBm high-level mixer (the MCL RAY-1) so that I would still have a strong signal to work with after carrier elimination. The low-pass filter is shown in Figure 9. The low-pass filter should be effective over the full range of LO/RF frequencies to be used. LO and RF feedthrough can upset the phase-lock and preamp circuits. Figure 10 shows the phase-lock dc amplifier with fre-

quency-response shaping and a means of adding bias for the electronic tuning. This circuit was designed for use with reference generators requiring zero bias at the tuning input. The battery shown provides a 4.5 V bias useful for many sources and their associated tuning networks. The amplifier is powered by batteries and housed in a die-cast aluminum box to reduce noise pick up.

The preamplifier in Figures 8 and 11 was included to increase the dynamic range of the measurements. It only needs to amplify signals up to the maximum frequency offset of interest, say 1 MHz, but it was built to serve this purpose and be useful around the
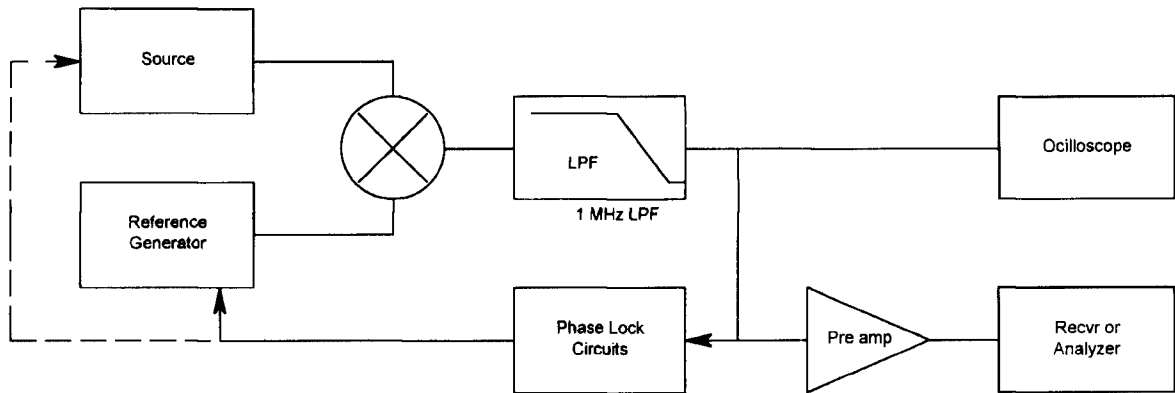

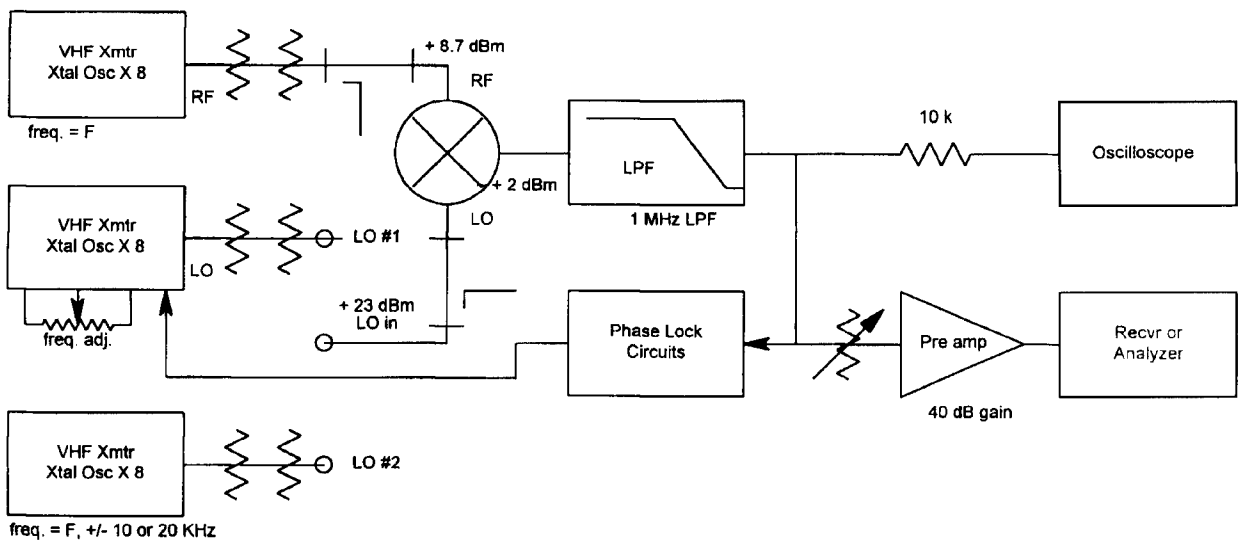Figure 7—Phase-locked carrier-elimination method.

Figure 8—Phase-locked carrier elimination with VHF sources using a high-level mixer, an MCL RAY-1. The dc-coupled modulation input on the transmitter was used to lock the two transmitters with the phase-lock circuit.

lab for other applications as well. The coupling capacitors are adequate to afford amplification from less than 500 Hz to 1300 MHz as shown in the response curve in Figure 12. This circuit is also housed in a die-cast aluminum box and powered with batteries to reduce noise pick-up and feedback. The variable attenuator in Figure 8 needs a range of 80 to 95 dB.

The third source shown in Figure 8 is used only for convenience in calibrating the test setup. It is set to a frequency of 10 or 20 kHz above or below the frequency of the sources under test. This third source is not required if the frequency of the one source under test can be easily changed to generate a calibration beat note.

The dc-coupled oscilloscope displays the quality of the mixer's output signal and, more importantly, shows the dc level and indicates phase lock with minimal offset. This keeps the mixer in the more-accurate, phase-sensitive area of operation near quadrature (within a few dozen millivolts of zero).

Operation is as follows: Turn off the preamp and set the attenuator to 90 dB, or so. The phase-locked loop is opened, or just pushed out of lock, and the sources are moved apart in frequency to provide a beat-note output from the mixer. The beat-note frequency can be anywhere near the offset frequencies of interest because the IF circuitry has a flat frequency response from near dc to the mega-hertz range. Set the beat note at 10 or 20 kHz and adjust for a convenient display or level on the analyzer or receiver. Record the attenuator and measuring-instrument settings and levels.

Alternatively, disconnect the source labeled LO #1 in the diagram (the middle one) and connect LO #2 and adjust for a convenient display. Exchanging the LO sources instead of the RF source causes less potential
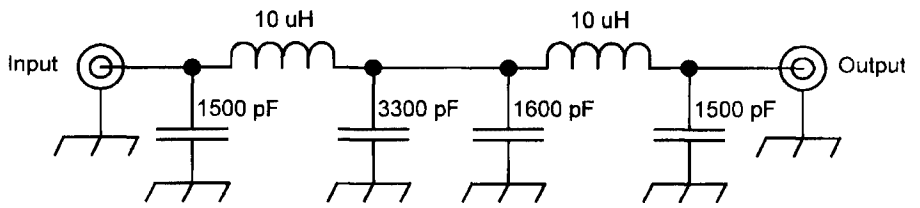

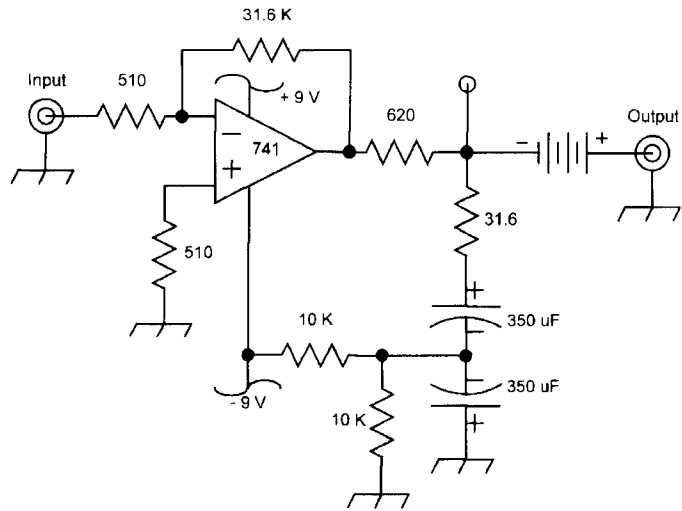Figure 9—1 MHz low-pass filter.

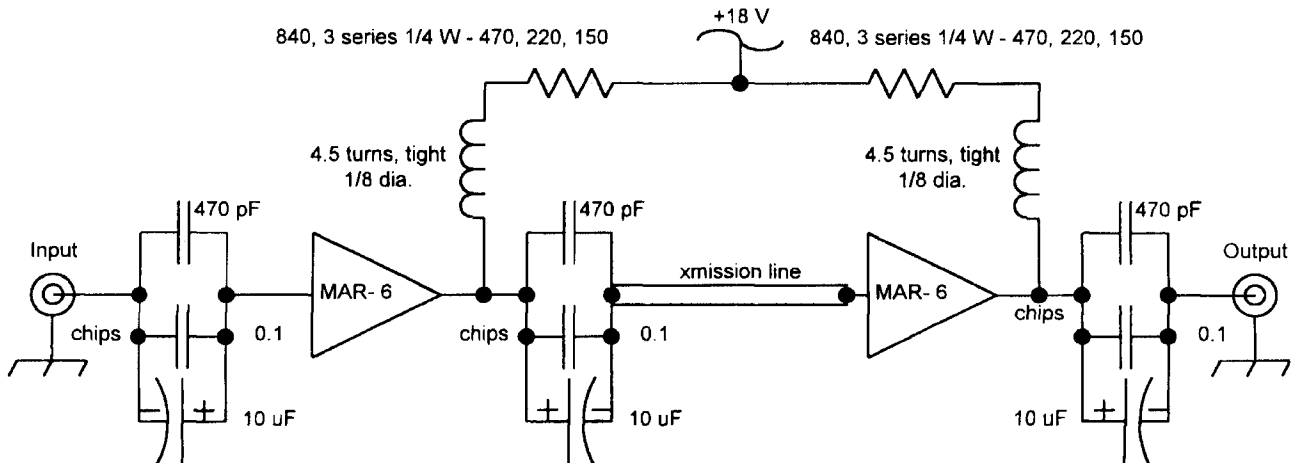Figure 10—Phase lock amplifier.

Figure 11—40 dB gain preamplifier.

variation in amplitude due to small differences in levels.

The power levels have all been previously set for +23 dBm from either LO source and up to +9 dBm on the RF source for the high-level +23 dBm mixer. (Use +17 dBm and +3 dBm for a +17-dBm-level mixer, etc).

Readjust the RF and LO sources to nearly the same frequency by observing the beat note on the oscilloscope (or put the LO #1 back on). When the sources are at or near zero beat, close the loop and observe phase-lock and the dc voltage at the mixer output on the oscilloscope. Remember—during adjustment to near zero beat—to keep the center-frequency bias on the tunable source very close to the same dc level as it will see when the loop is closed. There are different means for doing this depending on the sources in use. In the example under discussion here, the RF source and LO #1 remain close in frequency and the bias on LO #1's tuning input does not change during the calibration procedure. When it is reconnected, the loop locks up immediately, at about 0 V dc on the mixer output.

The calibration and adjustments are usually easier to do, than to describe. They really need to be done only once: to accommodate the gains and losses in the system, which will not change much over time. (Other cases using only two sources, where one is a tunable VCO, will be described later.) In most cases, the locking range is quite large and the dc offset can be lowered to zero after the loop is locked—by a mechanical adjustment on one or both sources. In experiments run with a 9 MHz VFO against a 9 MHz crystal oscillator, the amount of dc offset was up to maybe 75 mV at the mixer IF. This offset does not seem very critical to the measurement accuracy, but I like to "do it right."

Now the sources are locked together and noise measurements can be made down to the noise floor of the receivers in use. Turn on the preamp and switch out the attenuation. Noise at the reference level preset on the beat note (above) would be below the carrier by the amount of attenuation removed. For example: If we have +2 dBm coming out of the low-pass filter on the beat note and −92 dB in the attenuator, we have a −90 dBm CW signal.
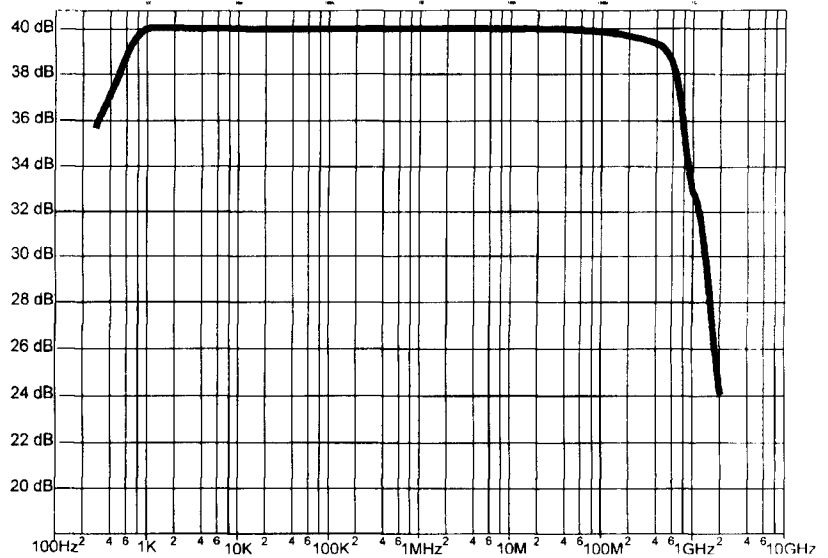

Figure 12—Frequency response of the 40 dB gain preamplifier. It only needs a few hundred kilohertz of bandwidth, but we made it a general purpose amplifier that can be used in other applications.

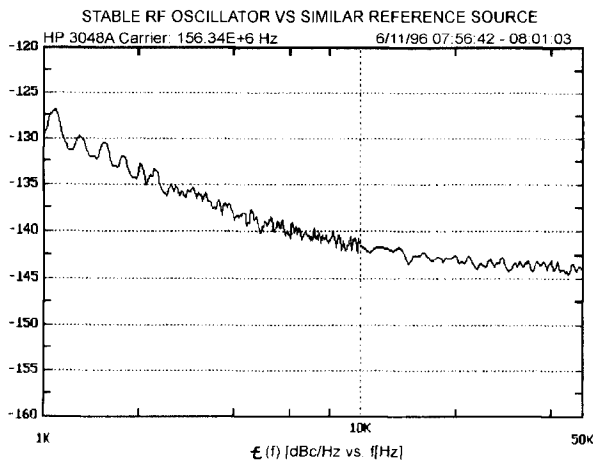Figure 13—A plot of the phase noise of the two VHF transmitters in Figure 8, as measured with the HP 3048A Phase Noise Measuring System.
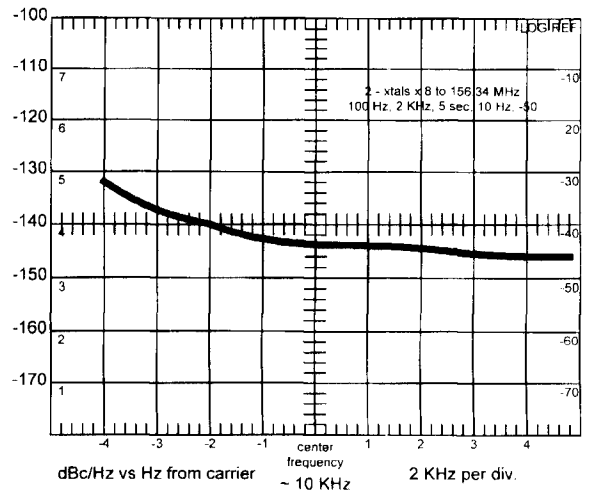

Figure 14—A plot of the phase noise of the two sources of Figures 8 and 13 taken with a spectrum analyzer for the detector. These are fundamental crystal oscillators multiplied by eight to the VHF range. The analyzer settings were 100 Hz BW, 2 kHz/division and 10 Hz video BW.

Then +40 dB gain in the preamp yields –50 dBm at the measuring detector/ display device, the receiver or analyzer. Minus 50 dBm is a good level to use in most cases. Now, if noise displays at the reference level when the attenuation is reduced to zero, the noise is 92 dB below the beat-note carrier reference, or –92 dBc, in whatever bandwidth is being used. Assuming a 1 kHz BW, the noise would be –92 dBc –30 dB = –122 dBc in a 1 Hz BW, or –122 dBc/Hz. If the measuring device records an average noise level, say 20 dB, lower than the reference level, then the phase noise would be –142 dBc/Hz.

Comparison of Data with Different Equipment

Now we can compare data.

Figure 13 is a plot from the HP 3048A Phase Noise Measurement
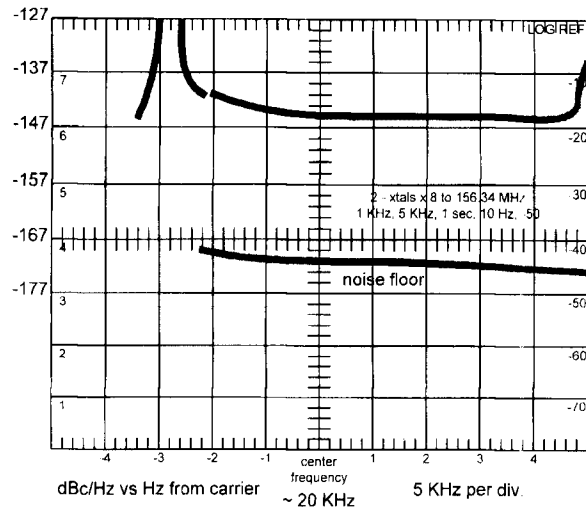
Figure 15—The same as Figure 14 except the analyzer settings were 1 kHz BW, 5 kHz/division and 10 Hz video BW. The "fuzz" in the noise display is only about 2 dB thick with these settings.

Figure 16—A frequency upconverter to allow use of a power meter or communication receiver for the display and measuring device in Figures 7 and 8.

Figure 17—An upconverter suitable for use with a communication receiver for measuring the noise output after the preamplifier.

System for the two sources in Figure 8. How closely did we duplicate this with lower-cost equipment? Figure 14 is a plot of the data taken using an HP 8556A/8552B/141T in the diagram of Figure 8. The numbers are almost the sam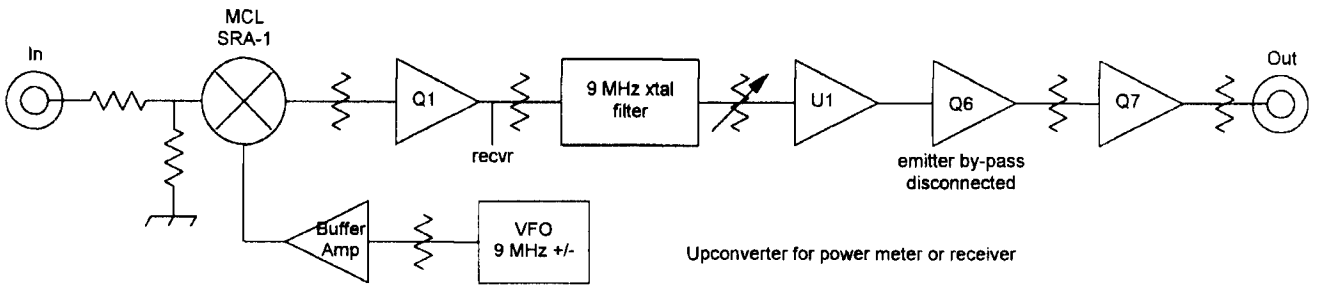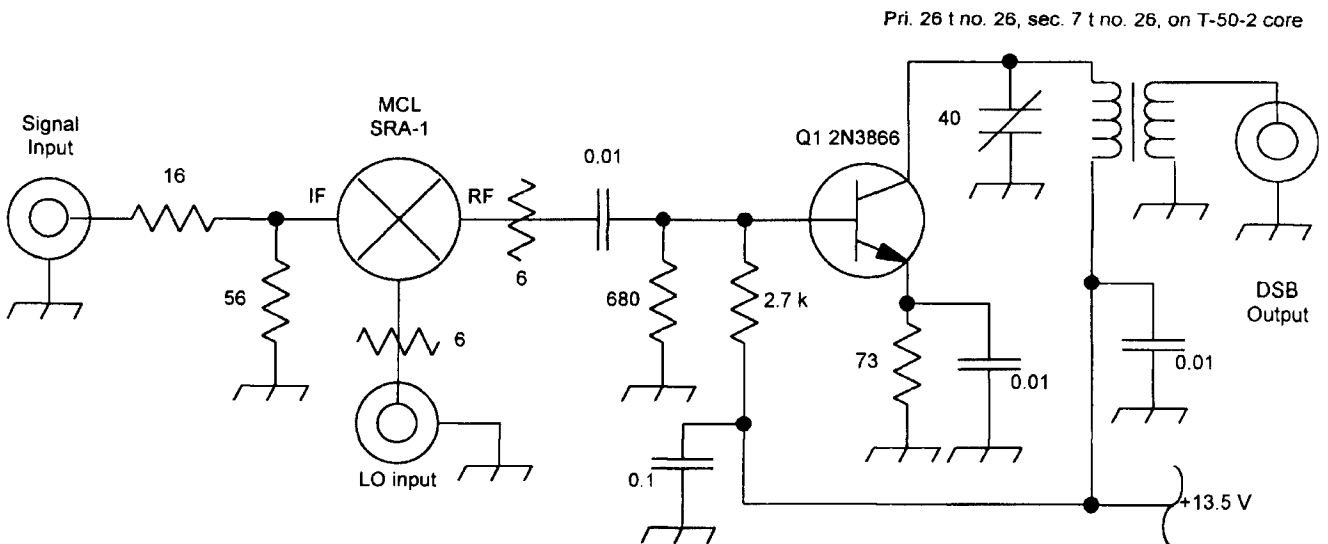e. In Figure 15, the sweep range was increased to observe higher offset frequencies. This was included to make a general observation of test setup techniques. Figure 15 has some spurious responses "blotching-up" the display. As we will see later, the HP 3048A system often records spurious responses (the software has a feature that will plot the data without the spurs). This type of spurious response can usually be eliminated by better isolating the sources from each other—and from everything else—through use of Class-A buffer amplifiers and attenuators after the sources and amplifiers. A short discussion of source isolation and attenuation/ amplification of signals under test is included in the Appendix. However, the data in Figure 15 is still useful from 5 kHz to over 40 kHz on this plot. There is close agreement between the two systems—within a few decibels.

The HP8556A low-frequency analyzer is a good unit that's not too expensive on the surplus market. When used with this phase-lock and preamp system, it can make measurements equivalent to the 3048A. However, a simple standard communication receiver or power meter can detect and display equivalent measurements, with some other circuits added. For example circuits, see the block diagram in Figure 16 and the schematics in Figures 17 and 18. This would bring the cost down significantly, especially if existing equipment is used.

Some communication receivers work down to a few kilohertz but most either don't work below 100 kHz or lack adequate sensitivity. The circuits in Figure 16 are from an AF spectrum analyzer described in an October 1993 *QEX* article.[1] They form an upconverter to place the energy in a range where a power meter or receiver will work better and crystal filters are available. These circuits partially replace the receiver or analyzer block in Figure 7 or 8.

When using a receiver, only the circuits up to the crystal filter are required. The receiver is attached

[1]B.Pontius, NØADL, "Narrowband Spectrum Analysis with High Resolution,"*QEX* Oct 1993, pp 3-12.

following Q1, at the point labeled **RECVR**, preceded by an adjustable attenuator. It can be operated as described previously in the discussion associated with the Figure 3 method.

Again, a calibration beat note is established and attenuators are adjusted for a reference level. The loop is locked, eliminating the carrier, and the noise is measured. If you use a tunable receiver with Figure 16, the VFO is not needed, and a fixed frequency LO can be used. Either sideband can be tuned across but watch out for spurious signals that can overload the receiver, or do worse damage. Find spur-free spots to take data. Leave in as much attenuation as possible, while looking for spurious ranges to stay clear of, then tune to clear spots before reducing attenuation. The appendix has a brief discussion of spurious-reduction techniques.

If using a power meter, the filter and the additional gain in Figure 16 are required. The meter is connected at the **OUT** connector. These circuits replace the receiver or analyzer block of Figure 7 or 8. The 40 dB gain block eliminates the need for the gain stages U1 through Q7 in this particular application. The emitter bypass capacitor on Q6 was disconnected and the overall gain was reduced by about 28 dB to a useful level. The gain from **IN** to **OUT** should be about 58 dB when the variable attenuator between the filter and U1 is at 3 dB (the minimum value). The signal on a beat note at the **IN** connector should be set to about −54 dBm. Then, when the sources are phase-locked, noise power is measured directly.

For example, assume a −5 dBm CW beat note out of the mixer/low-pass filter, a level typically realized using a +17 dBm mixer, like the MCL SRA-1H or JMS-5H. The variable attenuator in Figure 8 would be set to 89 dB for a level into the preamp of −5 −89 = −94 dBm. Add the 40 dB preamp gain and the signal to the **IN** connector is −54 dBm. An overall gain (from **IN** to **OUT**) of 58 dB results in a CW signal to the power meter (at the **OUT** connector) of about +4 dBm, a suitable maximum level for the amplifiers used in the circuits. Now the sources are locked and the attenuator set to zero as before, and the phase noise across the bandwidth of the filter used is displayed on the meter. The meter bounces a lot due to the combination of the noise-signal variations and the damping of the meter, but eyeballing
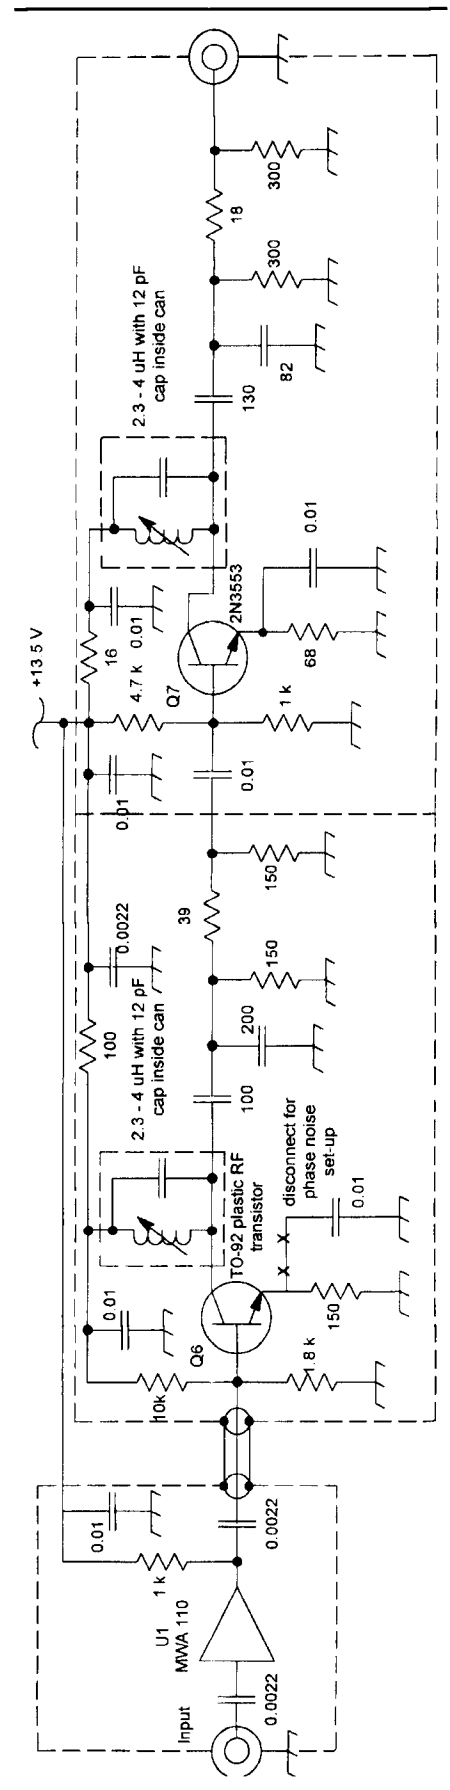
Figure 18—The remainder of the upconverter after the filter in Figure 16.
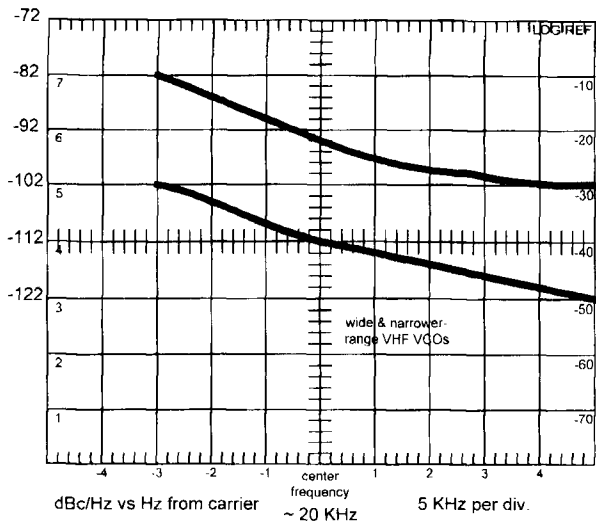
Figure 19—Some additional RF sources: two free-running VCOs. One with wide-tuning, the other with narrower-range tuning. Larger size, higher circuit Qs and narrower tune ranges lead to lower phase-noise levels.
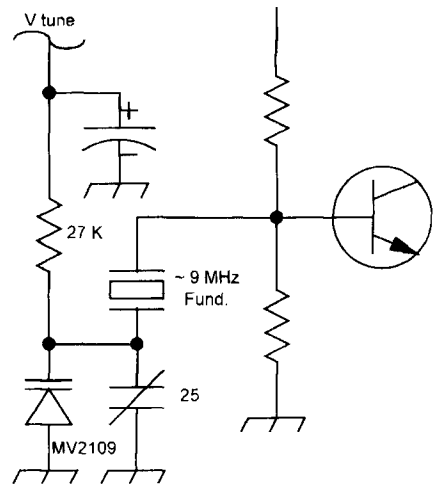
Figure 20—Typical fundamental crystal oscillator tuning-voltage input circuit.
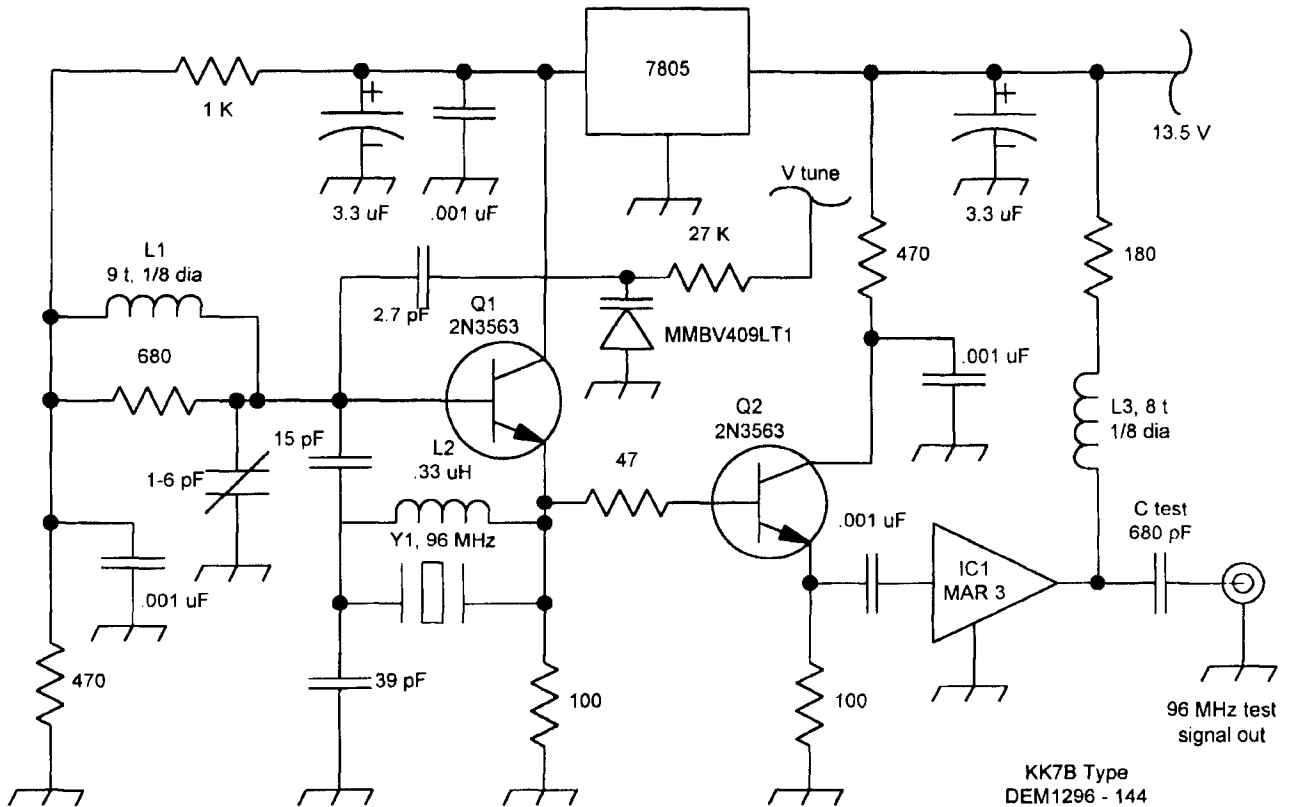
Fig 21—Fifth-overtone oscillator from a KK7B design as supplied by Down East Microwave in the DEM1296-144 LO board. A means of electronic tuning for phase locking has been added at V_{tune}.

the meter and judging an average level usually results in a satisfactorily accurate reading. Switching the power meter up one range often makes for easier readings. The warnings about avoiding spurious signals apply here, as well.

In the case of a +23 dBm mixer, signals would be 6 or 7 dB higher. This requires more attenuation (in the attenuator of Figure 8 or Figure 16) to keep the signal level of the CW beat note below +4 dBm on the power meter (for linear operation of the amplifiers).

In all cases, the input to the mixer in Figure 16 should be kept below −35 dBm.

The phase-noise numbers measured with receivers and power meter agreed very closely, usually within a decibel or two, of those measured with the 3048A and the 8556A analyzer.

Calibration Factors for Absolute Phase-Noise Numbers

Detailed calibration of the various instruments allows us to read absolute phase-noise numbers, but calibration

has almost no effect on comparative measurements between one source and another. Nonetheless, I will briefly summarize the calibration factors.

In general, a spectrum analyzer and a receiver S meter must be calibrated for noise response versus CW response. HP says the 8556A/8552B/141T requires a correction of −4.3 dB when reading noise levels for the phase-locked method and −1.3 dB when using noncorrelated noise-power methods, like the heterodyne/filter method or when measuring noise in general. This
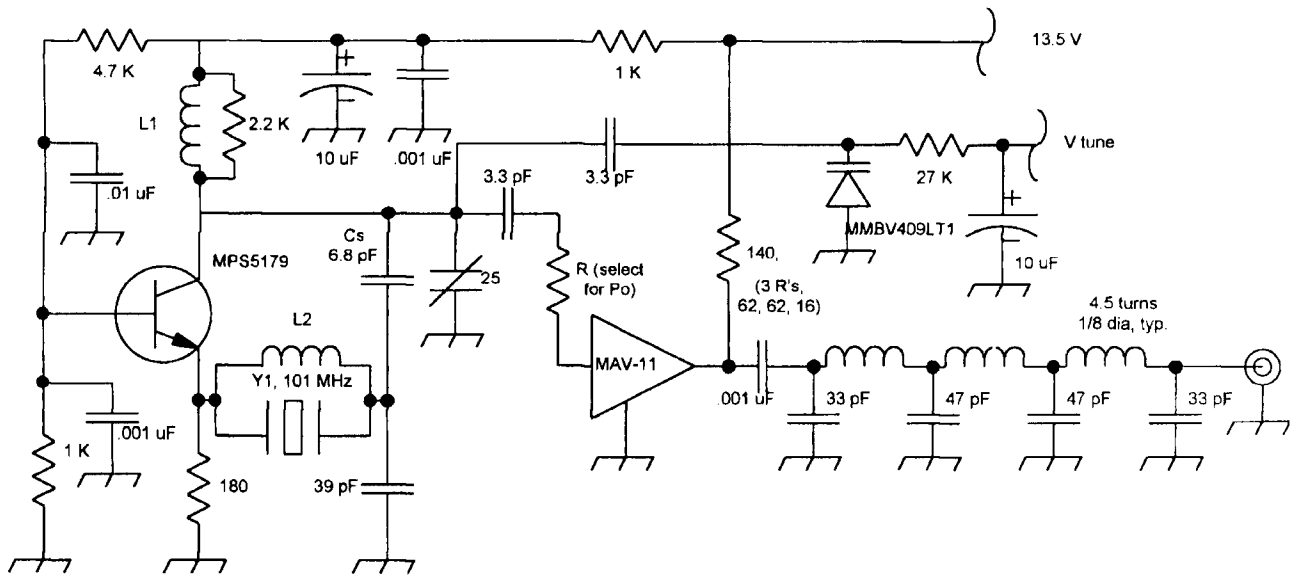
Figure 22—Fifth-overtone oscillator from *The ARRL Handbook*. Cs, 10 pF; L1 9½ turns on 0.10 inch ID, tight wound; L2—15 turns #28 enameled wire on T-25-6 core for 90 to 100 MHz range. Cs, 6.8 pF; L1 7½ turns and L2 13 turns for higher range. A means of electronic tuning for phase locking has been added at V_{tune}.
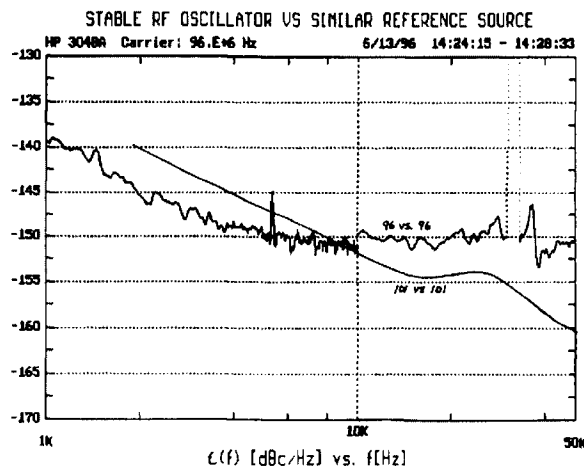
Figure 23—Phase-noise plots of the two fifth-overtone oscillators of Figure 21 (96 MHz) and Figure 22 (101 MHz) run against similar circuits on the HP 3048A Phase Noise Measurement System.
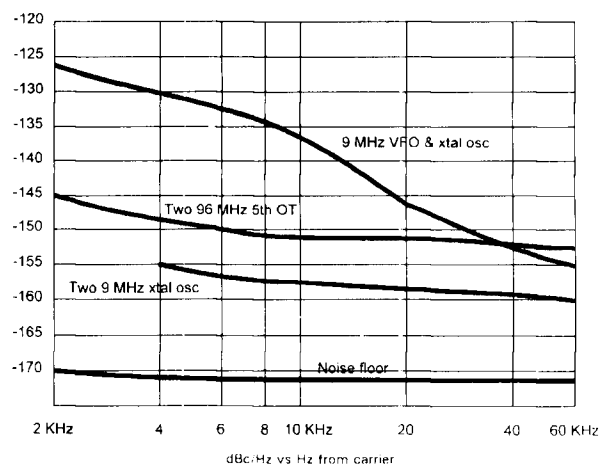
Figure 24—Measured phase noise of a 9 MHz VFO, 96 MHz fifth-overtone oscillator and fundamental crystal oscillators at 9 MHz.

takes into account the instruments' filter responses. (In the phase-locked case, the noise is correlated so the readings are 6 dB higher than for just one source, if the sources are similar. If not phase locked, the readings are 3 dB higher than for one source.) The receiver used must be individually calibrated, for CW versus noise response. I do this with a noise generator output heterodyned up to the 10 MHz region and read the power in a 500 Hz BW with a HP 432A/478A heat-sensitive power meter to get an absolute power number. The receiver or the analyzer can then be compared with the power-meter reading and calibrated.

Two Other Interesting Signal Sources:

Some other interesting sources and their data are shown in Figure 19. This data is shown as displayed on the analyzer, and the same numbers were measured with the receiver and the power meter. The top curve is for a wide-tuning-range (2.3:1) VHF VCO, and the lower curve is for a VHF VCO tuning over a 1.3:1 range. The lower-curve VCO is physically larger, with higher Q components, and with the narrower tuning range provides the lower phase-noise results.

With both VCO sources, the tests were run with the VCO on the RF port

and the VHF crystal transmitter (LO #1 in Figure 8) on the LO port. The LO can be mechanically tuned over a small range and is electronically tuned for phase locking. The loop can be left closed and the VCO tuned off frequency enough to break lock for calibrating on a beat-note tone. After calibration on the tone, the VCO can be slowly brought to the crystal transmitter's frequency, where the loop grabs hold and locks. Once locked, the LO source will follow the VCO around. Near-quadrature operation can be obtained by tuning the VCO or by adjusting the mechanical frequency adjustment on the LO.

Directional couplers were used to monitor frequencies with counters for experimenting with lock ranges.

Electronic Tuning Techniques for Various Sources

Electronic tuning for phase locking can use the dc-coupled FM input on a signal generator with bias at 0 V. A tuning-diode bias can be arranged for other sources' requirements. The circuits described in this article handle a wide range of conditions, but a few examples may be useful.

In the case of the VHF and UHF transmitters and other fundamental crystal oscillators used in these experiments, the dc-coupled modulation input is used as shown in Figure 20. A

9 MHz oscillator is shown there but the tuning part of the circuit can be similar for any fundamental oscillator. The tuning constant can be juggled around to get the desired results but usually only a few-hundred hertz of tuning range is needed to hold lock during measurements, even with wide-tuning-range VCOs.

Actually, it is often easier to lock the more-stable source to the VCO than the other way around. The phase-lock circuit in Figure 10 is capable of much more voltage swing than is usually necessary, so resistive dividers are often added along with the required dc bias voltage for the V_{tune} input. In general, use the least amount of tuning necessary to hold lock during a measurement.

Each application requires its own set of values determined by experimentation. The VHF and UHF transmitter circuits used here have 5 kHz/V tuning constants, referred to the output frequency, while the overtone-crystal circuits are in the range of 40 Hz/V.

The HP 3048A system is very fussy in the areas of tuning constants, linearity and distortion. Phase lock is often not accomplished without experimentation. In general (as discussed in the Appendix), sources should have low harmonic distortion and low inter-modulation distortion. The circuits
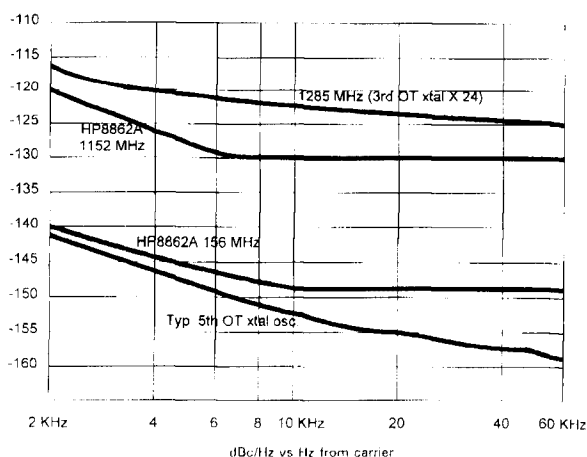
Figure 25—Additional data on 23 cm sources, the HP 8662A and a typical fifth-overtone oscillator for relative comparisons. Phase noise increases 6 dB each time the frequency is doubled. If the ×24 1285 MHz source started at –152 dBc/Hz at the crystal, it would be –124 dBc/Hz after frequency multiplication by 24. Third-overtone crystal oscillators are similar to their fifth-overtone cousins in phase-noise performance.
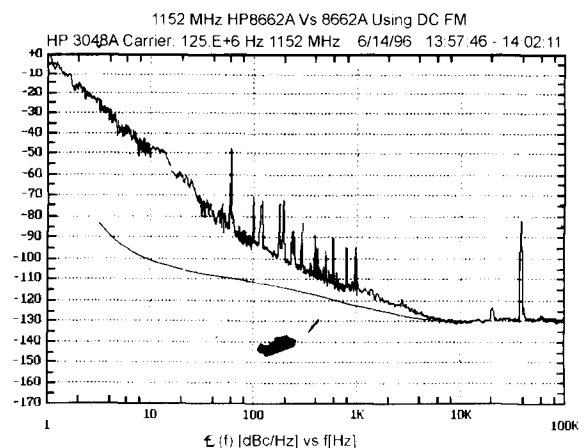
Figure 26—Plot of two HP 8662As phase locked at 1152 MHz with dc FM. The lower line shows phase noise with other locking means as discussed in the text.

described in this article, however, seem to be quite forgiving. The worst effect is generation of spurious signals, which we can usually work around.

Examples of fifth-overtone crystal oscillators with electronic tuning are shown in Figures 21 and 22. Each of these circuits was run against identical but untuned circuits for noise measurements. The 4.5 V built-in bias and the voltage directly out of the phase-lock circuit were adequate to phase lock and perform measurements, even when the crystal oscillators were locked to a VFO or VCO. The small values of tune-voltage coupling capacitance between the diode and the crystal can be adjusted to change the tune constant if necessary. Sources using third-overtone crystal oscillators multiplied by 24 to the 1152 to 1296 MHz range were also set up and tested with similar tuning arrangements. All tuning lines should use shielded cable for connections.

A 9 MHz crystal oscillator was phase locked to a 9 MHz VFO using the circuit of Figure 20 and the direct output of the phase-lock circuit, as with the fifth-overtone units discussed earlier.

Figure 23 is a printout from the 3048A system showing the phase noise on the two fifth-overtone sources. Figure 24 shows data taken using phase lock with the simpler equipment. Figures 24 and 25 provide some additional data for reference purposes.

Figure 26 shows what happens when a moderately wide-band modulation input to a stable oscillator is used. The close-in noise of the HP 8662A is much higher when using the dc FM modulation input than for the unmodulated mode. The generator can be locked through the back panel 10 MHz crystal reference-oscillator input, but only a very small tuning range is available. In this mode, the phase noise is the same as that for the basic generator. The lower curve shows the noise in that case. The same general principal

applies to other sources. Keep the tuning range as small as possible to minimize disturbances to the basic high-Q frequency-determining elements. In most cases, the effects of the tuning circuits disappear beyond a 10 kHz offset from the carrier, as is the case with the 8662.

We will experiment more in the future with close-in measurements (less than 2 kHz from the carrier) and with microwave sources through use

of multigigahertz-range frequency mixers for direct phase locking and down conversion prior to phase locking.

Conclusion

We have shown how low-cost equipment and circuits can be successfully used to accurately measure phase noise. The phase-lock method provides the greatest versatility, sensitivity and accuracy. ▫

Appendix

Spurious, harmonics and intermodulation products can cause measurement problems, so it is best to take precautions that avoid these problems from the beginning. The HP 3048A Phase Noise Measurement System sometimes refuses to lock two signal sources together. The software provides comments regarding the tuning linearity, tuning range, harmonic content or other comments to suggest improvements. The only problems I have seen with the lower-cost systems described in this article are spurious responses affecting a display or pegging a meter. The best ways to reduce spurious responses are to use:

- Filters on the sources such as the low-pass filter on the 101 MHz fifth-overtone oscillator in Figure 22
- Linear Class A isolation/buffer amplifiers
- Attenuation after the amplifier

MCL MAV-11 and ERA-4 are good LO or RF amplifiers up to a +17 dBm output level. I often use specially built amplifiers or something like the Motorola CA 2832C and MHW 593 for more power, up to 200 and 450 MHz, respectively. Amplifiers get expensive at higher frequencies, so you must accept more possibilities for spurs. A diplexer to properly terminate the mixer IF would probably be a good idea. With a high-pass/low-pass combination, the higher, undesired feedthrough energy would be terminated in a load, not reflected back to cause trouble in the mixer.

The signal levels must be strong to enable measurement of low noise values. Use no more gain than necessary in buffer amplifiers, if you have the flexibility. The noise measured from a low-noise source can be increased, if a lot of attenuation is followed by a high-gain amplifier. This happens because the amplifier noise is summed with the source noise. According to the HP 3048A manual, the following approximates the effect of an amplifier: Noise Out = –174 dB + Amplifier Noise Figure – Power into the Amplifier – 3 dB. If the power from the oscillator into the amplifier is high, the effect on phase-noise measure-ments is small. For example, if the power into the amplifier from the oscillator under test is +7 dBm and the amplifier noise figure is 5 dB, there will be no measurable increase in the noise level. When a source level of –21 dBm is raised to +15 dBm by a 36 dB fixed-gain amplifier, however, the thermal noise would exceed the phase noise. Keep signal levels as high as possible and attenuate after the amplifier to minimize additions to phase noise.
—*Bruce Pontius, NØADL*

The Flexible Frequency Generator

Do you need clean VHF LO? This system tunes from 120 to 150 MHz with very little phase noise.

By William Cross, KAØJAD

Theoretical Noise Increase in a PLL Synthesizer

When multiplying frequency upward or using a PLL synthesizer, the phase noise increases in proportion to the square of the multiplying ratio (20 dB/decade). Both the sloped noise (near the carrier) and the flat noise floor are raised. The raised noise extends out to the circuit bandwidth. The loop bandwidth of a PLL synthesizer should be chosen for minimum phase noise. This means the loop bandwidth frequency where the multiplied phase noise of the reference equals the phase noise of the VCO. (The noise of the

[1]Notes appear on page 53.

7100 E. Evans Ave, #411A
Denver, CO 80224

phase detector, divide-by-N counter and op amp should be added to the reference, if they are large enough to be significant.) The synthesizer output phase noise will follow the multiplied circuit, divider and reference noise inside the loop bandwidth, while it follows the VCO phase noise outside the loop bandwidth.[1]

Theoretical Noise Reduction in a Divider

When using division (a reference divider from a crystal oscillator or divide-by-N counter in a PLL synthesizer or a direct digital synthesizer) to reduce a frequency, the slope (near the carrier) noise level reduces 20 dB per decade of frequency ratio. The flat noise floor (further away from the carrier) and lower down behaves differ-

ently. Except for the direct digital synthesizer, the noise floor lowers 20 dB per decade, but it is also raised by sampling at the divider output frequency.[2] There is a divider-input effective noise bandwidth. The divider-input noise floor factor is:

$$\frac{2f_{out} \cdot \text{effective bandwidth}}{f_{in}^2} \qquad \text{(Eq 1)}$$

To determine the noise floor at the divider output, we must multiply this by the input phase-noise floor and add the divider noise. The effective bandwidth at the input of the divider is at least the input frequency. For an effective bandwidth of twice the input frequency, the noise factor is:

$$\frac{4f_{out}}{f_{in}} \qquad \text{(Eq 2)}$$

where f_{out} is less than f_{in}.

This is a smaller reduction than 20 dB/decade. In the case of a PLL synthesizer, the reference noise is at the practical circuit minimum, and dividing doesn't lower the noise level because of the divider's own noise level.
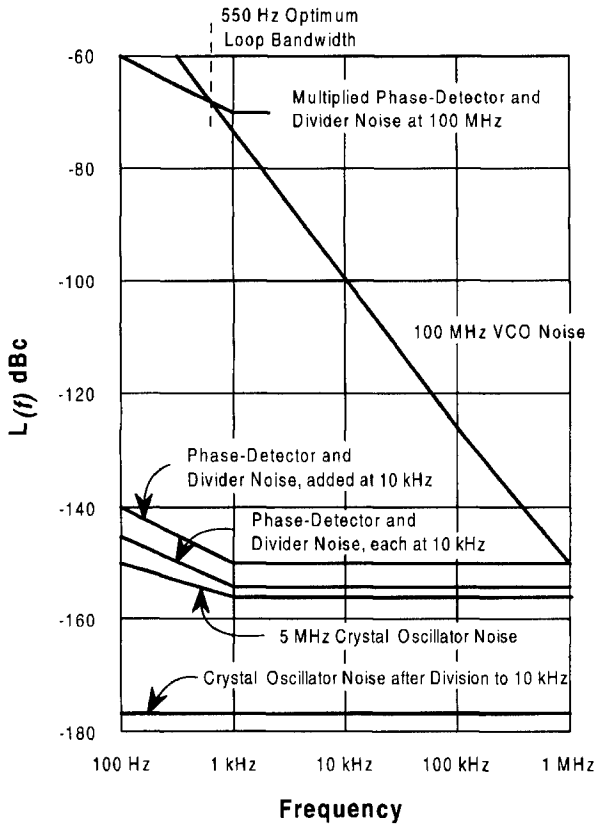
Synthesizer Example

Figure 1 is a noise graph and Figure 2 is a block diagram for a 100 MHz output synthesizer with a 5 MHz crystal oscillator divided to a 10 kHz reference. The left scale shows noise power in a 1 Hz bandwidth, in decibels below the carrier (dBc). The 5 MHz oscillator has a noise floor of –157 dBc. In the reference divider, this noise level is reduced by 21 dB, to –178 dBc, in going from 5 MHz to 10 kHz. This is much less than the circuit noise and is insignificant. The divide-by-N counter, reference divider and phase frequency detector are ECL circuits with low noise. The level 1 kHz from the carrier is –155 dBc for the dividers. Assuming the

Fig 1—Noise chart for a simple synthesizer. (Noise curves for the VCOs were taken from the Vari-L VCO catalog—see Note 4).

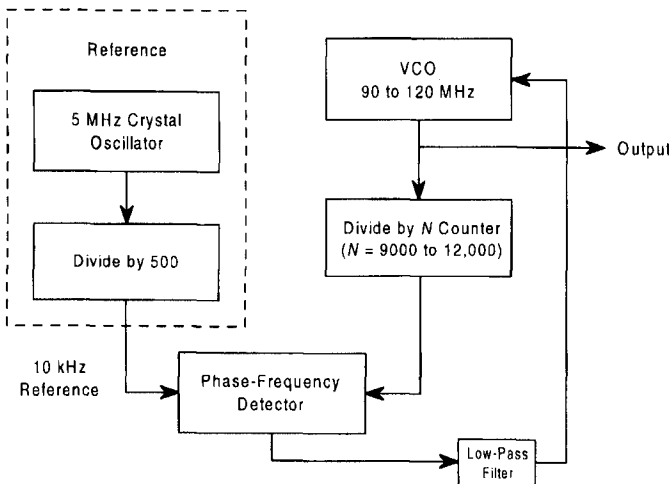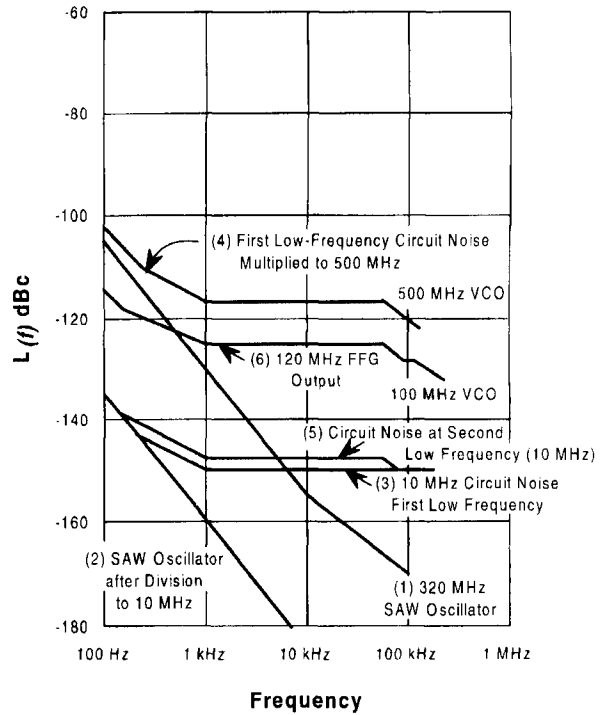

Fig 3—Noise chart for a digital, flexible frequency generator (FFG). (Noise curves for the VCOs were taken from the Vari-L VCO catalog—see Note 4).

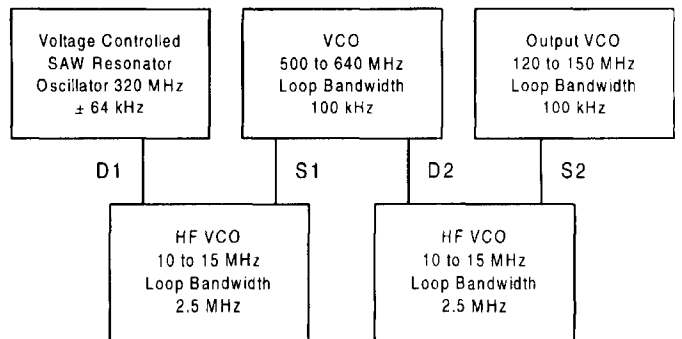Fig 2—Simple synthesizer block diagram.

Fig 4—Frequency diagram for an FFG.

same level in the phase detector gives –150 dBc for all three added up. The signal is then multiplied from 10 kHz to 100 MHz in the PLL. This raises the noise 80 dB, from –150 to –70 dBc. The 100 MHz VCO noise equals this level at 550 Hz, the optimum loop bandwidth.

From 0 to 550 Hz, the synthesizer output noise follows the multiplied phase-detector and divider noise. Above 550 Hz, the synthesizer noise follows the 100 MHz VCO noise curve.

A Flexible Frequency Generator

In the synthesizer, the low 10 kHz reference led to a lot of noise. The flexible frequency generator (FFG) uses a higher frequency than the usual PLL synthesizer reference, while providing densely packed channels. The channels are unevenly spaced but they cover the band, and the maximum gap between channels is small enough so that all frequencies are covered. It took my computer about an hour to test all output frequencies and verify this.

Figures 4 and 5 show an FFG for an up-conversion transceiver with a 120 MHz IF. Its SAW resonator oscillator reference suffers from high thermal drift but is quiet.[3] The tuning range is 400 parts per million. This works out to 128 kHz at the SAW oscillator frequency (320 MHz). Without changing any multiply or divide ratios, the tuning step is 48 kHz at the output of the FFG (120 MHz).

To cover a 120 to 150 MHz output tuning range (a span of 30 MHz) requires 625 channels if they are evenly spaced 48 kHz apart—as in a basic synthesizer. In an FFG the channel spacing is not even, so an FFG uses more channels. In this FFG the first divider produces 11 channels, covering 10 to 15 MHz. The first synthesizer produces 11 channels from each of the first divider's 11 channels. The second divider produces approximately 19 channels from each input channel. The second synthesizer produces three channels from each previous channel. If you multiply the channel numbers together, the total number is 6897, approximately 11 times the number of channels for the basic synthesizer. The basic synthesizer requires a reference of 48 kHz. The FFG's lowest frequency is 10 MHz, much higher than that of the basic synthesizer, so the noise from the dividers and phase detector is much less.

Figure 3 shows SSB phase-noise power for the FFG in dBc versus frequency away from the carrier. The first

point is the 320 MHz SAW oscillator in the lower right corner (1). The SAW oscillator is divided down to 10 MHz

(2). It gets noise the same as the synthesizer in Figure 1 at –150 dBc (3). Then the wave is multiplied up to
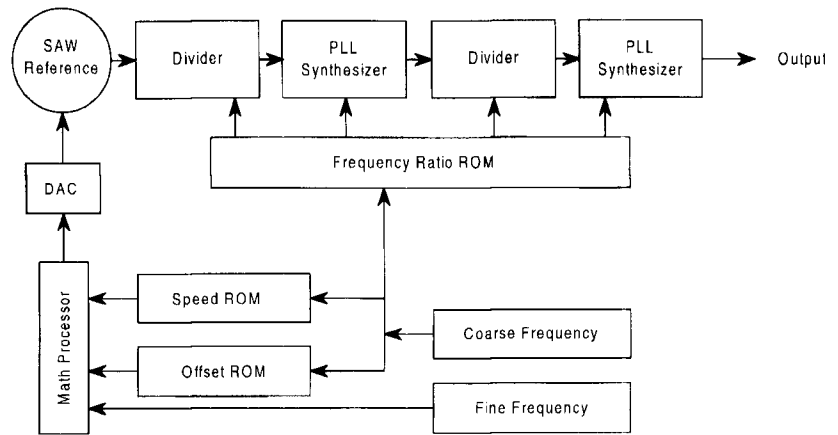

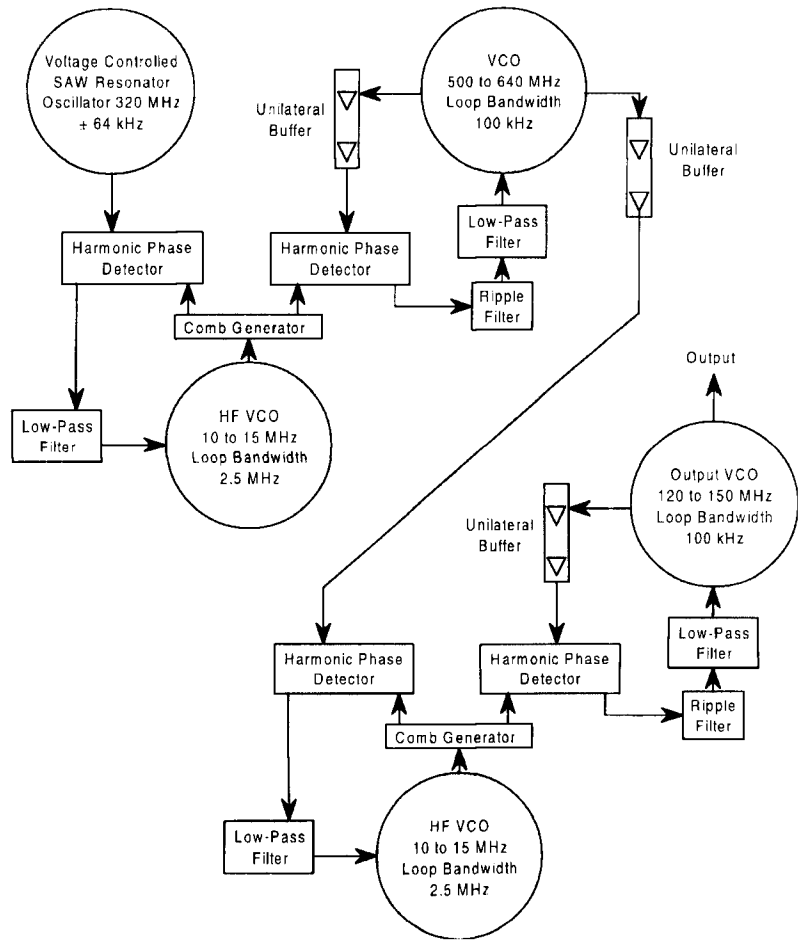Fig 5—Block diagram for a divider-controlled FFG.

Fig 6—Block diagram for an indirectly controlled FFG. (This shows no means of control.)

500 MHz (4). The signal is then divided to 10 MHz at the second low frequency, and gets more noise at –150 dBc for a total noise power 3 dB higher at –147 dBc (5). Finally, the wave is multiplied up to 120 MHz, adding 22 dB for an output noise level of –125 dBc (6).

Figure 6 shows a slightly different approach. For a lower noise level, the circuit could be built with comb generators, VCOs and harmonic phase detectors instead of digital dividers and synthesizers. This would take four loops, with four VCOs plus the SAW oscillator. The frequency diagram is the same as for the divider controlled FFG, but the block diagram is more complex. A single digital divider is replaced by a VCO, a comb generator and a harmonic phase detector. The bandwidth of the divider loops can be very wide, up to 1 MHz or more. This can be done because only the VCO's frequency and harmonics are present in the phase detector. There are no lower-frequency references that need to be filtered out. The other loops must filter out the low frequencies so their bandwidth is limited to 100 kHz, 1% of 10 MHz. Ripple filters are added to the low-pass filters for these two loops, to suppress the 10 to 15 MHz and harmonic energy from the harmonic phase detectors. Digital dividers or a processor can set the loops to the desired frequency.

Acknowledgement

Thanks to Dee and the group for their interest and support in writing this. *William Cross, KAØJAD, has a BSEE from Purdue and has been researching synthesizers for 15 years. He is a customer service engineer and fixes terminals at Trident Computer Services. William has been licensed for about 18 years and now has a General license. He recently got a commercial license.*

Notes

[1]V. F. Kroupa, "Noise Properties of PLL Systems," *IEEE Transactions on Communications* Vol Com-30, No. 10, October 1982.

[2]W. F. Egan, "Modeling Phase Noise in Frequency Dividers," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* Vol 37, No. 4, July 1990.

[3]M. M. Driscoll, "Linear Frequency Tuning of SAW Resonators," *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, Vol 38, No. 4, July 1991.

[4]Vari-L Company Inc, 4895 Peoria St, Denver, CO 80239; tel 303-371-1560, fax 303-371-0845; e-mail **sales@vari-l.com**; URL **http://www.vari-l.com/products.htm.**

RF

By Zack Lau, W1VT

Assembling a 10-Band 50 MHz to 10 GHz Station

As Jon Jones pointed out in his recent *NCJ* column, one of the keys to doing well in ARRL VHF contests is to operate as many bands as possible.[1] (See Fig 1.) I've done pretty well with my station; it's normally one of the top QRP stations in the New England area. To encourage more people to activate stations on the higher bands, I'm describing my particular station, which was initially assembled in the late 1980s.

The most important aspect is site location: When the access road is open, I

[1]Notes appear on page 58.

225 Main St
Newington, CT 06111-1494
e-mail **zlau@arrl.org**

operate from Mt Equinox, Vermont, the highest spot in southern Vermont. (See Fig 2.) It doesn't hurt that it is located in a relatively rare grid square, FN33.

Operating from the highest spot with access via paved roads is also a disadvantage: Just about everything from amateur repeaters to FM broadcast transmitters are also located there. This means a lot of interference to deal with—both in and out of band. A more recent problem is foliage blockage—the current political climate makes trimming the trees quite difficult in Vermont. The site is far enough north that it is only useable during warm weather, from June through October.

The foliage problem is especially bad for casual operators who don't scout their sites beforehand—they are virtually asking for foggy weather. Only ex-

perienced contesters seem able to cope with 20 foot visibility and significant tree blockage on the microwave bands.

Interestingly, microwave system improvements are quite noticeable from this location. While stations in Albany, New York, are within my line of sight and incredibly loud, communication with Boston, Rochester and Philadelphia involve tough, non-line-of-sight paths. Thus, degradations or improvements in conditions or my station are quite noticeable. Being able to work these areas at any time would certainly help my score—I could then save hours of better propagation for working more-distant multipliers.

My antennas are generally on the small side—I often find myself erecting them under less-than-ideal weather conditions. High winds and

Fig 1 — Schematic

4 Element Yagi — 3 W — ICOM IC-202S — 2 M SSB Liaison

4 Element Yagi — 9 W — R2/T2 Based XCVR — 6 M SSB/CW

6 Element Yagi — 9 W — 9 W Amp — 1.5 W — ICOM IC-3AT — 223.5 FM

10 Element Yagi — 4 W — Radio Shack HTX-404 — 446 FM

2 foot Dish — 3 W — 10 GHz XVTR — 10 GHz

Ground Plane — ICOM IC-20L Amp — 9.7 W

Sloping Dipole Hanging from Tree — 9.5 W — 50 MHz XVTR

Mast 1

Band: 2 M, 6 M

IF:
6 M FM/SSB/CW
2 M FM
903 - 10368 MHz FM/SSB/CW

Yaesu FT-290RII — 3 W

Microwave

22 Element Yagi — 9 W — 903 MHz XVTR — 903

38 Element Loop Yagi — 9 W — 1296 MHz XVTR — 1296

42 Element Loop Yagi — 9 W — 9 W Amp — 2 W — 2304 MHz XVTR — 2304

2 foot Dish — 5 W — 3456 MHz XVTR — 3456

2 foot Dish — 1.5 W — 5760 MHz XVTR — 5760

FO-22 22 Element Yagi — 9 W — 9 W Amp — 2 W — ICOM IC-402 — 432 MHz SSB/CW

Mast 2

IF Box — 10 mW — Band

Mast 3

KB6KQ Loop — Omni

Cushcraft 13B2 13 Element Yagi — Beam — 9 W — FET Amp — 3 W — ICOM IC-202 — 2 M SSB/CW

Homebrew FO-12 12 Element Yagi — 7 W — Homebrew 222 MHz XCVR — 222 MHz SSB/CW

Fig 1—A schematic of W1VT's portable VHF, UHF and microwave installation. For details of the loop on Mast 3, contact Norman Pederson, KB6KQ, 70 Arrowhead Dr, Carson City, NV89706; e-mail *kb6kqnorm@aol.com*.

wet weather aren't unusual on New England mountaintops during the warmer months. I've had a UT-141 jumper cable unsolder itself at both ends after being exposed to high winds for a couple of hours. I've since gone to UT-250 semi-rigid cable with compression-clamp connectors for greater reliability on 10 GHz.

Interestingly, notable VHF types didn't think too highly of my idea of running lots of bands at the 1989 Eastern VHF/UHF conference. Back then, the goal was to run eight bands, from 50 to 3456 MHz. True, there weren't that many stations to work back then, and the contacts were often challenging, but the multipliers certainly helped boost my score.

I decided that transverters for the four microwave bands would fit easily in a half-size rack. Instead of 19-inch-wide panels, I'd use 9.5×1.75 inch panels in front of 1.7×8×12 inch boxes. I'd
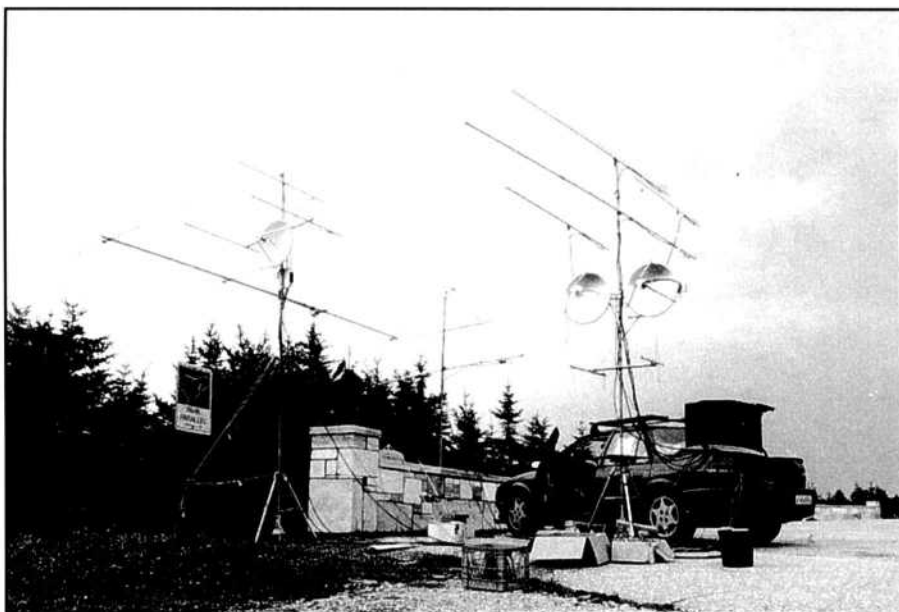
Fig 2—W1VT's portable VHF, UHF and microwave installation. The leftmost mast is Mast 1; the rightmost is Mast 2. Mast 3 is in the center, behind the stone wall.

do it differently today. I'd design all the transverters for either mast mounting or mounting in a rack. This would better accommodate sites with foliage problems. Rack mounting certainly has its advantages. I have a nice, neat package with the IF-switching and dc-power connections more or less permanently hooked up. This significantly reduces wear and tear on the cables, which reduces repair work required to keep the gear working.

I also decided to use 2 meters as the intermediate frequency for reliability. Two meter SSB/CW radios that operate off batteries are easy to find and relatively affordable. Most HF radios are power hogs; that makes them unsuitable for extended battery-powered operation. One disadvantage is the potential for IF interference. It is a significant problem with a 144 MHz IF and a multi-op station a few-hundred feet away. Moving the IF to 145 MHz neatly solves the problem—I've not had problems even with WB1GQR located nearby. They try to call CQ on 2 meters for almost the entire contest, with just a few breaks for sleep and operator fatigue. There doesn't appear to be much of a problem with a 142 MHz IF, but I rarely need to use the backup gear. I use 142 MHz with 222 and 432 MHz transverters, though removing spurious mixing products from the latter transverter required machining a rather large interdigital filter. The 8×8×1.5 inch filter allows tight filtration without excessive loss. 222 MHz performance is acceptable with commercial helical filters. On the plus side, I have a complete set of nine transverters with 2 meter IFs to cover 50 MHz to 10 GHz.

If you bring spares and can easily swap transverters, it makes sense to use slightly different IFs. This may allow you to swap transverters in case some local interference covers an important frequency. Most stations like to maintain a particular frequency offset from the band edge, such as 903.105, 1296.105, 2304.105 etc. This example is intentionally chosen to be too close to the calling frequency—in practice one might expect interference from stations calling CQ (it's not likely to be someone's favorite offset). The fixed offset makes it easier to set up schedules.

My preferred IF radio is currently the Yaesu FT-290RII, with its digital frequency display that allows the frequency to be set precisely. I've often heard WA1MBA on 5760 MHz without needing to tune around for him—sure makes contacts a lot quicker. The

analog system of the ICOM 202 is actually easier to tune, so I preferred that initially. As I got better at estimating where stations show up, the more precise frequency readout became more useful.

I briefly explored using 432 MHz as an IF, but cheap SSB/CW radios are too scarce. While it's rather attractive for 10 GHz and 5760 MHz designs, it would have created too many weak links in my station. Keeping the transverter IFs on 2 meters increases flexibility and reliability. Should anyone ask, I can try NBFM on any band from 6 meters to 10 GHz. I'm not locked into just SSB/CW.

I've been using an ICOM IC-202 or IC-202S and a 22 V FET power amplifier for 2M SSB/CW work. These are some of the cleaner transmitters I've seen, with regard to IMD.

CW key clicks were a problem until I modified the keying circuit to properly shape the signal. Fig 3 shows the schematic of the circuit I used. It increases the rise and fall times to 1.2 and 3 ms, respectively. While 3 ms on and off would be better, this modification can be added to a dedicated external keying unit, so one needn't modify the rig. Without the modification, rise and fall times were rather sharp.

One disadvantage of this portable radio is the imprecise frequency readout. It really needs a digital display, but I've not done this modification.

I use three antennas: a Cushcraft 13B2, a horizontally polarized omnidirectional loop and a four-element homebrew Yagi. I've also tried using two rigs, so I can listen to people making microwave schedules while making contacts on 2 meters. On FM, I use the FT-290RII, an IC-20L amplifier and a simple ground-plane vertical antenna.

I've been using a version of Rick Campbell's R2/T2 on 6 meters. Because it uses direct conversion, it needs only one high-quality RF bandpass filter ahead of the mixer—right after the transmit/receive switch. Superheterodyne receivers need two—one before and another after the RF preamplifier—for use in strong RF environments. The first keeps garbage out of the preamplifier, and the second improves image rejection. To gain about 10 dB more dynamic range, I use Mini-Circuits TAK-1H mixers,[2] which are pin compatible with the SBL-1s that are used on the board available from Bill Kelsey.[3] Initially, I set it up with a switch selectable preamplifier and variable attenuator, but found I really didn't need the extra dynamic
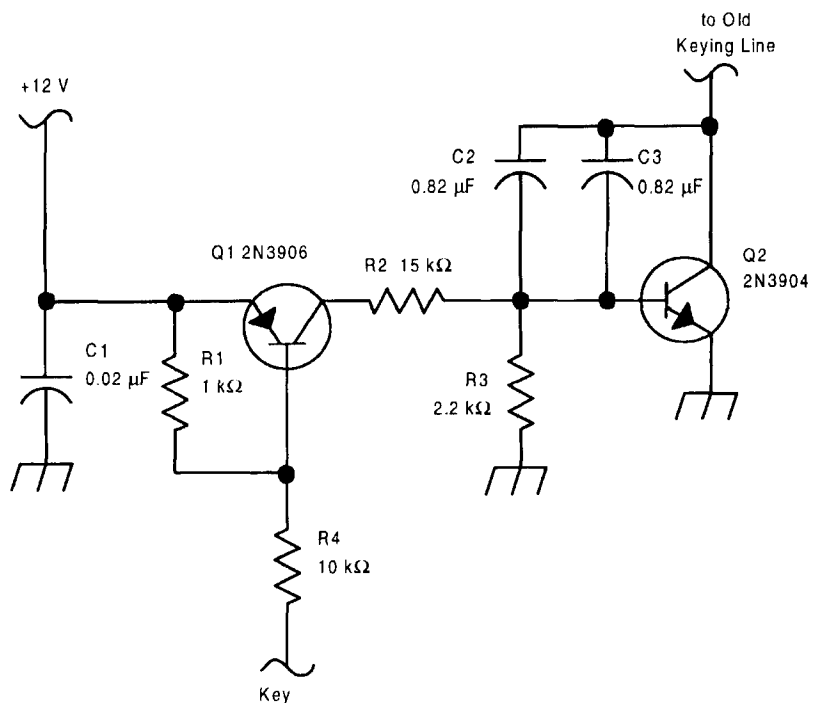

Fig 3—CW wave shaper for the ICOM IC-202.
C2, C3—0.82 mF high-quality capacitors for low-frequency timing circuits (Digi-Key[14] P4674-ND).

range. I took out the attenuator and switch to reduce the chance of operator error—it is too easy to leave switches in the wrong place late at night. The VFO is just an 8 MHz VFO mixed with a 42 MHz crystal oscillator—probably not an ideal choice of frequencies, but I had some high-quality 42 MHz crystals from a previous project.

The homebrew 6 meter transceiver works well during band openings, but I haven't spotted band openings as effectively with it as when I use an ICOM IC-502 with less dynamic range. Of course, the portable 502 was never meant to be a world-class receiver. It's a fun little go-anywhere rig for working occasional E-skip openings with its built-in whip antenna. A switch that significantly degrades dynamic range of a receiving system may have advantages. A low-dynamic-range system often generates many spurious responses that fill the band with ghost images and help you notice when the band is open. Once the band is open, though, those images make it difficult to find weaker stations in distant grids. Maybe that's the secret behind those cascaded preamps that make the S-meter read "9" on band noise. They are only used when the band appears to be dead, and are intelligently switched off-line as required.

The four-element 6 meter antenna works well for troposcatter, although it's more difficult to set up than the little two-element Yagi I've used previously.[4, 5] The big Yagi can take half an hour to set up, while it takes just minutes to add the element tips to the fully assembled yet transportable center section of the little Yagi. I like to have two Yagis during the June contest. I try to have one pointed south towards Florida and the other pointed west to the Midwest. Of course, I'd really like another pointed east towards Europe …

My 6 meter transverter is the one published in the September 1995 *QEX* and ARRL's *QRP Power*[6] book. Just like superhetrodyne receivers, it has a pair of bandpass filters for use in RF-intensive environments. While it doesn't offer the dynamic range of the R2/T2 transceiver, it's a useful spare that allows me to pick up a few FM contacts. I figure that if people want to work me in a contest on FM, why not accommodate them?

Similarly, I added an IF box to my system—this allows me to work W2SZ/1 microwave rover stations with a 5.595 MHz offset. They use 5.595 MHz SSB receivers hooked up to diode mixers with the crystal con-

trolled transmitter as the local oscillator—quite similar to Gunnplexer operation. One major difference is that Gunnplexer operation is mostly voice, while these CW transmitters have no provisions for voice transmission. Doug Sharp said I was the first one to ever bother to build equipment suitable for working their rover stations, despite his offers to provide technical details.[7]

An important aspect of my station is the ease of repair—it is relatively easy to field service my station by swapping out a box or antenna. I also have a spare IF box, in case the first one fails. The transverters can be field-modified to handle three watts of RF from the IF radio, should that be necessary. While there is a lot of duplicated hardware—for instance each transverter box has its own sequencer, this also has its advantages. I think you learn a lot more by designing several versions of something, if you incorporate improvements in each version. Then you can go back and scrap the initial version that didn't quite work as well as you had hoped. Of course, if you are making copies of a well-tested design, building half a dozen of something helps a great deal in meeting the minimum-order requirements of some mail-order companies.

The ease of swapping hardware also helps me find out what works on the mountaintop. I quickly learned that my 1296 MHz Rick Campbell based no-tune transverter really doesn't work as well as an older transverter that uses a pair of GaAs FET preamps feeding an interdigital filter in the front end.[8] I significantly improved receiver performance by reducing transverter gain and adding a bandpass filter ahead of the preamplifier, though it actually measures a little worse on the NF meter.

I find that it's very important to bandpass filter receivers on all bands through 1296 MHz. I use some sort of filter ahead of the preamplifier. It doesn't appear to be a problem on 2304 MHz, but perhaps a 2304 MHz loop Yagi doesn't respond well at the interfering frequencies. I don't expect problems on 3456 and higher because the horn dish feeds act as very good high-pass filters using the waveguide-below-cutoff effect. Fortunately, I don't think there are any strong transmitters above 2 GHz permanently located on the mountain.

Some issues are more difficult to decide. For instance, I tried a vertically stacked pair of 1296 MHz 6-foot-boom loop Yagis mounted on one side of an H

frame. These are similar to those used at W1XX/3 in FN00, but half the size. The array was certainly easier to point and didn't lose any gain, but I decided it is better to have an antenna that clears the trees in all directions. Thus, I use a single 38-element Yagi, the highest antenna on the microwave mast (see Fig 2). Pointing becomes a lot easier after going to the same spot over a dozen times—I know where most of the people are located. The stack might be useful for CQing, but this hasn't proven to be an effective technique for me.

Similarly, deciding what antennas to group together is also a challenge. I've always put the microwave antennas closest to the car, since this allows the shortest feed lines. The trees get just a little taller each year, so I've been forced to figure out how to get the microwave antennas higher. Thus, the H frame isn't quite as useable as such—the lower rungs of the frame aren't high enough to attach microwave antennas. This is why the 2 foot, 10 GHz dish ended up on the FM mast. There wasn't any place for it on the microwave mast that still clears the trees. In practice, locating 10 GHz separately with the liaison antenna seems to work out well. Many 10 GHz operators are specialists who don't use all of the microwave bands. Those contacts need only the liaison and 10 GHz station. On the other hand, there is usually plenty of time to get the 10 GHz antenna into position when working a station on lots of bands.

I strongly recommend standardizing on antenna and power connectors wherever possible, although mistakes can result. VE3ASO once swapped his 903 and 1296 Yagis, which resulted in challenging 903 contacts. I still managed to work him on both bands for VUCC. First, with weak signals on 903, then with the correct 1296 MHz antenna later in the contest. Incorrectly swapping cables is much tougher if you have a whole bunch of connectors that won't mate. I think you need the ability to improvise for forgotten and broken hardware. Radio Shack sells cable labels that work just fine. For last minute corrections, I use an indelible marker to prevent mistakes.

You might even hardwire pigtails to your gear, instead of relying on adapters that might be forgotten. Some ICOM portables use an obscure three-pin connector style. I hard-wired "ARRL standard" power connectors in parallel with those on my IC-402.[9] Some of my transverters and transceivers include five-way binding posts.

If you space these 0.75 inches apart, they accept two-conductor male banana plugs. Some sort of reverse-polarity protection is essential when using binding posts; it is much too easy to reverse the connections.

Even if you do forget something important, like a 3456 MHz antenna, it is sometimes possible to improvise. I remember telling Ed, K1TR what length of wire to stick into the N connector to use as a whip antenna. Despite a 191-km path, we made a rather easy voice contact with just 80 mW! I like to have a supply of #14 or #12 insulated, solid-copper house wire for a variety of temporary mechanical applications, like holding together an improvised H frame made out of Radio Shack antenna masts. (The usual pieces were left at home, 256 km away.)

My 222 MHz station is quite similar to my 6 meter setup. I use a high dynamic range, single-conversion transceiver and have a spare 2 meter IF transverter. The no-tune 222 MHz microstrip filters I designed aren't quite narrow enough to adequately reject the high-order mixing products that result when 80 and 142 MHz combine. I don't recommend that design approach. Use the more costly helical filters, instead. Dedicated homebrew transceivers offer superior dynamic range if properly designed; transverters are typically limited by the dynamic range of the IF radios. Even top-of-the-line HF radios often have insufficient dynamic range to effectively use a really high-dynamic-range converter, although such a combination is often superior to a multimode radio that isn't optimized for excellent receive performance.

The 222 MHz single-conversion superheterodyne uses a Mini-Circuits TAK-1H mixer for good dynamic range. The preamp's UT-141 semi-rigid coax input circuit seems to have adequate selectivity when combined with a low-pass filter. Therefore, the narrow 1 MHz helical bandpass filter is located after the preamplifier.[10] The FM station is just an ICOM IC-3AT and a brick amplifier feeding a six-element Yagi.[11]

It's tough to decide what antenna gain is appropriate. Steve, K1FO published an excellent little 10-element 432 MHz Yagi that I've lent to rovers. They were quite pleased by its performance.[12] When I have it on a switch

with an FO-22, however, I hardly ever use it.[13] I don't find the wider pattern to be an advantage. I do use a scaled version of the 10-element Yagi on 446 MHz FM. The *Scale.exe* program that comes with *The ARRL Antenna Book* does an excellent job of revising the element lengths.

On the other hand, I've found a 10-foot-boom 903 MHz loop Yagi to be mechanically unwieldy when mounted on my H frame, so I opted for a scaled FO-22. I've also used smaller Yagis, but the ones I built were sensitive to rain. There was a noticeable increase in signal strengths if I shook the water off them. The situation on 2304 MHz is similar. I can get a little more gain with a dish antenna, but not enough to justify the extra wind loading. Things are different on 3456 MHz: I notice a significant improvement when using a two-foot dish fed with a simple circular horn feed and another boost in performance when I added scalar rings.

For the ultimate in microwave performance, mast mount the RF hardware, as I've done on 6 and 10 GHz. This minimizes the losses involved in connecting the transverter to the antenna, which can be several decibels or more, otherwise. It is a bit more work to design a transverter to accommodate this, but the flexibility offered by such a unit may make it worthwhile. Stations for the lower bands would also benefit from this approach, particularly if they require long feed lines.

10 GHz WBFM operation is particularly difficult from Mt Equinox. You won't clear the trees in all directions unless you put up a tower. With the gear used by most WBFM operators, a single tree can block even the shorter 50 mile paths available from this location. Thus, people have hiked up tree covered summits and not yet lost contacts because they were unable to find a clear path through the trees. Hiking is very weather dependent in New England. Sometimes I don't know about potential 10 GHz contacts until after the contest starts. Then, it's a little late to move a mast-mounted transverter to a better location! My current compromise is a bulky hand-held unit with a pair of 20 dBi horns. To run a schedule, I must run out to a known good spot; this solution is not ideal, but it's better than being stuck with the antenna hiding behind a tree.

Notes

[1]Jones, Jon, NØJK, "VHF/UHF Contesting," *National Contest Journal*, Jan/Feb 1998, pp 18-19.

[2]Mini-Circuits Labs, PO Box 350166, Brooklyn, New York 11235; tel 718-934-4500; URL **www.minicircuits.com**.

[3]Kanga Kits, Bill Kelsey, N8ET, 3521 Spring Lake Dr, Findlay, OH 45840; URL **http://www.bright.net/~kanga/kanga/**

[4]Lau, Zack, "Rover Antennas," *Proceedings of the 19th Eastern VHF/UHF Conference* (Newington: ARRL, 1993; Order No. 4602), pp 50-57. ARRL publications are available from your local ARRL dealer or directly from ARRL. Mail orders to Pub Sales Dept, ARRL, 225 Main St, Newington, CT 06111-1494. You can call us toll-free at tel 888-277-5289; fax your order to 860-594-0303; or send e-mail to **pubsales@arrl.org**. Check out the full ARRL publications line on the World Wide Web at **http://www.arrl.org/catalog**.

[5]Lau, Zack, W1VT, "Homebrewing a 6 Meter Yagi," RF, *QEX*, Jan/Feb 1998, pp 52-57.

[6]Kleinman, Joel, N1BKE, and Lau, Zack, W1VT, editors, *QRP Power*, (Newington: ARRL, 1996, Order No. 5617). See Note 4 for ordering information.

[7]Sharp, Doug, WB2KMY, "VHF Contesting from Mount Greylock with W2SZ/1," *Proceedings of the 38th Annual West Coast VHF/UHF Conference* (Newington: ARRL, 1993; Order No.4327). See Note 4 for ordering information.

[8]Hinsaw, Jerry, N6JH, and Monemzadeh, Sahrokh, "Computer-aided Interdigital Bandpass Filter Design," *Ham Radio*, Jan 1985, pp 12-26.

[9]Danzer, Paul, N1II, "The ARRL-Recommended 12-V Power Connector," *The ARRL Handbook for Radio Amateurs*, (Newington: ARRL, 1997, 75th ed.; Order No. 1786) p 22.6. See Note 4 for ordering information.

[10]Lau, Zack, "A Collection of VHF Filters," *Proceedings of the 20th Eastern VHF/UHF Conference* (Newington: ARRL, 1994; Order No. 4858) pp 109-114. See Note 4 for ordering information.

[11]Lau, Zack, "Rover Antennas," *Proceedings of the 19th Eastern VHF/UHF Conference* (Newington: ARRL, 1993; Order No. 4602) pp 50-57. See Note 4 for ordering information.

[12]Polishen, Steve, K1FO, "Rear-Mount Yagi Arrays for 432-MHz EME: Solving the EME Polarization Problem," *ARRL Antenna Compendium* Vol 3 (Newington: ARRL, 1992; Order No. 4017) pp 79-98. See Note 4 for ordering information.

[13]Polishen, Steve, K1FO, "A High-Performance 432-MHz Yagi," *ARRL Antenna book* (Newington: ARRL, 1997, 18th ed.; Order No. 6133), pp 18-21 to 18-24. See Note 4 for ordering information.

[14]Digi-Key Corp, 701 Brooks Ave S, PO Box 677, Thief River Falls, MN 56701-0677; tel 800-344-4539 (800-DIGI-KEY), fax 218-681-3380; Web site **http://www.digikey.com/**.

Upcoming Technical Conferences

16th Space Symposium and AMSAT Annual Meeting— Call for Papers

AMSAT's 16th Annual Meeting and Space Symposium will be held October 16-18, 1998, at the Park Inn International in Vicksburg, Mississippi. This is the first call to authors who want to present papers at the Symposium and have them printed in the *Proceedings*. The subject matter of the papers should be topics of interest to the Amateur Radio satellite service. Key dates in calling for and submitting papers are:

• June 1, 1998: Final due date for one-page abstracts.

• June 15, 1998: Authors will be notified by e-mail or postal mail whether or not their paper has been accepted. (Postal notices will arrive shortly after this date.)

• August 15, 1998: Final due date for camera-ready copy of accepted papers. Papers will be edited only superficially; most will be printed as submitted. Authors should provide an electronic file (preferably in any version of *Word* or *WordPerfect*) for possible last-minute corrections if editing spots a major disaster.

Send abstracts to **w5xx@magnolia .net** or Malcolm Keown, W5XX, 14 Lake Circle Dr, Vicksburg, MS 39180. If you send an abstract by e-mail, it's wise to send a follow-up copy by postal-mail; things can get lost in cyberspace.

The *Proceedings* of the Symposium (printed by the ARRL) will be available at and after the meeting. If you do not wish to present a paper, but have a topic of interest, please submit the topic. Perhaps we can arrange for a presentation.

Information regarding Vicksburg-area attractions and details of arrange-ments for the 16th Space Symposium and AMSAT Annual Meeting can be found at **http://pages.prodigy.com/ DXHF93A**.

1998 International Microwave Symposium and Exhibition June 7-12, Baltimore, Maryland

The International Microwave Symposium (IMS) and Exhibition is the most significant event of the year for microwave and RF technologists. It is sponsored by the Institute of Electrical and Electronics Engineers (IEEE) Microwave Theory and Techniques Society (MTT-S). The Symposium is a busy week for an expected 9,000+ attendees. Technical activities include 24 workshops on Sunday, Monday and Friday, and five parallel Technical Sessions and lunchtime Panel Sessions on Tuesday, Wednesday and Thursday. The Radio Frequency Integrated Circuits (RFIC) conference runs Monday and Tuesday, sharing with IMS a special focus on "Wireless" on Tuesday. The Automatic Radio Frequency Techniques Group (ARFTG) concludes the week with their 51st conference on Friday.

In addition to its coverage of the state of the art in microwave technology, the recent surge of interest in 'personal communications' has caused IMS to expand its coverage to include subjects more traditionally of interest to amateurs. These include sessions and workshops on low noise-techniques, HF/VHF/UHF power amplifiers and ICs, filters and other passive components, RF measurements, FCC policy issues in microwave spectrum management, digital TV broadcasting and video-on-demand systems and the International Mobile Telecommunications-2000 ITU standard. More than 300 technical papers will be presented during the week.

Nearly 400 companies will be displaying goods and services in 153,000 square feet of commercial exhibits. These cover everything from components to manufacturing systems. Several vendors offer textbooks and software for sale on site. More than 35 vendors will present application-oriented seminars during the exhibition.

Full-conference registration is available at price levels starting at $165. Rates for IEEE nonmembers and on-site registration are higher. Full registration includes a printed three-volume digest of all papers presented. Students, retirees and one-day attendees can register for significantly lower cost. All IMS technical registrants receive the Digest in CD-ROM form. Exhibit-only registration is available for $10.

Complete registration information and the program schedule is available on our Web site: **http://estd-www.nrl .navy.mil/ims/1998ims.html**. On-line preregistration is available at **http://www.expo-intl.com/shows/ mtt-s/register.**

All technical events of IMS '98 will take place in the Baltimore Convention Center, located in downtown Baltimore, next door to the world famous Inner Harbor.

The social program includes the RFIC Conference Reception on Sunday evening, the Microwave Journal/MTT-S Reception on Monday evening, a Crab Feast on Tuesday evening, the Industry Reception and MTT-S Awards Banquet on Wednesday evening and Guests' Program tours of Washington, DC, Annapolis and Baltimore.—*Steven Stitzer Chairman IMS '98 Steering Committee* ☐☐