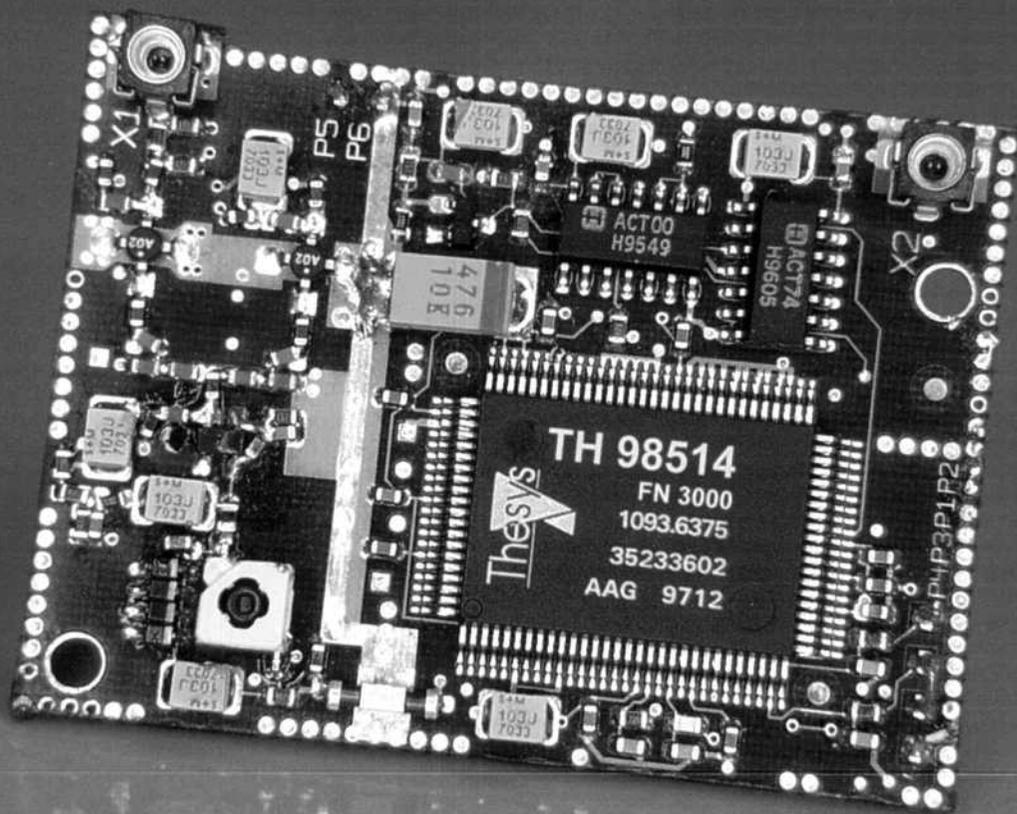


QEX

July/August 1998

Forum for Communications Experimenters

\$4



**Off-the-Shelf Fractional-N-Division Frequency Synthesis
by Dr. Ulrich L. Rohde, DJ2LR/HB9AWE/KA2WEU**

QEX: Forum for
Communications Experimenters
American Radio Relay League
225 Main Street
Newington, CT USA 06111-1494

QEX

QEX (ISSN: 0886-8093) is published bimonthly in January 98, March 98, May 98, July 98, September 98, and November 98 by the American Radio Relay League, 225 Main Street, Newington CT 06111-1494. Subscription rate for 6 issues to ARRL members is \$18; nonmembers \$30. Other rates are listed below. Periodicals postage paid at Hartford CT and at additional mailing offices.

POSTMASTER: Form 3579 requested. Send address changes to: QEX, 225 Main St, Newington CT, 06111-1494
Issue No. 189

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

Subscription rate for 6 issues:

In the US: ARRL Member \$18,
nonmember \$30;

US, Canada and Mexico by First Class Mail:
ARRL Member \$31, nonmember \$43;

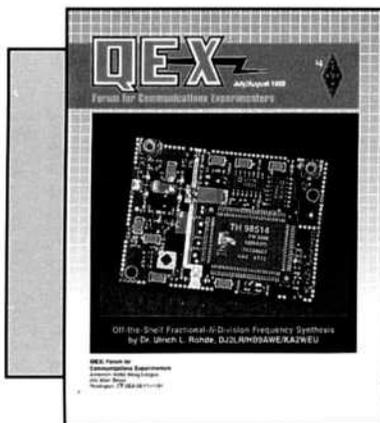
Elsewhere by Surface Mail (4-8 week delivery):
ARRL Member \$23,
nonmember \$35;

Elsewhere by Airmail: ARRL Member \$51,
nonmember \$63.

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

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.

Copyright © 1998 by the American Radio Relay League Inc. Material may be excerpted from QEX without prior permission provided that the original contributor is credited, and QEX is identified as the source.



About the Cover

A Synergy Microwave fractional-N-division synthesizer. It's only $2\frac{1}{8} \times 1\frac{5}{8}$ inches and offers outstanding performance. See Dr. Rohde's article for more information.



Features

3 A High-Performance Fractional-N Synthesizer

By Ulrich L. Rohde, KA2WEU

13 Signals, Samples, and Stuff: A DSP Tutorial (Part 3)

By Doug Smith, KF6DX/7

28 An Introduction to Microcontrollers: Ham Radio Style!

By Al Williams, WD5GNR

32 A Continuous Coverage HF VFO

By Francesco Morgantini, IK3OIL

38 An Updated Electronic Eyeball

By Bob Dildine, 7JIAFR, W6SFH

45 The Twin-T Filter

By Parker Cope, W2GOM/7

50 Multiple-Octave Bidirectional Wire Antennas

By Robert Zavrel, Jr, W7SX

Columns

53 RF

By Zack Lau, W1VT

57 Upcoming Technical Conferences

59 Letters to the Editor

July/Aug 1998 QEX Advertising Index

Alpha Delta Communications: Cov IV

American Radio Relay League: 49

ByteMark: 12

Communications Specialists, Inc: 49

Crestone Technical Books: Cov II

EM Scientific, Inc: 64

HAL Communications Corp: 61

Kachina Communications, Inc.: Cov III

PC Electronics: 12

SHF Microwave Parts Company: 62

T-Tech, Inc: 62

Timeline Inc: 63

Tucson Amateur Packet Radio

Corp: 61

Z Domain Technologies, Inc: 27



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

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

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

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

Membership inquiries and general correspondence should be addressed to the administrative headquarters at 225 Main Street, Newington, CT 06111 USA.

Telephone: 860-594-0200
Telex: 650215-5052 MCI
MCIMAIL (electronic mail system) ID: 215-5052
FAX: 860-594-0259 (24-hour direct line)

Officers

President: RODNEY STAFFORD, W6ROD
5155 Shadow Estates, San Jose, CA 95135

Executive Vice President: DAVID SUMNER, K1ZZ

Purpose of QEX:

- 1) provide a medium for the exchange of ideas and information between Amateur Radio experimenters
- 2) document advanced technical work in the Amateur Radio field
- 3) support efforts to advance the state of the Amateur Radio art

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

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

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

Empirically Speaking

This issue marks a full year since restarting *QEX*. For me, the year has been hectic but very rewarding. As you all can see, *QEX* is a going concern. We have received many messages and comments praising the new look and content. This has been very gratifying, and it would be easy to take the credit. But, as I said in my first editorial, the content and the success of *QEX* is largely dependent on the readers and the authors who send us the material, comments and suggestions that make up the magazine. Bob Schetgen, Maty Weinberg, the production crew and I have certainly worked hard, but without your input this magazine wouldn't exist. This magazine is what *you* make of it. As long as we get the enthusiastic support that we have enjoyed this past year, *QEX* will thrive.

I promised a series of articles on switch-mode converters last year, but I've not had time to write—being absorbed in the day-to-day business of *QEX*. I haven't forgotten, but it may be a while. Surely, we must have some volunteers out there? While on the subject of switching regulators, I received a note from AJ0J pointing out that there is an infinite supply of inexpensive, off-line switching converters available as PC power supplies. While not quite right for most ham projects, they can be modified. Have any of you modified one of these for ham applications? If you have, we would sure be interested in an article for *QEX*.

I received note from W7PMD about the inexpensive (relatively) RF power MOSFETs being made by APT. I took a quick look at their Web site (<http://www.advancedpower.com>). The ARF446 and '447 are particularly interesting. These give much better performance on 6 through 20 meters than the usual low-frequency MOSFETs intended for switching applications, and they don't cost hundreds of dollars as do most VHF power FETs—6 m solid-state, CW/FM, kilowatt amplifier anyone?

Amateur Television Quarterly has a new publisher: Harlan Technologies. Check them out at <http://www.cris.com/~Gharlan>.

In this QEX

Ulrich Rohde, KA2WEU, is back

with an article on high-performance synthesizers. As usual it's first-class stuff and very thought provoking.

Doug Smith's series on digital receivers continues this month with a closer look in several areas. You've had the introduction, now it's time to get into the nitty-gritty. To help with absorption, we have given you half of the conclusion (Part 3) in this issue, and the other half (Part 4) will be in the Sep/Oct issue.

Peripheral Interface Controllers (PIC) allow you to couple your computer to another piece of gear in the shack and reprogram that equipment at will. Al Williams, WD5GNR, shows you how to use a PIC to build a programmable keyer with all the bells and whistles. This is not only a useful piece of equipment, but it's a good introduction to PICs.

A wide-range VFO is a tricky part of any receiver or transceiver design. Francesco Morgantini, IK3OIL, gives you the lowdown on his partial-synthesis VFO. This VFO provides very good performance with only a moderate level of complexity.

Once you've had a panadapter to keep track of band activity, you really can't live without one. A good panadapter is expensive, however, and many homebrew designs are daunting. In many areas, inexpensive used oscilloscopes are available and a relatively simple RF section can use such 'scopes as a display. Bob Dildine, W6SFH, gives details of his panadapter design, using an old oscilloscope for the display. He also points out that monitors and even portable TV sets can be pressed into service for display units.

The twin-T is a very versatile network. Parker Cope, W2GOM/7, takes us back through the basics and shows how to use the circuit for selective amplifiers and low-distortion oscillators.

A simple antenna can still provide good DX performance as Bob Zavrel, W7SX, shows us this month.

Many readers would like to design circuits the way Zack Lau, W1VT, does. To help us learn, Zack is going to let us "inside his head" as he designs circuits in several future RF columns. We start this month, with a 6 meter band-pass filter.—73, Rudy Severns N6LF, rseverns@arrl.org

A High-Performance Fractional-N Synthesizer

A novel approach for high-performance synthesizers used in local-oscillator design for receivers and transmitters based on the fractional-N synthesis principle.

By Ulrich L. Rohde, KA2WEU
Chairman, Synergy Microwave Corporation

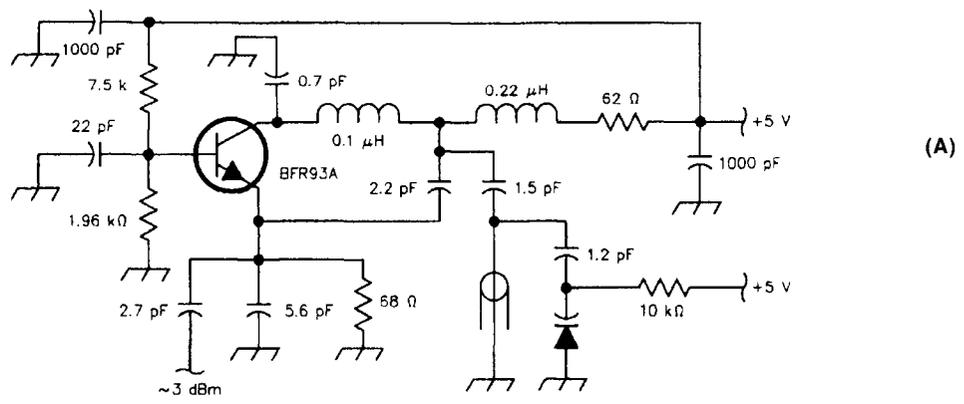
The design of high-performance frequency synthesizers has been discussed in great depth over many years. High performance means different things to different people. On one hand, portable cellular telephones need to be smaller, conserve energy, and be less costly. On the other hand, point-to-point stations need high performance in multifrequency, high-signal-level environments. Hand-held radios are seldom bombarded by as many signals as base stations. The trend in portable radios is toward single-chip solutions, which essentially

leave no room for high-Q circuits. Base stations have very stringent requirements for the spurious-free dynamic range, which is determined by the preamplifier, mixer and synthesizer. The amplifier and mixer technology is well understood, and the third-order and higher intercept points are determined by properties of devices, such as high-level diodes and transistors operated in their linear regions. In the case of the synthesizer, the ability to provide a good signal-to-noise ratio is highly dependent on the architecture of the synthesizer. This paper provides some application insight into modern solutions with fractional-N synthesis, rather than hybrid synthesizers—including DDS stages.^{12, 13, 14} It also looks at all angles of the synthesizer.

The VCO

Fig 1A shows the high-performance test VCO as well as a block diagram of a synthesizer using a dual-modulus prescaler. The prescaler is necessary because of the high input frequency. The tuning range of ceramic-resonator VCOs is about 5%; therefore, the influence of the tuning diode is not as drastic as experienced in wideband oscillators. On the other hand, microstrip-resonator VCOs can cover wider bandwidths with a slight degradation in phase noise, especially at frequencies above 400 MHz, while short-wave receivers need VCOs from 75 to 105 MHz with good performance. Fig 1B shows such a VCO.

The ceramic resonator is a high-Q



FROM FILTER

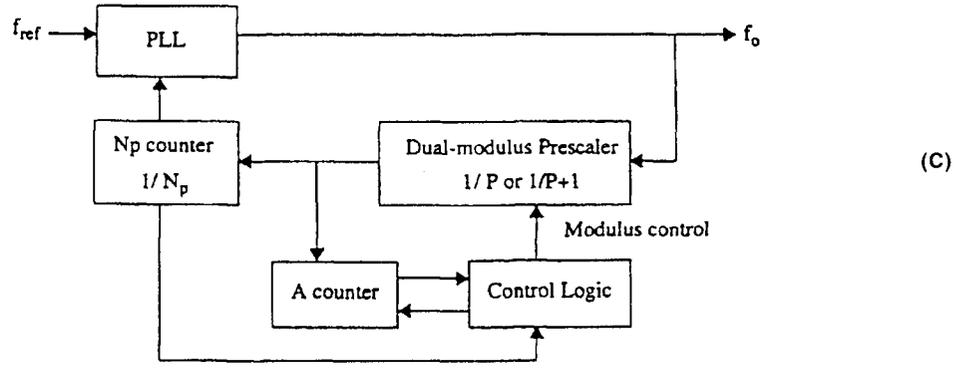
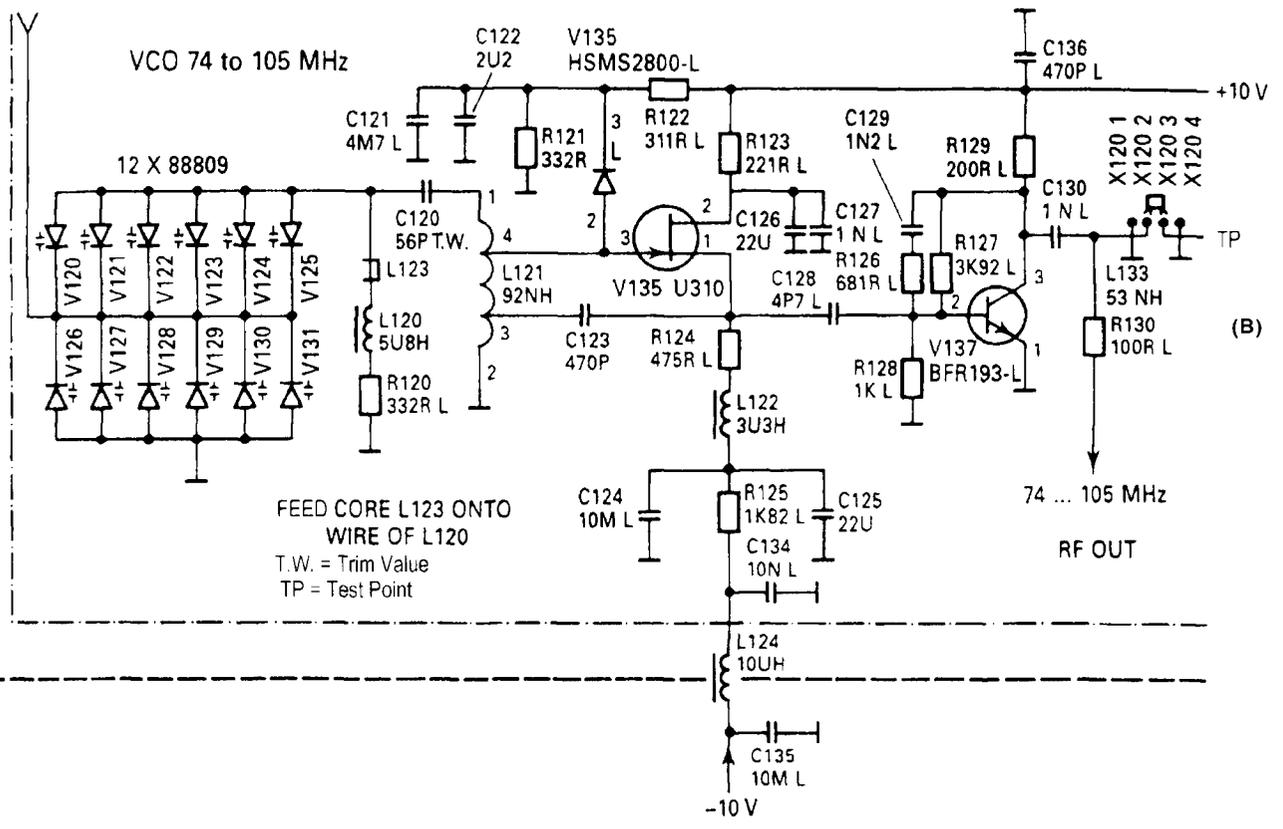


Fig 1—A shows a high-performance test circuit for use with a ceramic resonator. These resonators are available in the 500 MHz to 2 GHz range. For higher frequencies, dielectric resonators are recommended. B shows a wideband VCO with a large number of tuning diodes to improve phase noise. Note that the gate diode is biased in reverse and does not provide positive clipping as suggested by other authors. C is a frequency synthesizer with a dual-modulus prescaler using the VCO shown above.

element that determines the operating frequency. The phase noise of the VCO, without being locked by the synthesizer, is determined by the flicker corner frequency of the active device, the operating Q of the resonator and the noise floor of the transistor under large-signal conditions. While the correct method to determine phase noise of the free-running oscillator is the use of a nonlinear CAD tool, such as Ansoft's *Serenade 7.5* product, a good first-order approximation can be done by using an expanded form of Leeson's equation.^{1, 2} The original equation by Leeson did not consider the effect of the tuning diode, this enhancement was done later.² Considering the post amplifiers, the equation can be expanded further to calculate all other contributions.

The phase noise of a VCO is determined by

$$\mathcal{L}(f_m) = 10 \log \left\{ \left[1 + \frac{f_0^2}{(2f_m Q_{load})^2} \right] \left(1 + \frac{f_c}{f_m} \right) \frac{FKT}{2P_{sav}} + \frac{2KTRK_0^2}{f_m^2} \right\} \quad (\text{Eq 1})$$

where

$\mathcal{L}(f_m)$ = ratio of sideband power in a 1-Hz bandwidth at f_m to total power, in dB.

f_m = frequency offset.

f_0 = center frequency.

f_c = flicker frequency.

Q_{load} = loaded Q of the tuned circuit.

F = noise factor.

kT = 4.1×10^{-21} at 300 K₀ (room temperature).

P_{sav} = average power at oscillator output.

R = equivalent noise resistance of tuning diode (typically 200 Ω to 10 kΩ).

K = oscillator voltage gain.

When adding an isolating amplifier, the noise of an LC oscillator is determined by

$$S_{\phi}(f_m) = \left[a_R F_0^4 + a_E (F_0 / (2Q_L))^2 \right] / f_m^3 + \left[(2GFKT / P_0) (F_0 / (2Q_L))^2 \right] / f_m^2 + (2a_R Q_L F_0^3) / f_m^2 + a_E / f_m + 2GFKT / P_0 \quad (\text{Eq 2})$$

where

G = compressed power gain of the loop amplifier.

F = noise factor of the loop amplifier.

K = Boltzmann's constant.

T = temperature in kelvins.

P_0 = carrier power level (in watts) at the output of the loop amplifier.

F_0 = carrier frequency, in Hz.

f_m = carrier offset frequency in Hz.

$Q_L (= \pi F_0 \tau_g)$ = loaded Q of the resonator in the feedback loop.

a_R and a_E = flicker noise constants for the resonator and loop amplifier, respectively.

Fig 2A shows the noise spectrum that can be expected at this frequency range from a high-quality VCO. The performance of the oscillator, when part of the PLL, inside the PLL bandwidth, is determined by the noise contribution of the reference frequency, the phase detector, the dividers and other items in the loop. Outside the loop bandwidth, the phase noise is determined only by the quality of the oscillator, which explains the need to build a high-quality

VCO at the outset. The single-sideband phase noise is only one property to consider. In the PLL, the reference frequency (25 kHz, as an example) will show up on the output. This means that a discrete signal can appear at 25 kHz intervals. The loop filter and the board layout determine the suppression of these unwanted spurious signals. Fig 2B shows the noise spectrum for a 75 to 105 MHz VCO.

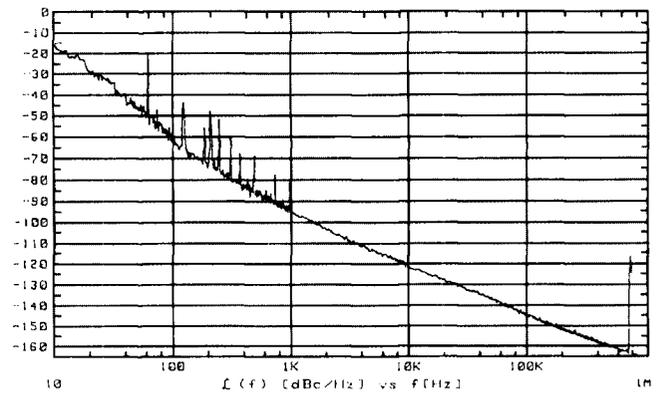
Fig 3 shows the single-sideband phase noise and spurious response of the VCO with the loop locked. As expected, the close-in phase noise improves, and the phase noise outside the loop bandwidth is equal to the oscillator's performance by itself.

Synthesizer System

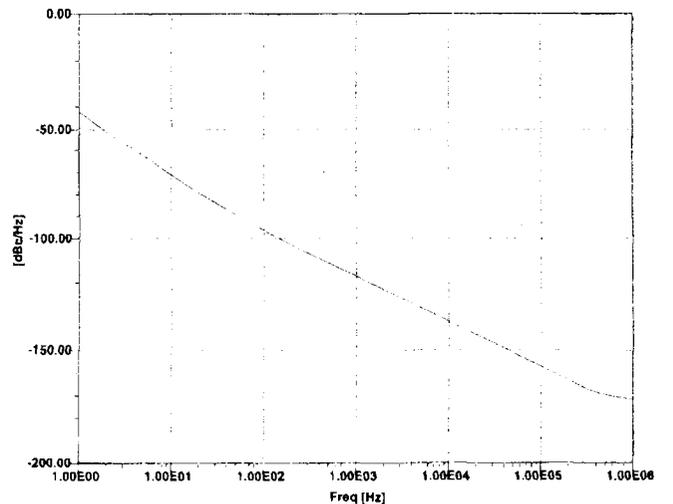
In order to close the loop, a loop filter must be determined and designed. From a user's point of view, there are several choices:

Passive Filter

The overall loop frequency possible must be calculated



(A)



(B)

Fig 2—A shows the measured phase noise of a "high end" ceramic-resonator-based VCO. B shows the predicted SSB phase noise of the multidiode oscillator operating from 75 to 105 MHz as shown above. Measurements and predictions agree within 1 dB. The flicker corner frequency can be observed. The good performance is due to the multiple parallel diodes and the limiting Schottky diode. This oscillator is used in the Rohde and Schwarz Model ESN test receiver.

¹Notes appear on page 11.

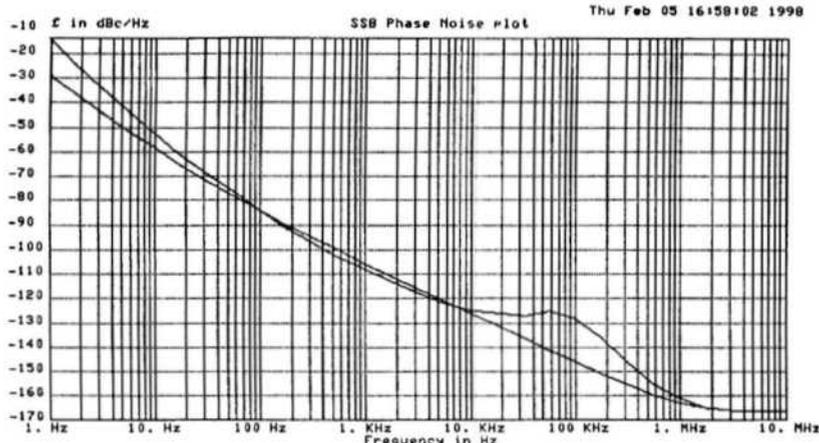


Fig 3—Open- and closed-loop phase noise of a “high-Q” oscillator. The loop bandwidth is set at about 80 kHz for high-speed switching. A reference frequency of 50 MHz was used.

first, and can be determined by:

$$\text{loop frequency} = \frac{K_{vco} K_{\phi}}{N} \quad (\text{Eq 3})$$

with K_{vco} being the oscillator sensitivity, K_{ϕ} = phase detector voltage gain and N being the total average division ratio. This type 1, second-order loop, however, has finite dc gain, which gives a large static phase error. To achieve a small phase error requires large dc gain.

Let's consider an example. A typical synthesizer using a standard chip set operates at a reference of 25 kHz ($N = 35200$), 880 MHz center frequency, a phase-detector voltage gain of 0.9 and an oscillator sensitivity of 5 MHz per volt. The resulting loop frequency equals 127.8 Hz. This also means that a loop filter bandwidth of less than 128 Hz is needed for acceptable stability. Needless to say, such a loop does not correct microphonic effects and other noise contributions. If we now take a step size of 200 kHz, or a reference frequency of 200 kHz (N becomes 4400), the loop frequency can change to about 1 kHz. Now, we can use a filter bandwidth of about 1 kHz, resulting in faster settling time and cancellation of microphonic effects. Fig 4 shows a passive (type 1) high-order filter.

Active Filters

If it is necessary to have zero phase error in response to step changes in the input frequency, the dc gain of a loop filter must be infinite. This can be realized by including in $F(s)$ a pole at the origin. An active loop filter implemented by an operational amplifier

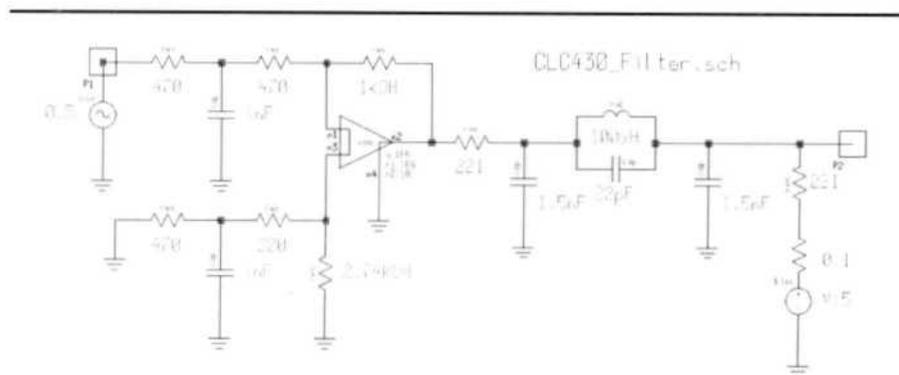


Fig 4—A type 1 high-order filter with a notch to suppress discrete reference spurs.

with large open-loop gain can accomplish this. The advantage of an active filter over its passive counterpart stems from the presence of a very high dc gain amplifier, which allows the realization of a near ideal integration in the loop filter. A filter with a pole at the origin helps reduce the static phase error θ_e to a very small residual value. The very high dc gain provided by the ideal integrator allows the PLL to track the input signal with very small phase error (almost zero) following a frequency step corresponding to a channel switch in a frequency synthesizer. While for an active loop filter with a very high gain amplifier, the dc gain $F(0)$ is very high, the static phase error θ_e becomes almost zero. A finite static phase error is undesirable, because it enhances the noise susceptibility of the system, especially at the high frequency range, as is the case with portable wireless applications.

A simple way to realize the same performance as an active filter, second-order PLL without incurring the penalty of using a noisy, power-consuming and offset-susceptible op amp is a charge-pump PLL. The charge-pump PLL offers two important advantages over analog PLLs: (1) The capture range is only limited by the VCO output frequency range. (2) The static phase error is zero if mismatches and offsets are negligible. In some charge pumps, the internal resistor with the capacitor introduces ripple in the control voltage even when the loop is locked. Upon each cycle of the PFD, the pump current, I_p , is driven into the loop filter, $F(s)$, which responds with an instantaneous voltage jump. At the end of the charging interval, the pump current switches off and a voltage jump of equal magnitude occurs in the opposite direction. This causes the frequency modulation

at the output of the VCO, generating sideband spurs. This effect is undesirable in frequency synthesis. To suppress the ripple, a capacitor, *C1*, can be connected from the output of the charge pump to ground, in parallel. This modification introduces a third pole in the PLL, requiring further study of stability issues. A fully symmetrical charge pump can reduce the magnitude of the ripple considerably, up to 70 dB.

Fig 5 shows the high-order active filter. PLL theory requires a 45° phase margin at 0 dB gain and a -6 dB-per-octave slope from 10 dB down to -10 dB gain, for stability.

Table 1 shows the phase noise of a crystal frequency standard operating at 50 MHz and its noise multiplied up to 880 MHz, dependent on the reference frequency selected. The examples are 25 kHz, 1.25 MHz and 50 MHz, derived from this 50 MHz high-performance standard.

If these values, which assume some linear scaling, are plotted against the phase noise of the free-running oscillator, Fig 6 results. This table also assumes that the phase frequency discriminator has a phase-noise floor of -160 dBc/Hz, while the noise of the actual standard, for more than 10 kHz away, went down to -165 dBc/Hz.

The phase noise of the reference, multiplied to the operating frequency, determines the loop-filter bandwidth. In the case of the 25 kHz reference, as can be seen from the turning point in F1, a loop frequency of 30 Hz should be chosen. When going to the 1.25 MHz

Table 1—Phase noise (in dBc/Hz) versus reference frequency derived from a high-performance 50 MHz crystal oscillator and multiplied to 880 MHz

Offset from Carrier	Reference Frequency			
	50 MHz Standard	25 kHz	1.25 MHz	50 MHz
10 Hz	-80	-55	-55	-55
100 Hz	-110	-69	-85	-85
1 kHz	-140	-69	-103	-115
10 kHz	-160	-69	-103	-135
100 kHz	-160	-69	-103	-135
1 MHz	-160	-69	-103	-135

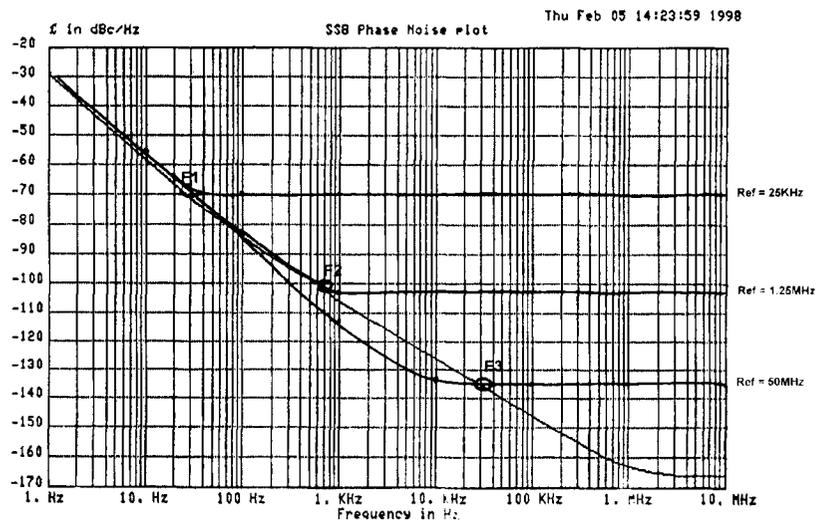


Fig 6—Noise contribution of a 50 MHz frequency standard divided down to the individual reference frequency and multiplied back up to 880 MHz. The continuing line below the breakpoint, F_3 , is the free running oscillator itself. $F_0 = 880$ MHz; $F_1 = 30$ Hz; $F_2 = 700$ Hz; $F_3 = 30$ kHz.

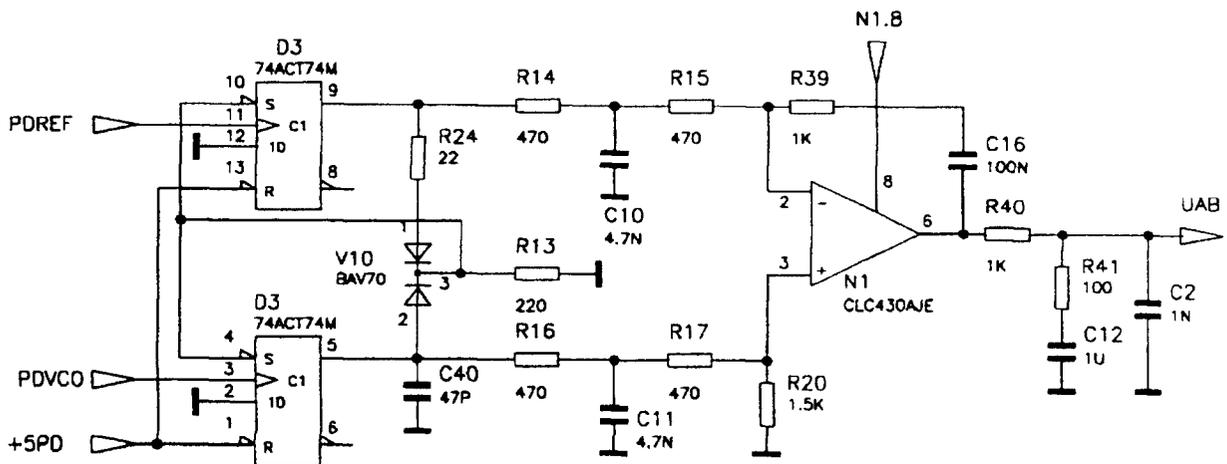


Fig 7—A phase/frequency discriminator including an active-loop filter capable of operating up to 100 MHz.

reference, the filter bandwidth can move out to 700 Hz, and for a reference of 50 MHz, 30 kHz is the appropriate loop frequency. Needless to say, references of 1.25 MHz and 50 MHz do not allow a 25 kHz step size. The higher reference frequencies also require more complex phase discriminators and filters. A good example is shown in Fig 7.

This problem is solved in a conventional synthesizer by resorting to multiple loops. In our case, we are going to apply a more clever scheme, the fractional- N synthesizer principle.

Fractional-Division Synthesizer

The principle of the fractional- N division synthesizer has been around for a while. In the past, it has been implemented in analog systems. The above mentioned single loop uses a frequency divider with an integer division ratio between 1 and some very large number, hopefully not as high as 50,000. It would be ideal to be able to build a synthesizer with the 1.25 MHz reference or 50 MHz reference and yet obtain the desired step-size resolution, such as 25 kHz. This would lead to much smaller division ratios; therefore, as pointed out, the results of Fig 6 would have much better phase noise performance.

An alternative would be for N to take on fractional values. The output frequency could then be changed in fractional increments of the reference frequency. Although digital dividers cannot provide fractional division ratios, there are ways to accomplish the task effectively.

The most frequently used method divides the output frequency by $N + 1$ once every M cycles and to divide by N the rest of the time. The effective division ratio is then $N + 1/M$, and the average output frequency is given by:

$$f_0 = \left(N + \frac{1}{M} \right) f_r \quad (\text{Eq 4})$$

This expression shows that f_0 can be varied in fractional increments of the reference frequency by varying M . The technique is equivalent to constructing a fractional divider, but the fractional part of the division is actually implemented using a phase accumulator. The phase accumulator approach is illustrated by the following example. This method can be expanded to frequencies much higher than 6 GHz using appropriate synchronous dividers.

Example: Consider the problem of generating 899.8 MHz using a fractional- N loop with a 50 MHz reference frequency:

$$899.8 \text{ MHz} = 50 \text{ MHz} \left(N + \frac{K}{F} \right)$$

the integral part of the division N must be 17 and the fractional part K/F needs to be $996/1000$; (the fractional part, K/F is not an integer) and the VCO output must be divided by 996 every 1000 cycles. This can be easily implemented by adding the number 0.996 to the contents of an accumulator once each cycle. Every time the accumulator overflows, the divider divides by 18 rather than 17. Only the fractional value of the addition is retained in the phase accumulator. If we move to the lower band or try to generate 850.2 MHz, N remains 17 and K/F becomes $4/1000$. This method of fractional division was first introduced using analog implementation and noise cancellation, but today it is implemented as a totally digital approach. The necessary resolution is obtained from the dual-modulus prescaling, which allows for a well-established method for achieving a high-performance frequency synthesizer operating at UHF and higher frequencies. Dual-modulus prescaling maintains system resolution where a simple prescaler

would not; it allows a VCO step equal to the value of the reference frequency. This method needs an additional counter and the dual-modulus prescaler then divides one or two values depending on the state of its control. The only drawback of dual-modulus prescalers is the minimum prescaler division ratio, approximately N^2 . The dual-modulus divider is the key to implementing fractional- N synthesizers.

Although the fractional- N technique appears to have good potential for solving the resolution limitation, it has its own complications. Typically, an overflow from the phase accumulator (which is the adder with the output feedback to the input after being latched) is used to change the instantaneous division ratio. Each overflow produces a jitter at the output frequency (caused by the fractional division) and is limited to the fractional portion of the desired division ratio.

In our case, we had chosen a step size of 200 kHz, and yet the discrete sidebands vary from 200 kHz for $K/F = 4/1000$ to 49.8 MHz for $K/F = 996/1000$. It is the task of the loop filter to remove those discrete spurious signals. While the removal of the

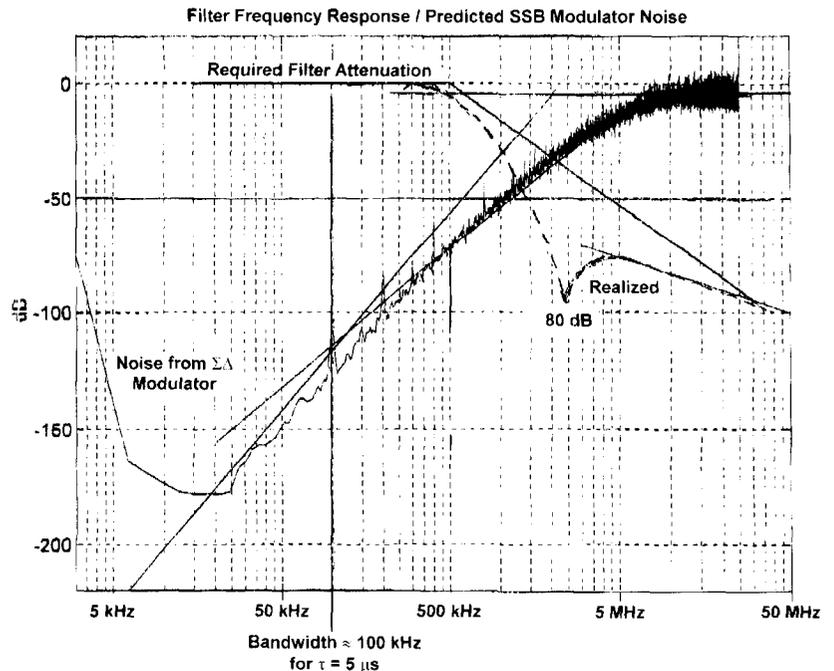


Fig 8—The filter-frequency-response/phase-noise-analysis graph shows the required attenuation for the reference frequency of 50 MHz and the noise generated by the $\Sigma\Delta$ converter (three steps) as a function of the offset frequency. It's apparent that the $\Sigma\Delta$ converter noise dominates above 80 kHz unless attenuated.

discrete spurs has been accomplished in the past by analog techniques, various digital methods are now available. The microprocessor has to solve the following equation:

$$N^* = \left(N + \frac{K}{F} \right) = [N(F - K) + (N + 1)K]$$

(Eq 5)

Table 2—Modern spur-suppression methods

Technique	Feature	Problem
DAC Phase Estimation	Cancel Spur by DAC	Analog Mismatch
Pulse Generation	Insert Pulses	Interpolation Jitter
Phase Interpolation	Inherent Fractional Divider	Interpolation Jitter
Random Jittering	Randomize Divider	Frequency Jitter
$\Sigma\Delta$ Modulation	Modulate Division Ratio	Quantization Noise

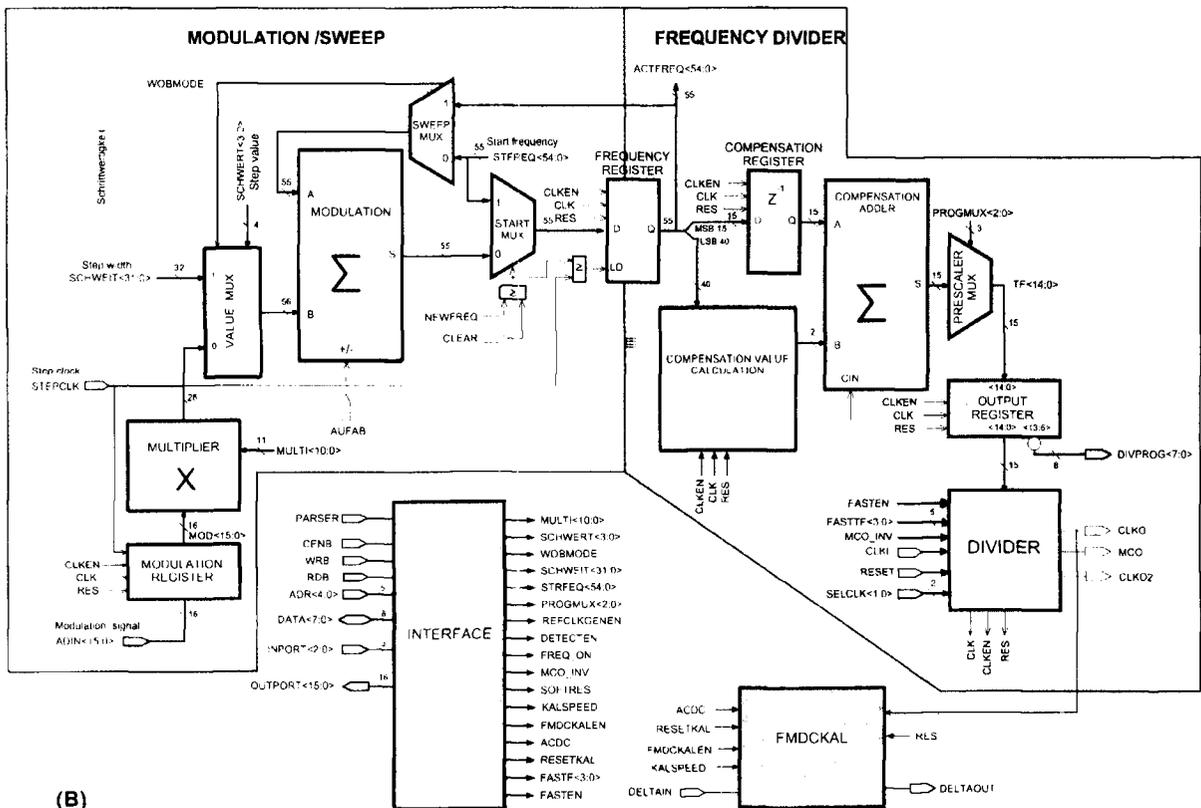
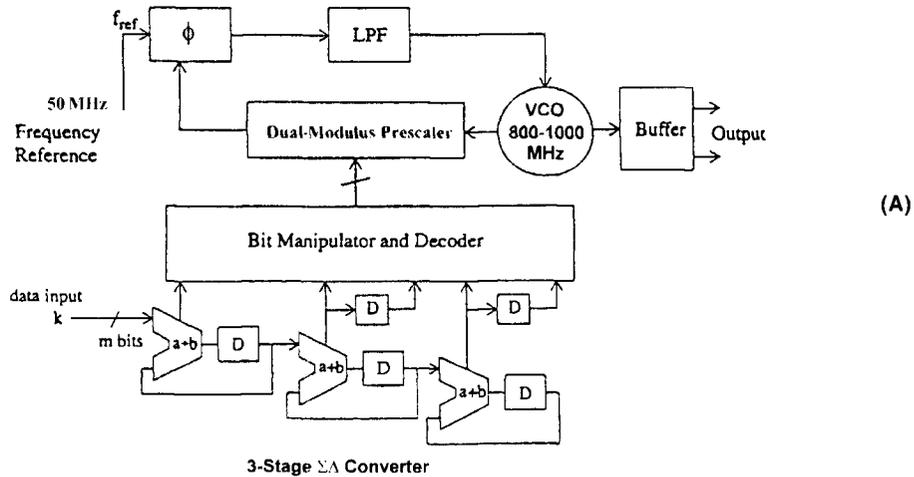


Fig 9—(A) Block diagram of the fractional-N-division synthesizer built using a custom IC capable of operation at reference frequencies up to 100 MHz. Its use of smaller modulus values is responsible for its frequency extension up to 3 GHz ($\pm 4 \pm 2$) with ripple or asynchronous counters, and allows the implementation of dual-modulus counts to $M/N+1$. It uses the VCO shown in Figure 1A, and the phase-frequency discriminator shown in Figure 7. **B** is a detailed block diagram of the fractional-N-division synthesizer chip.

Example

For $F_0 = 850.2$ MHz, we obtain:

$$N^* = \frac{850.2 \text{ MHz}}{50 \text{ MHz}} = 17.004$$

Following the formula above:

$$N^* = \left(N + \frac{K}{F} \right) = \frac{[17(1000 - 4) + (17 + 1) \times 4]}{1000}$$

$$= \frac{[16932 + 72]}{1000} = 17.004$$

$$F_{out} = 50 \text{ MHz} \times \frac{[16932 + 732]}{1000}$$

$$= 846.6 \text{ MHz} + 3.6 \text{ MHz}$$

$$= 850.2 \text{ MHz}$$

By increasing the number of accumulators, frequency resolutions much below 1 Hz step size are possible with the same switching speed.

Spur-Suppression Techniques

While several methods have been proposed in the literature (see patents in References 4 through 9), the method of reducing noise by using a $\Sigma\Delta$ (sigma-delta) modulator is most promising. The concept is to eliminate low frequency phase error by rapidly switching the division ratio to eliminate the gradual phase error at the discriminatory input. By changing the division ratio rapidly between different values, phase errors occur in both polarities, positive as well as negative, and at an accelerated rate that explains the phenomenon of high-frequency noise push-up. This noise is converted to a voltage by the phase/frequency discriminator and loop filter and attenuated by the low-pass filter. The main problem associated with this noise-shaping technique is that the noise power level rises rapidly with frequency. Fig 8 shows noise contributions with such a $\Sigma\Delta$ modulator in place.

On the other hand, we can now build a single-loop synthesizer with switching times as fast as 6 μs and very little phase-noise deterioration inside the loop bandwidth (see Fig 8). Since this system maintains the good phase-noise performance of a ceramic-resonator-based oscillator, the resulting performance is significantly better than the phase noise expected from high-end signal generators. This method does not allow us to increase the loop bandwidth beyond the 100 kHz limit, however, where the noise contribution of the $\Sigma\Delta$ modulator takes over.

Table 2 shows some of the modern spur suppression methods. These three-stage $\Sigma\Delta$ methods with larger

accumulators have the most potential.⁴⁻⁹

The power spectral response of the phase noise for the three-stage $\Sigma\Delta$ modulator is calculated from:

$$L(f) = \frac{(2\pi)^2}{12 \cdot f_{ref}} \cdot \left[2 \sin \left(\frac{\pi f}{f_{ref}} \right) \right]^{2(n-1)} \text{ rad}^2 / \text{Hz}$$

(Eq 6)

where n is the number of the stage of the cascaded $\Sigma\Delta$ modulator.¹⁰ Eq 6 shows that the phase noise resulting from the fractional controller is attenuated to negligible levels close to the center frequency, and further from the center frequency, the phase noise is increased rapidly and must be filtered out prior to the tuning input of the VCO to prevent unacceptable degradation of

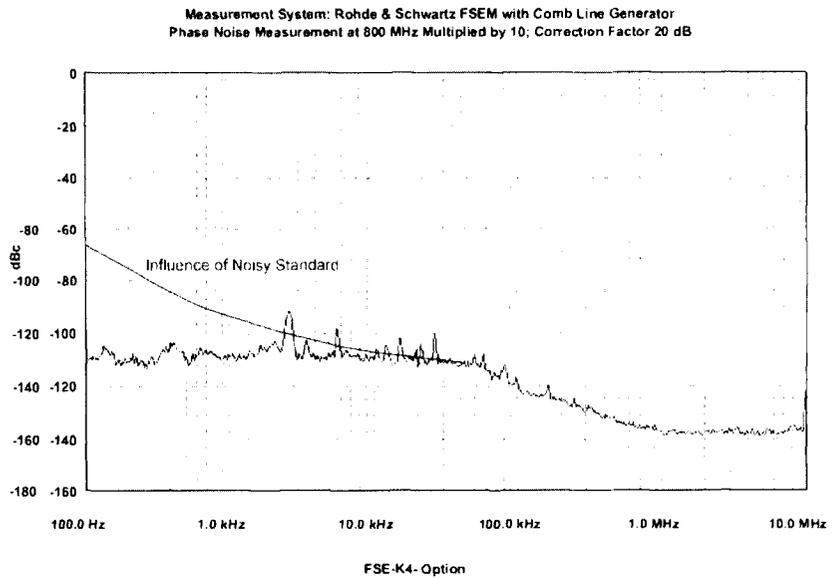


Fig 10—Measured phase noise of the fractional- N -division synthesizer using a custom-built, high-performance 50 MHz crystal oscillator as reference, with the calculated degradation due to a noisy reference plotted for comparison. Both synthesizer and spectrum analyzer use the same reference.

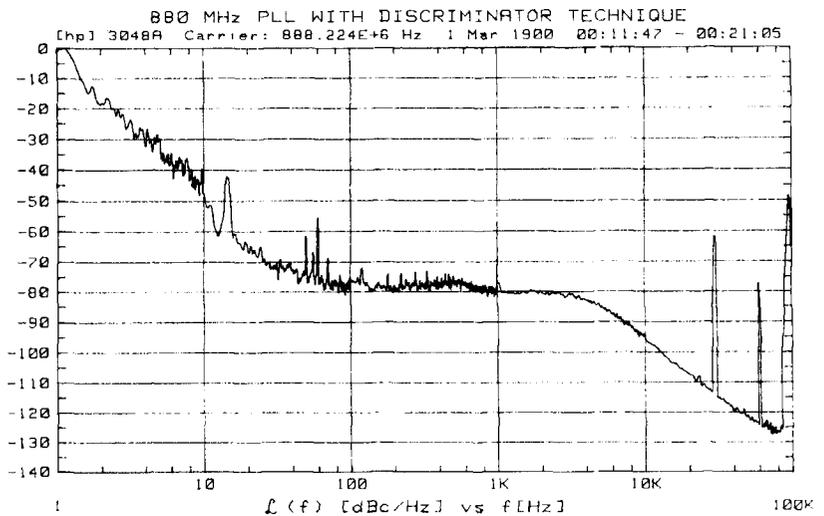


Fig 11—Measured phase noise of a 880 MHz synthesizer using a conventional synthesizer IC. Comparing this to Fig 10 shows the big improvement possible by fractional- N -division synthesizers as presented in this article.

spectral purity. A loop filter must be used to filter the noise in the PLL loop. Fig 8 is a plot of phase noise versus frequency offset from the center frequency. A fractional- N synthesizer with a three-stage $\Sigma\Delta$ modulator has been built (as shown in Fig 9A). The synthesizer consists of a phase/frequency detector, an active low-pass filter (LPF), a voltage-controlled oscillator (VCO), a dual-modulus prescaler, a three-stage $\Sigma\Delta$ modulator and a buffer.

After designing, building and predicting the phase-noise performance of this synthesizer, it becomes clear that the phase-noise measurement for such a system is tricky. Standard measurement techniques, with a reference synthesizer, would not provide enough resolution because synthesized signal generators on the market are not good enough to measure the phase noise as shown in Fig 2A. Therefore, we built a comb generator, which takes the output of the oscillator and multiplies it 10 to 20 times.

Passive phase-noise measurement systems based on delay lines are not selective, and the comb generator confuses them. Nonetheless, the Rohde and Schwarz FSEM spectrum analyzer with the K-4 option has sufficient resolution for such phase-noise measurements. All of the Rohde and Schwarz FSE-series spectrum analyzers use a somewhat more discrete fractional-division synthesizer with a 100 MHz reference. Based on the multiplication factor of 10, it turns out that there is enough dynamic range in the FSEM analyzer (with the K-4 option) for phase-noise measurement. The useful carrier-offset frequency range for the system is 100 Hz to 10 MHz, which is perfect for this measurement. Fig 10 shows the measured phase noise of the final frequency synthesizer.

During the measurements, it was also determined that the standard crystal oscillator proposed in Fig 6, and its data shown in Table 1, was not good enough. We found a need to develop a 50 MHz crystal oscillator with better phase noise. Upon examination of the measured phase noise shown in Fig 10, you can see that the oscillator used as the reference was significantly better. Otherwise, measurement of this phase noise would not have been possible. At 1 MHz and higher offset frequencies, the same phase noise as in Fig 2A was obtained. Also, the roll off in Fig 10 exhibits the loop-filter cutoff frequency of about 100 kHz. This fractional- N -

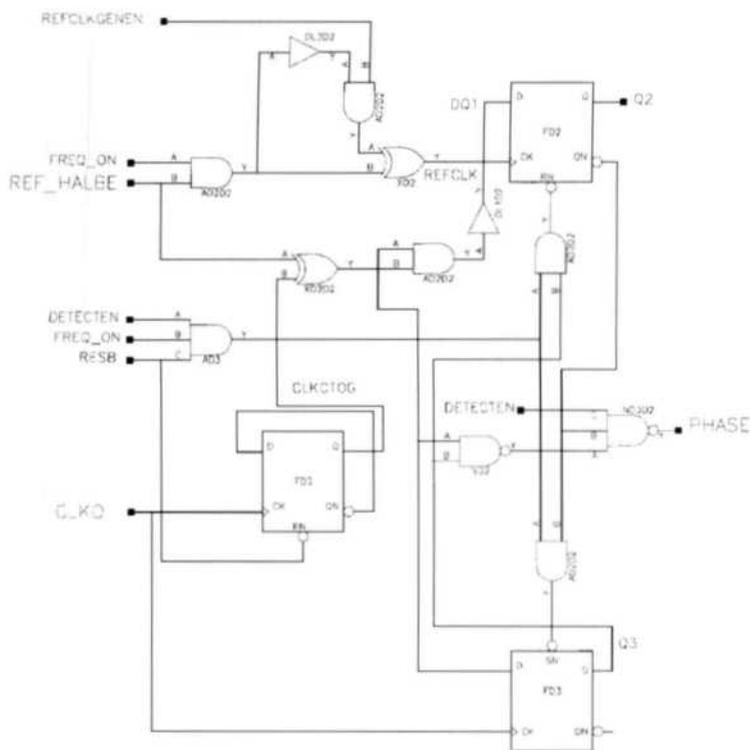


Fig 12—Custom-built phase detector with a noise floor of better than -168 dB. This phase detector shows extremely low phase jitter.

division synthesizer with a high-performance VCO has a significantly better phase noise than other example systems in this frequency range. In order to demonstrate the sufficient improvement, phase-noise measurements were made on standard systems, using typical synthesizer chips. While the phase noise and synthesizer design of those typical systems is quite good, it is no match for this new approach. See Fig 11.¹¹

Conclusion

It is possible to build an extremely high quality synthesizer system by combining the very best available technologies, such as:

- High-end VCOs with ceramic-resonator based tuned circuits
- High-Q LC arrangements including microstrips on Teflon material
- Modern fractional- N -division synthesizer blocks that can operate at 50 MHz reference frequencies

We have also learned that a synthesizer's limits are determined by the reference crystal oscillator and possibly by the phase detector. In our example, they required in-house designs to match our needs. The very

high bandwidth made switching speeds of 6 μ s possible. The resolution of the synthesizer itself depends on the accumulator size. A 25 kHz step size was successfully tested up to several megahertz. For wideband applications, some of the critical points are: 10 kHz where -120 dBc/Hz is desired, at 800 kHz better than -153 dBc/Hz, and at 3 MHz better than -155 dBc.

As validation, three types of synthesizers have been built: one from 75 to 105 MHz, with 1 Hz resolution, for an HF transceiver; another from 700 to 2000 MHz; and a third from 2700 to 3500 MHz, also with better than 1 Hz resolution.

I would like to acknowledge Department 1ES of Rohde and Schwarz, Munich, for their assistance with the design, specifically, the measurements and other useful suggestions. Without their support, some of this work would not have been possible.

Notes

- ¹D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, 1966, pp 329-330.
- ²Ulrich L. Rohde, Jerry Whitaker and T. N. N. Bucher, *Communications Receivers*, second edition, published by McGraw Hill,

New York, NY, January 1997, ISBN 0-07-053608-2, p 380.

³Ulrich L. Rohde, *Microwave and Wireless Synthesizers: Theory and Design*, by John Wiley and Sons, August 1997, ISBN 0471-52019-5, p 215.

⁴Alexander W. Hietala, Cary, IL; Duane C. Rabe, Rolling Meadows, IL Motorola, Inc., Schaumburg, IL, *Latched Accumulator fractional-N Synthesis With Residual Error Reduction*, United States Patent, Patent No. 5,093,632, March 3, 1992.

⁵Thomas A. D. Riley, Osgoode, Canada, Carleton University, Ottawa, Canada, *Frequency Synthesizers Having Dividing Ratio Controlled Sigma-Delta Modulator*, United States Patent, Patent No. 4,965,531, October 23, 1990.

⁶Nigel J. R. King, Wokingham, England, Racal Communications Equipment Limited, England, *Phase Locked Loop Variable Frequency Generator*, United States Patent, Patent No. 4,204,174, May 20, 1980.

⁷John Norman Wells, St. Albans, Hertfordshire (GB), Marconi Instruments, St. Albans, Hertfordshire (GB), *Frequency Synthesizers*, European Patent, Patent No. 0125790B2, July 5, 1995.

⁸Thomas Jackson, Twickenham, Middlesex (GB), Plessey Overseas Limited, Ilford, Essex (GB), *Improvement In Or Relating To Synthesizers*, European Patent, Patent No. 0214217B1, June 6, 1996.

⁹Thomas Jackson, Twickenham, Middlesex (GB), Plessey Overseas Limited, Ilford, Essex (GB), *Improvement In Or Relating To Synthesizers*, European Patent, Patent No. W086/05046, August 28, 1996.

¹⁰Byeong-Ha Park, Georgia Institute of Technology, "A Low Voltage, Low Power CMOS 900 MHz Frequency Synthesizer," a PhD Dissertation, December 1997.

¹¹Cris E. Hill, "All Digital Fractional-N Synthesizer for High Resolution Phase Locked Loops," *Applied Microwave and Wireless*, November/December 1997 and January/February 1998.

¹²Ulrich L. Rohde, *Microwave and Wireless Synthesizers: Theory and Design*, by John Wiley and Sons, August 1997, ISBN 0471-52019-5, p 489.

¹³Ulrich L. Rohde, "All About Phase Noise in Oscillators," *QEX*, Dec 1993, pp 3-6; Jan 1994 pp 9-16; Feb 1994 pp 15-24.

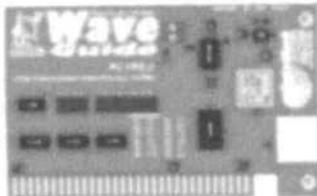
¹⁴Ulrich L. Rohde, "Designing low-Phase-Noise oscillators," *QEX*, Oct 1994, pp 3-12.

DDS from Your PC!

Announcing the new PC-VFO jr!

Digital VFO

Only \$139



PC plug-in card
DC to 54 MHz
.029 Hz Steps
+10 dBm output
Works on any ISA
motherboard

www.bytemark.com/waveguide



800 679-3184

AMIDON, INC.

Committed to Excellence Since 1963



Technical Specifications Online!
www.bytemark.com/amidon

W2FMI BALUNS and UNUNs

Exclusive discussion on the design and implementation of Transmission Line Transformers by the Legendary Dr. Jerry Sevick, W2FMI



www.bytemark.com/amidon/content6.htm

ByteMark Corporation

7714 Trent Street
Orlando, FL 32807
407 673-2083 FAX
407 679-3184 Voice

Amidon Master Distributor

AMATEUR TELEVISION

Web site: www.hamtv.com



SEE THE SPACE SHUTTLE VIDEO

With just our downconverter, antenna and TV set!

Many ATV repeaters and individuals are re-transmitting Space Shuttle Video & Audio from their TVRO's tuned to GE-2 (85W) transponder 9 vertical. Others may be re-transmitting weather radar during significant storms or home camcorder video. If it is being done in your area on 420 MHz - check page 607 in the 98-99 ARRL Repeater Directory or call us, ATV repeaters are springing up all over - all you need is the TVC-4G ATV 420-450 MHz downconverter, TV set to ch 2, 3 or 4, and 70cm antenna (you can use your 435 Oscar beam). We also have equipment for the 902-928 & 1240-1300 MHz bands. In fact we are your one stop for all your ATV info and equipment needs - antennas, all mode linear amps, transmitters, repeater modules, etc., for home, mobile, RACES, R/C, balloon and rocket ATV applications.

Hams, call for our complete 10 page ATV catalogue!

Low Cost Start



**Model TVC-4G
ATV Downconverter**
tunes 420-450 MHz to ch 3

Only \$89

TVC-9G 900 MHz \$99

TVC-12G 1200 MHz \$109

Get The **ATV Bug!**



**Companion TX70-1b
1.5Watt ATV
TRANSMITTER
only \$299**
or 10 Watt TX70-10
only \$439

Full color & sound
Plug in your camcorder,
antenna & reg. 13.8 Vdc

CALL (626) 447-4565 M-Th 8AM - 5:30 PM PST.

P. C. ELECTRONICS

2522 S. PAXSON Lane ARCADIA CA 91007



Email: tomsmb@aol.com

(626) 447-0489 24hr FAX



Tom (W6ORG) & Mary Ann (WB6YSS)

Signals, Samples, and Stuff: A DSP Tutorial (Part 3)

*Our tour continues with a foray
into advanced DSP techniques.*

By Doug Smith, KF6DX/7

Having learned about basic IF-DSP methods and their application in an actual transceiver, it's time to plunge into the truly magical stuff! In this third article in the series, we'll be looking at certain esoteric but extremely effective DSP techniques. By now, many of these concepts have found their way into production equipment, but they still are not generally well understood.

In this article, I will begin illustrating the underlying principles of current DSP noise reduction technology. Unfortunately, *QEX* space constraints

¹Notes appear on page 27.

require that Part 3 be broken across two issues. Therefore, Part 4 will complete the description of noise-reduction technology and the entire DSP series in the September/October issue.

DSP Noise Reduction Methods

Two noise reduction (NR) methods are prevalent in radio equipment today: the adaptive filtering method, and the Fourier transform method. We'll look at the theory behind each of these approaches, and discuss their implementation and performance. Then, a way of combining the two methods is considered. Along the way, I'll introduce a very fast way of calculating Fourier transforms—faster than the well-known “fast Fourier transform” algorithms.

Adaptive Filtering

In Part 1, we touched on the concept of an *adaptive interference canceler* and identified a design for an *adaptive notch filter* using the *least-mean-squares (LMS) algorithm*. These principles are explored in more detail here, as they apply to noise reduction systems. We'll find that it's possible to build an adaptive filter that accentuates the repetitive components of an input signal, and rejects the non-repetitive parts (noise). Further, we'll discover that the effectiveness of this technique depends on the characteristics of the input signals.

The nature of information-bearing signals is that they are in some way coherent; ie, they have some feature that distinguishes them from noise.

For example, voice signals have attributes relating to the pitch, syllabic content, and impulse response of a person's voice. CW signals are perhaps the simplest example because they constitute only the presence or absence of a single frequency.

Much research has been done about detection of a sinusoidal signal buried in noise.^{1,2} Adaptive filtering methods are based on the exploitation of the statistical properties of the input signal, specifically the *autocorrelation*. Simply put, autocorrelation refers to how recent samples of a waveform resemble past samples. We'll build an *adaptive predictor*, which actually makes a reasonable guess at what the next sample will be based on past input samples. This leads directly to an adaptive noise-reduction system. Later, we'll discover how this technology is applied to compression of voice and other signals for digital transmission over the telephone network.

The Adaptive Interference Canceler¹

Imagine that we have some input signal $x(k)$, and we want to filter it to enhance its sinusoidal content. The quantity $x(k)$ is just the discrete sample of continuous input signal x taken at time k . In the case of a CW signal, all that's required is a band-pass filter (BPF) centered at the desired frequency. We know the output will take the form of a sine wave, and that only its amplitude will change.

So we set up an FIR filter structure, and set the initial filter coefficients $h(k)$ to zero. Then we set up an error-measurement system to compare a sine wave $d(k)$ with the output of the filter, $y(k)$. See Fig 1. The reference input $d(k)$ is the same frequency we expect the CW input signal to be. The difference output $e(k)$ is known as the error signal. Then imagine we have some algorithm to adjust the filter coefficients so that the error $e(k)$ is

reduced at each sample time. Think of the algorithm as some person who is "eye-balling" the error signal on an oscilloscope and has their hands on the filter controls. If they can minimize the error, then the filter will have converged to a BPF centered at the frequency of $d(k)$.

We can already deduce that the speed and accuracy of convergence is going to depend on how well the person analyzes the error data. If it's difficult to tell that a sine wave is present, then adjusting the filter will be difficult, as well. Further, if the sampling rate is high enough, the person can't keep up; they can check the error only so often, or they can take long-term averages of the error.

Using the typical processes of the human mind, the person will soon discover that if they turn the controls the wrong way, the filter will diverge from the desired response; ie, the error increases. This information is used to reverse the direction of adjustment; the person will then turn the controls the other way. They will soon discover they are on a *performance surface* that has an "uphill" and a "downhill," and they know they want to go only downhill.

So they thrash about with the controls, sometimes making mistakes and heading the wrong way, but ultimately making headway overall down the hill. At some point, the error gets very small, and they know they're near the "bottom of the bowl." Once at the bottom, it's uphill no matter which way they go! So, they continue flailing about, but always staying near the bottom. They have successfully achieved the goal: Minimization of the total error $e(k)$. This story is analogous to aligning an analog BPF with an adjustment tool.

After doing this several times, the person finds that certain rules help them speed up the process. First, there is a relationship between the total error and the amount they need to tweak

the controls. If the total error is large, then a large amount of tweaking must be done; if small, then it's better to make small adjustments to stay near the bottom. Second, there is a correlation between the error signal $e(k)$, the input samples $x(k)$ and the filter coefficient set $h(k)$ they need to adjust.

Derivation of algorithms that provide for the quickest descent down the hill is a very long and tedious exercise in linear algebra. Let's just say the person goes to school, becomes an expert in matrix mathematics and discovers that one of the fastest ways down the hill is to make adjustments at sample time k according to:

$$h_{k+1} = h_k + 2\mu e_k x_k \quad (\text{Eq 1})$$

This is the LMS algorithm. It was developed by Widrow and Hoff³ in the late 1950s.

Properties of the Adaptive Interference Canceler

Now we have our adaptive interference canceler. See Fig 2. Note that both the desired output $y(k)$ and the undesired $e(k)$ are available. This is nice in case we want to exchange roles, to accept only the incoherent input signals and reject coherent ones. Such is the case for an adaptive notch filter, treated further below. This doesn't change the algorithm, however. Quantities of interest in this system are the adjustment error near the bottom of the performance surface, and the speed of adaptation.

One of the first things we discover about the LMS algorithm is that the speed of adaptation and the total misadjustment are both directly proportional to μ . We select its value, which ranges from 0 to 1, to set the desired properties. Note that there is a trade-off between speed and misadjustment. Large values of μ result in fast convergence, but large adjustment errors.

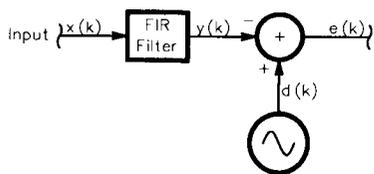


Fig 1—An adaptive modeling system.

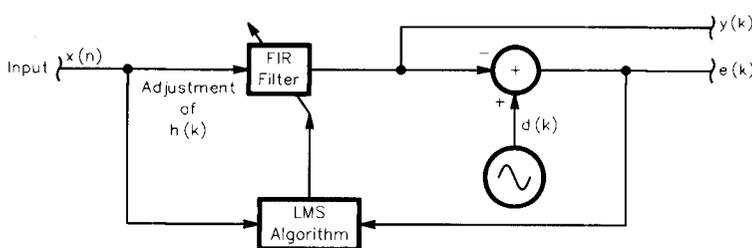


Fig 2—An adaptive interference canceler.

Also note that the number of filter coefficients $h(k)$ has a bearing on both of these performance parameters. It turns out that the total amount of misadjustment is directly proportional to the number of filter coefficients, and this places a limitation on the complexity of the filter. In addition, as the filter grows in length, the total delay through the filter grows proportionately. The delay through an FIR filter of length L is equal to:

$$T_{FIR} = \frac{LT_s}{2} \quad (\text{Eq 2})$$

where T_s is the sample time, and this may become unacceptable under certain conditions.

Attempts may be made to adjust the factor μ on an adaptive basis by using a value which changes in proportion to the total error $e(k)$. A large value is selected initially to obtain rapid convergence, then it's decreased to minimize the total misadjustment as we approach the steady-state solution. This works fine as long as the characteristics of the input signal don't change rapidly.

The Adaptive Interference Canceler Without an External Reference—The Adaptive Predictor

In the above example of a CW signal, we knew what to expect at the output: A sine wave of known frequency. What happens when we don't know much about the nature of the desired signal, except that it's coherent in the time domain? A number of circumstances arise wherein the only fact known about the desired signal is that it is distinguishable from noise in some way; ie, that it's periodic. It might seem at first that adaptive processing can't be applied. But if a delay, z^{-n} is inserted in the *primary input* $x(k)$ to create the reference input $d(k)$, periodic signals may be detected and, therefore, enhanced. See Fig 3. This delay is akin to an autocorrelation offset, and it represents the time differential used to compare past input samples with the present ones. The amount of delay must be chosen so that the desired components in the input signal autocorrelate, and the undesired components do not.

This system is an adaptive predictor. The predictable components are enhanced, while the unpredictable parts are removed. Fig 4 shows the result of an actual experiment using a sine wave buried in noise as the input. The input BW is 3 kHz, and the input SNR = 0 dB. For any given value of μ ,

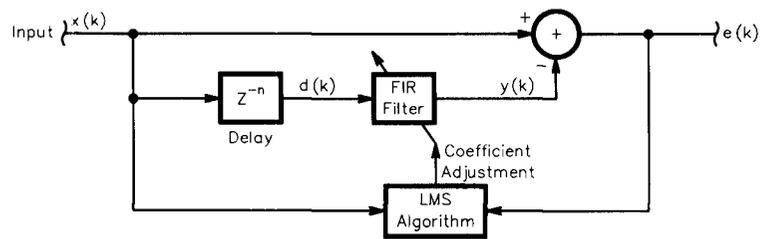
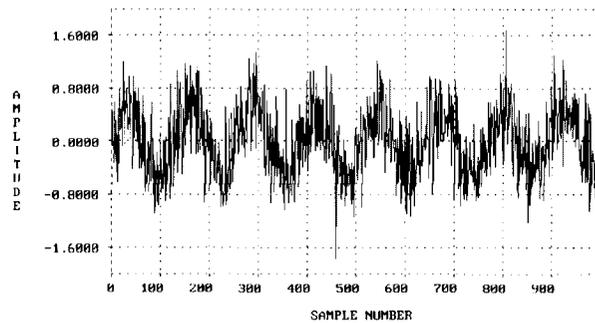
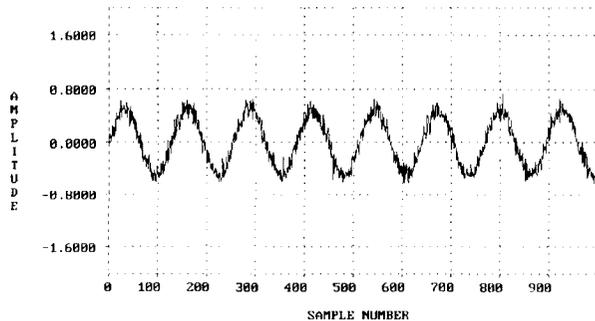


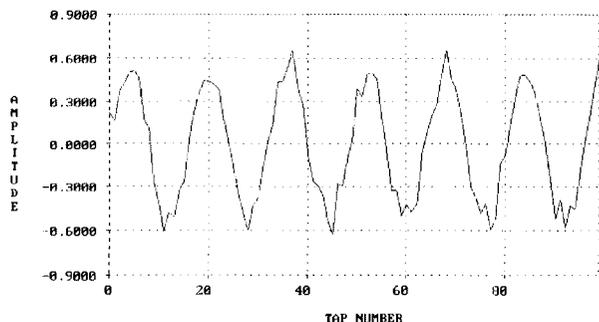
Fig 3—An adaptive predictor.



(A)



(B)



(C)

Fig 4—(A) A noisy sine wave. (B) Filtered output. (C) The adaptive filter's impulse response.

the filter converged on the optimal solution fastest when the delay was set roughly equal to the filter delay as defined in Eq 2. Note that the filter's impulse response is also a sinusoid. We find that the filter's BW¹ is:

$$BW = \frac{2\mu A^2}{T_s} \quad (\text{Eq 3})$$

where A is a long-term average of the amplitude of the input $x(k)$. So the speed of adaptation and the NR effectiveness are proportional to μ and to the amplitude of the input signal. In the example, $\mu = 0.005$, $A = 1$, and $T_s = (15 \times 10^3)^{-1}$. The SNR improvement is therefore:

$$\begin{aligned} \Delta SNR &= 10 \log \left(\frac{3 \text{ kHz}}{BW} \right) \\ &= 10 \log \left(\frac{(3 \times 10^3) T_s}{2\mu A^2} \right) \\ &\approx 13 \text{ dB} \end{aligned} \quad (\text{Eq 4})$$

Alternatively, the unpredictable components $e(k)$ may be taken as the output. This forms an adaptive notch filter. Say we have a desired voice signal corrupted by the presence of a single interfering tone or carrier. This is a very common situation on today's HF ham bands! We can set the autocorrelation delay and variable μ so that the steady tone is predictable, and the rapidly changing voice characteristics are not. The filter will converge to the solution that removes the tone and leaves the voice signal virtually unscathed.

The BW of the notch is the same as in Eq 3, but its depth is determined only by numerical-accuracy effects in the DSP system. When adaptive filters with many taps are used, multiple tones may be notched. See Fig 5. In this experiment, several nonharmonically related tones, plus noise, are used as the input. The filter's response converges to notch them all, leaving only noise at the output. In this case, the undesired components are large compared to the desired components. When the undesired signal level is low, there might not be enough thrashing about on the performance surface for us to find our way down the hill. Adding artificial noise to satisfy this condition is tempting, but it turns out that we can alter the algorithm to improve the situation without actually adding noise. Such additional terms in the algorithm are referred to as *leakage* terms.

"Leaky" LMS Algorithms

The unique feature of *leaky* LMS algorithms is a continual "nudging" of the filter coefficients toward zero. The effect of the leakage term is striking, especially when applied to NR of voice signals. The SNR improvement increases because the filter coefficients tend toward a lower throughput gain in the absence of desired input components. More significantly, the leakage helps the filter adapt under low SNR conditions—the very conditions when NR is needed most.

One way to implement leakage is to add a small constant of the appropriate sign to each coefficient at every sample time. This constant is positive for negative coefficients, and negative for positive coefficients:

$$h_{k+1} = h_k + 2\mu e_k x_k - \lambda [\text{sign}(h_k)] \quad (\text{Eq 5})$$

The value of λ can be altered to vary the amount of leakage. Large values prevent the filter from converging on *any* input components, and things get very quiet indeed! Small values are useful in extending the noise floor of the system. In the absence of coherent input signals, the coefficients linearly

move toward zero; during convergent conditions, the total misadjustment is increased to at least λ , but this isn't usually serious enough to affect received signal quality.

An alternate way to implement leakage is to scale the coefficients at each sample time by some factor, γ , thus also nudging them toward zero:

$$h_{k+1} = \gamma h_k + 2\mu e_k x_k \quad (\text{Eq 6})$$

For values of γ just less than one, leakage is small; values near zero represent large leakage and again prevent the filter from converging. This realization of the leaky LMS exhibits a logarithmic decay of coefficients toward zero, which may be advantageous under certain circumstances. It can be shown¹ that the leaky LMS is equivalent to adding normalized noise power to the input $x(k)$ equal to:

$$\sigma^2 = \frac{1-\gamma}{2\mu} \quad (\text{Eq 7})$$

Note that the leaky LMS algorithm must adapt to "survive." Were the factor μ suddenly set to zero, the coefficients would die away toward zero and never recover. Therefore, it's unwise to use these algorithms with adaptive

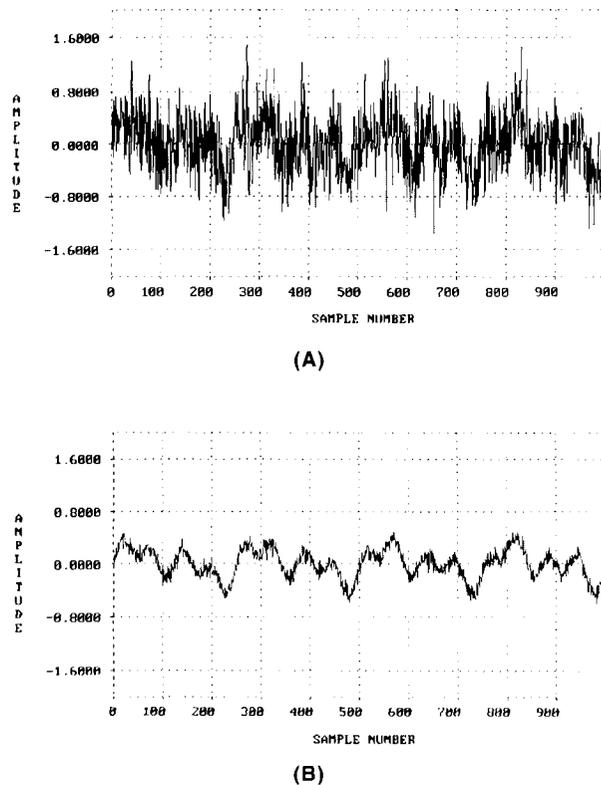


Fig 5— (A) Multiple noisy sine waves. (B) Filtered output.

values of μ . Although values for γ and μ greater than one have been tried, the inventors refer to these procedures as “the dangerous LMS algorithm.” Enough said.

Next, I’ll describe an even more powerful NR tool, the Fourier transform. Although it’s generally more processing-intensive than methods described thus far, the results can be impressive. First, however, let’s look at the great man whose work is so inexorably bound to modern electronic communications.

Joseph Fourier—A Legacy of Genius

Joseph Fourier (1768 to 1830) was born the son of a tailor.⁴ Educated by monks in a military school, Fourier apparently thought only the army or the church could provide him a career. Despite a recommendation from the famous mathematician Legendre, the army rejected his application to the artillery; he opted instead for a religious life. Mathematics had been his primary scientific study since an early age, and at the onset of the French Revolution, he was appointed by the monks to the principal chair of mathematics at Auxerre. There, Fourier met Napoleon, who often attended lectures at the major universities.

When Napoleon organized an expedition to Egypt in 1798, Fourier was asked to join. After three years there in the capacity of an engineer, he collaborated with Napoleon to produce the *Description of Egypt*, which established his literary prowess and eventually won him election to the French Academy.

On his return to France, he was granted a comfortable governmental position, which gave him free time to pursue his mathematical interests. In 1807, he submitted to the Academy his first paper describing the motion of heat in solid bodies. In 1812, he was awarded the prize for scientific accomplishment for his complete explanation of the effect—the judges were Laplace, Legendre and Lagrange! His place in history was thereby confirmed.

Fourier was elected a member of the Academy of Sciences in 1817, and to the French Academy in 1822. He died while still in government service in 1830. He could scarcely have imagined what impact his work has had in the field of electronics, especially DSP.

The Fourier Transform and Its Inverse

The relationship Fourier discovered between the application of heat to a solid body and its propagation has direct analogy to the behavior of electri-

cal signals as they pass through filters or other networks. The laws he found represent the connection between the time- and frequency-domain descriptions of signals. They form the basis for DSP spectral analysis, and therefore they are useful in digging signals out of noise, as we’ll see below.

The *Fourier Transform*⁵ of some continuous signal x_t is expressed as:

$$X_\omega = \int_{-\infty}^{\infty} e^{-j\omega t} x_t dt \quad (\text{Eq 8})$$

It’s obtained by making the following assumptions: x_t is a continuous, periodic function of time; and any continuous, periodic function of time can be expressed as the superposition (integral) of sines and cosines. Recall the Euler identity:

$$e^{j\omega t} = \cos \omega t + j \sin \omega t \quad (\text{Eq 9})$$

and observe that when the real and imaginary parts are separated, Eq 8 produces coefficients a_ω and b_ω :

$$a_\omega = \int_{-\infty}^{\infty} x_t \cos(\omega t) dt \quad (\text{Eq 10})$$

$$b_\omega = -\int_{-\infty}^{\infty} x_t \sin(\omega t) dt \quad (\text{Eq 11})$$

X_ω is just the sum of these terms as a complex pair:

$$X_\omega = a_\omega + j b_\omega \quad (\text{Eq 12})$$

The coefficients yield the amplitude and phase of the signal x_t at the frequency ω :

$$A_\omega = (a_\omega^2 + b_\omega^2)^{1/2} \quad (\text{Eq 13})$$

$$\phi_\omega = \tan^{-1} \left(\frac{b_\omega}{a_\omega} \right) \quad (\text{Eq 14})$$

Working in reverse, we can reconstruct x_t by integrating X_ω for all values of ω :

$$x_t = \frac{1}{2\pi} \int_{-\pi}^{\pi} X_\omega e^{j\omega t} d\omega \quad (\text{Eq 15})$$

This is known as the *Inverse Fourier Transform* of X_ω . The limits of integration in this case are finite because $e^{j\omega t}$ and X_ω are themselves continuous, periodic functions of ω , repeating with period 2π . This follows from Eq 8, since:

$$e^{-j(\omega+2\pi)t} = e^{-j\omega t} \quad (\text{Eq 16})$$

We had to use infinite limits in Eq 8, because we made no assumptions about the length of period of x_t . I.e., we didn’t know how far in time to look until x_t began repeating itself. If we had some knowledge about the periodicity of x_t , then we could restrict the integration limits without major deleterious effects.

Using infinite integration limits,

components in x_t not at frequency ω do not affect the result. In the DSP world, however, we deal with discrete samples of the amplitude of x_t , which we’ll refer to as $x(n)$. The discrete Fourier Transform (DFT) gives us an expression equivalent to Eq 8 for sampled signals. We’ll see that in this form, issues of frequency resolution arise because infinite evaluation intervals aren’t practical.

The Discrete Fourier Transform (DFT)

The discrete-sample equivalent of Eq 8 is expressed as:

$$X(\omega) = \sum_{n=-\infty}^{\infty} e^{-j\omega n} x(n) \quad (\text{Eq 17})$$

where n is the sample number. As alluded to above, we can’t compute this sum over an infinite number of samples. In discrete spectral analysis, many sums must be obtained; only a short time is available for these computations before the next iteration must be performed.

So let’s say that $x(n)$ has a period less than some number of samples N . We limit our summation to that number of samples, and as a consequence, the results are available only in integer multiples of the fundamental frequency $2\pi / N$:

$$X(k) = \sum_{n=0}^{N-1} e^{-\frac{2\pi j k n}{N}} x(n) \quad (\text{Eq 18})$$

This is the DFT for a frequency proportional to k . It turns out there are just N frequencies available that are integer multiples of $2\pi / N$. This is because $e^{-2\pi j k n / N}$ is periodic with period N . It’s a fact of DSP life that samples taken at discrete times transform to samples at discrete frequencies using the DFT.

DFTs are normally computed for N evenly spaced values of k . To relate the normalized analysis frequency k / N to the sampling frequency $1 / T_s$, we can write:

$$f_k = \frac{k}{NT_s}, \text{ for } k < \frac{N}{2} \quad (\text{Eq 19})$$

This is the actual frequency, in hertz, of the DFT represented by $X(k)$.

The idea is that if we can analyze our input signal at many frequencies, and exclude those results or *bins* not meeting certain criteria, we can eliminate undesired signals. Filters can be implemented by rejecting signals outside the frequencies of interest. And noise reduction can be accomplished by eliminating bins for which a preset amplitude threshold is not met.

The DFT In Noise-Reduction Systems

The efficacy of the DFT is that it evaluates the amplitude and phase of some particular frequency component to the exclusion of others. As far as we can reduce the *resolution bandwidth* (BW) of our frequency-specific measurements using the DFT, we can eliminate noise. Ie, the finer the frequency resolution, the less noise we are including in each bin, so any coherent signal in the measurement BW has an improved signal-to-noise ratio (SNR). Finer resolution is obtained by increasing the number of bins, N .

Shown in Fig 6 is the result of a DFT analysis of a low-level 1 kHz sine wave buried in noise. In the 3 kHz BW of interest, the noise power is just equal to the signal power at the input. The SNR is therefore 0 dB. The sampling frequency is 15 kHz, and $N = 1500$ for this DFT. Because the resolution BW of our DFT is:

$$f_{res\ BW} = \frac{1}{NT_s} = 10\text{ Hz} \quad (\text{Eq 20})$$

the SNR improves at the bin centered on 1 kHz by the factor:

$$SNR_{10\text{ Hz}} = 10 \log \left(\frac{3\text{ kHz}}{10\text{ Hz}} \right) \approx 24.8\text{ dB} \quad (\text{Eq 21})$$

The sine wave stands out clearly above the noise. We can modify the results of our spectral analysis by setting a threshold, below which we set the DFT results to zero. See Fig 7. We then convert this modified frequency-domain picture back to time-domain samples using the *inverse DFT* (DFT^{-1}):

$$x'(n) = \frac{1}{N} \sum_{k=0}^{N-1} e^{2\pi jkn} X'(k) \quad (\text{Eq 22})$$

Since the bins containing only noise have been zeroed in the modified transform samples $X'(k)$, the output signal $x'(n)$ now has an improved SNR of 24.8 dB! The remaining noise is centered in a 10 Hz BW around the 1 kHz tone. See Fig 8. Note that we have to compute 1500 DFT bins at each sample time to get this result. Any bins that are zeroed obviously make the conversion using the DFT^{-1} faster, since they need not be computed.

Setting the Threshold

Setting the cutoff threshold is critical to this noise-reduction method, because we may inadvertently exclude low-energy components, which are actually part of the desired signal. The simplest solution has the operator

set it manually. One just “mows the grass” to whatever depth produces a pleasing result.

An automatic system seems possible, but we would need to make some assumptions about the nature of the desired signal or signals. The requirements for a voice signal, for example, might be quite different from those for a RTTY or CW signal. Selective fading and multipath effects ultimately limit the usefulness of any automatic system.

Since we control which bins get ex-

cluded, based not only on their amplitudes, but also their frequencies, the DFT gives us a way to include custom hand-pass filter banks, similar to a graphic equalizer.

Computation of the DFT and Limitations on Accuracy

Components in $x(n)$ at frequencies other than k / N skew the result and significantly limit resolution BW. The limited summation range broadens the spectral line width, even for a single

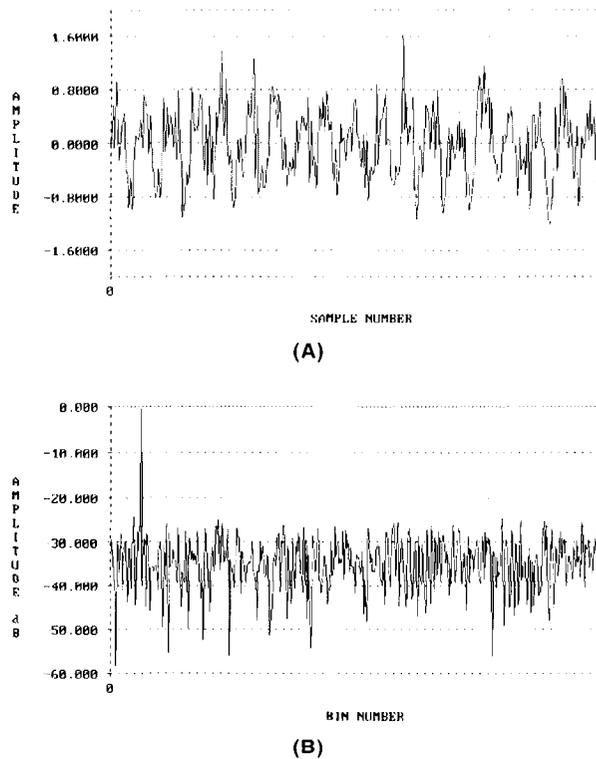


Fig 6—(A) A noisy sine wave. (B) Its DFT amplitude-versus-frequency.

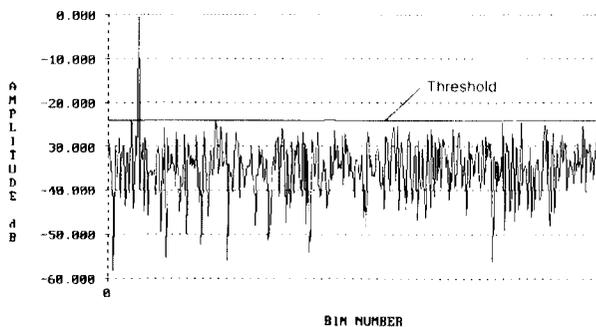


Fig 7—Applying a noise-reduction threshold.

input frequency as shown in Fig 9. The *side lobes* of the spectral broadening are evident in this diagram. Many advanced methods have been introduced to deal with this situation. They are based on either reduction of the computation time to obtain more bins, or modification of the data, or both.

Of help in reducing the broadening effect is the technique of *windowing* the input data. The data block is multiplied by a *window function*, then used as input to the DFT algorithm. Examples of window functions and their DFTs are shown in Fig 10. The rectangular window is equivalent to not using a window at all, since each input sample is multiplied by unity. The other window functions achieve various amounts of side-lobe reduction.

Also impacting the accuracy of our results are the familiar truncation and rounding effects previously discussed.⁶ Their influence on the DFT is treated further below.

Other advanced spectral-estimation techniques have been developed over the years,^{7, 8} some of which produce excellent results. Most use the DFT in some form. Because computation time is critical in embedded systems and Fourier analysis is so important to so many fields, much time and effort has been expended to find efficient DFT-computing algorithms.

In the years before computers, reduction of computational burden was extremely desirable, because computation was done by hand! Many excellent mathematicians, including Runge,⁹ applied their wits to the problem of calculating DFTs more rapidly than the direct form of Eq 18. They recognized that the direct form required N complex multiplications and additions per bin, and N bins were to be calculated, for a total computational burden proportional to N^2 . The first breakthrough was achieved when they realized that the complex exponential $e^{-2\pi jkn/N}$ is periodic with period N , so a reduction in computations was possible through the symmetry property:

$$e^{\frac{-2\pi jk(N-n)}{N}} = e^{\frac{2\pi jkn}{N}} \quad (\text{Eq 23})$$

This led to the construction of algorithms that effectively broke any N DFT computations of length, N , down into N computations of length $\log_2(N)$. Thus, the computational burden was reduced to $N \log_2(N)$. Because even this much work wasn't practical by hand, the usefulness of the discovery was largely overlooked until Cooley and Tukey¹⁰ picked up the gauntlet in the 1960s.

The Fast Fourier Transform (FFT)

Let's look at how FFT algorithms are derived from the repetitive nature of the complex exponential (Eq 23) above, and how they're implemented using *in-place* calculations. Since we're going to be dragging around a lot of complex exponentials, we'll adopt the simplified notation of Oppenheim and Schaffer⁵ where:

$$e^{\frac{-2\pi jkn}{N}} = W_N^{kn} \quad (\text{Eq 24})$$

To exploit the symmetry referred to, we have to break the DFT computation of length, N , down into successively smaller DFT computations. This is done by *decomposing* either the input or the output sequence. Algorithms wherein the input sequence $x(n)$ is decomposed into successively smaller subsequences are called *decimation-in-time* algorithms.

Let's begin by assuming that N is an integral power of two. That is, for a whole number p :

$$N = 2^p \quad (\text{Eq 25})$$

Next, we break the input sequence into two subsequences, one consisting of the even-numbered samples, and the other of the odd-numbered samples. For some index r , $n = 2r$ for n even, and $n = 2r + 1$ for n odd. Now, with Eq 18 in mind, we can write:

$$\begin{aligned} X(k) &= \sum_{r=0}^{\frac{N}{2}-1} W_N^{2rk} x(2r) + \sum_{r=0}^{\frac{N}{2}-1} W_N^{(2r+1)k} x(2r+1) \\ &= \sum_{r=0}^{\frac{N}{2}-1} (W_N^2)^{rk} x(2r) + W_N^k \sum_{r=0}^{\frac{N}{2}-1} (W_N^2)^{rk} x(2r+1) \end{aligned} \quad (\text{Eq 26})$$

To further simplify, we can use:

$$\begin{aligned} W_N^2 &= e^{\frac{(2)(-2\pi j)}{N}} \\ &= e^{\frac{-2\pi j}{\frac{N}{2}}} \\ &= W_{\frac{N}{2}} \end{aligned} \quad (\text{Eq 27})$$

and so now:

$$X(k) = \sum_{r=0}^{\frac{N}{2}-1} W_{\frac{N}{2}}^{rk} x(2r) + W_N^k \sum_{r=0}^{\frac{N}{2}-1} W_{\frac{N}{2}}^{rk} x(2r+1) \quad (\text{Eq 28})$$

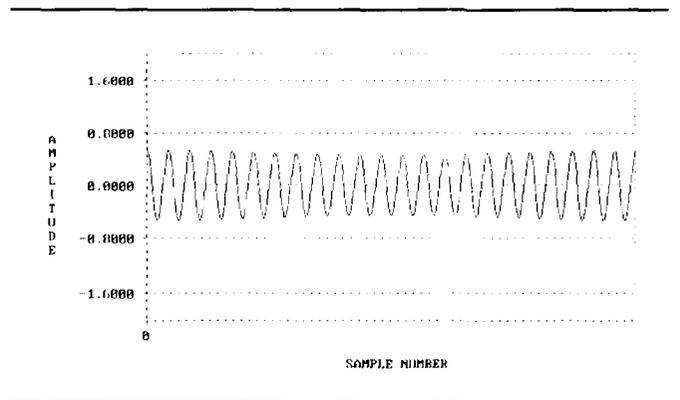


Fig 8—The reconstructed sine wave.

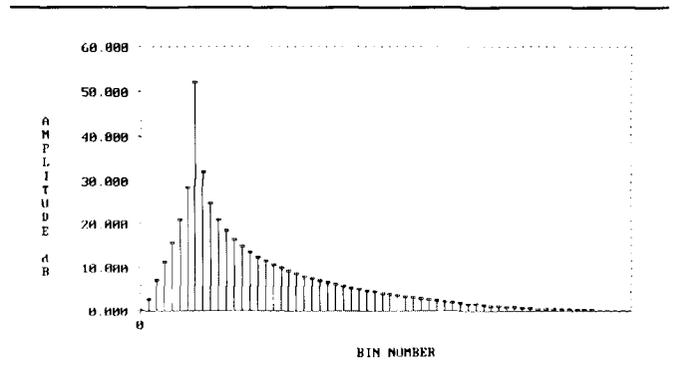
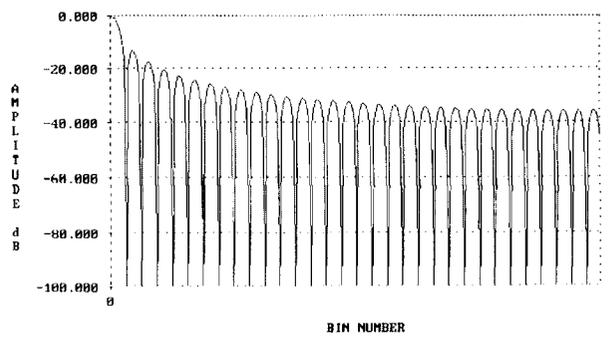
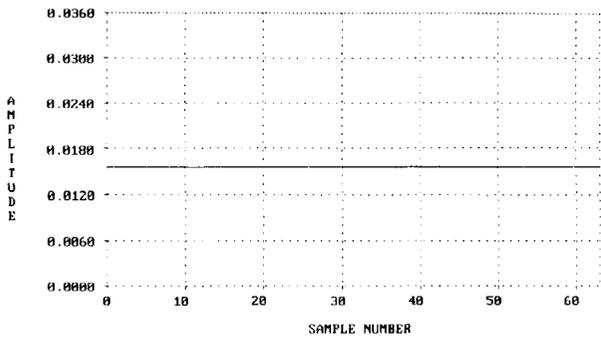
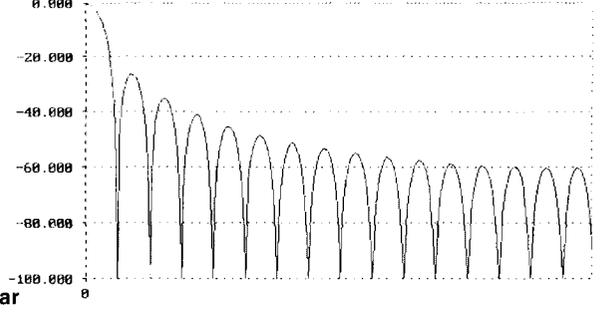
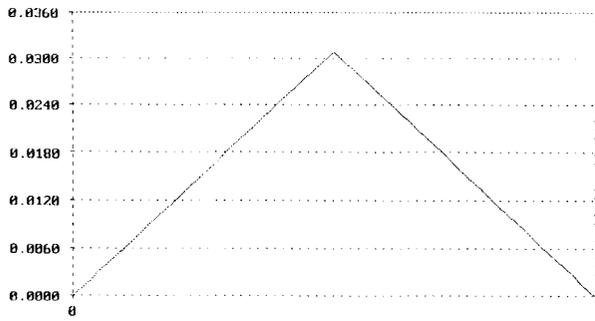


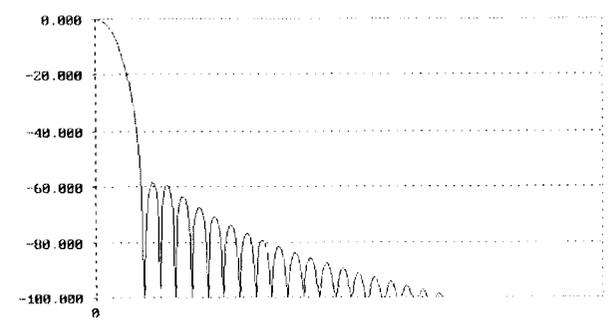
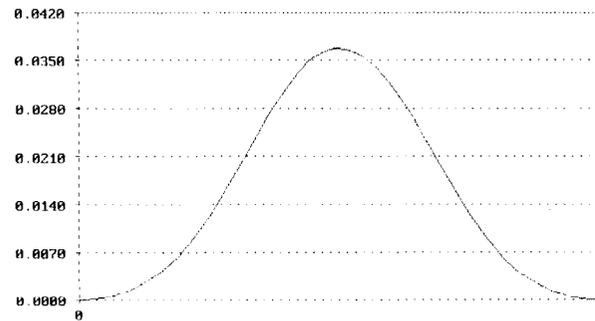
Fig 9—DFT of a noise-free sine wave, showing side lobes.



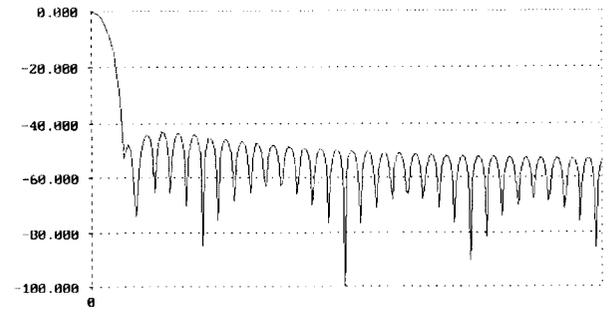
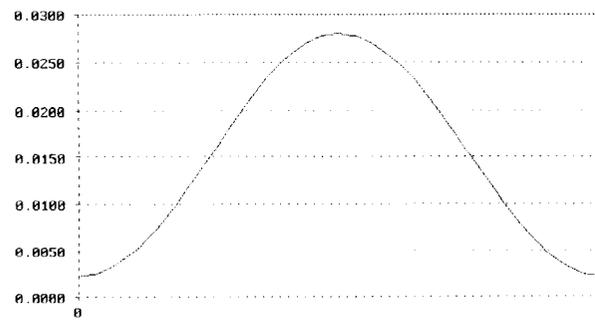
Rectangular



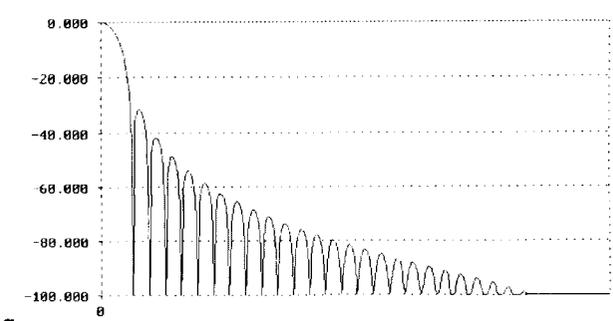
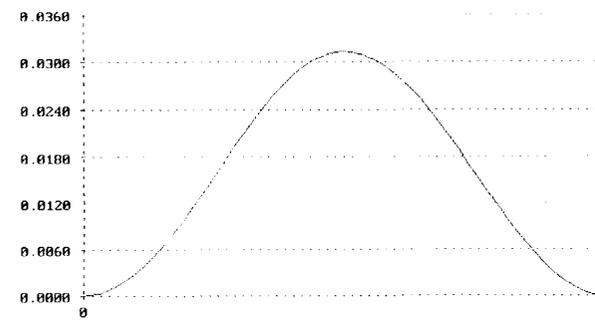
Triangular



Blackman



Hamming



Hanning

Fig 10—Various window functions and their Fourier transforms.

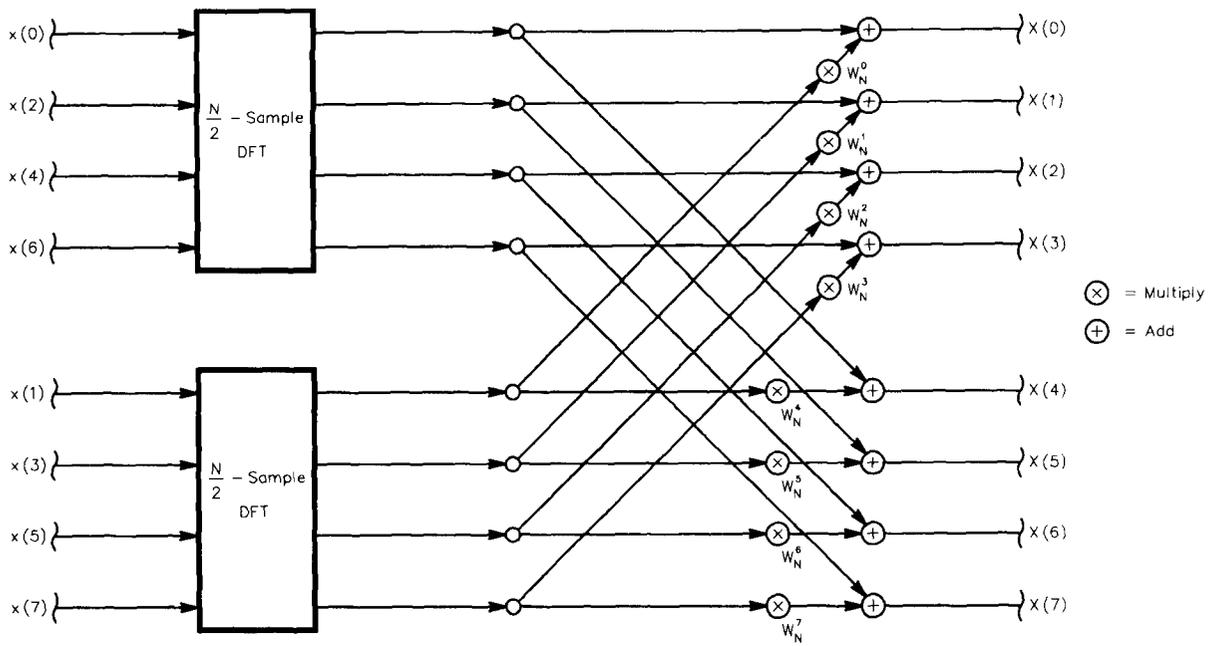


Fig 11—An eight-sample DFT as two four-sample DFTs.

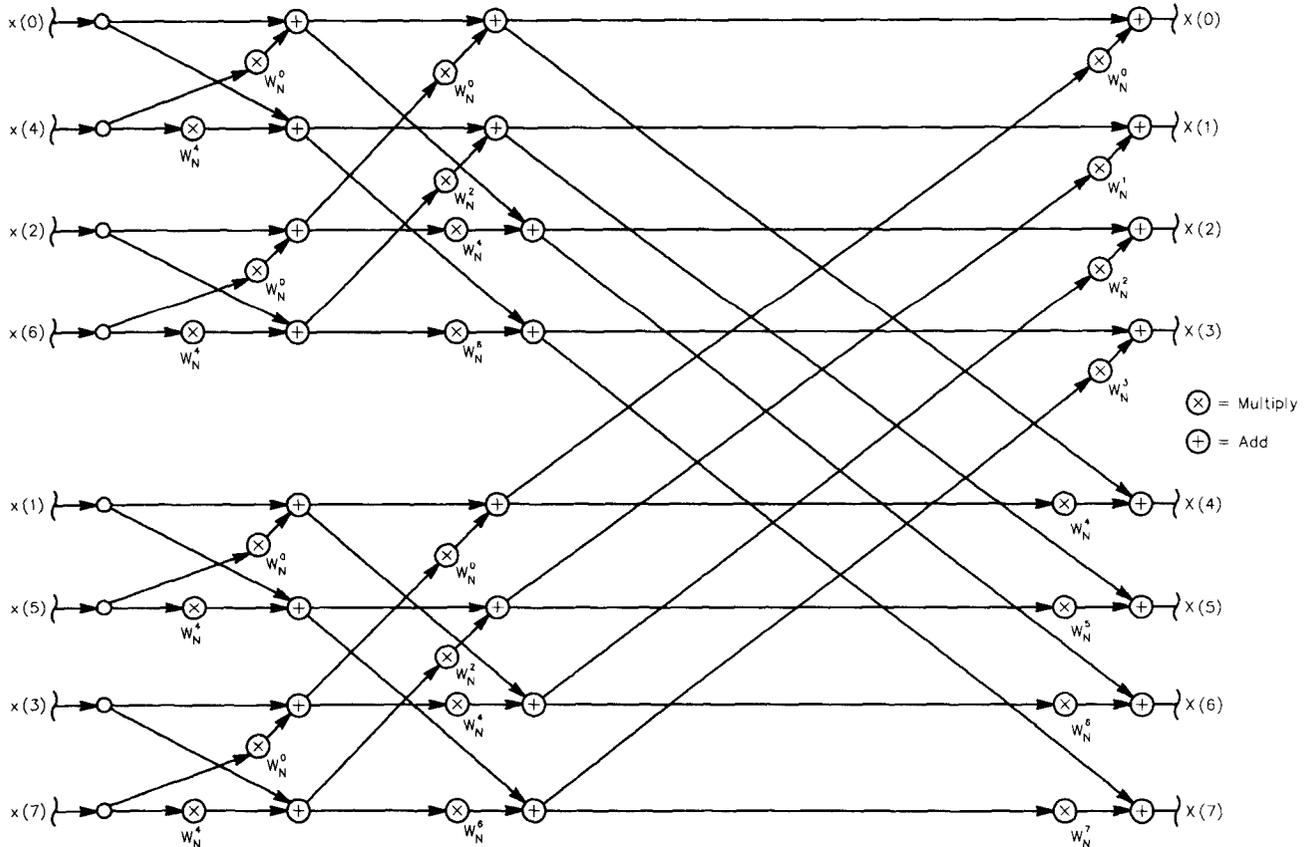


Fig 12—An eight-sample FFT.

It's evident that Eq 28 represents two $N/2$ -sample DFT calculations. It has in fact eased the computational load, since it requires:

$$N + 2\left(\frac{N}{2}\right)^2 = N + \frac{N^2}{2} \quad (\text{Eq 29})$$

complex multiplications and additions. Note that W_N^k is a function only of k , and is therefore a set of constants. We have reduced calculations by the factor:

$$\frac{N + \frac{N^2}{2}}{N^2} = \frac{1}{N} + \frac{1}{2} \quad (\text{Eq 30})$$

which for large N is nearly a two-fold reduction.

This is where flow charts become useful, so Eq 28 is used to produce Fig 11, an example of an eight-sample DFT calculation as broken into two four-sample calculations.

Carrying this idea further, since N is an integral power of two, we can break each of these $N/2$ -sample DFT calculations down into separate $N/4$ -sample calculations. Then we break each of those into separate $N/8$ -sample calculations, and so on, until

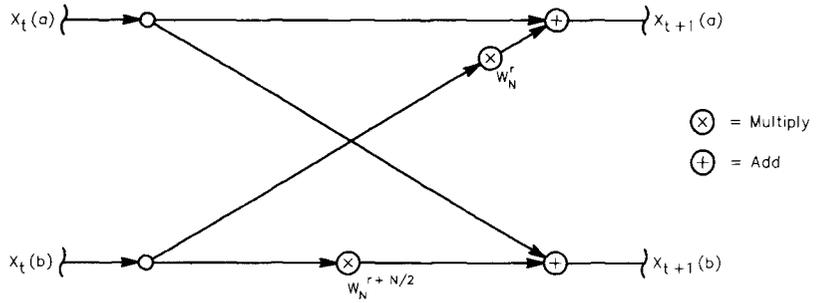


Fig 13—Butterfly calculation for a decimation-in-time FFT.

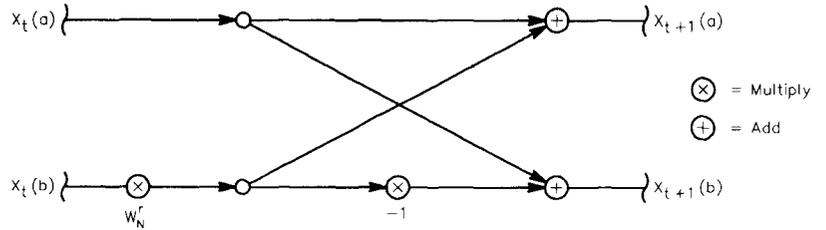


Fig 14—Modified butterfly.

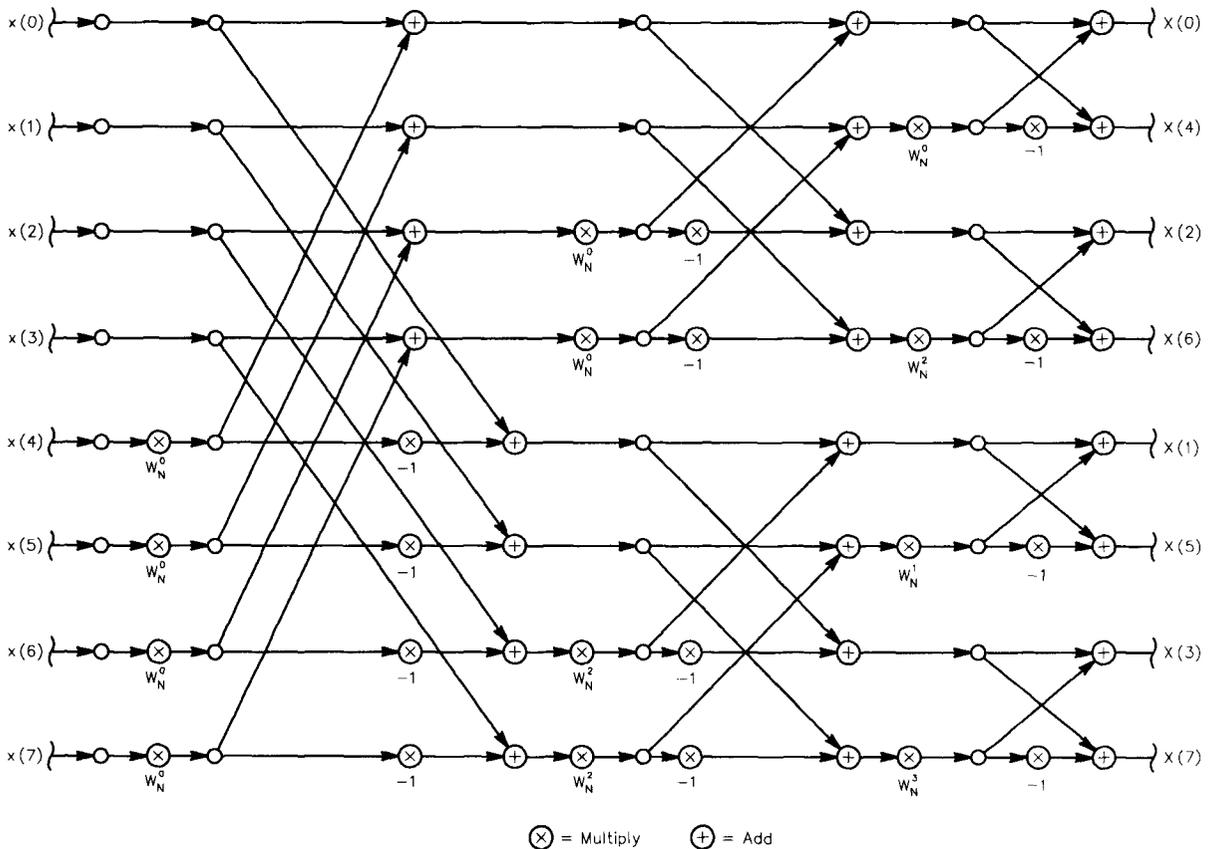


Fig 15—A decimation-in-time FFT with different input-output order and using modified butterflies.

we're left with only two-sample DFT calculations. From this discussion, it's obvious that we can break things down only $\log_2 N - 1$ times until we get to the two-sample level. When this has been done for our eight-sample DFT, the resulting flow chart is a complete FFT, as shown in Fig 12.

In-Place Calculation

The figure shows that starting with eight input samples, eight outputs are generated. Each stage requires N complex multiplications and additions, and there are $\log_2 N$ stages; hence, the total burden is $N \log_2 N$. Further, each stage transforms N complex numbers into another set of N complex numbers. This suggests we should use a complex array of dimension N to store the inputs and outputs of each stage as we go along. Finally, an examination of the branching of terms in the diagram reveals that pairs of intermediate results are linked by pairs of calculations like the one shown in Fig 13. Because of the appearance of this diagram, it is known as a *butterfly computation*.

This computational arrangement requires two complex multiplications and additions. For each butterfly, the intermediate results are in the form:

$$X_{t+1}(a) = X_t(a) + W_N^r X_t(b) \quad (\text{Eq 31})$$

$$X_{t+1}(b) = X_t(a) + W_N^{r+\frac{N}{2}} X_t(b)$$

where t represents the stage number of the calculations, and a and b are the branch numbers. Note that:

$$W_N^{r+\frac{N}{2}} = W_N^r W_N^{\frac{N}{2}} \quad (\text{Eq 32})$$

and:

$$\begin{aligned} W_N^{\frac{N}{2}} &= e^{-j\frac{2\pi}{N} \cdot \frac{N}{2}} \\ &= e^{-j\pi} \\ &= -1 \end{aligned} \quad (\text{Eq 33})$$

So now Eq 31 can be written:

$$\begin{aligned} X_{t+1}(a) &= X_t(a) + W_N^r X_t(b) \\ X_{t+1}(b) &= X_t(a) - W_N^r X_t(b) \end{aligned} \quad (\text{Eq 34})$$

The equivalent flow diagram is shown as Fig 14. Now that's slick, since we just reduced the total multiplications by an additional factor of two! The total burden is now proportional to $(N/2) \log_2 N$.

These calculations can be performed in place because of the one-to-one correspondence between the inputs and outputs of each butterfly. The nodes are connected horizontally on the flow diagrams. The data from locations a and b are required to compute the new data to be stored in those same locations, hence only one array is required during calculation.

An interesting result of our decomposition of the input sequence $x(n)$ is that in the FFT calculation of Fig 12, the input samples are no longer in ascending order. In fact, they are indexed in *bit-reversed* order. It turns out that this is a necessity for doing the computations in place. To see why this is so, let's review what we did in the derivation above.

First, we separated the input samples into evens and odds. So naturally, all the even samples appear in the top half, and the odds in the bottom half. The index n of an odd sample has its least-significant bit (LSB) set; an odd sample's LSB is cleared.

Next, we separated each of these sets into their even and odd parts, which can be done by examining the second LSB in the index n . This process was repeated until we had N subsequences of unity length. It resulted in the sorting of the input data in a bit-reversed way. This isn't very convenient for us in setting up the computation, but at least the output arrives in the correct order!

Alternative Forms

It's possible to rearrange things such that the input is in normal order, and the output is in bit-reversed order. See Fig 15. In-place computation is still possible. While it's even possible to arrange things such that the input *and* the output are in normal order, that makes in-place computation impossible.

We obtained a decimation-in-time algorithm by decomposing the input sequence $x(n)$ above. It's also possible to decompose the output sequence in the same way, with the same computational savings. Algorithms obtained in this way are called *decimation-in-frequency*.

To begin this derivation, we again separate the input sequence $x(n)$ into two parts, but this time, they are the first and second halves:

$$\begin{aligned} X(k) &= \sum_{n=0}^{\frac{N}{2}-1} W_N^{nk} x(n) + \sum_{n=\frac{N}{2}}^{N-1} W_N^{nk} x(n) \\ &= \sum_{n=0}^{\frac{N}{2}-1} W_N^{nk} x(n) + \left(W_N^{\frac{kN}{2}} \right)^{\frac{N}{2}-1} \sum_{n=0}^{\frac{N}{2}-1} W_N^{nk} x\left(n + \frac{N}{2}\right) \end{aligned} \quad (\text{Eq 35})$$

Using the relation:

$$W_N^{\frac{kN}{2}} = (-1)^k \quad (\text{Eq 36})$$

and combining the summations, we get:

$$X(k) = \sum_{n=0}^{\frac{N}{2}-1} W_N^{nk} \left[x(n) + (-1)^k x\left(n + \frac{N}{2}\right) \right] \quad (\text{Eq 37})$$

Now let's break the output sequence into even and odd parts, again using index $r < N/2$ as above:

$$\begin{aligned} X(2r) &= \sum_{n=0}^{\frac{N}{2}-1} W_N^{2rm} \left[x(n) + x\left(n + \frac{N}{2}\right) \right] \\ X(2r+1) &= W_N^r \sum_{n=0}^{\frac{N}{2}-1} W_N^{2rm} \left[x(n) - x\left(n + \frac{N}{2}\right) \right] \end{aligned} \quad (\text{Eq 38})$$

Since:

$$W_N^{2m} = W_N^m \quad (\text{Eq 39})$$

Eq 38 can be written:

$$\begin{aligned} X(2r) &= \sum_{n=0}^{\frac{N}{2}-1} W_N^{rm} \left[x(n) + x\left(n + \frac{N}{2}\right) \right] \\ X(2r+1) &= W_N^r \sum_{n=0}^{\frac{N}{2}-1} W_N^{rm} \left[x(n) - x\left(n + \frac{N}{2}\right) \right] \end{aligned} \quad (\text{Eq 40})$$

and the result is again two $N/2$ -sample DFT calculations.

Proceeding in direct analogy to the decimation-in-time algorithm above, decomposition continues until we have only two-sample DFT calculations left, and for the case $N = 8$, the flow diagram appears as shown in Fig 16.

Note that this algorithm can be performed as butterflies, and that calculation in place is possible just as before. The butterflies are a bit different now, however, as depicted in Fig 17. The corresponding equations are:

$$\begin{aligned} X_{i+1}(a) &= X_i(a) + X_i(b) \\ X_{i+1}(b) &= (X_i(a) - X_i(b))W_N^i \end{aligned} \quad (\text{Eq 41})$$

While it's a bit difficult to see at first, we can state the following: For every decimation-in-time algorithm there exists a decimation-in-frequency algorithm that is equivalent to swapping the input and the output, and reversing all the arrows in the flow diagram. This duality is useful as we consider computing the *inverse FFT* (FFT^{-1}).

Returning to the Time Domain

NR systems are typical in that after obtaining the FFT, we perform some modification of the frequency-domain data; we then transform the modified data back to the time domain. Since Eq 18 and 22 are so similar, the type of algorithms described above can be used to compute the FFT^{-1} . One way is to simply substitute $1/2 W_N^{-kn}$ for W_N^{kn} at each stage in Fig 12, and of course use $X(k)$ as the input to obtain $x(n)$ as the output. This results in the diagram of Fig 18.

Alternatively, we can compute the FFT^{-1} by using either form of FFT flow diagram, swapping the inputs and outputs and reversing direction of signal flow, as mentioned before. It's important to note that this is a consequence of the fact that we can rearrange the nodes of the flow diagrams however we want without altering the result. So, they work just as well in reverse as they do in the forward direction!

It's convenient to have the output order of the FFT the same as the input order of the FFT^{-1} , so we're wise to use decimation-in-time for one conversion, and decimation-in-frequency for the other.

General Computational Considerations

This business of bit-reversed indexing is what usually ties one's brain in knots during coding of these algorithms, and it's certainly one of the first things to be tackled—so let's have at it! Several approaches are feasible: a look-up table, the bit-polling method, reverse bit-shifting and the reverse-counter approach.

The look-up table is perhaps the most straightforward method. The table is calculated ahead of time, and the index used as the address to the table. See Fig 19. Most systems don't require extremely large values of N , so the space taken by the table isn't

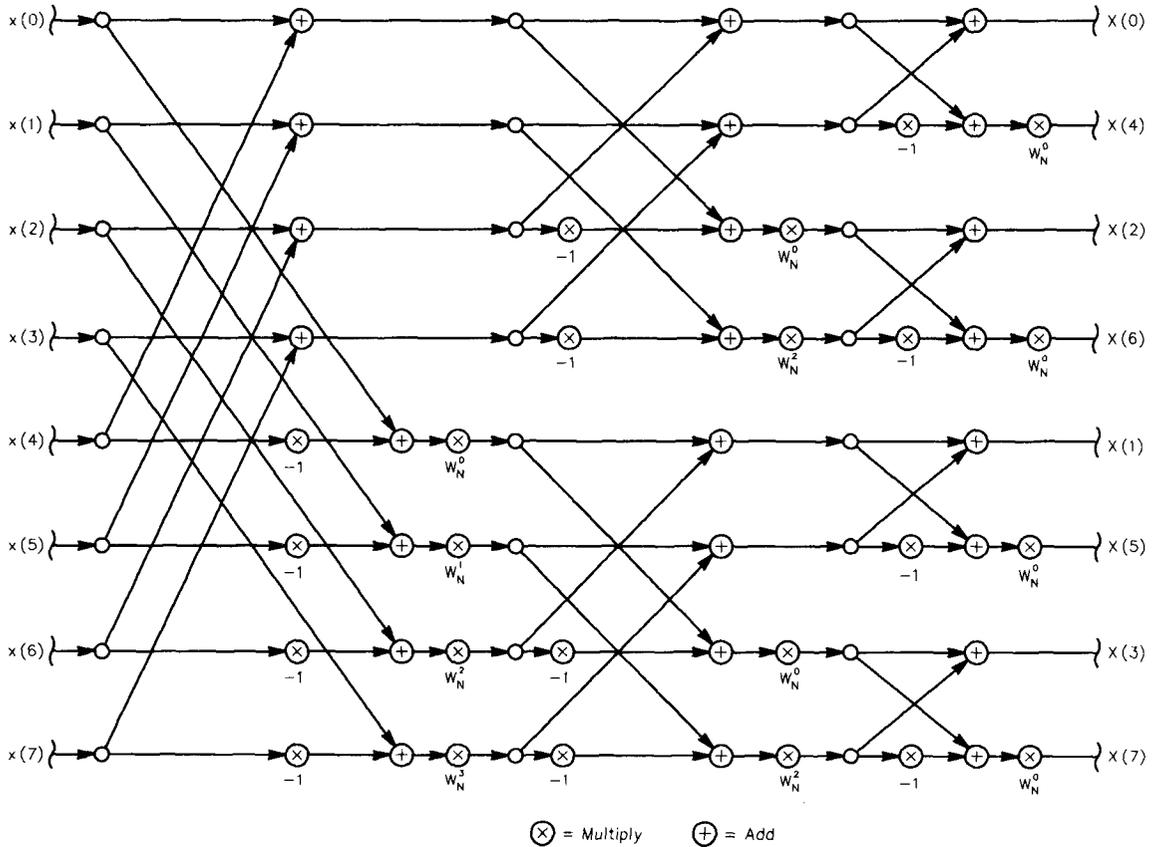


Fig 16—An eight-sample decimation-in-frequency FFT.

objectionable. For more space-sensitive applications, the bit-polling method may be attractive.

Since the bit-reversed indices were generated through successive divisions by two and determination of odd or even, a tree structure can be devised that leads us to the correct translation, based on bit-polling. See Fig 20. The algorithm examines the LSB, then branches either upward or downward in the tree based on the state of the bit. Then the second LSB is examined, a branch taken, and the procedure is repeated until all bits have been examined.

The bit-shifting method requires the same computation time. Two registers are used, one for the input index shifting right through the carry bit, the other shifting left through the carry. After all bits have been shifted, the left-shifting register contains the result. See Fig 21.

Finally, Gold and Rader¹¹ have described a flow diagram for a bit-reversal counter that can be "decremented"

each time the index is to change. If the data are actually to be moved during sorting, the exchange is made between data at input index n and bit-reversed index m , but only once! In other words, only $N / 2$ exchanges need to be performed.

During the actual calculations, the indexing of data and coefficients requires attention to many details. In particular, several symmetries about offsets of the index can be exploited. In

the case of a decimation-in-time FFT, at the first stage, all the multipliers are equal to $W_N^0 = 1$, so no actual multiplications need take place; all the butterfly inputs are adjacent elements of the input array $x(n)$. At the second stage, the multipliers are all either W_N^0 or integral powers of $W_N^{N/4}$, and the butterfly inputs are two samples apart. At the t th stage, the multipliers are all integral powers of $W_N^{N/2^t}$, and the butterfly inputs are separated by

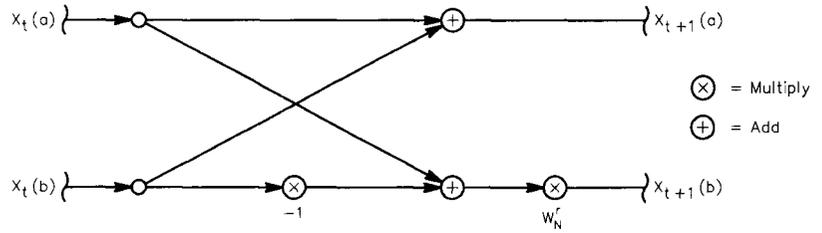


Fig 17—Butterfly calculation for decimation-in-frequency.

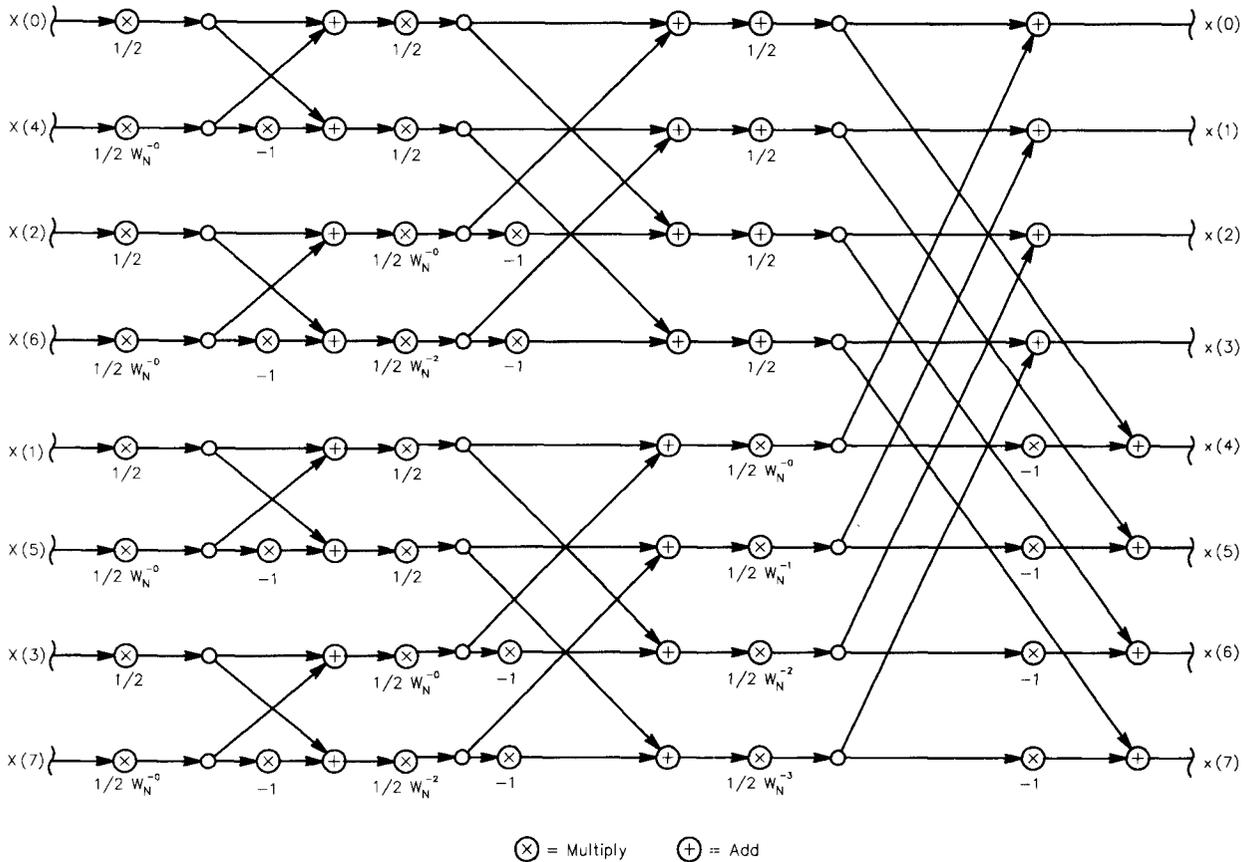


Fig 18—FFT⁻¹ implemented by interchange of input, output and coefficients.

2^{l-1} samples. Note that in most of the algorithms we've described, the coefficients are indexed in normal order.

Each FFT algorithm has its own unique indexing requirements. For our NR system, we don't care that the FFT result is stored in bit-reversed order, since we're not performing any frequency-specific operations. We'll use a decimation-in-time FFT to have the input sequence $x(n)$ in normal order, and a decimation-in-frequency FFT^{-1} to produce an output sequence also in normal order.

The coefficients of the complex exponential W_N can be obtained in various ways. The most common way is to generate them ahead of time and store them in a table. Another way is to use a recursion formula to generate them as needed. Since in general, the coefficients are all integral powers of W_N , we can use:

$$W_N^{kh} = W_N^k W_N^{k(h-1)} \quad (\text{Eq 42})$$

to get the h th coefficient from the $h-1$ th. However, errors will build up over time with this method because of the finite precision of our mathematics; each multiplication generates a rounding or truncation error that adds to the total. It's necessary to reset the value at periodic intervals to prevent divergence.

We'll see below that we always have to accept *some* error in our results because of accumulated rounding or truncation no matter how the DFT is calculated. We will analyze these *quantization* effects for a direct DFT calculation, and for a decimation-in-time FFT calculation.

Numerical-Accuracy Effects in DFT and FFT Calculations

In each complex multiplication, we must perform four real multiplications. Each of these contributes a round-off or truncation error to the output. We need to make some assumptions about the errors in order to do any analysis of them. Fixed-point, two's-complement arithmetic is assumed as well, as is typical in an embedded implementation.¹²

First, if b is the number of bits used to represent numbers, we'll assume the errors are uniformly distributed over the range:

$$-2^{-(b-1)} \leq \epsilon \leq 2^{-(b-1)} \quad (\text{Eq 43})$$

Each one of the errors therefore has variance:⁵

$$\sigma^2 = \frac{2^{(-2b)}}{12} \quad (\text{Eq 44})$$

Address	Data
000	000
001	100
010	010
011	110
100	001
101	101
110	011
111	111

Fig 19—Bit-reversal look-up table.

Also, we assume the errors are uncorrelated with each other, and also uncorrelated with the input and output.

Since noise powers add, the average value of the noise power is the *expected value* of four times the variance:

$$E[\epsilon^2] = 4 \left(\frac{2^{-2b}}{12} \right) = \frac{2^{-2b}}{3} \quad (\text{Eq 45})$$

At the output, the noise power is N times worse for the direct DFT calculation:

$$\sum_{n=0}^{N-1} E[\epsilon_n^2] = \frac{2^{-2b} N}{3} \quad (\text{Eq 46})$$

Just as in the case of FIR filter calculations,⁶ the noise at the output is directly proportional to N .

Also like the FIR analysis, the DFT calculations are subject to a dynamic-range limitation on the large-signal end of things. To prevent overflow, we require:

$$|X(k)| < 1 \quad (\text{Eq 47})$$

and this can be ensured if:

$$\sum_{n=0}^{N-1} |x(n)| < 1 \quad (\text{Eq 48})$$

We may need to scale the input by $1/N$ to prevent overflow. This scaling requirement has the effect of making the output noise worse, as will be discussed below.

For the decimation-in-time FFT calculation, the same assumptions about the nature of the noise are used. Referring to Fig 12, note that no more than one noise source is inserted at each node, because of the single complex multiplication there. The total noise at any node is the cumulative effect of all sources that propagate through to that node. Since we assume all the noise sources are uncorrelated, no more than $N-1$ noise sources propagate through to each output. So, the total output noise is, again, roughly proportional to N . When we take into account the requirement for data scaling, however, we'll see that the noise power must increase because the signal power at each node must decrease.

In fact, it can be shown⁵ that the SNR at the output—using optimum stage-by-stage scaling—cannot be better than:

$$SNR_{\text{OUTPUT}} = \frac{2^{(2b-2)}}{N} \quad (\text{Eq 49})$$

or a factor of 12 worse than the result given in Eq 46.

Until now, we've assumed absolute accuracy for the values of the coefficients W_N^{kn} . Whether these are held in a table or calculated "on the fly," they

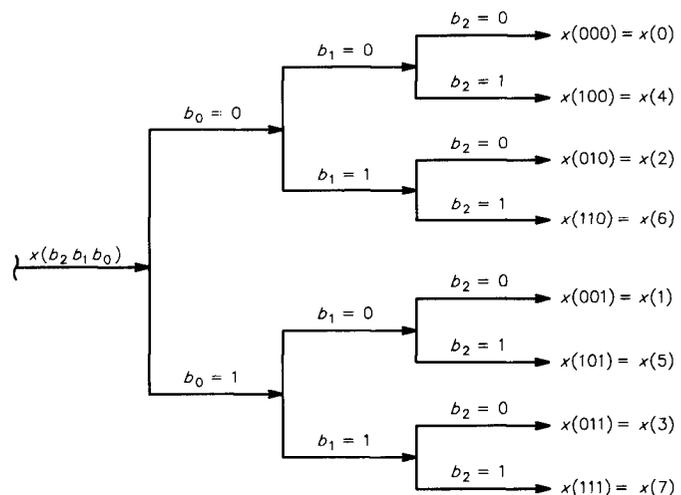


Fig 20—Bit-reversal tree.

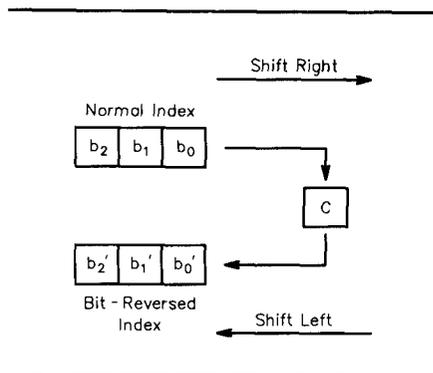


Fig 21—Bit-shifting register arrangement.

must be quantified to some number of bits, b ; this results in further cumulative error in the output. The analysis of the effects is much more difficult than that for data quantization above, since the nature of coefficient quantization is inherently nonstatistical.

Useful results have been obtained¹³ by introducing artificial noise or *jitter* into the coefficients, and analyzing the results for output error. The result obtained was that output SNR cannot exceed:

$$SNR_{OUTPUT} = \frac{3(2^{2b+1})}{p} \quad (\text{Eq 50})$$

where

$$p = \log_2 N \quad (\text{Eq 51})$$

The experimental results confirmed that noise from this effect increases in proportion to p , which means the increase with respect to N is slow. In other words, doubling N results in only a slight degradation in the SNR.

Engineering statistics is about the *grungiest* thing in the world, isn't it? Nevertheless, it sure is nice to know what to expect from these wonderful theories before committing to an implementation, because otherwise, *unexpected things can occur!*

There is a different calculation method for the DFT that ultimately dispenses with complex exponentials and improves speed and simplicity by a significant factor. Frerking touches on the idea, but provides no method for the control of its inherent divergence problem (see Note 12). I call it the "Damn-Fast Fourier Transform." We'll begin with that method in Part 4 of this series.

Doug Smith, KF6DX/7, is an electrical engineer with 18 years experience designing HF transceivers, control sys-

tems 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

- ¹Widrow, B. and Stearns, S. D., *Adaptive Signal Processing*, Prentice-Hall, Englewood Cliffs, New Jersey, 1985.
- ²Skolnik, M. I., *Introduction to Radar Systems*, McGraw-Hill, New York, NY, 1962.
- ³Widrow, B. and Hoff, M. E., Jr, "Adaptive Switching Circuits," *IRE WESCON Convention Records*, Pt 4, IRE, 1960.
- ⁴Hutchins, R. M., Ed-in-chief, *Great Books of the Western World*, 31st printing, Encyclopaedia Britannica Inc, Chicago, Illinois, 1989.
- ⁵Oppenheim, A. V. and Schaffer, R. W., *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, New Jersey, 1975.

- ⁶Smith, D., "Signals, Samples and Stuff: A DSP Tutorial" (Pt 1), *QEX*, Mar/Apr 1998, ARRL, Newington, Connecticut.
- ⁷Gardner, W. A., *Statistical Spectral Analysis: A Non-probabilistic Theory*, Prentice-Hall, Englewood Cliffs, New Jersey, 1987.
- ⁸Proakis, J. G., Rader, C. M., et al, *Advanced Digital Signal Processing*, Macmillan Publishing Company, New York, New York, 1992.
- ⁹Runge, C., *Z. Math. Physik*, Vol 48, 1903; also Vol. 53, 1905.
- ¹⁰Cooley, J. W. and Tukey, J. W., "An Algorithm for the Machine Calculation of Complex Fourier Series," *Math. Computation*, Vol 19, 1965.
- ¹¹Gold, B. and Rader, C. M., *Digital Processing of Signals*, McGraw-Hill, New York, New York, 1969.
- ¹²Frerking, M. E., *Digital Signal Processing in Communications Systems*, Van Nostrand-Reinhold, New York, New York, 1993.
- ¹³Weinstein, C. J., "Roundoff Noise in Floating Point Fast Fourier Transform Computation," *IEEE Transactions on Audio and Electroacoustics*, Vol AU-17, 1969. □

DSP

Without Tears ®

Multimedia Interactive CD-ROM Sale

Buy Before July, 4th 1998

Get all three CDs for \$999

Receive a free bonus Technical Book if order received before June 4th, 1998

This Could Be Your Last Chance!

Call Z Domain Technologies 1-800-967-5034 or 770-587-4812.

Hours: 9 - 5 EST. Or E-mail dsp@zdt.com

Ask for our CD sales price info, a free demo CD-ROM.

We also have live 3-day seminars: DSP Without Tears, Advanced DSP With a few Tears & Digital-Communications Without Tears.

<http://www.zdt.com/~dsp>

An Introduction to Microcontrollers: Ham Radio Style!

Curious about PICs? Learn the ropes with this simple keyer project built from a Basic Stamp.

By Al Williams, WD5GNR

For my 20th year as a ham, I talked my wife into letting me buy a new rig. After 20 years with a couple of Heathkits and a used Kenwood TS-120S that had seen better days, it was a big thrill to unpack a brand new Kenwood TS-570D. Not only did my shiny new toy have memories and DSP, but it also has a built-in memory keyer.

For 20 years, I've avoided using a keyer. Even when I handled a lot of traffic, I always did so with a straight key. I have a decent fist, and I've just never felt like I needed an electronic keyer. However, my new rig's memo-

ries require a paddle to program them. I didn't want to spend a lot of money on a paddle I didn't plan to use often, so my son Patrick and I built a paddle from his old Legos and two calculator keys (honest). After programming the keyer's memories, I discovered that I got a kick out of using my homemade paddle and started using it regularly.

The only problem is that I often work CW on homebrew QRP transmitters. These simple radios don't have built-in keyers (among other things), so I decided to build a keyer. At first, several designs came to mind, but I discovered that there is more to building a robust keyer than meets the eye. Of course, you can get keyers on a single chip, but they don't let you customize your keyer. You might as well buy one that's already built. Finally, I

concluded that I really want a microprocessor-based keyer.

In this article, I'll show you how to build a keyer using an inexpensive microprocessor called a Basic Stamp (from Parallax).¹ The keyer has many modern features, but the best feature is that you can program firmware for it using a language much like *BASIC* and download updated firmware to the keyer using an ordinary PC. You need not erase EPROMs or even remove the processor from the keyer. There's no easier way to get started programming microcontrollers.

About the Keyer

My prototype of the keyer has three front-panel switches and a potentiom-

¹Notes appear on page 31.

eter (see Figure 1). It has five jacks on the rear apron. One switch controls the power, and another enables or disables the internal sidetone speaker. The third button's function depends on the firmware. The basic firmware uses the button as a tune switch. The advanced firmware uses it to begin transmission of a stored message. The firmware you write can use the switch to do anything you like.

The same is true of the potentiometer. Most builders would use it to set the keying speed, but you could use it for anything. For example, you might use it to set the speed when the programmable switch is off, but set the dot-dash ratio (weight) when the switch is on.

Three of the rear-panel jacks connect to dc power (from a wall transformer), the radio and the paddle. A fourth connects to a PC and allows you to download firmware. The fifth jack provides a connection for a straight

key. The microprocessor can read the straight key, but the key directly switches the transmitter. This ensures that the straight key will operate even if a firmware error renders the keyer inoperative.

My final version of the firmware has several interesting features:

- Full iambic operation
- Dot and dash memories
- Sidetone
- Separate input for a straight key that overrides the keyer's output
- Analog speed control
- Prerecorded message playback

Inside the Basic Stamp

Parallax's Basic Stamp is a small PC board that looks like a 14 pin SIP (see Figure 2). It contains a voltage regulator, a PIC microprocessor and a 256 byte EEPROM. The PIC's (peripheral-interface controller) onboard ROM contains a simple run-time system that allows you to store and execute

programs written in a simple BASIC-like language (see Table 1).

You can supply the stamp with a regulated +5 V dc, or you can connect anything from 6 to 30 V and let the onboard regulator supply 5 V for your entire circuit. The chip has eight I/O lines that you can use to communicate with the outside world. It also has two pins that connect to a PC (via its parallel port) for programming.

You program the stamp with special software supplied by Parallax. You can download the software free of charge from the Internet (<ftp://ftp.parallaxinc.com/pub/parallax/stamp.zip>). You can also get an electronic version of the Basic Stamp manual at the same location (ftp://ftp.parallaxinc.com/pub/acrobat/stamp1_manual_v1.8.pdf).

How it Works

The keyer circuit is simple because the Basic Stamp microprocessor is

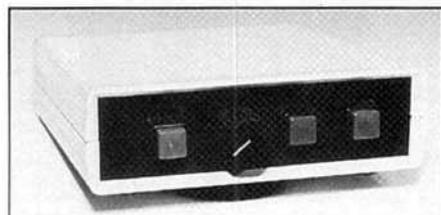


Figure 1—The keyer.

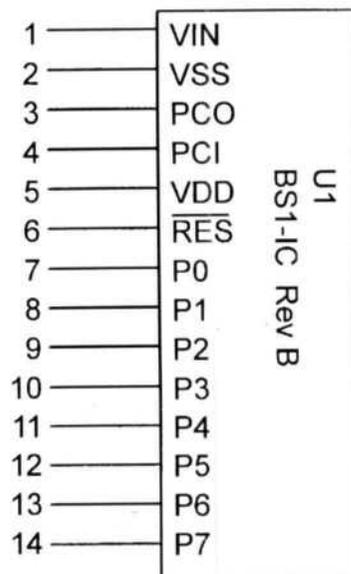


Figure 2—Pin diagram of U1.

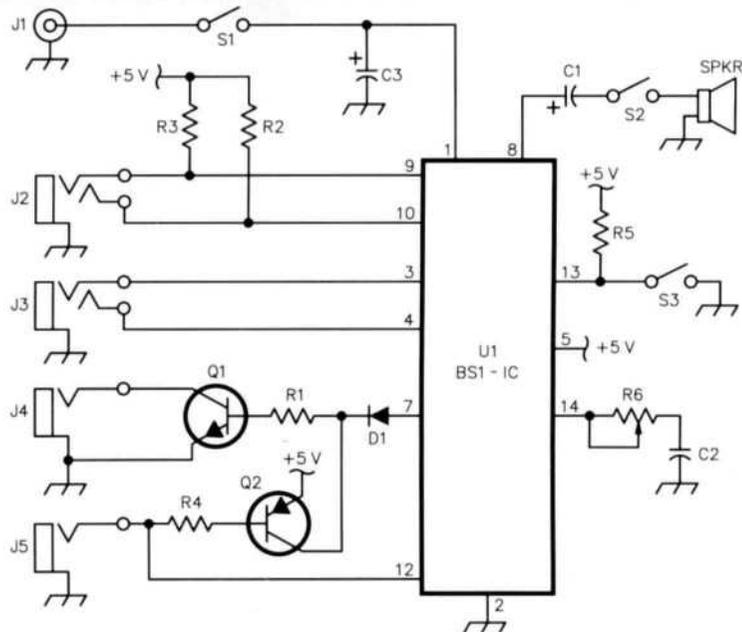


Figure 3—Schematic for the keyer. Unless otherwise specified, use $\frac{1}{4}$ W, 5%-tolerance carbon composition or film resistors. Pins not shown are unconnected.

C1—10 μ F, 16 V electrolytic capacitor
 C2—0.1 μ F capacitor
 C3—100 μ F, 16 V electrolytic capacitor
 D1—1N4148 diode
 U1—Basic Stamp BS1-IC microprocessor (see Note 1)
 J1—Power connector for 6 to 15 V dc
 J2—Paddle connector
 J3—Programming connector
 J4—Jack for transmitter keying line
 J5—Straight-key connector
 Q1—2N2222 NPN transistor
 Q2—2N3904 PNP transistor

R1—680 Ω , $\frac{1}{2}$ W resistor
 R2, R3, R5—10 k Ω , $\frac{1}{4}$ W resistor
 R4—1 k Ω , $\frac{1}{4}$ W resistor
 R6—10 k Ω potentiometer
 S1—Power switch, SPST
 S2—Sidetone switch, SPST
 S3—Function key, SPST momentary
 SPKR—32 Ω loudspeaker
 Enclosure
 Power supply or battery
 Printed Circuit Board
 14 Pin SIP Socket for U1

highly integrated. (See Figure 3.) The circuitry mainly consists of pull-up resistors for the inputs. The most complex part of the circuit is the keying control (Q1 and Q2). When the processor wants to key the transmitter, it brings pin 7 high, which saturates Q1 (through D1). You can also bring Q1 to saturation when Q2 turns on, because of a ground circuit closed by the straight key. The microprocessor can read the straight key, via pin 12, so that the keyer can provide a sidetone for the key, if desired.

The keyer creates a sidetone by generating a square wave on pin 8. The square wave drives a 32 Ω speaker via C1. On the prototype, a switch (S2) in series with the speaker allows you to mute the audio if you like.

The speed control is an ordinary potentiometer. How do you read a potentiometer with a digital input? The trick is to connect one side of the pot to a pin of the microprocessor and the other to a grounded capacitor. A special command in the firmware measures the time it takes to charge the capacitor through the pot. This time is proportional to the setting of the pot.

Building the Keyer

Constructing the keyer is straightforward with the supplied PC board pattern (see Figure 4). It's a good idea to use a socket for U1, but 14 pin SIP sockets are difficult to find. An easier (and often less expensive) answer is to use half of an open-frame 28 pin DIP socket. (Cut the socket in half, along its length.) This results in a SIP socket with excess plastic on one side. You can remove the excess plastic if you like, or simply be careful to install the socket so that any excess plastic doesn't obstruct mounting holes for other parts in the PC board.

There really aren't any other tricks to building the circuit. Pads for wires are larger than component pads. In addition, board connections for signal wires have designators (see Table 2). Unmarked wire pads connect to the ground plane, and you may use them as grounds for external components as needed.

S3 should be a momentary-contact switch. I used Radio Shack's RS275-1566 switch and the matching RS275-1565 for the other two switches. Speaking of Radio Shack, the board fits nicely inside an RS270-214 enclosure. R6 poses a problem if you are trying to buy parts from Radio Shack. They sell a 10 kΩ pot that has a very long shaft (RS271-1715), and another

one with a normally sized shaft (RS271-215). The short-shaft pot, however, has an integral SPST switch behind it. I used the shorter pot and carefully removed the tabs on the back that hold the switch in place. After removing the switch, the pot functions as you'd expect, although the housing remains a bit larger than necessary.

You can use practically any power supply that can produce between 6 and 30 V dc. A 9 V battery is perfect. If you decide to use one, use a 9 V clip for J1. I used a 9 V dc wall transformer and a coaxial connector for J1. It's a good idea to measure the voltage at pin 1 of U1's socket before installing the IC, to make sure you connected the power supply correctly.

Programming

Programming the keyer is simple,

but you need a special printer cable. The PC end of this cable has a male DB25 connector. The other cable end connects to J3. I used a ¼ inch phone plug for J3 because it's easy to mount and compact. A DB25 or DB9 connector would take up more room and require holes that are difficult to cut in the cabinet.

There is one confusing issue concerning the software. When programming in *PBASIC*, you use "pin" numbers from 0 to 7 to specify I/O ports. These "pin" numbers do not correspond to the hardware pin numbers, however. For example, Pin0 in software is actually pin 7 on the stamp itself. (See Figure 2.)

With that in mind, it is easy to create various firmware configurations to meet your needs. You simply bring Pin0 high to key the transmitter. Use

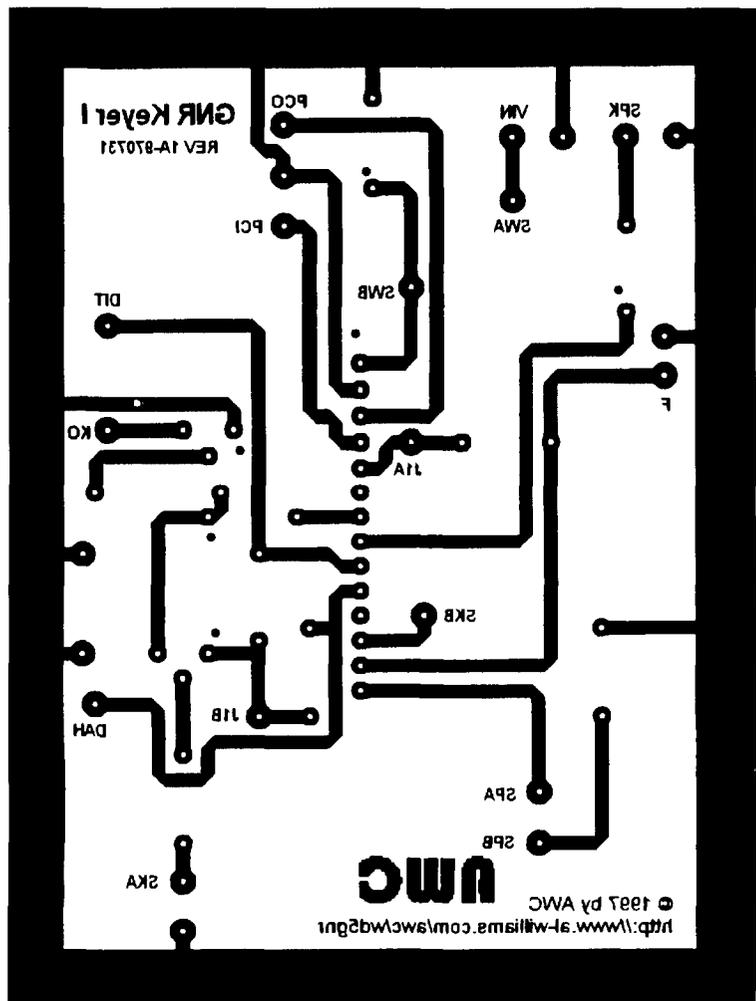


Figure 4—The PC board etching pattern (finished size 3×4 inches).

Table 1—Basic Stamp Commands

Command	Description
If... Then	Evaluate condition and jump if true
Branch	Jump to location based on offset
Goto	Jump to a new line of code
Gosub	Jump to a subroutine
Return	Return from a subroutine
For...Next	Execute statements in a loop
Let	Assign a value to a variable (optional)
Lookup	Look up value in a table
Lookdown	Convert value to an index in a table
Random	Generate a random number
Output	Make I/O pin an output
Low	Set output low
High	Set output high
Toggle	Reverse output pin's state
Pulsout	Output a pulse
Input	Make I/O pin an input
Pulsin	Measure a pulse input
Reverse	Make an output pin an input and vice versa
Button	Read button with debouncing and repeat functions
Serin	Read serial data
Serout	Write serial data
Pwm	Output a train of pulses; useful for generating an analog output with an external RC circuit
Pot	Read a potentiometer
Sound	Play a tone
Eeprom	Store data in the stamp's EEPROM at compile time
Read	Read EEPROM data
Write	Write EEPROM data at run time
Pause	Pause for a specified number of milliseconds
Nap	Briefly pause in low power mode
Sleep	Enter low power mode for a specified number of seconds
End	Enter low power mode until reset
Debug	Send data to host PC

the Pot command to read R6. You can also read the state of the switches by reading their corresponding input values (they are all active low; that is, they read logical 0 when you press them).

If you read the stamp manual, you'll notice that it has a Sound command specifically for generating tones. You probably can't use it for creating the keyer's sidetone however. Why? Because the stamp has no interrupts, it must poll the dot and dash paddles constantly to simulate dot and dash memories. The Sound command ties up the processor until it completes. You could use Sound to make other beeps (for example, to signify a mode change), but you won't want to use it for sidetone.

My firmware generates the square-wave sidetone from inside a loop. That means the sidetone frequency is fixed by the stamp's operating speed. You can, however, adjust the tone somewhat by altering the number of times you toggle the square wave during each iteration of the loop.

You can encode a short message to play back using the advanced firm-

ware. Encoding the letters is a bit tricky, however. The firmware uses a peculiar scheme to compactly represent the code stream. The program assumes that a zero is a dit, a one followed by a zero is a dah and a one followed by a one is a dit-sized pause. This allows for a crude form of Hoffman compression, which is important because the stamp doesn't have much memory. I've supplied a Windows program that will help encode any strings you might want to program.²

Conclusion

Although you can build a keyer for less, you can't get one that is more flexible. By using the Basic Stamp, you can make the keyer operate in any manner you wish. There is a spare I/O pin, so you could use that to control a T/R switch, for example. In fact, it is entirely possible to use the keyer circuit for completely different purposes. If you need a microprocessor with a few switch contact inputs, a speaker, a pot, and a keyed output, you could use this circuit just by reprogramming it. For example, you might wire up limit switches to J2

Table 2—Signals on the PC Board

Designator	Connection
VIN	Voltage In
SWA	One side of power switch
SWB	Other side of power switch
SPK	Speaker
PCO	Programming connection (pin 3)
PCI	Programming connection (pin 4)
F	Function key
J1A	Jumper (side A)
J1B	Jumper (side B)
DIT	Paddle dot switch
DAH	Paddle dash switch
KO	Keyer output
SKA	Straight key
SKB	Straight key (wired in parallel with SKA)
SPA	Speed control wiper
SPB	Speed control (C2-side)

and control an antenna rotator with J4.

Of course, the unit also makes a dandy keyer. With the Stamp's low power requirements, low parts count and small outline, you could easily build the keyer inside a rig, or in a very small case. The stamp can use practically any power supply. When you are done, you can program the keyer to suit your personal tastes. All of that while learning about embedded microcontrollers, too.

Al Williams, WD5GNR, has been an avid ham for over 20 years. Although Al used to design hardware, he now spends his time writing books and magazine columns about Windows programming. Al still keeps up with the hardware side of things on his own time, however. You can find out more at <http://www.al-williams.com/wd5gnr>.

Notes

¹The Basic Stamp described in this article is now called the Basic Stamp 1. Parallax also makes a Basic Stamp 2 that has 2048 bytes of EEPROM and a more powerful instruction set than the BS1. Parallax Inc, 3805 Atherton Rd, Ste 102, Rocklin, CA 95765; sales (888) 512-1024, office/technical support (916) 624-8333, fax (916) 624-8003, faxback (916) 624-1869, BBS (916) 624-7101 (300 - 14,400 bps, 8 data bits, 1 stop bit, no parity); general e-mail info@parallaxinc.com, technical e-mail stamptech@parallaxinc.com, pictech@parallaxinc.com and sxtech@parallaxinc.com; anonymous FTP <ftp://parallaxinc.com>; WWW <http://www.parallaxinc.com>.

²You can download the software and firmware at http://ourworld.compuserve.com/homepages/Al_Williams/keyer/keyer.htm. □□

A Continuous Coverage HF VFO

This circuit can help your next homebrew receiver roam the entire HF spectrum.

By Francesco Morgantini, IK3OIL

This project arises from the need to homebrew a variable local oscillator for a multiband transceiver. It is a partial-synthesis VFO intended for use with single conversion equipment using an IF near 9 MHz. Using such an IF, the circuit can cover the entire HF spectrum from 3.5 to 30 MHz (ie, 12.5 to 39 MHz LO output).

The VFO has been developed through several experiments based on PLL and crystal conversion circuits. I think it represents an acceptable compromise between simplicity (but not simple enough to be an elementary project)

and performance.

Some equipment is necessary for the alignment: An RF signal generator and a frequency meter are needed, but the use of an oscilloscope makes the job easier (especially in case of trouble).

How It Is Made

The circuit consists of two physically separated sections:

1. The VFO module is composed of three small PC boards and the variable capacitor. The three boards are stacked, and the overall cabinet dimensions depend essentially on the capacitor size. (See Fig 1.) This unit can be located behind the front panel of the rig.

2. The PLL Module contains the VCO PC board and the PLL PC board. This unit may be located anywhere in

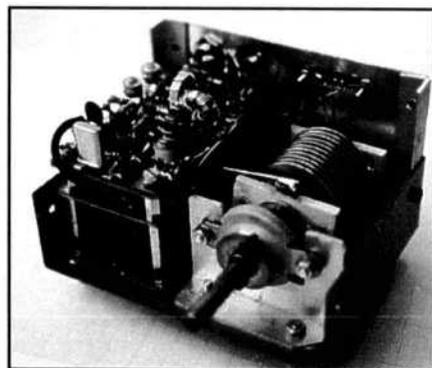


Fig 1—A photo of the VFO module with the cover removed. PCB1 through PCB3 are stacked at left; behind the capacitor.

the rig. The printed boards are double-sided fiberglass and components are soldered directly on the copper pattern without drilling. The lower side is

Via Leoncavallo, 21
35010 - VIGONZA (PD)
ITALIA
e-mail ik3oil@iol.it

an unetched ground plane. Grounded traces are connected to the ground plane by drilling through the board, installing a through wire and soldering on both sides. The two PC boards are contained in an aluminium cabinet 70x100x40 mm (approximately $2\frac{3}{4} \times 4 \times 1\frac{1}{2}$ inches).

The output phono jacks and a connector for the dc supply and band switching are located on a side of the cabinet. Fig 2 shows the overall arrangement.

How It Works

Fig 3 is a block diagram of the circuit. For this example we will assume you wish a VFO in the 5 to 5.5 MHz range with a 9 MHz IF in the rig, other frequency values near those indicated are also possible.

The unit described as the *VFO module* is a conversion VFO in the 41 MHz range. (See Fig 4.) It contains:

1. A Colpitts VFO ranging from 5 to 5.5 MHz (PCB1, Fig 5). The mechanical assembly and the component choice must be very accurate. Use a high-quality variable capacitor (ball bearing supported) and NP0 ceramic capacitors. An N150 capacitor can be used as a compensation element to reduce thermal drift. L1 is 24 turns of 0.8 mm (#20 AWG) wire on 13 mm ($\frac{1}{2}$ inch) Plexiglas core ($3.5\mu\text{H}$). The varactor allows a 20 kHz shift for fine tuning or **SPLIT** function. The VFO circuit could be replaced by a more sophisticated DDS unit.

2. Two buffers drive an external

frequency counter and the first mixer stage (PCB2, Fig 6). The output level to the mixer should be about $3 V_{p-p}$. PCB2 also contains the '7810 voltage regulator.

3. The first mixer and a related 41 MHz filter (PCB3, Fig 7). This board uses a BF960 MOSFET as a mixer and a 2N2222 as a crystal controlled oscillator to obtain the 36 MHz output from an 18 MHz crystal. L2 is made from 9 turns of 0.5 mm (#24 AWG) wire on a T-44-2 toroidal core ($0.42\mu\text{H}$). A 41 MHz output filter is formed by L4 and L5, 9 turns, 0.5 mm (#24 AWG) wire on a T-44-6 core ($0.34\mu\text{H}$ each). A 2N2222 buffer stage follows the filter.

L3 is 2 turns wound on L4.

Here's how to align the mixer:

- Remove the crystal to inhibit the oscillator.
- Inject a 41.2 MHz signal to the MOSFET's gate 1 and adjust the capacitors for maximum output.
- Insert the crystal and drive a 5 MHz signal into gate 1. Adjust the 60 pF capacitor for maximum output, which should be 0.7 to $1 V_{p-p}$.

The *VCO unit* (PCB4, Figs 8 and 9) contains a VCO that covers the range from 12.5 to 39 MHz using two distinct oscillators, which are switched by a relay that is driven from an

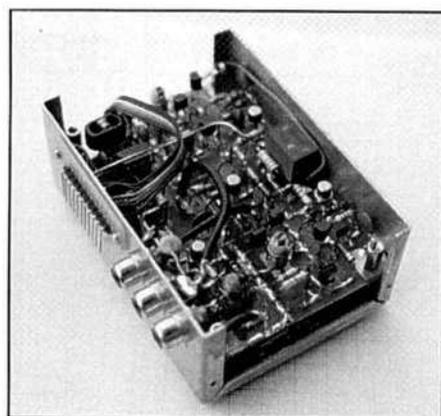
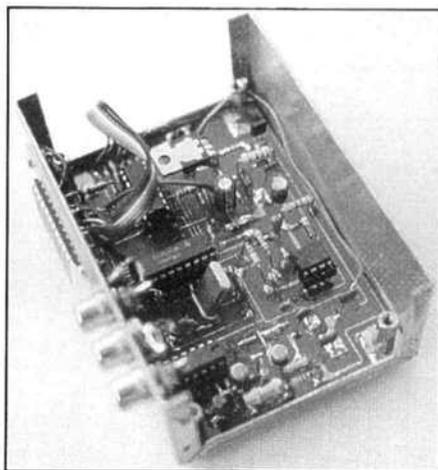


Fig 2—(A) photo of second cabinet with the PLL board in place. (B) shows the VCO board mounted above the PLL board.

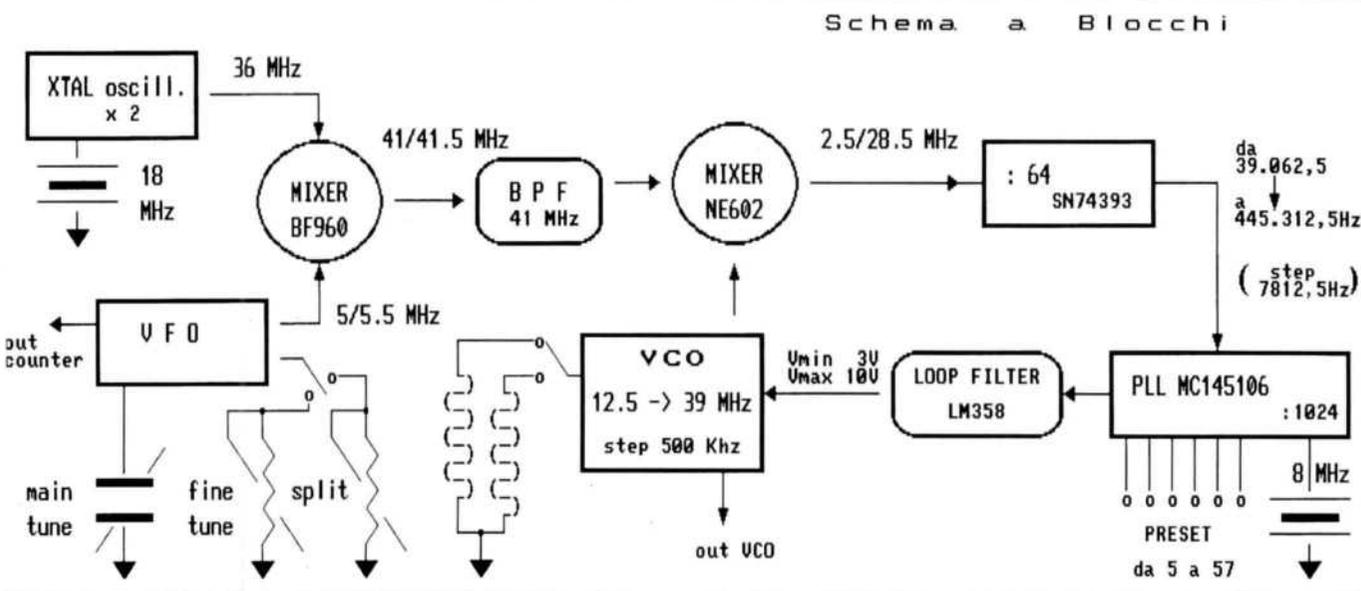


Fig 3—Circuit block diagram.

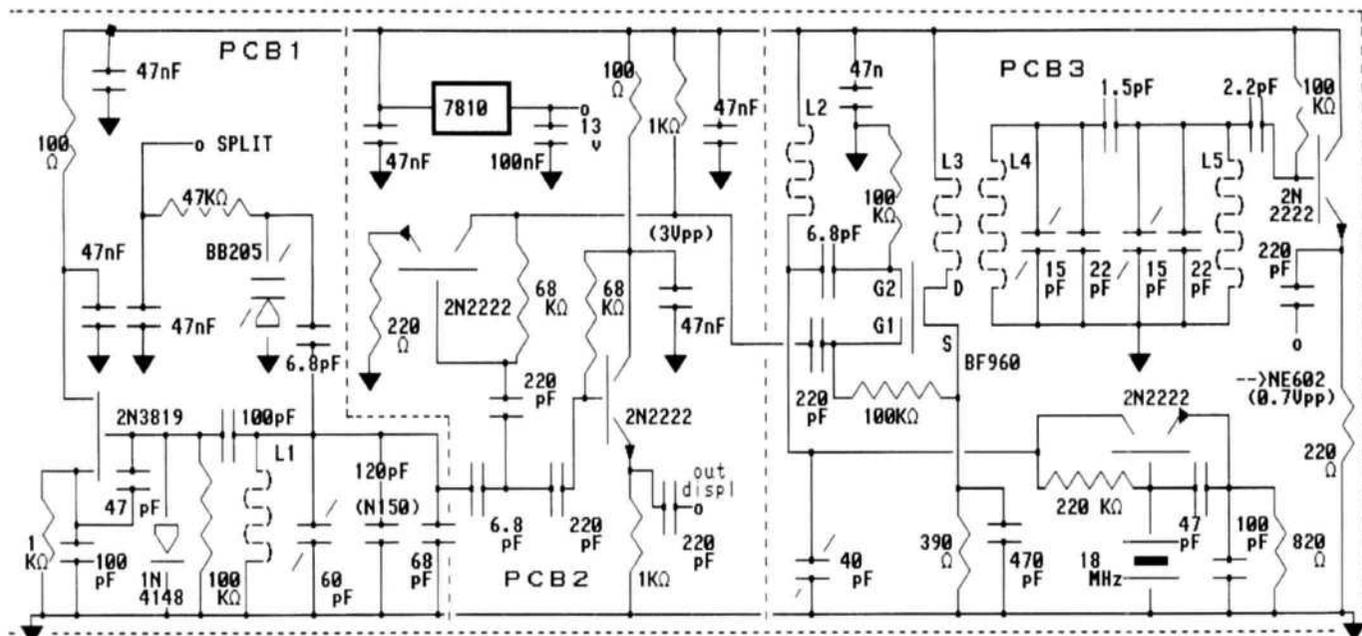


Fig 4—The VFO module schematic. All part values are labeled here and/or described in the text.

appropriate bandswitch section. The varactors are high-capacity devices for AM broadcast use (MVAM115, BB112 or equivalent). The circuit configuration of the oscillators and the use of compensation networks made it possible to obtain good quality, constant-level output over the entire frequency range. L6 is 7 turns of 0.8 mm (#20 AWG) wire on a 5 mm (³/₁₆ inch) plastic form with a ferrite variable core (0.18 to 0.48 μ H). L7 is 12 turns of 0.5 mm (#24 AWG) wire on a similar core (0.52 to 1.3 μ H).

The alignment can be performed in the following manner:

- Supply a 3.5 to 9.5 variable voltage to the varactors (do not exceed these limits).
- Tune the cores of L6 and L7 so that the varactor voltage variation produces frequency ranges from 22 to 39 MHz with L6 and from 12.5 to 22 MHz with L7. (Select each oscillator by keying or releasing the relay that controls the +10 V line at the top, left of center in Fig 8.—Ed.)

The output level across a 200 Ω load should be about 3 V_{P-P}.

The second mixer uses an NE602 IC. This device allows me to obtain the best linearity and balance over the entire frequency range. The input VCO signal is reduced with a capacitive divider and the two balanced inputs (pins 1 and 2) are driven with opposite phases using a

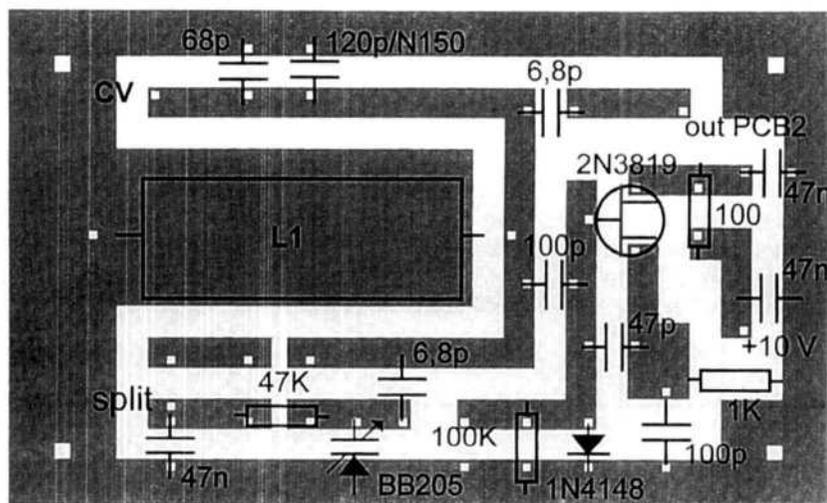


Fig 5—Parts placement details for the VFO module, PCB1 (actual dimensions 42x68 mm).

broadband transformer to limit spurious outputs.

T1 is wound with 5 bifilar 0.5 mm (#24 AWG) wires onto a TV balun ferrite core (14x8x8mm, type 43 material). Some adjustment of the value of the 2.2 pF capacitor may be required to obtain a level of 100-200 mV_{P-P} into pins 1 and 2 of the IC. A buffer stage equipped with two 2N2222 transistors and a compensation network on the second-stage emitter allowed me to obtain a

substantially constant output level over the entire frequency range covered by the mixer. This is very important to ensure reliable operation of the TTL 74LS393 divider (on PCB5—Ed.).

The PLL unit (PCB5, Figs 10 and 11) contains:

1. The frequency divider uses a TTL 74LS393, which divides the frequency coming from the mixer by 64. The output frequency is comparable to the internal reference of the PLL, and we can

An Updated Electronic Eyeball

This panadaptor features a calibrated logarithmic detector and linear, calibrated sweep. The panadaptor is powered from a single +12 volt power supply making it ideal for battery powered applications.

By Bob Dildine, 7J1AFR, W6SFH

The cover story¹ of my first issue of *QST* (that I received when I was in high school) described construction of a panoramic adapter (panadaptor) for use with 455-kHz-IF receivers of that time. The concept of a panadaptor has intrigued me ever since and that article was the inspiration for the panadaptor described here.

Following are the design goals for the panadaptor:

- Indicated signal level calibrated in decibels
- Linear, calibrated horizontal sweep
- Easily adaptable for different input frequencies
- Powered from +12 V
- Simple, straightforward, reliable design
- No exotic parts

For simplicity, I decided to use an oscilloscope for the

display and concentrate design effort on the RF portions of the panadaptor.

The first thought was to make use of one of the chip sets used for low-power receiver IF circuits. Several designs for low cost spectrum analyzers have been published using these parts,^{2, 3} and they provide mixer and oscillator functions along with potentially high gain and low parts count. The RSSI outputs from detectors such as the NE604 or MC3356 are proportional to the logarithm of input signal level and would meet the requirement for signal-level indication calibrated in decibels. The initial design for the panadaptor was based on the NE602/NE604 chip set, but the high input and output impedances made it difficult to test and troubleshoot filters and the finished circuits in a 50 Ω environment. For this reason and because low parts count was of secondary consideration, I decided to design the circuits of the panadaptor from discrete parts and for 50 Ω interfaces.

Block Diagram

The block diagram of the panadaptor is shown in Figure 1. I chose 455 kHz for the final IF because it is easy to

¹Notes appear on page 44.

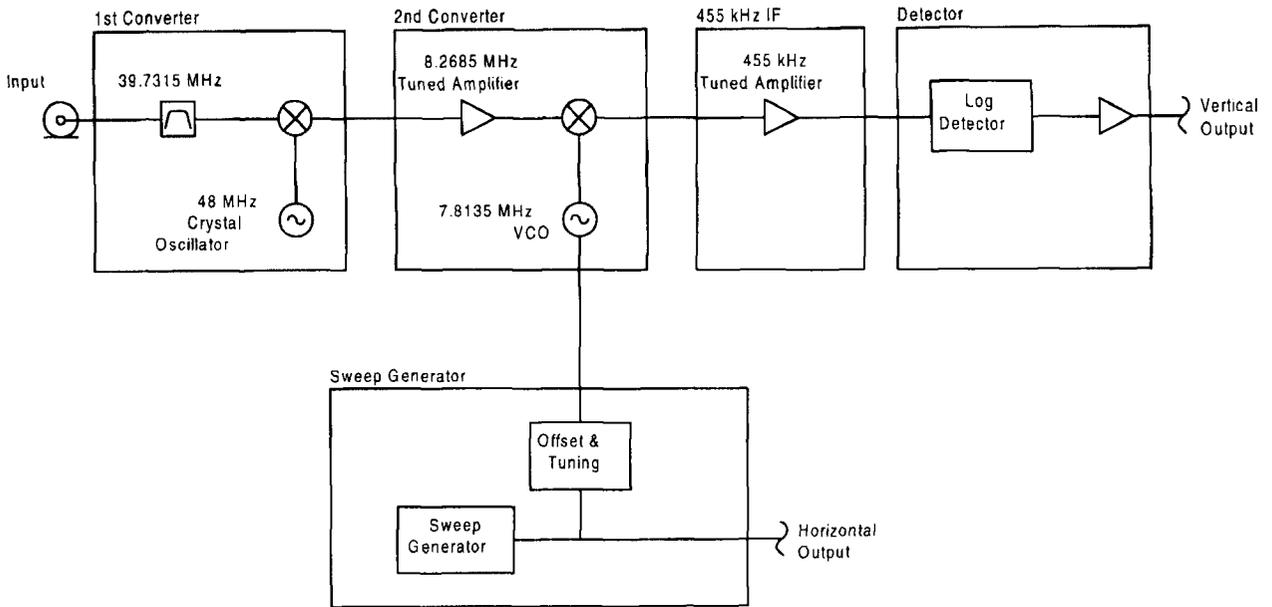


Fig 1—Block diagram.

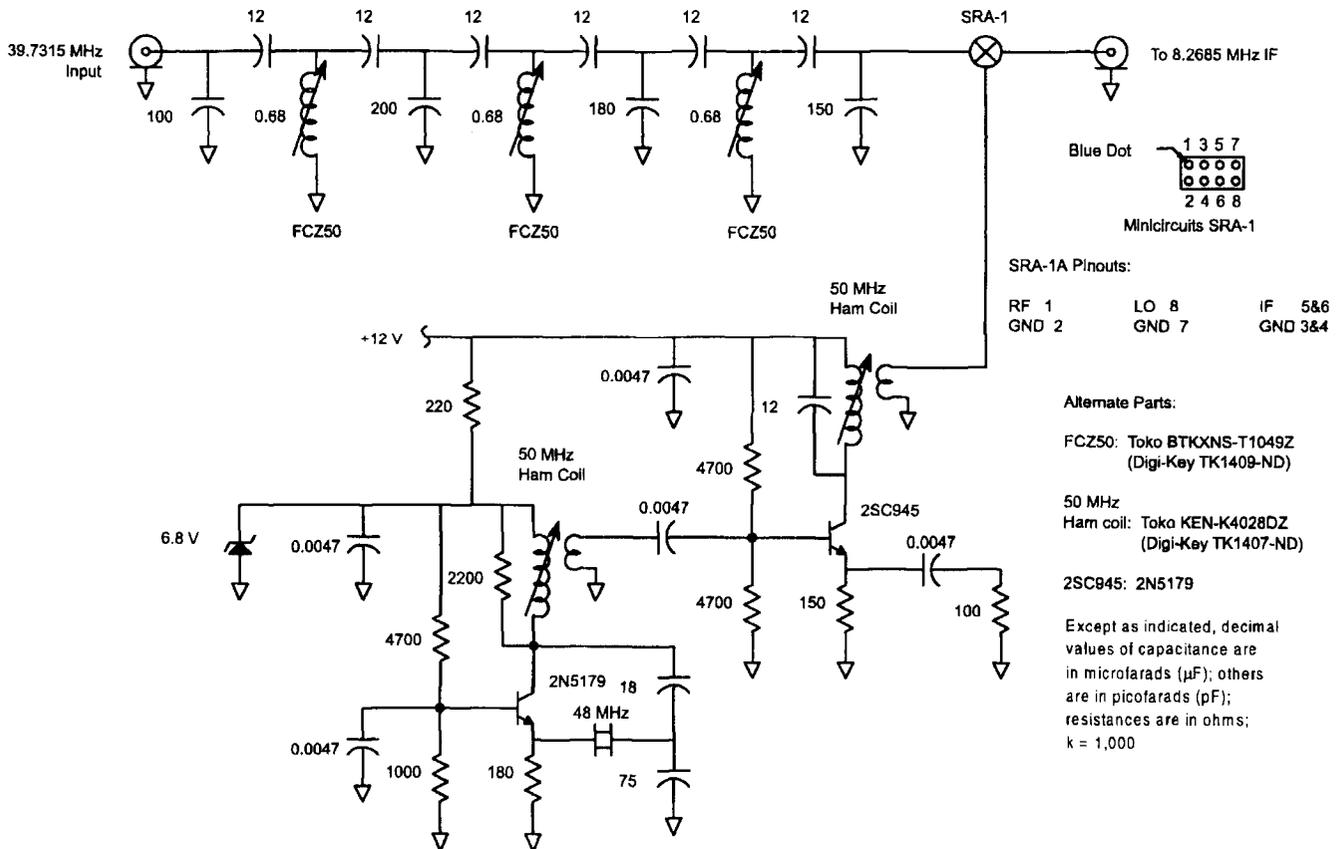


Fig 2—Schematic of first converter.

obtain good selectivity with tuned circuits and ceramic filters that are readily available at that frequency. A dual-conversion approach allows the 455 kHz IF to be used with much higher input frequencies and allows flexibility in modifying the panadaptor for use with a wide range of receiver IFs.

The log detector is based on an NE604 IF/Detector IC. The RSSI output of this device is proportional to the log of the input signal level. A 455 kHz tuned IF amplifier provides gain ahead of the detector as well as additional selectivity. The gain of this amplifier is controlled in steps of about 10 dB.

The second converter consists of an 8.2685 MHz tuned IF amplifier and a mixer. The LO signal is provided by a VCO that is tuned by the sweep voltage from the scan generator.

The first converter is crystal controlled and converts the 39.7315 MHz input frequency (the first IF of my ICOM 720A) to 8.2685 MHz. A band-pass filter before the mixer limits the input bandwidth of the panadaptor to

reduce spurious signals. Other input frequencies can be accommodated by changing the frequency of the band-pass filter and the crystal-controlled local oscillator.

The scan generator consists of a sawtooth oscillator and several op amp circuits. The op amps scale and offset the sawtooth signal to provide a tuning voltage for the VCO and a sweep voltage for the display.

First Converter

The first converter is shown in Figure 2. A six-pole band-pass filter with a bandwidth of about 2 MHz centered at 39.7 MHz helps reduce spurious responses. The filter passband can be adjusted with the aid of a spectrum analyzer and tracking generator or by tuning a signal generator across the input.

The mixer is a Mini-Circuits SRA-1A, which was used for simplicity.⁴ A similar doubly balanced mixer could be made from discrete diodes and transformers.

The crystal-controlled LO is based

on one shown in the 1989 *ARRL Handbook*.⁵ The 48 MHz LO frequency was chosen mainly because a 48 MHz crystal was available. For other input frequencies, use an LO frequency that gives a first IF somewhere in the 5 to 10 MHz range. If you change the first IF from what is shown here, change the tuned circuits in the second converter accordingly.

The first converter has an overall loss of about 11 dB, but this doesn't seem to affect performance on the HF bands, where sensitivity is limited by atmospheric noise.

Second Converter

The second converter consists of a two-stage tuned 8.2685 MHz amplifier as shown in Figure 3. The stages are stagger tuned for the flattest response over about a 200 kHz bandwidth. I adjusted them with a spectrum analyzer and tracking generator, but a signal generator can also be used by sweeping it across the passband. With the emitters of both stages bypassed, I noted some instability. It was probably due to

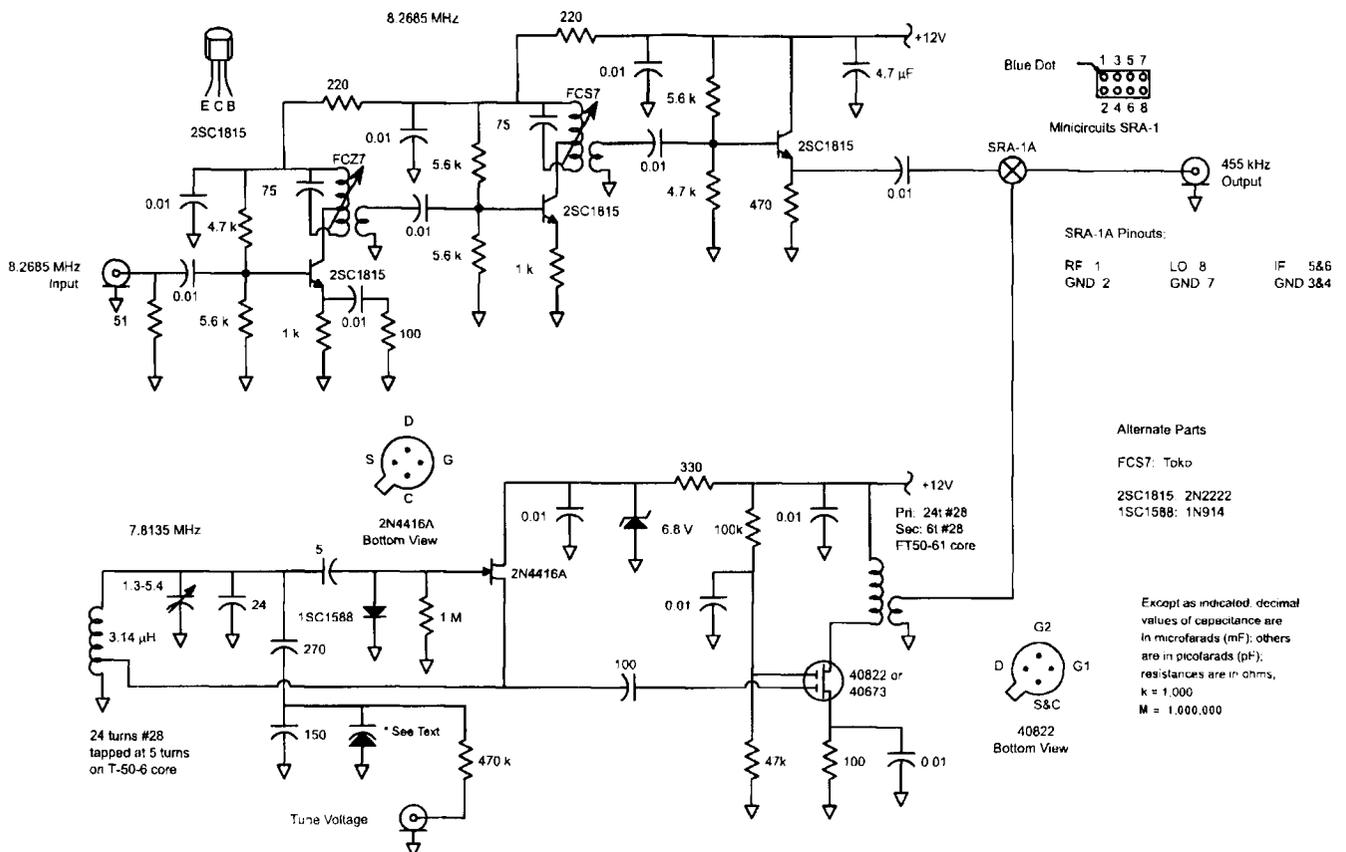


Fig 3—Schematic of second converter.

the compact layout. With only the emitter of the first stage bypassed, the amplifier is stable with about 23 dB gain.

The second mixer is Mini-Circuits SRA-1A and again, a similar doubly balanced mixer could be made from discrete diodes and transformers.

The VCO is copied from a VFO in the 1989 *ARRL Handbook*⁶ and modified for varactor tuning. The tuned circuits were scaled from their original values for 5 MHz to tune to 7.8135 MHz. The varactor is a small one of unknown type that was on hand. Coarsely adjust frequency by adding or removing turns from the inductor. Make fine frequency adjustments by changing the value of the fixed 24 pF capacitor across the inductor. A small 5 pF trimmer sets the final frequency after the VCO shield is

soldered in place. The trimmer has about a 150 kHz tuning range. After the coarse frequency was set, the VCO output frequency was measured as a function of tuning voltage to find the most-linear part of the tuning curve. In my prototype, linearity is best for tuning voltages between 3 and 13 V. Final frequency adjustment was done with the tune voltage set at 8 V. I then set values in the scan generator to provide the optimum tune voltage.

455 kHz IF Amplifier

The 455 kHz IF amplifier is shown in Figure 4. It provides enough gain to bring the output of the second converter up to a suitable level for the detector. The three stages are syn-

chronously tuned to give good selectivity. The IF transformers are from a line of adjustable coils available at many parts houses in Japan, but other 455 kHz IF transformers should work satisfactorily. (It might also be possible to substitute a salvaged 455 kHz IF strip from an old transistor radio.)

Gain control is accomplished by diode switching additional resistors in parallel with the emitter resistors of each stage. A single-pole, multiposition rotary switch and a simple diode matrix increases the gain of each stage sequentially. I found that any noise introduced into the emitters from the switching voltage is amplified and contaminates the output. This problem was cured by filtering the control voltage used to switch the diodes.

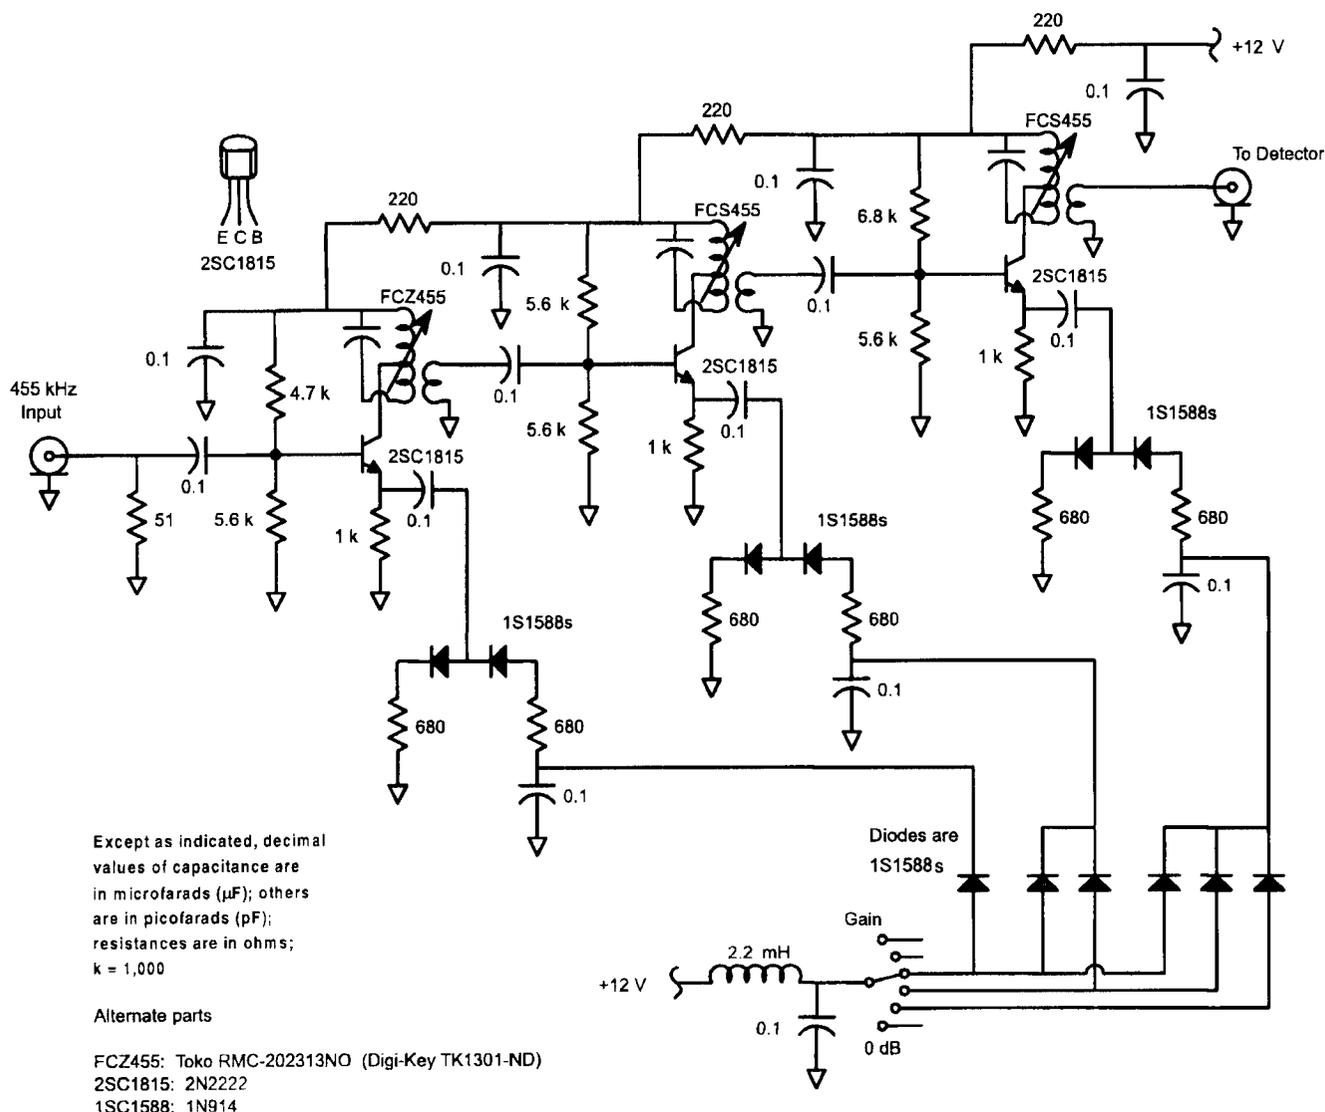


Fig 4—Schematic of 455 kHz IF amplifier.

Detector

The log detector is based on the NE604 and is shown in Figure 5. The circuit is based on the applications circuits shown for the NE604 in the Signetics data book.⁷ The NE604 is a high-gain device, so take care to use short leads on bypass capacitors and minimize feedback paths. A 455 kHz ceramic filter between the IF stages of the NE604 helps with selectivity. The inputs to each stage are loaded with 510 Ω to improve stability, as recommended by the Signetics data book.

The detector output is taken from the RSSI output (pin 5) of the NE604. This is a current source, terminated in a 100 k Ω resistor to give a voltage output of about 0.5 V per 10 dB. Signetics warns that an RSSI output greater than 250 mV with no signal input indicates possible regeneration or oscillation. After building the detector, ground the wiper of the 500 Ω vertical-offset potentiometer and measure the RSSI voltage at pin 5 of the NE604 with a high-impedance voltmeter while the detector input is disconnected. A voltage of more than 250 mV indicates possible stability problems in the detector.

An LM10 op amp in a voltage-follower configuration buffers the NE604 RSSI signal and provides a low-impedance output to drive the display. I chose the LM10 because it works well at input and output voltages down to its negative supply. This allows grounding the IC's negative supply, which eliminates the need for a negative power-supply voltage. A small, adjustable positive voltage summed into the LM10 input provides a vertical-offset adjustment to align the trace on the display with a convenient graticule.

A 6.2 V Zener diode supplies power to the NE604.

Scan Generator

The scan generator, shown in Figure 6, provides sweep voltage for the display and tune voltage for the VCO in the second converter. The sweep generator is based on a 555 timer and a current source consisting of a pair of PNP transistors. The 555 timer is connected in the usual manner for an astable multivibrator,⁸ except that the timing capacitor is charged from the current source instead of a resistor. The 0.47 μ F timing capacitor is charged until its voltage trips the threshold comparator in the 555 at $\frac{2}{3}$ Vcc. When that happens, the timer trips and discharges the timing capacitor through the 470 Ω resistor until its

voltage falls below the trigger level of $\frac{1}{3}$ Vcc. Thus, the voltage across the timing capacitor oscillates between $\frac{1}{3}$ Vcc and $\frac{2}{3}$ Vcc. Because the capacitor is charged from a constant current source, the rising part of this waveform is a linear sawtooth wave. Its period is determined by the current source and the capacitor value. When the timer trips, the output of the 555 could be used for blanking the trace if desired. Vcc for the 555 and current source is supplied by a 9V regulator to make the sweep period and magnitude of the sweep voltage, independent of variations in the panadaptor's supply voltage. One half of an LM10 op amp amplifies and offsets the voltage swing across the timing capacitor to provide a 0 to 10 V sweep signal for the display.

The tune voltage for the VCO in the second converter is provided by a 741 op amp. A portion of the 0 to 10 V sweep signal (determined by the setting of the **SPAN** control) is summed with an adjustable dc voltage (determined by the setting of the **CENTER FREQUENCY** control) and a fixed negative voltage generated by a second 741 op amp. The result is a sawtooth voltage that is adjustable from 0 to 4 V. The sawtooth is superimposed on a fixed voltage adjustable from about 4 to 12 V. With the **SPAN** control set at maximum, the panadaptor sweeps

± 100 kHz about the center frequency. The center frequency is adjustable by about ± 100 kHz, but sweep linearity begins to suffer at the edges.

Because the tuning voltage exceeds the +12 V supply voltage, the output of a small 9 V dc to dc converter is connected in series with the +12 V supply to obtain about 20 V, to power the op amp that supplies the tune voltage. Another 9 V dc to dc converter supplies the negative voltages in the scan generator. These dc to dc converters—about the size of postage stamps—were found in a surplus electronics store. Charge pumps based on the 555 timer or LTC1144 should work equally well. The dc to dc converters used here generate a considerable amount of electrical noise, and they require input filters to prevent their noise from contaminating the +12 V supply.

Construction

The RF circuits were all built "dead bug" or "ugly" style on unetched-copper PC stock. This method gives maximum flexibility in making changes and experimenting as the circuits are developed—as well as being an excellent construction method for RF circuits. The scan generator was built on perf board. The detector is susceptible to energy radiated from the LOs, so the first and second converters were en-

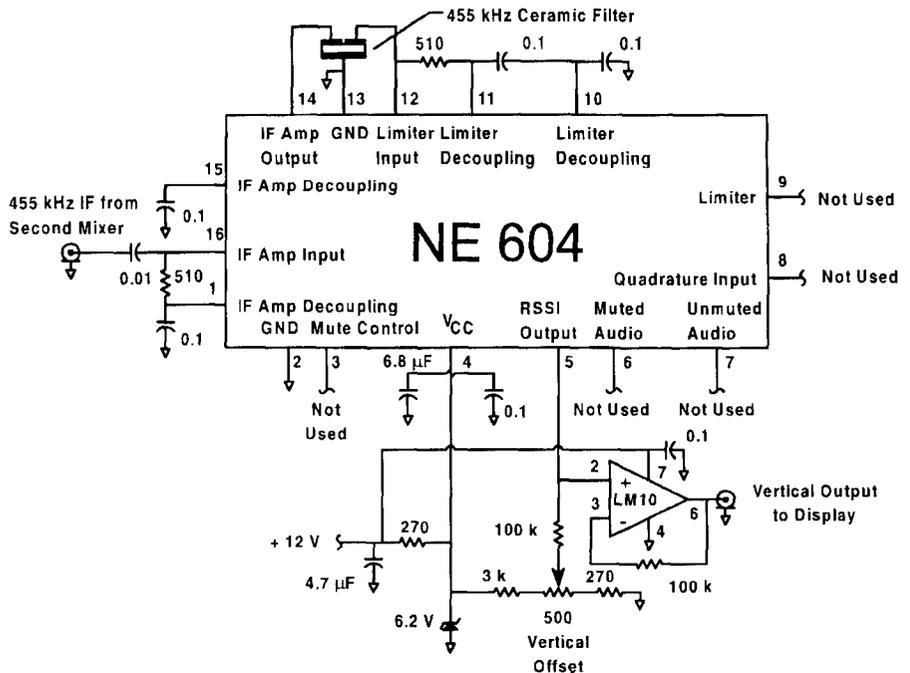


Fig 5—Schematic of detector.

closed in shielding cans. The detector circuit should probably be enclosed in a shield can also, but I haven't done this yet. Short lengths of miniature 50 Ω coax interconnect the stages.

Displays

Rather than design a display especially for the panadaptor, I use a standard oscilloscope. The panadaptor provides a horizontal sweep voltage of 0 to 10 V and a vertical output of 0 to 5 V.

Receiver Interface

The panadaptor should connect to the receiver at a point right after the first mixer and before any narrow-band filters. Some receivers and transceivers have a built in connector for a panadaptor. I found that the panadaptor output on my ICOM 720A was about 40 dB below the input signals from the antenna, even with the transceiver's preamp on. Figure 7 shows a FET buffer circuit that I constructed on a piece of PC board and placed inside the transceiver. This circuit presents a high impedance to the radio's mixer and provides a 50 Ω output for the panadaptor's input. The bandwidth of the tuned circuit is several megahertz. Change the values if you use other IF frequencies.

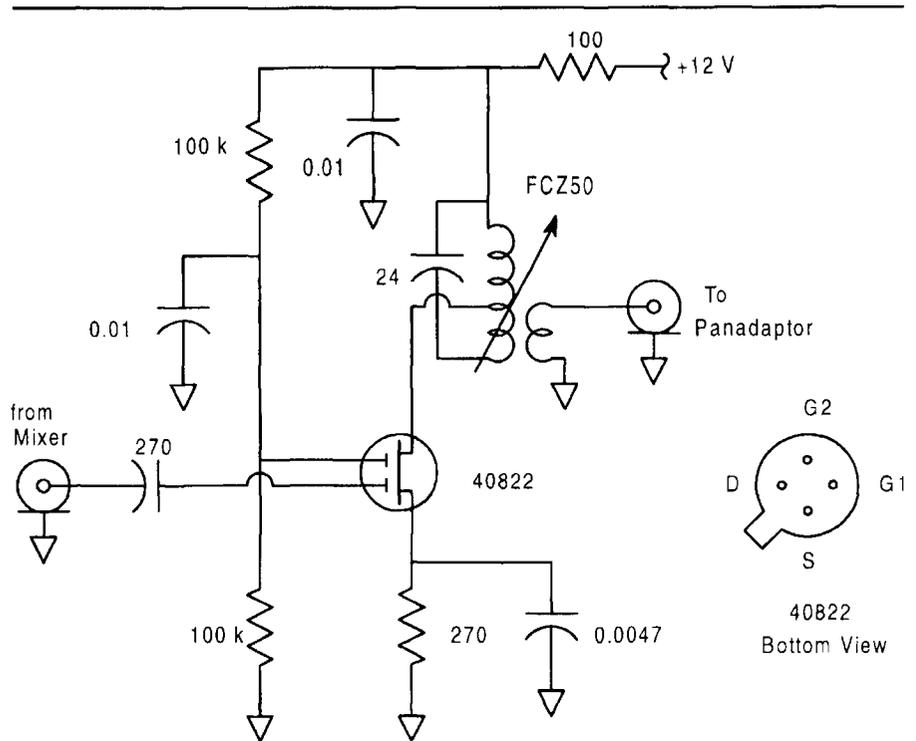
Performance and Use

The panadaptor has enough sensitivity to show atmospheric noise on the HF bands. For strong signals, such as international short-wave broadcasts, set the gain to minimum. On the amateur bands, set the gain to medium or maximum for good results. I find that leaving **SCAN** set to maximum (about ±100 kHz) gives a good view of the band in the general vicinity of the listening frequency. It's easy to spot signals as they come on the air, and you can recognize different types of modulation with a little practice. It's also interesting to watch occasional signals sweep across the band.

When there's a signal on the screen, a false response occurs at the beginning of each sweep, probably because the detector responds to signals during sweep retrace. Blanking the screen during retrace and for a short period as each sweep begins would eliminate this response, but it is not worth the effort for this application.

Future Improvements

All projects have room for improvement: The Signetics data book leads me to believe that the RSSI output of the NE604 has more dynamic range than I



Except as indicated, decimal values of capacitance are in microfarads (μF); others are in picofarads (pF); resistances are in ohms; k = 1,000

Alternate Parts:

FCZ50: Toko BTKXNS-T1049Z
(Digi-Key TK1409-ND)

Fig 7—Schematic of receiver mixer buffer.

have observed. More work could optimize the signal levels and gain distribution around the detector for better log fidelity and more dynamic range.

I didn't spend a lot of time optimizing the VCO. It should be possible to get a wider, more-linear tuning range.

On the VHF bands, there is much less atmospheric noise than at HF. It would be good to add a low-noise preamp to the panadaptor input, before the input band-pass filter. That would be something like one of the dual-gate FET preamps in common use on 6 meters.

Because it runs from +12 V, the panadaptor is well suited for portable work, but it requires a display that also runs from +12 V. One might construct a converter for use with a raster-scan display, such as a cheap video monitor or portable TV set. Computers also offer display possibilities. A desktop, notebook or palm-top PC running an oscilloscope program could easily serve as a display.

Notes

¹Louis Hutton, W0RQF, "The Electronic Eyeball," *QST*, Jan 1959, pp 37-40.

²Albert Helfrick, K2BLA, "A Simple Spectrum Analyzer," *RF Design*, Jan 1988, p 35.

³David Tomanek, "Inexpensive Spectrum Analyzer IF Uses Switchable Filtering," *RF Design*, Jun 1995, p 72.

⁴Mini Circuits Labs, PO Box 350166, Brooklyn, NY 11235-0003; tel 718-934-4500, 417-335-5935, fax 718-332-4661; URL <http://www.minicircuits.com/>.

⁵*The 1989 ARRL Handbook*, p 31-18.

⁶*The 1989 ARRL Handbook*, p 10-8.

⁷*RF Communications Data Handbook*, Signetics Co, 1990.

⁸Bob Marshall, WB6FOC, "Operational Characteristics of the 555 Timer," *Ham Radio*, Mar 1979, p 32.

First licensed in 1957 as WN6SFH, Bob Dildine has been primarily interested in the technical aspects of Amateur Radio. He received a BSEE from San Diego State College in 1967 and MSEE from the University of California at Berkeley in 1973. He then joined Hewlett-Packard as an R & D engineer. He has done precision analog design on microwave synthesizers and network analyzers. He is currently a Project Manager working on microwave test systems at Hewlett-Packard's Custom Solution Center in Tokyo, Japan. □□

The Twin-T Filter

*This classic circuit is often overlooked by builders.
Come see how it works and why it belongs in your repertoire.*

By Parker R. Cope, W2GOM/7

The twin-T is an RC frequency-selective network whose output has a null and zero phase shift at only one frequency. The twin-T is a three-terminal network with one input, one output, and one common terminal. As a result, it may be more conveniently used than a Wien bridge, which is the more common frequency-selective RC network. Also, the twin-T has a null, while the Wienbridge must have a separate resistive branch to complete the bridge and provide the null. Theoretically the output of the twin-T can be zero. The rate of change of phase shift with frequency at the null and the depth of the null is dependent on the tolerances of resistors and capacitors used in the network. While a rigorous analysis of the twin-T is rather complex, a qualitative description sheds some light on the "whys" and "hows" of component

selection. The appendix calculates how the null is developed and gives a feel for the sensitivity to component tolerances.

The classical twin-T is shown in Fig 1A. The twin-T has two branches, one provides a low-pass phase-lag characteristic, while the other provides a high-pass phase-lead characteristic. When the magnitudes of the signals from the lagging branch and the leading branch are equal and their phases are opposite (180° apart), the two signals cancel and the phase shift through the network is zero. In Fig 1B, the branch with R1, R2 and C3 provides the phase-lagging signal, while the branch with C1, C2 and R3 provides the phase-leading signal.

The null frequency of the twin-T is determined by two equations:

$$\omega_0^2 = \frac{1}{(R1+R2)R3 \cdot C1 \cdot C2} \quad (\text{Eq 1A})$$

and

$$\omega_0^2 = \frac{(C1+C2)}{C1 \cdot C2 \cdot C3 \cdot R1 \cdot R2} \quad (\text{Eq 1B})$$

The null condition as a function of component parameters is:

$$\frac{R1 \cdot R2 \cdot R3}{R1+R2} = \frac{(C1+C2)}{C3} = N \quad (\text{Eq 2})$$

N is a real (not imaginary) number that may vary from 0 to ∞, but the optimum value is N = 1. When N = 1, the phase shift for a given change in frequency is maximum.

From Eq 2, when R1 = R2 = R and C1 = C2 = C, then R3 = R/(2N) and C3 = 2C/N. When N = 1, R3 = R/2 and C3 = 2C, but if N is chosen to be 1/2, R1 = R2 = R3 and C1 = C2 = C3. When N = 1/2, the rate of phase change with frequency is 6% lower than when N = 1, but it has the advantage of equal-value components. In any case, the null occurs at:

$$\omega_0^2 = \frac{N}{R^2 C^2} \text{ or } f_0 = \frac{N}{2\pi RC} \quad (\text{Eq 3})$$

In the optimum twin-T (N = 1), R1 = R2 = R, R3 = R/2, C1 = C2 = C, and C3 = 2C. At null, $f_0 = 1/2\pi RC$ and $R = -jX_C$. The two branches are shown separated in Fig 2 with the normalized values of

R and $-jX_C$ shown. The lagging branch at null is shown in Fig 2A and the leading branch in Fig 2B. Appendix A calculates the lagging current in R2 at null to be $|0.353 \times E_{IN}/R| \angle -45^\circ$, or $0.249 - j0.249 \Omega$ and the leading current in C2 branch at null to be $|0.353/X_C| \angle +135^\circ$, or $-0.249 + j0.249 \Omega$. The current in the output is the sum of the lagging-branch current and the leading-branch current:

$$I_{out} = [0.249 - j0.249] + [-0.249 + j0.248] = 0 + j0$$

The twin-T can be used as a band-stop filter or—with an appropriate amplifier—a band-pass filter, or as the resonator in an RC oscillator. The frequency response of a band-pass frequency-selective amplifier is essentially the inverse of the rejection-network response. Fig 3 shows a simple single-stage frequency-selective amplifier. The amplifier has negative feedback provided through the twin-T; at the null, there is no negative feedback and the output is maximum. Off the null, where twin-T attenuation is small, the negative feedback is large, and the gain is low. In any amplifier with feedback, attention to the rate of change of total phase shift and gain is important. The high-frequency response of the amplifier used in a frequency-selective application can be neglected, but the low-frequency response is extremely important. The constants of the coupling network— C_C and R_g in Fig 3—must be chosen so that the cutoff frequency is much less than the null frequency of the twin-T. Enough so that additive phase shifts cannot cause the amplifier to oscillate. Conservative design criteria set the -3 dB-response frequency of the coupling network at approximately 0.1 of the twin-T network null frequency. That is, $R_g C_C = 10/(2\pi f_0)$, and the added phase shift is about 10° .

For any application, the basic twin-T has a major shortcoming in that its Q is fixed at 0.25. This limitation can be overcome by introducing positive feedback. Fig 4 shows a band-stop filter with positive feedback. Fig 4C shows a simplified configuration with the requirement that R1 of the twin-T is very much greater than $(1 - K)R$. Q can be increased by selecting a positive value of K near, but less than, unity:

$$Q = \frac{1}{4(1-K)} \text{ or } K = 1 - \frac{1}{4(Q)} \quad (\text{Eq 4})$$

For example, if a band-stop filter is required to have a center frequency of 1000 Hz and a 3 dB bandwidth, β_{3dB} , of

100 Hz, its Q must be 10 ($Q = f_0 / \beta_{3dB}$). From Eq 4, a Q of 10 requires a positive feedback of 0.975, $K = 1 - (1/40)$. In Fig 4, assume $KR = 1000 \Omega$ and $(1 - K)R = 25.6 \Omega$. With standard 1%-tolerance resistors, $(1 - K)R = 25.5 \Omega$ and $KR = 1000 \Omega$. With these nominal values, $K = 0.975$ and $Q = 10.05$. Of course, if $(1 - K)R$ and KR are parts of a potentiometer, Q can be varied.

The values used in the twin-T are calculated using Eq 3. For example, if 120 Hz is to be rejected, and C1 and C2 are chosen to be 2200 pF, R1 and R2 are calculated to be 602 k Ω . Tight tolerances are more readily obtained in small values of capacity, while precision resistors with 1% tolerances are more readily available. 1% RN60s are available to 1 M Ω . Therefore, C is chosen to be as small as practical and the resistors calculated for the frequency. Capacitors with $\pm 5\%$ tolerance are available, although they may not be normally stocked. For example, 5%-tolerance, 2200 pF units are available for C1 and C2, and two are connected in parallel for C3. The nearest standard resistors are 604 k Ω for R1 and R2, 301 k Ω for R3.

The simplified version in Fig 4C can use a source follower for the buffer amplifier as shown in Fig 5, but with limited Q enhancement. R_g provides the gate dc return and its value is not critical. The gate impedance can certainly meet the open-circuit requirements for the twin-T load, but the voltage gain of the source follower limits the maximum gains that can be realized.

The voltage gain, K, of the source follower is often assumed to be unity, even though it is always less than one. The voltage gain of a source follower is expressed as:

$$K_{sf} = \frac{g_{fs} \cdot R_k}{g_{fs} \cdot R_k + 1} \quad (\text{Eq 5})$$

Where g_{fs} is the forward transconductance of the JFET given in Eq 10, and R_k is the source resistance. The 2N5457 JFET data sheet shows that the gate/source voltage V_{gs} is 2.5 V when drain current I_D is 0.1 mA. In *Designing with Field Effect Transistors*,¹ Evans gives the basic interrelationships of JFET parameters as:

$$I_D = I_{DSS} \left(\frac{1 - V_{gs}}{V_{off}} \right)^n \quad (\text{Eq 6})$$

¹Notes appear on page 48.

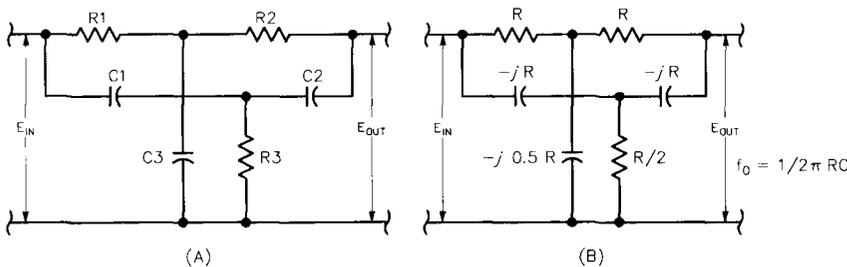


Fig 1—The classical twin-T notch filter. (A) The twin-T has a null in the output. (B) The twin-T is comprised of two branches.

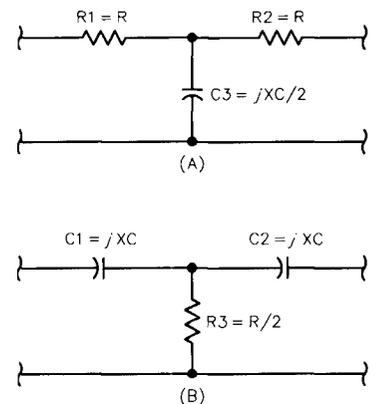


Fig 2—The twin-T has two branches: (A) lagging branch (B) leading branch.

Evans states: "Some texts indicate a value of 3/2 for n; however, experimental measurements on a number of N-channel geometries indicate the exponent, n, is close to 2."

Taking $n = 2$, Eq 6 can be rearranged to solve for V_{gs} and V_{off} :

$$\frac{V_{gs}}{V_{off}} = 1 - \sqrt{\frac{I_D}{I_{DSS}}} \quad (\text{Eq 7})$$

$$V_{gs} = V_{off} \left(1 - \sqrt{\frac{I_D}{I_{DSS}}} \right) \quad (\text{Eq 8})$$

$$V_{off} = \frac{V_{gs}}{\left(1 - \sqrt{\frac{I_D}{I_{DSS}}} \right)} \quad (\text{Eq 9})$$

$$g_{fs} = \frac{2I_D}{V_{gs} - V_{off}} = 2 \frac{\sqrt{I_D I_{DSS}}}{V_{off}} \quad (\text{Eq 10})$$

Where I_D = drain current for the particular value of V_{gs} .

I_{DSS} = drain current when V_{gs} is zero.

V_{gs} = gate to source voltage.

V_{off} = gate voltage required to reduce I_D to 1 μ A, essentially zero.

g_{fs} = the forward transconductance; the change in I_D for a change in V_{gs} ,

$$\frac{\Delta I_D}{\Delta V_{gs}}$$

Given I_D , I_{DSS} and V_{gs} , V_{off} can be calculated using Eq 9. With V_{off} known, g_{fs} can be calculated using Eq 10. For the 2N5457 with $I_D = 0.1$ mA and $V_{gs} = 2.5$ V, V_{off} calculates to be 3.06 V, and $g_{fs} = 3.57 \times 10^{-4}$. The 2.5 V V_{gs} can be developed across an R_k of 25 k Ω , and the gain, K_{sf} , of the source follower will be about 0.9. The maximum Q then is 2.5.

Higher values of gain K can be generated with the amplifier shown in Fig 6A. The JFET, a 2N5457, directly drives the base of a 2N3906 PNP transistor. The PNP multiplies the drain

current of the JFET and consequently multiplies the effective g_{fs} . The minimum h_{FE} of a 2N3906 is specified as 100 to 300 at 10 mA, so the collector current is $100 \times I_D$ to $300 \times I_D$. Since transconductance is defined as the change in output current for a change in input voltage, the effective transconductance is multiplied by h_{FE} . For $I_D = 0.1$ mA, the g_{fs} of 3.57×10^{-4} is multiplied by h_{FE} . When $h_{FE} = 100$, the effective g_{fs} is 3.57×10^{-2} . The 2.5 V V_{gs} is developed by 10 mA of collector current in an R_k of 250 Ω . The voltage gain is still 0.9, but a resistor RC can be placed between the collector and source to produce a gain greater than one. The gain, A_v , from the JFET gate to the PNP collector can be expressed as:

$$A_v = K_{sf} \left(1 + \frac{RC}{R_k} \right)$$

When $K_{sf} = 0.9$, $RC = 33$ Ω and $R_k = 250$ Ω , $A_v = 1.02$. A potentiometer between the collector and ground can vary the Q of the twin-T from 0.25 to infinity.

The twin-T can be used in a fre-

quency-selective band-pass amplifier as shown in Fig 6B. The open-loop amplifier must fulfill the following requirements:

- The band-pass characteristic must be essentially flat across the frequency range in which the selected frequency will lie.

- The amplifier must be essentially linear. The twin-T provides negative feedback around the amplifier. Low attenuation provides strong negative feedback for frequencies outside the notch, while frequencies within the notch are attenuated and produce no negative feedback. The high-frequency response of the amplifier can be neglected, but the low-frequency response is extremely important. The constants of the coupling network, C_C and R_k , must be chosen so that the network phase shift at the null frequency is very small and adds little to the overall phase shift at the null frequency. A conservative design would set the coupling circuit's cutoff frequency at approximately 0.1 of the null frequency. Then the phase lead introduced by the

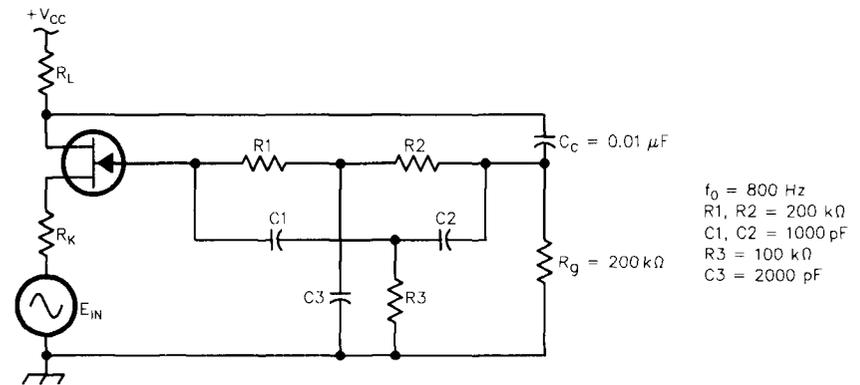


Fig 3—A frequency-selective amplifier has a peak at the twin-T null.

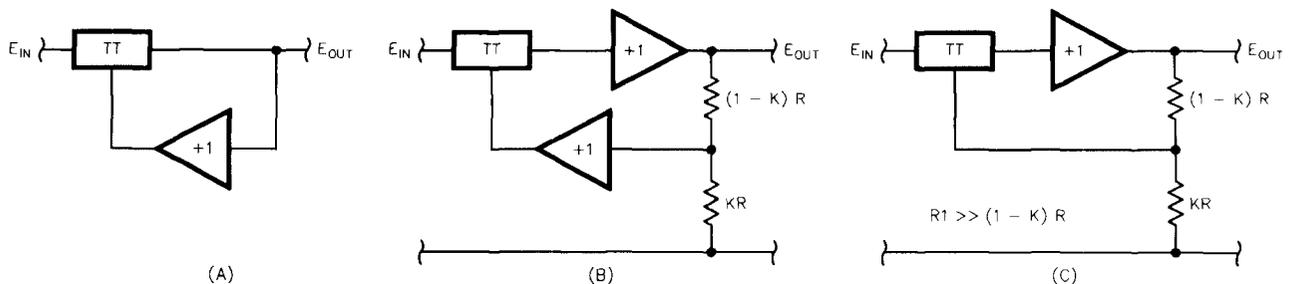


Fig 4—The Q of the twin-T is increased with positive feedback. (A) block diagram, (B) circuit realization, (C) simplified configuration with the requirement that R_1 of the twin-T is very much greater than $(1 - K)R$.

coupling network at the null frequency would be approximately 10° .

The dc from the gate can return to ground through the low-pass branch of the twin-T. In Fig 6B, the signal is shown introduced into the source of the amplifier to prevent the signal source from loading the twin-T. If the signal were introduced at the input of the twin-T, the signal would be coupled to the output through C_C . The signal can be introduced to the JFET's source through an emitter follower. The output impedance of an emitter follower is low, and it has negligible effects on the amplifier's ac characteristics, although it does affect the source bias.

The twin-T can be used as the resonator in an oscillator. In the frequency-selective amplifier the gain and phase shift are selected to avoid oscillation. In an oscillator, these characteristics are chosen to support oscillation. In the notch filter shown in Fig 4B, the circuit oscillates when $K > 1$. Unless steps are taken to limit the voltage swing, the amplifier will be driven into cutoff or saturation, and its output will be a square wave. If the output is taken from that of the twin-T, however, most harmonics are attenuated, and the waveform will be reasonably clean. The oscillator shown in Fig 7 uses two back-to-back diodes to limit the peak twin-T input to about 0.6 V, 1.2 V_{p.p.}. The filtered output can then be taken from the relatively low JFET source impedance. Be aware that the output loading can change the effective value of the source-to-ground impedance and, consequently, the gain of the amplifier.

The twin-T is a convenient RC filter, in some respects more convenient than the Wien bridge, because the T is a three-terminal circuit. Unfortunately, we must simultaneously vary the values of either three resistors or three capacitors to change the T null frequency, whereas the Wien bridge frequency depends on only two component values.

It is difficult to tune filters that use inductors (for example the bridged-T) to low or audio frequencies. Those filters also have relatively poor frequency stability; they are large, and their components are difficult to procure. For all of its disadvantages, the twin-T is a valuable filter that should be in any serious experimenter's bag of tricks.

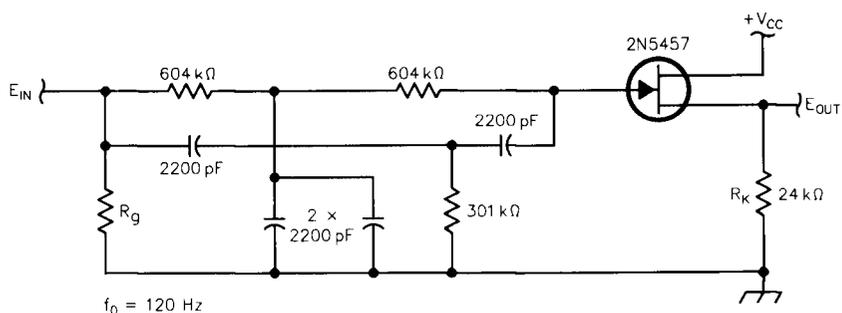


Fig 5—Q of the twin-T is limited by source-follower gain.

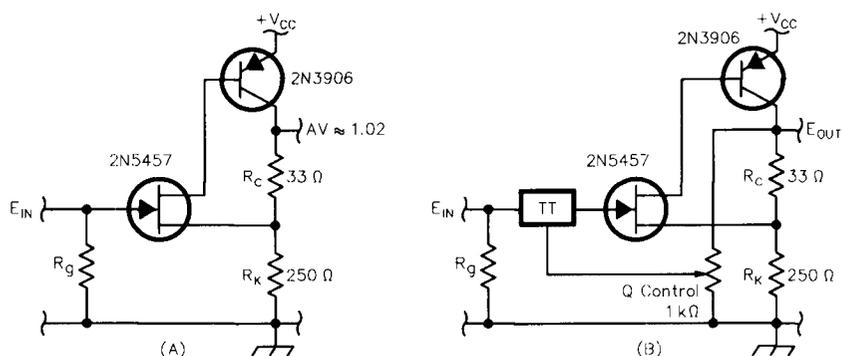


Fig 6—A simple amplifier provides positive feedback for high Q. (A) noninverting amplifier, (B) complete amplifier.

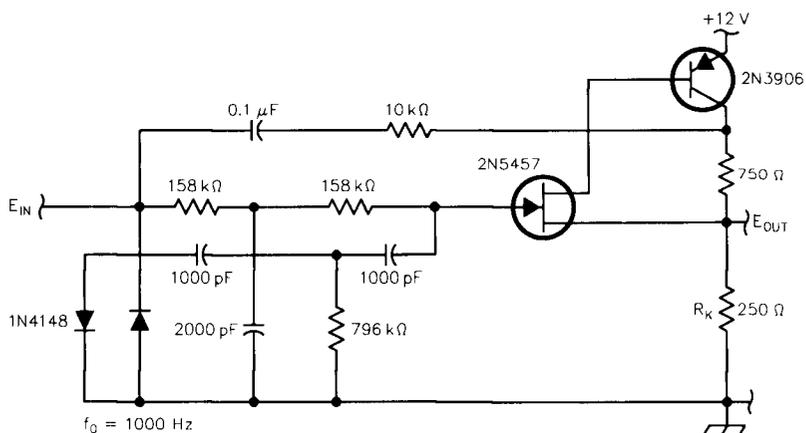


Fig 7—Diodes limit the voltages in a low-distortion oscillator.

Notes

¹Arthur D. Evans, Ed., *Designing with Field-Effect Transistors*. (New York: McGraw-Hill, 1981, 1990).

²Paul Danzer, N11I, Ed., *The ARRL Handbook for Radio Amateurs*, 75th edition (Newington: ARRL, 1997) p 26.19. 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>.

Appendix

The current in the output of the twin-T is the sum of the currents from the phase lagging and leading branches. In Fig A, the impedances, voltages and currents are shown for a null condition. At null, $R1 = R2 = R$, and $-jXC3 = -j0.5 R$. At null, the output voltage is zero, and the lagging circuit can be redrawn as Fig A1. The current in R2 is the output current of the lagging branch. Since $R2 = R$ the current in R2 is $I1$ ($I1 = E1/R$). The voltage across R2 is also the voltage across the admittance $Y1$ (R2 in parallel with $-jXC3$):

$$Y1 = G + j\beta = 1 + j2 = 2.236 \angle +63.43^\circ \quad (\text{Eq A1})$$

$$Z1 = \frac{1}{Y1} = 0.447 \angle -63.42^\circ = 0.2 - j0.4 \quad (\text{Eq A2})$$

$$\begin{aligned} \frac{E1}{E_{IN}} &= \frac{Z1}{Z_T} = \frac{Z1}{1 + 0.2 - j0.4} = \frac{Z1}{1.265 \angle -18.43^\circ} \\ &= \frac{0.447 \angle -63.43^\circ}{1.265 \angle -18.43^\circ} = 0.353 \angle -45^\circ \quad (\text{Eq A3}) \end{aligned}$$

$$I1 = \frac{E1}{R} = 0.353 \angle -45^\circ = 0.249 - j0.249 \quad (\text{Eq A4})$$

The leading branch is treated in a similar fashion. Figs A1C and A1D apply.

The current, $I2$, in $C2$ is $\frac{E2}{-jXC2}$.

$$Y2 = G + j\beta = 2 + j1 = 2.236 \angle +26.56^\circ \quad (\text{Eq A5})$$

$$Z2 = \frac{1}{Y2} = 0.447 \angle -26.56^\circ = 0.4 - j0.2 \quad (\text{Eq A6})$$

$$\begin{aligned} \frac{E2}{E_{IN}} &= \frac{Z2}{Z_T} = \frac{Z2}{0.4 - j1.2} = \frac{Z2}{1.265 \angle -71.56^\circ} \\ &= \frac{0.447 \angle -26.56^\circ}{1.265 \angle -71.56^\circ} = 0.353 \angle +45^\circ \quad (\text{Eq A7}) \end{aligned}$$

$$\begin{aligned} I2 &= \frac{E2}{-jXC2} = \frac{0.353 \angle +45^\circ}{1 \angle 90^\circ} = 0.353 \angle +135^\circ \\ &= -0.249 + j0.249 \quad (\text{Eq A8}) \end{aligned}$$

$$I1 + I2 = 0.249 - j0.249 + [-0.249 + j0.249] = 0 + j0$$

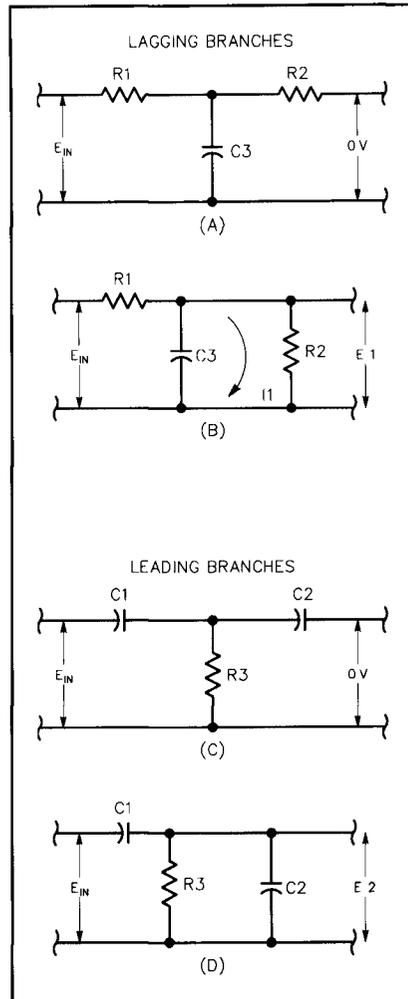


Fig A1—The equivalent circuits in the twin-T that produce a null at the output.

QEX Subscription Order Card

American Radio Relay League • 225 Main Street • Newington, CT 06111-1494 • USA
QEX, the Forum for Communications Experimenters, is available at the rates shown below.

For one year (6 bi-monthly issues) of QEX, ARRL Member* rates:

Subscribe toll-free with your credit card
1-888-277-5289

In the US, \$18.00

Renewal New Subscription

In Canada, Mexico and US by First Class mail, \$31.00

Name _____ Call _____

Elsewhere by Surface Mail (4-8 week delivery) \$23.00

Address _____

Elsewhere by Airmail \$51.00

City _____ State or Province _____ Postal Code _____

Payment Enclosed Charge:

*Non-members add \$12 to these rates

Remittance must be in US funds and checks must be drawn on a bank in the US. Prices subject to change without notice.

Account # _____

Good thru _____ Signature _____

4/98 QX

Surface Mount Chip Component Prototyping Kits—
Only \$49.95

INDIVIDUAL VALUES AVAILABLE

CC-1 Capacitor Kit contains 365 pieces, 5 ea of every 10% value from 1pf to 33µf. CR-1 Resistor Kit contains 1540 pieces, 10 ea of every 5% value from 10Ω to 10megΩ. Sizes are 0805 and 1206. Each kit is ONLY \$49.95 and available for Immediate One Day Delivery!

Order by toll-free phone, FAX, or mail. We accept VISA, MC, COD, or Pre-paid orders. Company PO's accepted with approved credit. Call for free detailed brochure.

COMMUNICATIONS SPECIALISTS, INC.
426 West Tait Ave. • Orange, CA 92665-4296
Local (714) 998-3021 • FAX (714) 974-3420
Entire USA 1-800-854-0547

Multiple-Octave Bidirectional Wire Antennas

*This simple, straight-line antenna works
across three harmonically related bands.*

By Robert Zavrel Jr, W7SX

With the sunspot cycle just leaving its theoretical minimum, the low bands are still experiencing unprecedented popularity. For several years, I have been using tree-supported wire antennas on 80, 40 and 20 meters with gratifying results. The total cost of a three-wire antenna farm suspended from high pine trees can be kept under \$500 (including tree climber's fees!). Yet with no towers, rotors, aluminum tubes, traps, or coax I managed 219 countries confirmed on 40 meters with 800 W and a casual operating style over the past four years. This represents impressive "bang for the buck." Furthermore, I accomplished this from California, where only five DXCC countries lie within 2000 miles!

Since moving to North Carolina, I've wanted to explore improving the multiband performance of a center-fed wire antenna. In particular, I've

wanted to achieve a bidirectional pattern on multiple bands, developing as much gain as possible by using the greatest possible wire length on each band. It is very convenient to know that a given antenna's performance will be optimum for Europe/ZL or South America/JA on several bands.

Dipoles using traps can provide multiband bidirectional operation with the added advantage of normalized feed-point impedances on multiple bands. However, trap dipoles are more expensive, heavier and provide no more gain than a dipole on any band of operation. Traps are associated with losses and can have problems dealing with high-power operation. Furthermore, coax-fed trap dipoles typically have narrow low-SWR bandwidths.

Although the feed-point impedance varies widely, a simple center-fed wire antenna exhibits a bidirectional pattern over 1.25 octave, if the wire is cut as a half-wave dipole on the lowest operating frequency. For example, an 80 meter dipole will function as a two-

element collinear on 40 meters and an extended double Zepp on 34 meters, which is just a bit too low for the 30 meter band. So, you get only two bands from this antenna. The use of open-wire feeders and an antenna tuner¹ solves the matching problem and leaves the entire antenna length active on all bands, with associated gain increases. I wanted to operate over two octaves (80, 40, 30 and 20 m), however, and achieve some gain over a dipole on the three higher bands. I also wanted to keep the design simple, lightweight, inexpensive and QRO friendly. (As an aside, it is often desirable to shorten the length of the wire on the lowest operating frequency with little or no sacrifice in performance.)

As the length of a center-fed $\lambda/2$ dipole increases, the broadside gain also increases. A 1λ center-fed dipole is also referred to as a "two element collinear" or "two half-waves in phase." The two-element collinear has a free-

8122 Reynard Road
Chapel Hill, NC 27516
e-mail rzavrel@vnet.ibm.com

¹Notes appear on page 52.

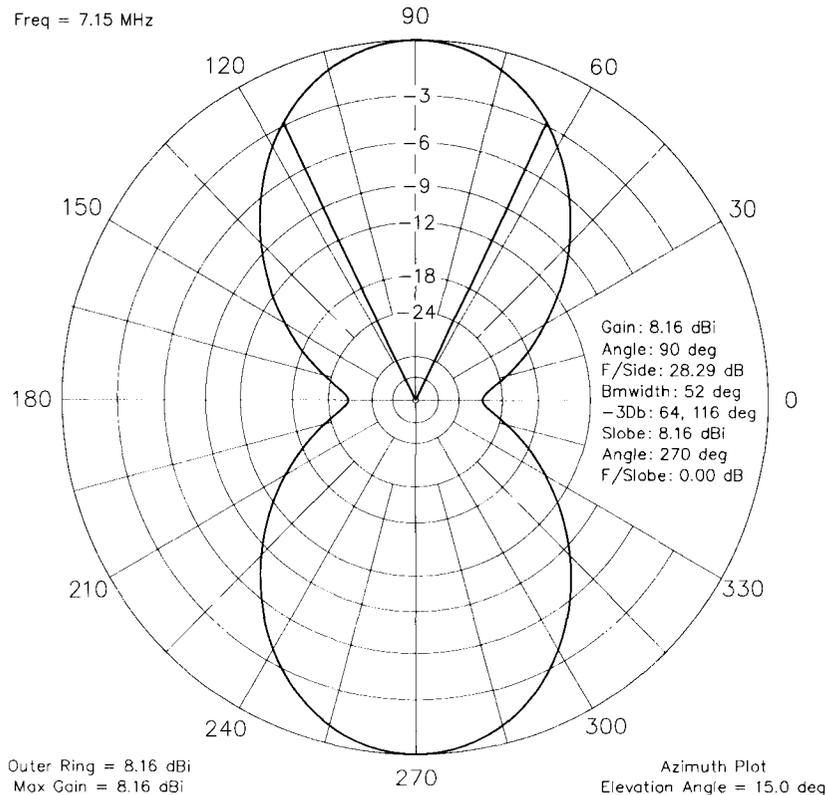


Fig 3—Azimuth plot (15° elevation) for the multi-octave antenna at 7.15 MHz (see Note 2 for modeling details).

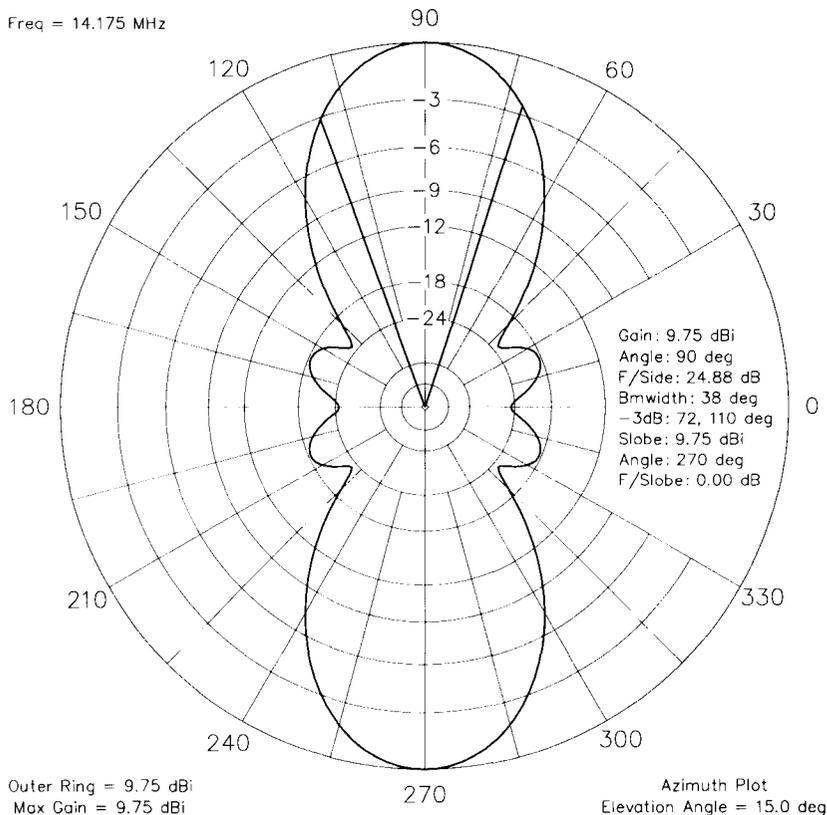


Fig 4—Azimuth plot (15° elevation) for the multi-octave antenna at 14.175 MHz (see Note 2 for modeling details).

America/JA. The 57 foot antenna is broadside to the east and west. This combination of antennas provides full-azimuth coverage on 80 meters, a slight north/south null on 40 meters and several nulls on 20 meters. My primary DX target areas are well covered, however.

Construction

There are many possible construction techniques for this antenna. I used 2x3x1/8-inch-thick Plexiglas plates for the five connection points along the antenna. I drilled holes in the plates that are just large enough to pass the antenna and stub wires; this provides some strain relief. The ladder line is fastened to the faces of the Plexiglas plates with 1/4 inch nylon nuts and bolts. Each bolt passes through the center of the ladder line's plastic web insulator and the Plexiglas. The non-conducting nuts and bolts help prevent arcing across the line.

I don't climb trees! I prefer to have an eyebolt secured to the tree trunk as high as the climber feels safe. The eyebolt holds a pulley, which holds the rope. [Editors note: By making the line twice as long—running from ground level to the block and back to form a continuous halyard—the upper end of the line can be retrieved from ground level just in case the antenna wire or one of the insulators breaks!] The rope connects to the end of the antenna and 35 pound counterweights.³ Keep the 114 foot antenna nearly horizontal. (The going rate for tree climbers is about \$40 per tree.)

Conclusion

This very simple, lightweight, inexpensive bidirectional antenna solution permits operation on 80 through 20 meters or 40 through 10 meters. These antennas have proven themselves well.

References

- ¹"Up Front in QST," QST, April 1995, p 11.
- ²Figures 2, 3 and 4 were modeled on EZNEC 1.0 for presentation here on the standard ARRL antenna grid. The maximum gains and general shapes conform well to the author's plots, while the nulls are somewhat (about 6 dB, maximum) deeper than on the plots shown here. The models use high accuracy modeling over average (0.005 S/m, $\epsilon = 13$) ground. You can download the EZNEC description files from the ARRL "Hiram" BBS (tel 860-594-0306), or the ARRL Internet ftp site: oak.oakland.edu (in the pub/hamradio/arrl/qex directory). In either case, look for the file ZAVREL.ZIP.
- ³"Technical Correspondence," QST, March 1992, p 84.

RF

By Zack Lau, W1VT

Designing Band-pass Filters

I've been asked to provide a little more detail about how one actually designs RF circuits. I'll attempt to describe how to design a 6 meter band-pass filter. I like to show actual circuits one can build, so I decided to let the bandwidth vary, while keeping the parts similar. A more theoretical presentation might fix the bandwidth and let the parts values fall where they may—even though those parts might not be available.

Perhaps the toughest step in designing LC band-pass filters is determin-

ing the inductor values, particularly in terms of the Q (quality factor) of what is available. It may be necessary to measure the value—I've not located or derived useful equations for calculating Q values of toroidal inductors, which are quite popular in modern designs. The self-shielding nature of toroids makes them much more convenient than solenoidal homebrew air-core inductors.

It can be rather time consuming to make your own shielded boxes; even though they can be made cheaply out of surplus double-sided, copper-clad circuit board. Commercial inductors mounted in little metal cans are convenient, but they seldom have enough Q for anything but casual applications. At 50 MHz, the Q is typically around 80. The Q of the coils in little metal cans is

proportional to frequency—it is even worse at lower frequencies. In contrast, the Q of toroidal inductors can be quite good at HF (often over 300), while the best I've seen at 6 meters is around 170.

The Q , more accurately the unloaded Q of the inductor, often sets a limit on how narrow you can make a band-pass filter without incurring excessive loss. While the capacitor loss also plays a factor, it may be small enough to be ignored in the initial design. Capacitor Q s around 1000 aren't unusual for high-quality silver-mica capacitors. Porcelain-chip capacitors have even higher Q s: A Q of 31,000 is estimated for a 470 pF chip capacitor at 3.5 MHz. For low-loss receiving circuits, you often want the ratio of unloaded to loaded Q s to be around 18, while a factor of five is often acceptable in transmitting applica-

¹Notes appear on page 56.

tions. These are just rules of thumb—in practice a designer will carefully model the filter with lossy components and see how the resultant filter affects actual system performance.

Solid State Design describes a simple method for measuring inductor Q.¹ If you have a computer program like *Amateur Radio Designer*, you can get even better measurements by correcting for the inaccuracy created by the test fixture. If you do the calculations by hand, I'd be inclined to skip the extra work, since the error usually isn't a problem if the Q is slightly underestimated—it rarely takes much work to add loss to a circuit.

A good source of measured values appears in the book *QRP Power*—it includes a "RF" column I wrote with five pages of toroidal-inductor measurements.² Another source is the *Amidon Tech Data Handbook*,³ which has graphs of Q versus frequency.

After winding many inductors, I found that eight turns of #22 or #24 enameled wire on a T-50-10 offers decent performance at 50 MHz. I typically achieved a Q around 170, though a sample on an old T-50-10 core had a Q of 140. Further investigation may yield better inductors, but this seemed like a reasonably good combination. The worse-than-HF results are easy to explain—materials with lower permeabilities require more turns that result in more copper loss, while materials with higher permeabilities have too much core loss to produce high-quality coils. However, low loss isn't everything in a narrow-band filter. You also need to consider temperature stability. This is why I didn't consider using ferrite material—a typical temperature coefficient is 0.15% / °C. Thus, over a 20° temperature variation, one might ex-

pect the inductance to vary by 3%. This is considerable for a LC filtering having a 3% bandwidth. In contrast, #10 iron-powder toroidal inductors typically vary by just 0.3% over the same temperature range, a factor small enough to safely ignore.

I also looked at what happens if you squeeze the turns together, rather than distribute the turns evenly over the core.

Core covered (%)	Apparent Inductance (μH)	Q
25	0.46	172
50	0.34	171
75	0.31	168
90	0.30	167

Thus, you can squeeze the turns together to get more inductance without

adversely affecting Q. If you have a frequency-sweep setup to test filter performance, it is a simple matter to squeeze and separate turns for best performance. (See Photo A.) A filter like that in Fig 1 might work well. I typically choose the value of the coupling capacitor, C_{coup} , first. It has the most effect on bandwidth. Then I design the capacitive divider, an impedance transformer, to efficiently match the filter to 50 Ω loads. There are usually several possible designs, each with a different amount of passband ripple. I like to keep the ripple to less than 1 dB within an amateur passband. On the other hand, it may be worthwhile to let the ripple of a multiband filter

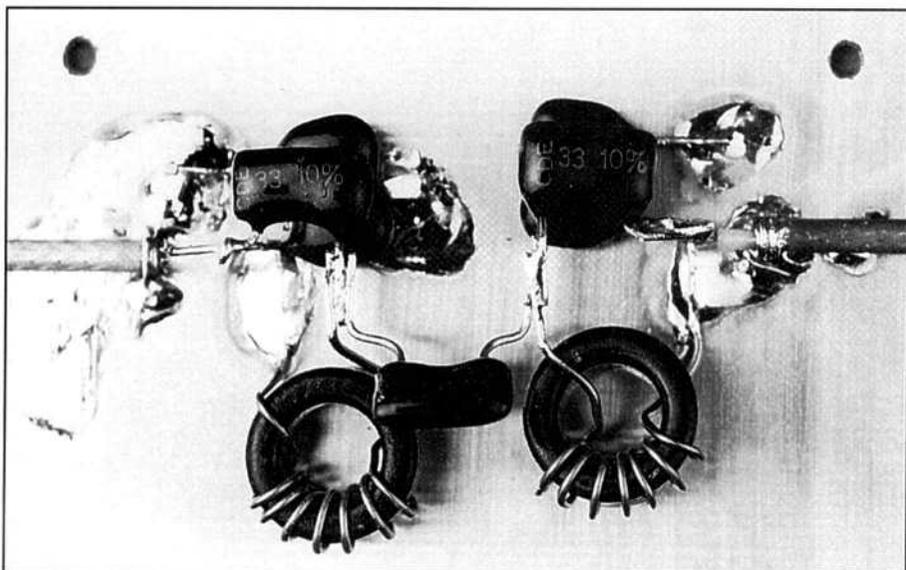


Photo A—A 6 meter band-pass filter using compressed toroidal windings for tuning.

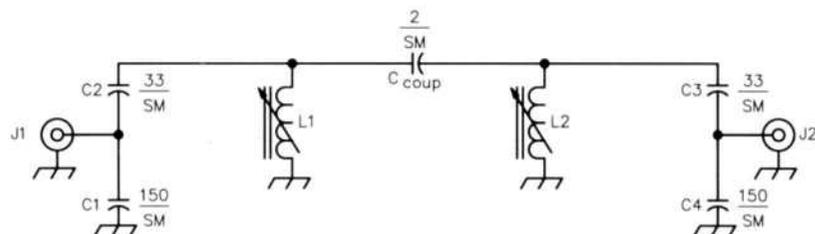


Fig 1—50-MHz band-pass filter with 1.5 dB insertion loss: -1 dB BW 4 MHz; -3 dB BW 5 MHz. All capacitor values are in pF.

J1, J2—Your favorite coaxial connectors
L1, L2—8 turns #22 or #24 enameled wire on T-50-10 iron-powder toroidal core. Inductance value is tweaked by

squeezing turns together (see text). *Amateur Radio Designer* simulation indicates a nominal value of 340 nH. The Q is approximately 170 at 50 MHz.

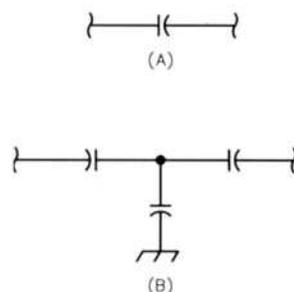


Fig 2—A is a T network to lightly couple circuits with practical capacitors. Circuit (B) simulates the low coupling of circuit (A).

rise to several decibels, if the peaks conveniently coincide with amateur bands. For narrow-bandwidth circuits, like many single-band RF filters, C_{coup} can become quite small. Wes Hayward suggests using a network of three capacitors to effectively obtain small, nonstandard values, as shown in Fig 2. The two circuits aren't interchangeable; circuit 2B loads the resonant circuit with more capacitance. It requires a bit of redesign to go from one to the other. Alternatively, you could choose a standard value like 1 or 2 pF and vary the inductors and matching circuitry. This works better if you have more flexibility with the inductor value. Air-core inductors are much better, in this regard, than toroids.

When choosing small standard-value capacitors, it is important to consider the available component tolerances. While you may be able to obtain small values, the variation may be too great to assure repeatable designs without hand selecting components. This might be used to your advantage—a big bag of capacitors may provide a sufficient range to design filters requiring a num-

ber of different but closely spaced values. Such a custom design might be difficult to reproduce, but it could yield that high-performance filter you need to complete a one-of-a-kind design. Of course, a bag of capacitors purchased for such a purpose will often yield a batch of virtually identical values, while someone hoping to get those identical capacitors for a batch of QRP kits will get the ones you wanted.

Creating a good design that uses standard values may be a little too challenging. It may make more sense to add a pair of tuning capacitors, as shown in Fig 3. This results in a lot more design flexibility and makes circuit alignment with a fixed-frequency signal generator much easier. Cost is a disadvantage here: Cheap variable capacitors tend to degrade performance. For example, ceramic trimmer capacitors may have a Q of only 300, which may significantly degrade the Q of the resonant circuits. I've minimized loss by making the trimmers as small as practical; yet trimmers that are too small may have insufficient range for circuit tuning. Expensive air-variable

capacitors have much higher Qs, sometimes as high as 5000, but they may not be practical for low-cost amateur projects. The latest Newark catalog lists prices for some Johanson air-variable trimmers in excess of \$10. Another problem with cheap trimmers is their susceptibility to mechanical movement: They might not be suitable for circuits used in mobile applications. One of the better deals is the Johnson T6-5, a leaded part designed for PC-board mounting. It has a Q of 2000 at 1 MHz.

You may need to consider how the filter interacts with the rest of the circuit. What does the filter do with the unwanted signals? For instance, the circuit of Fig 1 presents high impedance to low-frequency signals outside its passband. Suppose you were terminating a frequency multiplier that works much better with a low-impedance termination. A much better alternative might be the circuit of Fig 4, which uses a tapped inductor for impedance matching. In addition, while active circuits like MMICs like to see open circuits that don't interfere with dc bias, diode circuits often require a ground return. An elegant design eliminates extra coupling capacitors and RF chokes.

A disadvantage of the circuit of Fig 4 is that it is difficult to precisely model the tap point. A tapped-inductor model isn't always available, so it may be necessary to improvise an approximation. One possibility is an ideal transformer and an inductor, as shown in Fig 5. With low-permeability-core inductors the turns ratio is apt to be a little bit higher than predicted by an ideal model. It may be necessary to iterate a few times between measured and modeled cir-

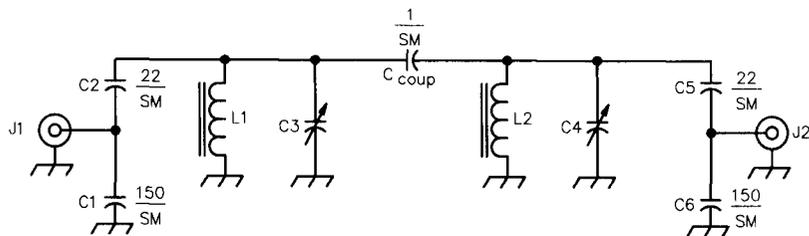


Fig 3—50-MHz band-pass filter with 2.5 dB insertion loss: -1 dB BW 1.2 MHz; -3 dB BW 1.6 MHz. All capacitor values are in pF.

L1, L2—8 turns #22 or #24 enameled wire on T-50-10 iron-powder toroidal core. 0.30 μ H with a Q of 170 at 50 MHz
C4, C5—4.5 to 65 pF polypropylene-film

trimmer capacitors. Sprague Goodman GYC65000. Digi-Key SG3009. Specified Q of 1000 minimum at 1 MHz.

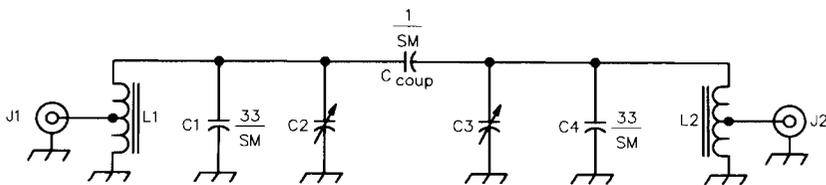


Fig 4—50 MHz band-pass filter with 2.9 dB insertion loss: -1 dB BW 1 MHz; -3 dB BW 1.6 MHz. All capacitor values are in pF.

C2, C3—4 to 20 pF. Sprague-Goodman GKG20011. Rated minimum Q of 300 at 10 MHz. Digi-Key SG1003-ND
L1, L2—8 turns #22 or #24 enameled

wire on T-50-10 iron-powder toroidal core. Tap 1 turn above ground. 0.30 μ H with a Q of 170 at 50 MHz.

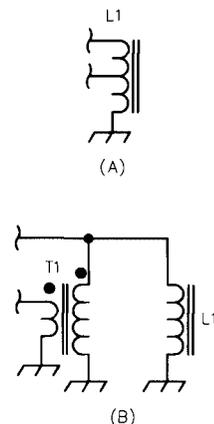


Figure 5—Modeling a tapped inductor with a transformer and inductor.

cuits to get an excellent design.

The next challenge is building the circuit. High-Q resonant circuits require low-loss connections to work properly. It's not quite so obvious that the ground leads are just as important as the "hot" leads—after all, just as much current flows out of an inductor as into the inductor. It helps to strap the resonating capacitors directly across their associated inductors. It may also help to mount the inductors orthogonal to each other, to minimize coupling between them. The coupling can help in some cases. The most common error with toroids is improperly counting turns—people forget to count that first turn. The *ARRL Handbook* has details. Finally, toroids can be wound in two different senses, as show in Photo B. Winding them in different senses helps when trying to mount them in a symmetrical layout. This may also explain why your friend's toroids drop neatly into a circuit board, while you can't get your board to look as neat.

Before building the circuit, you must also decide which parts should be adjustable. For casual applications, like the low-power stages of simple transmitters, the ability to peak each resonator for maximum output is often adequate. The loss will vary a bit, but a good design will accommodate small variations. Thus, tweaking the inductor winding spread or adjusting the capacitors for resonance are typically good enough. For more critical applications—such as the input filters of sensitive receivers—you want to adjust the coupling between the two resonators. The losses of the filter can vary considerably when dealing with low coupling values, such as 1 or 2 pF—tolerances as high as 50% aren't unusual. For mobile and portable applications, where resistance to vibration is important, I'd try a couple different

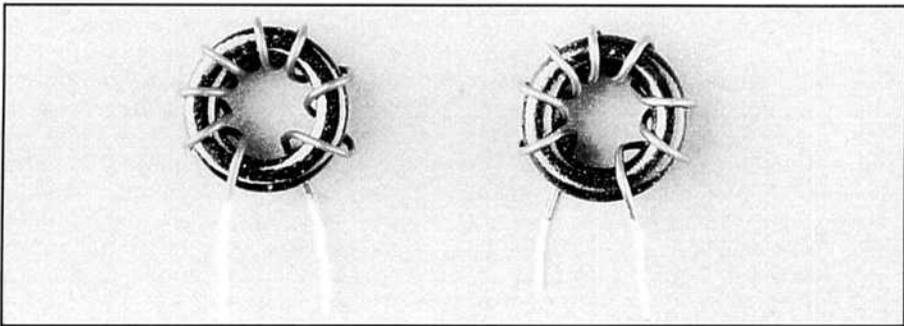


Photo B—Toroid cores wound in opposite senses.

capacitors and select the best one. The T-network idea might also work well; because the center capacitor is grounded, it is a convenient place to put a variable capacitor. Most trimmer capacitors allow you to ground the rotating element, so your screwdriver is close to ground potential when you adjust the trimmer. This significantly reduces the interaction between the screwdriver and your circuit. The bandwidth will vary a little with input and output coupling tolerances, but this is less critical because the capacitors are larger and the variation less.

For very critical circuits (like band-pass filters that determine a response of a spectrum analyzer), it makes sense to use a technique that properly sets the coupling between the filter sections. For example, the Dishal technique is described on page 241 of *Solid State Design*. It is very attractive for those with access to tracking generators and spectrum analyzers to just tweak stuff until it looks right. This is more effective with network analyzers that let you look at the return loss and transmission coefficient simultaneously.

While most of the design work is done with computer programs such as the ARRL's *Amateur Radio Designer*⁴ and the companion software to Wes

Hayward's *Introduction to RF Design*,⁵ it is by no means a necessity. It is also possible to do serious work using the tables in the *Handbook of Filter Synthesis*, by Anatol I. Zverev.⁶ Its price tag is also rather serious—\$192 the last time I looked.

Notes

¹D. DeMaw, W1FB, and W. Hayward, W7ZOI, *Solid-State Design for the Radio Amateur* (Newington: ARRL, 1986), Order No. 0402, page 240. 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>.

²J. Kleinman, N1BKE, and Z. Lau, W1VT, editors, *QRP Power* (Newington: ARRL, 1996), ARRL Order No. 5617. See Note A for ordering information.

³Bytemark Corporation, 7714 Trent St, Orlando, FL 32807; tel 800-679-3184, 407-673-3184, fax 407-673-2083; <http://bytemark.com/amidon/>

⁴*Amateur Radio Designer* (Newington: ARRL) Order No. 6796. See Note 1 for ordering information.

⁵W. Hayward, W7ZOI, *Introduction to RF Design* (Newington: ARRL, 1994), Order No. 4920. See Note 1 for ordering information.

⁶A. I. Zverev, *Handbook of Filter Synthesis* (New York: Wiley, 1967). □

Upcoming Technical Conferences

Central States VHF Conference

The 32nd annual Central States VHF Conference will be held in Kansas City, Missouri, on July 24-25, 1998. The conference will be held at the Adam's Mark Hotel, next to the Kansas City Sports Complex. Technical presentations for advanced and beginning operators interested in weak-signal modes of communication on the VHF and UHF bands, antenna-gain and noise-figure measuring contests will be featured. A wide range of activities is planned for hams, their family members and friends. For more information contact Denise, AJ0E, or Tom, K0TLM (e-mail aj0e@juno.com or k0tlm@juno.com). For hotel reservations, call 1-800-444-ADAM.

Microwave Update '98

This year's Microwave Update conference will take place on October 15-18, 1998, in Estes Park, Colorado

Key Contacts

Bill McCaa, K0RZ, Conference Coordinator and Registrations, PO Box 3214, Boulder, CO 80307-3214; tel 303-441-3069; e-mail WMcCaa@aol.com

John Anderson, WD4MUO, Program and Proceedings Coordinator, 3592 Ridge Rd, Nederland, CO 80466; tel 303-258-3711; e-mail wd4muo@webaccess.net

Don Nelson, N0UGY, Web page and Prize Coordinator, 3495 Longwood Ave, Boulder, CO 80303; tel 303-499-4990; e-mail n0ugy@worldnet.att.net

Conference Schedule

Thursday, October 15, 1998: 8 AM-5 PM, Surplus Tour (informal); 6 PM-Midnight, registration and informal get together.

Friday, October 16, 1998: 8:30 AM 5 PM, speaker presentations and

miscellaneous activities. 7 PM-12 AM, equipment swap, noise figure and impedance measurements. Saturday, October 17, 1998: 8:30 AM-4 PM, speaker presentations and miscellaneous activities. 6-9 PM, Conference Banquet and close.

Sunday, October 18, 1998: No formal conference activities.

Hotel Registration

The room rate for the conference is \$75 per night for single or double occupancy. A block of rooms has been set aside for conference attendees. Please register early with both the hotel and Microwave Update, so that we can increase the reservation block if necessary. Contact the Holiday Inn Resort to make reservations: Holiday Inn Resort, PO Box 1468, Estes Park, CO 80517; tel 1-800-803-7837 or 970-586-2333.

Call for Papers

Almost any paper dealing with technical, operating, editorial, humor, etc, aspects of Amateur Radio microwave communications. Paper deadline: Camera-ready hard copy must be received by September 4, 1998. Contact John, WD4MUO, if you wish to present a paper or contribute a paper to the *Proceedings*.

Surplus Tour

J. B. Saunders (electronic parts), Boulder
Western Test Systems (instruments), Broomfield
OEM (electronic parts), Colorado Springs
K & K (junk), Commerce City
C & M (scrap metal), Denver
Gateway Electronics (electronic parts), Denver
HRO (ham store), Denver
Tech America (electronic parts), Denver

T & H (electronic store), Ft Collins
Science & Technology (electronic parts and instruments), Longmont
Eagle Electronics (electronic parts), Longmont

Door Prizes and Auction

In addition to the usual and valuable donations that are made by many of the attendees and supporting commercial firms, we hope to have a surprise or two.

Points of Interest

Estes Park is adjacent to the Rocky Mountain National Park. The park entrances are nominally 4-5 miles west of the hotel. In the fall, the major added attractions in the park are bugling elk.

The Holiday Inn Resort is located within walking distance (0.5 miles) of downtown Estes Park. Estes Park is a resort/recreational town with a rich variety of boutiques. Within several blocks of the hotel, there is a strip mall with a major grocery store and typical related stores.

Amateur Radio Papers at URSI 1999?

I am looking for people who would be interested in presenting papers at an URSI Commission-F session. URSI is an international professional radio-science organization, and Commission F is primarily interested in VHF, UHF and microwave radio propagation issues. This covers any Amateur Radio frequency greater than 50 MHz. The amateur community has had a lot of experience with anomalous propagation, and I would like to see some of this experience presented in professional circles. Professionals are beginning to recognize the contribution of Amateur Radio.

The URSI Bolder meeting is scheduled for January 4-8, 1999 in Bolder, Colorado. Abstracts should be due some time around September 19, 1998. (Based on last year's deadline.) Ab-

stracts are one-page descriptions of the presentation.

For more information, please contact Dr. Dennis G. Sweeney, WA4LPR, Research Assistant Professor, Center for Wireless Telecommunications Department of Electrical Engineering, Virginia Tech, Blacksburg, VA 24061-0111; tel 540-231-2653, fax 540-231-3004; e-mail dsweeney@vt.edu, URL www.cwt.vt.edu

24th Eastern VHF/UHF Conference

The 24th Eastern VHF/UHF Conference will take place August 21-23, 1998 at the Harley Hotel in Enfield, Connecticut. The event is sponsored by the Eastern VHF/UHF Society and the North East Weak Signal Group (ARRL affiliated). The committee has worked diligently to provide an excellent, full program this year. They've improved our conference by carefully evaluating the critiques we received last year.

Speakers are coming from New England and as far south as Maryland. We'll have a technical laboratory that will operate during most of the conference. We encourage attendees to bring a project or test equipment to debug or tune with some of the newer test equipment. Watch the experts at work!

Gerry Rodski, K3MKZ, of SSB Electronic USA returns to manage the preamp noise-figure measurement workshop (50 MHz to 10 GHz). Joe

Reisert, W1JR, will again manage the antenna-gain-measurement range (222 MHz and higher, entrants receive a hardcopy plot of the results). This is a great opportunity to check out your preamps and antennas! The VHF-SHF Trivia Quiz will be hosted by Ernie Gray, W1MRQ.

We will also have a manufacturers' presentation session at the end, Amateur Radio manufacturers and distributors contact Bruce Wood, N2LIV, if interested. Bruce also needs nominations for the Tom Kirby award.

Registration includes entry for drawings and an ARRL published *Proceedings* book with the event schedule and many interesting articles. Registrants will receive their copies on arrival at the conference.

Call for Papers

The *Proceedings* Editor (Bruce Wood, see "Contacts" below) must receive all articles by July 1, 1998. Last years' *Proceedings* includes articles about operating, antennas, equipment design, interfacing and testing. This was made possible only by the generosity of VHFers who share their knowledge with other VHFers. Please contact Bruce if you'd like to contribute this year. Articles must be camera-ready with one-inch borders. Separate photos and text, and enclose a page layout.

Accommodations

Harley Hotel, 1 Bright Meadow Blvd

(off SR 5), Enfield, CT 06082; tel 800-321-2323, 860-741-2211. Call for reservations; mention the Eastern VHF-UHF Conference for \$59 overnight rates. Alternative dining is available nearby.

Contacts

Conference Chairman and *Proceedings* Editor: Bruce Wood, N2LIV, 3 Maple Glen Ln, Nesconset, NY 11767-1711; tel 516-225-9400 (w), 516-265-1015 (h); e-mail bdwood@erols.com

Swap and Sell Information: Mark Casey, N1LZC, 303 Main St, Hampden, MA 01036; tel 413-566-2445; e-mail map@map.com

Official Web page: <http://uhavax.hartford.edu/~newsvhf>; e-mail bdwood@erols.com

A Note to Conference Planners

QEX is prepared two months prior to the first month on its cover. (This July/August issue was prepared in May.) That means that press releases for events (or event deadlines, as in calls for papers) in August needed to reach the *QEX* Managing Editor in the first half of May, four months prior to the event. Send your press release (if by mail, include a 3 1/2-inch PC or Mac disk, with the file in *Word for Windows* or ASCII format) to Bob Schetgen, KU7G, *QEX* Managing Editor, 225 Main St, Newington, CT 06111; e-mail qex@arri.org. □□

Letters to the Editor

Discontinued Part for "The Belittler"

◇ I am sorry that the ZN416E IC is no longer manufactured by GEC-Plessey. After submitting the article,¹ I received the latest catalog from DC Electronics and found it was not listed by them anymore. Perhaps there are some elsewhere.

Perhaps the more-available ZN414s could be used and followed by micro- or mini-power amplifiers. The LM10 might be one prospect, as it is rated to work well at 1.1 V.

At any rate, I hope the article inspires more tinkerers like me. It's not difficult to build if a modular, building-block approach is used.—*William Latta, Jr, 7309 Greenlawn Rd, Louisville, KY 40222*

¹B. Latta, N4LH, "The Belittler," *QEX*, Nov 1997, pp 22-24.

QEX Improvements, and the Ionosphere

◇ I'm watching the exciting improvements in *QEX*. The articles are more varied, more interesting and instructive to me. Rudy, I still remember your great talks at Powercon, you really captured the minds of the audience. It is fun to see you doing it again through this publication.

I sure wish I knew more about propagation. I enjoyed the article by Eric Nichols.² However, I have never heard an actual theory of how "one way" propagation could occur until now. Nevertheless, I still don't see how the tilted ionosphere could do it. Maybe I'm missing something, but the way I see it, if I want to QSO single hop (to make it simple) to, say, Kansas, then if the ionosphere tilts so that it is lower near me, then the reflection point moves closer to me than halfway to Kansas. This results in the angle of incidence still being equal to the angle of reflection. Turning it around, if we assume that the angle of incidence is always equal to the angle of reflection, then the point of reflection has to move closer to me so that the beginning is at Reno and the end is at Kansas. Therefore, there is no difference in the critical angle with direction of propagation. And hence, no explanation for theory of "one way" skip.

Cheers, Bill Avery, K6GNX, 2900 Fantasy Ln, Sparks, NV 89436; e-mail bavery@telemetry.com

Dear Bill:

I'm glad to see someone's on their toes! Because of space limitations, I didn't give the full explanation of the ionosphere-tilt business. It's well worth a whole article in itself. What you say, Bill, is absolutely true—if we assume both parties are aiming at the same general location on the ionosphere. In real life, however, we generally aim our beams more or less horizontally. In this case, the party beneath the "lower end" of the ionosphere will be receiving the signal from a very high angle. (This is confirmed when, under these conditions, you can swing your Yagi 360° around the azimuth and see no change in received signal strength!) Going the other way (say from Alaska, where the ionosphere is low, to the Midwest, where the ionosphere is high) we can essentially "miss" the ionosphere altogether...we're basically shooting parallel to the thing.

Another assumption is that the angle of incidence is equal to the angle of reflection. This is not necessarily a valid assumption when you have a magnetic field parallel to the direction of radiation, as it is over Fairbanks! I was remiss in neglecting the magnetic, hence asymmetrical, aspects of the ionosphere at high latitudes. 'Nother subject in itself. I hope this helps some. Thanks for the feedback! Sincerely, Eric P. Nichols, KL7AJ, PO Box 56235, North Pole, AK 99701; e-mail enichols@gei.com

Additional note to all:

This tilted ionosphere also explains why many DXers "up here" have discovered NVIRs (near-vertical-incident radiators) to be very effective DX antennas. An NVIR is the last thing you want if you're trying to work DX from a "normal" place.—Eric Nichols, KL7AJ

Eric:

I am happy. Multiplying the angular response of the ionosphere with the H-plane radiation pattern of the Yagi would explain why there is less gain on one end than the other. Nevertheless, when combining the total of all factors, end to end, I still don't get any differ-

ence in path loss, A to B or B to A. I don't know anything about how the magnetic field affects the propagation of a plane wave parallel to a steady magnetic field. My uninformed guess would be nothing. Is there some interaction between the wave and the ions that is direction dependent? And then, if so, has different attenuation for two waves whose propagation vectors differ only in sign? If these are time-consuming questions, please just point me to some common text. I have several books from which I have difficulty extracting practical data, like Davies, Boithias and others. If there is one that you find instructive, please let me know. Thanks for your quick response too! Cheers, Bill Avery, K6GNX

No problem. I live for this! Also, I hope I made it sufficiently clear in my article that we don't come close to knowing it all. I just wanted to whet people's appetites, and let them know that not all is business as usual "up there"! However, one-way propagation is well beyond the stage of mere hypothesis. It's a reproducible phenomenon in plasma chambers, and there's plenty of evidence that we're seeing the same thing in nature. Eric Nichols, KL7AJ

²E. Nichols, KL7AJ, "How the Ionosphere REALLY Works." *QEX*, Mar/Apr 1998, pp 37-40.

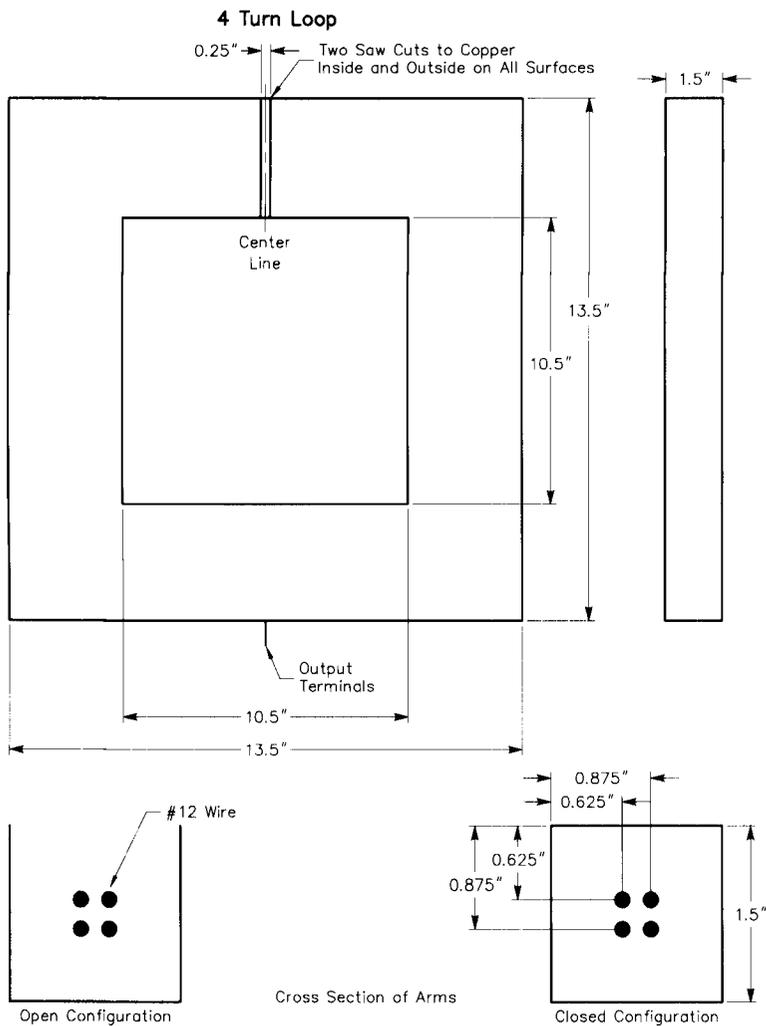
Q of Shielded Loop Antennas

◇ I am writing to provide some information I have not previously published regarding shielded receiving loop antennas.

The use of electrostatic shields to maintain balance in a receiving loop antenna is a common technique. Over the years I have seen various published designs, with a few that have shown the shield in the form of channel or trough (one side open), while others have shown a complete enclosure. I have never seen an explanation of the effects of the different configurations.

While building a loop for some personal experiments a few years ago I had a chance to make some measurements with laboratory equipment on the Q of the loop. The loop had two configurations as shown in Fig 1. The loop conductor was #12 THHN insulated electrical wire.

The measurement instrument was a Hewlett-Packard 4342A Q-meter. This instrument has a basic Q accuracy of approximately 7%. Please note that in order to resonate the loop below 2 MHz it was necessary to par-



Material: 0.062" Double Sided G10 Board with Both Copper Sides Electrically Connected.

allel the loop with a 150 pF mica cap. This cap would represent a Q of approximately 1200 at these frequencies. So, for points at 2 MHz and below in the figure, the values plotted are approximately 10% low. The inductance of the loop is 13 μ h at 2.52 MHz.

If you look at the figure you will see that the Q is in excess of 135 for the channel (open-sided) shield. If the loop is perfect (no proximity effect from multiple turns, etc) and we calculate unloaded Q, it will vary from 540 at 1.8 MHz to 1070 at 7 MHz. So the numbers we are seeing with the shield are reasonable.

When you close the shield around all four sides, the Q decreases to between 54% and 89% of the open (three-sided) values. This would have a measurable effect on the loop bandwidth, SNR and efficiency. This trade-off has to be weighed against the particular application.

For those using loops for VLF applications, I had a chance to measure Q at 100 kHz. In the open configuration the Q was 26; the closed configuration showed a Q of 9. (This was measured using a Hewlett-Packard 4274A LCR meter).

I hope this information will spark further experimentation in this area.—*Domenic M. Mallozzi, N1DM, 168 Speen St, Natick, MA 01760-2557; e-mail DMALLOZZI@AOL.COM*

Fig 1—Details of the small loop antenna described in N1DM's letter.

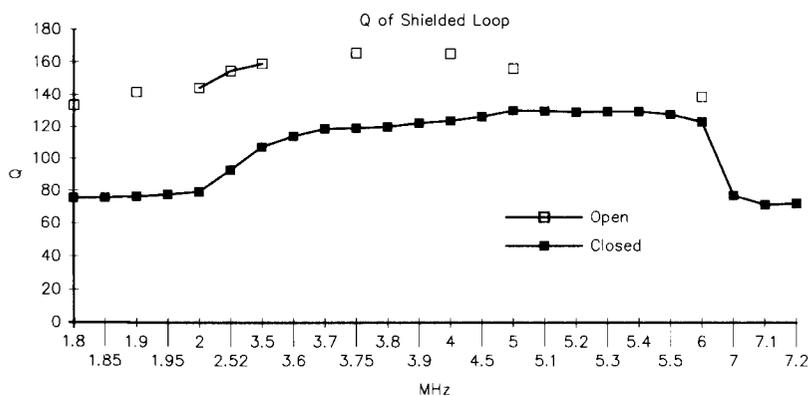


Fig 2—Measured Qs for open and closed shielding of a small loop antenna described in N1DM's letter.