KF6ZZ's "Souped-Up" ATX Switching Supply

QEX

INCLUDING: COMMUNICATIONS QUARTERLY

About the Cover

Revenge! When an ATX power supply killed his computer, KF6ZZ resurrected it as a 13.8-V, 20-A station supply. The story of Phil's adventures begins on p 36.

Features

Columns

Nov/Dec 2004 *QEX* Advertising Index

Empirical Outlook

On Journalistic Integrity and Other Observations

In late September of this year, the national news media made much ado about the reportage of a certain television broadcast network. The case in point dealt with whether the network had exercised due diligence in checking some documents they had received from an evidently unsolicited source. After they aired the documents and later admitted that they could not substantiate their authenticity, allegations of bias flew freely all around. With that much egg on their faces, the question might have been "Could I get some bacon and toast with that?"

Everyone has an opinion. If you ask for one, you are going to get it. To pretend that journalists do not have opinions is inane; to think they are always going to suppress them is naive. Yes, we are supposed to keep them out of our reporting, save in editorial columns, but we are seeing less of that self-restraint these days and the trend is not abating. In fact, all one need do is examine the responses from other television networks to that September scandal to see it. How ironic that seems.

Fortunately, in the scientific and engineering worlds, we have a neat procedure that allegedly keeps opinion on the sidelines in favor of what can be proved or disproved. It is funny, though. Sometimes we cannot agree on the best way of doing something, such as testing a transceiver. At other times, we each think we have the answer to some other issue, only to find out later that someone can disprove it.

Many years ago, a colleague declared that the square root of 2 was certainly an irrational number (cannot be written exactly because the digits go on forever) in base 10. Yet he claimed that in base square root of 2, the square root of 2 was a rational number because it could be written simply as 1. Then he saw Carl Sagan's *reductio ad absurdum* proof of the irrationality of that number. Find it in the back of this issue (p 62). It does not refer to number base anywhere.

Another colleague related how he met Carl at a planetarium. He argued against Carl's statement that the Big Dipper would look like a mirror image of its normal appearance if you traveled to the other side of it, at a straight-line distance equal to the average distance to the stars in it from the Earth side. The discussion was heated; but fortunately, the planetarium could be programmed to show the actual result. Carl lost.

And so it goes. Those examples allude to one reason we have *QEX*. Here you can put forth your theorems, proofs and disproofs, along with some good stuff that we know works—that readers can build. Complete parts lists and minute details are not always necessary, but we want to make sure readers can contact authors. While our staff does check for accuracy, authors are expected to defend their own assertions. We need your comments for our letters column as well as your articles. Keep those projects—and discussions—going!

In This Issue

John Champa, K8OCL, and John Stephensen, KD6OZH, describe their work with high-speed multimedia (HSMM) networking on the microwave bands using 802.11 and other equipment. Much of the work is associated with successes achieved by the ARRL HSMM Working Group. KD6OZH also contributes a separate piece on software-defined radio.

Karl-Otto Müller, DG1MFT, discusses coaxial traps for antennas. Unfortunately, we must hold the 2004 Index and Randy Evans, KJ6PO's PLL article for the next issue. Yet, for you synthesizer fans, Kjell Karlsen, LA2NI, contributes a piece on measuring phase noise in oscillators.

Robert LaFrance, N9NEO, brings a unique way of modulating a class-E transmitter in AM mode. Phil Eide, KF6ZZ, opens the mysteries of loop control and magnetics in switching power supplies as he tells us how to resurrect an ATX computer power supply as a main station (13.8 V, 20 A) supply.

On behalf of the staff of *QEX*, may your holiday season be merry and your outlook bright! *Doug Smith, KF6DX, kf6dx@arrl.org*. ☐☐

HSMM Radio Equipment

Readily available computer oriented Wi-Fi equipment can be used to form the basis for high speed data transport at 2.4 GHz and above. This article shows you how it's done and how it works.

By John Champa, K8OCL; and John B. Stephensen, KD6OZH;
With input from Dave Stubb, VA3BHF

Introduction

This is the first article to discuss what is known in Amateur Radio as High-Speed Multimedia (HSMM) radio in technical detail. HSMM Radio is a form of Amateur Packet Radio that starts at speeds of 56 kbps and goes up from there up to 5000 times faster than conventional packet radio. This capability enables multimedia, or simultaneous digital video, digital voice, data, and text. Initial HSMM Amateur Radio research has been based on readily available, inexpensive commercial gear designed for WiFi or wireless local area networking (WLAN). HSMM is not a specific mode—it is, instead, a direction or a driving force within Amateur Radio to develop high-speed digital networking capability under Part 97 regulations.

Military surplus radio equipment fueled Amateur Radio in the 1950s. Commercial FM radios and repeaters snowballed the popularity of VHF/UHF amateur repeaters in the 1960s and 70s. In the same way, current availability of commercial wireless LAN (WLAN) equipment is driving the direction and popularity of Amateur Radio use of spread spectrum in the early 2000s.

The Institute of Electrical and Electronics Engineers (IEEE) has provided the standards under which manufacturers have developed WLAN equipment for sale commercially and hams have adapted this equipment to outdoor use. The IEEE 802.11 series of standards defines a series of RF modems similarly to the way that the International Telecommunications Union (ITU) defined a series of telephone modems in the past. The term "WiFi" is short for wireless fidelity and indicates that the subject equipment has been tested to ensure that it fully complies with the applicable IEEE 802.11 standard.

Accordingly, the first part of this article describes existing 802.11 equipment for the 13-cm and 5-cm amateur bands. The second part of this article describes a proposed communication protocol for HSMM operation that will fit into the existing ARRL band plans from 219 to 2400 MHz. The initial implementation will make use of the DCP-1 hardware module described in an article by John Stephensen, KD6OZH.

Existing Products—High Speed Multimedia Radio

In early 2002 the ARRL Technology Task Force (TTF) established the High Speed Multimedia (HSMM) Working Group with John Champa, K8OCL, as its chairman. John moved quickly to identify two initial goals for the new working group to immediately begin the development of such high-speed digital Amateur Radio networks:

- Encourage the amateur adoption and modification of COTS IEEE 802.11 spread spectrum hardware and software for Part 97 uses.
- Encourage or develop other high-speed digital radio networking techniques, hardware, and applications.

These efforts were rapidly dubbed HSMM Radio. Although initially dependent on adaptation of COTS 802.11 gear to Part 97, the emphasis is on simultaneous voice, video, data, and text modes.

Applications

HSMM radio has some unique ham

2491 Itsell Rd
Howell, MI 48843-6458
k8ocl@arrl.net

3064 E Brown Ave
Fresno, CA 93703-1229
kd6ozh@verizon.net

radio networking applications and operational practices that differentiate it from normal WiFi hotspots at coffeehouses and airports as described in the popular press. HSMM radio techniques are often used for system RC (remote control) of Amateur Radio stations.

In this day of environmentally sensitive neighborhoods, one of the greatest challenges, particularly in high-density residential areas, is constructing ham radio antennas; particularly high tower-mounted HF beam antennas. Such amateur installations also represent a significant investment in time and resources. This burden could be easily shared among a small group of friendly hams, a radio club or a repeater group.

Implementing a link to a remote HF station via HSMM radio is easy to do. Most computers now come with built-in multimedia support. Most Amateur Radio transceivers are capable of PC control. Adding the radio networking is relatively simple. Most HSMM radio links use small 2.4 GHz antennas mounted outdoors or pointed through a window. These UHF antennas are relatively small and inconspicuous when compared to a full-size 3-element HF Yagi on a tall steel tower.

A ham does not have to have an antenna-unfriendly homeowners association or a specific deed restriction problem to put RC via HSMM radio to good use. This system RC concept could be extended to other types of Amateur Radio stations. For example, it could be used to link a ham's home to a shared high performance Amateur Radio DX station, EME station, or OSCAR satellite ground station for a special event or even on a regular basis.

Shared Internet Access

Sharing high-speed Internet access (Cable, DSL, etc.) with another ham is a popular application for HSMM radio. As long as it is not done for profit, it is entirely legal in the US under Part 97

rules. However be careful to read the terms of service supplied by your service provider. Many have restrictions against sharing your service with another party. If you violate the terms and conditions of your service agreement, the provider can (and will) disconnect your service. Pop-up ads, although a nuisance, are not illegal and can readily be controlled by the proper browser configuration. Just as on the Internet, it is possible to do such things as playing interactive games, complete with sound effects and full motion animation, with HSMM radio. This can be lots of fun for new and old hams alike, plus it can attract others in the "Internet Generation" to get interested in Amateur

Radio and perhaps become new radio club members. In the commercial world these activities are called "WLAN Parties". Such e-games are also an excellent method for testing HSMM radio link speed.

Emergency Communications

There are a number of significant reasons why HSMM radio is the wave of the future for many Emergency Communications Support (RACES, ARES, etc.) situations.
• The amount of digital radio traffic on 2.4 GHz is increasing and operating under low powered, unlicensed Part 15 limitations cannot overcome this noise.
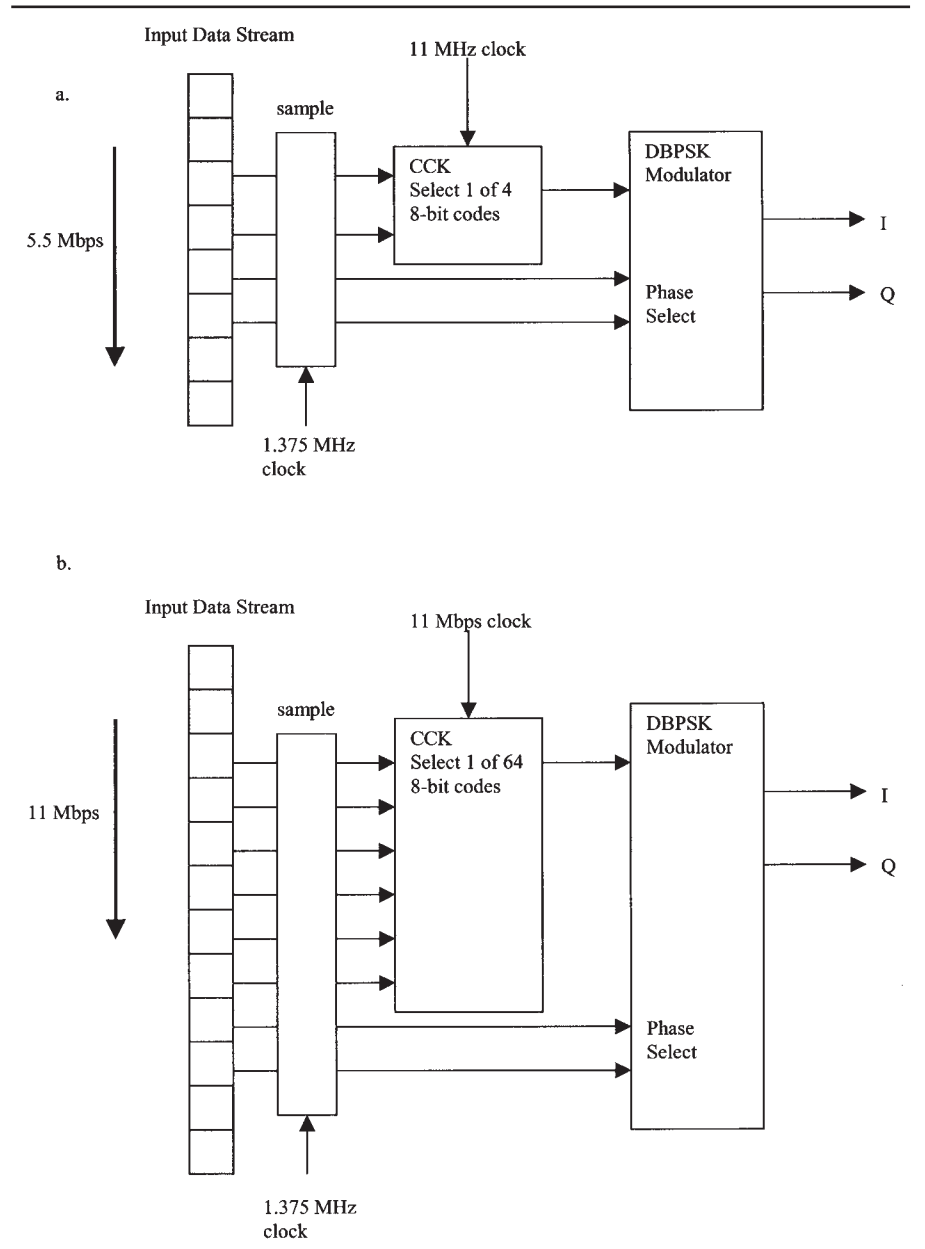
Fig 2—Encoding complementary code frequency sequence.

Fig 1—OFDM carrier interleaving.

- EmComm organizations increasingly need high-speed radio networks that can get out of the disaster area and into an area where ADSL, cable modem, satellite or other broadband Internet access is available.

With HSMM radio often all that would be needed to accomplish this in the field is a laptop computer with a headset, and perhaps an attached digital camera. The laptop must be equipped with a special wireless local area network card (PCMCIA) *with* an external antenna jack. In HSMM radio jargon such a card is simply called a RIC (radio interface card). Then connect the RIC to a short Yagi antenna (typically 18 inches of antenna boom length), or perhaps a small dish antenna mounted on a tripod weighted with a sandbag. Connection is established by pointing the antenna toward the HSMM repeater back at the EOC. More details are provided further into this paper.

Radio Relay for the 21st Century

There are a number of ways to extend the HSMM link. The most obvious means would appear to be to run higher power and place the antennas as high as possible, as is the case with VHF/UHF FM repeaters. In some densely populated urban areas of the country this approach with 802.11, at least in the 2.4 GHz band, may cause some interference with other users. Other means of getting greater distances using 802.11 on 2.4 GHz or other amateur bands should be considered. One approach is to use highly directive, high-gain antennas, or what is called the directive link approach. Another method used by some HSMM radio networks is what is called a low-profile radio network design. It depends on several low power sources and radio relays of various types. For example, two HSMM radio repeaters (known commercially as access points, or APs, about $100 devices) may be placed back-to-back in what is known as bridge mode. In this configuration they will simply act as an automatic radio relay for the high-speed data. It is possible to cover greater distances with relatively low power and yet still move lots of multimedia data.

A Basic HSMM Radio Station

How does one set up an HSMM radio base station? It is really very easy. HSMM radio amateurs will just need to go to any electronics outlet or office supply store and buy commercial off-the shelf (COTS) Wireless LAN gear, either IEEE 802.11b or IEEE 802.11g. They then connect external outdoor antennas. That is all there is to it.

There are some purchasing guidelines to follow. First, decide what interfaces you are going to need to connect to your computer. Equipment is available for all standard computer interfaces: Ethernet, USB, and PCMCIA. If you use a laptop in your station, get the PCMCIA card. Make certain it is the type with an external antenna connection. If you have a PC, get the Wireless LAN adapter type that plugs into either the USB port or the RJ45 Ethernet port. Make certain it is the type that has a removable rubber duck antenna or external antenna port! Finally, compare the RF performance of the devices you are contemplating. Unfortunately, there is little performance consistency across brands. Better cards can be purchased with up to 200-mW power output and –97 dBm receive sensitivity. Poor performers (while useful for covering a room in a home or office) have power outputs of less than 30 mW and receive sensitivities in the mid–80s dBm range. Buying the best performing card you can afford will assure the best performance. Also make sure the hardware selected for both ends of the link have equivalent performance. The overall link will be limited by the worst performing device. The included directions will explain how to accomplish the installation of these devices in your computer or network. These devices have two operating modes: *ad-hoc* and *Infrastructure*. Infrastructure mode is used to communicate with an access point (AP—more on this later). Ad-hoc mode allows these client cards to communicate together, associate and form an "ad-hoc" network (thus the name). Setting two or more cards into ad-hoc mode is the easiest way to get started experimenting with HSMM.

These client devices are the core of any HSMM radio station. They become a computer-operated HSMM 2.4 GHz radio transceiver and will probably cost about $20 to $80, depending on the performance of the hardware (better cards cost more). Start off your experimentation by teaming up with a nearby ham radio operator and setting each device in the ad-hoc mode and on a common channel. Channels 1 through 6 fall inside the Part 97 frequency allocation. However, channel 1 has output that falls within the AO-40 channel assignment, and channel 6 is commonly used by part 15 devices as the default channel. Using channels 2 through 5 limits the interference you may cause to other operators or have caused to you. Do your initial testing in the same room together. Then as you increase distances going toward your separate station locations, you can coordinate using a suitable local FM simplex frequency. Frequently hams will use 146.52 MHz or 446.00 MHz, the National FM Simplex Calling Frequencies for the 2-m and 70-cm bands, respectively, for voice coordination. More recently, HSMM radio operators have tended to use 1.2 GHz FM transceivers and handheld transceivers. The 1.2 GHz amateur band more closely mimics the propagation characteristics of the 2.4 GHz amateur band. The rule of thumb being, if you can not hear the other station on the 1.2 GHz FM radio, you probably will not be able to link up the HSMM radios.

HSMM Repeaters

What hams would call a repeater, and in the wired LAN world, computer buffs would call a hub, the WiFi industry refers to as a radio access point, or simply AP. This is a device that allows several Amateur Radio stations to share the radio network and all the devices and circuits connected to it.

An 802.11b AP will sell for about $80 and an 802.11g AP for about $100. The AP acts as a central collection point for digital radio traffic, and can be connected to a single computer or to another radio or wired network. Remember to select an AP with performance similar to the performance of the other 802.11 hardware you're using.

The AP identifies itself to its users by means of a station ID or *SSID*. Each AP is provided with an SSID, which is the station identification it constantly broadcasts. For ham purposes, the SSID can be set to your call sign, thus providing automatic, and constant station identification. To use an AP in a radio network the wireless computer users have to exit ad-hoc mode and enter what is called the infrastructure mode, in their operating software.

Infrastructure mode requires that you specify the radio network your computer station is intended to connect to, so set your computer station to recognize the SSID you assigned to the AP (yours or another ham's AP) to which you wish to connect.

Point-To-Point Links: The AP can also be used as one end of a radio point-to-point network. If you wanted to extend a radio network connection from one location to another, for example in order to remotely operate an HF station, you could use an AP at the network end and use it to communicate to a computer at the remote station location.

An AP allows for more network fea-

tures and improved information security than provided by ad-hoc mode. Most APs provide DHCP service, which is another way of saying they will automatically assign an Internet (IP) address to the wireless computers connected to the radio network. In addition, they can provide MAC address filtering which allows only known users to access the network.

Mobile Operating

When hams use the term mobile HSMM station what they are normally talking about is a wireless computer set-up in their vehicle to operate in a stationary portable fashion. Nobody is suggesting that you try to drive a vehicle and look at a computer screen at the same time. That could be very dangerous, and is illegal in some states. So unless you have somebody else to drive the vehicle keep your eyes on the road and not on the computer screen. Additionally, 802.11 was not designed for mobile use and is intolerant of the Doppler shift and signal fades associated with mobile operations.

• What sort of equipment is needed to operate an HSMM mobile station? Some type of portable computer, such as a laptop. Some hams use a PDA, notebook, or other small computing device. The operating system can be Microsoft Windows, Linux, or Mac OS, although Microsoft XP offers some new and innovative WLAN functionality. Some type of radio software hams would call an automatic monitor, and computer buffs would call a sniffer utility. The most common type being used by hams is Marius Milner's Network Stumbler for Windows, or "NetStumbler." All operating systems have monitoring programs available. Linux has Kismet; MAC OS has MacStumbler. Marius Milner has a version for the pocket PC called "MiniStumbler."

• A RIC (Radio Interface Card or PCMCIA WiFi computer adapter card with external antenna port) supported by the monitoring utility you are using. The most widely supported RIC is the Orinoco line. The Orinoco line is inexpensive and fairly sensitive.

• An external antenna attached to your RIC. This is often a magnetically mounted omnidirectional vertical antenna on the vehicle roof, but a small directional antenna pointed out a window or mounted on a small tripod are also frequently used. Be aware of the length and type of cable used to connect the antenna. The small diameter flexible coax often used can exhibit 6 dB of loss per 10 feet! If the antenna needs to be mounted more than 5 feet from the receiver, use LMR 400 or bet-

ter coax so as to minimize line losses. A pigtail or short strain relief cable will be needed to connect from the RIC antenna port to the N-series, RP/TNC or other type connector on the external antenna.

• A GPS receiver that provides NMEA 183 formatted data and computer interface cable will allow the monitoring utility to record where HSMM stations are located on a map just as in APRS. GPS capability is optional, but just as with APRS, it makes the monitored information much more useful since the station's location is provided.

While operating your HSMM mobile station, if you monitor an unlicensed Part 15 station (non-ham), some types of WiFi equipment will automatically associate or link to such stations, if they are not encrypted, and many are not (i.e., WEP is not enabled). Although Part 15 stations share the 2.4 GHz band on a non-interfering basis with hams, they are operating in another service. In another part of this section we will provide various steps you can take to prevent Part 15 stations from automatically linking with HSMM stations. So in like manner, except in the case of a communications emergency, we recommend that you do not use a Part 15 station's Internet connection for any ham purpose.

Area Surveys

Both licensed amateurs and unlicensed (Part 15) stations share the 2.4 GHz band. To be a good neighbor, find out what others are doing in your area before designing your community HSMM radio network. This is easy to do using IEEE 802.11 modulation. Unless it has been disabled, an active repeater (AP) is constantly sending out an identification beacon known as the SSID. In HSMM practice this is simply the ham station call sign (and perhaps the local radio club name) entered into the software configuration supplied with the CD that comes with the repeater. So every HSMM repeater is also a continuous beacon.

A local area survey using appropriate monitoring software, for example free NetStumbler software downloaded and running on your PC (**www.netstumbler.com/index.php**) is recommended prior to starting up any HSMM operations. Slew your station's directional antenna through 360°, or drive your HSMM mobile station (as described earlier) around your local area.

This HSMM area survey will identify and automatically log most other 802.11 station activity in your area.

There are many different ways to avoid interference with other users of the band when planning your HSMM operating. For example, moving your operating frequency 2-3 channels away from the other stations is often sufficient. Why several channels and not just one? Because the channels as named (1 through 11) are only 5 MHz wide each. The 802.11 carrier is 22 MHz wide, so a single 802.11 carrier occupies multiple numeric channels. Because of this, there is considerable overlap of occupied spectrum if you move only by a single 5 MHz channel. Why this situation exists is because the channel spacing was determined and allocated before the 802.11 standard was promulgated. Since other devices like video transponders, cordless phones, baby monitors, etc. also coexist in the band; it was not necessary or reasonable to change the channel allocations to support the unique behavior of 802.11. So, while there are 11 numeric channels in the Part 15 band, there are only three: 1, 6, and 11 that can support a non overlapped 802.11 carrier. Commercial users often recommend moving 5 channels away from the nearest AP to completely avoid interference. There are six channels within the amateur 2.4 GHz band, but there are problems for hams with two of them. Channel 1 centered on 2412 MHz overlaps into OSCAR satellite downlink frequencies. Channel 6 centered on 2437 MHz is by far the most common out-of-the-box default channel for the majority of WLAN equipment sold in the US, so that often is not the best choice. Subsequently, most HSMM radio groups end up using either channel 3 or channel 4, depending on their local situation. Again, an area survey is recommended before putting anything on the air.

Because of the wide sidebands generated by these inexpensive broad banded 802.11 devices, even moving 2 or 3 channels away from such activity may not be enough to totally avoid interference, especially if you are running what in HSMM is considered high power (typically 1800 mW RF output—more on that subject later). You may have to take other steps. For example, you may use a different polarization with your antenna system. Many HSMM stations use horizontal polarization because much of the non-ham 802.11 activity in their area is primarily vertically polarized.

Special Antenna Systems

There are a number of factors that determine the best antenna design for a specific HSMM radio application.

Most commonly, HSMM stations use horizontal instead of vertical polarization. Furthermore, most HSMM stations use highly directional antennas, instead of omnidirectional antennas. Directional antennas provide significantly more gain and thus better signal-to-noise ratios, which in the case of 802.11 modulation, means higher rate data throughput. Higher data throughput, in turn, translates into more multimedia radio capability.

Highly directional antennas also have many other advantages. Such antennas can allow two hams to "shoot over" or "shoot around" or even "shoot between" other wireless stations on the band. However, the nature of 802.11 modulation coupled with the various configurations of many COTS devices allows hams to economically experiment with many other fascinating antenna designs. Such unique antenna system designs can be used to simply help avoid interference, or to extend the range of HSMM links, or both.

Some APs and some RICs have space diversity capability built-into their design. However, it is not always operated in the same fashion, so check the literature or the Web site of your particular devices to be certain how the dual antenna ports are used. For example, many APs come equipped with two rubber ducky antennas and two antenna ports. One antenna port may be the primary and the other port the secondary input to the transceiver. Which signal input is used may depend on which antenna is providing the best S/N ratio at that specific instant. Experimentation using two outside high-gain antennas spaced 10 or more wavelengths apart (that is only about one meter on the 2.4 GHz band) may be very worthwhile in improving data throughput on long links. Such extended radio paths tend to experience more multi-path signal distortion. This multi-path effect is caused by multiple signal reflections off various objects in the path of the linking signal. The use of space diversity techniques may help reduce this effect and thus improve the data rate throughput on the link. Again, the higher the data rates the more multimedia radio techniques that can be used on that network.

Circular polarization can be considered as linear polarization with the angle of polarization rotating at the same frequency as the transmitted signal. The phase reversal in the electric field when the wave is reflected by a conductive surface causes the rotation sense to reverse. This is an improvement over linear polarization because, for example, right-hand circular polar-ization (RHCP) changes to left-hand circular polarization (LHCP) on the first reflection, which is usually the strongest reflection. An RHCP antenna at the receiver will then reject the strongest multi-path component with the reversed sense causing the unwanted multi-path component to be down around 20 dB. With linear polarization, although the electric field rotates 180° when the wave is reflected by a conductive surface, the resulting polarization is the same as the incident wave. This does nothing to help reject multi-path distortion at the receiver.

Circular polarization may be created by using helical antennas, patch feed-points on dish antennas, or other means and warrants further study by radio amateurs. Remember this is high-speed digital radio. To avoid symbol errors, circularly polarized antennas should be used at both ends of the link. Also, be certain that the antennas are of the same handedness, for example right hand circular polarization (RHCP). The ability of circular polarization to enhance propagation of long-path HSMM radio signals should not be overlooked.

A combination or hybrid antenna design combining both circularly polarized antennas and space diversity could yield some extraordinary signal propagation results. For example, it has been suggested that perhaps using RHCP for one antenna and LHCP for the other antenna, especially using spacing greater than 10 wavelengths, in such a system could provide a nearly "bullet-proof" design. Only actual field testing of such designs under different terrain features would reveal such potential.

High Power Operation

Hams often ask why operate 802.11 modes under licensed Part 97 regulations when we may also operate such modes under unlicensed Part 15 regulations, and without the content restrictions imposed on the Amateur Radio service? A major advantage of operating under Amateur Radio regulations is the feasibility of legally operating with more RF power output and larger, high-gain directive antennas. These added capabilities enable hams to increase the range of their operations. The enhanced signal-to-noise ratio provided by running high power would also allow better data packet throughput. This enhanced throughput, in turn, enables more multimedia experimentation and communication capability over such increased distances.

Increasing the effective radiated power (ERP) of an HSMM radio link also provides for more robust signal margins and consequently a more reliable link. These are important considerations in providing effective emergency communications services and accomplishing other important public service objectives in a band increasingly occupied by unlicensed stations and other noise sources.

It should be noted that the existing FCC Amateur Radio regulations covering *spread spectrum* (SS) at the time this is being written were implemented prior to 802.11 being available. The provision in the existing regulations calling for automatic power control (APC) for RF power outputs in excess of 1 W is not considered technologically feasible in the case of 802.11 modulation for various reasons. As a result the FCC has communicated to the ARRL that the APC provision of the existing SS regulations are therefore not applicable to 802.11 emissions under Part 97.

Using higher than normal output power in HSMM radio, in the shared 2.4 GHz band, is also something that should be done with considerable care, and only after careful analysis of link path conditions and the existing 802.11 activity in your area. Using the minimum power necessary for the communications has always been a good operating practice for hams as well as a regulatory requirement.

There are also other excellent and far less expensive alternatives to running higher power when using 802.11 modes. For examples, amateurs are also allowed to use higher gain directional antennas. Such antennas increase both the transmit and the receive effectiveness of the transceiver. Also, by placing equipment as close to the station antenna as possible, a common amateur OSCAR satellite and VHF/UHF DXing technique, the feed line loss is significantly reduced. This makes the HSMM station transceiver more sensitive to received signals, while also getting more of its "barefoot" transmitter power to the antenna. Only after an HSMM radio link analysis (see the link calculations portion of **www.arrl.org/hsmm/** or go to **logidac.com/gfk/80211link/pathAnalysis.html**) clearly indicates that additional RF output power is required to achieve the desired path distance, should more power output be considered.

At that point in the situation analysis, if higher power is required, what is needed is called a bi-directional amplifier (BDA). This is a super fast switching pre-amplifier / amplifier combination that is usually mounted at the end of the antenna pigtail near

the top of the tower or mast. As mentioned before, this is a two-way system, and the link will communicate only as far as the weakest link direction. A BDA needs to be used on both ends of the link in order to achieve greater communication distances. A system with a BDA on only one end may be heard by the far end station, but the BDA equipped station will probably not hear the weaker signal of the "barefoot" far end station. A reasonably priced 2.4 GHz 1800 mW output BDA is available from the FAB Corporation (**www.fab-corp.com**). It is specifically designed for amateur HSMM radio experimenters. Be certain to specify "HSMM" when placing your order. Also, to help prevent unauthorized use by unlicensed Part 15 stations, the FAB Corp may request a copy of your amateur license to accompany the order, and they will only ship the BDA to your licensee address as recorded in the FCC database.

This additional power output of 1800 mW should be sufficient for nearly all amateur operations. Even those supporting EmComm, which may require more robust signal margins than normally needed by amateurs, seldom will require more power output than this level. If still greater range is needed, there are other less expensive ways to achieve such ranges (see the section HSMM Radio Relays).

When using a BDA and operating at higher than normal power levels on the channels 2 through 5 recommended for Amateur Radio use (these channels are arbitrary channels intended for Part 15 operation and are not required for Amateur Radio use, but they are hard-wired into the gear so we are stuck with them). You should also be aware of the sidebands produced by 802.11 modulation. These sidebands are in addition to the normal 22 MHz wide spread spectrum signal. Accordingly, if your HSMM radio station is next door to an OSCAR ground station or other licensed user of the band, you may need to take extra steps in order to avoid interfering with them. The use of a tuned output filter may be appropriate in order to avoid causing QRM. Even when operating on the recommended channels in the 2-5 range, whenever you use higher than normal power, some of your now amplified sidebands may go outside the amateur band, which stops at 2450 MHz. So from a practical point of view, whenever the use of a BDA is required to achieve a specific link objective, it is a good operating practice to install a tuned filter on the BDA output. Such filters are not expensive and they are readily available from several commer-

cial sources. It should also be noted that most BDAs currently being marketed, while suitable for 802.11b modulation, are often not suitable for the newer, higher speed 802.11g modulation.

There is one further point to consider. Depending on what other 802.11 operating may be taking place in your area, it may be a good practice to only run higher power when using directional or sectional antennas. Such antennas allow hams to operate "over and around" other licensed amateur stations and unlicensed Part 15 activity in your area which you may not wish to disrupt (a local school WLAN, WISP, etc). Again, before running high power, it is recommended that an area survey be conducted using a mobile HSMM rig as described earlier to determine what other 802.11 activity is in your area and what channels are in use.

Information Security

An HSMM radio station could be considered a form of software defined radio. Your computer running the appropriate software combined with the RIC makes a single unit which is now your station HSMM transceiver. However, unlike other radios, your HSMM radio is now a networked radio device. It could be connected directly to other computers and to other radio networks, and even to the Internet. So each HSMM radio (PC + RIC + software) needs to be protected. There are at least two basic steps that should be taken for secure use of all HSMM radios:

The PC should be provided with an anti-virus program**.** This anti-virus must be regularly updated to remain effective. Such programs may have come with the PC when it was purchased. If that is not the case, reasonably priced anti-virus programs are readily available from a number of sources.

Secondly, it is important to use a firewall software program on your HSMM radio. It is recommended that the firewall be configured to allow no outgoing traffic unless it is coming from a known program, and to restrict all incoming traffic without specific authorization. Commercial personal computer firewall products are available from Symantec, Zone Labs and MCA Network Associates. Check this URL for a list of freeware firewalls for your personal computer: **www.webattack. com/freeware/security/ fwfirewall.shtml** and this one for a list of shareware firewalls for your personal computer: **www.webattack. com/Shareware/security/ swfirewall.shtml**.

Once a group of HSMM stations has

set up and configured a repeater (AP) into a radio local area network (RLAN) then addition steps may need to be taken to restrict access to the repeater. Only Part 97 stations should be allowed to associate with the HSMM repeater. Remember, in the case of 802.11 modulation, the 2.4 GHz band is shared with Part 15 unlicensed 802.11 stations. How do you keep these unlicensed stations from automatically associating (auto-associate) with your licensed ham radio HSMM network?

Many times the steps taken to avoid interference with other stations also limit those other stations' capability to auto-associate with the HSMM repeater, and improve the security of the HSMM station. For example, operating with a directional antenna oriented toward the desired coverage area rather than using an omnidirectional antenna, etc.

The most effective method to keep unlicensed Part 15 stations off the HSMM repeater is to simply enable the Wired Equivalent Protection (WEP) already built into the 802.11 equipment. The WEP encrypts or scrambles the digital code on the HSMM repeater based on the instruction or "key" given to the software. Such encryption makes it impossible for unlicensed stations not using the specified code to accidentally auto-associate with the HSMM repeater.

The primary purpose of this WEP implementation in the specific case of HSMM operating is to restrict access to the ham network by requiring all stations to authenticate themselves. Ham stations do this by using the WEP implementation with the appropriate ham key. Hams are permitted by FCC regulations to encrypt their transmission in specific instances; however, ironically at the time of this writing, this is not one of them. Accordingly, for hams to use WEP for authentication and not for encryption, the key used to implement the WEP must be published. The key must be published in a manner accessible by most of the Amateur Radio community. This fulfills the traditional ham radio role as a self-policing service. The current published ham radio WEP key is available at the home page of the ARRL Technology Task Force High Speed Multimedia Working Group: **www.arrl.org/hsmm/**.

Before implementing WEP on your HSMM repeater be certain that you have checked the Web site (**www.arrl. org/hsmm/**) to ensure that you are using the current published WEP key. The key may need to be changed occasionally.

The HSMM Working Group is cur-

rently investigating the feasibility of obtaining a waiver or station temporary authorization (STA) for selected Amateur Radio HSMM experimental stations. The purpose of the waiver would be to allow us to experiment with various wireless content security measures such as *virtual private networking* (VPN). Our research would be restricted to frequencies above 50 MHz and apply only to domestic amateur digital computer-to-computer networking experiments.

Commercial Part 15 Equipment

The IEEE standards for WLAN equipment have evolved from low speeds to high speeds, increasing the spectrum efficiency with each new version. IEEE 802.11 standardized *frequency-hopping spread spectrum* (FHSS) and *direct-sequence spread spectrum* (DSSS) for the 2.4 GHz ISM band to operate at data rates of 1 and 2 Mbps. Next came the release of 802.11b which provided the additional data rates of 5.5 and 11 Mbps but only for DSSS. The purpose of using FHSS and DSSS modulation techniques is to avoid inter-symbol interference (ISI) due to multipath propagation. In FHSS the receiver is on the next frequency when the delayed version of the last symbol arrives on the previous frequency. In DSSS the delayed version no longer matches the spreading code.

This was followed by 802.11g which provided standardization using *Orthogonal Frequency Division Multiplexing* (OFDM) for data rates of 6, 9, 12, 18, 24, 36, 48 and 54 Mbps as well as backward compatibility with 802.11b. As of this writing the most recent release of the standard is 802.11a. This release addresses the use of OFDM in the 5 GHz ISM and UNII bands. It provides the same data rates as 802.11g. The currently unreleased 802.11n standard promises data rates in excess of 108 Mbps.

Of course, none of these increases in capacity come for free. With each increase in capacity comes the need for more complex modulation to support it. As Claude Shannon theorized in 1948, increasing the bandwidth of a fixed size channel leads to the need for more power in order to discern the intelligence from the channel noise. In other words, increasing modulation complexity reduces receiver sensitivity. For example, an 802.11b link operating at 1 MBPS uses BPSK and has a receive sensitivity of around –94 dBm. For an 802.11g link operating at 54 Mbps the modulation is 64QAM, and the receive sensitivity drops to –68 dBm because of the addi-

tional signal to noise ratio required to retrieve the information from 64 possible modulation points rather than the 2 points associated with BPSK.

Note that the power increase is nonlinear as doubling the number of states per transmitted symbol increases the number of bits transmitted by an ever-decreasing amount.

Frequency Hopping Spread Spectrum

FHSS radios, as specified in 802.11, hop among 75 of 79 possible non-overlapping frequencies in the 2.4 GHz band. A complete hop sequence occurs approximately every 400 ms with a hop time of 224 µs. Since these are Part 15 devices the radios are limited to a maximum peak output power of 1 W and a maximum bandwidth of 1 MHz (at –20 dB) at any given hop frequency. The rules allow using a smaller number of hop frequencies at wider bandwidths (and lower power: 125 mW) but most manufacturers have opted not to develop equipment using these options. Consequently, off-the-shelf equipment with this wider bandwidth capability is not readily available to the amateur.

The hopping sequences are well defined by 802.11. There are three sets of 26 such sequences (known as channels) consisting of 75 frequencies each. The ordering of the frequencies is designed as a pseudo-random sequence hopping at least 6 MHz higher or lower than the current carrier frequency such that no two channels are on the same frequency at the same time. Channel assignment can be coordinated among multiple collocated networks so that there is minimal interference among radios operating in the same band.

The FHSS radio can operate at data rates of 1 and 2 Mbps. The binary data stream modulates the carrier frequency using frequency shift keying. At 1 Mbps the carrier frequency is modulated using 2-Level Gaussian Frequency Shift Keying (2GFSK) with a shift of +/-100 kHz. The data rate can be doubled to 2 Mbps by using 4GFSK modulation with shifts of +/-75 kHz and +/-225 kHz.

Direct Sequence Spread Spectrum

DSSS uses digital modulation to accomplish signal spreading. That is, a well-known pseudo-random digital pattern of ones and zeros is used to modulate the data at a very high rate. In the simplest case of DSSS, defined in 802.11, an 11-bit pattern known as a Barker sequence (or Barker code) is used to modulate every bit in the input data stream. The Barker sequence is 10110111000. Specifically, a "zero" data bit is modulated with the Barker sequence resulting in an output sequence of 10110111000. Likewise, a "one" data bit becomes 01001000111 after modulation (the inverted Barker code). These output patterns are known as "chipping" streams; each bit of the stream is known as a "chip". It can be seen that a 1 Mbps input data stream becomes an 11 Mbps output data stream.

The DSSS radio, like the FHSS radio, can operate at data rates of 1 and 2 Mbps. The chipping stream is used to phase modulate the carrier via phase shift keying. Differential Binary Phase Shift Keying (DBPSK) is used to achieve 1 Mbps and Differential Quadrature Phase Shift Keying (DQPSK) is used to achieve 2 Mbps.

Table 1
Bit encoding as a function of data rate

Data Rate, Mbps	CCK encoded bits	DQPSK encoded bits
5.5	2	2
11	6	2

Table 2
Modulation methods and coding rates

Data Rate, Mbps	Modulation	Coding Rate, (R)
6	BPSK	$^1/_2$
9	BPSK	$^3/_4$
12	QPSK	$^1/_2$
18	QPSK	$^3/_4$
16	QAM	$^1/_2$
36	16QAM	$^3/_4$
48	64QAM	$^2/_3$
54	64QAM	$^3/_4$

The higher data rates specified in 802.11b are achieved by using a different pseudo-random code known as a Complimentary Sequence. Recall the 11 bit Barker code can encode one data bit. The 8 bit Complimentary Sequence can encode 2 bits of data for the 5.5 Mbps data rate or 6 bits of data for the 11 Mbps data rate. This is known as Complimentary Code Keying (CCK). Both of these higher data rates use DQPSK for carrier modulation. DQPSK can encode 2 data bits per transition. Table 1 shows how 4 bits of the data stream are encoded to produce a 5.5 Mbps data rate and 8 bits are encoded to produce an 11 Mbps data rate. There are 64 different combinations of the 8 bit Complimentary Sequence that have the mathematical properties that allow easy demodulation and interference rejection. At 5.5 Mbps only four of the combinations are used. At 11 Mbps all 64 combinations are used. See Fig 2.

As an example, for an input data rate of 5.5 Mbps, four bits of data are sampled at the rate of 1.375 million samples per second. Two input bits are used to select 1 of 4 eight-bit CCK sequences. These 8 bits are clocked out at a rate of 11 Mbps. The two remaining input bits are used to select the phase at which the 8 bits are transmitted.

Orthogonal Frequency Division Modulation

OFDM transmits data simultaneously on multiple carriers. 802.11g and 802.11a specify 20 MHz wide channels with 52 carriers spaced every 312.5 kHz. Of the 52 carriers, four are non-data pilot carriers that carry a known bit pattern to synchronize demodulation. The remaining 48 carriers are modulated at 250 kbaud. The state of all 48 data carriers is known as a symbol. Thus, at any given instant in time 48 bits, or more, of data are being transmitted.

The term "orthogonal" is derived from the fact that these carriers are positioned such that they do not interfere with one another. The center frequency of one carrier's signal falls within the nulls of the signals on either side of it. Figure 1 illustrates how the carriers are interleaved to prevent intercarrier interference. OFDM avoids ISI by making the symbol period much longer than the multi-path delay. A gap is then placed between each symbol to occupy the time consumed by multipath reflections. The gap is 0.8 microseconds in 802.11a & g.

OFDM radios can be used to transmit data rates of 6, 9, 12, 18, 24, 36, 48 and 54 Mbps as specified by both 802.11a and 802.11g. In order to transmit at faster and faster data rates in the same 20 MHz channel different modulation techniques are employed: BPSK, QPSK, 16QAM and 64QAM. In addition, some of the bits transmitted are used for error correction so the raw data rates could be reduced by up to half of what they would be without error correction. For instance, assuming BPSK (1 bit per carrier) and assuming ½ the bits are used for error correction (known as the coding rate, R); the resulting data rate would be 6 Mbps.

48 carriers × 1 bit per carrier × 1/2 R = 24 bits (effective)

24 bits × 250 kilo transitions per second = 6 Mbps

Table 2 shows a complete list of the modulation methods and coding rates employed by 802.11 OFDM. The higher data rates will require better signal strength to maintain error free reception due to using few error correction bits and more complex modulation methods.

Frequencies for HSMM

Up to this point all the discussion has been regarding HSMM radio operations on the 2.4 GHz amateur band. However, 802.11 modulation can be used on any amateur band above 902 MHz, so we can research each of these options.

AM ATV on the 902-928 and 1240-1300 MHz bands is very susceptible to interference (–50 dBc can be seen) so it is would probably be difficult to find a good spot for 802.11 operation in major cities on either of these bands. The 902 MHz band is just 26 MHz wide so 802.11 modulation would occupy almost the entire band. The 1240 MHz band has ATV channels every 12 MHz so it is impossible to avoid interference. Luckily, ATV at 2400 MHz and above is 16 MHz wide FM and is much more immune to interference.

The 3.3-3.5 GHz band offers some real possibilities for 802.11, or the newer 802.16 standard. Activity is centered in three bands at 3.37-3.39 (FM ATV), 3.4-3.41 GHz (European weak-signal modes and U.S. satellite sub-band), 3.456-3.458 (U.S. weak-signal modes) and 3.47-3.49 GHz (FM ATV). There is lots of unused spectrum and frequency transverters could be used to get to this band from 2.4 GHz. Development in Europe of 802.16 with 108 Mbps data throughput may make 3.5 GHz gear available for amateur experimentation in the U.S. In the U.S. the 802.16 development is above the amateur 3.5 GHz band, while the European frequencies used are within the US amateur band. Hams are investi-

gating the feasibility of using such gear when it becomes available in the US for providing a RMAN or radio metropolitan area networks. The RMAN would be used to link the individual HSMM repeaters (AP) or RLANs together in order to provide countywide or regional HSMM coverage, depending on the ham radio population density.

The 5.65-5.925 GHz band is also being investigated. The COTS 802.11a modulation gear has OFDM channels that operate in this Amateur Radio band. The 802.11a modulation could be used in a ham RLAN operating much as 802.11g is in the 2.4 GHz band. It is also being considered by some HSMM groups as a means of providing RMAN links. This band is also being considered by AMSAT for what is known as a C-N-C transponder. This would be an HSMM transponder onboard a Phase 3 high-altitude OSCAR with uplink and downlink pass-band in the satellite sub-bands at 5.65-5.67 and 5.83-5.85 GHz. Some other form of modulation other than 802.11 would likely have to be used because of timing issues and other factors, but the concept is at least being seriously discussed.

The 10 GHz band could also host HSMM activity via transverters. Activity is currently limited to the 10.22-10.28 (WBFM), 10.368-10.37 (weak-signal) and 10.39-10.41 (FM ATV) GHz segments and 10.45-10.5 GHz is reserved for amateur satellites. The bottom 200 MHz of the band would be ideal for HSMM, perhaps in conjunction with ICOM DSTAR systems.

Other RMAN link alternatives are also being tested by hams. One of these is the use of wired networks for linking and the technique known as virtual private networks (VPN). This is similar to the method currently used to provide worldwide FM voice repeater links via the Internet, except that it would be broadband and multimedia. Mark Williams, AB8LN, of the HSMM Working Group is leading a team to test the use of various VPN technologies for linking HSMM repeaters. Mark recently made a presentation on this research at the 2004 Dayton Hamvention during the Technology Task Force (TTF) Forum. This forum is an annual event conducted by the ARRL TTF Chairman, Howard "Howie" Huntington, K9KM. The forum also involves our brothers in the two other TTF working groups: The Software Defined Radio (SDR) Working Group and the Digital Voice (DV) Working Group.

There are also commercial products being developed such as the ICOM D-STAR system which could readily be integrated into a RMAN infrastruc-

Table 3. OFDM Broadcasting Standards

Standard	Digital Radio Mondiale (DRM)			IBOC AM	Digital Audio Broadcasting (DAB)				IBOC FM
Frequency	150 kHz-30 MHz			0.5-1.7MHz	47-3000MHz	47 - 1492 MHz	47-230MHz		88-108MHz
Signaling Rate	37.5 Baud	37.5-60 Baud	37.5-60 Baud	172.3 Baud	6422 Baud	3211 Baud	1605 Baud	803 Baud	344.5 Baud
Carrier Spacing	42/47 Hz	42-107 Hz	42-107 Hz	181.7 Hz	8 kHz	4 kHz	2 kHz	1 kHz	363.4 Hz
Inter-Symbol Gap	2.7/5.3 ms	2.7-7.3 ms	2.7-7.3 ms	300 μs	31 μs	62 μs	123 μs	246 μs	151 μs
Path Differential	250/500 mi.	250-700 mi.	250-700 mi.	28 mi.	3 mi.	6 mi.	12 mi.	23 mi.	14 mi.
FFT Sample Rate	6 kSPS	12 kSPS	24 kSPS	≈24 kSPS	2.05 MSPS	2.05 MSPS	2.05 MSPS	2.05 MSPS	≈750 kSPS
Carriers	113/103	229-89	461-179	105	192	384	768	1536	1093
Bandwidth	4.9 kHz	9.8 kHz	19.4 kHz	18.9 kHz	1.54 MHz	1.54 MHz	1.54 MHz	1.54 MHz	397 kHz
Modulation	DQPSK	DQPSK	DQPSK	64QAM	DQPSK	DQPSK	DQPSK	DQPSK	B/QPSK

ture, especially with their ATM approach on 10 GHz.

HF frequencies are not being ignored. Neil Sablatzky, K8IT, is leading a team of ham investigators on the HF bands. Digital voice at 2400 BPS has been used on HF so it is possible that fast data rates will become available to efficiently handle e-mail type traffic on the HF bands while still occupying appropriate bandwidth. This would be helpful in an emergency by providing an e-mail outlet for HSMM RMAN e-mail traffic.

John Stephensen, KD6OZH, is leading the HSMM Working Group RMAN-UHF team. John has been investigating HSMM on the UHF amateur bands. 802.11 provides effective communication over short distances with omnidirectional antennas and can be extended to longer distances with highly directional antennas. However, it does not fit within most of the UHF bands and is not efficient at covering wide areas with omnidirectional antennas and may be limited in its HSMM applications to the 2.4 GHz bands and above.

The 802.11 OFDM standards do suggest a solution. OFDM with a slower symbol rate, narrower bandwidth and larger inter-symbol guard band would allow the use of omnidirectional antennas over long paths. In addition, the reduced path loss at lower frequencies will allow coverage of wide areas.

OFDM in Broadcasting

The latest additions to the 802.11 series standardize RF modems using orthogonal frequency division multiplexing or OFDM. This technology provides a bandwidth-efficient method of transmitting digital signals over long distances. OFDM is not only being applied to wireless computer networking but also to broadcasting on frequencies from 150 kHz to 3 GHz. The basic technology is the same, but certain parameters are modified to fit the characteristics of the radio channel. Table 3 shows several OFDM broadcasting standards.

Digital Radio Mondiale (DRM) is a standard for the long wave, medium wave and short wave broadcasting bands[1]. It is designed to tolerate the long multi-path delays caused by ionospheric propagation and therefore uses very low symbol rates. The inter-symbol guard band is 2.7-7.3 ms long and can therefore tolerate multi-path delays due to path length differences of up to 700 miles.

The Digital Audio Broadcasting (DAB) standard is a European standard for terrestrial broadcasting on VHF and UHF bands[2]. Multi-path delays are 1/10th to 1/100th of those in short wave radio and 4 modes of operation are specified. 1 kHz carrier spacing is used for VHF broadcasting and 2 or 4 kHz spacing is used for broadcasting up to 1500 MHz. 8 kHz spacing may be used up to 3 GHz. This system is designed for bands that have no existing analog broadcast stations.

IBOC (in-band on channel) AM and IBOC FM are systems marketed by Ibiquity that have been accepted by the FCC for use in the U.S. medium wave and VHF broadcast bands. Multi-path tolerance is similar to DAB but the bandwidth is narrower. This system is optimized to fit in bands with existing analog broadcasting.

OFDM in Amateur Radio

In the amateur bands, OFDM is being used in the HF bands for digital voice transmission. A 36-carrier OFDM modem was developed by G4GUO and is being marketed by AOR. It has characteristics similar to DRM but uses less than half the bandwidth. When used with an AMBE vocoder with rate 2/3 convolutional coding it has a throughput of 2400 BPS.

For high-speed data transmission, amateurs have been using IEEE 802.11 compliant products in the 13-cm band.[3,4] Most activity has been with DSSS equipment at a data rate of 11 MBPS. However OFDM modems are now available which operate in the 13-cm and 9-cm amateur bands with data rates up to 54 MBPS. This series of standards were designed for short-range (a few thousand feet) use and therefore tolerate a multi-path differential of only 400 feet. However, they can and are being used over longer distances by using directional antennas to suppress multi-path propagation. Transverters can be constructed to convert 13 cm 802.11 equipment to the 9, 3 and 1.2-cm bands.

There is a gap between the capabilities of the 802.11 and G4GUO modems that needs to be filled. The VHF and UHF amateur bands are ideal for multi-point local communication, as path losses are low with omnidirectional antennas. An OFDM RF modem with high data rates and longer multi-path delay tolerance would allow operation in urban areas over both line-of-sight (LOS) and non-line-of-sight (NLOS) paths.

In the effort to research various alternatives to linking Amateur Radio 802.11-based repeaters together, the HSMM Working Group has established several Radio Metropolitan Area Network (RMAN) project teams

[1]Notes appear on page 17.

Table 4. OFDM Modems for the Amateur Radio Service

Standard	G4GUO	RMAN-UHF Draft Standard						IEEE 802.11
Signaling Rate	50 baud	937.5 baud			7500 baud			250 kbaud
Carrier Spacing	62.5 Hz	1171.875 Hz			9375 Hz			312.5 kHz
IS Gap	4 ms	213.3 µs			26.7 µs			0.8 µs
Multi-path	750 miles	40 miles			5 miles			800 feet
Frequency (MHz)	1.8–30	219–450 (50-450*)	420–450 (222-450*)	420–450 (222-450*)	902–2400 (222-2400*)	902–2400 (222-2400*)	902–2400 (222-2400*)	2,400-10,500
FFT Sample Rate	4 ksps	150	300	1200	1200	2400	9600	20,000
Pilot Carriers	0	1	1	1	1	1	1	4
Data Carriers	36	64	128	512	64	160	512	48
Chan. Spacing	4	100	200	750	750	2000	6000	25,000
Bandwidth (kHz)	2.3	78	153	603	620	1520	4820	17,000
Low Rate (ksps)	2.4	120	240	960	960	2400	7680	6000
Modulation	DQPSK	D8PSK	D8PSK	D8PSK	D8PSK	D8PSK	D8PSK	BPSK
FEC Rate	2/3	2/3	2/3	2/3	2/3	2/3	2/3	1/2
High Rate (ksps)	-	240	480	1920	1920	4800	15,360	54,000
Modulation	-	64QAM	64QAM	64QAM	64QAM	64QAM	64QAM	64QAM
FEC Code Rate	-	2/3	2/3	2/3	2/3	2/3	2/3	3/4

*Under ARRL proposed regulations based on signal bandwidth.

lead by experts in their respective fields. These teams currently consist of the RMAN-VPN Project lead by Mark Williams, AB8LN; the RMAN-DSTAR and AMSAT C&C Project lead by John Champa, K8OCL; the RMAN-802.16 and Mesh Networking Project lead by Gerry Creager, N5JXS; the RMAN-UHF Project lead by John Stephensen, KD6OZH and the HSMM-HF Project (for e-mail) lead by Neil Sablatzky, K8IT.

John Stephensen, KD6OZH, as RMAN-UHF Project Leader, has been researching various alternatives for digital metropolitan area networks in the UHF amateur bands. The IEEE has developed the 802.16 WMAN standard, but this is for operation above 2 GHz and the bandwidth required is more than can be made available in the UHF amateur bands. Consequently, we need to develop an amateur standard for data transmission in the UHF bands. The HSMM group is tasked with developing links at data rates above 56 kbps and operation at 384 kbps or above is desirable as this supports full-motion compressed video.

OFDM Modem Physical Layer

The UHF amateur bands fall into 2 categories. The FCC limits the bandwidth available for data transmission in the 219-220 MHz and 420-450 MHz bands to 100 kHz, but there is no limit for the bands at 902 MHz and above. There is a practical limit of 6 MHz in the 902-928 MHz, 1240-1300 MHz, 2300-2305 MHz and 2390-2450 MHz bands because they are shared with existing users of analog modes. The goal is to develop a series of modems that operate above 56 KBPS and span the range of bandwidths available within the ARRL band plans. Table 4 shows the characteristics of OFDM modems being used in the amateur bands today

and the proposed standard described in this document. The bandwidths for the modems were chosen to fit off-the-shelf SAW filters used in GSM, CDMA and cable TV equipment.

The modem design was strongly influenced by the DAB standard as it operates in the same frequency range and supports both mobile and fixed users. Radio propagation in an urban area is characterized by strong multi-path propagation. Propagation measurements indicate that multi-path delay ranges from 0.4 to 10 µs typically and up to 90 µf worst case for LOS and NLOS paths in an urban environment. The modems defined in the middle seven columns of Table 4 use either 7500-Baud symbol rates with 9.375 kHz carrier spacing or 937.5-baud symbol rates with 1172-Hz carrier spacing. This results in an active symbol time of Ts = 106.7 or 853.3 µs with a guard band of Tg = 26.7 or 213.3 µs between adjacent symbols. The guard band is filled with a copy of the last ¼ of the OFDM symbol as shown in Figure 3.

The lowest speed modem is designed to fit in a 100-kHz channel and uses 64 data carriers plus a pilot carrier as shown in Figure 4. The pilot carrier is transmitted at 3 dB above the level of the data carriers and is placed in the center of the channel. Half of the data carriers are placed on each side of the pilot carrier and enumerated 1 through 64 from the lowest frequency to the highest frequency. The major lobes of the data carriers occupy 78 kHz. Extending beyond that limit on either side are the minor lobes of these carriers. Since the first minor lobe is at –13 dBc and the amplitude decreases at only 6 dBc/octave, additional filtering is required. A FIR filter with flat group delay must be used to attenuate minor lobes to –34 dBc at ±50 kHz.

Eight-phase differential phase shift keying (8DPSK) is used for the low data rate to allow mobile operation. As the station moves, the absolute phase varies as the strength and delay of multi-path rays vary so a fixed phase reference cannot be used. In-
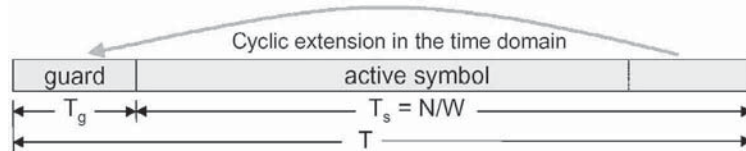
Fig 3—Guard Band
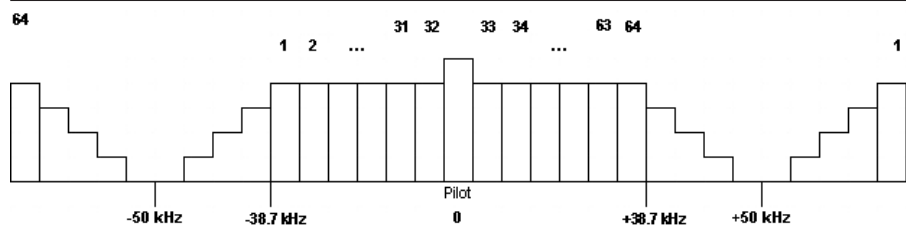
Fig 4—Format for 125 kHz Channel Spacing

stead the difference between the phase of the current symbol and the previous symbol is used to determine the value transmitted. Three coded bits are transmitted per symbol per carrier as shown in Table 5.

Trellis-Coded Modulation

Since the transmission channel will corrupt the transmitted data due to noise and fading, a forward error correcting (FEC) code must be used to provide adequate performance at reasonable signal to noise ratios (SNR). A plain block convolutional code could be used for FEC but it is much more efficient to use an error correcting code that is integrated with the modulation method. This is called trellis-coded modulation or TCM[5] and we will use a rate 2/3 trellis-code where 2 data bits (x^1 and x^2) are converted into a 3-bit code word (y^0, y^1 and y^2).

In TCM the signal constellation is partitioned into subsets as shown in Figure 5. Each partitioning increases the distance between constellation points. A convolutional coding of xn1, as shown in Figure 6, generates y0 and y1, which are used to select between the subsets C0, C1, C2, and C3 at the receiver. The data bit x2 = y2 then selects the final value.

The coding decreases the error rate because it increases the sequential distance between codes. The coded bits, y_{0-2}, may assume only certain sequences of values that are dependent on the state of the convolutional encoder, S_{0-1}, and the input, x_1, as shown in Figure 6.

Viterbi Decoding

The receiver can use this information to find the allowed sequence of symbols that is closest in Euclidean distance to the received sequence of symbols and determine the state of the convolutional coder in the transmitter. This is usually done using the Viterbi algorithm[6] with a soft-decision input.

The input is not a 3-bit vector, but a set of eight probabilities that the transmitted signal matches each of the eight signal constellation points shown in Fig 7. The algorithm associates a distance metric with each possible sequence of received signals and selects the maximum-likelihood path. The selection is made by tracing back the possible signal sequences and detecting segments that are common, as shown in Figure 8 (ML segment).

After determining the transmitter's state, the uncoded bit, x_2/y_2, is decoded by selecting the closest point in the remaining subset of the signal constellation. This is equivalent to decoding a BPSK data stream so the ultimate error rate for trellis-coded 8PSK is the same as for BPSK data. This results in a considerable coding gain, as the number of data bits actually received is double what BPSK would deliver. Figure 9 shows the gain provided by trellis-coded 8PSK compared to QPSK.

Since the outer Reed-Solomon code works on symbols, the event error rate curve is the one that is relevant.

Higher Data Rate

When the SNR is high, and the transmission path characteristics are stable, transmitting 4-bits per carrier results in a rate twice the basic data rate. This can be done in fixed stations where the phase of the received signal does not change rapidly. 64QAM modulation is used with a rectangular constellation as shown in Figure 10. The in-phase (I) and quadrature (Q) components of the signal are orthogonal and are treated separately in the encoding and decoding process. Two data bits are converted to three coded bits as was done for 8DPSK. One set of bits modulates the I carrier and another modulates the Q carrier as shown in Table 6. The maximum I and Q amplitude is limited to 0.7 so that the vector sum will not exceed 1.0.
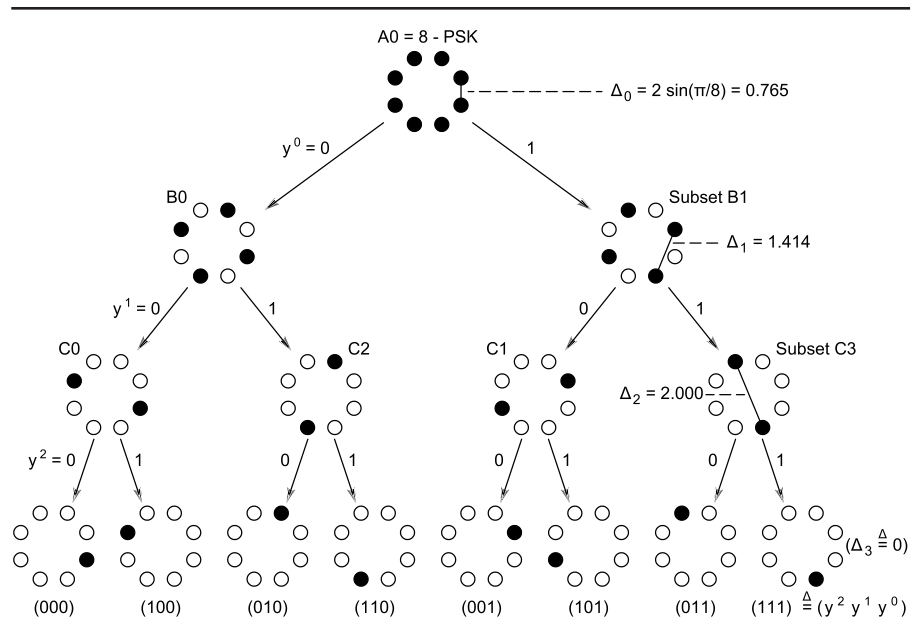
Fig 5—Signal Constellation Partitioning

Table 5
D8PSK Encoding

Tribit $y_2\, y_1\, y_0$	Carrier Phase Shift
0 0 0	0°
0 0 1	45°
0 1 0	90°
0 1 1	135°
1 0 0	180°
1 0 1	225°
1 1 0	270°
1 1 1	315°
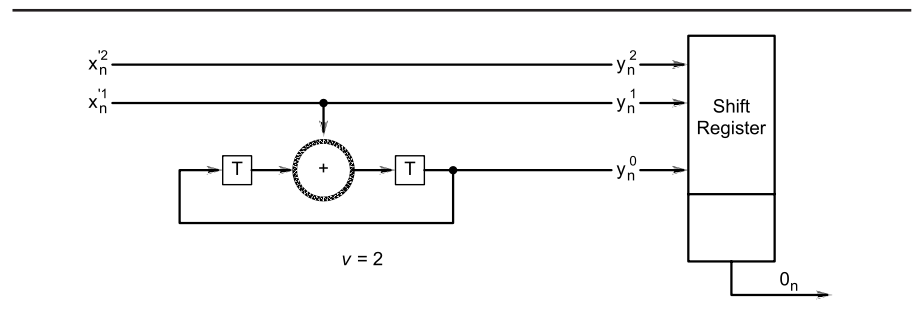
Fig 6—Convolutional Coding

Symbol Synchronization

To properly demodulate the 8DPSK or 64QAM encoded information, the receiver must maintain proper symbol synchronization as shown in Figure 11. This causes the inter-symbol interference (ISI) to be ignored when the fast Fourier transform (FFT) is calculated to demodulate the individual carriers.

Two special symbols are used for synchronization. Since the phase of the incoming carriers is in flux during the first part of the OFDM symbol period, the total amplitude of all carriers is used to delimit the symbol period. A special maximum amplitude symbol, called the reference (REF) symbol is defined, where the absolute phase of each carrier is set according to the formula:

$$q = 3.6315 \ k^2$$

where k is the carrier index by frequency[7]. This pattern minimizes amplitude distortion due to selective fading. In addition, the crest factor of the REF symbol waveform is less than 5 dB so that the reference symbol can be transmitted at 3 dB above normal power levels to improve amplitude and phase estimation.

The second special symbol is the null (NUL) symbol, which consists of the pilot carrier and no data carriers. The sequence REF-NUL-REF is present at the beginning of each data frame. The receiver normally uses a moving average filter with a time constant of one symbol period to detect end of the NUL symbol, as shown in Figures 12 and 13.

The REF-NUL-REF sequence is inserted into the transmitted data stream every 125 symbols. This reduces the required symbol clock accuracy to ±100 PPM. The REF symbol after the NUL is then used as an amplitude and phase reference for demodulating the following symbols. The format of a complete physical layer protocol data unit (PHY-PDU) is shown in Figure 14.

Protocol Control Information

The PHY-PDU begins with 8 PIL symbols. The PIL symbol is a full amplitude pilot carrier with no data carriers.

The high amplitude single carrier (PIL symbols) allows the receiver to acquire carrier frequency lock more easily. This is followed by the REF-NUL-REF sequence and a 1 to 125-symbol data block. If more than 125 symbols are to be transmitted, all blocks but the last have 125 data symbols. The PHY-PDU ends with a PIL symbol.

MAC Sublayer Error Correction

The physical layer provides forward error correction to compensate for errors due to Gaussian noise. However, the radio communications channel is also subject to fading and/or impulse noise that may introduce errors in bursts. The error correction provided in the physical layer may be overwhelmed and bytes containing errors may be delivered to the MAC sublayer. Reed-Solomon codes are particularly good at correcting bursts of errors and one is used in the MAC sublayer to alleviate this problem. This type of code operates on symbols m-bits wide, taking a block of k symbols and adding parity bits to form a block of n symbols where n = 2m–1. The encoded block consists of the k original symbols plus n–k parity symbols, as shown in Figure 15, and is capable of correcting t = (n-k)/2 symbol errors.

The code used is an RS (255,223) code that operates on 8-bit symbols and will correct errors in up to 16 symbols per block with an overhead of 12.6%. When 223 data bytes are available for transmission, an encoded block of 255 bytes is generated. The parity symbols are created by dividing a polynomial represented by the k data symbols by the RS generator polynomial. The symbols in the remainder are the parity symbols. If the end of the PHY-SDU is reached and the number of data bytes to be transmitted is less than 223, a shortened code block is generated.

At the receiver, the process of detecting an error is fairly simple, but correcting errors requires a lot of computation, as shown in Figure 16. As
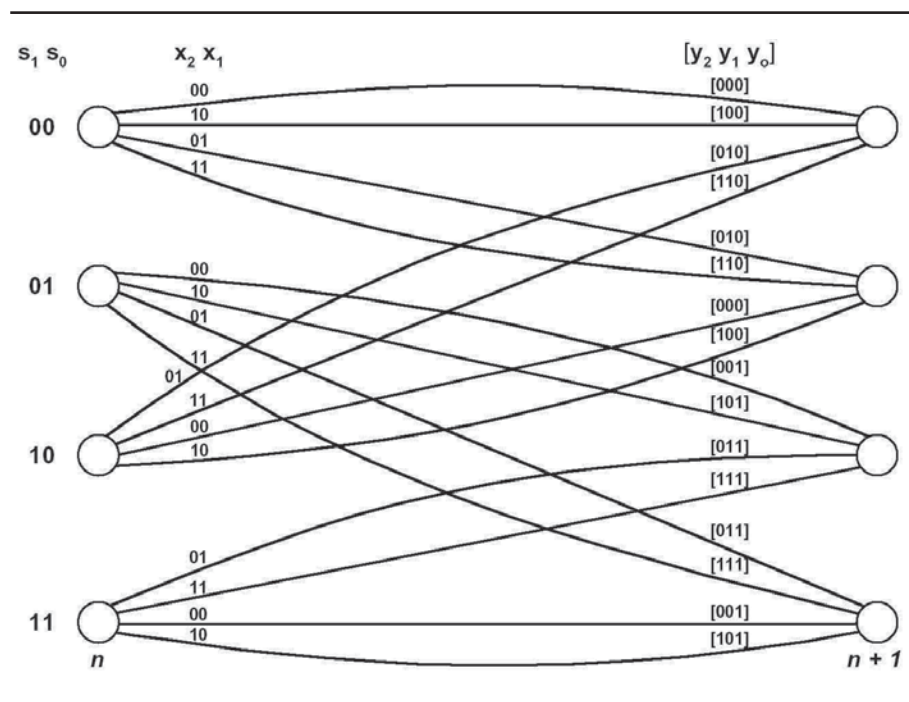
Fig 7—Allowed State Transitions
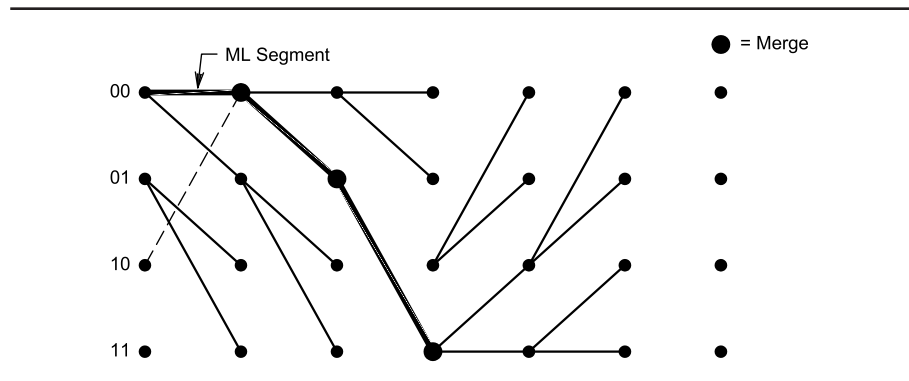
Fig 8—Viterbi Decoding

the data and parity symbols are received, they are divided by the generator polynomial and the remainder, called the syndrome, is zero if there are no errors. If the syndrome is not zero, the syndrome is processed to locate the errors. There are 2t simultaneous equations to be solved with the unknowns being the t locations of errors. The solution can be found in two steps. First the equations are solved using an iterative algorithm, such as Euclid's algorithm or the Berlekamp-Massey algorithm. This generates an error polynomial whose roots are the locations of the corrupted symbols. The error polynomial is then evaluated to find its roots using an exhaustive search, such as the Chien search. The error values are then calculated using the syndromes and the error polynomial roots. This is usually done using the Forney algorithm, which performs a matrix inversion. The error values are then exclusive-ORed with the received data to correct the errors.

MAC Service

This MAC entity is designed to provide a standard IEEE 802.3-style MAC service to the user. The user sends and receives service data units up to 1,536 bytes in length. The sender is identified by the source MAC address and the receiver is identified by the destination MAC address. Addresses are 48 bits in length and may be either individual or group addresses. Individual addresses consist of a six-character amateur-radio-service call sign plus a one-character extension. Group addresses are arbitrary 7-character strings. Characters are encoded in 6-bit ASCII.

Since the physical layer transmits up to 128 bytes per OFDM symbol, each station will accumulate multiple MAC protocol data units (MPDUs) for transmission in one PHY-SDU whenever possible. Each MPDU consists of MAC protocol control information (MPCI) and, optionally, a MAC service data unit (MSDU). Figure 17 shows an example with five MPDUs with three containing MSDUs. The maximum PHY-SDU length is 5,184 bytes.

MPDU Formats

There are three types of MPDUs defined. A Data MPDU transports a complete MSDU. It consists of 21-bytes of MPCI containing the address and type fields followed by a variable-length user-data field as shown in Figure 18. The MPCI fields are the intermediate address (IA), destination address (DA), source address (SA) and length (L). DA, SA and L are obtained from the MAC service user while IA is generated by the MAC entity. IA is the next destination address while DA is
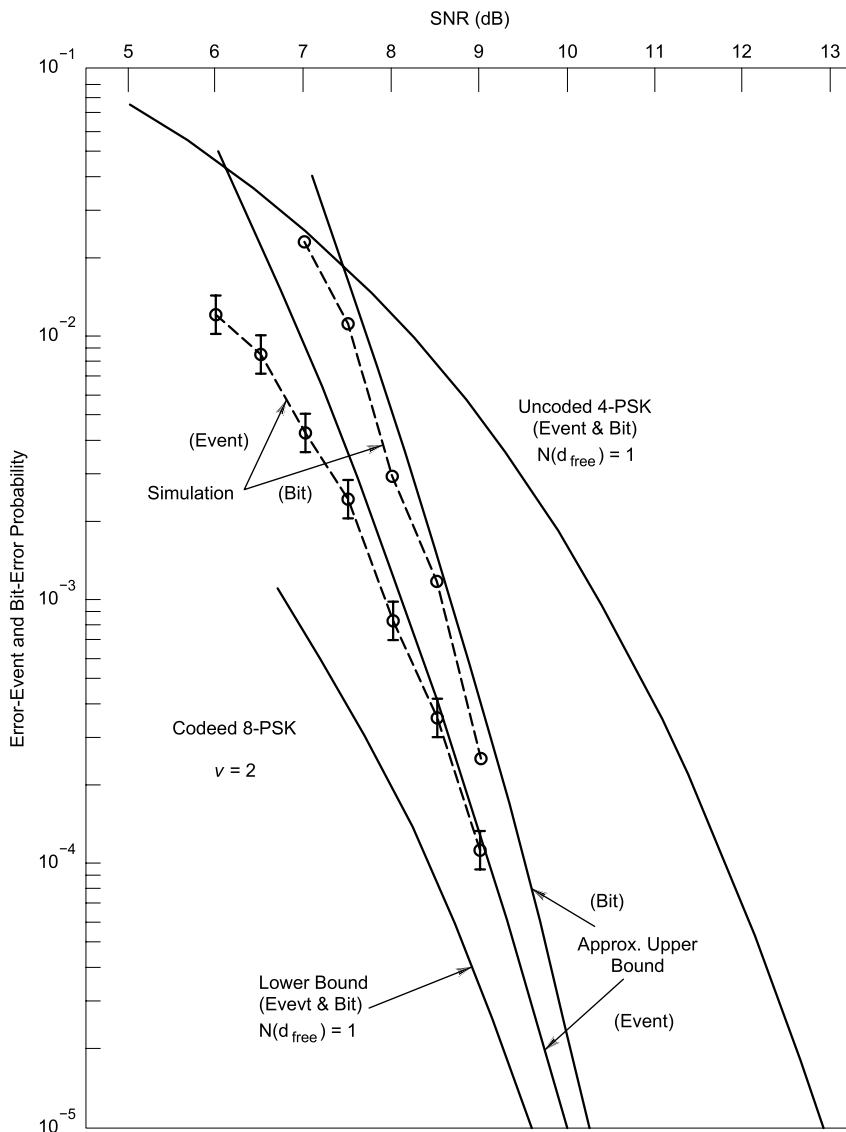
Fig 9—TC8PSK vs. QPSK

Fig 10—64QAM Signal Constellation with Coded I and Q Tribits shown in Octal

Table 6
64QAM Encoding

Tribit $y_2\,y_1\,y_0$	I & Q Amplitude
0 0 0	−0.7
0 0 1	−0.5
0 1 0	−0.3
0 1 1	−0.1
1 0 0	+0.1
1 0 1	+0.3
1 1 0	+0.5
1 1 1	+0.7

the ultimate destination address. A secondary station may be set IA to the primary station address to cause it to forward data to another secondary station that it cannot reach directly.

Access is controlled by a primary station that polls multiple secondary stations for traffic. It transmits a token that confers the right to transmit to the addressed secondary station. The secondary station then transmits any accumulated traffic and gives the token back to the primary station. The Token MPDU contains the address of the primary station (PA) and the next secondary station (SA) to transmit as shown in Figure 19.

Stations exchange received signal strength indication (RSSI) reports to determine what other stations are reachable in a network. The primary station periodically transmits an RSSI MPDU to each secondary station and the secondary stations respond by broadcasting RSSI MPDUs. Each station builds up a database of neighbor stations with the strength of its signal at each neighbor and the transmission capabilities of each neighbor. This can be used to select the modulation method and number of carriers to use when transmitting to adjacent stations.

The RSSI MPDU reports the received signal strength (SNR) for one or more transmitting stations (TA) at a particular receiving station (RA) as shown in Figure 20. The TA and RSSI fields are repeated N times. The C and M fields indicate the transmitter capabilities at the reporting station. C is the maximum number of data carriers supported divided by 4. M is the maximum number of bits transmitted per data carrier.

OFDM Modem Hardware

The OFDM modem described here is being implemented on the DCP-1 digital signal processing board. The DCP-1 uses an Xilinx Spartan-3 FPGA to implement the physical layer of the modem and an Oki Semiconductor ML67Q5000 MCU to implement the MAC sublayer. The received signal is digitized at the IF frequency by a 14-bit ADC at 19.2 Msps and transmitter I and Q baseband signals are generated by a dual 14-bit DAC at 9.6 Msps. The DCP-1 connects to its host via RS-232 or RS-485 up to 230.4 kbps or via USB at 12 Mbps or 480 Mbps. This hardware will be made available to amateurs by one of the authors, KD6OZH.

To allow the widest possible software compatibility, the modem will emulate an IEEE 802.11 LAN controller. For point-to-point operation, individual LAN addresses would be Amateur Radio service call signs. The net control stations call sign or an alphanumeric multicast address could be used for multi-point operation. Station identification is automatic, as the transmitting station's call sign is always the source address. Since the DCP-1 has an RS-232 interface, an interface that would emulate either a dial-up modem or a TNC in KISS mode is also being considered. This may be useful at low data rates for compatibility with older computers or legacy software is also being considered.

Conclusion

We expect to have OFDM modems using 8DPSK modulation operational and being tested in the field this year.
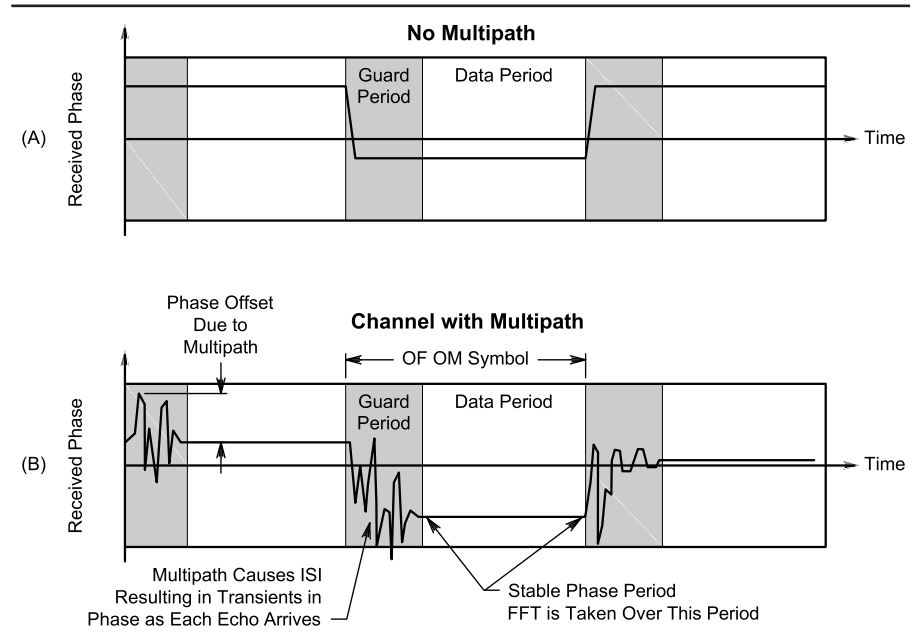
Fig 11—ISI Rejection using Guard Band
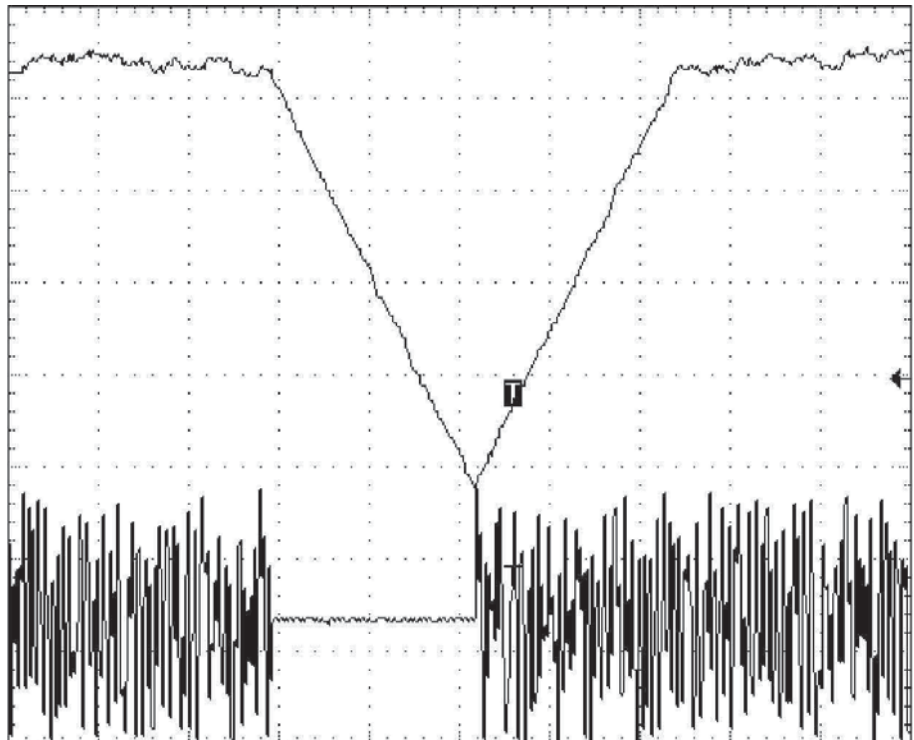
Fig 12—Synchronization with Normal REF Symbol

The source code will be made public and interested amateurs can modify or improve it as desired. The full specification is available from the HSMM working group. OFDM will allow a number of new applications in the UHF amateur bands and should provide much higher reliability than older technologies such as AFSK and FSK without FEC.

A 937.5-Baud OFDM modem can support data rates up to 240 Kbps in the 219-220 MHz band and up to 1.92 Mbps in the 420-450 MHz band. Operation in the 219 and 420 MHz bands is interesting because 100-W and higher-power amplifiers are available and the path loss is low. This allows the exploitation of paths with high losses. The low baud rate allows a wide (213.3-ms) guard band to suppress ISI for operation over NLOS paths that would be impossible with other types of equipment. The transmission characteristics should be ideal for mobile operation. Pure data transmission would be limited to 240 kbps by the current FCC regulations, but compressed video could be transmitted at the higher data rates as it would be classified as digital ATV. Since the modems will be implemented in software, the occupied bandwidth and data rate can be changed to accommodate FCC regulations and the current propagation conditions.

OFDM could enable mobile ATV, which is impossible with analogue techniques. You could turn on a video camera and the operator on the "talk-in" frequency at a hamfest could see where you are and provide better directions. This capability could prove invaluable in emergency situations. At the home station, applications such as NetMeeting could be used to allow radio club meetings to take place over the air. Presentations could be relayed over the video link. This would be very useful for members with physical disabilities. These data rates are also useful for linking clusters of computers where the main source of traffic is email or facsimile or where the Internet connection is a temporary dial-up link. The 120 or 240 kbps data rates should be adequate for forwarding health and welfare traffic at Red Cross shelters. This type of link also provides some level of security, as the general public won't have compatible equipment.

The 7500-baud modems can provide higher data rates and may be used on the 33, 23 and 13-cm bands. They are ideal for 802.11 AP linking as they can accommodate T1 data rates for applications that require database access or video. The ability to operate in the 23-cm band eliminates the interference present on 13-cm and allows longer paths to be accommodated when compared to 802.11 modems. This is ideal for linking 802.11 APs serving clusters of computers at remote sites.

There are probably many other applications that haven't been thought of yet—Amateur Radio operators will always find new uses for state-of-the-art technology.

Notes

[1]"Digital Radio Mondiale (DRM); System Specification," ETSI TS 101 980 v1.1.1 (2001-09).
[2]"Radio Broadcasting Systems; Digital Audio Broadcasting (DAB) to mobile, portable and fixed receivers," ETS 300 401, second edition, May 1997.
[3]"Supplement to IEEE Standard for informa-
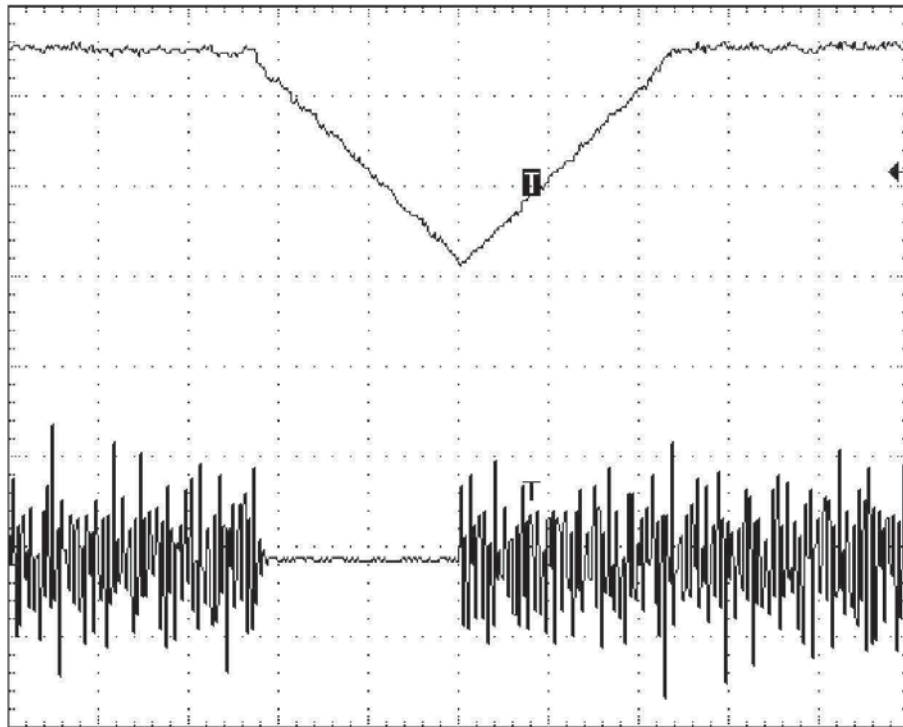
Fig 13—Synchronization with Selective Fading
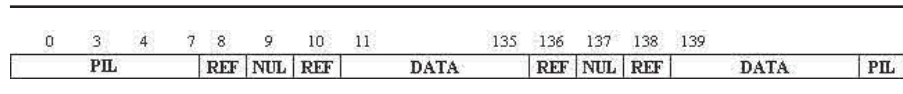
Fig 14—PHY-PDU Format (2 data blocks)
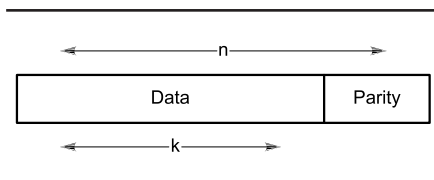
Fig 15—RS(n,k) Encoded Block
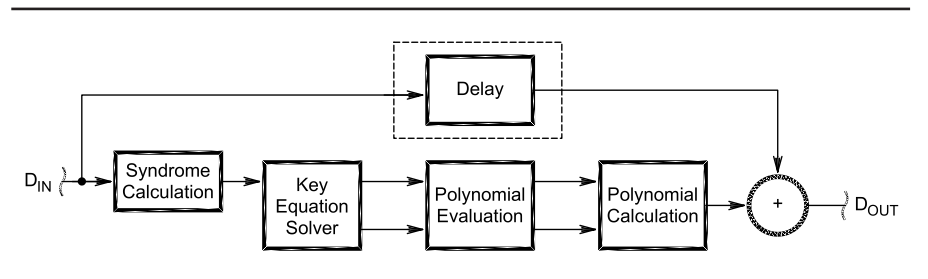
Fig 16—RS Error Correction Process

MPDU 1	MPDU 2	MPDU3	MPDU 4	MPDU 5
$MPCI_1$ | $MSDU_1$	$MPCI_2$	$MPCI_3$	$MPCI_4$ | $MSDU_2$	$MPCI_5$ | $MSDU_3$
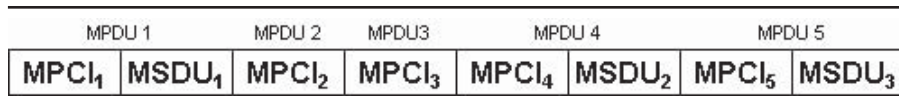
Fig 17—PHY-SDU with Multiple MPDUs

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	21+
0		IA				DA				SA								L			MSDU	

Fig 18—Data MPDU

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
2		RA						C	M	N		TA_1				SN
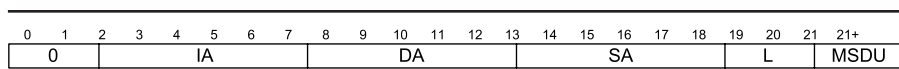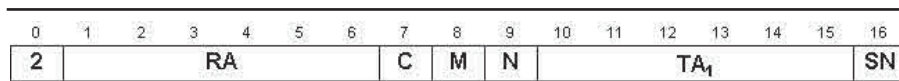
Fig 20—RSSI MPDU with one signal report

0	1	2	3	4	5	6	7	8	9	10	11	12	13
1		PA					1			SA			

Fig 19—Token MPDU

tion technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications—High-Speed Physical Layer in the 5 GHz Band," ISO/IEC 8802-11:1999/Amd 1:2000 (E).

[4]"IEEE Standard for information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment 4 - Further Higher Rate Data Extension in the 2.4 GHz Band," IEEE Std 802.11g-2003.

[5]"Trellis Coded Modulation with Redundant Signal Sets," Gottfried Ungerboeck, *IEEE Communications Magazine*, February 1987 – Vol. 25, No. 2.

[6]"Self-Correcting Codes Conquer Noise, Part 1: Viterbi CODECs," Syed Sabzad Shah, *EDN*, February 15, 2001.

[7]"Adaptive Techniques for Multi-user OFDM," Eric Phillip Lawrey, James Cook University, December 2001.

Bibliography

M. Burger, AH7R, and J. Champa, K8OCL, "HSMM in a Briefcase," *CQ VHF*, Fall 2003, p. 32.

J. Champa, K8OCL, and R. Olexa,KA3JIJ, "How To Get Into HSMM," *CQ VHF*, Fall 2003, pp. 30-36.

T. Clark, W3IWI, "C-C RIDER, A New Concept for Amateur Satellites", *Proceedings of the AMSAT-NA 21st Space Symposium*, November 2003, Toronto, Ontario, Canada (this book is available from the ARRL Book Store).

G. Cooper and C. McGillem, *Modern Communications and Spread Spectrum*, New York, McGraw-Hill, 1986.

J. Duntemann, K7JPD, *Jeff Duntemann's Wi-Fi Guide*, Paraglyph Press, 2003.

H. Feinstein, WB3KDU, "Spread Spectrum: Frequency Hopping, Direct Sequence and You," *QST*, June 1986, pp. 42-43.

R. Flickenger, *Building Wireless Community Networks* – 2nd Edition, O'Reilly, 2003. (This book is available from the ARRL Book Store).

R. Flickenger, *Wireless Hacks*, O'Reilly, 2003.

S. Ford, WB8IMY, "VoIP and Amateur Radio," QST, February 2003, pp. 44-47.

S. Ford, WB8IMY, *ARRL's HF Digital Handbook*, American Radio Relay League, 2001.

M. Gast, *802.11 Wireless Networks, The Definitive Guide*, O'Reilly, 2002. (This book is available from the ARRL Book Store).

J. Geier, *Wireless LANs, Implementing High Performance IEEE 802.11 Networks*, Second Edition, SAMS, 2002.

G. Held, *Building a Wireless Office*, Auerbach, 2003.

C. Holmes, *Coherent Spread Spectrum Systems*, New York, NY. Wiley Interscience, 1982

K. Husain, and T. Parker, Ph.D., et al. *Linux Unleashed*, SAMS, 1995.

A. Kesteloot, N4ICK, "Practical Spread Spectrum: An Experimental Transmitted-Reference Data Modem," *QEX*, July 1989, pp. 8-13.

T. McDermott, "Wireless Digital Communications: Design and Theory", *TAPR*, 1996.

K. Mraz, N5KM, "High Speed Multimedia Radio," *QST*, April 2003, pp. 28-34.

R. Olexa, KA3JIJ, "Wi-Fi for Hams Part 1: Part 97 or Part 15," *CQ*, June 2003, pp. 32-36.

R. Olexa, KA3JIJ, "Wi-Fi for Hams Part 2: Building a Wi-Fi Network," CQ, July 2003, pp. 34-38.

R. Olexa, KA3JIJ, *Implementing 802.11, 802.16, and 802.20 Wireless Networks Planning, Troubleshooting, and Operations*, Elsevier, 2004

B. Patil, et. al. IP in Wireless Networks, Prentice Hall, 2003.

B. Potter, and B. Fleck, 802.11 Security, O'Reilly, 2003.

H. Price, NK6K, "Spread Spectrum: It's Not Just for Breakfast Any More!" (*Digital Communications*), *QEX*, June 1995, pp. 22-27.

T. Rappaport, N9NB, "Spread Spectrum and Digital Communication Techniques: A Primer," *Ham Radio*, December 1985, pp. 13-16, 19-22, 24-26, 28.

J. Reinhardt, AA6JR, "Digital Hamming: A Need for Standards," *CQ*, January 2003, pp. 50-51.

P. Rinaldo, W4RI, and J. Champa, K8OCL, "On The Amateur Radio Use of IEEE 802.11b Radio Local Area Networks," *CQ VHF*, Spring 2003, pp. 40-42.

D. Rotolo, N2IRZ, "A Cheap and Easy High-Speed Data Connection," *CQ*, February 2003, pp. 61-64.

N. Sablatzky, K8IT, "Is (sic) All Data Acceptable Data," *CQ VHF*, Fall 2003, pp. 48-49.

M. Simon, J. Omura, R. Scholtz, and K. Levitt, *Spread Spectrum Communications* Vol I, II, III, Rockville, MD. Computer Science Press, 1985

D. Torrieri, "Principles of Secure Communication Systems," Boston, Artech House, 1985.

B. Wyatt, K6WRF, "Remote-Control HF Operation over the Internet," *QST*, November 2001, pp. 47-48.

R, Ziemer, and R. Peterson, "Digital Communications and Spread Spectrum Systems," New York, Macmillan, 1985.

John Champa, K8OCL, is Chairman of the ARRL HSMM Working Group. John B. Stephensen, KD6OZH, is the RMAN-UHF Project Leader of the ARRL HSMM Working Group □□

Coaxial Traps for Multiband Antennas, the True Equivalent Circuit

A new perspective on the analysis and design of this popular antenna element.

By Karl-Otto Müller, DG1MFT

Multiband Antenna Design

Parallel-resonant circuits (called traps) are widely used to isolate parts of multiband antennas to make the antenna resonant on different frequencies (see Fig 1). For more than 20 years these circuits have been implemented as coils wound from coaxial cables.[1,2,3] As shown in Fig 2, the inner conductor of the coil end is connected to the outer conductor at the beginning. Therefore the current is going around the core two times the number of turns. The coaxial cable capacitance represents the capacitor

[1]Notes appear on page 22.

Watzmannstr, 24A
D-85586 POING
Germany

of this parallel-resonant circuit. For an easy design of coaxial traps, VE6YP offers a program in the internet.[4]

In order to design multiband antennas with programs such as EZNEC,[5] traps must be modeled as "loads," defined by their equivalent circuit as shown in Fig 3. The easiest way to determine the values of this circuit is to measure C, L and R_s. C may also be calculated from the coaxial cable length and the capacitance per unit length (reasonable estimate if L < 1/10 λ), but L has to be measured by an appropriate inductance meter. To find out the series resistance R_s, the 3-dB bandwidth of the trap must be measured as described in Fig 4.

The Surprise

Insertion of the measured values of C and L into Thomson's formula

$$f_{res} = \frac{1}{2 \bullet \pi \bullet \sqrt{L \bullet C}}$$

gives exactly half the frequency value which was used in the coaxial trap program of VE6YP to get the number of turns of the trap.

An example: The VE6YP calculation of a coaxial cable trap for 9.5 MHz using RG58 with a core diameter of 35 mm yields 10 turns. The resonance check using a network analyzer results in 9.262 MHz, which is close. EZNEC asks for C, L and R_s and we have to determine these three values before we can start an EZNEC simulation.

Assuming that the resonant frequency is measured correctly, either the value of L or C is only a quarter of

the measured and calculated value or both are half the value. Only one of the following formulas is valid, but which?

$$f_{res} = \frac{1}{2 \bullet \pi \bullet \sqrt{\frac{L}{4} \bullet C}}$$

or

$$f_{res} = \frac{1}{2 \bullet \pi \bullet \sqrt{L \bullet \frac{C}{4}}}$$

or

$$f_{res} = \frac{1}{2 \bullet \pi \bullet \sqrt{\frac{L}{2} \bullet \frac{C}{2}}}$$

For a decision, the impedance versus frequency of the resonant circuit is calculated for all three cases, and compared with the measured values as shown in Fig 5.

It can be seen clearly that the parallel combination L and C/4 is correct. Now somebody may argue that it makes no difference which combination is used for the antenna design as long as the resonance frequency is the same. But there is a significant difference:

The impedance of the three parallel-resonant circuits differs by the factor two or four respectively. The impedance of the correct combination L, C /4 is four times higher than the impedance of the non-correct parallel combination of L/4 and C, which is given as a result of the VE6YP calculation. Thus, the inductive load of the correct combination, L and C/4, has a lengthening effect on the antenna below the first resonance (half the resonant frequency). As a result, the *EZNEC* antenna design, based on the correct equivalent circuit, results in a physically shorter antenna and therefore comes closer to reality.

The Explanation

Three steps are used to show, why the parallel combination of L and C/4 is correct.

Step 1: Symbolical reduction of the number of turns to one, see Fig 6.

Step 2: The winding is cut at the opposite side and connected "crossover" as in Fig 7. The inputs are connected in series.

Step 3: As can be seen from Fig 8, now the two capacitances, C/2 are connected in series, resulting in an effective capacitance of C/4.

Influence of the cable length

Looking again at Fig 5, we find a significant difference between measured and calculated value, based on L in parallel with C, around 60 MHz. It is suggested that this is caused by the cable length. Fig 9 shows the equivalent circuit of our symbolic "one-turn-coil" for frequencies much higher than the resonant frequency. Fig 10 shows the voltage distribution at these frequencies. At the input port, half the

voltage is across each of the coaxial cables. However, at the cross-over connection, both voltages are in phase and have the same amplitude. Therefore there is no current here as illustrated in Fig 10. Consequently, the cross-over connection can be opened without changing the behaviour at high frequencies, see Fig 11. For lower frequencies, up to approximately four times the resonant frequency, the coil inductance can be simulated by a con-
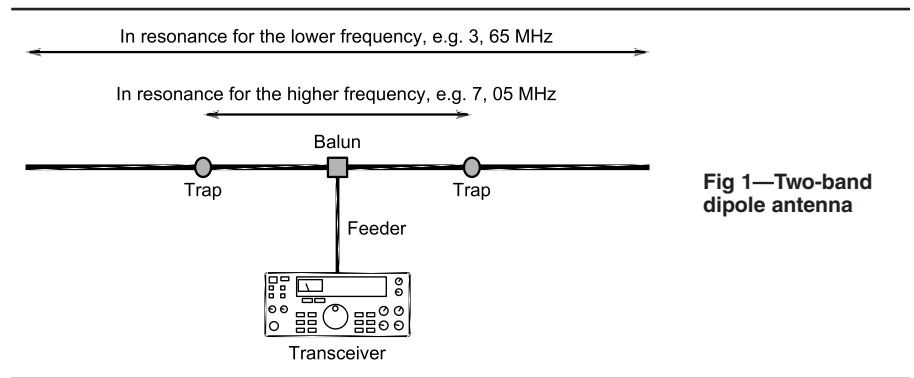
Fig 1—Two-band dipole antenna
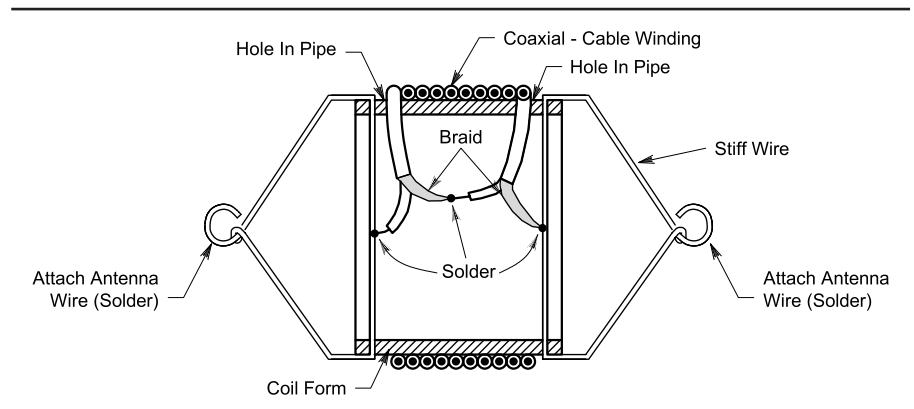
Fig 2—Typical coaxial cable trap (*QST*, Dec 1984)
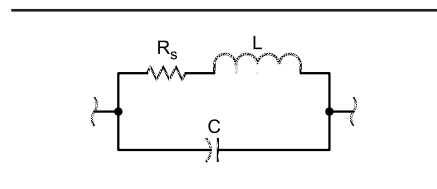
Fig 3—Equivalent circuit of a trap
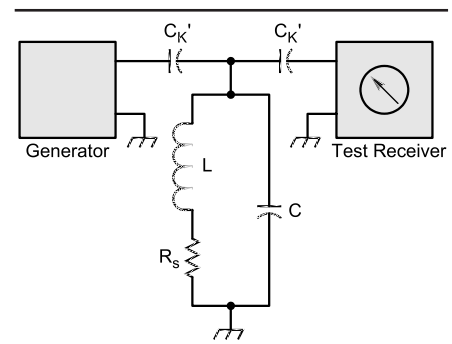
Fig 4—Measurement of the 3-dB bandwidth for calculation of R_s
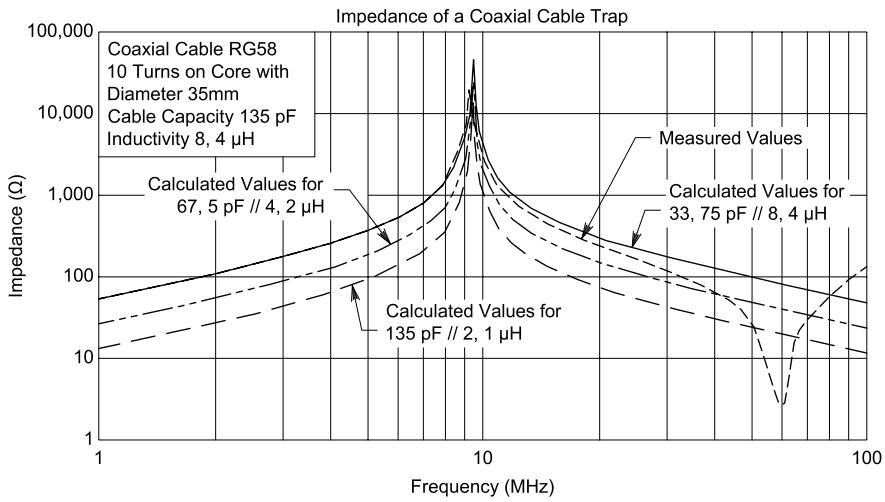
Fig 5—Impedance comparison

Coaxial Cable RG58
10 Turns on Core with
Diameter 35mm
Cable Capacity 135 pF
Inductivity 8, 4 µH

Calculated Values for
67, 5 pF // 4, 2 µH

Measured Values

Calculated Values for
33, 75 pF // 8, 4 µH

Calculated Values for
135 pF // 2, 1 µH

Impedance of a Coaxial Cable Trap
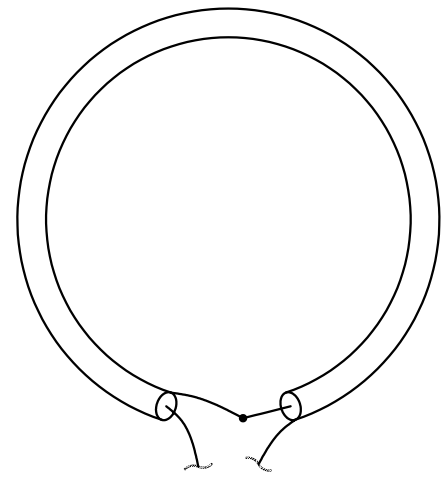
Fig 6—For an easy explanation, number of turns is reduced to one.
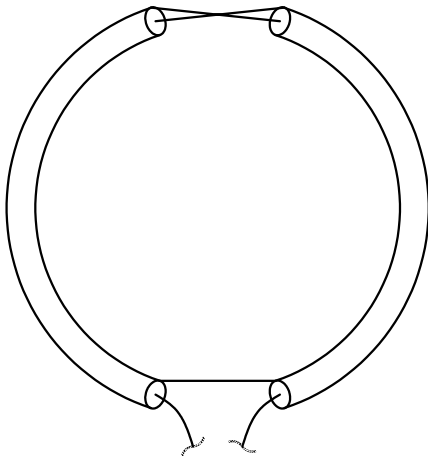
Fig 7—The winding is cut at the opposite side and connected "cross-over". The function of the coil remains totally unchanged.
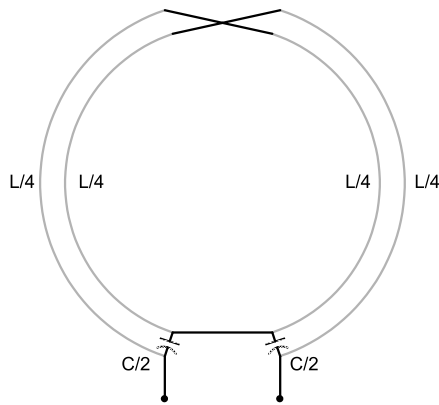
Fig 8—Distribution of capacitance and inductance of the coil
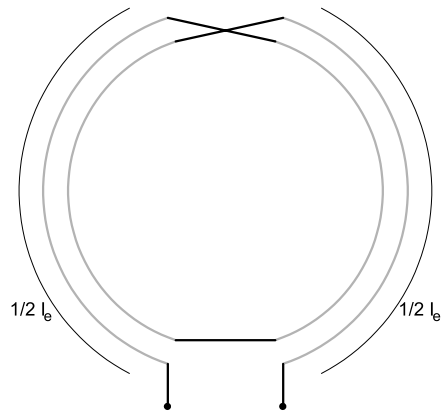
Fig 9—For higher frequencies the electrical length l_e of the coaxial cable is paramount.
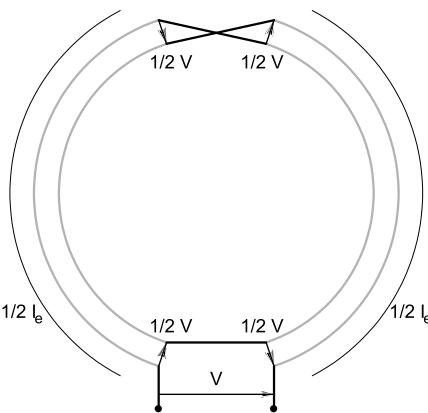
Fig 10—At the cross-over connection both voltages are in phase and have the same amplitude.
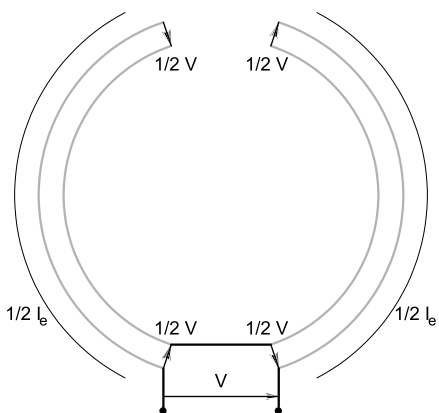
Fig 11—The cross-over connection can be opened without changing the behaviour for high frequencies.
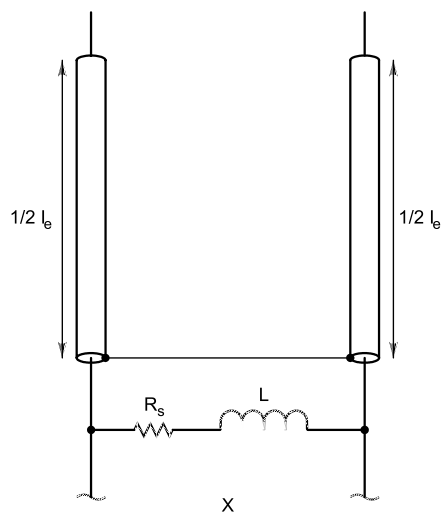
Fig 12—The complete equivalent circuit of a coaxial cable trap with electrical cable length l_e and coil inductance L with losses, represented by R_s

centrated inductance L in parallel with the input port. In series with this inductance we can insert the resistance representing the losses of the trap, as measured by the method of Fig4. Now, Fig 12 shows the complete equivalent circuit of a coaxial cable trap. The measured impedance over a wide frequency range (1 to 500 MHz) is given in Fig 13, showing minima where the total cable length $l_e = 1/2·n·\lambda$ (for odd n only) and maxima, where l_e = n λ (for arbitrary n).

Conclusion

It has been shown that the coaxial cable trap (electrical length l_e of the cable) behaves as a parallel resonant circuit, where $\lambda = (1/n) l_e$ (arbitrary n) and for

$$\frac{1}{2 \bullet \pi \bullet \sqrt{L \bullet \dfrac{C}{4}}} = fres$$

and as a series resonance circuit at all frequencies where $\lambda = (2/n) l_e$ (for odd n only).

Consequences

The correct higher impedance of the coaxial traps, compared to the now-in-use impedance values according to the VE6YP software has two consequences.
• The antenna length is more realistic (i. e. shorter) than predicted by the design software.
• The trap losses are significantly different than predicted and should be considered.
Both are illustrated in Fig 14.

Acknowledgements

I would like to thank Hartwig, DH2MIC, for helpful discussions and the Rohde & Schwarz company, Munich, for providing me with valuable test equipment.

Notes

Notes
[1]R. Johns, W3JIP, "Coaxial Cable Antenna Traps", *QST*, May 1981, pp 15–17.
[2]R. Sommer, N4UU, "Optimizing Coaxial Cable Traps", *QST*, Dec 1984, pp 37–42.
[3]*The ARRL Antenna Book*, 1988, Chapter 7, pp 8-9.
[4]T. Field, VE6YP, *Coaxial Trap Design*, (Freeware, *CoaxTrap.zip*), **www.members. shaw.ca/VE6YP**.
[5]EZNEC is available from Roy Lewallen, W7EL, at **www.eznec.com**.
[6]K. Müller, DG1MFT, "Ersatzschaltbild für Koaxiale Sperrkreise", *Funkamateur 53* (2004) Jan, pp 60-61 (in German).
[7]K. Müller, DG1MFT, "Koaxiale Traps für Multiband-Antennen, Das korrekte Ersatzschaltbild", paper, presented at the DARC Radio Amateur Meeting in Munich, Mar 13/14, 2004 (in German), **www. amateurfunktagung.de**.

CH1 S11   Z   LOG MAG ᵧ1 Ω   ⊿100 kΩ

CAL

CPL

FIL
10 k

Start 1 MHz   50 MHz/div   Stop 500 MHz

Comment 1: Koaxtrap 10 Windungen RG68 C/U auf Kern 35mm Durchmesser
Date: 27.Jan.04 13:03:39

Fig 13—Impedance minima and maxima of the coaxial cable trap from 1 MHz to 500 MHz; vertical log scale from 1 Ω to 100 k Ω

Antenna Length = 37m

Fictive Efficiency
η = 93% at 3.7 MHz
η = 64% at 7.06 MHz

Trap Losses and Voltages
P = 3.4 W V = 137 V at 3.7 MHz
P = 18.2 W V = 518 V at 7.06 MHz

EZNEC Design Based on the
Non - Correct Equivalent Circuit

Antenna Length = 33m

Efficiency
η = 76% at 3.7 MHz
η = 87% at 7.06 MHz

Trap Losses and Voltages
P = 12.6 W V = 531 V at 3.7 MHz
P = 6.2 W V = 607 V at 7.06 MHz

EZNEC Design Based on the
Correct Equivalent Circuit

Fig 14—Errors, caused by the use of the wrong (¼ L // C, upper picture, printed in red) equivalent circuit. The antenna below (green) is calculated on basis of the correct equivalent circuit (L // ¼ · C). The differences in trap losses and voltages are not negligible! Example is a dipole antenna for 40/80 m, applied power is 100 W, wires are lossless, 10 m above ground, traps made from RG58 C/U, Q = 100.

Karl-Otto Müller, DG1MFT, was a development engineer at Rohde & Schwarz in Munich until his retirement. For more than 40 years he was responsible for all EMI test instrumentation with a specialization in test receivers. □□

Software Defined Radios for Digital Communications

An Open Platform for SDR Development with Free Development Software

By John B. Stephensen, KD6OZH

O ne of my interests in Amateur Radio is building the equipment that I use to communicate. In the early '90s I constructed a VHF-UHF Amateur Radio station that was entirely home-made and in the late '90s I did the same with an HF station.[1] Both have computer-controlled tuning but are essentially analog designs. They operate well but are not up to today's state of the art.

In the past I've experimented with DSP evaluation cards and FPGAs[2], but the available hardware did not provide the bandwidth and processing power necessary for a modern soft-

3064 E. Brown Ave
Fresno, CA 93703
kd6ozh@verizon.net

ware-defined radio. Today's radio must support high-speed digital communication for new multi-media applications. In the past year, several new processors, programmable logic devices and data converters have appeared using 90 to 180-nm feature sizes to provide a low cost solution. This article describes the DSP card that I have developed (dubbed the DCP-1) to take advantage of this technology.

Software-Defined Receiver Architecture

The first thing that I did was to examine the available technology to determine the proper architecture for the new radios. There are many high-speed analog to digital converters (ADCs) available with 90 to 100 dB dynamic range. For ADCs, dynamic range is defined as the ratio of the

maximum signal level that can be digitized by the ADC and the minimum level of distortion products generated during conversion at any signal level. It is much like blocking dynamic range for analog radios.

Recently, there has been experimentation with direct digitization of the RF signal.[3] This works well for wide-band signals, such as a 76.8 kbps FSK terrestrial data link, where dynamic range is limited:

−174	dBm/Hz	290 K thermal noise
+2	dB	2 dB NF
+52	dB-Hz	140 kHz bandwidth
−120	dBm	MDS
−33	dBm	S9 + 60 dB (ADC full scale)
−120	dBm	MDS
87	dB	Required dynamic range

Narrow-band modes like PSK-31 require a much higher dynamic range. The minimum discernable signal (MDS) for a PSK31 signal on a satellite downlink versus the maximum signal level is:

–174	dBm/Hz	290 K thermal noise
–9	dB	0.5 dB NF
+15	dB-Hz	30 Hz bandwidth
–168	dBm	MDS
–33	dBm	S9 + 60 dB
		(ADC full scale)
–168	dBm	MDS
135	dB	Required dynamic range

This ADC needs to be preceded by an analog filter that attenuates interfering signals by 35-45 dB. Here's an exercise that brings ADC dynamic range into focus. Calculate the blocking dynamic range for the lowest performance analog mixer IC available:

–174	dBm/Hz	290 K thermal noise
+5	dB	NE602 45 MHz NF
+15	dB-Hz	30 Hz bandwidth
–154	dBm	MDS
–30	dBm	NE602 1 dB compression level
–154	dBm	MDS
124	dB	NE602 blocking dynamic range

The NE602 mixer has been used in SSB receivers that cost less than a single high-speed high-performance ADC. A high-performance microwave mixer provides much more dynamic range:

–174	dBm/Hz	290 K thermal noise
+11	dB	SYM-30DHW conversion loss + IF amplifier NF
+15	dB-Hz	30 Hz bandwidth
–148	dBm	MDS
+14	dBm	SYM-30DHW 1 dB compression level
–148	dBm	MDS
162	dB	SYM-30DHW blocking dynamic range

Another approach that has been used lately is a direct-conversion receiver[4] in which the RF or IF is heterodyned to dc in a quadrature mixer. The resulting in-phase (I) and quadrature (Q) signals are then digitized and processed. Sigma-delta ADCs designed for high-quality audio systems have dynamic ranges exceeding 120 dB and bandwidths up to 70 kHz. Unfortunately, the poor opposite-sideband suppression, which rarely exceeds 50 dB, wastes the dynamic range.

The best receiver architecture is still a superheterodyne with analog filters as shown in Fig 1. The last IF amplifier is followed by an ADC and software signal demodulation. The requirements placed on the analog filters are very much relaxed when compared to an all-analog design. They are now present only to increase dynamic range and can be very inexpensive. The DSP provides the steep-skirted filters. For example, a $5, 4-pole monolithic filter can replace a $200, 10-pole filter in an SSB receiver.

Software-Defined Transmitter Architecture

Since 1960, transmitters for amateur bands have tended to copy receiver architecture in order to share expensive analog filters and provide better frequency stability. The analog filters are now inexpensive and frequency is controlled to fine tolerances by a PLL. There is no longer a need to share components in a transceiver and it often would increase costs to do so.

Here is where a direct-conversion design can be used. Transmitter dynamic-range requirements are minimal, compared to a receiver, and signal levels are high, so there is no problem with 1ow-frequency noise. DSP can generate I and Q base-band signals for any desired modulation and simple analog low-pass filters suppress any spurious signals. Two matched digital-to-analog converters (DACs) are required to support this architecture.

The main issue in the past has been the generation of quadrature RF local oscillator signals. At HF and below, digital dividers can create accurate signals from an oscillator at two or four times the carrier frequency. In the VHF and UHF range there are IC quadrature modulators with integrated wide-band polyphase 90° phase-shift networks. See Fig 2.

Digital Down-Converters

Many designs have used digital down-converters (DDCs) to process the output of a high-speed ADC. This works well when the signal of interest is narrow compared to the IF. When wide-band signals are digitized at low IFs most DDCs cannot be used. This is because cascaded integrator
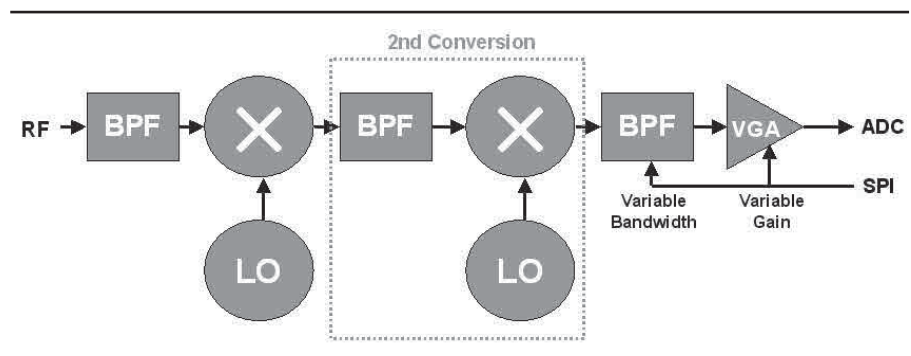
Fig 1—Superheterodyne receiver block diagram.
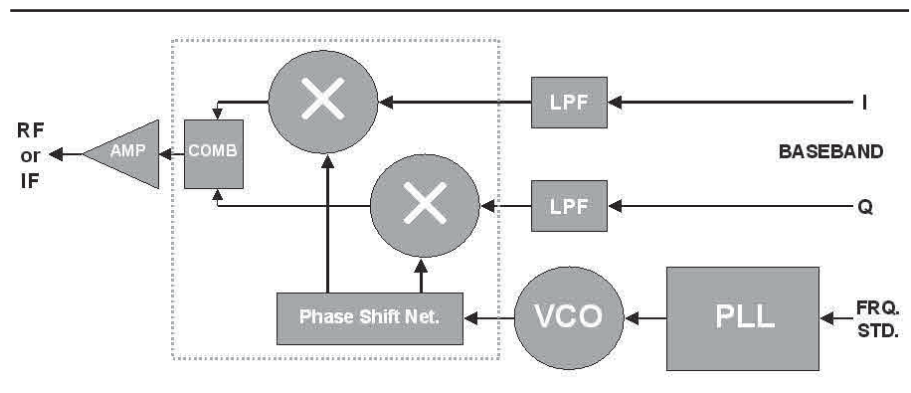
Fig 2—Direct-Conversion Transmitter.
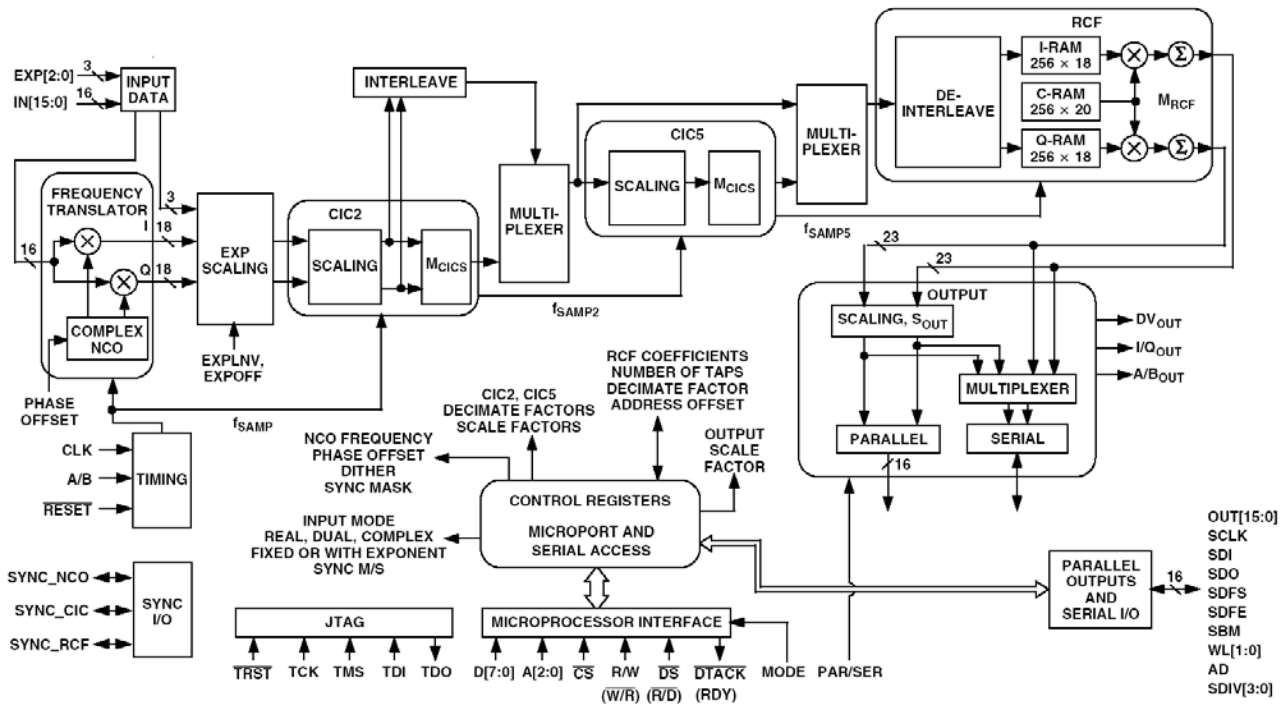
Fig 3—AD6620 DDC with CIC2 and CIC5 Filters.

comb (CIC) filters are used to perform the initial filtering and they support only very narrow pass-bands. For example, in the AD6620 DDC, the usable CIC2 filter output is 0.18% of the sample rate for 90 dB alias rejection. The CIC5 filter is better at 3% bandwidth. Yet, if the IF is 10 MHz and the signal bandwidth is 2 MHz, CIC filters cannot be used. See Fig 3.

More flexibility is needed in the circuitry following the ADC but DSP chips do not provide the necessary processing power at a reasonable cost. Today's FPGAs combine power and flexibility with low cost because they include dedicated multipliers and larger amounts of block RAM. The FPGA can easily be configured to implement FIR filters immediately following the ADC and can also perform operations such as fast Fourier transforms (FFTs) that DDCs do not support.

PC Interface

Any new radio should be able to be closely coupled to a personal computer. The PC is often the source and destination of the data, voice or video being exchanged. In addition, PCs now have 2 GHz processors and signal-processing instructions so they will be used for source coding and decoding. See Fig 4.

Traditional radios have used an
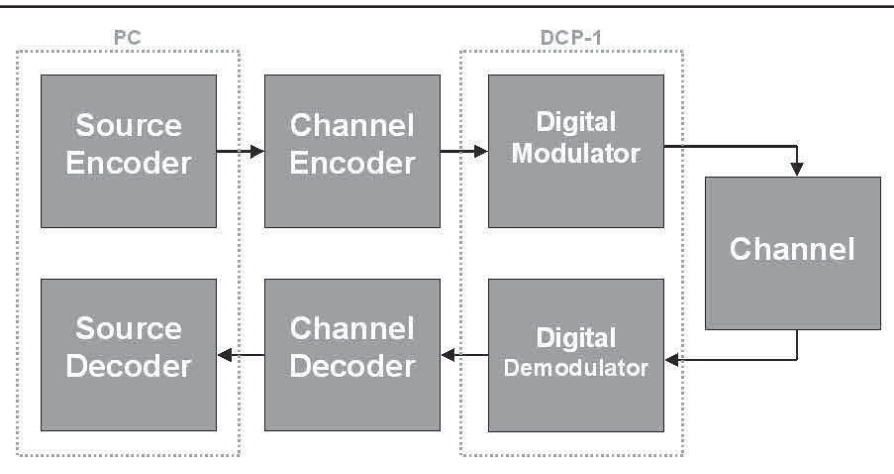
Fig 4—Digital Communications Link.

RS-232 interface to PCs. Even with the latest enhancement to 1 Mbps data rates, this is inadequate for multimedia applications. The Universal Serial Bus (USB) is the best interface to use on modern PCs. Full-speed USB (12 Mbps) is adequate for today's applications, but High-Speed USB (480 Mbps) may be needed in the future.

USB is not appropriate in an interface between high-speed digitizers and modulation or demodulation software in the PC. USB has a built-in 1 to 2-millisecond latency that results in

overly large buffers and the inability to control timing closely. Consequently, a USB radio needs an internal processor to perform real-time tasks.

Channel encoding and decoding to implement error detection and correction algorithms may be done in the PC or the external processor depending upon complexity.

DCP-1 Digital Communications Processor

The DCP-1 is contained on a 3.5-inch square PCB and contains all necessary data converters and signal

processing for amateur transceivers. It digitizes the receiver IF at 19.2 Msps. This sample rate was chosen to support analog and digital modes in use on the amateur bands today and allow high-speed modes such as OFDM.

The 19.2 Msps sampling rate supports IF bandwidths up to 6 MHz. The UHF front end module uses a 330 MHz SAW filter and an image-reject down-converter to obtain a 5.4 MHz -0.5 dB bandwidth. This translates to a 2.1-7.5 MHz second IF with over 90 dB of image rejection. A low final IF frequency maximizes the dynamic range of the ADC. Other front-end modules have narrower roofing filters to match the widest signal bandwidths in their frequency ranges. At lower RF frequencies, direct conversion to a 6 MHz IF is used. This IF frequency was chosen to allow the optional use of narrow-band ceramic and quartz crystal filters to obtain narrower sampling bandwidths and increase dynamic range. The sampling rate is maintained within +/-2.5 PPM by a low-cost TXCO. An accurate clock is necessary for IF sampling for OFDM modems with over 100 subcarriers. It is also necessary when the FPGA and one DAC are used to implement a low-spur DDS. The TCXO output is buffered and made available to external RF modules.

The design is built around a Xilinx Spartan-3 FPGA and Oki ML67Q5003 MCU as shown in the DCP-1 block diagram. The FPGA and MCU are the 144-pin TQFP packages in the center of the PCB photograph, Fig 5.

The FPGA functions as a highly programmable high-speed DSP coprocessor. The XC3S200 contains twelve 5.8 ns 18×18 multipliers, 216 k of 2.4 ns dual-port RAM and 4320 logic elements (200,000 gates) with 750 ps propagation delays.

An FPGA slice, consisting of two logic elements, is shown in Fig 7. Each logic element contains a four-input look up table (LUT) which can generate any arbitrary logic function. Multiple LUTs can be combined to create functions with more inputs. Half of the LUTs on the chip can also be reconfigured as 16-bit shift registers or 16-bit RAMs. To the right of the LUTs is dedicated carry logic to speed up arithmetic functions. Following that are storage elements that may be configured as D-type flip-flops or level-sensitive latches. The flip-flops toggle at 500 MHz.

Four slices are grouped into a configurable logic block (CLB) that can process one byte of data as shown in Fig 8. The CLBs may exchange data directly with their neighbors or con-nect to chip-wide busses via the switch matrix. The CLBs are combined with other components on the chip as shown in Fig 9.

Also contained in the FPGA are digital clock managers (DCMs) to multiply and divide clocks and generate multiphase clocks for the logic elements. Around the periphery of the chip are input/output blocks (IOBs) that contain registers and three-state drivers. Various logic families from 1.2 V to 3.3 V are supported plus LVDS.

The FPGA may be programmed to provide FIR signal filters, FFT engines and perform other signal processing tasks. The ML67Q5003 MCU, shown in Fig 10, controls the configuration
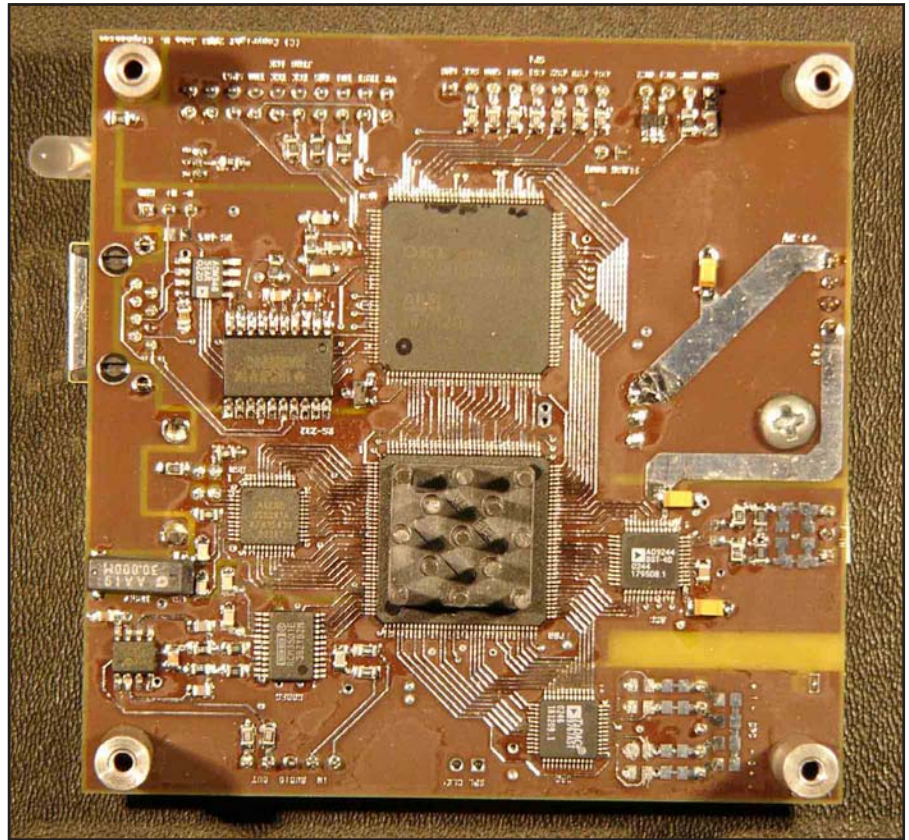
Fig 5—DCP-1 PCB. The ADC and DAC are in the lower left corner. Working clockwise are the FPGA (under heat sink), audio CODEC, USB transceiver, RS-232/485 interface and 32-bit RISC MCU.
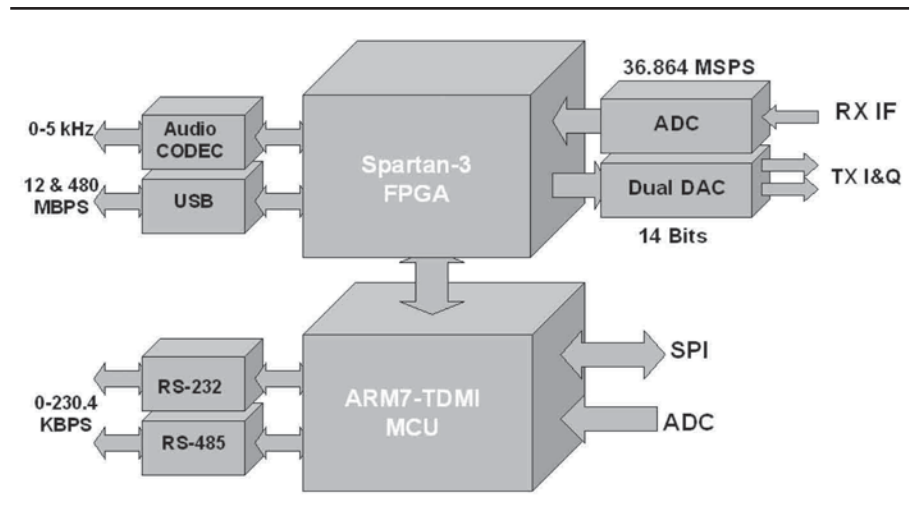
Fig 6—Digital Communications Processor Module.

Fig 7—Spartan-3 FPGA Slice.
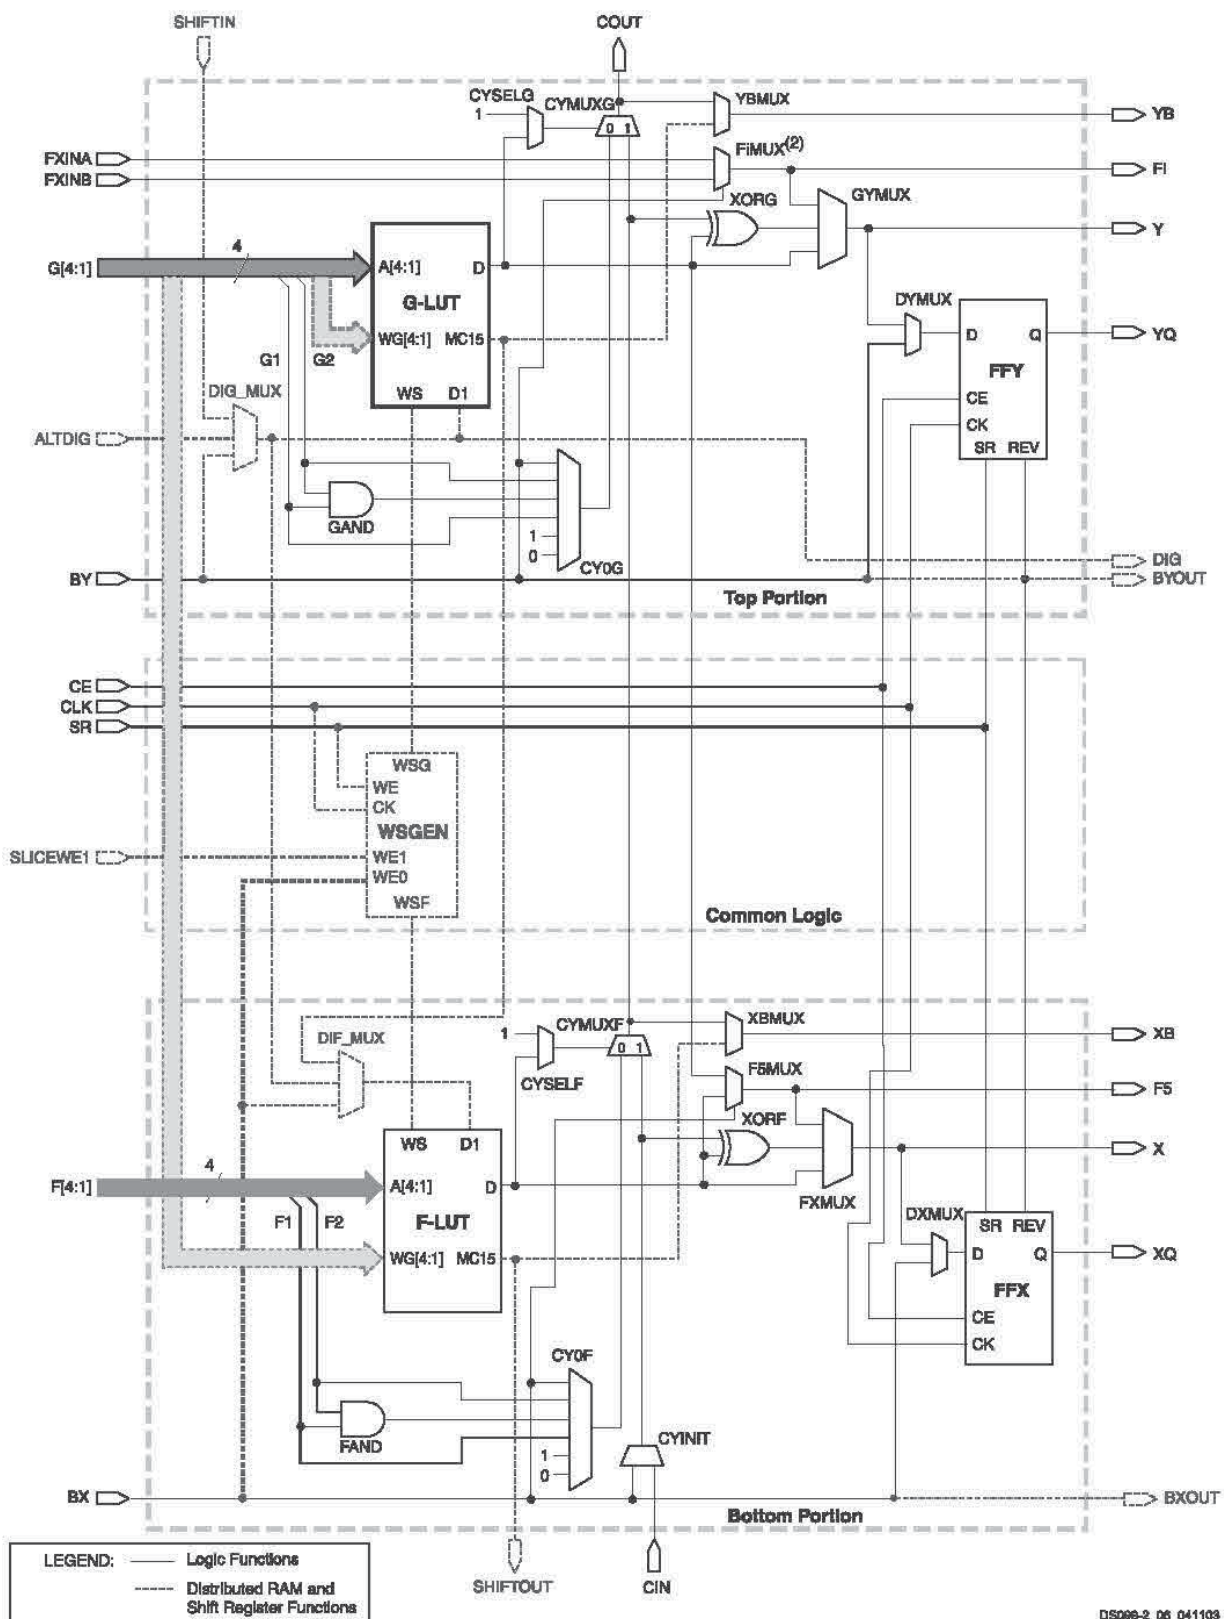
Notes:
1. Options to invert signal polarity as well as other options that enable lines for various functions are not shown.
2. The index i can be 6, 7, or 8, depending on the slice. In this position, the upper right-hand slice has an F8MUX, and the upper left-hand slice has an F7MUX. The lower right-hand and left-hand slices both have an F6MUX.

of the FPGA and provides that programming.

This MCU is based on an ARM7-TDMI processor that is clocked at 58.9824 MHz. This is a classic RISC CPU that uses 3-address arithmetic and logic instructions operating on 31 general-purpose 32-bit-wide registers. Arithmetic operations include multiply-accumulate for signal processing. Memory is accessed via load and store instructions that may move one or multiple words. Each instruction can be made conditional on various status flags and the results of arithmetic and logic operations.

The FPGA JTAG interface is connected to five MCU PIO port bits to allow the MCU to program the FPGA. After programming, the MCU has access to the FPGA logic and RAM via the 16-bit MCU data bus (XD0-15) and 5 address lines (XA1-5). The two direct memory access controller (DMAC) channels are also connected to the FPGA to allow data transfers without processor intervention.

The MCU has 32 kB of RAM and 512 kB of flash ROM to hold software for the MCU and configuration data for the FPGA. The MCU also has a mask ROM that contains a bootstrap loader to load the flash ROM via the 16550-compatible UART. This UART connects to the outside world via an RS-232 or RS-485 interface at up to 230.4 kbps. This serial interface uses the RJ-45 connector shown in the PCB photograph.

After programming, the serial port may be used by the MCU to control other devices. A Serial Peripheral Interface (SPI) port (labeled SSIO) is also provided via the connector at the bottom of the PCB. This is commonly used to control PLL chips and configure analog hardware via shift registers or CPLDs. One ADC port is also made available on the bottom connector along with two high-voltage high-current open-collector drivers. Another ADC port is used to monitor the USB bus voltage.

The FPGA also provides the necessary glue logic to interconnect the MCU with the high-speed ADC and DACs, the audio CODEC and the USB controller. The ADC and DAC connect to a Spartan-3 FPGA via two 14-bit parallel busses and the FPGA provides all clocking for those devices.

The high-speed ADC is an Analog Devices AD9244-40 (shown in Figure 11), which is a 14-bit device. It contains a fast sample-and-hold amplifier (SHA) capable of sampling inputs up to 240 MHz. The ADC uses internal error-correction logic in the 10-stage

pipeline to ensure maximum linearity. It has a 90-dB dynamic range up to 45 MHz (as shown in Fig 12), which degrades to 75 dB at higher frequencies. This allows sampling of 10.7 MHz IFs from VHF or UHF radios or the 2-meter IF from a microwave transverter.

The high-speed DAC is an AD9767, which is a dual 14-bit device with a common internal voltage reference that is capable of running at 125 Msps. In this case, data for both DACs is multiplexed onto port 1. The DACs are highly

linear, as shown in Fig 14, so that the transmitted signal can occupy less bandwidth than the transmitter low-pass filters without generating excessive spurs within the passband.

A Texas Instruments PCM3501 single-channel 16-bit audio CODEC (Figure 15) is provided so that analog voice modes may be used independently of the PC. It uses a serial interface and is capable of operating at 8, 12, 16 or 24 ksps with an 88-dB dynamic range.

An Agere USS2X1A UTMI chip,
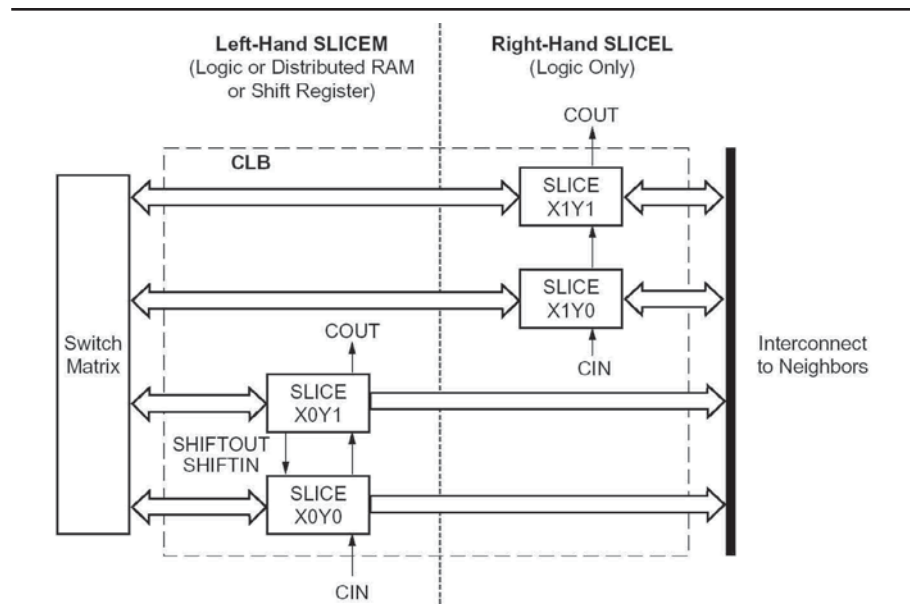
Fig 8—Spartan-3 FPGA CLB.
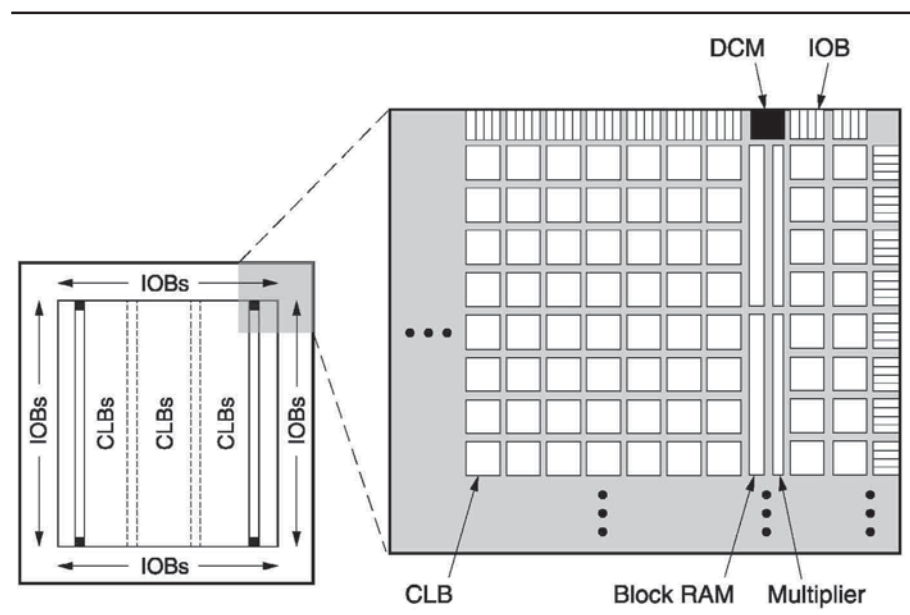
Fig 9—Spartan-3 FPGA Die Layout.

shown in Fig 16, provides the USB interface. It performs all serial-to-parallel and parallel-to-serial conversions, clock and data recovery, bit stuffing and unstuffing and data-rate buffering. It operates at either 12 or 480 Mbps and provides a byte-wide interface to the FPGA.

Development Tools

Three types of development tools are used to program the DCP-1. They

Table 1: Commonly used Sampling Rates and Corresponding Decimation Factors

Application	Sample Rate	Decimation
7.68 Mbps OFDM	19.2 Msps	1
240 kbps OFDM	600 ksps	32
128 kbpsS GMSK	256 ksps	75
76.8 kbps FSK	153.6 ksps	125
9.6 KBPS FSK	19.2 ksps	1000
5 kHz Audio	12 ksps	1600
3 kHz Audio	8 ksps	2400
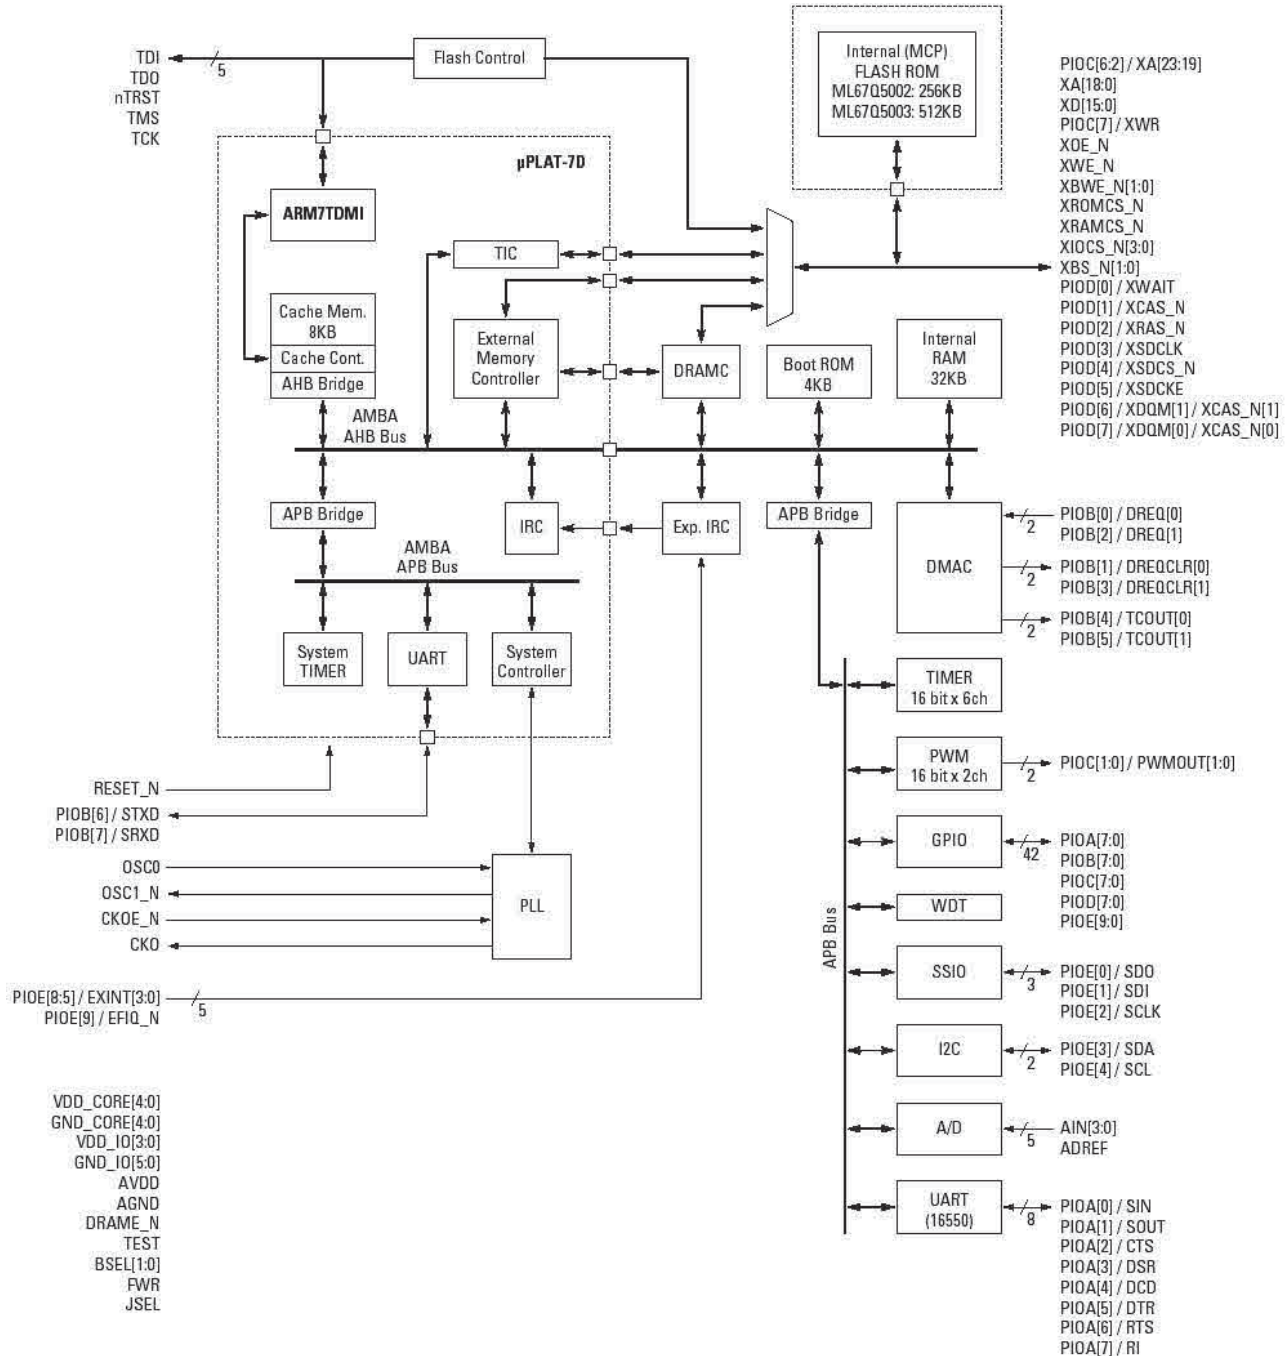
Fig 10—ML67Q500x Series MCU Block Diagram.

are used to configure the FPGA, program the MCU and generate digital filter coefficients. All three are available for download at no charge from the Internet.

FPGA development is done with the *Xilinx ISE 6.2i* development software. The *ISE 6.2i WebPACK* is available for download at no charge from the Xilinx Web site. It supports design entry in schematic diagram form, VHDL, Verilog and state-transition diagrams. The system then synthesizes the necessary logic, lays it out on the chip, routes interconnections and produces a configuration file. Free tools are also available for design simulation, timing analysis and test bench generation. The user interface is shown in figures 17 through 20.

ARM7 CPU software development is supported by the GNU Development Environment (*GNUDE*), which is available for download at no charge from the Free Software Foundation. It includes a CPU simulator, debugger, assembler, linker, and compilers for *C, C++, Ada, Java* and *Fortran. C* language utilities for embedded systems are also available in source code form.

FIR filter development can be done using various free tools available on the Internet. One Web site, **www.nauticom.net/www/jdtaft**, created by J. D. Taft, contains *Java* applets for designing most types of digital filters. These include FIR and IIR low-pass, high-pass, band-pass and band-reject filters, plus Hilbert transformers, differentiators, notch filters and comb filters.

Additional development tools are also available for a fee from many suppliers including Xilinx, Nohau and Momentum Data Systems. Figure 21 shows the MDS filter development software. The DCP-1 includes a connector for in-circuit emulators.

Conclusion

The DCP-1 provides a much better base for software-defined transceiver design than commonly available development boards. The board may used as an add-on to existing transceivers or form the basis for developing a new state-of-the-art radio. It provides the necessary analog and digital hardware for both narrow and wide-band transceivers in one package and the development tools are free.

As this goes to press, the author is readying an improved version of the DCP-1 that increases the ADC dynamic range to 96 dB and includes a more powerful and easier-to-use USB interface. The new board fits standard extruded-aluminum enclosures and

provides fully filtered and shielded I/O connectors. The author will make PCBs and parts kits available. Pricing is expected to be below $200. Assembly services will also be provided.

Future articles will describe the analog front-end modules that combine with the DCP-1 to make a complete radio.

John Stephensen, KD6OZH, has been interested in radio communications
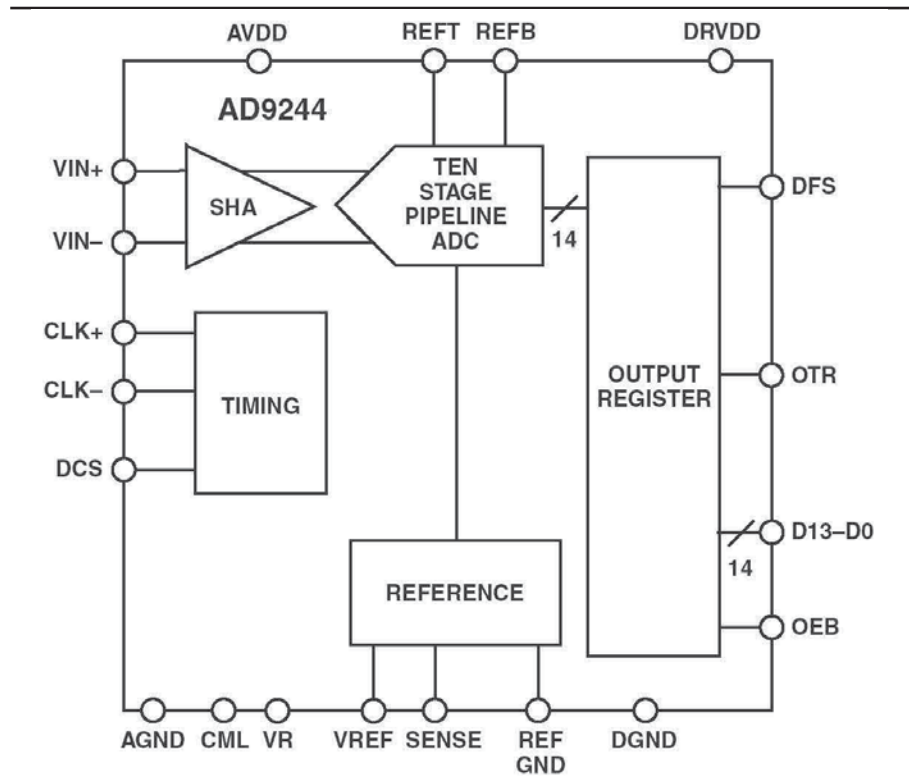
Fig 11—AD 9244 ADC block diagram.
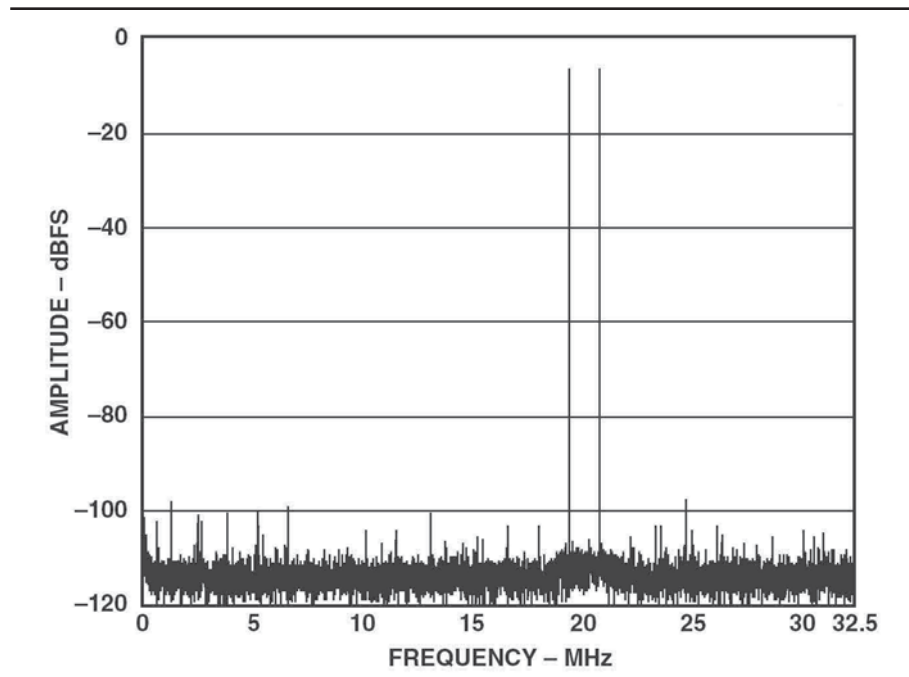
Fig 12—AD9244 ADC output for 2-tone Input.

since building a crystal radio kit at age 11. He went on to study electronic engineering at the University of California and has worked in the computer industry for almost 30 years in engineering development and management positions. He was a founder of PolyMorphic Systems, which started manufacturing personal computers in 1975, a founder of Retix, a communications software and hardware manufacturer, and Vice President of Technology at ISOCOR, which developed messaging and directory software. John received his amateur radio license in 1993 and has been active on amateur bands from 7 MHz to 24 GHz. His interests include digital and analog amateur satellites, VHF and microwave contesting, HF DXing and designing and building
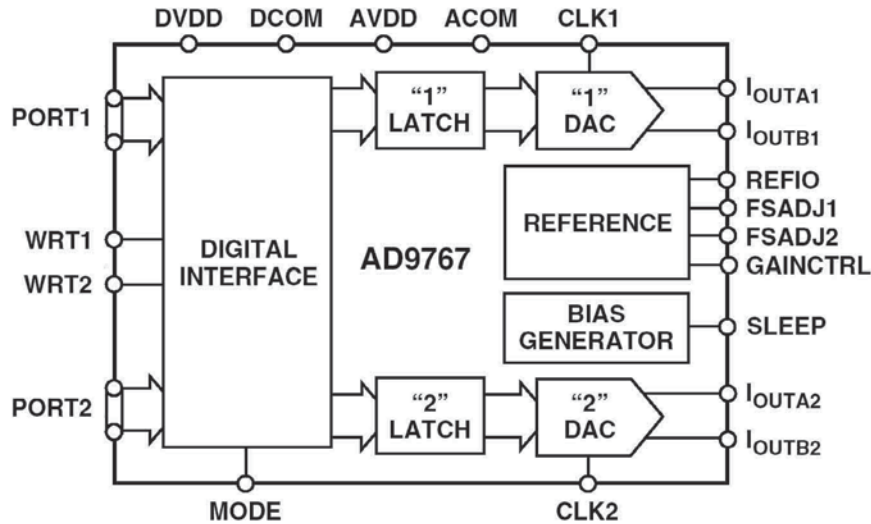
Fig 13—AD9767 dual 14-bit DAC.
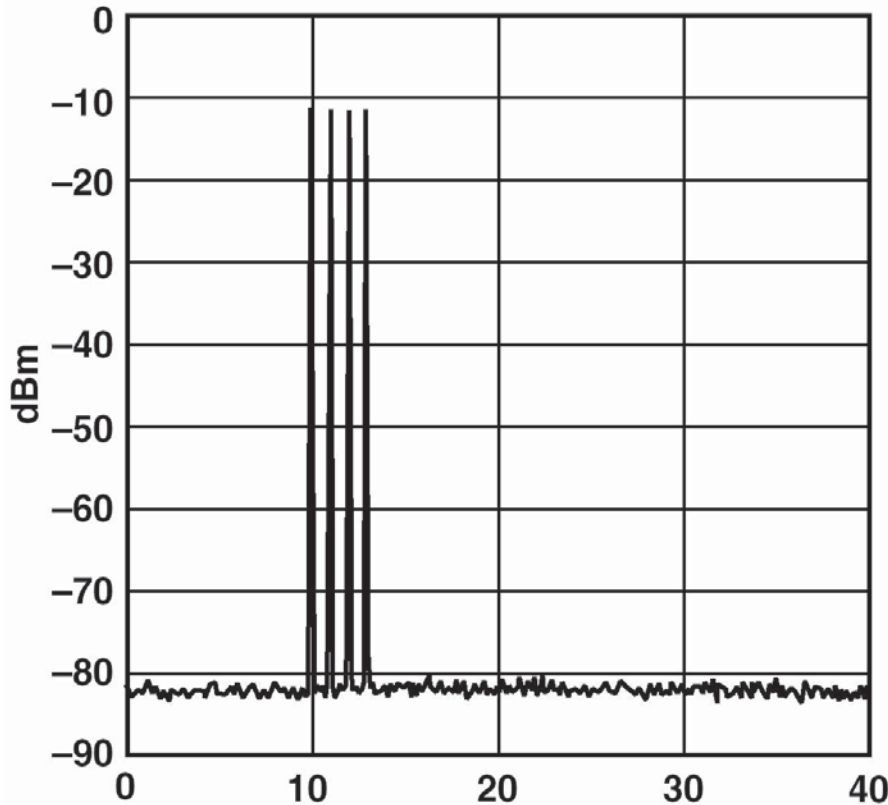
Fig 14—AD9767 Output for four tones.

amateur radio gear. Recently, he has been experimenting with FPGA-based software defined radios and applying DSP to high-speed digital communication. John serves as the RMAN-UHF project leader for the ARRL HSMM Working Group.

Notes

[1] J. B. Stephensen, KD6OZH, "The ATR-2000: A Homemade, High-Performance HF Transceiver—Part 1", QEX Mar/Apr 2000, pp 3-15; Part 2, May/Jun 2000 pp 39-51; Part 3, Mar/Apr 2001 pp 3-17.

[2] J. B. Stephensen, KD6OZH, "Software-Defined Hardware for Software-Defined Radios", *QEX*, Sep/Oct 2002, pp 41-50.

[3] G. Youngblood, AC5OG, "A Software-Defined Radio for the Masses", QEX, Jul/Aug 2002. pp 13-21

[4] J. Scarlett, KD7O, "A High-Performance Digital Transceiver Design", QEX, Jul/Aug 2002, pp 35-44.
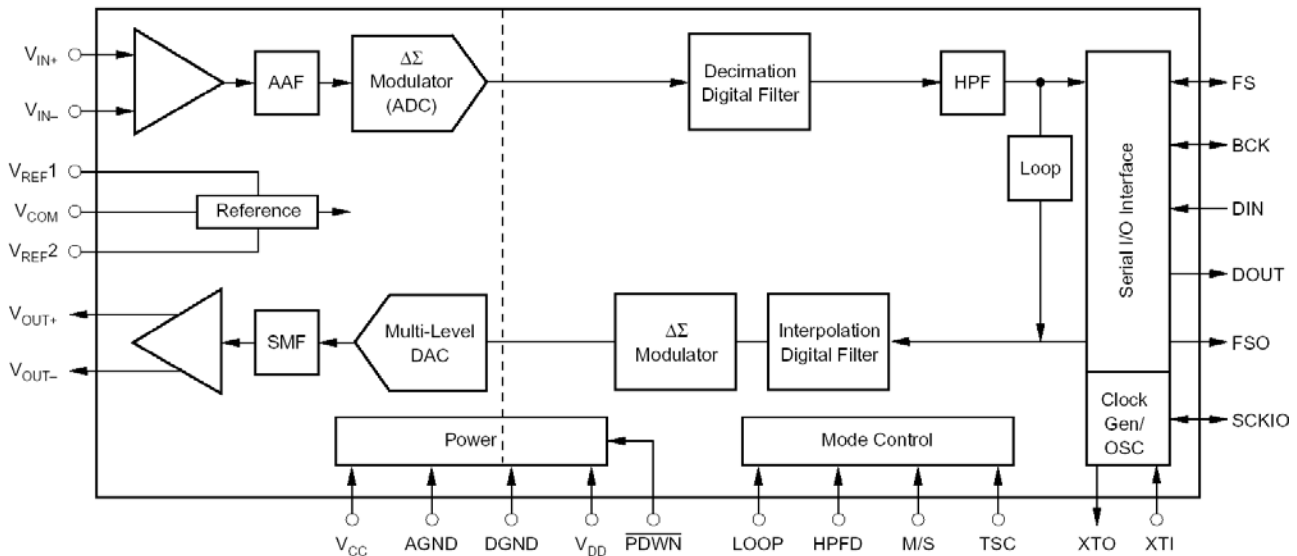
Fig 15—PCM3501 audio CODEC.
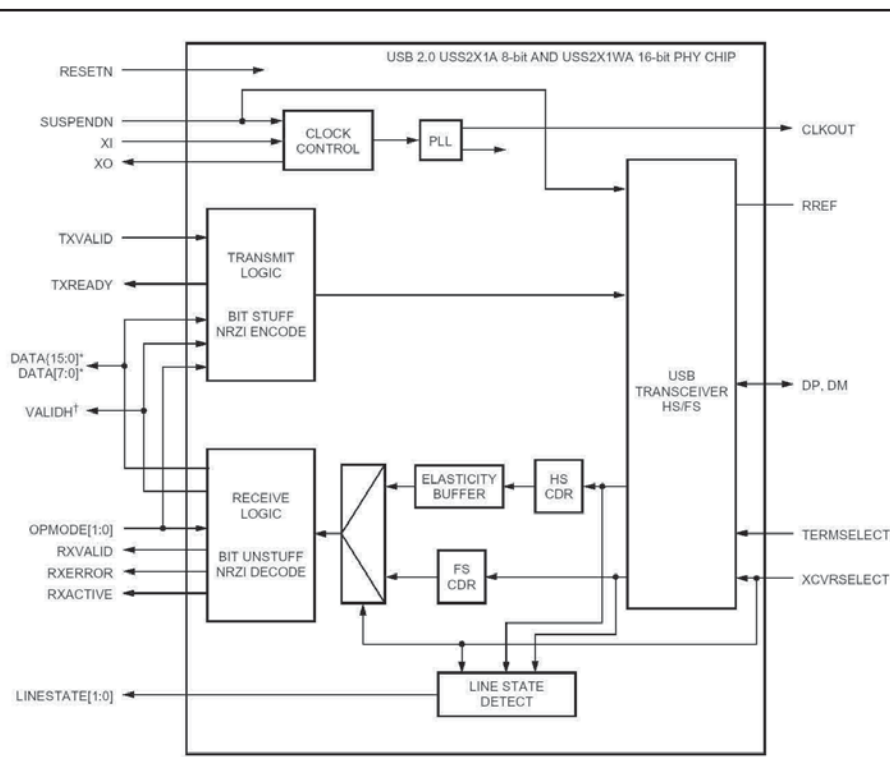
Fig 16—USS2X1(W)A USB interface block diagram.
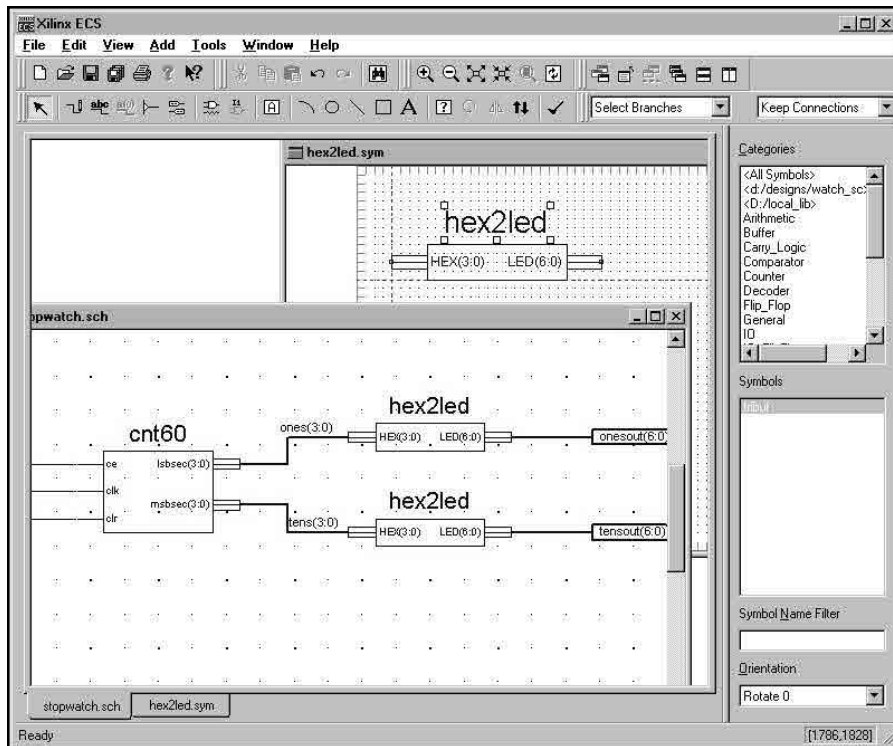
Fig 17—Xilinx ISE schematic entry.
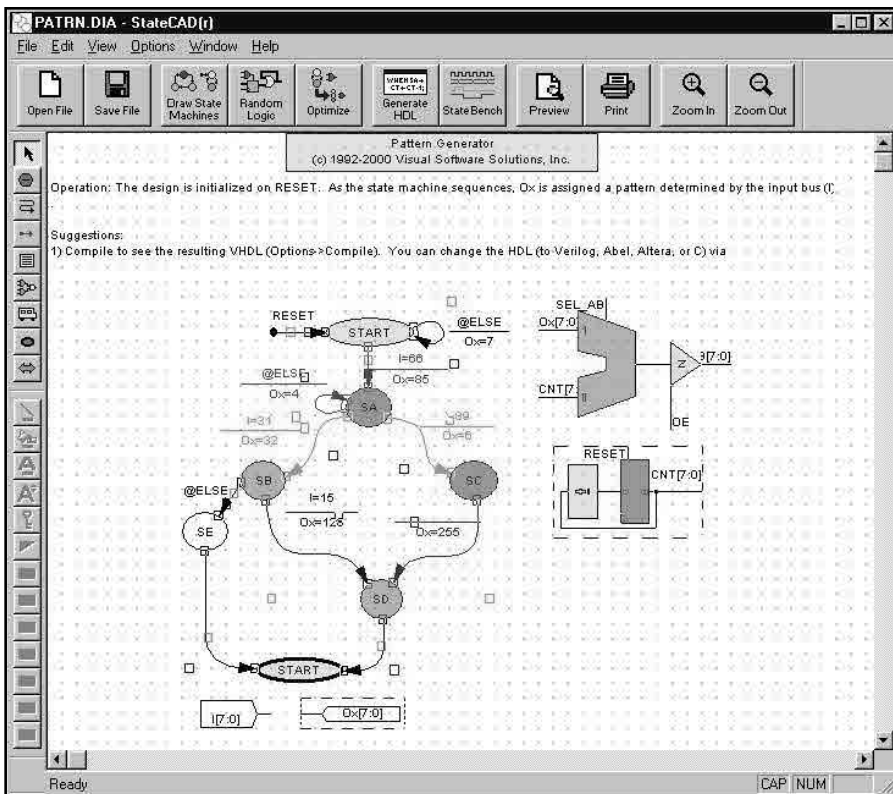
Fig 18—Xilinx ISE State Machine entry.
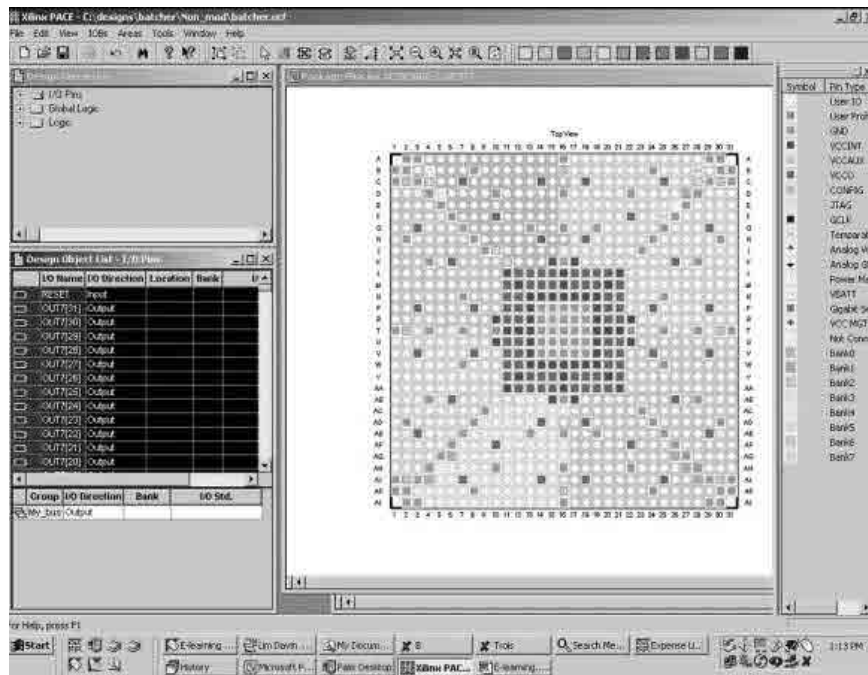
Fig 19—Xilinx ISE pin and constraint entry.
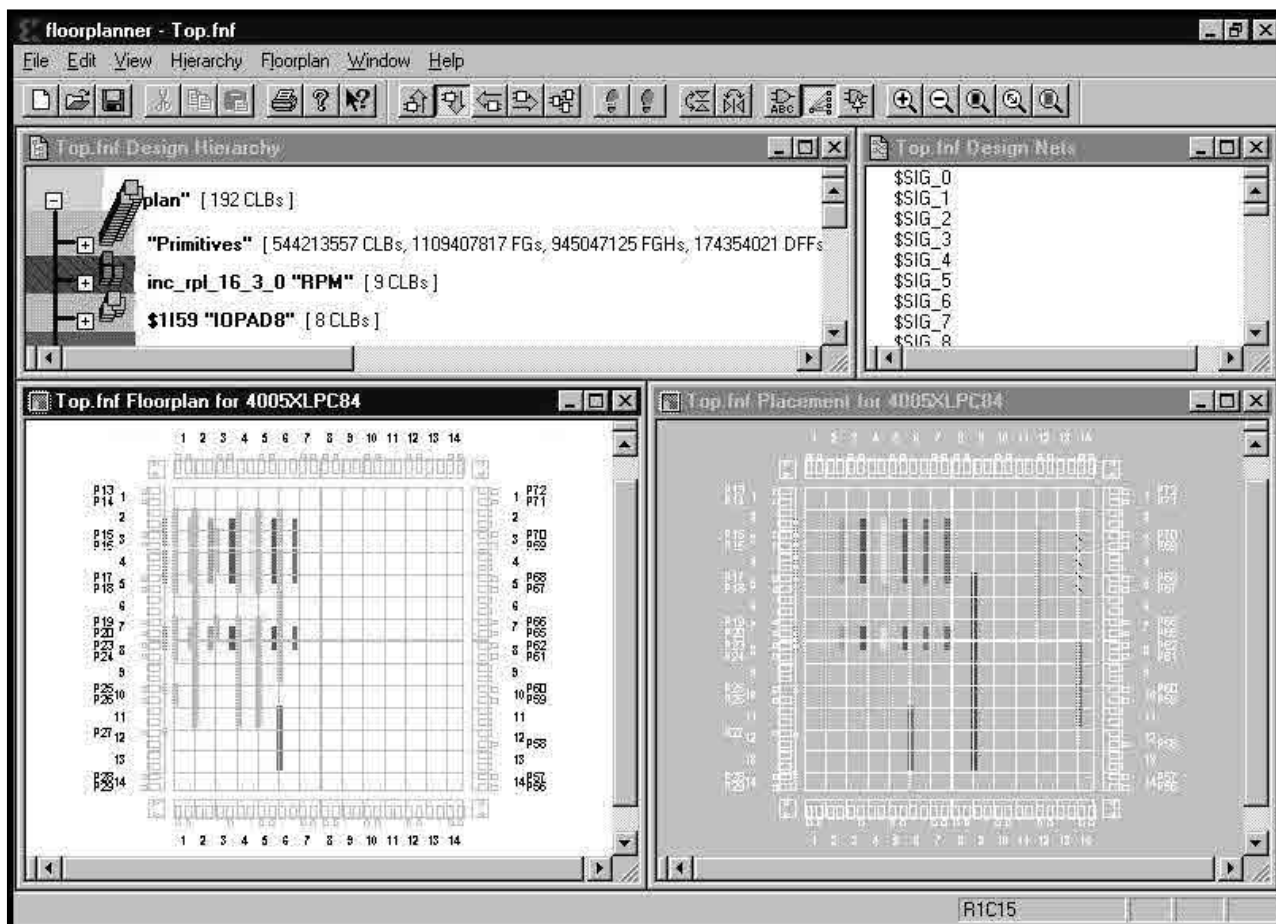
Fig 20—Xilinx ISE floor planner.
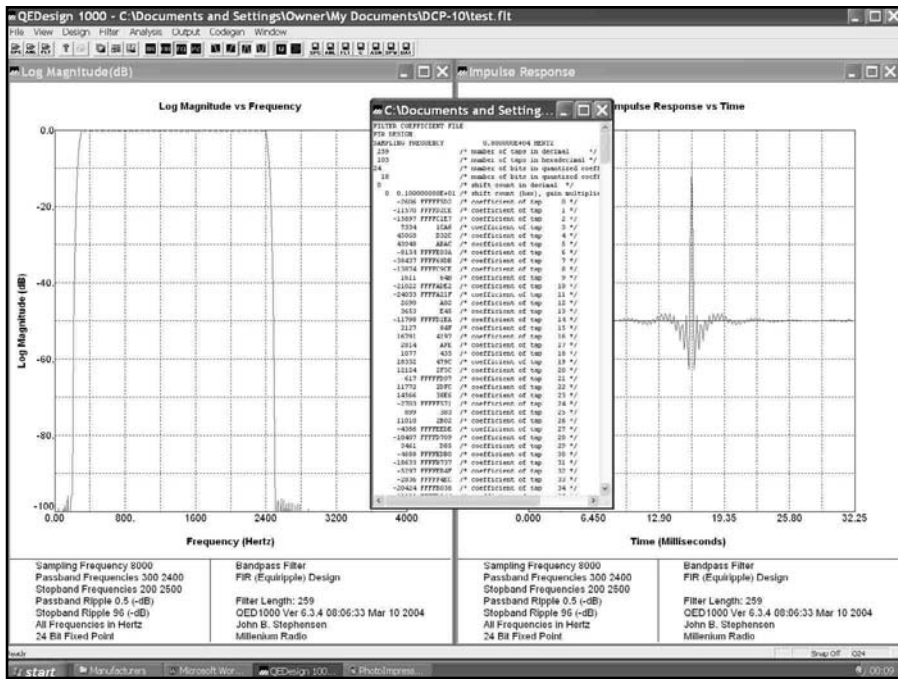
Fig 21—MDS *QED1000* filter design software. □□

ATX Adventures

Phil describes how a surplus PC ATX power supply
can be transformed into a 20 A 13.8 V supply
suitable for transceiver use.

By Phil Eide, KF6ZZ

Last December I cast a cold eye at the dead ATX switching supply that had destroyed my computer. It seemed fitting revenge to convert it to 13.8 V dc and put it back online powering my 100 W HF transceiver. This article details the long, long train of hurdles on the road to victory.

I set up the ATX on the workbench and removed the cover. The box was stuffed full of components, packed in like sardines. Half the board area was a dense forest of electrolytics and power inductors, all pretty typical of an ATX supply. There were two large heatsinks, two 470 µF 200 V electrolytics, one large ferrite transformer and two tiny ones. On the first heatsink were a pair of 2SC4107 high-voltage switching transistors, and off to the side the famous TL494 PWM controller IC. Excellent! This was exactly what I hoped to find!

The presence of two high-voltage bipolar transistors combined with the '494 controller meant I had a push-pull half-bridge converter topology. This is a good basic design approach, well suited to modification. In contrast, some ATX power supplies are designed with a single-ended forward converter topology, driven by a single transistor switch. Single-ended designs can be converted, but such an effort is not covered in this presentation.

All the fuss over "switchers" boils down to high power density and light weight. You rectify 120 V ac into high voltage dc, chop it up at around 33 kHz, step that down through a transformer, then rectify and filter to create the desired dc output. Now that the power transformer is running at an ultrasonic frequency instead of 60 Hz, you can shrink a 300 W transformer from the

18140 Eccles St
Northridge CA 91325
818-885-6960
zz@kj6eo.com
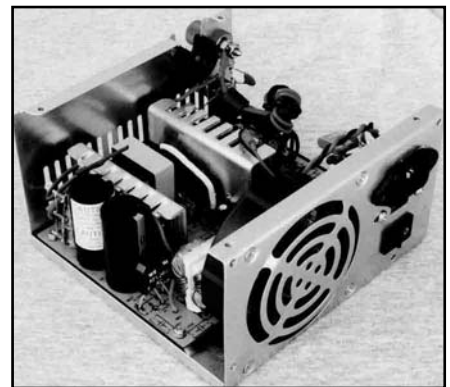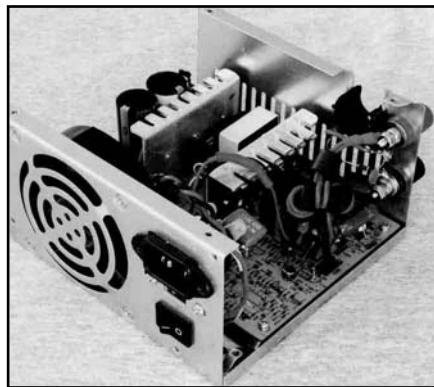
size of a ripe grapefruit down to the size of an apricot. The whole box ends up a lot smaller, and weighs much less.

This paper will review step-by-step all the major circuit functions from one end of the ATX to the other. Keep in mind this is an adaptation and simplification of a common design approach and not a brand new engineering effort from scratch. I changed many things but retained the basic factory circuit topologies. Many topics are given only a cursory review, as exhaustive detail would easily double the length of this presentation and most likely bore you to death in the process. I wanted to keep it inter-esting.

The Journey Begins...

The first glaring item on the circuit board was a horribly burned 1/8 W resistor over a large black patch; it had even unsoldered itself from the copper traces! I replaced it with a fresh 1-W resistor, loaded the 5-V dc output and slowly cranked the Variac up to 120 V ac. The ATX came back to life!

All outputs were now alive and well: 5 V, 3.3 V, 12 V, and –12 V; but at this point I began to notice some problems. There were jumper wires where critical components belonged, and it was beginning to look like most of the major components were over-stressed. Just for starters: the 120-V ac 1-A line rectifier diodes were running at around 2 A, and there were burned spots scattered in a half-dozen places from hot resistors. The heavy output rectifiers were running double the rated current—the claimed 5 V output spec was 25 A and they used 10 A rectifiers. The input RFI/EMI chokes and capacitors were missing. The locations for the small RFI clean-up chokes on all outputs were jumpered over. I realized that I was looking at a low-cost build of a fundamentally good electrical design. Using undersized parts seriously compromised reliability. No wonder it had failed!

In spite of all this, I still saw real potential in this ATX. The circuit topology was fundamentally sound, and the new application would draw low power 99% of the time. On receive, a typical 100 W HF transceiver pulls no more than 2 A; so the ATX will be delivering less than 30 W, an easy ride compared to its previous duty. On transmit, it must deliver 15 to 20 A (maximum) for about 30 % of the time. As long as the switches and rectifiers are robust enough to handle the heavy currents, the only issue left is overheating. If the fan keeps both heatsinks below 160° F we are safe. I speculated that this would be easy to do

since the RX power drain will be a fraction of the original load. Typical ATX duty in a desktop PC is over 150 W continuous; we will not even get close to that level, so thermal stress management is predictably a piece of cake. Later, I reduced the fan speed to whisper quiet. The ATX would still supply 11 A dc continuous hour after hour, and the hottest temperature in the unit was 138° F on the rectifier heatsink. Not bad!

With that encouragement in mind, it was time to press on. The ultimate goal was to convert the ATX to provide a single 14 V dc output, fix the design flaws and go for a solid 20 A key-down output capability.

Thus Began the Reconstruction

I began the project by removing the unnecessary circuit functions: the flea-power flyback oscillator that supplied standby power, the quad LM339 comparator circuit and most significantly, all output rectifiers and filters. This cleared the circuit board of over two-thirds of its components. The multiple-output filter section alone had consumed nearly half the board area. It was surprising how simple things were becoming.

Switch-mode power supply design always begins at the output terminals. You start from the output, then back up and design one section at a time until you arrive at the 120 V ac input. We will follow the agenda in our ATX adventure. It's all very simple. The basic switching regulator (the academic world has named this topology a *buck regulator*) is a low-pass LC filter fed by a pulse train. Put a zero-to-28 V square-wave into the filter and you get out 14 V dc.

Notice in Fig 2A how the low-pass LC filter simply extracts the dc average value of the pulse train. Output ripple is minimal, provided the LC corner frequency is low enough—typically about 1/20th of the switching frequency. A good basic design guideline is—whatever voltage you want out, double it on the input. Since we want about 14 V out, we need 28 V in. The output voltage is directly proportional to the pulse width. Varying the width of the 28 V pulse during the fixed switching period is defined as pulse width modulation. This is also known as duty ratio control, where D is the duty ratio defined as:

$$D = {t_{on}}/{T_s}$$

Typical values are t_{on} = 8 μs and t_s = 15 μs. As D goes from zero to unity, the dc output will go from zero to 28 V. Since the goal is constant output voltage, PWM control is used to overcome line and load disturbances to maintain a steady 14 V dc out. In a perfect world, this conversion is lossless—100% efficient. In this world we can't quite achieve that, but we can achieve 85 % without too much effort.

Practical realities demand that we consider a few other loss factors in this otherwise ideal circuit! Added to the 28 V requirement is the diode drop of the output rectifier. This bumps the voltage requirement up to about 29 V; on top of that are resistive losses in transformer windings, resistive drop in the output choke and primary winding input depression from 120 Hz ripple. All of these contribute to degrade the voltage level going into the LC output filter, demanding a longer duty ratio. Since it is desirable to keep the duty ratio at a nominal 50%, we need higher voltage at the secondary winding than the 28 V we first envisioned. Add all these voltage drops and the requirement for secondary voltage rises to around 31 V—it's all a big numbers game.

The ATX factory design implemented the buck regulator as shown in Fig 2—center-tapped transformer secondary windings into a full wave rectifier, then into the LC filter. Since the buck regulator (Fig 2B) is fed by the main power transformer, we now examine the previous stage.

Power Transformer Rebuild

In my experience, many ATX computer supplies do not supply a solid 12 V dc to the PC motherboard. It is usually low by 0.5 V or so—they really only deliver about 11.5 V dc. Based on this, I strongly suspected the original factory turns ratio would not supply a voltage close enough to 31 V to permit re-using this transformer without modifications.

Prior to dismantling the ATX, I had taken the time to measure the voltage of the various windings and noted that the winding associated with the 12 V output had a peak voltage of only about 25 V. Since this low voltage would widen out the nominal pulse width and degrade the low ac line limit; the existing turns ratio was marginal at best. It was really beginning to look like a new design was essential.

It was time for deep surgery—the power transformer had to come out! After ten minutes of wicking solder, the transformer was on the bench. Continuity check revealed that the factory
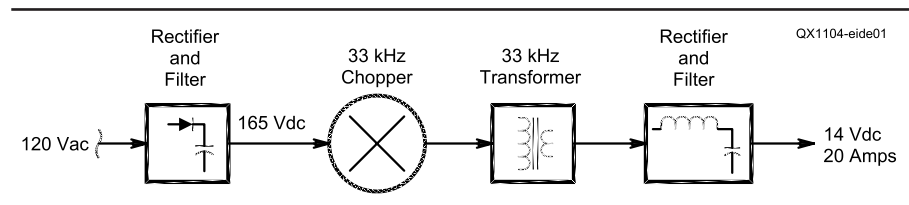
Fig 1—Basic switch-mode power supply functions.

(A)

(B)

(C)
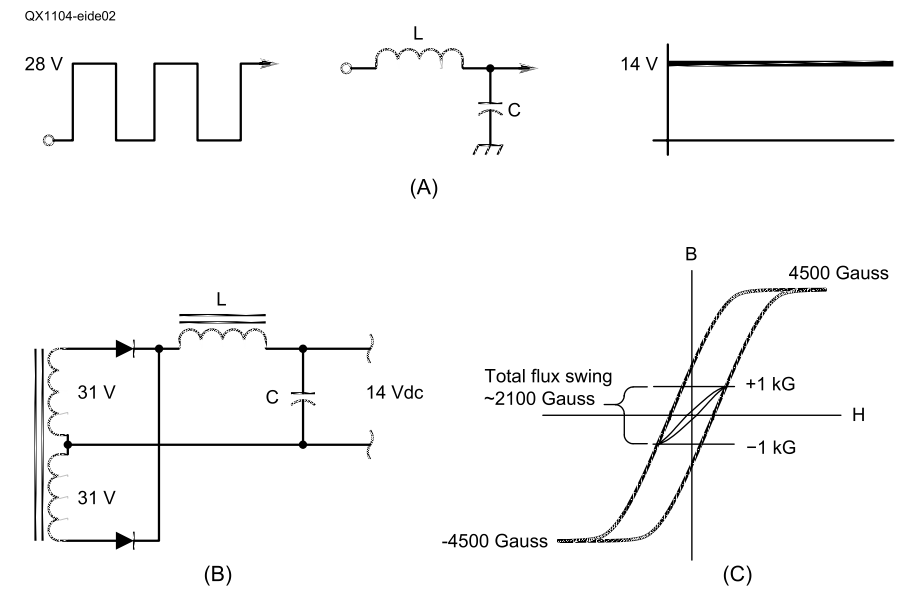
Fig 2—The buck regulator output stage. A shows a buck regulator voltage waveforms, while B shows a transformer coupled buck regulator. C shows the main power transformer flux trajectory.

had wired the secondary windings in a rather bizarre, illogical manner. This was extremely untidy. It was obvious that this arrangement was unsuitable to create the 14 V dc output. I definitely needed a new transformer design.

The next hurdle was to dismantle the transformer. Ferrite cores are routinely cemented together with epoxy, and since epoxy tensile strength evaporates at high temperature, I wrapped the transformer in aluminum foil to avoid noxious fumes and cooked it at 400° F for about an hour. I then peeled off the foil, and as I shoved an Exacto knife between the core pieces, they just fell apart. Then I pushed the ferrite out of the plastic bobbin. After it all cooled down I peeled the off the windings. First I unwound the secondaries, discovered an electrostatic shield of about 3 mil copper foil buried underneath, removed the foil and unwound the primary. Next we move on to calculate the new design details.

Faraday's Law and T1 redesign

Five factors determine transformer operation: voltage, frequency, core cross-sectional area, core material flux capacity and the number of turns. The relation between these is known as Faraday's law and is expressed as:

$$E \cdot \Delta T = N \cdot \text{Core Area} \cdot \Delta B \cdot 1E{-}08$$
(cgs units)

This relationship is used to calculate the new design parameters. The expected primary voltage is 165 V dc. The cross sectional area of the ferrite core center leg was measured at 1.35 cm². I chose B_{max} to be around 1000 Gauss, based on the saturation issues covered in the next few paragraphs. This is about 25% of typical power ferrite saturation flux density of about 4500 Gauss. As a bonus, choosing this low flux level keeps the saturation time out near 40 µs, more than double that of base drive transformer T2. This time disparity is extremely important for safe start up, as we will find out later. Now just re-arrange Faraday's Law to calculate the number of turns:

$$N = (E \cdot \Delta T) / (\text{Core Area} \cdot \Delta B \cdot 1E{-}08),$$

All parameters except N are known: E = 165 V, Core Area = 1.35 cm², ΔT = 9 µs (expected pulse width at full load), $\Delta B = 2 \cdot 1050 = 2100$ Gauss (from push-pull drive).

Solving for N, we get 53 turns on the primary side and with the ratio of 31 to 165, the secondary needs to be 10 turns.

The transformer was rebuilt as follows. The first layer onto the bare bobbin was the primary winding, 53 turns of #22 magnet wire. The heavy-duty high current secondary winding was wound directly over that with 10 turns six-filar #22 to provide the required the 31 V. Finally, 6 turns bifilar #32 were added to provide the 17 V dc for the flea-power housekeeping function. I omitted the copper shield; I saw no reason for it. The heavy enamel insulation is quite robust, with a rated breakdown of over 600 V.

Now with the bobbin rewound, it was time to reassemble the transformer. I coated the pole piece faces with five-minute epoxy, reassembled the core into the bobbin and pressed the ferrite pieces together until I felt the epoxy oozing out between the surfaces. Then I slathered all the edges around the ferrite with more epoxy and bound it together with yellow Mylar tape. The results are shown in Fig 3. After the epoxy set, the measured primary inductance was 16 mH. Perfect—even one-tenth of that inductance would have been enough. Now that the windings of the power transformer have been determined, it is time to examine the primary side switching waveforms and circuit considerations.

Primary Converter Topology

This circuit section (Fig 4) chops up the high voltage dc at 33 kHz and applies it to the primary winding. The primary winding of T1 is driven in a push-pull fashion with an ultrasonic 165 V quasi-square wave. This is nicely accomplished with the venerable half-bridge topology. Standard 120 V ac 60 Hz input power is rectified by a voltage doubler to create raw unregulated +165 V dc and –165 V dc. These two voltages are stored across C1 and C2, and are used as the primary energy reservoir for the ATX supply, implementing a bipolar version of the classic capacitor input filter.

The 33 kHz drive of the transformer primary is as follows: C1 and C2 pin one end of the primary at 0 V. Q1 switches the other end of the primary winding from 0 V up to 165 V for 8 µs; then back to zero for 7 µs. Then Q2 switches on and pulls the primary down to negative 165 V for 8 µs, then back to zero for 7 µs again, and so on. This is repeated at a 33 kHz rate. This push-pull action applies the quasi-square-wave voltage to the power transformer primary, and the transformer lowers it to around 31 V on the secondary. The diodes full-wave rectify the 33 kHz into a 66 kHz pulse train. Then the LC filter extracts the 14 V dc component. During all this, the maximum expected switch currents are just under 4 A, less than 40 % of the device rating. Lots of margin here! I made no changes in the Q1 and Q2 half-bridge circuit topology, the factory approach was fine.

Continuing on our journey through the ATX, we now arrive at the stage that drives the high voltage transistor switches.

Base Drive Circuit

The base drive circuit in conjunction with power transformer T1 is a current–driven variant of the famous Royer Oscillator (circa 1954, see Fig 5), a free-running magnetic oscillator that displays an ingenious combination of proportional base drive, current sense imaging and input-output ground isolation by virtue of transformer action. It also has the unique and valuable property that the free-running mode can be slaved to a PWM drive with just a few parts. When the ATX is first energized and as C1 and C2 charge up, the 330 kΩ collector-base leakage resistors start both Q1 and Q2 conducting. One transistor always conducts ever so slightly more than the other, due to different betas.

For the moment, assume that Q1 conducts more. This forces a small current through the single-turn primary of T2,
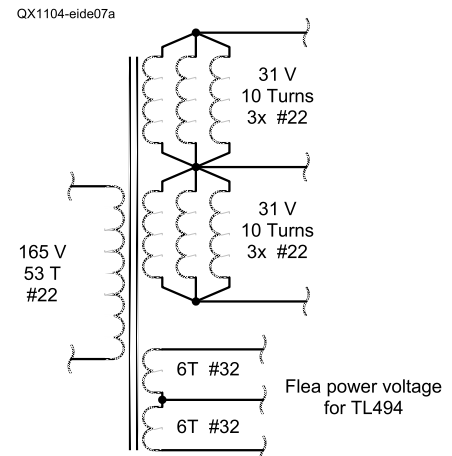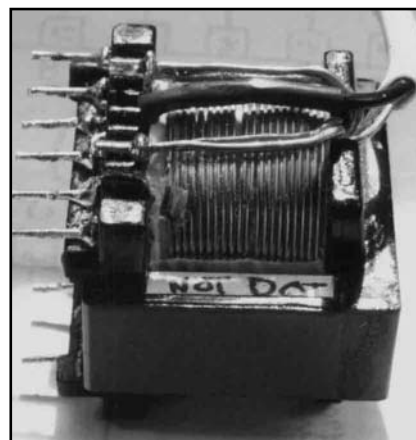
Figure 3—Rewound power transformer T1.

and by transformer action $1/5$ of that current is driven into the base of Q1 providing positive feedback. In a fraction of a microsecond Q1 snaps on and Q2 is cut off. Now there is a full 165 V across the T1 primary and T1's core flux starts its long slow climb. Simultaneously, the magnetic flux in the base drive transformer (BDT) T2 is integrating up to the saturation level. For a properly designed BDT, this should take about 19 µs then the field collapses and all windings reverse their polarity. Once again, from positive feedback, Q2 snaps on and Q1 is cut off, and the flux integrates down the other side of the BH curve, until it saturates and flips again—flip-flop, flip-flop, flip-flop, and so on.

After several dozen cycles the housekeeping voltage is up to a nominal 17 V and the PWM controller takes over. The PWM forces the flux trajectory in little T2 to operate on a minor loop well within the saturation limits of the core material. This is essential. The PWM must run at a shorter on-time, around 13 µs max, to be able to reverse T2's polarity prior to saturation to achieve pulse width control. Normal PWM control will never exceed 13 µs on-time.

Big Trouble On The ATX— Power Transformer Saturation During Start-Up!

During startup it is essential that the saturation time of base drive transformer T2 is considerably shorter than power transformer T1; otherwise if T1 saturates first, Q1 and Q2 will be conducting into a dead short across one of the filter caps and the collector currents will be huge. This is a serious potential failure mode.

This is exactly what happened in my ATX! It was caused by a sloppy design of T2. The factory base drive transformer had way too much flux capacity. Notice in Fig 6 that Q1 and Q2 collector currents hit 35 A peaks from this revolting development! Keep in mind that we are viewing an envelope of dozens of switching cycles, on the order of ten times the normal switching period of the regulator.

Obviously, the main power transformer is saturating before the base drive transformer does. Note that this phenomenon is not the inrush current into C1 and C2. The power transformer must never saturate, at start-up or at any other time.

The large start-up hump occurs as the main power stage oscillates in the free-running mode before the PWM circuit can take control. Once the PWM comes alive, the switching period and duty ratio are controlled by the feedback loop and the current envelope stabilizes at a nice safe low level as shown in Figure 8. The BDT must saturate first, so that the switches are only conducting the magnetizing current of power transformer T1, instead the uncontrolled dead short current of the primary dc resistance during saturation. The main power transformer must not saturate at all. Ever!

The factory BDT was of E-E ferrite core construction with too much core area and twice the number of turns it should have had. Tests revealed saturation time in excess of 70 µs. I had to redesign it for shorter saturation time, target 19 µs, to insure that it would saturate first and prevent huge switch currents. The new saturation time would have to exceed the longest expected on time during TL494 control (around13 µs). I wound the new T2 on a little cheerio-size ferrite toroid from the junkbox (actually, I cut it out of the
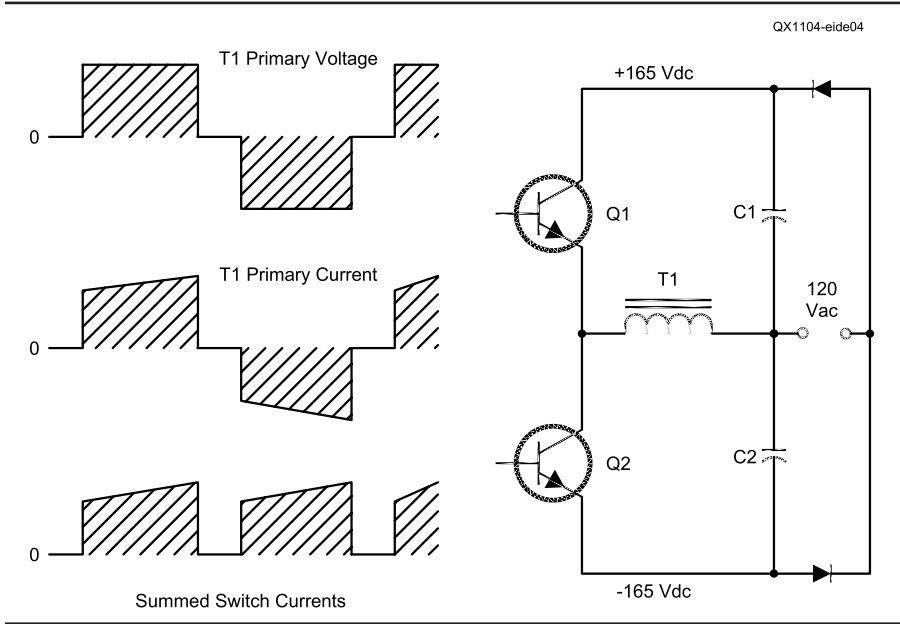
Fig 4—Half-bridge converter topology.
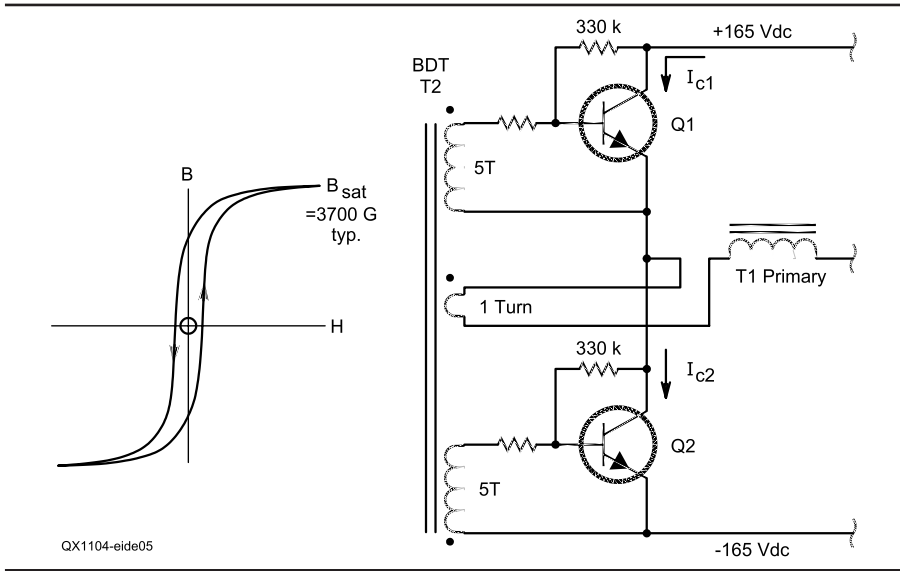
Fig 5—Simplified circuit of the Royer magnetic power oscillator.
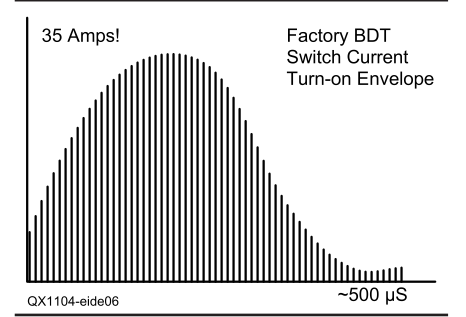
Fig 6—Switch current turn-on envelope.

kitchen telephone!). This new BDT has $\frac{1}{2}$ the core area and $\frac{1}{2}$ the number of turns of the factory BDT. Hence it will saturate in about around $\frac{1}{4}$th the time. The new turns ratio is 18:5:1—perfect for this application. Now for the pertinent calculations.

Once again, we invoke Faraday's law, except this time we calculate the saturation time instead of the turns:

$$\Delta T = N \cdot \text{Core Area} \cdot \Delta B \cdot 1E\text{-}08 / E,$$
all parameters except ΔT are known

E was measured around 2.7 V, Core Area = .118 cm², N = 5 turns, ΔB = 7500 Gauss (measured)

Solving for ΔT, we get about 19 µs. Since the controlled pulse width will never exceed 13 µs, there is adequate margin.

With the new BDT (Fig 7) in place, the start-up current envelope amplitude is much lower. The remaining hump is probably the output filter charging up, and since it was no longer a reliability issue, I did not investigate further.

These fascinating switch current profiles (Fig 8) were observed by installing a pair of 100:1 current-sense transformers in the collectors of Q1 and Q2. The transformers were diode-ORed into a single 100 Ω resistor to sum the alternating images of the switch currents together and produce the well known pulsating input current profile of the buck regulator. Since each core resets its flux against the PRV of the 1N914 diode on a cycle-by-cycle basis, the resultant voltage across the 100 Ω is truly a dc image of the switch currents. This cannot be done with a single core. You need two separate cores to create a dc image as shown in Fig 9. This technique was invented by Dr. Loman Rensink in 1979. When you get the ATX running okay, look for the image of the input current at full load shown in Fig 10.

Note that the DCCT is not required for ATX operation. I used it strictly for diagnostic purposes. The image of the summed switch currents is the heartbeat of the buck regulator and is without question the single most useful waveform to judge the health of regulator operation. Switching times, output load and flux balance of T1 can be instantly evaluated at a glance.

Slaving the Base Drive Transformer to the PWM

Once the power oscillator is free running, the next step is to take control of the switching to achieve pulse width control. Fortunately, the Royer oscillator can easily be slaved to an external

waveform. Synchronizing the BDT to a pulse width modulated waveform simply boils down to overpowering the net magnetizing current in the core, which reverses the magnetic flux trajectory in the core material prior to saturation. During the conduction pulse, the BDT primary current amp-turns equals I_c times one-turn plus the magnetizing current I_m ; subtract $I_c/5$ coming out of the dot times 5 turns. The remaining current is simply I_m referenced to one turn. Since this is a high permeability core, the magnetizing current is just a few percent of the switch current I_c. All of this is just another example of how for any core and coil construction, the core material "sees" only the magnetizing current applied and nothing else.

The primary winding of 18 turns provides an 18-to-1 leverage to stop and reverse the direction of magnetic flux in the core. This causes T2 to operate on a minor BH loop, forcing the magnetic flux to reverse on a cycle-by-cycle basis before the core saturates, taking only about 20 mA to flip the core. This is really slick. Whoever invented this circuit (see Fig 11) is a genius. This is a truly elegant design! I have seen its widespread use in the world of off-line switchers. Only the base drive transformer T2 was modified in this section, the rest of the factory circuit was unchanged.

Incidentally, if you get the phase of
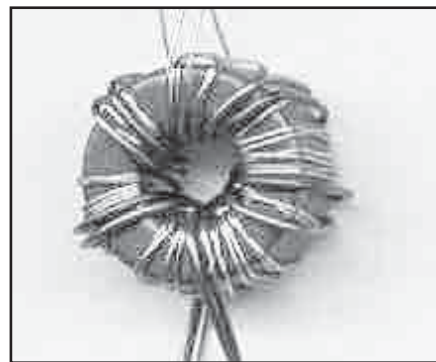
Fig 7—New base drive transformer T2 (around 0.5 inches in diameter).
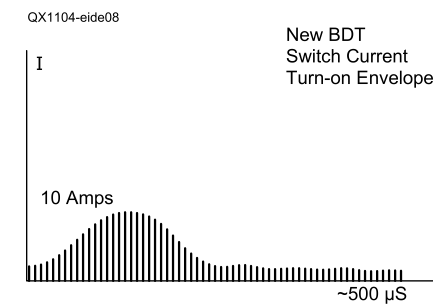
Fig 8—Switch current envelope with new T2.

the one-turn loop reversed, the oscillator will not start and it will not respond to external PWM drive. I found that out! Continuing our journey, the next stage is the PWM controller.

The TL494 PWM Controller IC

The ATX came with the Texas Instruments TL494, a push-pull voltage mode PWM controller that includes two genuine op amps, a voltage reference and digital logic for the A and B drive. A well designed, reliable controller, it has enjoyed widespread use for several decades.

The '494 is used here to generate the active-low PWM pulses that drive the base drive circuit. Please refer to the Fig 17 schematic for more detail. I will not review the detailed timing diagrams of the TL494 beyond that of the basic operation of a generic pulse width modulator. The datasheet can be downloaded from the TI Web site. Intimate details of the controller function are beyond the scope of this presentation.

Fig 12 provides the timing diagram of the pulse width modulator function. A periodic 66 kHz sawtooth ramp is fed into the negative input of the compara-
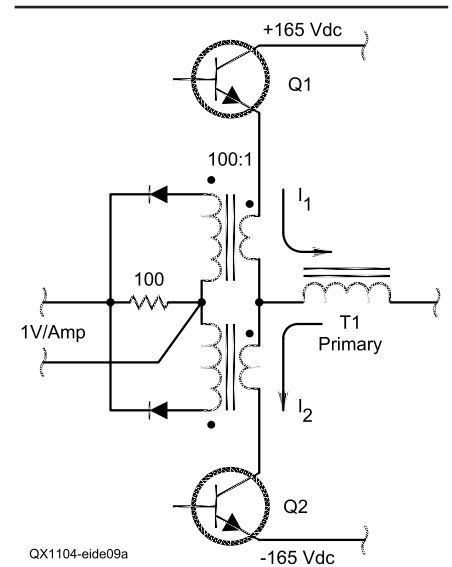
Fig 9—Dc current transformer assembly. Each toroid core is only 375 mils in diameter.
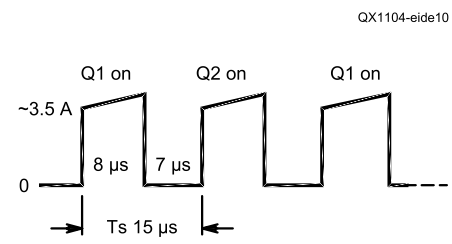
Fig 10—Buck regulator input current image.

tor. Control voltage V_c is fed into the positive input of the comparator. The comparator output goes high when the control voltage exceeds the ramp voltage. The high state corresponds to a conduction pulse of one of the transistor switches Q1 or Q2. Notice how the pulse width decreases as the control voltage V_c increases. This negative sloped PWM introduces a 180° phase inversion into the control loop; and forces the opamp to be configured in the non-inverting topology. This crucial detail is included in the control loop analysis by putting a negative sign in front of the modulator slope k_m.

Internal steering logic alternates the conduction pulses between the A and B outputs to implement the push-pull drive required by the half-bridge power stage. If you want to run the pulse width to zero just pull the comparator output (pin 3) above 3.5 V.

ATX current limiting was implemented by sampling the image of the switch currents that appear at the center tap of T2, and feeding that voltage to the second opamp inside the '494. When it exceeds 5 V, the second op-amp takes over control of the PWM comparator to reduce the duty ratio and limit the output current.

R_{10} adjusts the current limit.

Control Loop Considerations

The objective of the feedback control loop is to maintain constant 13.8 V dc output under all conditions. Since we expect variations of the nominal 120 V ac 60 Hz input power from 100 V ac up to as high as 140 V, and load variations from zero to 20 A maximum, the control loop must adjust the pulse width to any value necessary to keep the output constant. In perfect world, there would be zero error in the 13.8 V output, no matter what the demands were.

Since we are not in a perfect world, we must compromise, but we still can

set a realistic goal for loop performance that will deliver superb dc regulation and fast correction in response to load or line disturbances. That demands high loop gain at dc, and in the frequency range of the expected disturbances, the low audio frequencies. The worst offender is the 120 Hz ripple that appears across T1 primary caused by droop in the main storage caps C1 and C2.

This ripple is about 25 V peak-to-peak at maximum load, and if the feedback control loop did not correct for it, this 120 Hz ripple would be transformed down to 4 V p-p riding on the 13.8 V dc output level.

You can demonstrate this by controlling pin 3 on the TL494 with a dc bench supply set to about 2.5 V, to manually control the pulse width, and the 120 Hz

ripple appears on the output. When you disconnect the clip lead, the control loop automatically takes over and the ripple vanishes!

The ac circuit model for a buck regulator is simply the LC output filter preceded by a linear gain. The linear gain is the product of the input voltage, turns ratio, duty ratio (D) and k_m, the pulse-width-modulator slope. This is the general ac model for any transformer coupled buck regulator. All that one need do is fill in the appropriate constants.

The buck regulator control loop (see Fig 13) falls into the category of a sampled data system, and as such, is limited by Shannon's sampling theorem, which states, among other things, that the maximum bandwidth is one
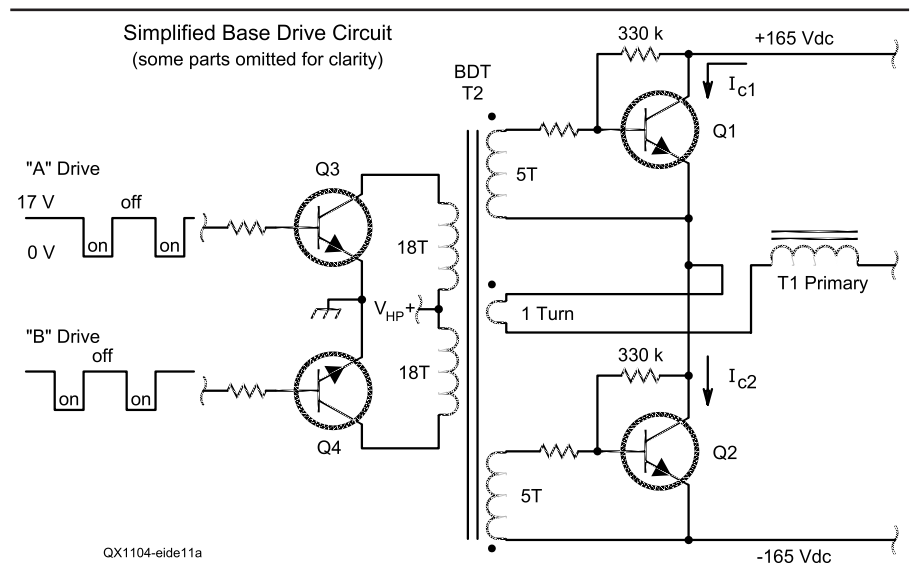
Fig 11—Base drive circuit (simplified).
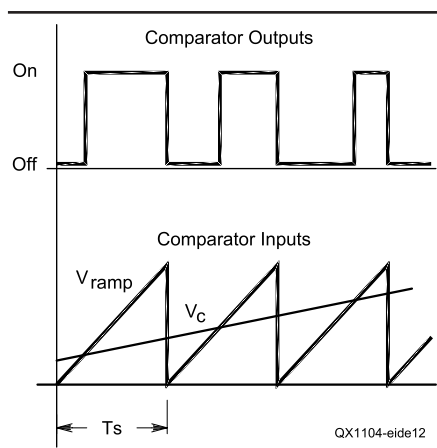
Fig 12—TL494 PWM timing diagram.
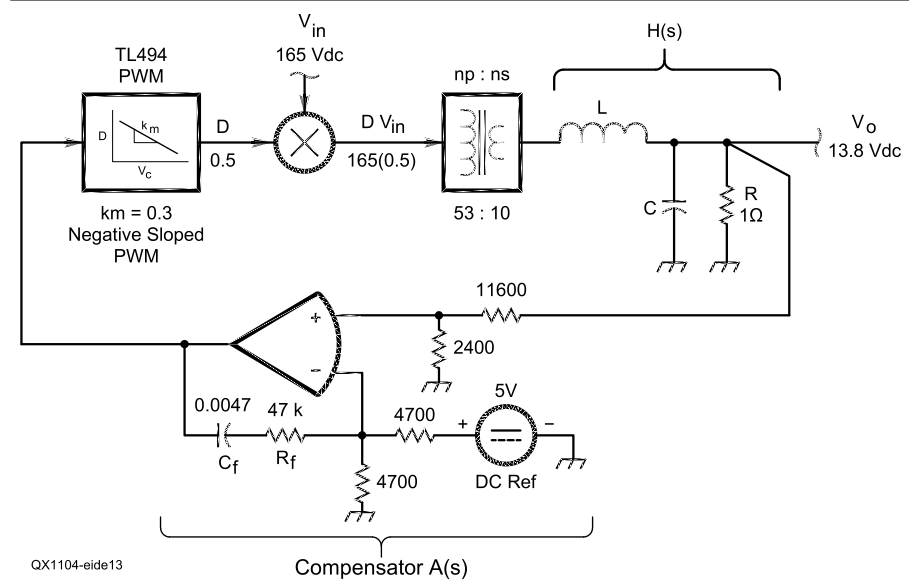
Fig 13—ATX control loop.

half the switching frequency. It all boils down to a rather simple criterion: if there is adequate phase and gain margin below 33 kHz, the loop will be stable. This is exactly the situation we have here.

$V_o / V_c = k_m V_m D N_s / N_p \bullet H(s)$, where
Vc is the control voltage.
$= -.3 \bullet 165 \bullet .5 \bullet 31 / 165 \bullet H(s)$
$= -4.7 \bullet H(s)$

As shown, the transformer-coupled buck regulator forward transfer function has a linear gain of negative 4.7. The 180° phase inversion is caused by a negative PWM slope, followed by an LCR network. The feedback compensator is the op amp in the TL494; configured as an integrator and used to close the loop and deliver optimized performance. The compensator has a zero placed at the corner frequency of the LC filter, yielding superb dc accuracy and excellent transient response, equal to anything a typical HF transceiver can throw at it.

Backing up for a minute, let's take a closer look at the output filter as shown in Fig 14. In the real world you need to add a second smaller LC clean-up filter to clobber nasty little switching spikes that sneak through the winding capacitance of the main filter inductor. The corner frequency of the second LC is chosen to be ten times higher than the first one, to keep total phase shift manageable. It turns out this is quite adequate. Now let's run a computer model of the dual-section output filter including the inductor dc winding resistance and capacitor equivalent series resistance. Note that this is only the two-stage output filter, the negative 4.7 scalar function is added later, in the complete control loop model.

In any control loop, the mere presence of cascaded LC sections always conjures up the specter of instability—360° phase shift can make for a control system nightmare. It didn't happen. Even with the two LC sections, the total phase shift is less than 135° out to the Nyquist limit, and beyond. What a delightful turn of events! It means that stabilizing the loop will be easy. The equivalent series resistances of the filter capacitors limit the ultimate phase shift to considerably less than 360°, so we could close the loop with a simple linear gain if so desired.

The problem is, if we close the loop with a gain of one, there would only be 13 dB of loop gain in the low audio frequencies, where most expected disturbances will occur. The 120 Hz spectral line is the big troublemaker and 13 dB is just not enough. Even though the loop would be stable, its error-correcting performance would be woefully inadequate.

The whole purpose of feedback control is to correct for all disturbances and maintain a controlled output. In this instance, we need a lot more loop gain in the low audio frequencies.

What would be the best approach to conquer this problem? Suppose there were a way to tip up the flat slope of the amplitude curve below the first LC corner at 700 Hz to match the *single-pole* slope above 700 Hz without degrading the phase margin ? We need more gain, and we sure don't need more phase lag to bugger up the phase margin. The solution is one often used in the nether world of classical control theory:

Why Not Artificially Increase The Order Of The System By One?

Yes. Make the compensator an integrator with a well-planted zero. Just plant the zero right on top of the corner frequency of the first LC section right at around 700 Hz, as shown in Fig 15. In this case, the negative 4.7 factor is modeled with an inverting amp even though the second op-amp is not present in the circuit.

Now we have lots of loop gain (> 40 dB) in the low audio region, along with superb dc accuracy, and adequate phase margin out to and beyond one-half the switching frequency. Fig 15 displays a zero-dB crossover at 5600 Hz

with a phase margin of 67°. With the control loop configured as shown, performance is impressive! The closed loop totally tracks out the nasty 120 Hz ripple; with a 20 A load, I measured the output noise and ripple at down around 50 mV. Classical control theory predicts that error correction is just one divided by the loop gain at the frequency in question. So then, since the loop gain is 40 dB at 120 Hz, we divide by 100, which means that 4 V of output ripple is reduced to 40 mV. This loop also exhibits excellent transient performance—keying my Kenwood TS-430 to full power CW steps the load current from 1.5 A to 17 A dc, and the output dips less than 100 mV, with no undershoot or ringing.

This is impressive performance by any standard! Keep in mind all this mathematical machinating is only a computer simulation; to actually measure the loop dynamics would require a $50,000 network analyzer.

Imagine that; $50K in test equipment to test a one dollar junkbox ATX. Only in ham radio!

Output Filter Magnetics Design Et. Al.

The original factory output filter relied on multiple windings on a single toroid core to satisfy the 5 V, 3.3 V, and 12 V output filter inductor requirement. Although I could probably get away with re-using the choke with no changes, I
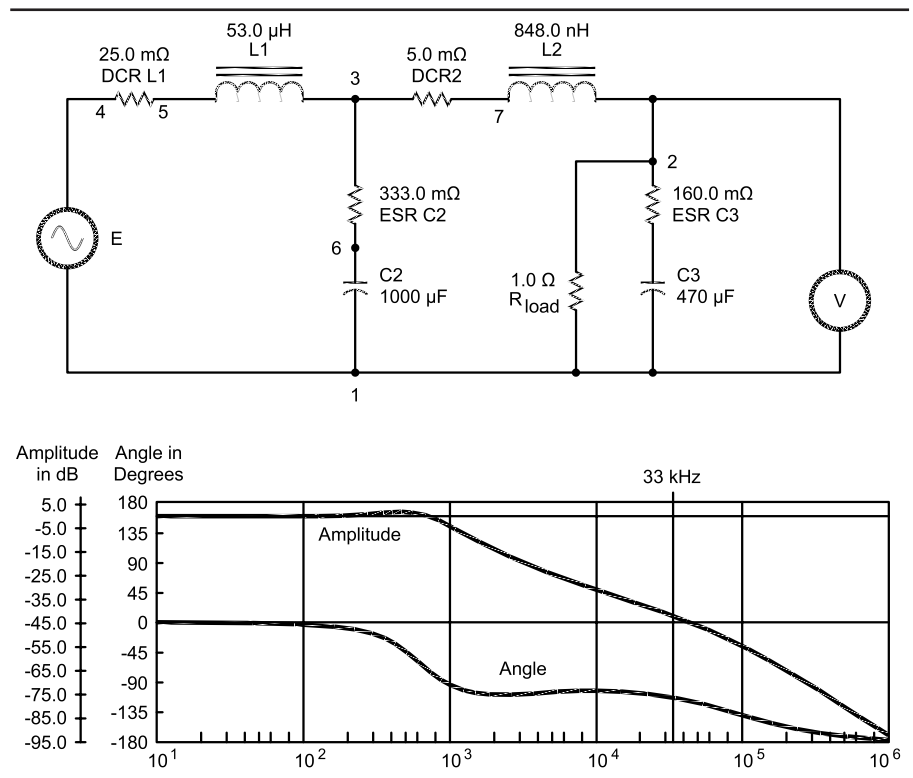
Fig 14—AC model of the dual-section output filter.

wanted a close look at optimizing this part for the new 14 V, 20 A application. Before we get into all the close detail on inductor redesign, it is prudent to discuss the subject of ripple current in the output inductor.

Ripple current occurs when an inductor is subjected to an ac voltage. When it is a square wave, the resulting current is simply a triangular waveform. Following the relation of $V = L \bullet di/dt$, the current is simply the integral of the voltage over time, divided by the inductance value, plus the constant of the dc output current.

Notice in Fig 16 that the output choke ripple current and capacitor ripple current are identical but inverted from one another. The capacitor is forced by Kirchoff's current law to submit to the inverse of the inductor current image in order to maintain the output current at a flat dc value. As we will find out later, it is wise to choose an inductance large enough to keep the ripple current about one-tenth of the maximum expected dc load value.

Of all the constraints that affect the output filter, ripple current is always the most stringent. The casual observer would never expect this, but it turns out that ripple current is an extremely important consideration. It impacts control loop dynamics, inductor core loss, output voltage ripple, minimum load cri-

teria, and determines the selection of the filter capacitor connected to it. Minimum load is *the* criterion that sets the inductance value.

I chose 1.2 A dc as the minimum load value, since that is what my TS-430 pulls on receive.

Ripple current is double the dc minimum load:

$$\Delta I = 1.2 \bullet 2 = 2.4 \text{ A p-p}$$

Now calculate L:

$$L = V \bullet dt/di = 14 \text{ V} \bullet 9 \text{ μs} / 2.4 \text{ A} = 52 \text{ μH}$$

The next step is to see if we can first achieve a 52 μH build on the original factory core at the low current level of 1.5 A, then re-calculate the inductance at the full 20 A load.

First, I cut off all the windings, revealing a yellow-white core; this color coding identified it as a Micrometals powdered iron #26 material, widely used for dc inductor applications. Lots of flux capacity; B_{sat} almost 13 kilogauss, more than three times the saturation level of ferrites. We can make a nice compact choke with this!

After measuring the core dimensions, a quick check of the Micrometals catalog identified the core as T90-26. Micrometals specifies the T90-26 core A_L value at 70 μH per turns squared. Now armed with all core parameters

and material characteristics, it was time to design a new output choke. The objective was to get 52 μH at low current, then recalculate the inductance at a dc bias of 20 A; and see if we still have enough left to have adequate voltage attenuation, and that we haven't moved the LC corner frequency too high to jeopardize the loop stability.

$$L = A_L \bullet N^2 = (31 \text{ turns})^2 \bullet 70 \text{ nH}/T^2 = 53 \text{ μH, this is close enough.}$$

Now we calculate the reduction in inductance at the 20 A level using Amperes law:

$$H_c = .4\pi \bullet N \bullet I_{dc} / l_c$$

where I_{dc} is expected max dc current and l_c is the path length of the core.

$$H_c = 0.4\pi \bullet 31 \text{ turns} \bullet 20 \text{ A} / 5.8 \text{ cm} = 134 \text{ Oersteds}$$

Micrometals' permeability chart reveals that the -26 material retains 24% of original permeability at 134 Oersteds, leaving us only around 13 μH at full load. Two things occur in this condition: the ripple current goes up by a factor of four and the corner frequency moves up by a factor of two. Turns out that both of these are of little consequence and can be easily tolerated. Although not presented here, the loop dynamics change slightly but there is still adequate phase margin, and the ripple current, although
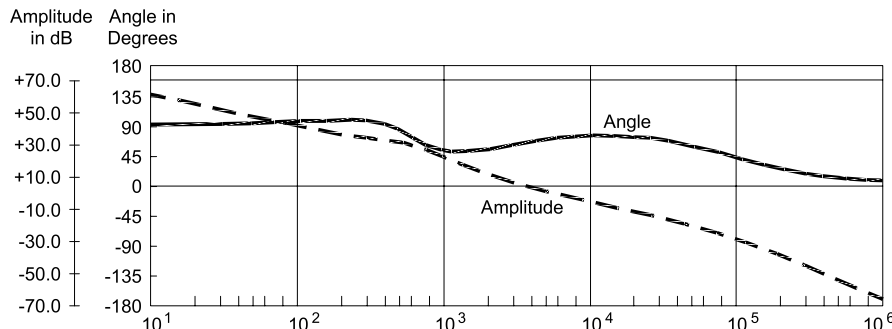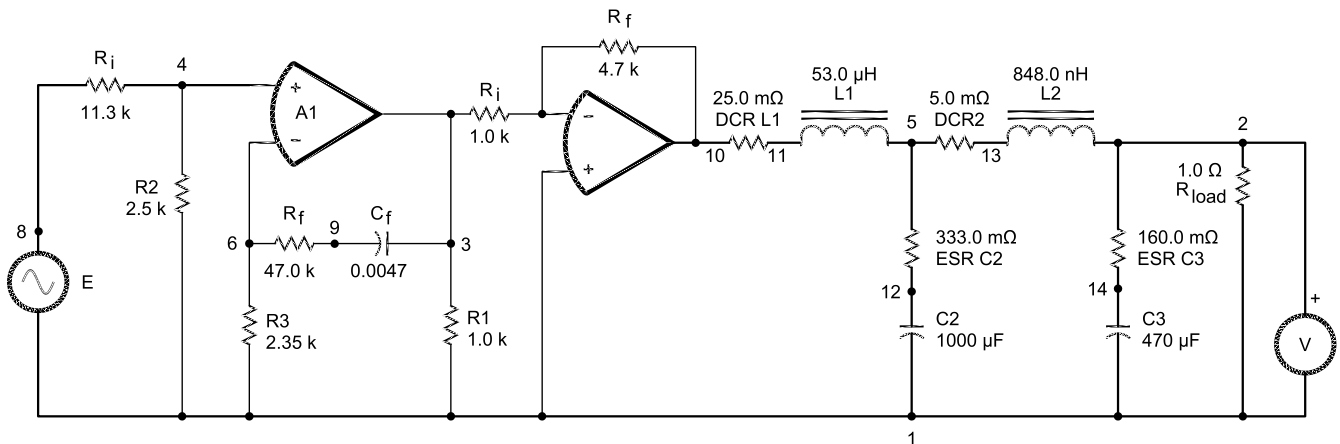
Fig 15—Total control loop model.

now up to 10 A p-p occurs only for short duty cycles and will not cause excessive heating in the windings or the ESR of the output filter capacitor.

Output Filter Capacitor Considerations

The next step is to select the output capacitor—why not just re-use the 1000 µF capacitor left over from the 12 V output circuit ? For a quick check of attenuation—53 µH and 1000 µF yield a corner frequency of around 700 Hz. This is greater than 90 times lower than the 66 kHz PRF—we've got plenty of voltage attenuation. Even at maximum load the corner frequency will go up to 1400 Hz, and this is still low enough. A typical ripple current spec for this part (1000 µF / 16 V dc) in aluminum electrolytic is about 2 A p-p. I had to wonder—can the output capacitor with-stand the 2.4 A p-p ripple current ? Well, so far it has.
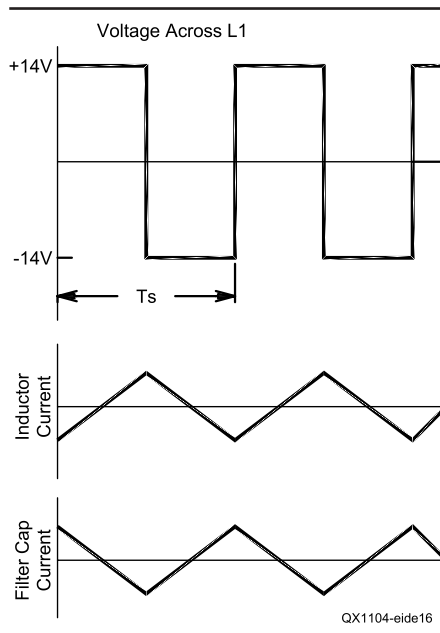
Fig 16—Ripple current in the output filter components.
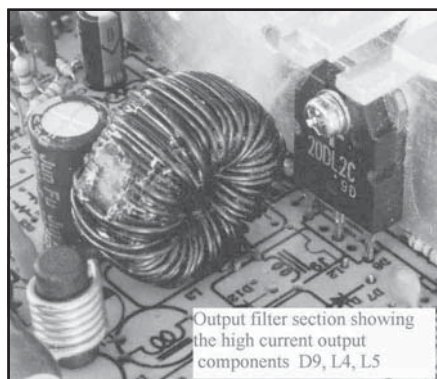
Fig 17—Output filter section.
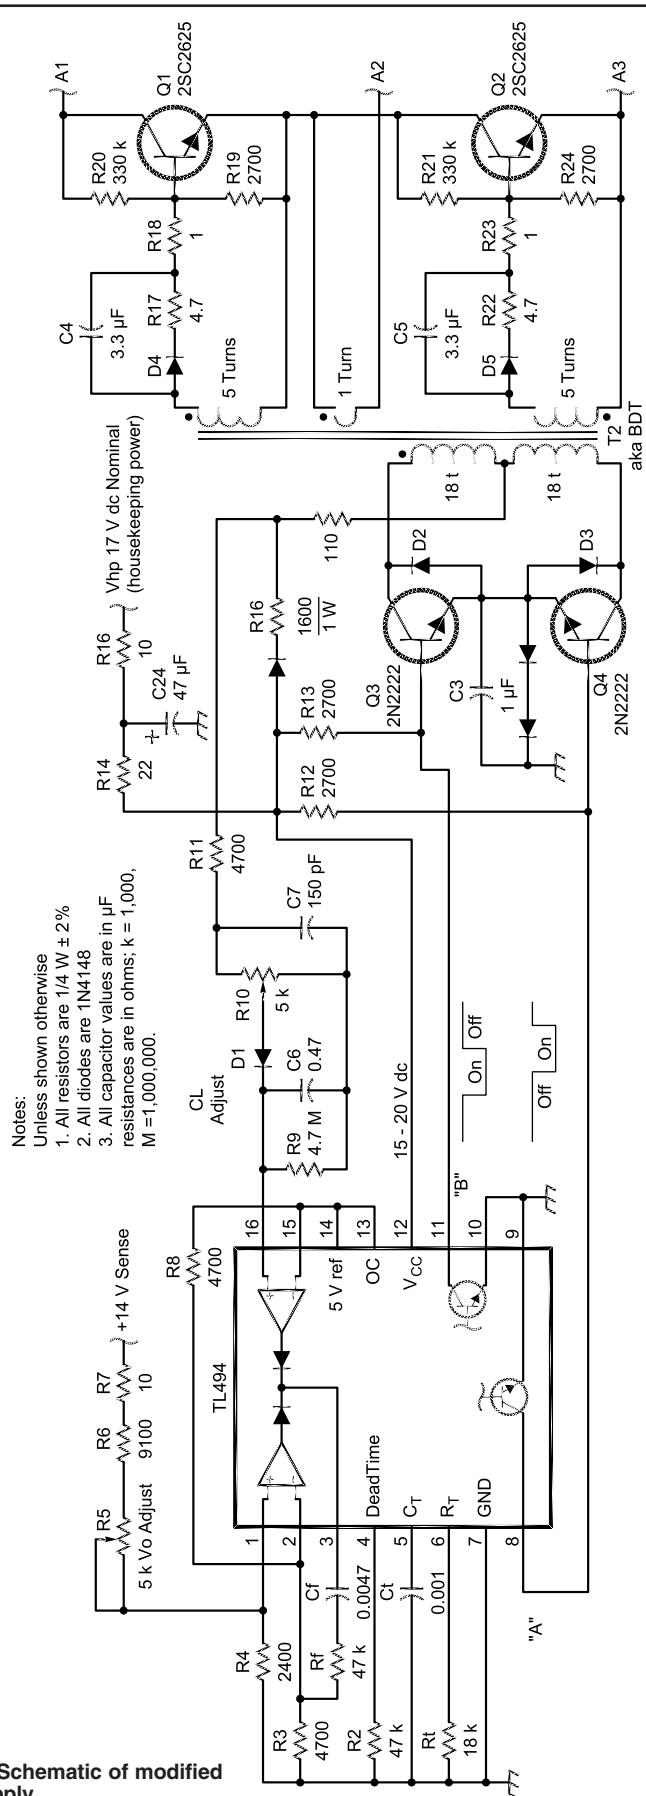
Fig 18A—Schematic of modified power supply.
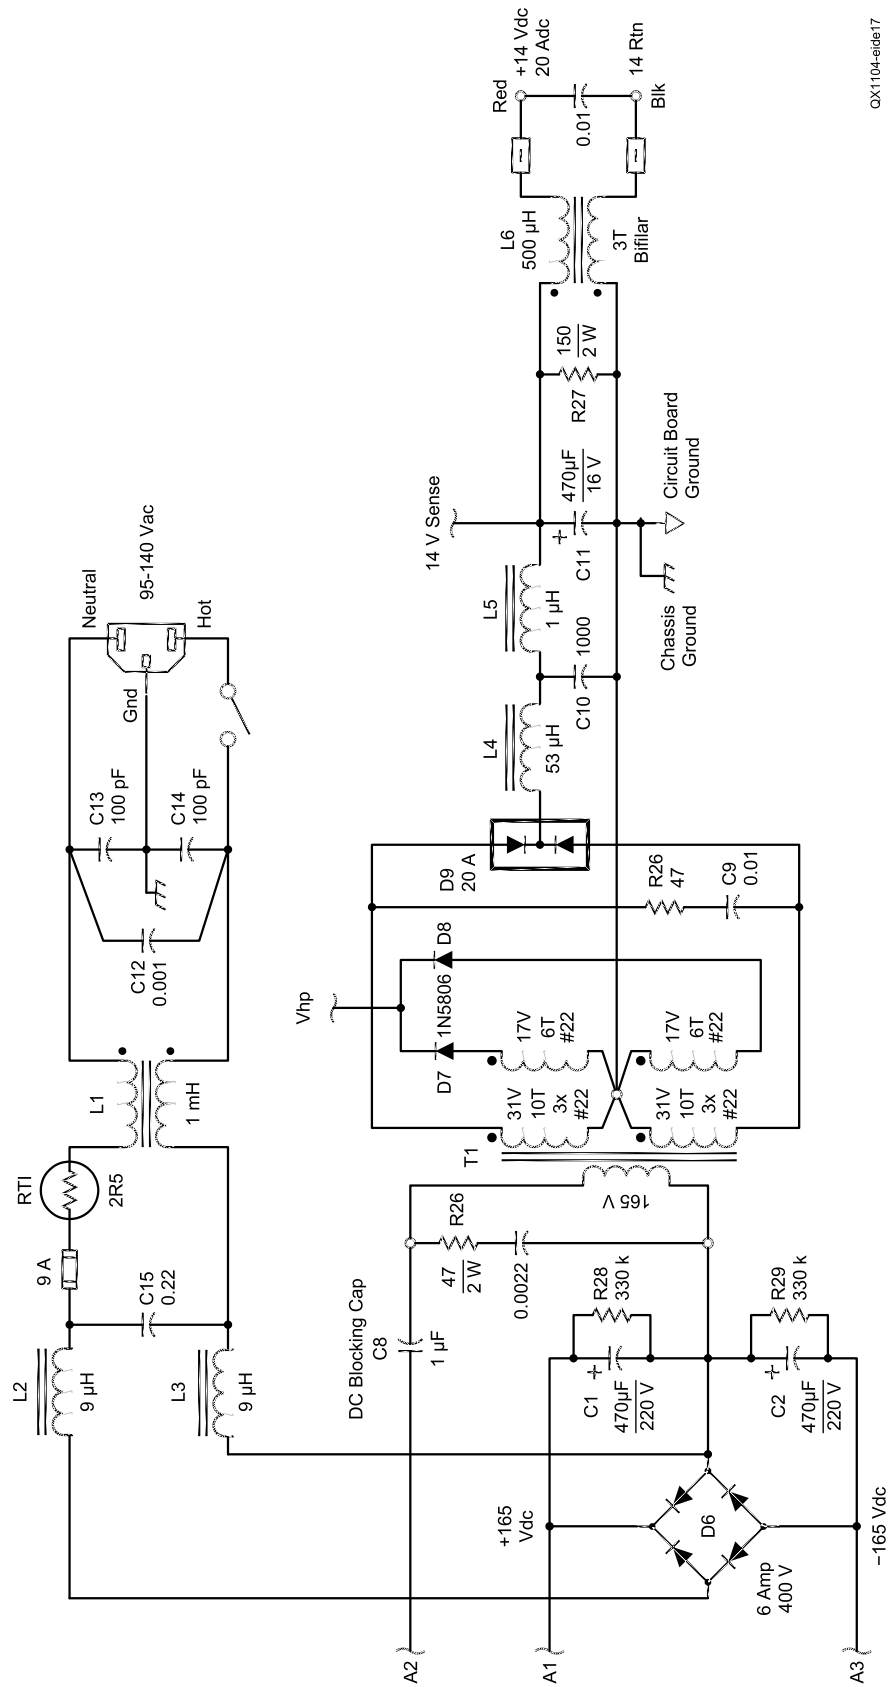
Several hundred hours of operation and it doesn't even get warm. Nor has it failed. It seems to me that excessive ripple current would overheat the capacitor, however, we are not experiencing any overheating. It does not even get warm to the touch. So it looks like the 2.4 A of ripple current will cause no real harm.

For those interested, you can view the image of the inductor ripple current that passes through the ESR of the big filter capacitor by simply looking at the miniscule ripple voltage across it! Set your Scope to ac coupled at about 100 mV per division and clip the probe across the capacitor. Now load the ATX to 3 to 4 A, and observe a triangular voltage waveform shaped much like those in Fig. 16, plus the inevitable switching spikes that lurk inside all switching supplies. The amplitude will be about 100 mV or so. Now step the load up to 18 to 20 A. The peak-to-peak amplitude will triple, demonstrating how the choke drops in value under heavy dc current; causing the ripple current to increase by the same factor. As stated before, this phenomenon will cause no real degradation in ATX performance.

Test Results

The output voltage was adjusted to 13.8 V dc via R5. The output maintained a constant 13.8 V as the ac input was cranked up from 95 to 140 V ac. The lowest I could go and still maintain regulation at 20 A load was 95 V ac. This is more than adequate and will outperform a linear supply "hands down." Below 95 V, the duty ratio maxed out, and 120 Hz ripple began to show up on the output. Try to pull 20 A out of your RS-20 at 95 V ac input and see what you get! Noise and ripple at a 20 A load was less than 100 mV p-p. Transient response to a step load was less than 100 mV change for a step load of 1.5 to 17 A. In my opinion, this performance is far and away more than adequate for any modern 100 W HF transceiver.

Thermal Overload Considerations

Even though the electrical design of the ATX can accommodate a 20 A load, it will overheat if loaded to 20 A continuously due to heat transfer limitations. The heatsinks and fan are inadequate to get the heat out of the box. To make the fan noise tolerable, I slowed it down to whisper-quiet (7 V dc) with a series resistor. The factory heatsinks are low-cost aluminum stampings that require lots of air movement to dissipate heat. Extruded heatsinks are great performers, but cost a lot more than el-cheapo stampings. My ATX had the stampings. Bench testing the ATX with

Fig 18B—Schematic of modified power supply.

an 11 A dc resistive load resulted in the output rectifier heatsink rising to and stabilizing at 140° F. The output rectifier heatsink is the ATX hot spot! Running all day long at 11 A, the temperature never exceeded 140° F. The typical 100 W HF Transceiver draws roughly 10 A average during SSB TX. Current drain varies in the 3 to 20 A range, but averages typically around 10 A. With the possible exception of 75-meter phone nets, no one engages the PTT button 100% of the time, hour after hour. So then, for the sake of argument, let's say the most extreme load is 10 A, 50% of the time—5 minutes TX, 5 minutes RX. This is still a lighter load than a continuous 11 A load, and consequently, the ATX will not overheat in normal transceiver use.

I tested three 100 W HF transceivers—TS-120, TS-430 and Corsair II—for their current drain while running full power into a 50 Ω dummy load. Even for a windbag like me with speech processing engaged, the average current level hovered around 10 A. I placed thermocouples inside the ATX to monitor the most critical components' temperatures and found that the heatsink for the output rectifiers exhibited the greatest temperature rise.

Even during effusive monologues, it never got above 120° F. This is quite safe. I have not tested the ATX at 15 A full time—it will never be stressed to that level while powering a 100 W SSB transmitter. Admittedly, this is an area that needs more consideration, but the ATX performance as described is totally adequate to meet my original design objective. At a later date I may add a simple thermostat circuit to switch the fan to high speed if the box gets too hot.

EMI / RFI Victory

During the long weeks of bench testing, I would monitor 40 meters on my HF radio across the room. The receiver would howl and screech from the harmonics and spurs radiated by the ATX. The original factory circuit had no RFI/EMI filters on the ac input. As mentioned earlier, choke locations were jumpered over with bare wires. First I installed separate 11 μH chokes (these were the former 12 V output chokes) in each leg of the ac just before the 60 Hz mains rectifier, and shunted a .47 μF capacitor across the ac line. There was no improvement, so on a whim I added a junkbox common-mode choke in the ac line. The RF hash nearly vanished. I was astounded! I never expected such dramatic improvement. I added another common-mode choke in the 14 V output lines and all RF hash disappeared.

I'm still testing, but methinks we got this one clobbered. There is a lot more to common-mode filtering than meets the eye. I also added ferrite beads right at the red and black binding posts for good measure. In recent days I have had occasion to inspect several commercial grade off-line switch-mode power supplies and they had two cascaded common-mode filter stages on the ac input. Looks like someone else has been down this road before.

Failure Department

The biggest pothole on the ATX Victory Road was Schottky rectifier failure. With the first snap-on start up with the new transformer, the main ac fuse blew. The Schottky output rectifiers were dead shorted, and as a bonus Q1 and Q2 were blown open circuit. It turned out that re-using the original factory Schottkys for the main high current output rectifiers over-stressed their reverse voltage rating. For a 5 V output configuration, a common PRV value for high current Schottky rectifiers is around 40 V. With our 31 V secondary windings the diodes must withstand 62 V (plus a bit more for inevitable switching spikes). A 100 V rating would be reasonable. I substituted 20 A, 200 V fast-recovery silicon rectifiers from my junkbox. They exhibit a bit more forward voltage drop, but this failure mode has not re-occurred.

Conclusions and Recommendations

This ATX journey has been an engrossing adventure. What a kick it was to take a crumb from the table and make a unique project out of it. This has been one of my most fun projects in a long, long time. There were so many surprises! I never expected that the control loop would be so fast that the 120 Hz ripple is effectively cancelled out. I was really impressed that the output voltage drop from a 17 A step load was less than 100 mV. Offhand, one would not expect such excellent performance from an LC output filter with small energy storage. In fact, the diminutive size of the two output capacitors really amazed me. I still don't see how such small parts can deliver such superb filtering. This is a marvelous demonstration of a fast, wideband control loop doing all the hard work, proving that you don't need a lot of stored energy in the output filter to respond to fast and heavy current demands.

Another pleasant surprise had to do with the housekeeping power. The factory design powered the TL494 with a nominal 15 to 20 V dc from flea power windings on the main power transformer. Those windings were simply peak rectified and run into the '494 with a minimal RC decoupling filter. I would have expected logic disruption from switching spikes getting into the PWM controller, but that has not happened. Since it worked, I left it alone. I am most impressed with the stout behavior of the TL494 in this regard.

As to RFI / EMI issues, the calming effect of common-mode chokes on both the inputs and outputs was truly amazing, I never would have expected such superb filtering. This phenomenon demands more research!

The overall ATX performance is most impressive and exceeds the requirements of any modern 100-W HF transceiver. This is not the ultimate power supply nor was it ever intended to be. This project was a compromise! I started with a junkbox ATX, incorporated a few minor 5-minute mods and ended up with a switching supply that will hold its own with any commercially available unit. The ATX transient response far and away outperforms my Astron SS-18 and SS-30 boxes.

What do I recommend ? Well let's see—if you have an interest in pursuing an ATX conversion, I recommend finding a candidate with the earmarks of quality engineering and construction. Look for a unit with the TL494 controller, 470 μF main filter capacitors, an ac rectifier bridge of at least 4 A or so, and RFI filter chokes on inputs and outputs. If they bothered to put in a common-mode choke on the ac input, then the rest of the ATX will be of good quality. Look for 2SC4107 or 2SC2625 switching transistors. If you have gotten this far, you are on the right track.

Dead ATX boxes can be obtained for pennies from computer repair shops or hamfests. Three or four of them should provide all the parts you need. Don't forget safety. There are lethal voltages on the circuit board! You can really get "fried" with 340 V dc! I taped a 5-mm plastic sheet to the foil side to limit the possibility of electrocution. I only got bit a couple of times while poking around!

It is also prudent to run the TL494 from a bench supply until the loop is stable and working correctly before you snap on the ac switch! I found a Variac to be invaluable.

Keep on homebrewin'!

Old 'ZZ wound coils in the magnetics lab of a now-defunct aerospace company for many long years, exploring hidden worlds, desperately seeking to part the veil of ignorance held by a coercive force. He is now retired with his memories, his cigars, and his shortwave radio. You can contact the author at the addresses shown at the beginning of this article, or find him on 7198.6 kHz during most daylight hours in Southern California.

□□

A New Approach to Modulating the Class E AM Transmitter

Homebrew a high performance modulator using switch-mode technology. There's a movement going on to populate the AM bands using low-cost technologies. Get on board with a tall *signal for* short *money.*

By Bob LaFrance, N9NEO

T he method commonly used to modulate a class E/F AM transmitter requires a dc power supply and an open-loop audio modulator. While this method is easily implemented and can produce quality audio, it is an expensive proposition. The expense resides in the power supply transformer and filtering capacitors. A large 60-Hz transformer is used to provide isolation and reduce the voltage to a level required by the class-E RF deck. The reduced voltage is then rectified and filtered. The filtering requirements are quite stringent in that any 60- or 120-Hz components of power supply ripple on the dc

21 Moorland Drive
Uxbridge, MA01569
yzordderrex@verizon.net

bus will be observed in the demodulated audio. This filtering requirement translates into cost, volume and weight issues related to both the filter capacitors and power transformer.

The new approach is to use high frequency switching technology with a closed feedback loop. The need for a large 60-Hz transformer and a large bank of filter capacitors disappears when this approach is chosen. In addition to these benefits, the modulating voltage is easily selectable by simple turns ratio adjustment controlled by the builder. The high frequency transformer can be wound in a step-down or step-up configuration. Class-E mobile operation can be easily implemented with a step-up design.

The full bridge, phase shifted, ZVS resonant supply topology was chosen to implement the modulation in a

300-W (PEP) push-pull class-F transmitter. This supply can easily be scaled to provide full legal limit power capability. The cost of putting a high frequency modulator together should be

QX1104-Lafrance01
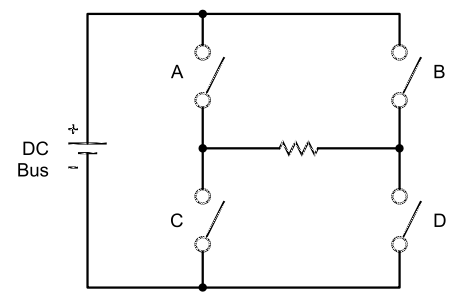
Fig 1—Schematic of simplified full-bridge switching arrangement.

near $100, regardless of the power level chosen. Fortunately for us there exists a class of power supply control chips required to implement this modulation strategy, so our job becomes that much easier.

Theory of Operation

Fig 1 depicts a full bridge switching arrangement. It consists of a dc source and four switches wired in an H configuration with a resistor load. A square-wave output can be easily produced across the resistor. If we turn on switches A and D we will apply the full dc voltage across the resistor in one direction. If we turn on switches B and C we will apply the full dc voltage across the resistor in the opposite direction. If we toggle between these two states we will effectively produce an alternating square wave voltage across the resistor. Both of these states are called active states in that power is being transferred to the load. There are also two states the switches may be in where no power is being transferred to the load. These states are with switches A and B both on, or switches C and D both on. These are called zero states. It is with a precise combination of active states and zero states that we are able to create the necessary modulating voltages to

drive the class-E transmitter.

It is worth mentioning that switches A and C or switches B and D should never be on at the same time. These states are to be avoided, as they will cause a large current to flow through the pair and will certainly destroy the switches. The particular
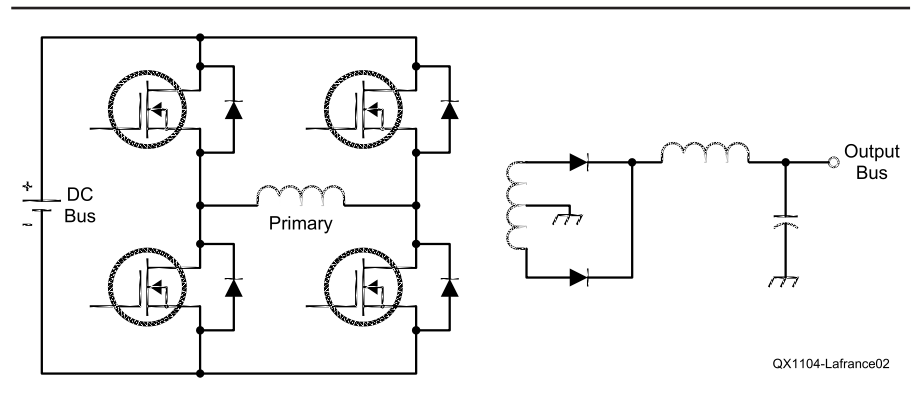
Fig 2—Full bridge with MOSFET switches and transformer primary load.
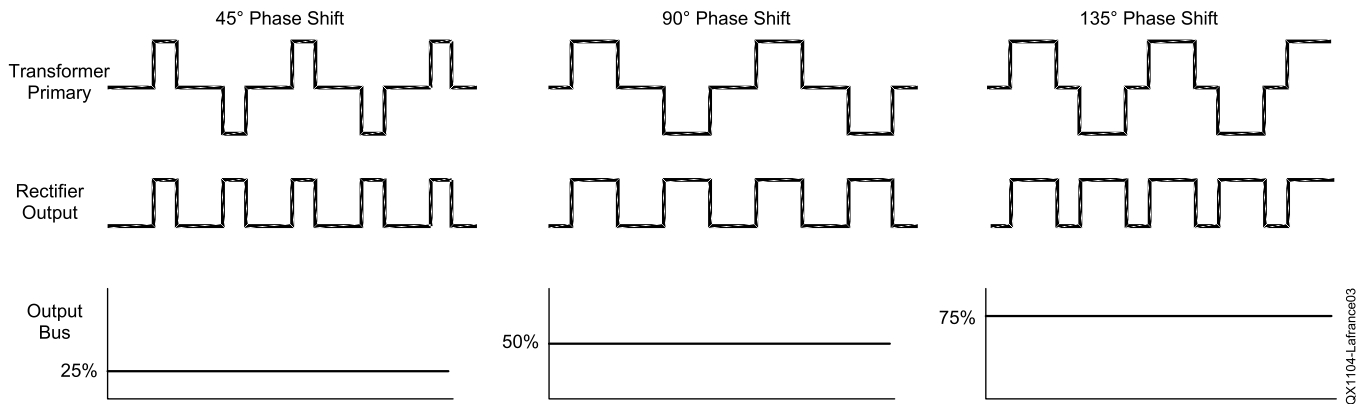
Fig 3—Key voltage waveforms.
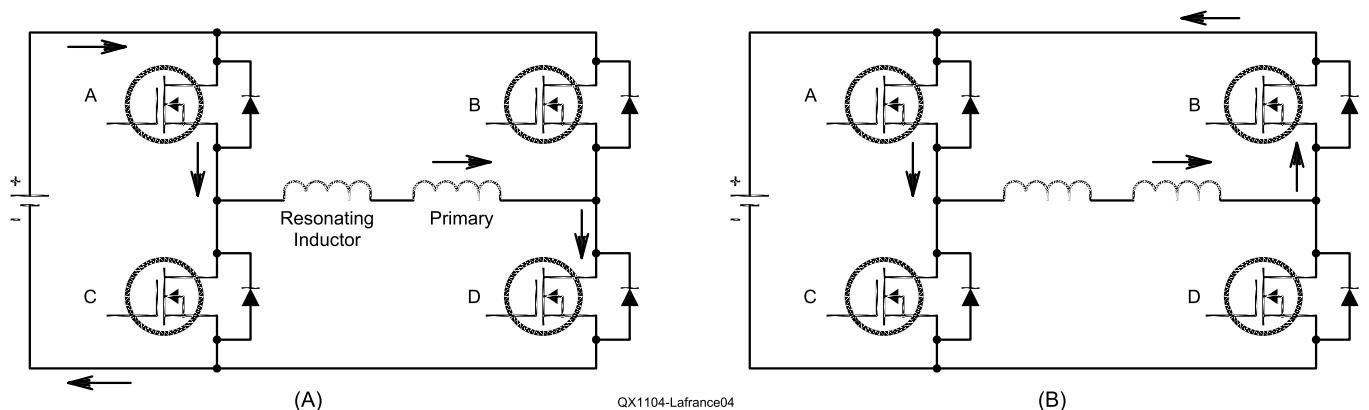
Fig 4—Current directions. 4A—in active state, 4B in zero state.

implementation chosen uses pulse transformers to drive the MOSFETs. This method will ensure these unwanted states are avoided.

Armed with both active states and zero states, we have all that is necessary to create any voltage we desire. The manner in which we create a particular voltage is by changing the phase relationship between the switch pairs. As previously described we can create a square wave of maximum amplitude by toggling between the two active states. Similarly we can create zero voltage across the resistor by toggling between the two zero states. Notice that when we toggle between the two active states the switching between the left pair of switches and the right pair of switches is 180° out of phase. When we toggle between the two zero states the phase relationship between the two pairs of switches is zero. If we control the switching so that the phase relationship between the two switch pairs is 90°, we will still see the full bus voltage across the resistor, but for only half of the time. We are effectively controlling the voltage across the load by controlling the phase relationship between the switch pairs. It is this method of controlling the voltage that gives us the name Full Bridge Phase Shifted power supply topology.

In Fig 2 we have replaced the resistor load with the primary of a transformer, the switches have been replaced with MOSFETs and the secondary side of the transformer is rectified and filtered before being connected to the load. Notice the diodes connected anti-parallel to the MOSFETs. These diodes, called body diodes, are intrinsic to the MOSFET manufacturing process. That is, we get them for free. You will soon see that these diodes are very useful.

Fig 3 shows the voltage waveforms observed across the transformer primary, rectifier output and filter output with varying phase relationships. It should now be evident that with rectification and filtering we can create any voltage that we wish.

There are some interesting subtleties concerning switching from one state to another which you should be aware of. Zero Voltage Switched (ZVS) converters, much like a class-E RF deck, depend on the voltage across the transistor being zero when switched on. This condition reduces switching losses and makes high frequency operation possible. Another benefit of the resonant switching is a reduction of spurious noise. This noise can cause havoc with control circuits.

Refer to Fig 4A. Assume that we are in an active state with MOSFETs A and D on, and current is flowing in the primary path as shown. We will enter a zero state by turning MOSFET D off. Since the transformer path is of an inductive nature, the current through it must continue to flow. The current that was flowing through MOSFET D will now commutate to the body diode across MOSFET B. It is the body diodes that clamp the transformer voltage to the dc bus. Without them MOSFET destruction is assured. There are parasitic capacitances related to both the transformer and MOSFETs B and D that must be charged before the diode across MOSFET B will become forward biased and conduct. The resonating inductor helps to store the energy required to charge these capacitances. A finite amount of time is required to charge the parasitic capacitances and this time varies based upon the current level. Resonant switching requires that we
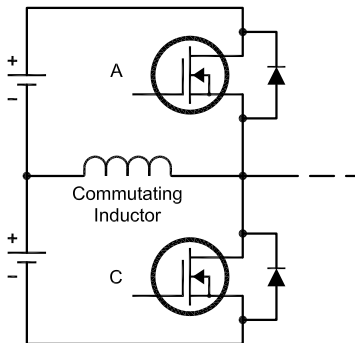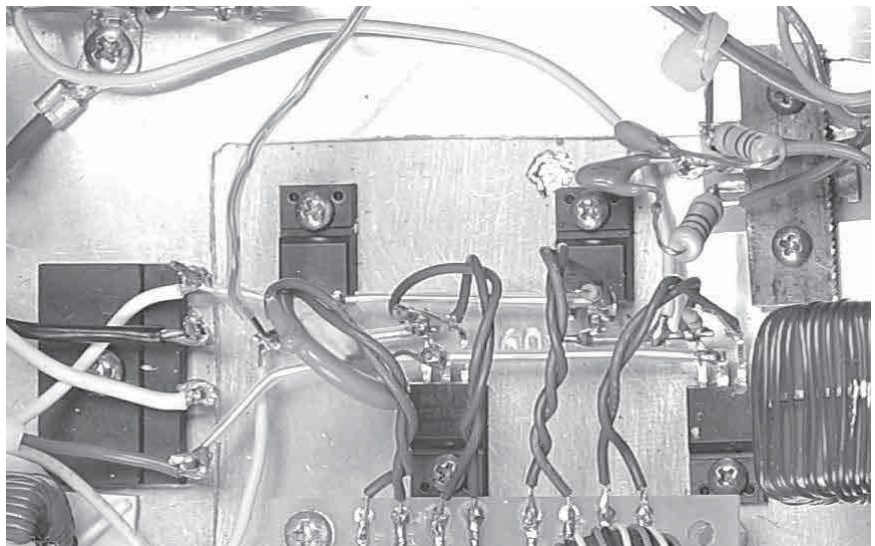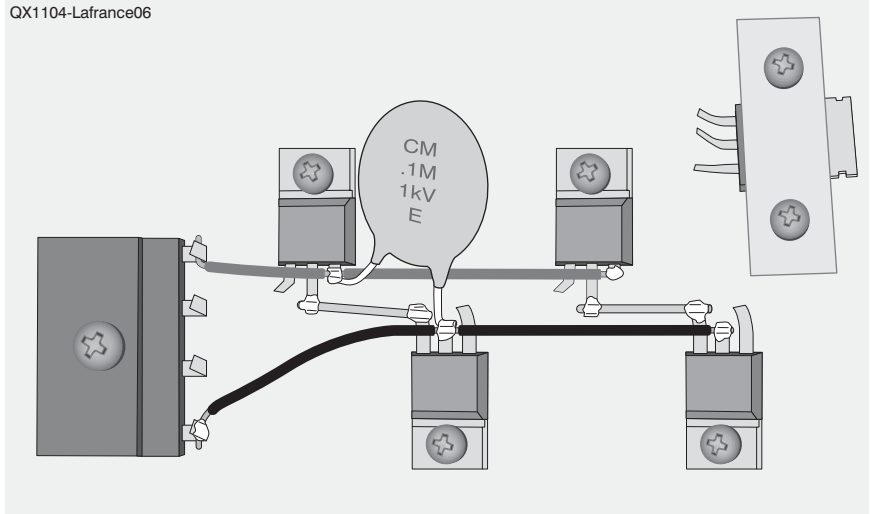
Fig 5—Circuit with auxiliary commutating inductor.

Fig 6—Bus structure layout (drawn by Robin LaFrance).
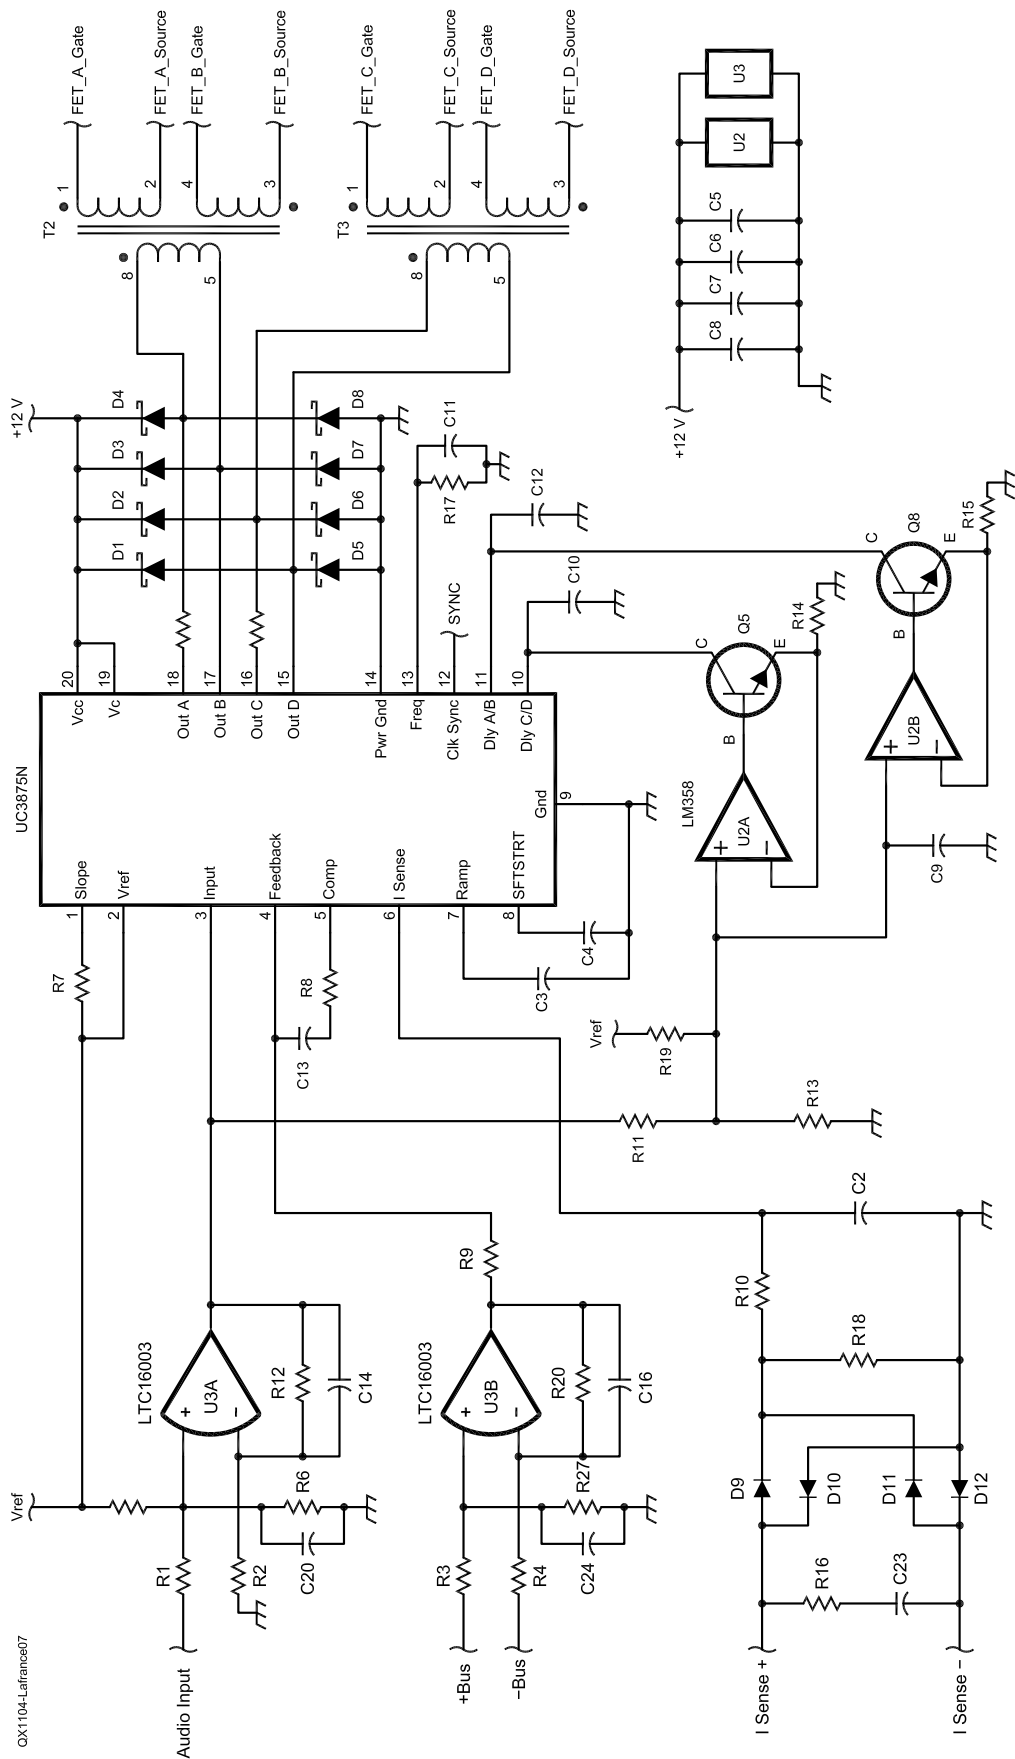
Fig 7—Schematic, resonant modulator control board.

wait until the diode across MOSFET B is conducting, and the voltage across the MOSFET is very small, before we turn on MOSFET B.

Attention must be paid to these timing considerations, called commutation delays, in order to successfully implement resonant switching. Luck is once again on our side and this is a relatively easy task to manage. Since the RF deck appears as a resistive load to the modulator, the output current of the modulator is then proportional to the audio input command. We can then directly modulate the commutation delays with the audio signal. Assume that we are now in a zero state with MOSFETs A and B on—no energy is being transferred to the output. The current will circulate along the path of MOSFET A, diode B, and the transformer primary. When it is time to enter the other active state MOSFET A will turn off, and after the programmed commutation delay, MOSFET C will turn on. MOSFETs B and C are now supplying energy to the output. Notice that the zero states are entered when MOSFET D or MOSFET B turns off, and active states are entered when MOSFETs A or C are turned off.

There is always more primary current circulating in the transformer when the active state is left, rather than when it is entered. The commutation delays can be adjusted for this differ-

QX1104-Lafrance08

L1 - 100µH, T130-15 core

L2 40µH, T108-15 core

C1 - 1 µF, 200 V, Poly

C15, C17 - 2000 µF/100V

Q1, Q2, Q3, Q4 - IRFI 644

Pulse XFMRS - FT87-75, 14 turns Trifilar

C25 - 0.2 µF, 500 V Ceramic

L3 - 2µH, T68-2, 15 turns

D14, D15 - 600V/16A Ultrafast

L4 - 300µH 8A, T130-15, two cores stacked

C19 - 1 µF, Poly

R24 - 6 Ohm

C18 - 0.1 µF ceramic

R25, R26 - 220 Ω/2 W

C21, C22 - 0.001 µF, 500 V ceramic
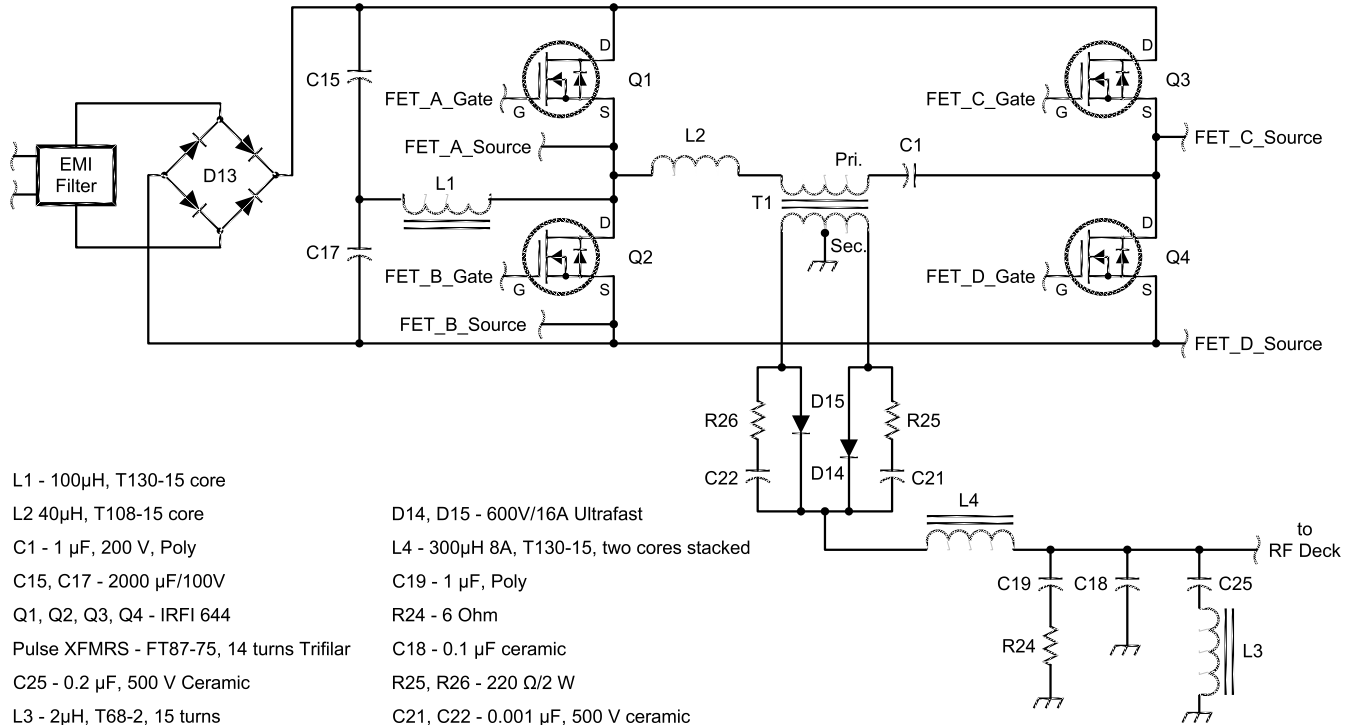
Fig 8—Schematic, resonant modulator power section.
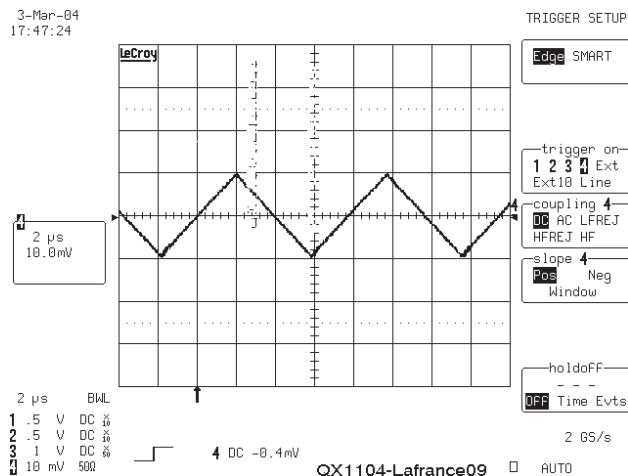
Fig 9—Commutating inductor current, 1A/div.
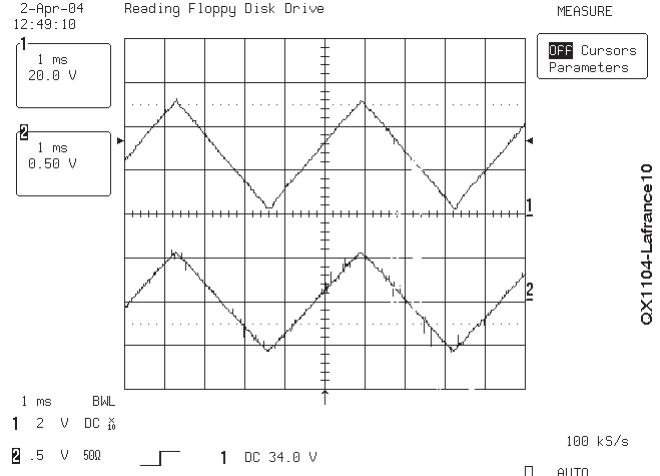
Fig 10—Triangle wave—upper is output, lower is reference.
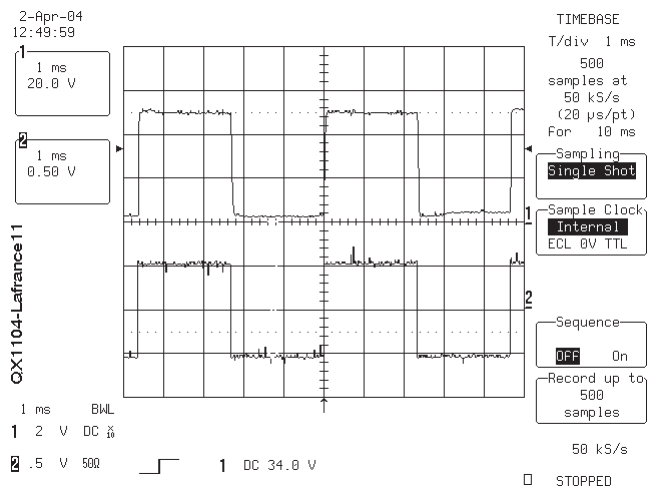
Fig 11—Square wave—upper is output, lower is reference.
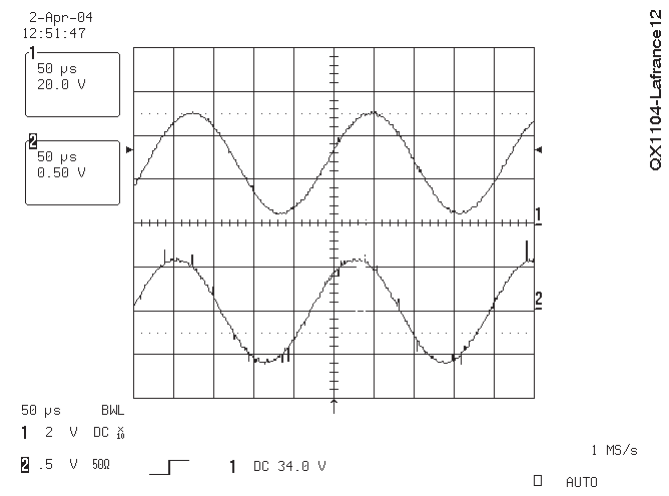
Fig 12—4 kHz sine wave—upper is output, lower is reference.

Fig 13—Spectrum analyzer output. Switching noise at 270 kHz is 67 db down.

ence in resonating current. The commercially available control chips are capable of programming these delays independently. At very low current levels there may not be enough energy stored in the magnetic components to properly commutate the switches.

The argument has been made that the switching losses will be low at low currents so it becomes less of a concern. I disagree with this philosophy—as noise is as much a concern as MOSFET losses, especially at the relatively low powers associated with ham radio communications. This problem can be easily remedied by splitting the bus into two series capacitors. A commutating inductor is then connected between the center point of these two capacitors and the drain of MOSFET C, as shown in Fig 5. This inductor insures there will always be enough energy to commutate the passive to active transitions. The inductor should be chosen to provide about 1 A of commutation current.

The output filter design is worth a few words. The filter capacitance is chosen based on both the audio modulating frequency and the input impedance of the RF deck. It should be understood that the modulator is not capable of driving the voltage down on the output capacitor. It is the input impedance of the RF deck that discharges the capacitor. The time constant of the output capacitor and RF deck impedance must be fast enough to prevent distortion at the bottom of the modulating sine wave. A time constant between 10 μs and 20 μs, based on the RF deck input impedance, seems to provide a good balance between audio quality and switching noise. If switching noise is allowed to pass to the RF deck, there will be unwanted sidebands produced at mul-

tiples of the switching frequency. The filter inductor is then chosen to roll off the switching noise. The inductor should be chosen so that the filter poles are placed near 10 kHz. Do not place the filter poles any lower. This, along with a strategically placed resistor in series with the output filter capacitor, will prevent the filter phase angle from approaching 180° of phase shift and causing modulator oscillations.

A series trap filter on the modulator output, tuned to the switching frequency, is used to further reduce the switching noise. Noise levels approaching 68 dB down at the switching frequency have been observed with the trap in use.

Circuit Description

U3-A is a differential audio amplifier that takes a line-level input and

provides both gain and offset to drive the modulator input pin. The gain is set to about 2.5, while the offset is set to 2.5 V by R15. A line level audio input signal will drive the modulator input 0 to 5 V. Other strategies are easily implemented and may prove more attractive based upon your particular audio processing methods. C14 and C20 provide a low pass filter function to roll off higher frequencies and provide immunity to switching noise.

U3-B is a differential amplifier that provides output voltage feedback to the controller. The pulse width modulator uses this feedback signal to keep the output voltage tracking the audio command signal. It is precisely this function which filters out the 120-Hz ripple voltage on the dc bus. The gain of this circuit determines the output voltage swing. I've chosen a 0-5 V input signal

to command a 0-60 V modulator output. The gain necessary is then 60/5 or 12. C16 and C24 provide switching noise immunity. Care must be taken that these capacitors are not so large as to roll the feedback off early. If we prematurely lose the feedback, the modulator will increase the output voltage to compensate for the reduced gain—we will end up with an unwanted high frequency boost in the audio. While both the input and output circuits are at earth ground and theoretically a simple voltage divider would have worked as well, I've chosen to use a differential amplifier in an effort to eliminate any potential noise problems. R8, R9, and C13 provide compensation to prevent the system from oscillating.

U2-A, U2-B, Q5, Q6 and the associated circuitry modulate the commutation delay set pins on the controller. At low output currents, the delays are longer to provide the MOSFETs enough time to commutate.

Component selection

Control chip

I've chosen to use the Unitrode/TI UC3875N power supply controller as it has integrated the necessary functions, with the least amount of peripheral components. One precaution that should be taken is to insure that the amplifiers interfacing with the control chips can operate to near zero on both inputs and outputs. This will minimize any distortion at the bottom of the modulating envelope. Choose a rail-to-rail type of op amp. In an effort to make the design suitable for mobile operation, I've chosen to use a single 12 V supply.

MOSFETs

When operating from 120 V mains, choose a MOSFET with a minimum voltage rating of 250 V. The current rating shouldn't be larger than necessary. A smaller part will have less capacitance and be easier to commutate. I've had success running the International rectifier IRFI644 parts. These parts are in a plastic overmold type package and require no insulating pad. The IRFI644 part should do full legal power using a heatsink with less than 300 square inches area.

Rectifiers

Choose an ultra-fast rectifier with a peak voltage rating of 4 times the maximum modulator output voltage. The rectifier current rating should be near the maximum modulator output current. The rectifier losses can be substantial, so heatsinking is necessary.

Transformer

The transformer design is straightforward, and almost any platform will do. I've run the ETD-44 platform using 3F3 material, and toroids using both J material, and an EMI core of unknown permeability. Keep the flux swing to a
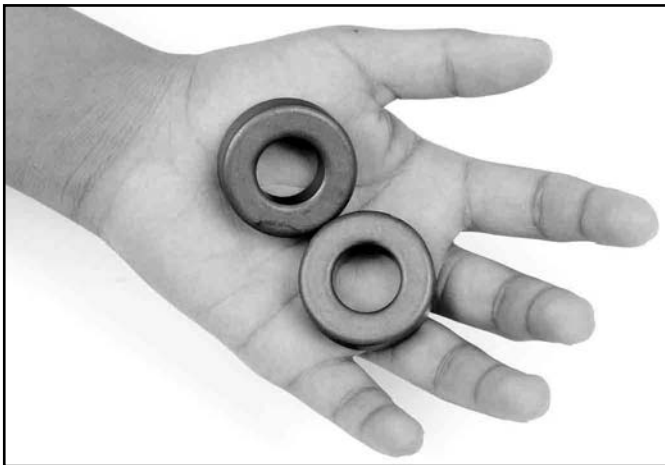
Photo 1—Two little cores will support full PEP out.
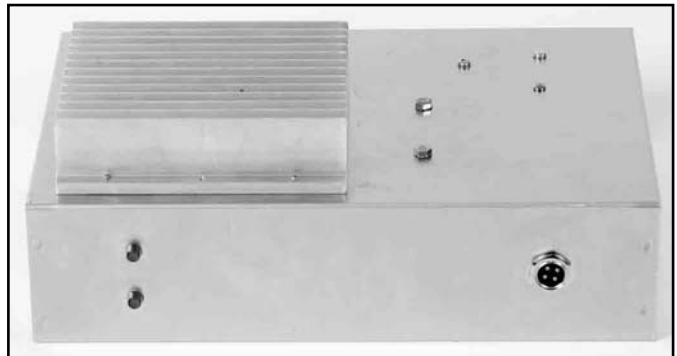
Photo 2—Front panel showing indicators for bus voltage and modulation.
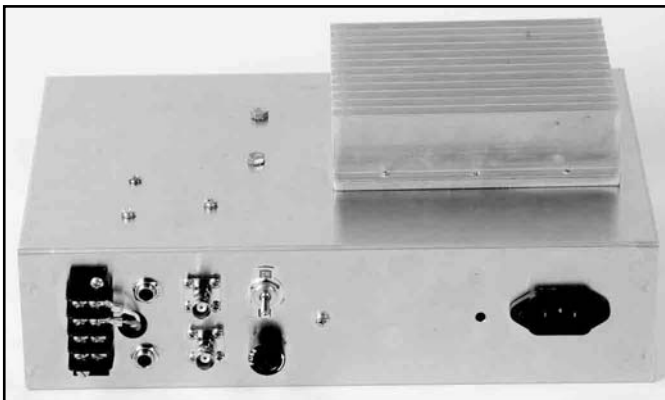
Photo 3—Back view showing audio input and output jacks for nanocompressor input and output. The BNC connector is the output to RF deck and monitor scope. The wires are to TB for transmitter keying.
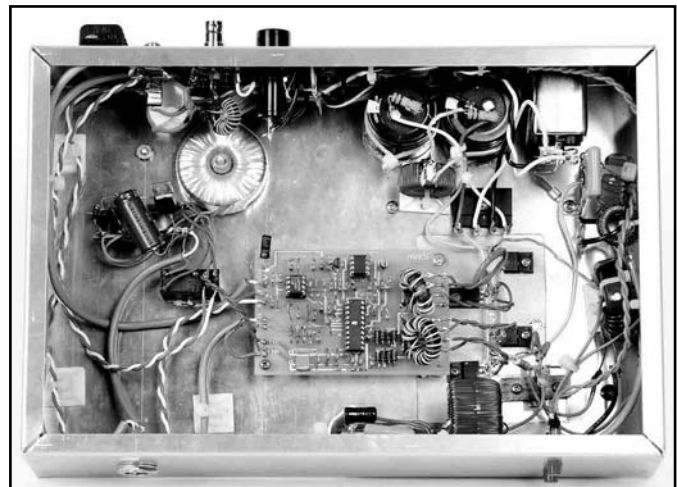
Photo 4—Bottom view. The main transformer is on the right side with the white dot on top

reasonable level and all is well. Some designers will insert a small gap in the core to prevent *flux-walking*. This occurs when there is a dc offset in the core due to an asymmetrical driving of the core. The same can be accomplished with a small film cap in series with the primary. Adjust the turns ratio so that the peak output voltage can be reached under low line conditions with maximum bus ripple voltage. Skin effect requires us to pay attention to the wire gauge we choose when operating at high switching frequencies. One solution is to use multiple strands of magnet wire. Another solution is to wind with a large gauge wire if the inner core of the wire is left unused. A 16- or 14-gauge wire should work well at full power.

Resonating inductor and commutation inductor

Gapped ferrite or powder iron rings would be appropriate. Choose the inductances so that under low current conditions there is enough energy stored to resonate the MOSFET capacitances. I've chosen 40 µH for the resonating inductor, and 100 µH for the commutating inductor. Some experimenting into the merits of eliminating the resonating inductor in favor of two commutating inductors may be worthwhile.

Current sense transformer

A small ferrite ring is the best platform for the current sensor. The exact ratio is not critical, but it must be known that the smaller the turns ratio, the larger will be the loss in the terminating resistor. The CT is used for protection only, not for controlling the output voltage. This makes the design of the part less critical.

Pulse drive transformers

A ferrite ring will make a good pulse transformer. I've chosen a small ring core with 14 turns wound trifilar. The leakage inductance needs to be minimized, and a trifilar winding style is appropriate. A small gauge telephone type wire might be good for this. Insure that the insulation is in good condition.

Bus capacitors

The value of capacitance is only critical in that there must be enough voltage headroom to drive the transformer primary. Choose a value of capacitance that will give you 10 to 20 V of ripple. A full legal limit modulator will need about 5000 µf. A series parallel combination of 4700 µf 100 V capacitors may work well.

Conclusion

The intent of this paper is to introduce the technology and provide a basic understanding of the full bridge topology as applied to modulating the class-E transmitter. Each transmitter design will have its own unique set of requirements, but the principles outlined here apply to any set of operating conditions. I've had good reports on the audio quality when using an inexpensive homebrew audio chain while driving a class F push-pull RF deck. There is an abundance of material available on this particular topology, as it is very popular with the power supply crowd. I hope this platform will provide a basis for continued experimentation with both modulation strategies and resonant MOSFET transmitters in general.

A Method of Measuring Phase Noise in Oscillators

How to dig deep for phase noise measurements with an easy-to-find test setup

By Kjell Karlsen, LA2NI

I have tried to find a suitable method to measure phase noise in oscillators over many years as a constructor of HF and VHF amateur equipment. The equipment normally used by professionals for this measurement is too expensive for many amateurs. I have an elderly spectrum analyzer at home, and a more modern one that I can borrow from my workplace; however, the dynamic range is too limited on both of them. When it comes to measuring noise only a few hundred hertz away from the carrier with a level 130-150 dB higher, a dynamic range of 80-90 dB is not useful. This is a problem that's very hard to solve.

PAØJOZ wrote an article[1] for the Jan/Feb 1999 issue of *QEX* describing a method for measuring phase noise using a known-clean crystal oscillator or a good signal generator as a reference. He mixes this with the oscillator under test by phase locking them both to the same frequency. This results in a zero IF (baseband) signal consisting of only the noise. With filters of 1 kHz, 10 kHz and 100 kHz we can measure the phase noise with

[1]Notes appear on page 59.

Petersborggt. 6
Tromso, Norway 9009
la2ni@online.no

quite good results. He also suggests using a PC sound card as a spectrum analyzer for this measurement, but I had not found any software that could do the job until recently.[2] The measuring bandwidth must be 1 to 10 Hz and the dynamic range must be in the 140-160 dB range. Without usable software for a PC sound card analyzer, I had to find another means to do phase measurements.

Making It Happen With the Help of a Parts Bin

One day when I was looking through some old parts in one of my scrap boxes, I found a crystal filter on 38 MHz with a bandwidth of 5.4 kHz and 50-Ω impedance in and out. Could this be used to expand the dynamic range of my old Hewlett-Packard 8558B Spectrum Analyzer?

I have a Marconi 2019A signal generator with specifications that should make it a candidate for phase-noise measurement. On 90 MHz Marconi claims –110 dBc at 1 kHz and –135 dBc (measured in a bandwidth of 1 Hz) at 10 kHz. Could this be measured by using the crystal filter to get rid of the carrier and letting the noise go through to be displayed on the analyzer? See Fig 1.

I calibrated the setup by injecting the signal in the middle of the filter passband, adjusted the attenuator on the generator to have an indication at the top of the display. Then the frequency was moved up until there was a drop in level of 3 dB and then back to the top edge of the filter. This frequency is used as the reference, and then the generator was moved 10 kHz further up. The output reading dropped around 90 dB. The output of the generator was then increased 30 dB to the maximum output of the 2019A. I could now see the carrier near the –60 dB line on the analyzer and 10 to 16 kHz lower; the noise through the filter is visible at –75 dB. See Fig 2. As the output has been increased by 30 dB after calibration, the noise is –105 dB below the carrier. This is measured in a bandwidth of 1 kHz, and by subtracting 30 dB we get the noise in 1 Hz BW. Marconi claims –135 dBc/Hz for the 2019A. After that I moved the carrier down so that the slope of the carrier just straddles the noise. Now we can see the noise from
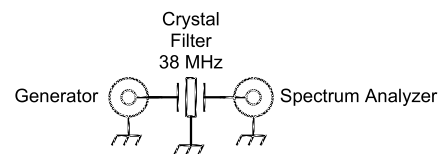
Fig 1—Initial noise measurement using Marconi generator with carrier 10 kHz from filter center frequency.
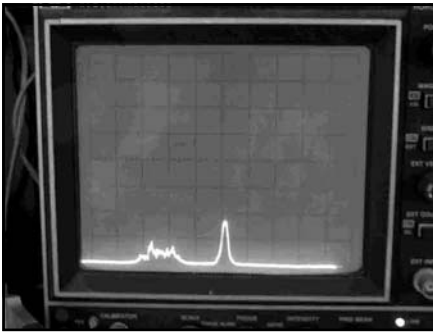
Fig 2—Measurement with 38-MHz filter and carrier moved up 10 kHz from the filter passband, level increased by 30 dB. The noise is at –105 dBc/kHz (–135 dBc/Hz) from 10 to 16 kHz, recorded using a HP 8558B spectrum analyzer with vertical resolution of 10 dB/div. The reference line is at –30 dBc. Horizontal resolution is 5 kHz/div.
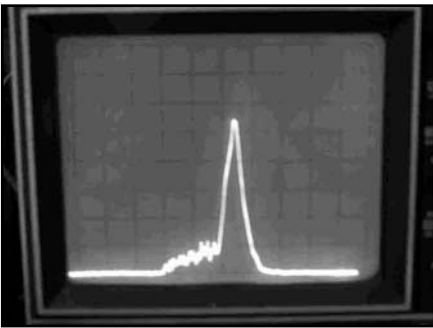
Fig 3—Measurement with a 38-MHz filter and carrier moved as close to the filter pass band as possible. We can see the noise from 5 to 10 kHz at –100 to –105 dBc/kHz (–130 to –135 dBc/Hz). Measurement parameters as those in Fig 2.
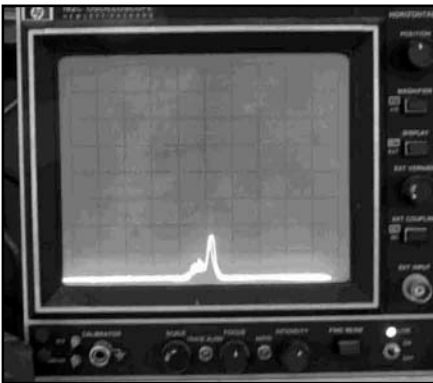
Fig 4— Measurement made with the 5.2-MHz filter. The oscillator frequency has been moved so near that we can observe the noise from 2 to 4 kHz. This picture also shows the limitation using the HP-8558B.

–5 to –10 kHz at –100 to –110 dB (–130—135 dBc/1Hz). This is also in accord with the specifications. See Fig 3.

This was very promising, but I am
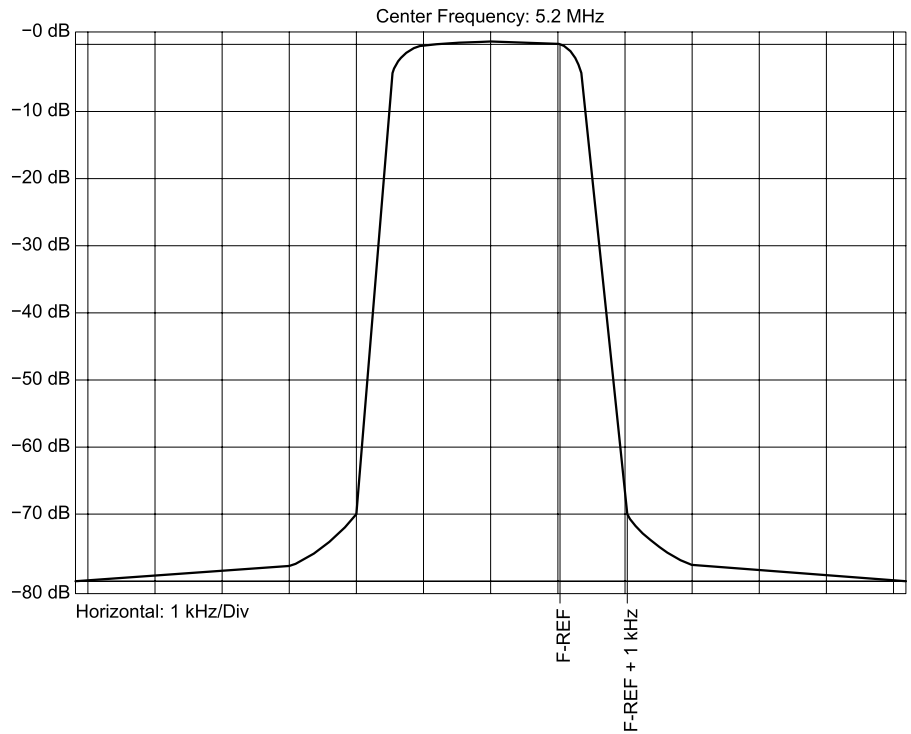
Fig 5—The attenuation of the carrier by the filter. Here the carrier is attenuated 70 dB, and we may increase the level the same amount without overdriving the spectrum analyzer. This results in a theoretical dynamic range of 140 dB.

also interested in measuring the noise from 100 Hz to 1 kHz away from the carrier. The 38 MHz crystal filter is not steep enough to get the necessary attenuation so near the carrier. Using the Tektronix 492AP with a 100 Hz bandwidth, I can measure down to around 600-700 Hz.

I then dug further down into my scrap boxes and discovered two LSB filters from a commercial HF transceiver that I worked on 30 years ago. By connecting them in series and matching them to 50 ohms, I achieved the result I wanted. Even using only one filter can be sufficient, as they are symmetrical with 30 dB of attenuation 300 Hz away from the pass-band edge on the steepest side. As I had two filters available, I decided to use both in my setup.

I now have a filter with a bandwidth of 2.4 kHz at –3 dB, and with a shape factor of 1:1.66 from –3 dB to –100 dB! The attenuation outside ±5 kHz is better than 120 dB with the lids on the box. This is sufficient to measure at a distance of 2 kHz from the carrier with the HP-8558B, and down to around 100 Hz with a better spectrum analyzer. See Fig 4.

Until now, we could only measure oscillators at the same frequency as the filter. To make this method usable on any frequency we have to add a

mixer into the system. Another search in my old spare parts bin turned up some mixers taken from old OMEGA receivers. These were Mini-Circuits SRA-3H mixers designed to operate at the 17 dBm level and covering the frequency range from 50 kHz to 200 MHz, just right for my project.

To be able to measure noise levels down to –135 to –50 dBc/Hz, you may need amplifiers ahead of the mixer both in the signal and LO path. Today you will find excellent low-noise amplifiers from several manufacturers. I use Mini-Circuits ERA-5 in the signal path and ERA-6 for the LO. A step attenuator is also necessary in both paths to calibrate the levels into the mixer. Use the recommended drive level for your particular mixer, and if you have to buy a new one, the SRA-3H is priced at $25. A 13-dBm level mixer such as the TUF-3MH is available for $10. If you want to operate at an even higher level, the RAY-6 is a 23-dBm mixer priced at $41, but it can be worth the money if you do experiments with new direct digital synthesizers. They may have phase-noise levels down to –140 dBc at 1 kHz and –150 dBc and better at 10 kHz. See Fig 6 for a schematic.

Not every ham experimenter has a spectrum analyzer, but a good receiver might also be used. Nearly every re-
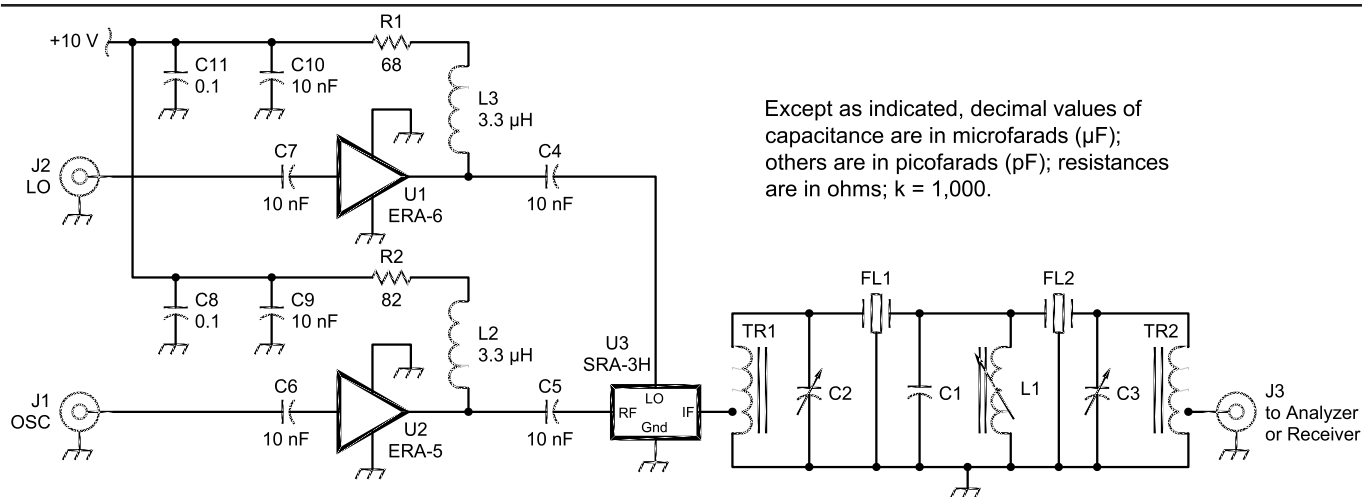
Fig 6—Schematic of the low-noise amplifiers, the mixer and the filters. Attenuators may be needed ahead of amplifiers to adjust for optimum levels to the mixer.

ceiver today can receive on all HF frequencies. I tried using my ICOM IC-756, and I got the same results, but make sure to take into account the bandwidth used. My narrowest filter is 500 Hz wide, but some of the newer rigs go down to 100 Hz, so then you must correct the readings by 20 dB instead of 30 dB when the bandwidth is 1 kHz. Remember that the filters must be in the IF, not LF.

Procedure for Using a Receiver Instead of a Spectrum Analyzer

Use a receiver with a narrow bandwidth to improve frequency resolution and it becomes a manually tuned spectrum analyzer. You must then adjust your results for the noise bandwidth of the receiver. Usually the 3-dB bandwidth points of a filter provide a good approximation. For example, the 500 Hz CW filter in my receiver requires a correction of 26 dB to convert the measurements to dBc/Hz. Using the RF amplifier provides as much sensitivity as possible. Set the frequency of the generator to the upper −3 dB point of the filter, and then down to maximum again. This is the reference frequency. Set the output level to the maximum minus 30 dB, (in my generator maximum is +13 dBm, so I use −17 dBm) then decrease the output until the S-meter shows S1. In my case it is −100 dBm. The difference is 83 dB. Let the receiver stay on the reference frequency, move the generator 1 kHz up and increase level to −17 dBm. If your filters and oscillator are good, you should have a reading of less than S1 on the meter. Increase the level (or decrease, if the oscillator is noisy) until the meter shows S1. For
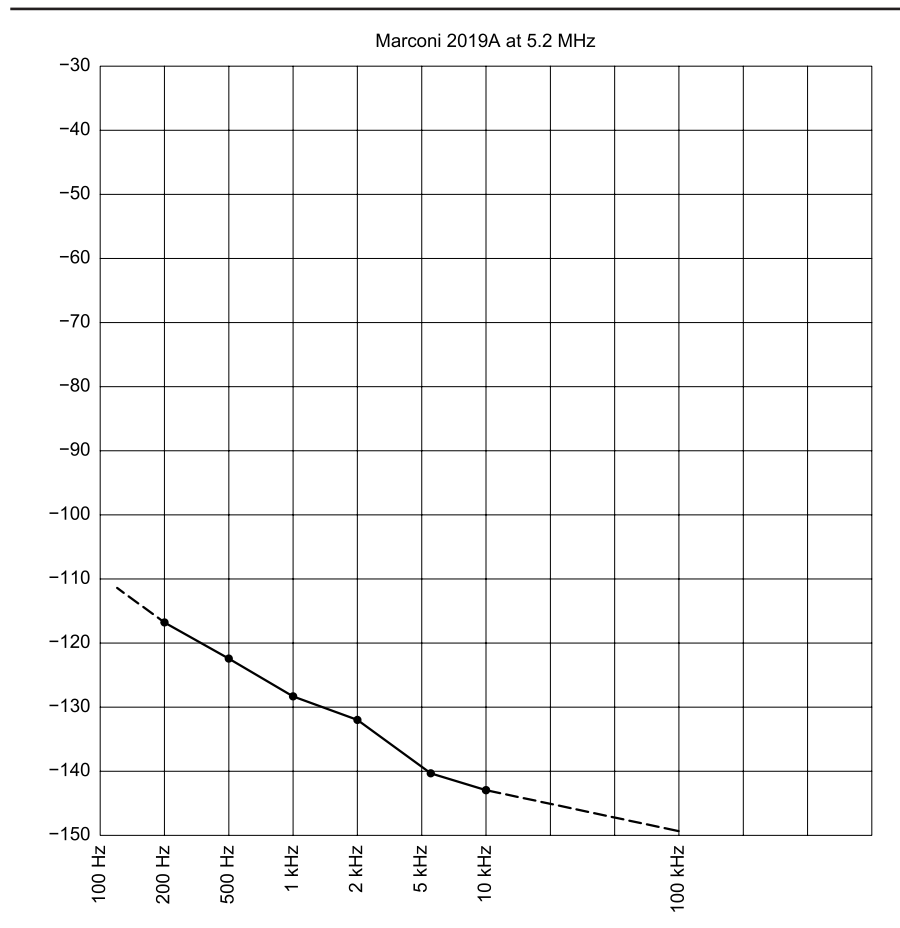
Fig 7—Marconi 2019A measured at 5.2 MHz using the ICOM IC-756.

example, I read −7 dBm giving a difference of −93 dB −26 dB=−119 dBc/Hz, then go up to +5 kHz. Increase the generator level until the S-meter reads S1. I now had +4 dBm from the generator giving a difference of −104 dB,

−26 dB = −130 dBc/Hz. At + 10 kHz. I had to increase the level to −111 dBm, −26 dB = −137 dBc/Hz. All results are the same as specifications given for the 2019A signal generator ±3 dB, except at 1 kHz off, where they are

9 dB better than specified. That is probably due to the frequency. On 5.2 MHz the noise is much lower than on 90 MHz. See Figures 7 and 8 for results using this method.

One limitation of the mixing method is that you are really measuring the noise of both the oscillator under test and the local oscillator. When you are measuring oscillators with noise levels higher than the noise in your signal generator you can use the generator as the LO, but if you are testing a really quiet oscillator you must use an even quieter crystal oscillator as the LO.

Remember that you always will have the sum of the noise from both oscillators as a result. If the two oscillators are equal, the noise from each of them is 3 dB lower than the measured value. If the difference is more than 10 to 15 dB, the additional noise from the best can be neglected.

I have not made any PC boards for this project because of the old parts I used. I think most people will also use parts they already have. If you must buy the filters, for a narrow filter, you may find that two cascaded Murata CFJ455K5 may do the job. If you can get a pair of the old XF-9B filters from KVG, they are perfect. Also filters for modern Japanese-made transceivers are very good.

The matching to 50 Ω in and out can be done with toroid transformers wound with 2, 3 or 4 twisted parallel wires connected in series to get impedance transformation factors of 4, 9 or 16 (50 Ω to 200, 450 or 800 Ω). If I remember right, the filter I use has an impedance of 640 Ω in parallel with 20-30 pF. I use the 800 Ω tap and a trimmer and the filter response is with

less than 2-dB variation in the passband. Between the filters I use direct coupling with a parallel LC circuit.

This technique is not perfect for absolute measurements of phase noise, but it has made me able to com-

pare different oscillators and tell if one is better than the other. I tried to measure the noise in an old Wavetek Model 3002 and compare it with the 2019A. I knew that the 3002 was quite noisy, but a difference of 30 dB was more
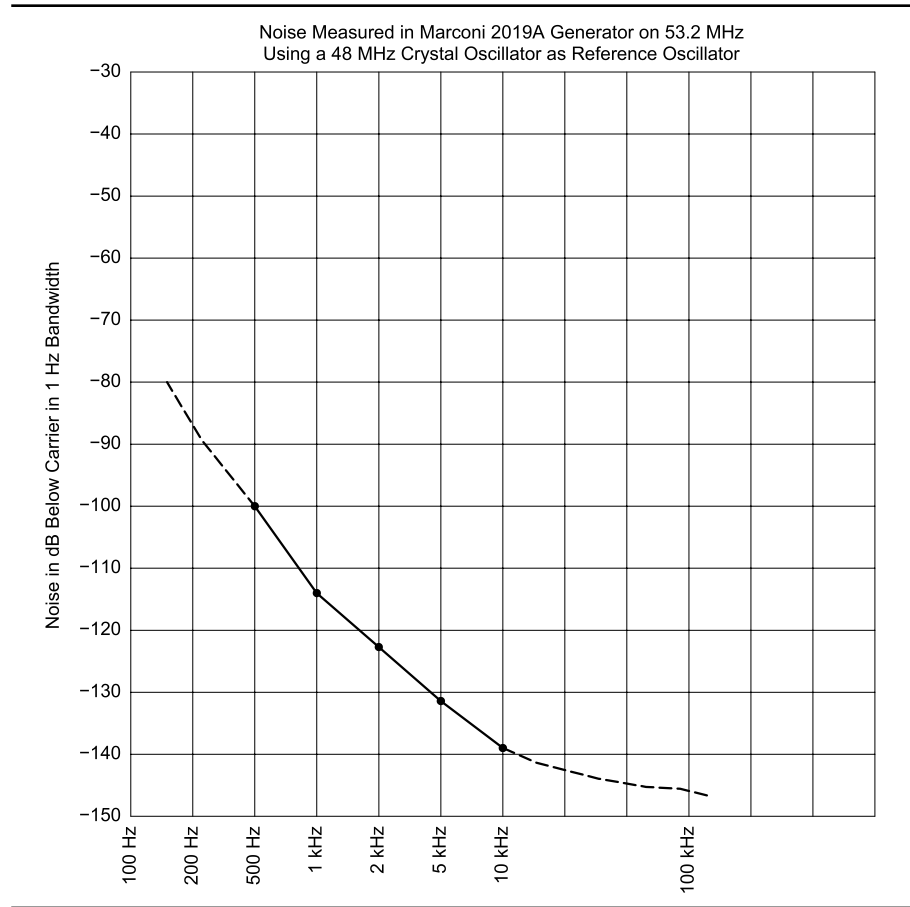
Fig 8—Marconi 2019A at 53.2 MHz mixed down to 5.2 MHz with a crystal oscillator on 48 MHz. Measured using an ICOM IC-756.
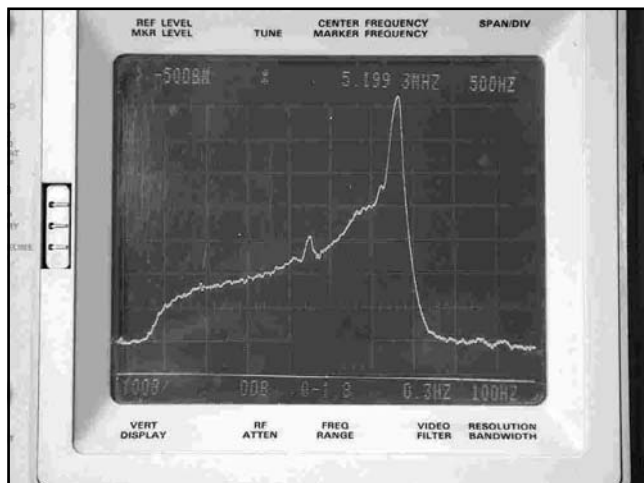
Fig 9—Measurements as described in text. Carrier frequency 53.2 MHz to 0.5 kHz. Noise level at 0.5 kHz, −115 dBc/Hz, at 2.0 kHz, −122 dBc/Hz.
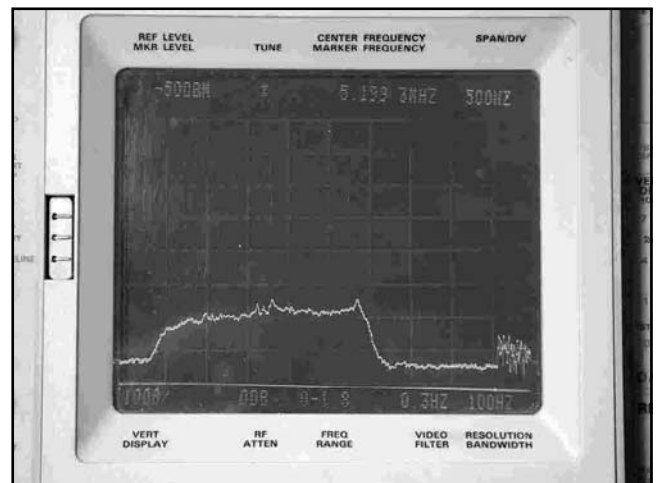
Fig 10—Measurements as described in text. Carrier frequency 53.2 MHz + 2.5 kHz. Noise level at 2.5 kHz −125 dBc/Hz, at 4.5 kHz, −132 dBc/Hz.
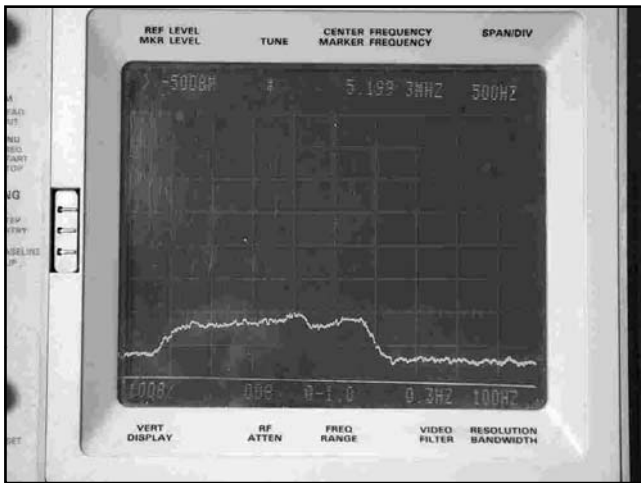
Fig 11—Measurements as described in text. Carrier frequency 53.2 MHz + 4.5 kHz. Noise level at 4.5 kHz, –132 dBc/Hz, at 6.5 kHz, –134 dBc/Hz.
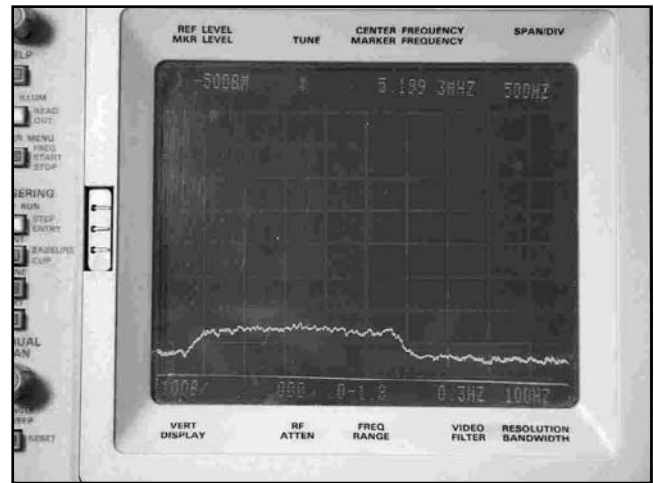
Fig 12—Measurements as described in text. Carrier frequency 53.2 MHz + 6.5 kHz. Noise level from 6.5 kHz to 8.5 kHz around –135 dBc/Hz.
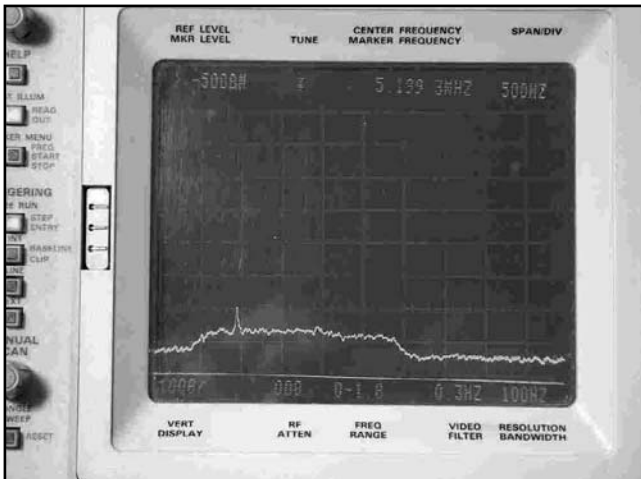
Fig 13—Measurements as described in text. Carrier frequency 53.2 MHz + 8.5 kHz. Noise level from 8.5 kHz to 10.5 kHz around –138 dBc/Hz. We have about reached the limit of the setup.
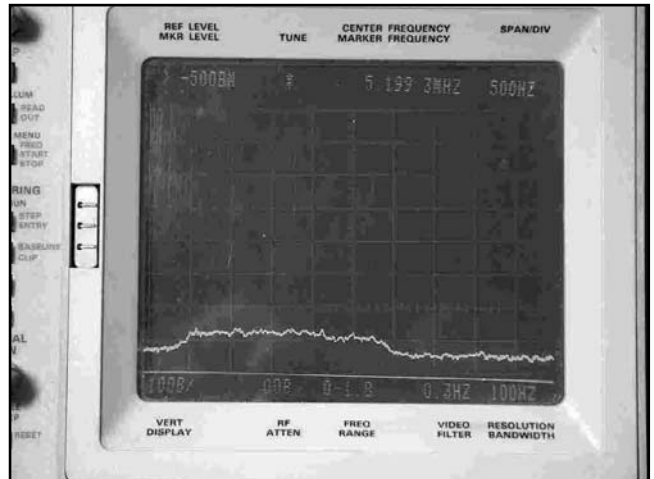
Fig 14—Measurements as described in text. Carrier frequency 53.2 MHz+10.5kHz. Noise level from 10.5 kHz to 12.5 kHz around –140 dBc/Hz. We have reached the limit of the setup since an additional offset of 100 Hz will not result in further decrease.

than I expected. I also compared the original PLL synthesizer I made in 1995 for a HF transceiver with a new DDS driven PLL synthesizer constructed recently. The first measurement showed more noise in the new one, but after correcting some problems, the new is as good as it can be with the commercial VCO I used. After replacing this VCO with one made after the description in an article by Ulrich Rohde, KA2EUW, some years ago, I got much better results with the new one. Now I can experiment with different solutions, measure and compare and have the result without guessing as I did before.

Figs 9 through 14 illustrate the results achieved using a Tektronix 492 AP spectrum analyzer, 48-MHz crystal oscillator with the 2019A as an RF source on 53.2 MHz. The reference line on the analyzer is at –50 dBc. The analyzer bandwidth is 100 kHz, meaning that we need to add 20 dB to get result in dBc/Hz. On the first picture we see the carrier near the pass band, but on the rest, the carrier is outside the picture and even invisible due to the attenuation as we move the analyzer up in frequency.

The last two pictures show that we have reached the measuring limit for this setup. If we try to move the oscillator a further 100 Hz up, we will see that the noise does not decrease further.

Kjell Karlsen, LA2NI, has been a ham since 1962, but began his interest years before as a 12-year old when a neighbor built a receiver from a kit and made it work. Kjell has worked as an *avionics engineer for the last 20 years, installing and maintaining electronic equipment in helicopters and fixed-wing aircraft. His primary amateur interest is in constructing VHF and HF equipment, including his first digital synthesizer in 1968. In 1995, he built a small 25 W HF transceiver that was used by Norwegian adventurer Børge Ousland during a solo Antarctic crossing. His current project is a direct digital synthesizer that he hopes to offer as a kit in the near future.*

Notes

[1]van der List, J.F.M., PAØJOZ, "Experiments with Phase-Noise Measurement," *QEX* Jan/Feb 1999, p 31.

[2]After I wrote this article I found software on the Internet that may do the job, and as soon as I get time I will try it out and possibly make a PCB and publish it in *QEX*. □□

Letters to the Editor

Digital Voice Articles
(QST, *Jan-Feb 2002*)

Doug,

Nice articles in *QST* Jan/Feb '02. You did an excellent job of walking the reader through the fundamentals of PSTN codecs through APCO-25 and up to the current challenges of digital voice over an HF link.

I have been trying to overcome inertia to get my license and have been around ham radio folks since I was *this* high. My father's interest in Amateur Radio contributed to my becoming an EE. Now it seems that it is coming full circle. I work constantly in RF engineering projects with digital networks from VHF up to 39 GHz. One of my current projects is the design, installation and startup of a FHSS telemetry system in Sevier County (your neighborhood) for the electric utility.

HF has always fascinated me, and digital voice modes on narrow-bandwidth links, across a fading propagation path would be *really* cool work. Your articles have given me further incentive to get off my butt and get the Morse code stuff done. I grew up with a purist ham and heard *so* much grief about "no-code tickets" that if I am going to do it, I am going to do it "right." Besides, Amateur Radio needs more women in its ranks.

Thanks—*Tisha A Hayes, Senior Communications Systems Engineer, Edison Automation Inc, Nashville,* Tennessee; **thayes@ edisonautomation.com**

Resistance—The Real Story (Jul/Aug 2004)

Hi Doug,

I enjoyed your article in the July/August *QEX*. The only part of the article that is difficult for me is the expression near the end of the third column on page 51. The expression in question, "(charge/cm^2)/(charge/cm^2 – s)=s!", is shown as printed. There seems to be a conflict of units. Specifically, subtracting time from charge per centimeter squared would have gotten me a red check mark in my freshman physics class, as an invalid use of units, as would the alternative arrangement of parentheses. I suppose another way to explain my question would be to answer how time can be subtracted from field strength or area? By the way, I'm interpreting it with either all parentheses in place, "(charge/(cm^2 – s))" [time subtracted from area] or "(charge/cm^2) – s" [time subtracted from field strength].

Thanks again for the article Doug. I enjoy *QEX* immensely. I'm still waiting for the article about "spooky action at a distance" and photon entanglement!— *Sincerely, Wayne Quernemoen, KØRCH;* **wpq@rea-alp.com**

Author's reply:

Hi Wayne,

That is supposed to be a hyphen, not a subtraction sign. We had a bit of last-minute confusion during editing that resulted in inconsistent representations like that. I'm sorry for the confusion. I suppose we could have just written cm^{2}s without too much trouble.—*73, Doug Smith, KF6DX,* QEX *Editor;* **kf6dx@arrl.org**

Networks for 8-Direction 4-Square Arrays (Sep/Oct 2004)

Hello Doug,I just received my complimentary copies of the Sept/Oct issue of *QEX*, which contains my article. I'm pleased with the way it turned out, but I noticed a couple of small typographical errors.

1. On page 35, in the final paragraph in the left-hand column, nine lines up from the bottom, it should say, "Equalizing the input resistances *is* accomplished. . ."

2. There are several places scattered throughout the text where I mention parallel combinations of impedances. For lack of an appropriate symbol, the notation "11" was used to represent "in parallel with." Thus, Za in parallel with Zb would be represented as "(Za)11(Zb)" within the body of the text.

3. I chose to use the prime (′) notation in conjunction with some of my input parameters and this notation is used correctly in most of the text. On several occasions, however, an apostrophe was used instead. The first place I noticed this is on page 40, for Vss′, and it appears on page 43 for Vff″.

4. On page 44, in the right-hand column of text, two lines above Table 6, it should say "0.25-WL phasing lines..." Somehow, an upper-case Greek letter Omega crept in there.

5. Also on page 44, in Fig 11, the "Z Match" heading should be placed beneath the impedance-matching network, which consists of the 1611-pF capacitor and the 0.33-µH inductor.

6. In Fig 15, in the small box that lists the relays, the first line should read, " K1, K2, K4, K5, K7-K10 = DPST"

Thanks and 73—*Al Christman, K3LC;* **amchristman@gcc.edu**

Letters (Sept/Oct 2004)

Dear Mr. Smith,

I was *most* impressed by your response to Mr. Czuhajewski's letter. Some time ago the then *QEX* editor posed a question about what could be done to reenergize Amateur Radio— was that you? I could not bring myself to write a non-critical answer. So it is probably good I did not write.

Your answer suggests that maybe a word or two might be helpful.

From what I can see, *QST* has two types of articles: "Gee this commercial radio is great" and "let me tell you how to solder the wires on this connector." Gone are the "you can build it" and "this is how it works" stuff.... Why is this? It seems to me that perhaps all organizations get pretty comfortable with the status quo: If we just don't rock the boat, we will continue to get paid and get lots of free lunches and other perks. Then, one day we will retire and live happily ever after.

Thus, when I first became aware of it, I particularly appreciated *QEX*. It has real content about real stuff. So much so, that I bought the CDs of back issues only to learn that unless I loaded the right version of Microsoft software, I could not meaningfully read them!

So, who let that happen? Oh—And yes, it would be nice if that problem were fixed!

Anyway, that is why I decided to write you a thank-you for your simple reply. I hope that was followed by an unwritten [reply]: "I'll try not to let this happen again in my magazine." Thanks for your hard work and dedication.

Regards—*John Harrison, NI1B;* **jmh5@nei.mv.com**

Improved Remote Antenna Impedance Measurement (Jul/Aug 2004)

Doug:

Thank you for promptly starting my subscription to *QEX*. Issues for Jul/Aug 2004 and Sept/Oct 2004 arrived yesterday. I'm enjoying reading them.

In *QEX* for Sept/Oct 2004, on page 59, Ron Barker indicates that the spreadsheet described in his article appearing in the previous issue is posted on the ARRL Web site **www. arrl.org/qexfiles**. My check of this site this morning didn't turn up Ron's spreadsheet. Has the posting not yet taken place or is it in another location?

Thanks—*Dale Covington;* **dwcov @bellsouth.net**

Hi Doug,

Letters to the Editor of the Sept/Oct 2004 *QEX* indicated that Ron Barker's spreadsheet for his article "Improved Remote Antenna Impedance Measurement' in the Jul/Aug would be in the *QEX* files section of the ARRL Web page. I checked today and I could not find it. Am I looking in the wrong place?

I enjoy your magazine! Please include information on where we can get parts or kits for the projects mentioned in the articles. It's frustrating to read an interesting article and find that the parts mentioned are not longer available.—*73, Mike St. Angelo, N2MS;* **mstangelo@comcast.net**

Hi guys,

Sorry, we got behind on our postings but it is now there. Regarding kits and parts, we do make sure you have contact information for the author, and your best bet is to contact Ron directly. He'll likely be glad to help you. —*73, Doug*

On Signal-to-Noise Ratio and Decision-Making

Hi Doug,

I have been musing over the application and value of SNR in aural CW and in the decision-making taking place in data communications. SNR is defined in the texts as Eb/No, where Eb is the energy per bit and No is the noise power density. Energy in joules =watt-seconds [*a hyphen, not a subtraction sign—Ed.*] with dimensions of power times time. Noise power density is watts/hertz. As Hz has the dimension of 1/time, the units cancel and thus SNR is dimensionless. This is all very well for textbooks, but it doesn't seem to me to fit well with practical communications.

Take the case of CW. For a fixed bandwidth, what I hear with my ears is an improved SNR when the signal energy is increased by lengthening the dots. The noise appears bandwidth-dependent and not time-dependent as it would by obeying the SNR formula. Somehow, our acoustic powers differentiate between the coherent signal and the random noise in a way not suggested by Eb/No.

I also understand that some DSP "de-noise" filtering algorithms also exploit the different character of signals and noise. Doesn't Shannon have something to say about entropy? If any of this makes sense perhaps it is worth some discussion from you in *QEX*?— *73, Ron Skelton, W6WO;* **ron-skelton @charter.net**

Hi Ron,

Yes, that is a fascinating subject. I went into it a little bit in an article a few years back during my research into human hearing: "PTC: Perceptual Transform Coding...." in *QEX*, May 2000. I guess the main obstacle to quantifying those things is that you can only ask questions of the listener and try to glean something from the responses. CW as received by ear may be an exception because the listener is required to copy the code. Then one would need to be sure that the listener's basic code-copying ability with strong signals was beyond reproach to remove bias. It would be interesting to see the spread in noisy-signal copying ability from one listener to the next.

I notice that EME (moonbounce) fans are up against some of the toughest conditions in this area and I guess they do indeed slow their code speeds quite a bit. On a path where fading or multipath is present, the SNR may be changing and then the difficulty with CW becomes one of judging the duration of each element. Therefore, it seems that reducing speed beyond the fading rate produces another issue.

It is unclear to me whether the ear-brain combination does anything like what we do in DSP noise-reduction algorithms, but I suspect that it does. The basic idea is that noise does not repeat itself exactly over relatively long time frames and the desired signal does. Coherent CW fans have overcome some of the general difficulties, I think. Thanks for the suggestion! —*73, Doug* □□

Reductio Ad Absurdium and the Square Root of Two

The following is a partial reproduction of the proof from Appendix 1 of Carl Sagan's book *Cosmos* (Random House, 1983, ISBN 0-39471-596-9).

"We assume $\sqrt{2}$ is a rational number: $\sqrt{2}=p/q$, where p and q are integers, whole numbers. They can be as big as we like and can stand for any integers we like. We can certainly require that they have no common factors. If we were to claim $\sqrt{2}=14/10$, for example, we would of course cancel out the factor 2 and write $p=7$ and $q=5$, not $p=14$. $q=10$. Any common factor in the numerator or denominator would be canceled out before we start. There are an infinite number of ps and qs we can choose. From $\sqrt{2}=p/q$, by squaring both sides of the equation, we find that $2=p^2/q^2$, or, by multiplying both sides of the equation by q^2, we find $p^2=2q^2$.

p is then some number multiplied by 2. Therefore, p^2 is an even number. But the square of any odd number is odd ($1^2=1$, $3^2=9$, $5^2=25$, $7^2=49$, etc). So p itself must be even, and we can write $p=2s$, where s is some other integer. Substituting for p we find $p^2=(2s)^2=4s^2=2q^2$. Dividing both sides of the last equality by 2, we find $q^2=2s^2$.

Therefore q^2 is also an even number, and, by the same argument as we just used for p, it follows that q is even too. But if p and q are both even, both divisible by 2, then they have not been reduced to their lowest common factor, contradicting one of our assumptions. *Reductio ad absurdium*. But which assumption? The argument cannot be telling us that reduction to common factors is forbidden, that 14/10 is permitted and 7/5 is not. So the initial assumption must be wrong; p and q cannot be whole numbers; and $\sqrt{2}$ is irrational.

What a stunning and unexpected conclusion! How elegant the proof! But the Pythagoreans felt compelled to suppress this great discovery."